

8149938



THE UNITED STATES OF AMERICA

TO ALL TO WHOM THESE PRESENTS SHALL COME:

UNITED STATES DEPARTMENT OF COMMERCE

United States Patent and Trademark Office

August 24, 2021

THIS IS TO CERTIFY THAT ANNEXED IS A TRUE COPY FROM THE RECORDS OF THIS OFFICE OF THE FILE WRAPPER AND CONTENTS OF:

APPLICATION NUMBER: 16/278,107
FILING DATE: February 17, 2019
PATENT NUMBER: 10484510
ISSUE DATE: November 19, 2019



Certified by

Under Secretary of Commerce
for Intellectual Property
and Director of the United States
Patent and Trademark Office

MAY PATENTS LTD.
Yehuda BINDER, U.S. Patent Agent
B.Sc.E.E.; M.Sc.E.E; M.B.A

February 15, 2019

U.S. Patent and Trademark Office (USPTO)
Customer Service Window
Mail Stop Patent Application
401 Dulany Street
Alexandria, VA 22314

Re: **New Utility Patent Application in U.S.**
Applicant(s): **WEB SPARK LTD.**
**Title: SYSTEM PROVIDING FASTER AND MORE EFFICIENT DATA
COMMUNICATION**
Attorney Docket No.: HOLA-005-U10

Sir/Madam,

Attached herewith is the above-identified application for Letters Patent including:

1. Applicant asserts small entity status. See 37 CFR 1.27.
2. Application Data Sheet (PTO/AIA/14);
3. Specification (33 pages), Claims 1-24 (4 pages) and abstract (1 page).
4. Fifteen (15) sheets of Drawings (Figures 1-15).
5. Declaration
[X] Newly executed [] Copy from prior application no. _____
6. Assignment submitted through EPAS
[X] Newly executed [] Copy from prior application no. _____
7. Power of Attorney
8. Information Disclosure Statement (PTO/SB/08).

Certain documents were previously filed or cited to the USPTO in the prior application 15/957,945, which is relied upon under 35 U.S.C. § 120. Applicant(s) identify these documents by attaching an Information Disclosure Statement listing these documents and request that they be considered and made of record in accordance with 37 CFR § 1.98(d). Per Section 1.98(d), copies of these documents need not be filed in the application.

e-mail: yehuda@maypatents.com; Mobile: +972-54-4444577

.....
MAY PATENTS LTD.
Yehuda BINDER, U.S. Patent Agent
B.Sc.E.E.; M.Sc.E.E; M.B.A
.....

9. Electronic Payment in the amount of US\$ 985 is being made by deposit account no. 506726 to cover filing fee calculated as follows:

Application as Filed				Fee (US\$)
Basic Filing Fee				75.00
Search Fee				330.00
Examination Fee				380.00
Total Sheets	Extra Sheets		Rate	
54/100	--		200.00	--
Claims	# Filed	# Extra	Rate	
Total Claims	1-24	4	50.00	200.00
Independent Claims	3	--	230.00	--
Multiple Dependent Claim			410.00	--
Total Filing Fee				985.00

1. As in the prior application 15/957,945, please associate the above referenced application with **Customer No. 131926**.
2. The Correspondence Address associated with **Customer No. 131926**.

Submitted by,
May Patents Ltd.

By: /Yehuda Binder/
Yehuda Binder
Registration No. 73,612

e-mail: yehuda@maypatents.com; Mobile: +972-54-4444577

SYSTEM PROVIDING FASTER AND MORE EFFICIENT DATA COMMUNICATION

CROSS-REFERENCE TO RELATED APPLICATIONS

The present application is a continuation application of U.S. non-provisional patent application no. 15/957,945, filed Apr. 20, 2018, which is a continuation application of U.S. non-provisional patent application no. 14/025,109, filed Sep. 12, 2013 and issued as U.S. Patent No. 10,069,936 on Sep. 04, 2018, which is a divisional application of U.S. non-provisional patent application entitled "SYSTEM AND METHOD FOR PROVIDING FASTER AND MORE EFFICIENT DATA COMMUNICATION" having Ser. No. 12/836,059, filed Jul. 14, 2010 and issued as U.S. Patent No. 8,560,604 on Oct. 15, 2013, and claims priority to U.S. provisional patent application entitled "SYSTEM AND METHOD FOR REDUCING INTERNET CONGESTION," having Ser. No. 61/249,624, filed Oct. 8, 2009, which are hereby incorporated herein by reference in their entirety.

FIELD OF THE INVENTION

The present invention is related to Internet communication, and more particularly, to improving data communication speed and bandwidth efficiency on the Internet.

BACKGROUND OF THE INVENTION

There are several trends in network and Internet usage, which tremendously increase the bandwidth that is being used on the Internet. One such trend is that more and more video is being viewed on demand on the Internet. Such viewing includes the viewing of both large and short video clips. In addition, regular shows and full-featured films may be viewed on the Internet. Another trend that is increasing the traffic on the Internet is that Web sites (such as shopping portals, news portals, and social networks) are becoming global, meaning that the Web sites are serving people in many diverse places on the globe, and thus the data is traversing over longer stretches of the Internet, increasing the congestion.

The increase in bandwidth consumption has created several major problems, a few of which are described below:

The problem for users – the current Internet bandwidth is not sufficient, and thus the effective ‘speed’ experienced by users is slow;

The problem for content owners – the tremendous amount of data being viewed by users is costing large amounts of money in hosting and bandwidth costs; and

The problem for Internet Service Providers (ISPs) – the growth in Internet traffic is requiring the ISPs to increase the infrastructure costs (communication lines, routers, etc.) at tremendous financial expense.

The need for a new method of data transfer that is fast for the consumer, cheap for the content distributor and does not require infrastructure investment for ISPs, has become a major issue which is yet unsolved.

There have been many attempts at making the Internet faster for the consumer and cheaper for the broadcaster. Each such attempt is lacking in some aspect to become a widespread, practical solution, or is a partial solution in that it solves only a subset of the major problems associated with the increase in Internet traffic. Most of the previous solutions require billions of dollars in capital investment for a comprehensive solution. Many of these attempts are lacking in that much of the content on the Internet has become dynamically created per the user and the session of the user (this is what used to be called the “Web2.0” trend). This may be seen on the Amazon Web site and the Salesforce Web site, for example, where most of the page views on these Web sites is tailored to the viewer, and is thus different for any two viewers. This dynamic information makes it impossible for most of the solutions offered to date to store the content and provide it to others seeking similar content.

One solution that has been in use is called a “proxy”. FIG. 1 is a schematic diagram providing an example of use of a proxy within a network 2. A proxy, or proxy server 4, 6, 8 is a device that is placed between one or more clients, illustrated in FIG. 1 as client devices 10, 12, 14, 16, 18, 20, that request data, via the Internet 22, and a Web server or Web servers 30, 32, 34 from which they are requesting the data. The proxy server 4, 6, 8 requests the data from the Web servers 30, 32, 34 on their behalf, and caches the responses from the Web servers 30, 32, 34, to provide to other client devices that make similar requests. If the proxy server 4, 6, 8 is geographically close enough to the client devices 10, 12, 14, 16, 18, 20, and if the storage and bandwidth of the proxy server 4, 6, 8 are large enough, the proxy server 4, 6, 8 will speed up the requests for the client devices 10, 12, 14, 16, 18, 20 that it is serving.

It should be noted, however, that to provide a comprehensive solution for Internet surfing, the proxy servers of FIG. 1 would need to be deployed at every point around the world where the Internet is being consumed, and the storage size of the proxy servers at each location would need to be near the size of all the data stored anywhere on the Internet. The abovementioned would lead to massive costs that are impractical. In addition, these proxy solutions cannot deal well with dynamic data that is prevalent now on the Web.

There have been commercial companies, such as Akamai, that have deployed such proxies locally around the world, and that are serving a select small group of sites on the Internet. If all sites on the Web were to be solved with such a solution, the capital investment would be in the range of billions of dollars. In addition, this type of solution does not handle dynamic content.

To create large distribution systems without the large hardware costs involved with a proxy solution, “peer-to-peer file sharing” solutions have been introduced, such as, for example, BitTorrent. FIG. 2 is a schematic diagram providing an example of a peer-to-peer file transfer network 50. In the network 50, files are stored on computers of consumers, referred to herein as

client devices 60. Each consumer can serve up data to other consumers, via the Internet 62, thus taking the load of serving off of the distributors and saving them the associated costs, and providing the consumer multiple points from which to download the data, referred to herein as peers 70, 72, 74, 76, 78, thus increasing the speed of the download. However, each such peer-to-peer solution must have some sort of index by which to find the required data. In typical peer-to-peer file sharing systems, because the index is on a server 80, or distributed among several servers, the number of files available in the system is not very large (otherwise, the server costs would be very large, or the lookup time would be very long).

The peer-to-peer file sharing solution is acceptable in file sharing systems, because there are not that many media files that are of interest to the mass (probably in the order of magnitude of millions of movies and songs that are of interest). Storing and maintaining an index of millions of entries is practical technically and economically. However, if this system were to be used to serve the hundreds of billions of files that are available on the Internet of today, the cost of storing and maintaining such an index would be again in the billions of dollars. In addition, these types of peer-to-peer file sharing systems are not able to deal with dynamic HTTP data.

In conclusion, a system does not exist that enables fast transmission of most of the data on the Internet, that does not incur tremendous costs, and/or that provides only a very partial solution to the problem of Internet traffic congestion. Thus, a heretofore unaddressed need exists in the industry to address the aforementioned deficiencies and inadequacies.

SUMMARY OF THE INVENTION

The present system and method provides for faster and more efficient data communication within a communication network. Briefly described, in architecture, one embodiment of the system, among others, can be implemented as follows. A network is provided

for accelerating data communication, wherein the network contains: at least one client communication device for originating a data request for obtaining the data from a data server; at least one agent communication device which is assigned to the data server for receiving the data request from the client communication device, wherein the agent keeps track of which client communication devices have received responses to data requests from the assigned data server; at least one peer communication device for storing portions of data received in response to the data request by the at least one client communication device, wherein the portions of data may be transmitted to the at least one client communication device upon request by the client communication device; and at least one acceleration server for deciding which agent communication device is to be assigned to which data server and providing this information to the at least one client communication device.

The present system and method also provides a communication device within a network, wherein the communication device contains: a memory; and a processor configured by the memory to perform the steps of: originating a data request for obtaining data from a data server; being assigned to a data server, referred to as an assigned data server; receiving a data request from a separate device within the network, and keeping track of which client communication devices within the network have received responses to data requests from the assigned data server; and storing portions of data received in response to the originated data request, wherein the portions of data may be transmitted to communication device upon request by the communication device.

Other systems, methods, features, and advantages of the present invention will be or become apparent to one with skill in the art upon examination of the following drawings and detailed description. It is intended that all such additional systems, methods, features, and advantages be included within this description, be within the scope of the present invention, and be protected by the accompanying claims.

BRIEF DESCRIPTION OF THE DRAWINGS

Many aspects of the invention can be better understood with reference to the following drawings. The components in the drawings are not necessarily to scale, emphasis instead being placed upon clearly illustrating the principles of the present invention. Moreover, in the drawings, like reference numerals designate corresponding parts throughout the several views.

FIG. 1 is a schematic diagram providing a prior art example of use of a proxy within a network.

FIG. 2 is a schematic diagram providing a prior art example of a peer-to-peer file transfer network.

FIG. 3 is a schematic diagram providing an example of a communication network in accordance with the present invention.

FIG. 4 is a schematic diagram further illustrating a communication device of the communication network of FIG. 3.

FIG. 5 is a schematic diagram further illustrating the memory of FIG. 4.

FIG. 6 is a schematic diagram further illustrating elements of the acceleration application of FIG. 5, as well as communication paths of the acceleration application.

FIG. 7 is a chart further illustrating two of the main databases utilized within the communication network.

FIG. 8 is a flowchart illustrating operation of the acceleration system initializer module.

FIG. 9 is a flowchart further illustrating communication between different elements of the communication network.

FIG. 10 is a flowchart continuing the flowchart of FIG. 9 and focused on agent response to the HTTP request.

FIG. 11 is a flowchart continuing the flowchart of FIG. 10, which illustrates actions taken upon receipt of the list of peers, or single peer listing, from the agent.

FIG. 12 is a flowchart illustrating steps taken by an agent, client, or peer to determine whether a certain HTTP request is still valid.

FIG. 13 is a flowchart outlining operation of the acceleration server.

FIG. 14 is a flowchart further illustrating TCPIP acceleration in accordance with an alternative embodiment of the invention.

FIG. 15 is a flowchart further illustrating TCPIP acceleration in accordance with an alternative embodiment of the invention, detailing the communication between the client and the TCPIP server (read and write commands) after the connect phase has completed successfully.

DETAILED DESCRIPTION

The present system and method provides for faster and more efficient data communication within a communication network. An example of such a communication network 100 is provided by the schematic diagram of FIG. 3. The network 100 of FIG. 3 contains multiple communication devices. Due to functionality provided by software stored within each communication device, which may be the same in each communication device, each communication device may serve as a client, peer, or agent, depending upon requirements of the network 100, as is described in detail herein. It should be noted that a detailed description of a communication device is provided with regard to the description of FIG. 4.

Returning to FIG. 3, the exemplary embodiment of the network 100 illustrates that one of the communication devices is functioning as a client 102. The client 102 is capable of communication with one or more peers 112, 114, 116 and one or more agents 122. For exemplary purposes, the network contains three peers and one agent, although it is noted that a client can communicate with any number of agents and peers.

The communication network 100 also contains a Web server 152. The Web server 152 is the server from which the client 102 is requesting information and may be, for example, a typical HTTP server, such as those being used to deliver content on any of the many such servers on the Internet. It should be noted that the server 152 is not limited to being an HTTP server. In fact, if a different communication protocol is used within the communication network, the server may be a server capable of handling a different protocol. It should also be noted that while the present description refers to the use of HTTP, the present invention may relate to any other communication protocol and HTTP is not intended to be a limitation to the present invention.

The communication network 100 further contains an acceleration server 162 having an acceleration server storage device 164. As is described in more detail herein, the acceleration server storage device 164 has contained therein an acceleration server database. The acceleration server database stores Internet protocol (IP) addresses of communication devices within the communication network 100 having acceleration software stored therein. Specifically, the acceleration server database contains stored therein a list of communication devices having acceleration software stored therein that are currently online within the communication network 100. For each such agent, the acceleration server assigns a list of IP addresses.

In the communication network 100 of FIG. 3, the application in the client 102 is requesting information from the Web server 152, which is why the software within the communication device designated this communication device to work as a client. In addition, since the agent 122 receives the request from the client 102 as the communication device closest

to the Web server 152, functionality of the agent 122, as provided by the software of the agent 122, designates this communication device to work as an agent. It should be noted, that in accordance with an alternative embodiment of the invention, the agent need not be the communication device that is closest to the Web server. Instead, a different communication device may be selected to be the agent.

Since the peers 112, 114, 116 contain at least portions of the information sought by the client 102 from the Web server 152, functionality of the peers 112, 114, 116, as provided by the software of the peers 112, 114, 116, designates these communication devices to work as peers. It should be noted that the process of designating clients, agents, and peers is described in detail herein. It should also be noted that the number of clients, agents, peers, acceleration servers, Web servers, and other components of the communication network 100 may differ from the number illustrated by FIG. 3. In fact, the number of clients, agents, peers, acceleration servers, Web servers, and other components of the communication network 100 are not intended to be limited by the current description.

Prior to describing functionality performed within a communication network 100, the following further describes a communication device 200, in accordance with a first exemplary embodiment of the invention. FIG. 4 is a schematic diagram further illustrating a communication device 200 of the communication network 100, which contains general components of a computer. As previously mentioned, it should be noted that the communication device 200 of FIG. 4 may serve as a client, agent, or peer.

Generally, in terms of hardware architecture, as shown in FIG. 4, the communication device 200 includes a processor 202, memory 210, at least one storage device 208, and one or more input and/or output (I/O) devices 240 (or peripherals) that are communicatively coupled via a local interface 250. The local interface 250 can be, for example but not limited to, one or more buses or other wired or wireless connections, as is known in the art. The local interface 250

may have additional elements, which are omitted for simplicity, such as controllers, buffers (caches), drivers, repeaters, and receivers, to enable communications. Further, the local interface 250 may include address, control, and/or data connections to enable appropriate communications among the aforementioned components.

The processor 202 is a hardware device for executing software, particularly that stored in the memory 210. The processor 52 can be any custom made or commercially available processor, a central processing unit (CPU), an auxiliary processor among several processors associated with the communication device 200, a semiconductor based microprocessor (in the form of a microchip or chip set), a macroprocessor, or generally any device for executing software instructions.

The memory 210, which is further illustrated and described by the description of FIG. 5, can include any one or combination of volatile memory elements (*e.g.*, random access memory (RAM, such as DRAM, SRAM, SDRAM, *etc.*)) and nonvolatile memory elements (*e.g.*, ROM, hard drive, tape, CDROM, *etc.*). Moreover, the memory 210 may incorporate electronic, magnetic, optical, and/or other types of storage media. Note that the memory 210 can have a distributed architecture, where various components are situated remote from one another, but can be accessed by the processor 202.

The software 212 located within the memory 210 may include one or more separate programs, each of which contains an ordered listing of executable instructions for implementing logical functions of the communication device 200, as described below. In the example of FIG. 4, the software 212 in the memory 210 at least contains an acceleration application 220 and an Internet browser 214. In addition, the memory 210 may contain an operating system (O/S) 230. The operating system 230 essentially controls the execution of computer programs and provides scheduling, input-output control, file and data management, memory management, and communication control and related services. It should be noted that, in addition to the

acceleration application 220, Internet browser 214, and operating system 230, the memory 210 may contain other software applications.

While the present description refers to a request from the client originating from an Internet browser, the present invention is not limited to requests originating from Internet browsers. Instead, a request may originate from an email program or any other program that would be used to request data that is stored on a Web server, or other server holding data that is requested by the client device.

Functionality of the communication device 200 may be provided by a source program, executable program (object code), script, or any other entity containing a set of instructions to be performed. When a source program, then the program needs to be translated via a compiler, assembler, interpreter, or the like, which may or may not be included within the memory 210, so as to operate properly in connection with the operating system 230. Furthermore, functionality of the communication device 200 can be written as (a) an object oriented programming language, which has classes of data and methods, or (b) a procedure programming language, which has routines, subroutines, and/or functions.

The I/O devices 240 may include input devices, for example but not limited to, a keyboard, mouse, scanner, microphone, *etc.* Furthermore, the I/O devices 240 may also include output devices, for example but not limited to, a printer, display, *etc.* Finally, the I/O devices 240 may further include devices that communicate via both inputs and outputs, for instance but not limited to, a modulator/demodulator (modem; for accessing another device, system, or network), a radio frequency (RF) or other transceiver, a telephonic interface, a bridge, a router, *etc.*

When the communication device 200 is in operation, the processor 202 is configured to execute the software 212 stored within the memory 210, to communicate data to and from the memory 210, and to generally control operations of the communication device 200 pursuant to

the software 212. The software 212 and the O/S 230, in whole or in part, but typically the latter, are read by the processor 202, perhaps buffered within the processor 202, and then executed.

When functionality of the communication device 200 is implemented in software, as is shown in FIG. 4, it should be noted that the functionality can be stored on any computer readable medium for use by or in connection with any computer related system or method. In the context of this document, a computer readable medium is an electronic, magnetic, optical, or other physical device or means that can contain or store a computer program for use by or in connection with a computer related system or method. The functionality of the communication device 200 can be embodied in any computer-readable medium for use by or in connection with an instruction execution system, apparatus, or device, such as a computer-based system, processor-containing system, or other system that can fetch the instructions from the instruction execution system, apparatus, or device and execute the instructions. In the context of this document, a "computer-readable medium" can be any means that can store, communicate, propagate, or transport the program for use by or in connection with the instruction execution system, apparatus, or device.

The computer readable medium can be, for example but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, device, or propagation medium. More specific examples (a non-exhaustive list) of the computer-readable medium would include the following: an electrical connection (electronic) having one or more wires, a portable computer diskette (magnetic), a random access memory (RAM) (electronic), a read-only memory (ROM) (electronic), an erasable programmable read-only memory (EPROM, EEPROM, or Flash memory) (electronic), an optical fiber (optical), and a portable compact disc read-only memory (CDROM) (optical). Note that the computer-readable medium could even be paper or another suitable medium upon which the program is printed, as the program can be electronically captured, via for instance optical scanning of the paper or other medium, then

compiled, interpreted or otherwise processed in a suitable manner if necessary, and then stored in a computer memory.

In an alternative embodiment, where the functionality of the communication device 200 is implemented in hardware, the functionality can be implemented with any or a combination of the following technologies, which are each well known in the art: a discrete logic circuit(s) having logic gates for implementing logic functions upon data signals, an application specific integrated circuit (ASIC) having appropriate combinational logic gates, a programmable gate array(s) (PGA), a field programmable gate array (FPGA), *etc.*

The at least one storage device 208 of the communication device 200 may be one of many different categories of storage device. As is described in more detail herein, the storage device 208 may include a configuration database 280 and a cache database 282. Alternatively, the configuration database 280 and cache database 282 may be located on different storage devices that are in communication with the communication device 200. The description that follows assumes that the configuration database 280 and cache database 282 are located on the same storage device, however, it should be noted that the present invention is not intended to be limited to this configuration.

The configuration database 280 stores configuration data that is common to all elements of the communication network 100 and is used to provide set up and synchronization information to different modules of the acceleration application 220 stored within the memory 210, as is described in further detail herein. The cache database 282 stores responses to HTTP requests that the communication device 200 has dispatched, either for its own consumption or on behalf of other elements of the communication network 100. As is explained in additional detail herein, the responses to HTTP requests are stored within the cache database 282 for future use by this communication device 200, or for other communication devices within the

communication network 100 that need to retrieve this information and will use this communication device as either a peer or an agent.

In addition to the abovementioned, as is explained in further detail herein, the cache database 282 has stored therein a list of URLs that the communication device is aware of (i.e., has seen requests for). For each URL, the cache database 282 has stored therein the URL itself, HTTP headers returned by the Web Server for this URL, when the last time was that the contents of this URL was loaded directly from the Web Server, when the contents of the URL had last changed on the Web Server, as well as a list of chunks that contain the contents of this URL, and the chunks of data themselves. Chunks in the present description are defined as equally sized pieces of data that together form the whole content of the URL. It should be noted that while the present description provides for chunks being equally sized pieces of data, in accordance with an alternative embodiment of the invention, the chunks may instead be of different size.

FIG. 5 is a schematic diagram further illustrating the memory 210 of FIG. 4. As shown by FIG. 5, the memory 210 may be separated into two basic levels, namely, an operating system level 260 and an application level 270. The operating system level 260 contains the operating system 230, wherein the operating system 230 further contains at least one device driver 262 and at least one communication stack 264. The device drivers 262 are software modules that are responsible for the basic operating commands for various hardware devices of the communication device 200, such as the processor 202, the storage device 208 and the I/O devices 240. In addition, the communication stacks 264 provide applications of the communication device 200 with a means of communicating within the network 100 by implementing various standard communication protocols.

The application level 270 includes any application that is running on the communication device 200. As a result, the application level 270 includes the Internet browser 214, which is used to view information that is located on remote Web servers, the acceleration application 220,

as described in more detail below, and any other applications 216 stored on the communication device 200.

As is explained in additional detail below, the acceleration application 220 intercepts the requests being made by applications of the communication device (client) that use the Internet, in order to modify the requests and route the requests through the communication network. There are various methods that may be used to intercept such requests. One such method is to create an intermediate driver 272, which is also located within the memory 210, that attaches itself to all communication applications, intercepts outgoing requests of the communication applications of the communication device 200, such as the Internet browser 214, and routes the requests to the acceleration application 220. Once the acceleration application 220 modifies the requests, routes the requests to other system elements on the communication network 100, and receives replies from other system elements of the communication network 100, the acceleration application 220 returns the replies to the intermediate driver 272, which provides the replies back to the requesting communication application.

FIG. 6 is a schematic diagram further illustrating elements of the acceleration application 220, as well as communication paths of the acceleration application 220. The acceleration application 220 contains an acceleration system initializer module 222, which is called when the acceleration application 220 is started. The acceleration system initializer module 222 is capable of initializing all elements of the communication device 200. The acceleration application 220 also contains three separate modules that run in parallel, namely, a client module 224, a peer module 226, and an agent module 228, each of which comes into play according to the specific role that the communication device 200 is partaking in the communication network 100 at a given time. The role of each module is further described herein.

The client module 224 provides functionality required when the communication device 200 is requesting information from the Web server 152, such as, for example, but not limited to,

Web pages, data, video, or audio. The client module 224 causes the communication device 200 having the client module 224 therein to intercept the information request and pass the information request on to other elements of the communication network 100, such as, servers, agents or peers. This process is further described in detail herein.

The peer module 226 provides functionality required by the communication device 200 when answering other clients within the communication network 100 and providing the other clients with information that they request, which this communication device 200, having this peer module 226 therein, has already downloaded at a separate time. This process is further described in detail herein.

The agent module 228 provides functionality required when other communication devices of the communication network 100 acting as clients query this communication device 200, having this agent module 228 therein, as an agent, to obtain a list of peers within the communication network 100 that contain requested information. This process is further described in detail herein.

The acceleration application 220 interacts with both the configuration database 280 and the cache database 282 of the storage device 208. As previously mentioned herein, the configuration database 280 stores configuration data that may be common to all communication devices of the communication network 100 and is used to provide setup and synchronization information to different modules 222, 224, 226, 228 of the acceleration application 220 stored within the memory 210.

The cache database 282 stores responses to information requests, such as, for example, HTTP requests, that the communication device 200 has dispatched, either for its own consumption or on behalf of other elements of the communication network 100. The responses to HTTP requests are stored within the cache database 282 for future use by this communication device 200, or for other communication devices within the communication network 100 that

need to retrieve this same information and will use this communication device 200 as either a peer or an agent. This process is described in detail herein.

Information stored within the cache database 282 may include any information associated with a request sent by the client. As an example, such information may include, metadata and actual requested data. For example, for an HTTP request for a video, the metadata may include the version of the Web server answering the request from the client and the data would be the requested video itself. In a situation where there is no more room for storage in the cache database, the software of the associated communication device may cause the communication device to erase previous data stored in order to clear room for the new data to store in the cache database. As an example, such previous data may include data that is most likely not to be used again. Such data may be old data or data that is known to no longer be valid. The communication device may choose to erase the least relevant data, according to any of several methods that are well known in the art.

FIG. 7 is a chart further illustrating two of the main databases utilized within the communication network 100, namely, the acceleration server database 164 and the cache database 282. As previously mentioned, the acceleration server database 164 stores IP addresses of communication devices located within the communication network 100, which have acceleration software stored therein. Specifically, the acceleration server database 164 contains stored therein a list of communication devices having acceleration software stored therein that are currently online within the communication network 100. The acceleration server assigns a list of IP addresses to each communication device functioning as an agent. Each communication device will be the agent for any Web servers whose IP address is in the range 'owned' by that communication device. As an example, when a first ever communication device goes online, namely, the first communication device as described herein having the acceleration application 220 therein, the acceleration server assigns all IP addresses in the world to this communication device, and this communication device will be the agent for any Web server. When a second

communication device goes online it will share the IP address list with the first communication device, so that each of the communication devices will be responsible for a different part of the world wide web servers.

The cache database 282 of the communication device 200 has stored therein a list of URLs 286 of which the communication device 200 is aware. The communication device 200 becomes aware of a URL each time that the communication device 200 receives a request for information located at a specific URL. As shown by FIG. 7, for each URL 288 within the list of URLs 286, the cache database 282 stores: the URL itself 290; HTTP headers 292 returned by the Web Server 152 for this URL; when the last time 294 was that the contents of this URL were loaded directly from the Web Server 152; when the contents of the URL last changed 296 on the Web Server 152; and a list of chunks 298 that contain the contents of this URL, and the content of the chunk. As previously mentioned, chunks, in the present description, are defined as equally sized pieces of data that together form the entire content of the URL, namely, the entire content whose location is described by the URL. As a non-limiting example, a chunk size of, for example, 16KB can be used, so that any HTTP response will be split up into chunks of 16KB. In accordance with an alternative embodiment of the invention, if the last chunk of the response is not large enough to fill the designated chunk size, such as 16KB for the present example, the remaining portion of the chunk will be left empty.

For each such chunk 300, the cache database 282 includes the checksum of the chunk 302, the data of the chunk 304 itself, and a list of peers 306 that most likely have the data for this chunk. As is described in additional detail herein, the data for the chunk may be used by other clients within the communication network 100 when other communication devices of the communication network 100 serve as peers to the clients, from which to download the chunk data.

For each chunk, a checksum is calculated and stored along side of the chunk itself. The checksum may be calculated in any of numerous ways known to those in the art. The purpose of having the checksum is to be able to identify data uniquely, whereas the checksum is the “key” to the data, where the data is the chunk. As an example, a client may want to load the contents of a URL, resulting in the agent that is servicing this request sending the checksums of the chunks to the client, along with the peers that store these chunks. It is to be noted that there could be a different peer for every different chunk. The client then communicates with each such peer, and provides the checksum of the chunk that it would like the peer to transmit back to the client. The peer looks up the checksum (the key) in its cache database, and provides back the chunk (data) that corresponds to this checksum (the key). As shown by FIG. 7, for each peer 308 within the list of peers 306, the cache database 282 includes the peer IP address 310, as well as the connection status 312 of the peer, which represents whether the peer 308 is online or not.

In accordance with one embodiment of the invention, the cache database 282 may be indexed by URL and by Checksum. Having the cache database indexed in this manner is beneficial due to the following reason. When the agent is using the cache database, the agent receives a request from a client for the URL that the client is looking for. In such a case the agent needs the cache database to be indexed by the URL, to assist in finding a list of corresponding peers that have the chunks of this URL. When the peers are using this cache database, the peers obtain a request from the client for a particular checksum, and the peers need the database to be indexed by the checksum so that they can quickly find the correct chunk. Of course, as would be understood by one having ordinary skill in the art, the cache database may instead be indexed in any other manner.

Having described components of the communication network 100, the following further describes how such components interact and individually function. FIG. 8 is a flowchart 300 illustrating operation of the acceleration system initializer module 222 (hereafter referred to as the initializer 222 for purposes of brevity). It should be noted that any process descriptions or

blocks in flowcharts should be understood as representing modules, segments, portions of code, or steps that include one or more instructions for implementing specific logical functions in the process, and alternative implementations are included within the scope of the present invention in which functions may be executed out of order from that shown or discussed, including substantially concurrently or in reverse order, depending on the functionality involved, as would be understood by those reasonably skilled in the art of the present invention.

The initializer 222 is the first element of the communication device 200 to operate as the communication device 200 starts up (block 302). As the initializer 222 starts, it first communicates with the acceleration server 162 to sign up with the acceleration server 162. This is performed by providing the acceleration server 162 with the hostname, and all IP addresses and media access control (MAC) addresses of the interfaces on the communication device 200 having the initializer 222 thereon.

In accordance with an alternative embodiment of the invention, as shown by block 304, the initializer 222 checks with the acceleration server 162 whether a more updated version of the acceleration application software is available. This may be performed by any one of many known methods, such as, but not limited to, by providing the version number of the acceleration application software to the acceleration server 162. The message received back from the acceleration server 162 indicates whether there is a newer version of the acceleration application software or not. If a newer version of the acceleration application software exists, the initializer 222 downloads the latest version of the acceleration application software from the acceleration server 162, or from a different location, and installs the latest version on the communication device 200. In addition to the abovementioned, the initializer 222 may also schedule additional version checks for every set period of time thereafter. As an example, the initializer 222 may check for system updates every two days.

As shown by block 306, the initializer 222 then redirects outgoing network traffic from the communication device 200 to flow through the acceleration application 162. As previously mentioned, one way to redirect the outgoing network traffic is to insert an intermediate driver 212 that intercepts and redirects the traffic. It should be noted that there are many other ways to implement this redirection, which are well known to those having ordinary skill in the art.

As shown by block 308, the initializer 222 then launches the client module 224 of the communication device 200, and configures the client module 224 of the communication device 200 to intercept to all outgoing network communications of the communication device 200 and route the outgoing network communications to the client module 224, from the intermediate driver 272 or other routing method implemented. This is performed so that the client module 224 is able to receive all network traffic coming from the network applications, modify the network traffic if necessary, and re-route the traffic. As is known by those having ordinary skill in the art, in order to re-route the traffic, the traffic needs to be modified, as an example, to change the destination of requests.

As shown by block 310, the initializer 222 then launches the agent module 228 and the peer module 226 to run on the communication device 200. The agent module 228 and peer module 226 listen on pre-determined ports of the communication device 200, so that incoming network traffic on these ports gets routed to the agent module 228 and peer module 226. As is explained in further detail herein, the abovementioned enables the communication device 200 to function as an agent and as a peer for other communication devices within the communication network 100, as needed.

FIG. 9 is a flowchart 350 further illustrating communication between different elements of the communication network 100, in accordance with the present system and method for providing faster and more efficient data communication.

As shown by block 352, an application running on the client 200 initiates a request for a resource on a network. Such a request may be, for example, “GET http://www.aol.com/index.html HTTP/1.1”. The request may come from an Internet browser 214 located on the client 200, where the Internet browser 214 is loading a page from the Internet, an application that wants to download information from the Internet, fetch or send email, or any other network communication request.

Through the intermediate driver 272, or other such mechanism as may be implemented that is re-routing the communication to the client module 224 of the client 200, the resource request is intercepted by the client module 224 that is running on the client 200 (block 354). The client module 224 then looks up the IP address of the server 152 that is the target of the resource request (e.g., the IP address of the Web server that is the host of www.aol.com in the example above), and sends this IP address to the acceleration server 162 (block 356) in order to obtain a list of communication devices that the client 200 can use as agents (hereafter referred to as agents). It should be noted that the process of performing an IP lookup for a server is known by one having ordinary skill in the art, and therefore is not described further herein.

In response to receiving the IP address of the server 152, the acceleration server 162 prepares a list of agents that may be suitable to handle the request from this IP address (block 358). The size of the list can differ based on implementation. For exemplary purposes, the following provides an example where a list of five agents is prepared by the acceleration server 162. The list of agents is created by the acceleration server 162 by finding the communication devices of the communication network 100 that are currently online, and whose IP address is numerically close to the IP of the destination Web server 152. A further description of the abovementioned process is described here in.

As shown by block 360, the client module 224 then sends the original request (e.g., “GET http://www.aol.com/index.html HTTP/1.1”) to all the agents in the list received from the

acceleration server 162 in order to find out which of the agents in the list is best suited to be the one agent that will assist with this request.

It should be noted that, in accordance with an alternative embodiment of the invention, the communication device 200 may be connected to a device that is actually requesting data. In such an alternative embodiment, the communication device would be a modular device connected to a requesting device, where the requesting device, such as, for example, a personal data assistant (PDA) or other device, would request data, and the communication device connected thereto, either through a physical connection, wireless connection, or any other connection, would receive the data request and function as described herein. In addition, as previously mentioned, it should be noted that the HTTP request may be replaced by any request for resources on the Web.

FIG. 10 is a flowchart continuing the flowchart 380 of FIG. 9 and focused on agent response to the request. As shown by block 382, upon receiving the request from the client 200, each agent that received the request from the client responds to the client 200 with whether it has information regarding the request, which can help the client to download the requested information from peers in the network. Specifically, each agent responds with whether the agent has seen a previous request for this resource that has been fulfilled. In such a case, the agent may then provide the client with the list of peers and checksums of the chunks that each of them have.

As shown by block 384, the client then decides which of the agents in the list to use as its agent for this particular information request. To determine which agent in the list to use as its agent for the particular information request, the client may consider multiple factors, such as, for example, factoring the speed of the reply by each agent and whether that agent does or does not have the information required. There are multiple ways to implement this agent selection, one practical way being to start a timer of a small window of time, such as, for example, 5ms, after receiving the first response from the agents, and after the small window, choosing from the list

of agents that responded, the agent that has the information about the request, or in the case that none of the agents responded, to choose the first agent from the list received from the acceleration server 162.

As shown by block 386, after selecting an agent, the client notifies the selected agent that it is going to use it for this request, and notifies the other agents that they will not be used for this request. The client then sends the selected agent a request for the first five chunks of data of the original information request (block 388). By specifying to the selected agent the requested chunks by their order in the full response, the client receives the peer list and checksums of the requested chunks from the selected agent. As an example, for the first five chunks the client will ask the selected agent for chunks one through five, and for the fourth batch of five chunks the client will ask the agent for chunks sixteen through twenty. As previously mentioned, additional or fewer chunks may be requested at a single time.

As shown by block 390, after receiving the request from the client, the selected agent determines whether it has information regarding the requested chunks of data by looking up the request in its cache database and determining if the selected agent has stored therein information regarding peers of the communication network that have stored the requested data of the request, or whether the selected agent itself has the requested data of the request stored in its memory. In addition to determining if the selected agent contains an entry for this request in its database, the selected agent may also determine if this information is still valid. Specifically, the selected agent determines whether the data that is stored within the memory of the selected agent or the memory of the peers, still mirrors the information that would have been received from the server itself for this request. A further description of the process utilized by the selected agent to determine if the information is still valid, is described in detail herein.

As shown by block 392, if the information (requested data of the request) exists and is still valid, then the agent prepares a response to the client, which includes for each of the chunks:

(i) the checksum of the chunk; (ii) a list of peers that according to the database of the selected agent contains these chunks; and (iii) if these are the first five chunks of the information, then the selected agent also provides the specific protocol's headers that would have been received from the server, had the initial request from the client been made directly to the server.

As shown by block 394, the list of peers for each chunk is sorted by geographical proximity to the requesting client. In accordance with the present example, only the five closest peers are kept in the list for every chunk, and the rest of the peers are discarded from this list. As shown by block 396, the prepared response, namely, the list of closest peers, is sent back to the client. It should be noted that, if this were the last set of chunks to be provided for this request, then it would be beneficial to include information about this to the client.

If the selected agent discovers that it does not have information about this request, or if the selected agent discovers that the information it has is no longer valid, the selected agent needs to load the information directly from the server in order to be able to provide an answer to the requesting client. As shown by block 400, the selected agent then sends the request directly to the server. The selected agent then stores the information it receives from the server (both the headers of the request, as well as chunks of the response itself) in its database, for this particular response to the client, as well as for future use to other clients that may request this data (block 402). The selected agent then prepares a response (list) for the client, where the response includes the protocol headers (if these are the first five chunks), and the checksums of the five chunks, and provides itself as the only peer for these chunks (block 404). This list is then sent back to the client (block 406).

FIG. 11 is a flowchart 420 continuing the flowchart of FIG. 10, which illustrates actions taken upon receipt of the list of peers, or single peer listing, from the agent. As shown by block 422, the client receives the response from the agent (including the list of chunks and their corresponding data, including peers and other information previously mentioned) and, for each

of the five chunks, the client sends a request to each of the peers listed for the chunk to download the chunk. The chunk request that the client sends to each of the peers is the checksum of the data that the client seeks to receive, which is the key (identifier) of the chunk.

As shown by block 424, the peers then respond regarding whether they still have the data of the chunk. As an example, some of the peers may not currently be online, some may be online but may have discarded the relevant information, and some may still have the relevant information, namely, the chunk. As shown by block 426, the client then selects the quickest peer that responds with a positive answer regarding the requested information, the client lets that peer know that it is chosen to provide the client with the chunk, and the client notifies the other peers that they are not chosen.

As shown by block 428, the chosen peer then sends the chunk to the client. It should be noted that if no peers answer the request of the client, the client goes back to the agent noting that the peers were all negative, and the agent either provides a list of 5 other agents, if they exist, or the agent goes on to download the information directly from the Web server as happens in the case where no peers exist as described above.

The client then stores the chunks in its cache for future use (block 430), when the client may need to provide the chunks to a requesting communication device when acting as a peer for another client that is looking for the same information. As shown by block 432, if some of the chunks were not loaded from any of the peers, the client requests the chunks again from the agent in a next round of requests, flagging these chunks as chunks that were not loadable from the client list of peers. In this situation, the agent will load the data directly from the server and provide it back to the client.

The client then acknowledges to the agent which of the chunks it received properly (block 434). The agent then looks up these chunks in the database of the agent, and adds the

client to the list of peers for these chunks, specifically, since this client is now storing these chunks, and can provide these chunks to other clients that turn to it as a peer (block 436).

As shown by block 438, the client then passes the data on to the Web browser or other application of the client that made the original request, for it to use as it had originally intended. The client then checks whether all of the chunks for this request were received (block 440), by checking the flag set by the agent. Specifically, when the agent is providing the list of the last 5 chunks, the agent includes that information as part of its reply to the client, which is referred to herein as a flag. This information is what enables the client to know that all information has been received for a particular resource request.

If the last received chunks were not the last chunks for this request, the processing flow of the client continues by returning to the functionality of block 384 of FIG. 10, but instead sending the chosen agent a request for the next five chunks of data of the original information request. Alternatively, if all chunks for this request were received, the request is complete, and the flow starts again at block 352 of FIG. 9.

FIG. 12 is a flowchart 500 illustrating steps taken by an agent, client, or peer to determine whether a certain HTTP request is still valid. Specifically, the following provides an example of how the agent, client, or peer can determine whether particular data that is stored within the memory of the agent, or the memory of a peer or client, still mirrors the information that is currently on the Web server. As shown by block 502, the HTTP request is looked up in the cache database of the agent, client or peer that is checking the validity of the HTTP request. As an example, the HTTP protocol, defined by RFC 2616, outlines specific methods that Web servers can define within the HTTP headers signifying the validity of certain data, such as, but not limited to, by using HTTP header information such as “max age” to indicate how long this data may be cached before becoming invalid, “no cache” to indicate that the data may never be cached, and using other information.

As shown by block 504, these standard methods of validation are tested on the HTTP request information in question. As shown by block 506, a determination is made whether the requested information that is stored is valid or not. If the requested information is valid, a "VALID" response is returned (block 508). Alternatively, if the requested information is not valid, an HTTP conditional request is sent to the relevant Web server, to determine if the data stored for this request is still valid (block 510). If the data stored for this request is still valid, a "VALID" response is returned (block 508). Alternatively, if the data stored for this request is not valid, an "INVALID" response is returned (block 514). It should be noted, that the abovementioned description with regard to FIG. 12 is an explanation of how to check if HTTP information is still valid. There are similar methods of determining validity for any other protocol, which may be utilized, and which those having ordinary skill in the art would appreciate and understand.

FIG. 13 is a flowchart 550 outlining operation of the acceleration server, whose main responsibility in the present system and method is to provide clients with information regarding which agents serve which requests, and to keep the network elements all up to date with the latest software updates. As shown by block 552, the acceleration server sends "keep alive" signals to the network elements, and keeps track within its database as to which network elements are online. As shown by block 554, the acceleration server continues to wait for a client request and continues to determine if one is received.

Once a request is received, the acceleration server tests the type of request received (block 556). If the client request is to sign up the client within the network, an event that happens every time that the client starts running on its host machine, then that client is added to the list of agents stored on the acceleration server, sorted by the IP address of the client (block 558).

If the request is to find an agent to use for a particular request, the acceleration server creates a new agent list, which is empty (block 560). The acceleration server then searches the

agent database for the next 5 active agents whose IP address is closest to the IP address of the server who is targeted in the request (block 562). In this context, 192.166.3.103 is closer to 192.166.3.212 than to 192.167.3.104. The acceleration server then sends this agent list to the client (block 564).

If instead, the request is to check the version of the latest acceleration software then the acceleration server sends that network element (client, peer or agent) the version number of the latest existing acceleration software version, and a URL from where to download the new version, for the case that the element needs to upgrade to the new version (block 566).

While the abovementioned example is focused on HTTP requests for data, as previously mentioned, other protocol requests are equally capable of being handled by the present system and method. As an example, in separate embodiments the acceleration method described may accelerate any communication protocol at any OSI layer (SMTP, DNS, UDP, ETHERNET, etc.). In the following alternative embodiment, it is illustrated how the acceleration method may accelerate TCPIP. As is known by those having ordinary skill in the art, TCPIP is a relatively low-level protocol, as opposed to HTTP, which is a high level protocol. For purposes of illustration of TCPIP communication, reference may be made to FIG. 3, wherein the Web server is a TCPIP server.

In TCPIP there are three communication commands that are of particular interest, namely, connect, write, and read. Connect is a command issued by an application in the communication device that is initiating the communication to instruct the TCPIP stack to connect to a remote communication device. The connect message includes the IP address of the communication device, and the port number to connect to. An application uses the write command to instruct the TCPIP stack to send a message (i.e., data) to a communication device to which it is connected. In addition, an application uses the read command to ask the TCPIP stack to provide the message that was sent from the remote communication device to which it is

connected. A communication session typically exists of a connect, followed by a read and write on both sides.

FIG. 14 is a flowchart 600 further illustrating TCPIP acceleration in accordance with this alternative embodiment of the invention. As shown by blocks 601 and 602 when an application of the communication device makes a request to the communications stack to connect with the TCPIP server, that communication is intercepted by the acceleration application.

To find an agent, upon receiving that connect message from the communication device application, which includes the IP address of the TCPIP server and the port to connect to, the acceleration application in the client makes a request to the acceleration server to find out who the agent for the communication with the TCPIP server is. This step is performed in a similar manner to that described with regard to the main HTTP embodiment of the invention (block 604). As shown by block 606, the server then provides the client with a list of agents, for example, a primary agent and four others.

To establish a connection, as shown by block 608, the client issues a TCPIP connect with the primary agent or one of the other agents if the primary agent does not succeed, to create a connection with the agent. The client then sends to the agent the IP address of the TCPIP server and connection port that were provided by the communication device application (block 610). As shown by block 612, that agent in turn issues a TCPIP connect to the TCPIP server to the port it received from the client, to create a connection with the agent.

FIG. 15 is a flowchart 800 further illustrating TCPIP acceleration in accordance with this alternative embodiment of the invention, detailing the communication between the client and the TCPIP server (read and write commands) after the connect phase has completed successfully.

As shown by block 802, if the network application within the client wants to send a message to the TCPIP server, the network application within the client writes the message to the TCPIP stack in the operating system of the client. This WRITE command is received by the acceleration application of the client and handled in the manner described below. If the TCPIP server wants to send a message to the client, the TCPIP server writes the message to the TCPIP stack of TCPIP operating system, on the connection to the agent, since this agent is where the server received the original connection. This WRITE command is received by the acceleration application of the agent and handled in the manner described below.

When the acceleration application of the client receives a message from the network application of the client to be sent to the agent, or when the acceleration application of the agent receives a message from the connection to the TCPIP server that is to be sent to the client, the acceleration application proceeds to send the message to the communication device on the other side. For instance, if the client has intercepted the message from the communication application, the client sends the message to the agent, and if it is the agent that intercepted the message from the connection to the TCPIP server, such as the TCPIP server sending a message that is intended for the communication with client, the agent sends the message to the client in the following manner:

As shown by block 804, the acceleration application breaks up the content of the message to chunks and calculates the corresponding checksums, in the same manner as in the main embodiment described herein. The acceleration application then looks up each checksum in its cache database (block 806). As shown by block 808, the acceleration application checks if the checksum exists in the cache database. If it does, then, as shown by block 810, the acceleration

application prepares a list of peers that have already received the chunk of the checksum in the past (if any), and adds the communication device of the other side to the list of communication devices that have received this chunk (adds it to the peer list of the checksum in its database), to be provided to other communication devices requesting this information in the future. As shown by block 812, the list of peers is sent to the receiving communication device, which, as shown by block 814 retrieves the chunks from the peers in the list received, in the same manner as in the main embodiment.

If the checksum does not exist within the cache database of the sending communication device then, as shown by block 820, the acceleration application adds the checksum and chunk to its cache database, sends the chunk to the communication device on the other side, and adds the other communication device to the list of peers for that checksum in its database.

As shown by block 816, a determination is then made as to whether all chunks have been received. If all chunks have not been received, the process continues on again from block 806.

Once all data has been received, as shown by block 818, the acceleration application passes the data on to the requester. Specifically, in the client, the acceleration application passes on the complete data to the communication application, and in the agent, the acceleration application passes on the complete data to the requesting TCP/IP server.

It should be emphasized that the above-described embodiments of the present invention are merely possible examples of implementations, merely set forth for a clear understanding of the principles of the invention. Many variations and modifications may be made to the above-described embodiments of the invention without departing substantially from the spirit and principles of the invention. All such modifications and variations are intended to be included

Attorney Docket No. 19459-6105P

herein within the scope of this disclosure and the present invention and protected by the following claims.

Claims

1. A method for use with a web server that responds to Hypertext Transfer Protocol (HTTP) requests and stores a first content identified by a first content identifier, the method by a first client device comprising:

 establishing a Transmission Control Protocol (TCP) connection with a second server;

 sending, to the web server over the Internet, the first content identifier;

 receiving, the first content from the web server over the Internet in response to the sending of the first content identifier; and

 sending the received first content, to the second server over the established TCP connection, in response to the receiving of the first content identifier.

2. The method according to claim 1, further comprising receiving, by the first client device from the second server over the established TCP connection, the first content identifier.

3. The method according to claim 1, wherein the sending of the first content identifier to the web server over the Internet comprises sending a Hypertext Transfer Protocol (HTTP) request that comprises the first content identifier.

4. The method according to claim 1, further comprising storing, by the first client device in response to the receiving from the web server, the first content.

5. The method according to claim 1, wherein the second server is a Transmission Control Protocol/Internet Protocol (TCP/IP) server that communicates over the Internet based on, or according to, using TCP/IP protocol or connection, and wherein the first client device is a Transmission Control Protocol/Internet Protocol (TCP/IP) client that communicates

with the second server over the Internet based on, or according to, TCP/IP protocol.

6. The method according to claim 1, wherein the first client device communicates over the Internet based on, or according to, one out of UDP, DNS, TCP, FTP, POP#, SMTP, or SQL standards.

7. The method according to claim 1, wherein the first content comprises web-page, audio, or video content, and wherein the first content identifier comprises a Uniform Resource Locator (URL).

8. The method according to claim 1, further comprising executing, by the first client device, a web browser application or an email application.

9. The method according to claim 1, for use with a third server that comprises a web server that is Hypertext Transfer Protocol (HTTP) server, the third server responds to HTTP requests and stores a second content identified by a second content identifier, the method by the first client device further comprising:

receiving the second content identifier;

sending, to the third server over the Internet in response to the receiving, the second content identifier; and

receiving the second content from the third server over the Internet in response to the sending.

10. The method according to claim 9, further comprising executing, by the first client device, a web browser application or an email application.

11. The method according to claim 1, further comprising periodically communicating over the TCP connection between the second server and the first client device.

12. The method according to claim 11, wherein the periodically communicating comprises exchanging 'keep alive' messages.

13. The method according to claim 1, wherein the first client device is identified by a Media Access Control (MAC) address or a hostname, and wherein the method further comprising sending, by the first client device, during, as part of, or in response to, a start-up or power-up of the first client device, a first message to the second server, and wherein the first messages comprises the first client IP address, the MAC address, or the hostname.

14. The method according to claim 13, for use with a first application stored in the first client device and associated with a first version number, wherein the first message comprises the first version number.

15. The method according to claim 14, for use with a second application that is a version of the first application, is stored in the second server, and is associated with a second version number, wherein the method further comprising receiving, by the first client device from the second server, in response to the first message, a second message that comprises the second version number.

16. The method according to claim 15, wherein the method further comprising downloading over the Internet, by the first client device from the second server, in response to the first message, the second application from the second server, and installing the second application in the first client device as a replacement for the first application.

17. The method according to claim 1, further comprising determining, by the first client device, that the received first content, is valid.

18. The method according to claim 17, wherein the determining is based on the received HTTP header according to, or based on, IETF RFC 2616.

19. The method according to claim 17, further comprising:

sending, a message over the Internet in response to the determining that the received first content, is not valid; and

receiving, over the Internet in response to the sending of the message, from the second server or from a second client device selected from a plurality of client devices, the first content.

20. The method according to claim 1, further comprising storing, operating, or using, a client operating system.

21. The method according to claim 1, wherein the steps are sequentially executed.

22. The method according to claim 1, for use with a software application that includes computer instructions that, when executed by a computer processor, cause the processor to perform the sending of the Hypertext Transfer Protocol (HTTP) request, the receiving and storing of the first content, the receiving of the first content identifier, and the sending of the part of, or the whole of, the stored first content, the method is further preceded by:

downloading, by the first client device from the Internet, the software application; and

installing, by the first client device, the downloaded software application.

23. The method according to claim 22, wherein the software application is downloaded from the second server.

24. A non-transitory computer readable medium containing computer instructions that, when executed by a computer processor, cause the processor to perform the method according to claim 1.

Abstract

A system designed for increasing network communication speed for users, while lowering network congestion for content owners and ISPs. The system employs network elements including an acceleration server, clients, agents, and peers, where communication requests generated by applications are intercepted by the client on the same machine. The IP address of the server in the communication request is transmitted to the acceleration server, which provides a list of agents to use for this IP address. The communication request is sent to the agents. One or more of the agents respond with a list of peers that have previously seen some or all of the content which is the response to this request (after checking whether this data is still valid). The client then downloads the data from these peers in parts and in parallel, thereby speeding up the Web transfer, releasing congestion from the Web by fetching the information from multiple sources, and relieving traffic from Web servers by offloading the data transfers from them to nearby peers.




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6

ACCELERATION DATABASE 164			
166	AGENT IP A ONLINE/OFFLINE		
>>> INDEXED BY: AGENT IP ADDRESS			
CACHE DATABASE 282			
286	LIST OF URLS:		
288	URL 1		
	290 URL		
	292 URL HTTP HEADERS		
	294 LAST CHECKED ON SERVER		
	296 LAST CHANGED ON SERVER		
	298 LIST OF CHUNKS FOR THIS URL:		
	300 CHUNK 1		
		302	CHUNK CHECKSUM
		304	CHUNK DATA
		306	LIST OF PEERS:
		308	PEER 1
		310	PEER 1 IP ADDRESS
		312	PEER 2 CONNECTION STATUS




FIG. 7

300




FIG. 8

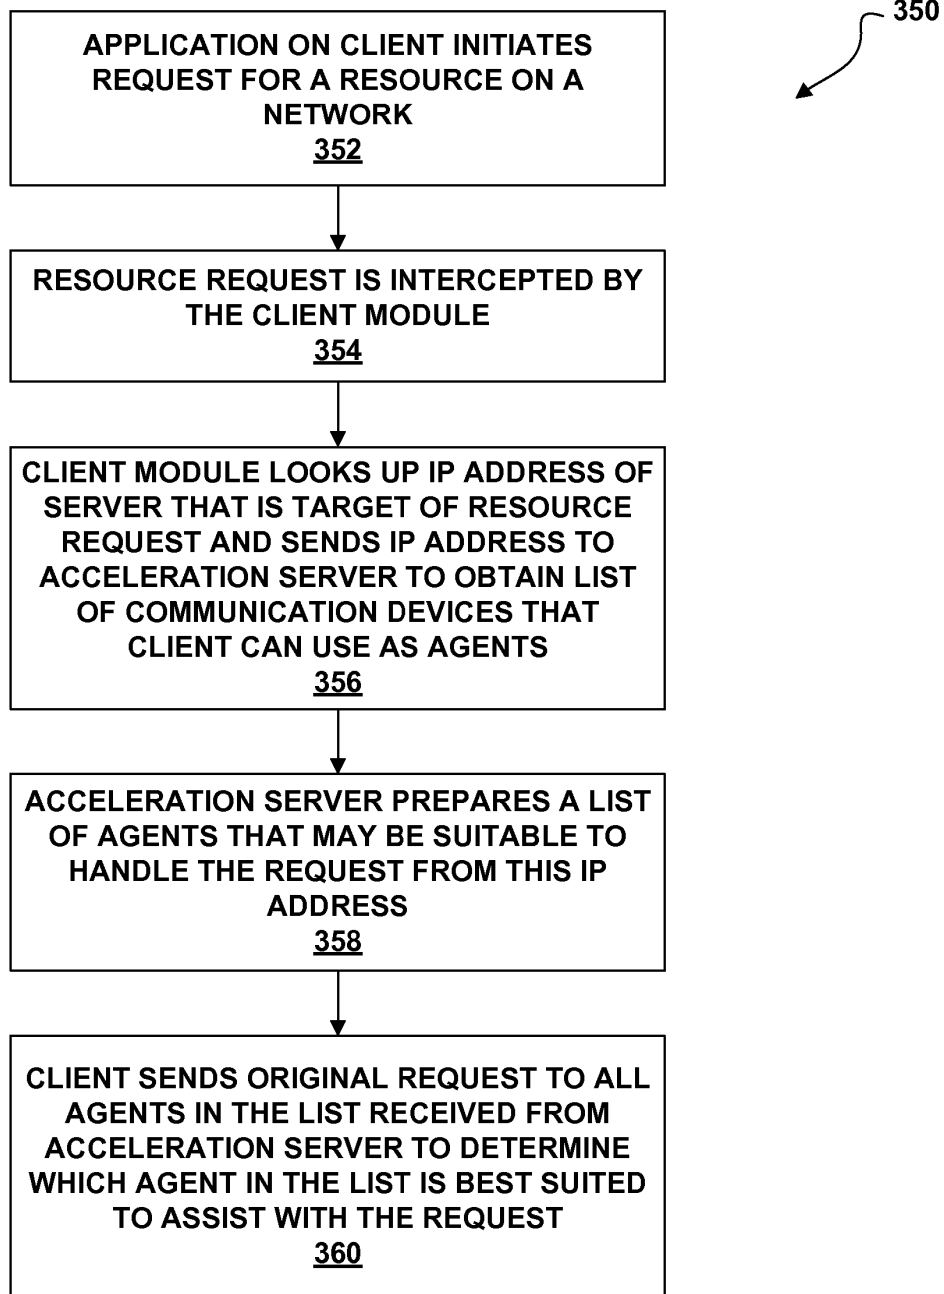


FIG. 9

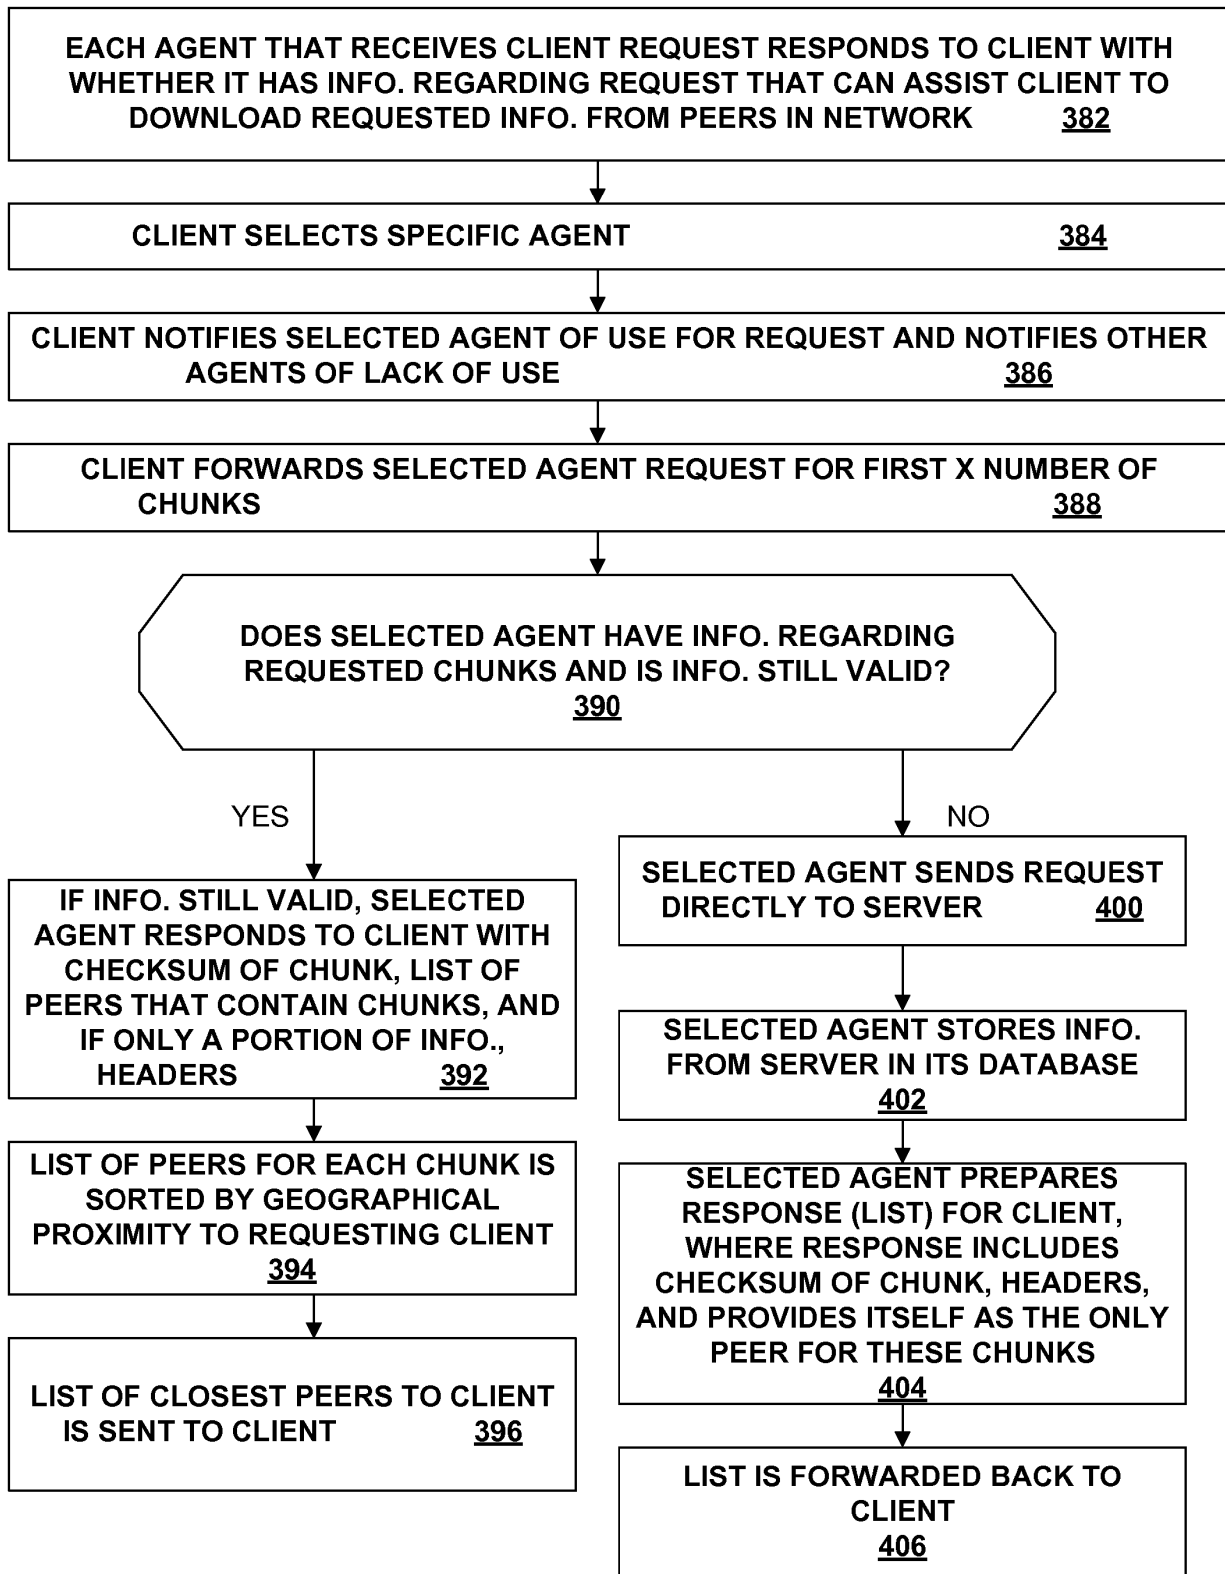


FIG. 10




FIG. 11

↖
420




FIG. 12




FIG 13

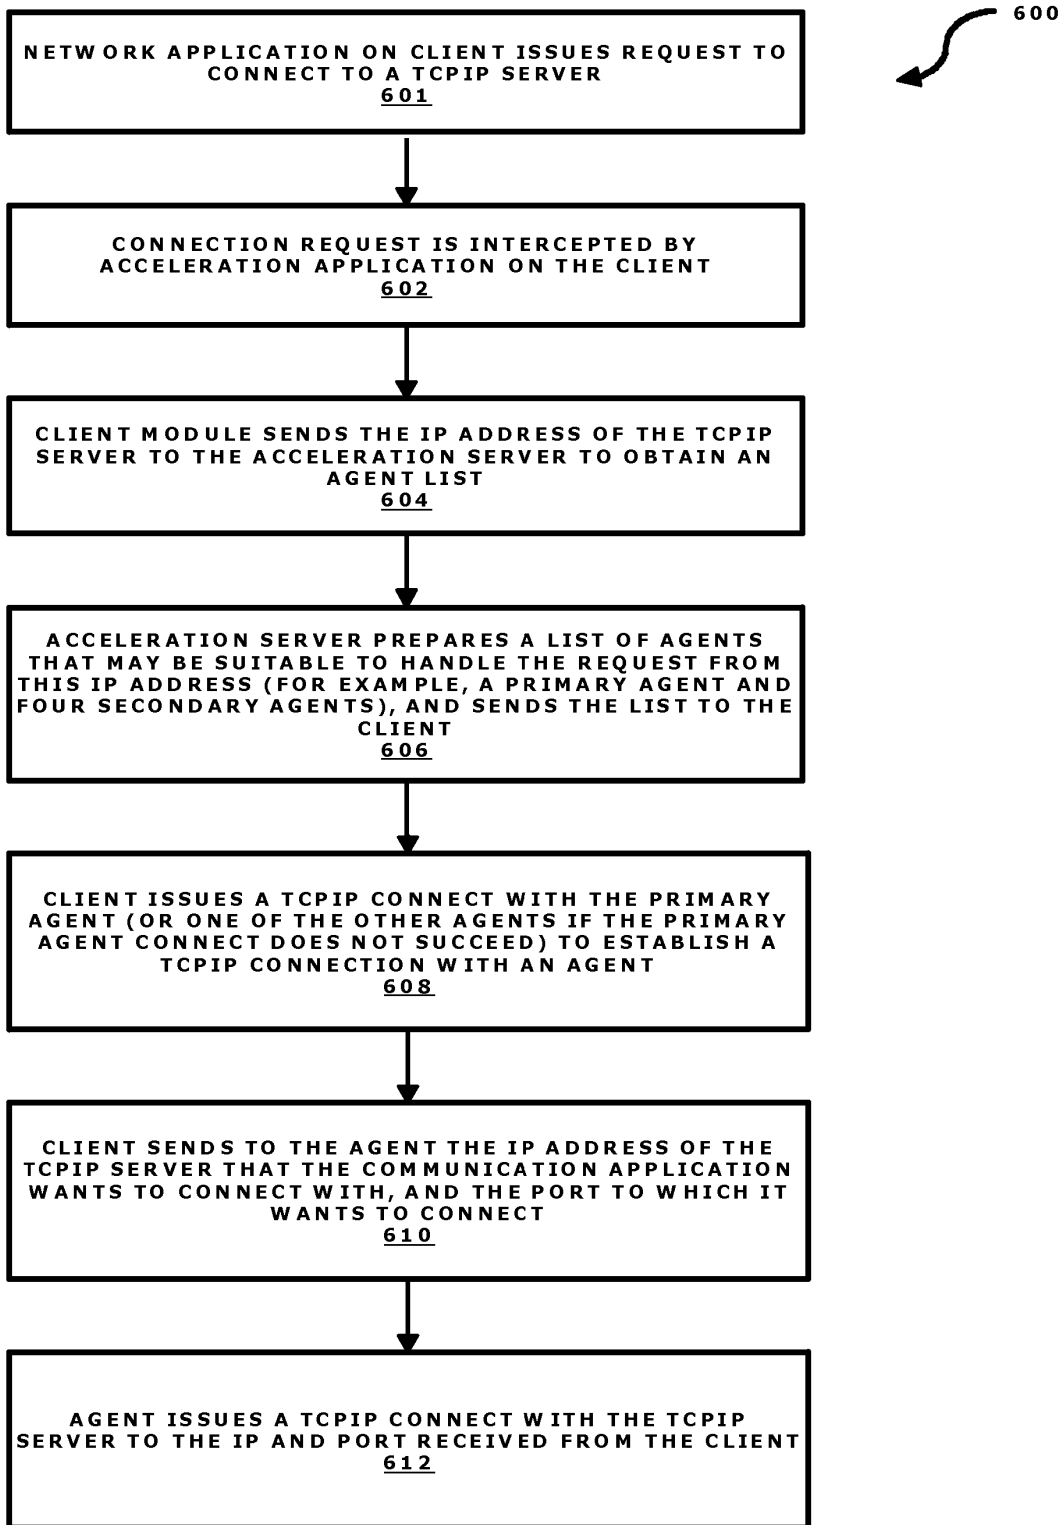


FIG. 14

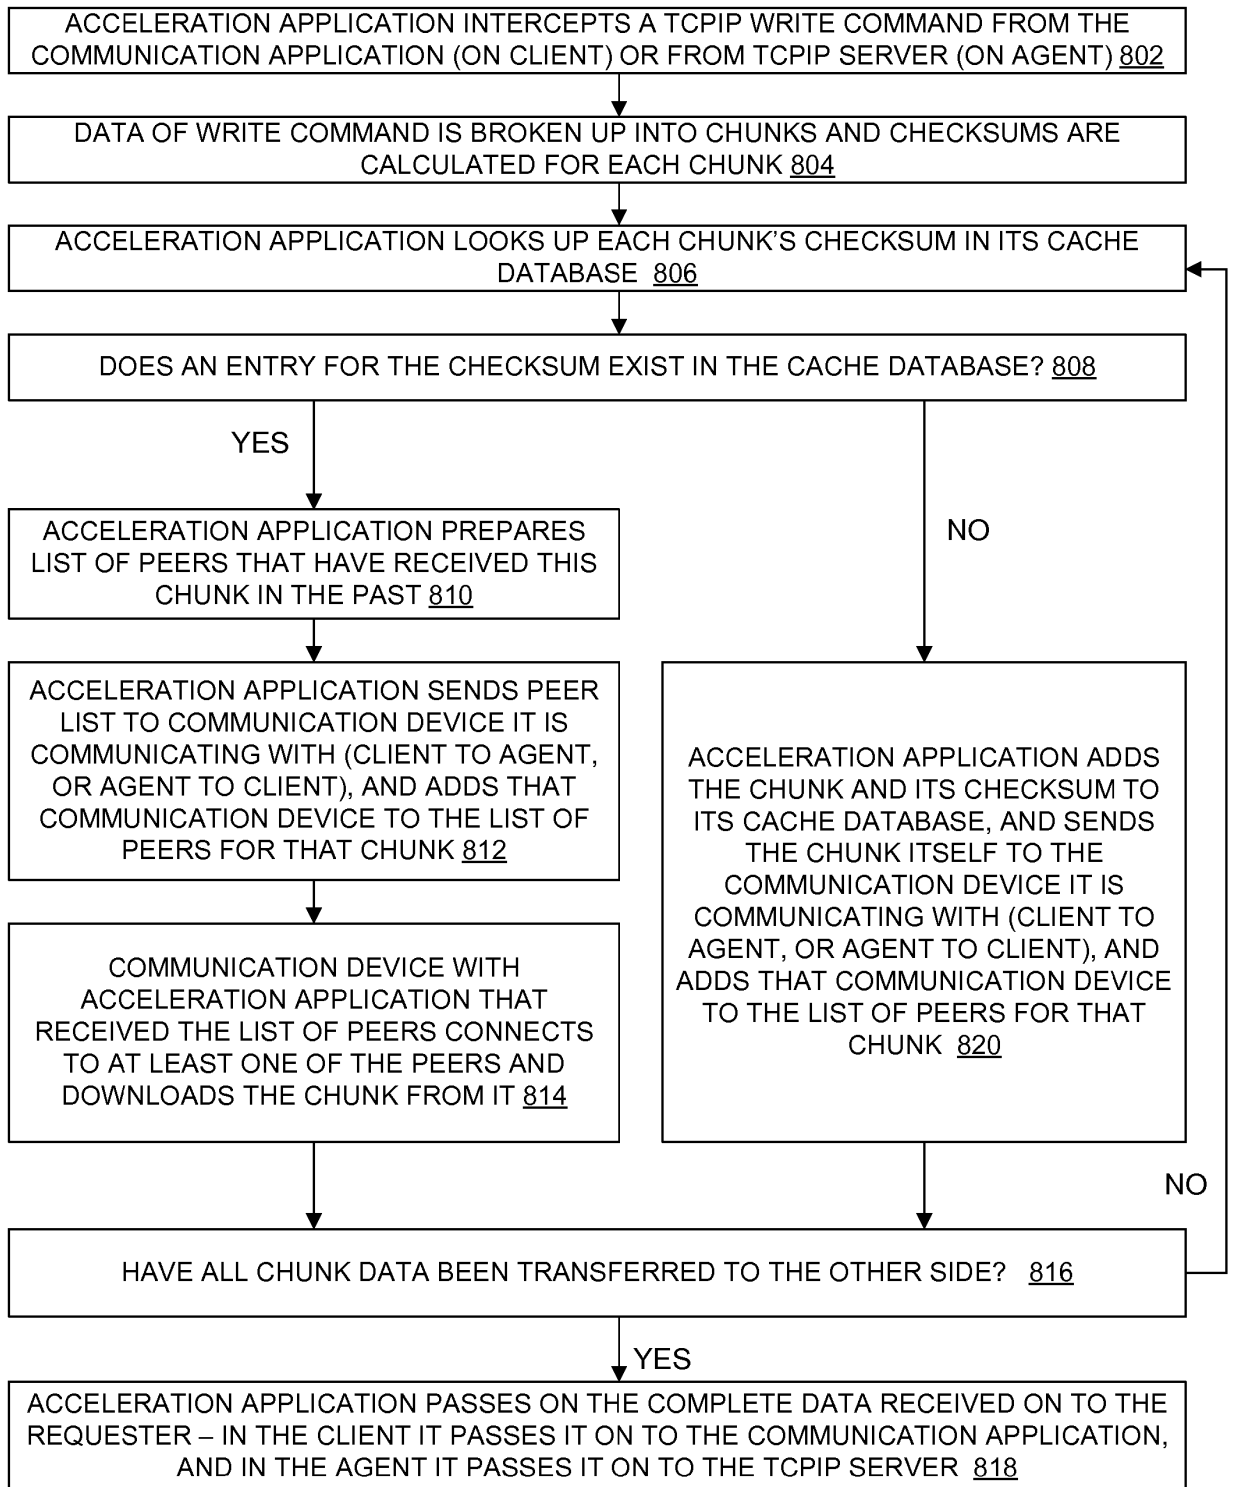


FIG. 15

800

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76		Attorney Docket Number	HOLA-005-US10
		Application Number	
Title of Invention	SYSTEM PROVIDING FASTER AND MORE EFFICIENT DATA COMMUNICATION		
The application data sheet is part of the provisional or nonprovisional application for which it is being submitted. The following form contains the bibliographic data arranged in a format specified by the United States Patent and Trademark Office as outlined in 37 CFR 1.76. This document may be completed electronically and submitted to the Office in electronic format using the Electronic Filing System (EFS) or the document may be printed and included in a paper filed application.			

Secrecy Order 37 CFR 5.2:

Portions or all of the application associated with this Application Data Sheet may fall under a Secrecy Order pursuant to 37 CFR 5.2 (Paper filers only. Applications that fall under Secrecy Order may not be filed electronically.)

Inventor Information:

Inventor	1				Remove
Legal Name					
Prefix	Given Name	Middle Name	Family Name	Suffix	
	Derry		Shribman		
Residence Information (Select One) US Residency <input type="radio"/> Non US Residency Active US Military Service					
City	Tel Aviv		Country of Residence ⁱ	IL	
Mailing Address of Inventor:					
Address 1	9/6 Beylinson St.,				
Address 2					
City	Tel Aviv		State/Province		
Postal Code	6356709		Country ⁱ	IL	
Inventor	2				Remove
Legal Name					
Prefix	Given Name	Middle Name	Family Name	Suffix	
	Ofer		Vilenski		
Residence Information (Select One) US Residency <input checked="" type="radio"/> Non US Residency Active US Military Service					
City	Moshav Hadar Am		Country of Residence ⁱ	IL	
Mailing Address of Inventor:					
Address 1	8 Hahollandim Street				
Address 2					
City	Moshav Hadar Am		State/Province		
Postal Code	42935		Country ⁱ	IL	
All Inventors Must Be Listed - Additional Inventor Information blocks may be generated within this form by selecting the Add button.					
					Add

Correspondence Information:

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76		Attorney Docket Number	HOLA-005-US10
		Application Number	
Title of Invention	SYSTEM PROVIDING FASTER AND MORE EFFICIENT DATA COMMUNICATION		

Enter either Customer Number or complete the Correspondence Information section below.
For further information see 37 CFR 1.33(a).

An Address is being provided for the correspondence information of this application.

Customer Number	131926		
Email Address		Add Email	Remove Email

Application Information:

Title of the Invention	SYSTEM PROVIDING FASTER AND MORE EFFICIENT DATA COMMUNICATION		
Attorney Docket Number	HOLA-005-US10	Small Entity Status Claimed	<input checked="" type="checkbox"/>
Application Type	Nonprovisional		
Subject Matter	Utility		
Total Number of Drawing Sheets (if any)	15	Suggested Figure for Publication (if any)	

Filing By Reference:

Only complete this section when filing an application by reference under 35 U.S.C. 111(c) and 37 CFR 1.57(a). Do not complete this section if application papers including a specification and any drawings are being filed. Any domestic benefit or foreign priority information must be provided in the appropriate section(s) below (i.e., "Domestic Benefit/National Stage Information" and "Foreign Priority Information").

For the purposes of a filing date under 37 CFR 1.53(b), the description and any drawings of the present application are replaced by this reference to the previously filed application, subject to conditions and requirements of 37 CFR 1.57(a).

Application number of the previously filed application	Filing date (YYYY-MM-DD)	Intellectual Property Authority or Country

Publication Information:

Request Early Publication (Fee required at time of Request 37 CFR 1.219)

Request Not to Publish. I hereby request that the attached application not be published under 35 U.S.C. 122(b) and certify that the invention disclosed in the attached application **has not and will not** be the subject of an application filed in another country, or under a multilateral international agreement, that requires publication at eighteen months after filing.

Representative Information:

Representative information should be provided for all practitioners having a power of attorney in the application. Providing this information in the Application Data Sheet does not constitute a power of attorney in the application (see 37 CFR 1.32). Either enter Customer Number or complete the Representative Name section below. If both sections are completed the customer number will be used for the Representative Information during processing.

Please Select One:	<input checked="" type="radio"/> Customer Number	US Patent Practitioner	<input type="radio"/> Limited Recognition (37 CFR 11.9)
Customer Number	131926		

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76		Attorney Docket Number	HOLA-005-US10
		Application Number	
Title of Invention	SYSTEM PROVIDING FASTER AND MORE EFFICIENT DATA COMMUNICATION		

Domestic Benefit/National Stage Information:

This section allows for the applicant to either claim benefit under 35 U.S.C. 119(e), 120, 121, 365(c), or 386(c) or indicate National Stage entry from a PCT application. Providing benefit claim information in the Application Data Sheet constitutes the specific reference required by 35 U.S.C. 119(e) or 120, and 37 CFR 1.78.

When referring to the current application, please leave the "Application Number" field blank.

Prior Application Status	Pending		<input type="button" value="Remove"/>		
Application Number	Continuity Type		Prior Application Number	Filing or 371(c) Date (YYYY-MM-DD)	
	Continuation of		15/957945	2018-04-20	
Prior Application Status	Patented		<input type="button" value="Remove"/>		
Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)	Patent Number	Issue Date (YYYY-MM-DD)
15/957945	Continuation of	14/025109	2013-09-12	10069936	2018-09-04
Prior Application Status	Patented		<input type="button" value="Remove"/>		
Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)	Patent Number	Issue Date (YYYY-MM-DD)
14/025109	Division of	12/836059	2010-07-14	8560604	2013-10-15
Prior Application Status	Expired		<input type="button" value="Remove"/>		
Application Number	Continuity Type		Prior Application Number	Filing or 371(c) Date (YYYY-MM-DD)	
12/836059	Claims benefit of provisional		61/249624	2009-10-08	
Additional Domestic Benefit/National Stage Data may be generated within this form by selecting the Add button.					<input type="button" value="Add"/>

Foreign Priority Information:

This section allows for the applicant to claim priority to a foreign application. Providing this information in the application data sheet constitutes the claim for priority as required by 35 U.S.C. 119(b) and 37 CFR 1.55. When priority is claimed to a foreign application that is eligible for retrieval under the priority document exchange program (PDX)ⁱ the information will be used by the Office to automatically attempt retrieval pursuant to 37 CFR 1.55(i)(1) and (2). Under the PDX program, applicant bears the ultimate responsibility for ensuring that a copy of the foreign application is received by the Office from the participating foreign intellectual property office, or a certified copy of the foreign priority application is filed, within the time period specified in 37 CFR 1.55(g)(1).

<input type="button" value="Remove"/>			
Application Number	Country ⁱ	Filing Date (YYYY-MM-DD)	Access Code ⁱ (if applicable)
Additional Foreign Priority Data may be generated within this form by selecting the Add button.			
<input type="button" value="Add"/>			

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76		Attorney Docket Number	HOLA-005-US10
		Application Number	
Title of Invention	SYSTEM PROVIDING FASTER AND MORE EFFICIENT DATA COMMUNICATION		

Statement under 37 CFR 1.55 or 1.78 for AIA (First Inventor to File) Transition Applications

<p>This application (1) claims priority to or the benefit of an application filed before March 16, 2013 and (2) also contains, or contained at any time, a claim to a claimed invention that has an effective filing date on or after March 16, 2013.</p> <p><input type="checkbox"/> NOTE: By providing this statement under 37 CFR 1.55 or 1.78, this application, with a filing date on or after March 16, 2013, will be examined under the first inventor to file provisions of the AIA.</p>
--

Application Data Sheet 37 CFR 1.76		Attorney Docket Number	HOLA-005-US10
		Application Number	
Title of Invention	SYSTEM PROVIDING FASTER AND MORE EFFICIENT DATA COMMUNICATION		

Authorization or Opt-Out of Authorization to Permit Access:

When this Application Data Sheet is properly signed and filed with the application, applicant has provided written authority to permit a participating foreign intellectual property (IP) office access to the instant application-as-filed (see paragraph A in subsection 1 below) and the European Patent Office (EPO) access to any search results from the instant application (see paragraph B in subsection 1 below).

Should applicant choose not to provide an authorization identified in subsection 1 below, applicant **must opt-out** of the authorization by checking the corresponding box A or B or both in subsection 2 below.

NOTE: This section of the Application Data Sheet is **ONLY** reviewed and processed with the **INITIAL** filing of an application. After the initial filing of an application, an Application Data Sheet cannot be used to provide or rescind authorization for access by a foreign IP office(s). Instead, Form PTO/SB/39 or PTO/SB/69 must be used as appropriate.

1. Authorization to Permit Access by a Foreign Intellectual Property Office(s)

A. Priority Document Exchange (PDX) - Unless box A in subsection 2 (opt-out of authorization) is checked, the undersigned hereby **grants the USPTO authority** to provide the European Patent Office (EPO), the Japan Patent Office (JPO), the Korean Intellectual Property Office (KIPO), the State Intellectual Property Office of the People's Republic of China (SIPO), the World Intellectual Property Organization (WIPO), and any other foreign intellectual property office participating with the USPTO in a bilateral or multilateral priority document exchange agreement in which a foreign application claiming priority to the instant patent application is filed, access to: (1) the instant patent application-as-filed and its related bibliographic data, (2) any foreign or domestic application to which priority or benefit is claimed by the instant application and its related bibliographic data, and (3) the date of filing of this Authorization. See 37 CFR 1.14(h)(1).

B. Search Results from U.S. Application to EPO - Unless box B in subsection 2 (opt-out of authorization) is checked, the undersigned hereby **grants the USPTO authority** to provide the EPO access to the bibliographic data and search results from the instant patent application when a European patent application claiming priority to the instant patent application is filed. See 37 CFR 1.14(h)(2).

The applicant is reminded that the EPO's Rule 141(1) EPC (European Patent Convention) requires applicants to submit a copy of search results from the instant application without delay in a European patent application that claims priority to the instant application.

2. Opt-Out of Authorizations to Permit Access by a Foreign Intellectual Property Office(s)

A. Applicant **DOES NOT** authorize the USPTO to permit a participating foreign IP office access to the instant application-as-filed. If this box is checked, the USPTO will not be providing a participating foreign IP office with any documents and information identified in subsection 1A above.

B. Applicant **DOES NOT** authorize the USPTO to transmit to the EPO any search results from the instant patent application. If this box is checked, the USPTO will not be providing the EPO with search results from the instant application.

NOTE: Once the application has published or is otherwise publicly available, the USPTO may provide access to the application in accordance with 37 CFR 1.14.

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76		Attorney Docket Number	HOLA-005-US10
		Application Number	
Title of Invention	SYSTEM PROVIDING FASTER AND MORE EFFICIENT DATA COMMUNICATION		

Applicant Information:

Providing assignment information in this section does not substitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.			
Applicant	1	<input type="button" value="Remove"/>	
If the applicant is the inventor (or the remaining joint inventor or inventors under 37 CFR 1.45), this section should not be completed. The information to be provided in this section is the name and address of the legal representative who is the applicant under 37 CFR 1.43; or the name and address of the assignee, person to whom the inventor is under an obligation to assign the invention, or person who otherwise shows sufficient proprietary interest in the matter who is the applicant under 37 CFR 1.46. If the applicant is an applicant under 37 CFR 1.46 (assignee, person to whom the inventor is obligated to assign, or person who otherwise shows sufficient proprietary interest) together with one or more joint inventors, then the joint inventor or inventors who are also the applicant should be identified in this section.			
<input type="button" value="Clear"/>			
<input checked="" type="radio"/> Assignee	Legal Representative under 35 U.S.C. 117		Joint Inventor
Person to whom the inventor is obligated to assign.		Person who shows sufficient proprietary interest	
If applicant is the legal representative, indicate the authority to file the patent application, the inventor is:			
<input type="button" value="Add"/>			
Name of the Deceased or Legally Incapacitated Inventor: <input type="text"/>			
If the Applicant is an Organization check here. <input checked="" type="checkbox"/>			
Organization Name	WEB SPARK LTD.		
Mailing Address Information For Applicant:			
Address 1	3 Hamahshev St.,		
Address 2			
City	Netanya	State/Province	
Country	IL	Postal Code	42507
Phone Number		Fax Number	
Email Address			
Additional Applicant Data may be generated within this form by selecting the Add button. <input type="button" value="Add"/>			

Assignee Information including Non-Applicant Assignee Information:

Providing assignment information in this section does not substitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76		Attorney Docket Number	HOLA-005-US10
		Application Number	
Title of Invention	SYSTEM PROVIDING FASTER AND MORE EFFICIENT DATA COMMUNICATION		

Assignee 1				
Complete this section if assignee information, including non-applicant assignee information, is desired to be included on the patent application publication. An assignee-applicant identified in the "Applicant Information" section will appear on the patent application publication as an applicant. For an assignee-applicant, complete this section only if identification as an assignee is also desired on the patent application publication.				
<input type="button" value="Remove"/>				
If the Assignee or Non-Applicant Assignee is an Organization check here. <input type="checkbox"/>				
Prefix	Given Name	Middle Name	Family Name	Suffix
Mailing Address Information For Assignee including Non-Applicant Assignee:				
Address 1				
Address 2				
City		State/Province		
Country i		Postal Code		
Phone Number		Fax Number		
Email Address				
Additional Assignee or Non-Applicant Assignee Data may be generated within this form by selecting the Add button. <input type="button" value="Add"/>				

Signature:

NOTE: This Application Data Sheet must be signed in accordance with 37 CFR 1.33(b). **However, if this Application Data Sheet is submitted with the INITIAL filing of the application and either box A or B is not checked in subsection 2 of the "Authorization or Opt-Out of Authorization to Permit Access" section, then this form must also be signed in accordance with 37 CFR 1.14(c).**

This Application Data Sheet **must** be signed by a patent practitioner if one or more of the applicants is a **juristic entity** (e.g., corporation or association). If the applicant is two or more joint inventors, this form must be signed by a patent practitioner, **all** joint inventors who are the applicant, or one or more joint inventor-applicants who have been given power of attorney (e.g., see USPTO Form PTO/AIA/81) on behalf of **all** joint inventor-applicants.

See 37 CFR 1.4(d) for the manner of making signatures and certifications.

Signature	/Yehuda Binder/		Date (YYYY-MM-DD)	2019-02-03
First Name	Yehuda	Last Name	BINDER	Registration Number
				73612
Additional Signature may be generated within this form by selecting the Add button. <input type="button" value="Add"/>				

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76		Attorney Docket Number	HOLA-005-US10
		Application Number	
Title of Invention	SYSTEM PROVIDING FASTER AND MORE EFFICIENT DATA COMMUNICATION		

This collection of information is required by 37 CFR 1.76. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 23 minutes to complete, including gathering, preparing, and submitting the completed application data sheet form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number		
	Filing Date		
	First Named Inventor	Derry Shribman	
	Art Unit		
	Examiner Name		
	Attorney Docket Number	HOLA-005-US10	

U.S.PATENTS						Remove
Examiner Initial*	Cite No	Patent Number	Kind Code ¹	Issue Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear
	1	7788378		2010-08-31	Ravi T. Rao	
	2	9253164		2016-02-02	Christopher S. Gouge	
	3	7890547	B2	2011-02-15	Timo Hotti	
	4	8832179	B2	2014-09-09	Owen , et al.	
	5	7818430	B2	2010-10-19	Gal Zuckerman	
	6	6154782	A	2000-11-28	NAOHISA KAWAGUCHI	
	7	5577243	A	1996-17-11	Sherwood , et al.	
	8	8135912	B2	2012-13-03	Shribman , et al.	

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number		
Filing Date		
First Named Inventor	Derry Shribman	
Art Unit		
Examiner Name		
Attorney Docket Number	HOLA-005-US10	

9	8719505	B2	2014-06-05	Shribman , et al.
10	9201808	B2	2015-01-12	Shribman , et al.
11	9990295	B2	2018-06-05	Shribman , et al.

If you wish to add additional U.S. Patent citation information please click the Add button.

Add

U.S.PATENT APPLICATION PUBLICATIONS

Remove

Examiner Initial*	Cite No	Publication Number	Kind Code ¹	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20080109446	A1	2008-05-08	Matrix Xin Wang	
	2	20110066924	A1	2011-03-17	Gregory Dorso	
	3	20110128911	A1	2011-06-02	Kamel M. Shaheen	
	4	20130157699	A1	2013-06-20	Mohit Talwar	
	5	20130326607	A1	2013-12-05	Liang Feng	
	6	20030204602	A1	2003-30-10	Hudson, Michael D. ; et al.	

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number		
Filing Date		
First Named Inventor	Derry Shribman	
Art Unit		
Examiner Name		
Attorney Docket Number	HOLA-005-US10	

7	20120124173	A1	2012-17-05	De; Pradipta ; et al.
8	20020069241	A1	2002-06-06	Narlikar, Girija ; et al.
9	20130201316	A1	2013-08-08	BINDER; Yehuda ; et al.
10	20120099566	A1	2012-26-04	Laine; Tuomas ; et al.
11	20120254370	A1	2012-10-04	Utz BACHER
12	20080125123	A1	2008-05-29	Jheroen P. Dorenbosch
13	20140301334	A1	2014-10-09	Miguel Labranche
14	20070239655	A1	2007-10-11	Masakuni Agetsuma
15	20070226810	A1	2007-09-27	Timo Hotti
16	20100094970	A1	2010-04-15	Gal Zuckerman
17	20130007253	A1	2013-01-03	Guohuai Li

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number		
Filing Date		
First Named Inventor	Derry Shribman	
Art Unit		
Examiner Name		
Attorney Docket Number	HOLA-005-US10	

18	20090037529	A1	2009-02-05	Gilad Armon-Kest
19	20090182843	A1	2009-07-16	Michael G. Hluchyj
20	20060036755	A1	2006-02-16	Ibrahim S. Abdullah
21	20140376403	A1	2014-12-25	Wenqi Shao
22	20050228964	A1	2005-13-10	Sechrest, Stuart ; et al.
23	20080086730	A1	2008-10-04	Vertes; Marc
24	20060259728	A1	2006-16-11	Chandrasekaran; Sashikanth ; et al.
25	20040254907	A1	2004-16-12	Crow, Preston F. ; et al.
26	20050015552	A1	2005-20-01	So, Kimming ; et al.
27	20050022236	A1	2005-01-27	Akihiko Ito; et al.

If you wish to add additional U.S. Published Application citation information please click the Add button.

FOREIGN PATENT DOCUMENTS

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number	
Filing Date	
First Named Inventor	Derry Shribman
Art Unit	
Examiner Name	
Attorney Docket Number	HOLA-005-US10

Examiner Initial*	Cite No	Foreign Document Number ³	Country Code ²ⁱ	Kind Code ⁴	Publication Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear	T ⁵
	1	2597869	EP	A1	2013-18-12	Sharp Kk		
	2	2010090562	WO	A1	2010-12-08	Telefonaktiebolaget L M Ericsson (Publ)		
	3	2011068784	WO	A1	2011-09-06	Azuki Systems, Inc		

If you wish to add additional Foreign Patent Document citation information please click the Add button

NON-PATENT LITERATURE DOCUMENTS

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T ⁵
	1	Screen captures from YouTube video clip entitle "nVpn.net Double your Safety and use Socks5 + nVpn" 38 pages, last accessed 11/20/2018 < https://www.youtube.com/watch?v=L0Hct2kSnn4 >	
	2	Screen captures from YouTube video clip entitle "Andromeda" 47 pages, publicly known and available as of at least 2011 < https://www.youtube.com/watch?v=yRRYpFLbKNU >	
	3	SpyEye, https://www.symantec.com/security-center/writeup/2010-020216-0135-9 ; http://seuresql.info/riskyclouds/spyeye-user-manual ; known as of at least 2010 (13 pages)	
	4	Screen captures from YouTube video clip entitle "Change Your Country IP Address & Location with Easy Hide IP Software" 9 pages, publicly known and available as of at least 2011, < https://www.youtube.com/watch?v=ulwkf1sOfdA and https://www.youtube.com/watch?v=iFEMT-o9DTc >	
	5	CoralCDN ("CoralCDN"), https://pdos.csail.mit.edu/6.824/papers/freedman-coral.pdf (14 PAGES)	

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number		
Filing Date		
First Named Inventor	Derry Shribman	
Art Unit		
Examiner Name		
Attorney Docket Number	HOLA-005-US10	

6	European Search Report for EP 14182547.1, dated July 30, 2015
7	R. Fielding et al, RFC 2616: Hypertext Transfer Protocol -- HTTP/1.1, June 1999, retrieved from the Internet http://rfc-editor.org [retrieved Apr. 15, 2002]
8	"On the leakage of personally identifiable information via online social networks", Wills et al. AT&T, Apr. 2009 http://www2.research.att.com/-bala/papers/wosn09.pdf
9	"Slice Embedding Solutions for Distributed Service Architectures" - Esposito et al., Boston University, Computer Science Dept., 10/2011 http://www.cs.bu.edu/techreports/pdf/2011-025-slice-embedding.pdf
10	International Search Report of PCT/US2010/034072 dated July 01, 2010
11	YouTube video clip entitled "nVpn.net Double your Safety and use Socks5 + nVpn" < https://www.youtube.com/watch?v=L0Hct2kSnn4 >
12	YouTube video clip entitled "Andromeda" < https://www.youtube.com/watch?v=yRRYpFLbKNU >
13	YouTube video clip entitled "Change Your Country IP Address & Location with Easy Hide IP Software" < https://www.youtube.com/watch?v=ulwfk1sOfdA and https://www.youtube.com/watch?v=iFEMT-p9DTc >

If you wish to add additional non-patent literature document citation information please click the Add button

EXAMINER SIGNATURE

Examiner Signature		Date Considered	
--------------------	--	-----------------	--

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

¹ See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. ² Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). ³ For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. ⁴ Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. ⁵ Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number		
	Filing Date		
	First Named Inventor	Derry Shribman	
	Art Unit		
	Examiner Name		
	Attorney Docket Number	HOLA-005-US10	

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.

The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Yehuda Binder/	Date (YYYY-MM-DD)	2019-02-03
Name/Print	Yehuda Binder	Registration Number	73,612

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. **DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number		
	Filing Date		
	First Named Inventor	Derry Shribman	
	Art Unit		
	Examiner Name		
	Attorney Docket Number	HOLA-005-US10	

U.S.PATENTS						Remove
Examiner Initial*	Cite No	Patent Number	Kind Code ¹	Issue Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear
	1	8479251	B2	2013-07-02	Feinleib et al	
	2	8499059	B2	2013-07-30	Stoyanov	
	3	7970835	B2	2011-28-01	Xerox Corporation	
	4	8832179	B2	2014-09-09	Owen , et al.	
	5	6173330	B1	2001-09-01	Guo , et al.	
	6	8769035	B2	2014-01-07	Resch , et al.	
	7	8171101	B2	2012-05-01	Gladwin , et al.	
	8	7558942	B2	2009-07-07	Chen , et al.	

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number		
Filing Date		
First Named Inventor	Derry Shribman	
Art Unit		
Examiner Name		
Attorney Docket Number	HOLA-005-US10	

9	4937781	A	1990-06-26	Lee , et al.
10	7970835	B2	2011-06-28	Robert St. Jacques

If you wish to add additional U.S. Patent citation information please click the Add button.

Add

U.S.PATENT APPLICATION PUBLICATIONS

Remove

Examiner Initial*	Cite No	Publication Number	Kind Code ¹	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20150067819	A1	2015-03-05	Hola Networks Ltd.	
	2	20120254456	A1	2012-10-04	Misharam Zubair et al.	
	3	20080222291	A1	2008-09-11	Weller et al.	
	4	20100235438	A1	2010-09-16	Narayanan et al.	
	5	20120124239	A1	2012-05-17	Shribman et al.	
	6	20130166768	A1	2013-06-27	Thomson Licensing	
	7	20020065930	A1	2002-30-05	Rhodes, David L.	

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number		
Filing Date		
First Named Inventor	Derry Shribman	
Art Unit		
Examiner Name		
Attorney Docket Number	HOLA-005-US10	

8	20030204602	A1	2003-10-30	Hudson Michael D.
9	20120099566	A1	2012-04-26	Laine; Tuomas ; et al.
10	20130201316	A1	2013-08-08	BINDER; Yehuda ; et al.
11	20080125123	A1	2008-05-29	Dorenbosch; Jheroen P. ; et al.
12	20140301334	A1	2014-10-09	Labranche; Miguel ; et al.
13	20070239655	A1	2007-10-11	Agetsuma; Masakuni ; et al.
14	20070226810	A1	2007-09-27	Hotti; Timo
15	20100094970	A1	2010-04-15	Zuckerman; Gal ; et al.
16	20020120874	A1	2002-29-08	Shu, Li ; et al.
17	20100115063	A1	2010-06-05	GLADWIN; S. CHRISTOPHER ; et al.
18	20100154044	A1	2010-17-06	Manku; Tajinder

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number		
Filing Date		
First Named Inventor	Derry Shribman	
Art Unit		
Examiner Name		
Attorney Docket Number	HOLA-005-US10	

19	20100293555	A1	2010-15-11	VEPSALAINEN; Ari M.
20	20130272519	A1	2013-17-10	Huang; Lawrence P.
21	20030115364	A1	2003-06-19	Shu Li et al.
22	20090217122	A1	2009-27-08	Yokokawa; Takashi ; et al.
23	20010033583	A1	2001-25-10	Rabenko, Theodore F. ; et al.
24	20080109446	A1	2008-05-08	Wang Matrix XIN
25	20020133621	A1	2002-09-19	Talmon Marco et al
26	20040107242	A1	2004-06-03	John Vert et al
27	20070073878	A1	2007-03-29	Alfredo C. Issa
28	20090319502	A1	2009-12-24	Olivier Chalouhi et al
29	20060212584	A1	2006-09-21	Mingjian Yu et al

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number		
Filing Date		
First Named Inventor	Derry Shribman	
Art Unit		
Examiner Name		
Attorney Docket Number	HOLA-005-US10	

If you wish to add additional U.S. Published Application citation information please click the Add button.

FOREIGN PATENT DOCUMENTS

Examiner Initial*	Cite No	Foreign Document Number ³	Country Code ² i	Kind Code ⁴	Publication Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear	T ⁵
	1	2015034752	WO	A1	2015-03-12	Akamai Technologies INC		
	2	2000/018078	WO	A1	2000-03-30	Sopuch David. J		
	3	0948176	EP	A2	1999-10-06	Siemens Inf &Comm Networks		
	4	2597869	EP	A1	2015-05-29	Sharp Kabushiki Kaisha Osaka-shi		
	5	2010090562	WO	A1	2010-08-12	Telefonaktiebolaget L M Ericsson		
	6	2007280388	JP		2007-25-10	Xerox Corporation		
	7	1020090097034	KR		2009-15-09	KT Corporation		
	8	2343536	RU	C2	2009-10-01	Microsoft Corporation		
	9	101075242	CN	A	2007-11-21	TENGXUN SCIENCE & TECHNOLOGY		

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number		
	Filing Date		
	First Named Inventor	Derry Shribman	
	Art Unit		
	Examiner Name		
	Attorney Docket Number	HOLA-005-US10	

	10	101179389	CN	A	2008-05-14	Wang Matrix XIN	
--	----	-----------	----	---	------------	-----------------	--

If you wish to add additional Foreign Patent Document citation information please click the Add button

NON-PATENT LITERATURE DOCUMENTS

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T ⁵
	1	R. Fielding et al, RFC 2616: Hypertext Transfer Protocol -- HTTP/1.1, June 1999, retrieved from the Internet http://rfc-editor.org [retrieved Apr. 15, 2002] (114 pages)	
	2	"On the Leakage of Personally Identifiable Information via Online Social Networks"-Wills et al, AT&T, Apr. 2009 http://www2.research.att.com/~bala/papers/wosn09.pdf .	
	3	Notice of Preliminary Rejection in KR Application No. 10-2012-7011711 dated July 15, 2016	
	4	KEI SUZUKI, a study on Cooperative Peer Selection Method in P2P Video Delivery, Vol. 109, No. 37, IEICE Technical Report, The Institute of Electronics, Information and Communication Engineers, May 14, 2009	

If you wish to add additional non-patent literature document citation information please click the Add button

EXAMINER SIGNATURE

Examiner Signature		Date Considered	
--------------------	--	-----------------	--

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

¹ See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. ² Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). ³ For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. ⁴ Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. ⁵ Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number		
	Filing Date		
	First Named Inventor	Derry Shribman	
	Art Unit		
	Examiner Name		
	Attorney Docket Number	HOLA-005-US10	

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.

The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Yehuda Binder/	Date (YYYY-MM-DD)	2019-02-03
Name/Print	Yehuda BINDER	Registration Number	73612

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. **DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number		
	Filing Date		
	First Named Inventor	Derry Shribman	
	Art Unit		
	Examiner Name		
	Attorney Docket Number	HOLA-005-US10	

U.S.PATENTS						Remove
Examiner Initial*	Cite No	Patent Number	Kind Code ¹	Issue Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear
	1	3922494	A	1975-11-25	Cooper , et al.	
	2	5758195	A	1998-05-26	Balmer; Keith	
	3	6061278	A	2000-05-09	Kato , et al.	
	4	6466470	B1	2002-10-15	Houn Chang	
	5	7865585		2011-01-04	Allen Samuels, et al.	
	6	7120666		2006-10-10	Steven McCanne, et al.	
	7	7203741		2007-04-10	Talmon Marco, et al.	
If you wish to add additional U.S. Patent citation information please click the Add button.						Add
U.S.PATENT APPLICATION PUBLICATIONS						Remove

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number		
Filing Date		
First Named Inventor	Derry Shribman	
Art Unit		
Examiner Name		
Attorney Docket Number	HOLA-005-US10	

Examiner Initial*	Cite No	Publication Number	Kind Code ¹	Publication Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear
	1	20030009518	A1	2003-01-09	Harrow, Ivan P. ; et al.	
	2	20030074403	A1	2003-04-17	Harrow, Ivan P. ; et al.	
	3	20140082260	A1	2014-03-20	OH; HakJune ; et al.	
	4	20110314347	A1	2011-12-22	NAKANO; Rikizo ; et al.	
	5	20100329270	A1	2010-12-30	Asati; Rajiv ; et al.	
	6	20100085977	A1	2010-04-08	Khalid; Mohamed ; et al.	
	7	20100066808	A1	2010-03-18	Tucker; Curtis E. ; et al.	
	8	20090279559	A1	2009-11-12	Wong; Yuen Fai ; et al.	
	9	20080025506	A1	2008-01-31	Muraoka; Jochiku	
	10	20040264506	A1	2004-12-30	Furukawa, Rei	

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number		
Filing Date		
First Named Inventor	Derry Shribman	
Art Unit		
Examiner Name		
Attorney Docket Number	HOLA-005-US10	

11	20020123895	A1	2002-09-05	Sergey Potekhin
12	20150033001	A1	2015-01-29	Ivanov; Vladimir
13	20150358648	A1	2015-12-10	Limberg; Allen LeRoy
14	20160021430	A1	2016-01-21	LaBosco; Mark ; et al.
15	20110087733	A1	2011-04-14	Derry Shribman; et al.
16	20030174648	A1	2003-09-18	Mea Wang; et al.
17	20080008089	A1	2008-01-10	Claudson F. Bornstein; et al.
18	20040088646	A1	2004-05-06	William J. Yeager; et al.
19	20030009583	A1	2003-01-09	Chung Chan; et al.
20	20080235391	A1	2008-09-25	Christopher Painter; et al.
21	20070156855	A1	2007-07-05	Moses Johnson

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number	
Filing Date	
First Named Inventor	Derry Shribman
Art Unit	
Examiner Name	
Attorney Docket Number	HOLA-005-US10

22	20020007413	A1	2002-01-17	JJ Garcia-Luna-Aceves, et al.
23	20030210694	A1	2003-11-13	Suresh Jayaraman, et al.
24	20030200307	A1	2003-10-23	Jyoti Raju, et al.

If you wish to add additional U.S. Published Application citation information please click the Add button.

FOREIGN PATENT DOCUMENTS

Examiner Initial*	Cite No	Foreign Document Number ³	Country Code ²	Kind Code ⁴	Publication Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear	T ⁵
	1							

If you wish to add additional Foreign Patent Document citation information please click the Add button.

NON-PATENT LITERATURE DOCUMENTS

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T ⁵
	1	International Search Report issued in PCT Application No. PCT/US2010/051881 dated 09 December 2010	
	2	Supplementary European Search Report issued in EP Application No. 10822724 dated 24 April 2013	

If you wish to add additional non-patent literature document citation information please click the Add button.

EXAMINER SIGNATURE

Examiner Signature		Date Considered	
--------------------	--	-----------------	--

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number	
Filing Date	
First Named Inventor	Derry Shribman
Art Unit	
Examiner Name	
Attorney Docket Number	HOLA-005-US10

¹ See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. ² Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). ³ For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. ⁴ Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. ⁵ Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number		
	Filing Date		
	First Named Inventor	Derry Shribman	
	Art Unit		
	Examiner Name		
	Attorney Docket Number	HOLA-005-US10	

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.

The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Yehuda Binder/	Date (YYYY-MM-DD)	2019-02-03
Name/Print	Yehuda BINDER	Registration Number	73612

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. **DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number	
	Filing Date	
	First Named Inventor	Derry Shribman
	Art Unit	
	Examiner Name	
	Attorney Docket Number	HOLA-005-US10

U.S.PATENTS						Remove
Examiner Initial*	Cite No	Patent Number	Kind Code ¹	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	6868453	B1	2005-03-15	Mitsuhiro Watanabe	
	2	8595786	B2	2013-11-26	In Hwan Choi	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS						Remove
Examiner Initial*	Cite No	Publication Number	Kind Code ¹	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20030097408	A1	2003-05-22	Masahiro Kageyama	
	2	20070100839	A1	2007-05-03	Deok-ho Kim	
	3	20080256175	A1	2008-10-16	Sang-kwon Lee	
	4	20060212542	A1	2006-09-21	Han Fang	

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number		
Filing Date		
First Named Inventor	Derry Shribman	
Art Unit		
Examiner Name		
Attorney Docket Number	HOLA-005-US10	

5	20110035503	A1	2011-02-10	SAM ZAID
6	20050097441	A1	2005-05-05	Jonathan D. Herbach

If you wish to add additional U.S. Published Application citation information please click the Add button.

FOREIGN PATENT DOCUMENTS

Examiner Initial*	Cite No	Foreign Document Number ³	Country Code ² i	Kind Code ⁴	Publication Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear	T ⁵
	1							

If you wish to add additional Foreign Patent Document citation information please click the Add button.

NON-PATENT LITERATURE DOCUMENTS

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T ⁵
	1		

If you wish to add additional non-patent literature document citation information please click the Add button.

EXAMINER SIGNATURE

Examiner Signature		Date Considered	
--------------------	--	-----------------	--

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

¹ See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. ² Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). ³ For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. ⁴ Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. ⁵ Applicant is to place a check mark here if English language translation is attached.

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number	
Filing Date	
First Named Inventor	Derry Shribman
Art Unit	
Examiner Name	
Attorney Docket Number	HOLA-005-US10

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.

The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Yehuda Binder/	Date (YYYY-MM-DD)	2019-02-03
Name/Print	Yehuda BINDER	Registration Number	73612

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. **DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number		
	Filing Date		
	First Named Inventor	Derry Shribman	
	Art Unit		
	Examiner Name		
	Attorney Docket Number	HOLA-005-US10	

U.S.PATENTS						Remove
Examiner Initial*	Cite No	Patent Number	Kind Code ¹	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	7742485	B2	2010-06-22	Xinyan Zhang	
	2	7831720	B1	2010-11-09	Wael Noureddine	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS						Remove
Examiner Initial*	Cite No	Publication Number	Kind Code ¹	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20140359081	A1	2014-12-04	Mattijs Oskar Van Deventer	
	2	20090010426	A1	2009-01-08	Scott D. Redmond	
	3	20130007232	A1	2013-01-03	Wei Wang	
	4	20150206197	A1	2015-07-23	Assaf Toval	

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number		
Filing Date		
First Named Inventor	Derry Shribman	
Art Unit		
Examiner Name		
Attorney Docket Number	HOLA-005-US10	

5	20150206176	A1	2015-07-23	Assaf Toval
6	20170221092	A1	2017-08-03	Assaf Toval
7	20070174246	A1	2007-07-26	Johann Tomas Sigurdsson
8	20100262650	A1	2010-10-14	Abhishek Chauhan
9	20060047844	A1	2006-03-02	Li Deng
10	20130171964	A1	2013-07-04	Sumeet Singh Bhatia
11	20130219458	A1	2013-08-22	Vasudevan Ramanathan

If you wish to add additional U.S. Published Application citation information please click the Add button.

FOREIGN PATENT DOCUMENTS

Examiner Initial*	Cite No	Foreign Document Number ³	Country Code ² i	Kind Code ⁴	Publication Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear	T ⁵
	1							

If you wish to add additional Foreign Patent Document citation information please click the Add button.

NON-PATENT LITERATURE DOCUMENTS

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number		
Filing Date		
First Named Inventor	Derry Shribman	
Art Unit		
Examiner Name		
Attorney Docket Number	HOLA-005-US10	

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T ⁵
	1		

If you wish to add additional non-patent literature document citation information please click the Add button

EXAMINER SIGNATURE

Examiner Signature		Date Considered	
--------------------	--	-----------------	--

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

¹ See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. ² Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). ³ For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. ⁴ Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. ⁵ Applicant is to place a check mark here if English language translation is attached.

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number	
Filing Date	
First Named Inventor	Derry Shribman
Art Unit	
Examiner Name	
Attorney Docket Number	HOLA-005-US10

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.

The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Yehuda Binder/	Date (YYYY-MM-DD)	2019-02-03
Name/Print	Yehuda Binder	Registration Number	73,612

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. **DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Patent Application Fee Transmittal

Application Number:				
Filing Date:				
Title of Invention:	SYSTEM PROVIDING FASTER AND MORE EFFICIENT DATA COMMUNICATION			
First Named Inventor/Applicant Name:	Derry Shribman			
Filer:	Yehuda Binder/Dorit Binder			
Attorney Docket Number:	HOLA-005-US10			
Filed as Small Entity				
Filing Fees for Utility under 35 USC 111(a)				
Description	Fee Code	Quantity	Amount	Sub-Total in USD(\$)
Basic Filing:				
UTILITY FILING FEE (ELECTRONIC FILING)	4011	1	75	75
UTILITY SEARCH FEE	2111	1	330	330
UTILITY EXAMINATION FEE	2311	1	380	380
Pages:				
Claims:				
CLAIMS IN EXCESS OF 20	2202	4	50	200
Miscellaneous-Filing:				
Petition:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD(\$)
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				
Miscellaneous:				
Total in USD (\$)				985

Electronic Acknowledgement Receipt

EFS ID:	35173725
Application Number:	16278107
International Application Number:	
Confirmation Number:	4936
Title of Invention:	SYSTEM PROVIDING FASTER AND MORE EFFICIENT DATA COMMUNICATION
First Named Inventor/Applicant Name:	Derry Shribman
Customer Number:	131926
Filer:	Yehuda Binder/Dorit Binder
Filer Authorized By:	Yehuda Binder
Attorney Docket Number:	HOLA-005-US10
Receipt Date:	17-FEB-2019
Filing Date:	
Time Stamp:	04:49:10
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	DA
Payment was successfully received in RAM	\$985
RAM confirmation Number	021919INTEFSW00007557506726
Deposit Account	
Authorized User	

The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows:

File Listing:					
Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Transmittal of New Application	Transmittal-letter.pdf	77426	no	2
			48e302932a13fd2e0c3bd415b864cee54c02d36		
Warnings:					
Information:					
2	Power of Attorney	Signed-PoA-US7-US8-US9-US10-US11.pdf	1625826	no	4
			11fe2ecf081b56b974a10c73d6c43a200cf1bb20		
Warnings:					
Information:					
3	Specification	Spec.pdf	3001588	no	33
			65902c440cef2c477c0b3de0ee9540abcfd0c6f8		
Warnings:					
Information:					
4	Claims	Claims.pdf	21671	no	4
			d7632fbc8489248c892f3b6b58d47b1177e1eb73		
Warnings:					
Information:					
5	Abstract	Abstract.pdf	11557	no	1
			25e11eb59e98b498aa94b914157c1c71edffe04f4		
Warnings:					
Information:					
6	Drawings-only black and white line drawings	Drawings-14025109.pdf	258016	no	15
			fa77c1f6c6442922c86a9e390fa2a1ba55639fe3		
Warnings:					
Information:					

7	Application Data Sheet	ADS.pdf	1823564	no	9
			4789dd60bb4377bac0a900e526ae7ab96227bb8f		
Warnings:					
Information:					
8	Oath or Declaration filed	Signed-oath-Derry.pdf	862135	no	2
			6a2192d0180f49adf2112bf067a659cb3f1f8144		
Warnings:					
Information:					
9	Oath or Declaration filed	Signed-oath-Ofer.pdf	871562	no	2
			164ffbbaa7776b17398cb08bf75bfa2cac4e62d20		
Warnings:					
Information:					
10	Information Disclosure Statement (IDS) Form (SB08)	005-US6-003-004-OL.pdf	1036993	no	8
			54a54196daeecb34f39145a5813556709cfe4da52		
Warnings:					
Information:					
11	Information Disclosure Statement (IDS) Form (SB08)	IDS.pdf	1037469	no	8
			8ad956f54af8c1a942fe5ba533c27122893e77fd		
Warnings:					
Information:					
12	Information Disclosure Statement (IDS) Form (SB08)	IDS2.pdf	1036390	no	7
			b466a030f400a9005c3cb756fc7910bf7badfed8		
Warnings:					
Information:					
13	Information Disclosure Statement (IDS) Form (SB08)	IDS3.pdf	1034627	no	4
			b306fe151e2a171b8c8073add9e0011abd52f86e		
Warnings:					
Information:					

14	Information Disclosure Statement (IDS) Form (SB08)	IDS4.pdf	1034788	no	5
			315693975b8ed7c34d572a61cf348e6c736e023c		
Warnings:					
Information:					
15	Fee Worksheet (SB06)	fee-info.pdf	36689	no	2
			24548292b6367c9da4930f137e4dd3d3bdf8ef79		
Warnings:					
Information:					
Total Files Size (in bytes):			13770301		
<p>This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.</p> <p><u>New Applications Under 35 U.S.C. 111</u> If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.</p> <p><u>National Stage of an International Application under 35 U.S.C. 371</u> If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.</p> <p><u>New International Application Filed with the USPTO as a Receiving Office</u> If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.</p>					

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

TRANSMITTAL FOR POWER OF ATTORNEY TO ONE OR MORE REGISTERED PRACTITIONERS

NOTE: This form is to be submitted with the Power of Attorney by Applicant form (PTO/AIA/82B) to identify the application to which the Power of Attorney is directed, in accordance with 37 CFR 1.5, unless the application number and filing date are identified in the Power of Attorney by Applicant form. If neither form PTO/AIA/82A nor form PTO/AIA82B identifies the application to which the Power of Attorney is directed, the Power of Attorney will not be recognized in the application.

Application Number	
Filing Date	
First Named Inventor	Derry Shribman
Title	SYSTEM PROVIDING FASTER AND MORE EFFICIENT DATA COMMUNICATION
Art Unit	
Examiner Name	
Attorney Docket Number	HOLA-005-US7, US8, HOLA-005-US9, US10

SIGNATURE of Applicant or Patent Practitioner			
Signature	/Yehuda Binder/	Date (Optional)	2019-02-03
Name	Yehuda Binder	Registration Number	73,612
Title (if Applicant is a juristic entity)			
Applicant Name (if Applicant is a juristic entity)			
<p>NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4(d) for signature requirements and certifications. If more than one applicant, use multiple forms.</p>			
<input type="checkbox"/> *Total of _____ forms are submitted.			

This collection of information is required by 37 CFR 1.131, 1.32, and 1.33. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 3 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number

POWER OF ATTORNEY BY APPLICANT

I hereby revoke all previous powers of attorney given in the application identified in either the attached transmittal letter or the boxes below.

Application Number	Filing Date

(Note: The boxes above may be left blank if information is provided on form PTO/AIA/82A.)

- I hereby appoint the Patent Practitioner(s) associated with the following Customer Number as my/our attorney(s) or agent(s), and to transact all business in the United States Patent and Trademark Office connected therewith for the application referenced in the attached transmittal letter (form PTO/AIA/82A) or identified above: 131926
- OR
- I hereby appoint Practitioner(s) named in the attached list (form PTO/AIA/82C) as my/our attorney(s) or agent(s), and to transact all business in the United States Patent and Trademark Office connected therewith for the patent application referenced in the attached transmittal letter (form PTO/AIA/82A) or identified above. (Note: Complete form PTO/AIA/82C.)

Please recognize or change the correspondence address for the application identified in the attached transmittal letter or the boxes above to:

- The address associated with the above-mentioned Customer Number
- OR
- The address associated with Customer Number:
- OR

<input type="checkbox"/> Firm or Individual Name			
Address			
City	State	Zip	
Country			
Telephone	Email		

I am the Applicant (if the Applicant is a juristic entity, list the Applicant name in the box):

WEB SPARK LTD.

- Inventor or Joint Inventor (title not required below)
- Legal Representative of a Deceased or Legally Incapacitated Inventor (title not required below)
- Assignee or Person to Whom the Inventor is Under an Obligation to Assign (provide signer's title if applicant is a juristic entity)
- Person Who Otherwise Shows Sufficient Proprietary Interest (e.g., a petition under 37 CFR 1.46(b)(2) was granted in the application or is concurrently being filed with this document) (provide signer's title if applicant is a juristic entity)

SIGNATURE of Applicant for Patent

The undersigned (whose title is supplied below) is authorized to act on behalf of the applicant (e.g., where the applicant is a juristic entity).

Signature	Date (Optional)
	6 - Feb - 2019
Name	Derry Shnbman
Title	CEO of Web Spark Ltd.

NOTE: Signature - This form must be signed by the applicant in accordance with 37 CFR 1.33. See 37 CFR 1.4 for signature requirements and certifications. If more than one applicant, use multiple forms.

Total of _____ forms are submitted.

This collection of information is required by 37 CFR 1.131, 1.32, and 1.33. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 3 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number

POWER OF ATTORNEY BY APPLICANT

No more than ten (10) patent practitioners total may be appointed as set forth below by name and registration number. This page need not be submitted if appointing the Patent Practitioner(s) associated with a Customer Number (see form PTO/AIA/82B):

Name	Registration Number

Privacy Act Statement

The **Privacy Act of 1974 (P.L. 93-579)** requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

DECLARATION (37 CFR 1.63) FOR UTILITY OR DESIGN APPLICATION USING AN APPLICATION DATA SHEET (37 CFR 1.76)

Title of Invention SYSTEM PROVIDING FASTER AND MORE EFFICIENT DATA COMMUNICATION

As the below named inventor, I hereby declare that:

This declaration is directed to: The attached application, or
 United States application or PCT international application number _____
filed on _____.

The above-identified application was made or authorized to be made by me.

I believe that I am the original inventor or an original joint inventor of a claimed invention in the application.

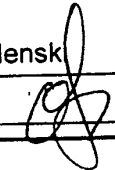
I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

WARNING:

Petitioner/applicant is cautioned to avoid submitting personal information in documents filed in a patent application that may contribute to identity theft. Personal information such as social security numbers, bank account numbers, or credit card numbers (other than a check or credit card authorization form PTO-2038 submitted for payment purposes) is never required by the USPTO to support a petition or an application. If this type of personal information is included in documents submitted to the USPTO, petitioners/applicants should consider redacting such personal information from the documents before submitting them to the USPTO. Petitioner/applicant is advised that the record of a patent application is available to the public after publication of the application (unless a non-publication request in compliance with 37 CFR 1.213(a) is made in the application) or issuance of a patent. Furthermore, the record from an abandoned application may also be available to the public if the application is referenced in a published application or an issued patent (see 37 CFR 1.14). Checks and credit card authorization forms PTO-2038 submitted for payment purposes are not retained in the application file and therefore are not publicly available.

LEGAL NAME OF INVENTOR

Inventor: Ofer Vilensk _____ Date (Optional) : _____

Signature:  _____

Note: An application data sheet (PTO/SB/14 or equivalent), including naming the entire inventive entity, must accompany this form or must have been previously filed. Use an additional PTO/AIA/01 form for each additional inventor.

This collection of information is required by 35 U.S.C. 115 and 37 CFR 1.63. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 1 minute to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Privacy Act Statement

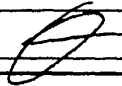
The **Privacy Act of 1974 (P.L. 93-579)** requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether disclosure of these records is required by the Freedom of Information Act.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspection or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

DECLARATION (37 CFR 1.63) FOR UTILITY OR DESIGN APPLICATION USING AN APPLICATION DATA SHEET (37 CFR 1.76)

Title of Invention	SYSTEM PROVIDING FASTER AND MORE EFFICIENT DATA COMMUNICATION
<p>As the below named inventor, I hereby declare that:</p> <p>This declaration is directed to: <input checked="" type="checkbox"/> The attached application, or <input type="checkbox"/> United States application or PCT international application number _____ filed on _____.</p> <p>The above-identified application was made or authorized to be made by me.</p> <p>I believe that I am the original inventor or an original joint inventor of a claimed invention in the application.</p> <p>I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.</p> <p align="center">WARNING:</p> <p>Petitioner/applicant is cautioned to avoid submitting personal information in documents filed in a patent application that may contribute to identity theft. Personal information such as social security numbers, bank account numbers, or credit card numbers (other than a check or credit card authorization form PTO-2038 submitted for payment purposes) is never required by the USPTO to support a petition or an application. If this type of personal information is included in documents submitted to the USPTO, petitioners/applicants should consider redacting such personal information from the documents before submitting them to the USPTO. Petitioner/applicant is advised that the record of a patent application is available to the public after publication of the application (unless a non-publication request in compliance with 37 CFR 1.213(a) is made in the application) or issuance of a patent. Furthermore, the record from an abandoned application may also be available to the public if the application is referenced in a published application or an issued patent (see 37 CFR 1.14). Checks and credit card authorization forms PTO-2038 submitted for payment purposes are not retained in the application file and therefore are not publicly available.</p>	
LEGAL NAME OF INVENTOR	
Inventor: <u>Derry Shribman</u> Date (Optional) : _____	
Signature: 	
<p>Note: An application data sheet (PTO/SB/14 or equivalent), including naming the entire inventive entity, must accompany this form or must have been previously filed. Use an additional PTO/AIA/01 form for each additional inventor.</p>	

This collection of information is required by 35 U.S.C. 115 and 37 CFR 1.63. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 1 minute to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Privacy Act Statement

The **Privacy Act of 1974 (P.L. 93-579)** requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether disclosure of these records is required by the Freedom of Information Act.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspection or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

Table with 7 columns: APPLICATION NUMBER, FILING or 371(c) DATE, GRP ART UNIT, FIL FEE REC'D, ATTY. DOCKET NO, TOT CLAIMS, IND CLAIMS. Row 1: 16/278,107, 02/17/2019, 2447, 985, HOLA-005-US10, 24, 1

CONFIRMATION NO. 4936

FILING RECEIPT

131926
May Patents Ltd. c/o Dorit Shem-Tov
P.O.B 7230
Ramat-Gan, 5217102
ISRAEL



Date Mailed: 03/06/2019

Receipt is acknowledged of this non-provisional patent application. The application will be taken up for examination in due course. Applicant will be notified as to the results of the examination. Any correspondence concerning the application must include the following identification information: the U.S. APPLICATION NUMBER, FILING DATE, NAME OF APPLICANT, and TITLE OF INVENTION. Fees transmitted by check or draft are subject to collection. Please verify the accuracy of the data presented on this receipt. If an error is noted on this Filing Receipt, please submit a written request for a Filing Receipt Correction. Please provide a copy of this Filing Receipt with the changes noted thereon. If you received a "Notice to File Missing Parts" for this application, please submit any corrections to this Filing Receipt with your reply to the Notice. When the USPTO processes the reply to the Notice, the USPTO will generate another Filing Receipt incorporating the requested corrections

Inventor(s)

Derry Shribman, Tel Aviv, ISRAEL;
Ofer Vilenski, Moshav Hadar Am, ISRAEL;

Applicant(s)

WEB SPARK LTD., Netanya, ISRAEL;

Power of Attorney: The patent practitioners associated with Customer Number 131926

Domestic Priority data as claimed by applicant

This application is a CON of 15/957,945 04/20/2018
which is a CON of 14/025,109 09/12/2013 PAT 10069936
which is a DIV of 12/836,059 07/14/2010 PAT 8560604
which claims benefit of 61/249,624 10/08/2009

Foreign Applications for which priority is claimed (You may be eligible to benefit from the Patent Prosecution Highway program at the USPTO. Please see http://www.uspto.gov for more information.) - None.

Foreign application information must be provided in an Application Data Sheet in order to constitute a claim to foreign priority. See 37 CFR 1.55 and 1.76.

Permission to Access Application via Priority Document Exchange: Yes

Permission to Access Search Results: Yes

Applicant may provide or rescind an authorization for access using Form PTO/SB/39 or Form PTO/SB/69 as appropriate.

If Required, Foreign Filing License Granted: 03/05/2019

The country code and number of your priority application, to be used for filing abroad under the Paris Convention, is **US 16/278,107**

Projected Publication Date: 06/13/2019

Non-Publication Request: No

Early Publication Request: No

**** SMALL ENTITY ****

Title

SYSTEM PROVIDING FASTER AND MORE EFFICIENT DATA COMMUNICATION

Preliminary Class

709

Statement under 37 CFR 1.55 or 1.78 for AIA (First Inventor to File) Transition Applications: No

PROTECTING YOUR INVENTION OUTSIDE THE UNITED STATES

Since the rights granted by a U.S. patent extend only throughout the territory of the United States and have no effect in a foreign country, an inventor who wishes patent protection in another country must apply for a patent in a specific country or in regional patent offices. Applicants may wish to consider the filing of an international application under the Patent Cooperation Treaty (PCT). An international (PCT) application generally has the same effect as a regular national patent application in each PCT-member country. The PCT process **simplifies** the filing of patent applications on the same invention in member countries, but **does not result** in a grant of "an international patent" and does not eliminate the need of applicants to file additional documents and fees in countries where patent protection is desired.

Almost every country has its own patent law, and a person desiring a patent in a particular country must make an application for patent in that country in accordance with its particular laws. Since the laws of many countries differ in various respects from the patent law of the United States, applicants are advised to seek guidance from specific foreign countries to ensure that patent rights are not lost prematurely.

Applicants also are advised that in the case of inventions made in the United States, the Director of the USPTO must issue a license before applicants can apply for a patent in a foreign country. The filing of a U.S. patent application serves as a request for a foreign filing license. The application's filing receipt contains further information and guidance as to the status of applicant's license for foreign filing.

Applicants may wish to consult the USPTO booklet, "General Information Concerning Patents" (specifically, the section entitled "Treaties and Foreign Patents") for more information on timeframes and deadlines for filing foreign patent applications. The guide is available either by contacting the USPTO Contact Center at 800-786-9199, or it can be viewed on the USPTO website at <http://www.uspto.gov/web/offices/pac/doc/general/index.html>.

For information on preventing theft of your intellectual property (patents, trademarks and copyrights), you may wish to consult the U.S. Government website, <http://www.stopfakes.gov>. Part of a Department of Commerce initiative, this website includes self-help "toolkits" giving innovators guidance on how to protect intellectual property in specific countries such as China, Korea and Mexico. For questions regarding patent enforcement issues, applicants may call the U.S. Government hotline at 1-866-999-HALT (1-866-999-4258).

LICENSE FOR FOREIGN FILING UNDER
Title 35, United States Code, Section 184
Title 37, Code of Federal Regulations, 5.11 & 5.15

GRANTED

The applicant has been granted a license under 35 U.S.C. 184, if the phrase "IF REQUIRED, FOREIGN FILING LICENSE GRANTED" followed by a date appears on this form. Such licenses are issued in all applications where the conditions for issuance of a license have been met, regardless of whether or not a license may be required as set forth in 37 CFR 5.15. The scope and limitations of this license are set forth in 37 CFR 5.15(a) unless an earlier license has been issued under 37 CFR 5.15(b). The license is subject to revocation upon written notification. The date indicated is the effective date of the license, unless an earlier license of similar scope has been granted under 37 CFR 5.13 or 5.14.

This license is to be retained by the licensee and may be used at any time on or after the effective date thereof unless it is revoked. This license is automatically transferred to any related applications(s) filed under 37 CFR 1.53(d). This license is not retroactive.

The grant of a license does not in any way lessen the responsibility of a licensee for the security of the subject matter as imposed by any Government contract or the provisions of existing laws relating to espionage and the national security or the export of technical data. Licensees should apprise themselves of current regulations especially with respect to certain countries, of other agencies, particularly the Office of Defense Trade Controls, Department of State (with respect to Arms, Munitions and Implements of War (22 CFR 121-128)); the Bureau of Industry and Security, Department of Commerce (15 CFR parts 730-774); the Office of Foreign Assets Control, Department of Treasury (31 CFR Parts 500+) and the Department of Energy.

NOT GRANTED

No license under 35 U.S.C. 184 has been granted at this time, if the phrase "IF REQUIRED, FOREIGN FILING LICENSE GRANTED" DOES NOT appear on this form. Applicant may still petition for a license under 37 CFR 5.12, if a license is desired before the expiration of 6 months from the filing date of the application. If 6 months has lapsed from the filing date of this application and the licensee has not received any indication of a secrecy order under 35 U.S.C. 181, the licensee may foreign file the application pursuant to 37 CFR 5.15(b).

SelectUSA

The United States represents the largest, most dynamic marketplace in the world and is an unparalleled location for business investment, innovation, and commercialization of new technologies. The U.S. offers tremendous resources and advantages for those who invest and manufacture goods here. Through SelectUSA, our nation works to promote and facilitate business investment. SelectUSA provides information assistance to the international investor community; serves as an ombudsman for existing and potential investors; advocates on behalf of U.S. cities, states, and regions competing for global investment; and counsels U.S. economic development organizations on investment attraction best practices. To learn more about why the United States is the best country in the world to develop technology, manufacture products, deliver services, and grow your business, visit <http://www.SelectUSA.gov> or call +1-202-482-6800.

PATENT APPLICATION FEE DETERMINATION RECORD Substitute for Form PTO-875	Application or Docket Number 16/278,107
---	--

APPLICATION AS FILED - PART I			SMALL ENTITY		OR	OTHER THAN SMALL ENTITY	
	(Column 1)	(Column 2)					
FOR	NUMBER FILED	NUMBER EXTRA	RATE(\$)	FEE(\$)		RATE(\$)	FEE(\$)
BASIC FEE <small>(37 CFR 1.16(a), (b), or (c))</small>	N/A	N/A	N/A	75		N/A	
SEARCH FEE <small>(37 CFR 1.16(k), (j), or (m))</small>	N/A	N/A	N/A	330		N/A	
EXAMINATION FEE <small>(37 CFR 1.16(o), (p), or (q))</small>	N/A	N/A	N/A	380		N/A	
TOTAL CLAIMS <small>(37 CFR 1.16(i))</small>	24	minus 20 = *	x 50 =	200	OR		
INDEPENDENT CLAIMS <small>(37 CFR 1.16(h))</small>	1	minus 3 = *	x 230 =	0.00			
APPLICATION SIZE FEE <small>(37 CFR 1.16(s))</small>	If the specification and drawings exceed 100 sheets of paper, the application size fee due is \$310 (\$155 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).			0.00			
MULTIPLE DEPENDENT CLAIM PRESENT <small>(37 CFR 1.16(j))</small>				0.00			
* If the difference in column 1 is less than zero, enter "0" in column 2.			TOTAL	985		TOTAL	

APPLICATION AS AMENDED - PART II					SMALL ENTITY		OR	OTHER THAN SMALL ENTITY		
	(Column 1)	(Column 2)	(Column 3)							
AMENDMENT A		CLAIMS REMAINING AFTER AMENDMENT	HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA	RATE(\$)	ADDITIONAL FEE(\$)		RATE(\$)	ADDITIONAL FEE(\$)	
	Total <small>(37 CFR 1.16(i))</small>	*	Minus	**	x	=	OR	x	=	
	Independent <small>(37 CFR 1.16(h))</small>	*	Minus	***	x	=	OR	x	=	
	Application Size Fee <small>(37 CFR 1.16(s))</small>							OR		
	FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM <small>(37 CFR 1.16(j))</small>							OR		
					TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE		
AMENDMENT B		CLAIMS REMAINING AFTER AMENDMENT	HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA	RATE(\$)	ADDITIONAL FEE(\$)		RATE(\$)	ADDITIONAL FEE(\$)	
	Total <small>(37 CFR 1.16(i))</small>	*	Minus	**	x	=	OR	x	=	
	Independent <small>(37 CFR 1.16(h))</small>	*	Minus	***	x	=	OR	x	=	
	Application Size Fee <small>(37 CFR 1.16(s))</small>							OR		
	FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM <small>(37 CFR 1.16(j))</small>							OR		
					TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE		
<p>* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.</p> <p>** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".</p> <p>*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".</p> <p>The "Highest Number Previously Paid For" (Total or Independent) is the highest found in the appropriate box in column 1.</p>										

ACCELERATION DATABASE 164					
166	AGENT IP A	ONLINE/OFFLINE			
>>> INDEXED BY: AGENT IP ADDRESS					
CACHE DATABASE 282					
286	LIST OF URIS:				
288	URL 1				
	290	URL			
	292	URL HTTP HEADERS			
	294	LAST CHECKED ON SERVER			
	296	LAST CHANGED ON SERVER			
	298	LIST OF CHUNKS FOR THIS URL:			
		300	CHUNK 1		
			302	CHUNK CHECKSUM	
			304	CHUNK DATA	
			306	LIST OF PEERS:	
				308	PEER 1
					310 PEER 1 IP ADDRESS
					312 PEER 2 CONNECTION STATUS




FIG. 7

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number	16278107
	Filing Date	2019-02-17
	First Named Inventor	Derry Shribman
	Art Unit	
	Examiner Name	
	Attorney Docket Number	HOLA-005-US10

U.S. PATENTS						Remove
Examiner Initial*	Cite No	Patent Number	Kind Code ¹	Issue Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear
	1	7673048		2010-03-02	James W. O'Toole	
	2	7783777		2010-08-24	Kuldipsingh A. Pabla	
	3	8719430		2014-05-06	Michel Van Ackere	
	4	8838811		2014-09-16	Songqing Chen	
	5	7751628	B1	2010-07-06	Richard R. Reisman	
	6	5519693	A	1996-05-21	ROBERT J. GALUSZKA	
	7	6519693	B1	2003-02-11	HENRY C. DEBEY	
	8	7234059	B1	2007-06-19	Cheryl L. Beaver	

If you wish to add additional U.S. Patent citation information please click the Add button. [Add](#)

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number	16278107
Filing Date	2019-02-17
First Named Inventor	Derry Shribman
Art Unit	
Examiner Name	
Attorney Docket Number	HOLA-005-US10

U.S.PATENT APPLICATION PUBLICATIONS						Remove
Examiner Initial*	Cite No	Publication Number	Kind Code ¹	Publication Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear
	1	20150206176	A1	2015-07-23	Assaf Toval	
	2	20020091760	A1	2009-09-08	John Rozen	
	3	20060224687	A1	2006-10-05	Laird Alexander Popkin	
	4	20090248793	A1	2009-10-01	Sanny Jacobsson	
	5	20110035503	A1	2011-02-10	SAM ZAID	
	6	20110087733	A1	2011-04-14	Derry Shribman	
	7	20120124239	A1	2012-05-17	Derry Shribman	
	8	20120166582	A1	2016-06-28	Yehuda BINDER	
	9	20130064370	A1	2013-03-14	Christopher S. Gouge	

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number	16278107
Filing Date	2019-02-17
First Named Inventor	Derry Shribman
Art Unit	
Examiner Name	
Attorney Docket Number	HOLA-005-US10

10	20130080575	A1	2013-03-28	Matthew Browning Prince
11	20060039352	A1	2006-02-23	Christopher K. Karstens
12	20080222291	A1	2008-09-11	Timothy N. Weller
13	20100235438	A1	2010-09-16	Kumar Narayanan
14	20150067819	A1	2015-03-05	Derry Shribman
15	20120254456	A1	2012-10-04	Zubair Visharam
16	20150189401	A1	2015-07-02	Donghoon YI
17	20150341812	A1	2015-11-26	Gino Louis Dion
18	20110264809	A1	2011-10-27	Robert P. Koster

If you wish to add additional U.S. Published Application citation information please click the Add button.

FOREIGN PATENT DOCUMENTS

Examiner Initial*	Cite No	Foreign Document Number ³	Country Code ²ⁱ	Kind Code ⁴	Publication Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear	T ⁵
-------------------	---------	--------------------------------------	----------------------------	------------------------	------------------	---	--	----------------

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number		16278107		
	Filing Date		2019-02-17		
	First Named Inventor	Derry Shribman			
	Art Unit				
	Examiner Name				
	Attorney Docket Number		HOLA-005-US10		

	1	2004094980	WO	2004-11-04	FONTIJN, Wilhelmus, F., J. et al		
--	---	------------	----	------------	-------------------------------------	--	--

If you wish to add additional Foreign Patent Document citation information please click the Add button

NON-PATENT LITERATURE DOCUMENTS

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	
	1		T ⁵

	1		
--	---	--	--

If you wish to add additional non-patent literature document citation information please click the Add button

EXAMINER SIGNATURE

Examiner Signature		Date Considered	
--------------------	--	-----------------	--

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

¹ See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. ² Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). ³ For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. ⁴ Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. ⁵ Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number	16278107
	Filing Date	2019-02-17
	First Named Inventor	Derry Shribman
	Art Unit	
	Examiner Name	
	Attorney Docket Number	HOLA-005-US10

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.

The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Yehuda Binder/	Date (YYYY-MM-DD)	2019-03-11
Name/Print	Yehuda Binder	Registration Number	73,612

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. **DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
4 November 2004 (04.11.2004)

PCT

(10) International Publication Number
WO 2004/094980 A2

- (51) International Patent Classification⁷: **G01M 11/00**
- (21) International Application Number:
PCT/IB2004/050491
- (22) International Filing Date: 22 April 2004 (22.04.2004)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
03101131.5 24 April 2003 (24.04.2003) EP
- (71) Applicant (for all designated States except US): **KONINKLIJKE PHILIPS ELECTRONICS N.V.** [NL/NL]; Groenewoudseweg 1, NL-5621 BA Eindhoven (NL).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **FONTIJN, Wilhelmus, F., J.** [NL/NL]; c/o Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL). **LAMBERT, Nicolaas** [NL/NL]; c/o Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL).
- (74) Agent: **GROENENDAAL, Antonius, W., M.**; Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM,

AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.


(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declaration under Rule 4.17:

— as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii)) for the following designations AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM,

[Continued on next page]

(54) Title: PEER TO PEER TRANSFER OF CONTENT



(57) Abstract: This invention relates to a method, a device, a server and a system of / for peer to peer transfer of content. Said method includes the steps of receiving and transmitting, from a first device (11), a first request with a first selection criterion for a first content to a server (18) or to a second device (12); transferring the first content satisfying said first selection criterion to said first device from the server, when said server previously has acknowledged said first device as a legal recipient of said first content and in case said first content is available only on said server, and noting that said first device subsequently has the requested first content available for other devices (14, 15, 16, 17); or re-directing said first request to a third device (13) on which the server knows that the requested first content is still available and transferring said first content satisfying said first selection criterion to said first device from the third device; or transferring the first content satisfying said first selection criterion to said first device from the second device, when said first content is available on said second device, and informing the server that said first content has been transferred to said first device from said second device; and rewarding the one of said second or third device from which said first content was transferred to said first device, when content was transferred from one of these; and charging said first device for reception of said first content. This enables for download, upload and sharing of legally protected paid-for content.

WO 2004/094980 A2



PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW, ARIPO patent (BW, GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG)

Published:

— without international search report and to be republished upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

Peer to peer transfer of content

This invention relates to a method of peer to peer transfer of content.

The present invention also relates to a computer system for performing the method.

5 The present invention further relates to a computer program product for performing the method.

This invention further relates to a device on which parts of said method is executed.

10 This invention further relates to a server on which remaining parts (not run on the device) of said method is executed.

This invention further relates to a system on which said method is executed.

15 EP 1229443 discloses a system and a method for providing advertisements in a peer to peer networking environment. Each of the advertisements is defined as a structured, language neutral metadata structure. This is used to name, describe and publish an existence of a peer to peer platform resource, such as the peer itself, a pipe or a service. The advertisements are subsequently available to other peers in the networking environment.

20 From the art it is known that Peer-to-peer is a communications model in which each party (i.e. each peer) has the same capabilities and either party can initiate a communication session. Other models with which the pure Peer-to-peer communications model might be contrasted include the *client / server* model and the *master/slave* model, both also known in the art. In some cases, peer-to-peer communications is implemented by giving
25 each communication node both server and client capabilities. In recent usage, peer-to-peer has come to describe applications in which users can use the Internet to download or upload multimedia content or simpler content in form of files with and to each other directly or through a mediating server.

On the Internet, peer-to-peer (referred to as P2P) is a type of transient Internet network that allows a group of computer users (peers) with the same networking program to connect with each other and directly access files from one another's hard drives. Napster and Gnutella are examples of this kind of peer-to-peer software.

5 When the Internet P2P is applied, it is known in the art that the user must first download and execute a peer-to-peer networking program, e.g. Gnutella-net is currently one of the most popular of these decentralized P2P programs because it allows users to exchange all types of files.

10 As discussed later, it is a problem that the files may represent a stolen property right, such as music, a movie, etc, and/or the files may have a poor quality and / or said files may contain virus.

15 After launching the program, the user enters the IP address of another computer belonging to the network, typically, the Web page where the user obtained the download will list several IP addresses as places to begin. Once the computer finds another network member on-line, it will connect to that user's connection, which has obtained their IP address from a connection of another user, and so on.

20 It is however, a problem especially for un-experienced, unaware users that downloadable content typically available in a peer-to-peer network may be legally protected and thus it is illegal do download it and use it, i.e. play back or view said content. In other words, many users – except for the few who know they are deliberately infringing rights of the owner of copy protected content when downloaded – prefer to apply a method and device where they are secure that downloaded content is legal so that they subsequently can play back or view being sure that no rightful proprietor (of said content) is being infringed.

25 It is a further problem for users that downloadable available 'for-free' (in fact stolen from a legal point of view) content can contain virus, i.e. when said content is subsequently played back or viewed on the user's device, the virus may also get life, i.e. it may be executed simultaneous with the playing back or viewing of content on the user's device. Said virus can then consequently harm the file or operating system of the device of the user thus making the user device malfunction or lose previously downloaded content as well.

30 It is a further problem for users that downloadable available 'for-free' content may be in a poor quality, since the content is illegal recorded during a concert, in a cinema or recorded from the original content by means of poor quality recording equipment, thus content in this case is in fact illegal obtained and in a poor quality.

In other words, it is a problem that user are uncomfortable with 'for-free' content since it may contain virus and / or make the user a thief, if the unaware user downloads such content.

5 Additionally, it is a problem that users are reluctant in sharing (i.e. transmitting to others) copy protected content since they risk being caught in infringement of the rightful proprietor, if sharing, especially when using Web-pages (to download content) not telling that eventually provided material (content) from a legal point of view is in fact a violation of copyright laws.

10 It is a further problem that currently the real cost of a downloaded item of content is not transparent to the user.

The present invention enables users to download legally protected content when applying the method according to the invention and / or when using the device according to the invention which communicates with the server in the peer-to-peer network. Subsequently, it is legal to play it back, view it and share it with others. This is possible, since the method (and the device and server applying the method) handles the property rights and the payments in a legal manner, which both the users and content providers are comfortable with, i.e. the user is assured that he does not make a thief out of himself, and the content providers (artist, singer, movie manufacturer, etc) are assured that their content is not being stolen, but paid for.

20 Further, it is assured that the downloaded content is virus-free and in an approved quality.

25 Additionally, users can - when applying the method by means of their device - easily and legally share (i.e. transmit to others) copy protected content since some steps of the methods ensure that the proprietor of content gets paid for his content, since users are charged for downloads. Further, users (of said devices), themselves can obtain a reward for sharing, this further expands sharing.

Further, it is an advantage of the invention that the real cost of a downloaded item of content is transparent to the user.

30 Said device and server, in combination and the system provide the same advantages and solve the same problem(s) for the same reasons as described previously in relation to the method.

The invention will be explained more fully below in connection with preferred embodiments and with reference to the drawings, in which:

fig. 1 shows a network of devices and a server; and

fig. 2 shows a method of peer to peer transfer of content.

5 Throughout the drawings, the same reference numerals indicate similar or corresponding features, functions, etc.

Figure 1 shows a network of devices and a server. Said network of devices
10 with the server are illustrated by means of reference numeral 10. As will be explained more detailed in the next figure, a first device, reference numeral 11, or its user is looking for certain content (a video film as an example), the user will then try to find out from where the video film can be obtained, i.e. downloaded. He will use a specific selection criterion for the video film content. In technical terms, his device (first device) will receive the selection
15 criterion e.g. movie name, genre, etc, which it then will send to another device (a second device. reference numeral 12,)) and to a server, reference numeral 18, since his own device (said first device) cannot know whether the server or another peer to peer device, has the requested content available. If the server has the content satisfying the selection criterion, it will provide it to the requesting device, i.e. to said first device. However, in order to offload and distribute network usage more efficient – if the server knows that another peer (device)
20 has the requested content available, the server will redirect the transfer of content to this device which then will provide the content satisfying the selection criterion, i.e. transfer it to the requesting device equalling said first device. In the last case, the server is informed – by the actual device transferring content that content has been transferred to said first device,
25 which then can be accordingly charged for receiving the requested content. Hereby, the first device (and its user) is comfortable with content charged for, since it is virus free and has been legally bought, i.e. the user is sure that he did not make a thief out of himself; further the user can rely on that the content has an approved quality level, since it comes form the legal owner or an administrator of network, he can trust.

30 In the first case, i.e. the server supplied directly the requested content, the server typically previously acknowledged that said first device is in fact a subscribing or paying (or one who later will pay) rightful recipient of the content, i.e. said video film. The content, in general, can be uploaded to or downloaded from more devices, e.g. reference numerals 13 and 14. In the network further devices may be present, e.g. reference numerals

15, 16 and 17. Generally, the server has to be accessible to and in the network of devices, i.e. to all devices, either for transfer of content the first time, and/ or subsequently for charging and rewarding, this is illustrated by means of the arrows connecting the server to the devices.

5 A requester needs not register or be registered to the server. There may be a third party that certifies the requester to the server. The server trusts the certifier and assumes the requester is allowed to receive. Or the requester pays 'on-the-spot' using virtual tokens or a mediation service (Pay-Pal).

10 The network is shown for illustrative purposes, any other dynamic or static topology or arrangement of peers or devices and one or more additionally servers may also be applied in the present invention.

Any of said devices may be a video cassette recorder (VCR), a personal digital assistant (PDA), a mobile phone, a television, a radio, a DVD player, a CD player, an information panel, a web tablet, a smart remote, a peer or a personal computer.

15 The device alternatives as mentioned may be understood as corresponding peers in a peer-to-peer type of transient network similar to the type found on the Internet, that allows a group of computer users (with access to their corresponding peers or devices) with the same or similar networking program or protocol to connect with each other and directly access content, e.g. in the form of files, etc to/from one another's hard drives, memories, etc.

20 A peer-to-peer network is simply a network of peers, the Internet, Gnutella software, computers are all just examples of aspects of specific implementations, however the present invention applies said server for rewarding direct peer to peer content sharing, and said server is furthermore applied to charge peers for download of content. Since content typically is copy protected content, at least one of said servers tracks, charges and rewards peers (devices) for down and upload, respectively of copy protected content.

25 In a preferred embodiment of the invention said content comprises one or more selected from the group:

- a DVD picture and sound signal;
- a CD sound signal;
- a given digital audio format (e.g. MP3, WMA, Real Audio, WAV, etc);
- 30 - a given digital movie format (e.g. DivX, DVD/MPEG2, Avi, wmf, MOV, Real Video, etc);
- a given picture format (e.g. JPEG, GIF, BMP, TIFF, etc); and/or
- any such format that is capable of causing the device to emit a picture and/or sound signal, e.g. G72x, aiff, real.

This is possible since said device can be a CD player, a DVD player, a radio, a mobile phone, etc. as discussed, accordingly content can be presented, i.e. shown and/ or played back on said device.

5 In other words, the above content combinations are copy protected content, which, generally, are in the form of numerical, textual information, picture, video, sound and / or any combination(s) thereof, and which, generally, also are being free from virus and in an approved quality.

Figure 2 shows a method of peer to peer transfer of content. The content is transferred among device in the peer to peer network, in initial situations, i.e. the first times
10 content gets available, by means of the server.

Prior to the following steps, it is assumed that - as a starting point - that only the server can provide content; later on content can be distributed (or spread) to various devices (second, third, etc.) i.e. at later occasions these devices can provide content without directly involving the server, however, still devices requesting and receiving content are
15 charged accordingly regardless from where (i.e. from the server or from the peer to peer device) said content is being transferred.

Further, content is copy protected content, i.e. legal content being free from virus and in an approved quality. The server is in all cases – also when content is transferred directly between devices – responsible for that the copy protected content is legal, free from
20 virus and in the proper quality, this is possible since – from the starting point - content can only be introduced into the network via the server. The actual (content) data does not have to originate from the server. The server just needs to certify it. Any user may offer a piece of content to the server for certification. On the server side the content will be checked and when it is found to be acceptable, the content is certified, for download, redirection, etc. In
25 step 100, a first request may be received on a first device. The request typically comprises a first selection criterion for a first content, the user of the first device can e.g. key in his selection criterion for the content by means of a keyboard or by means of any common user interface know in the art, e.g. a GUI like windows, soft-keys, menu driven, click by means of a mouse, etc. The content may reside on the server and / or another second device, i.e. said
30 second device. Therefore, subsequently the request is transmitted from the first device to the server or to said second device, since said first device cannot know whether the server or another peer to peer (second) device has said requested first content available.

Said first selection criterion may be composed by means of one or more combinable items, e.g. program, channel, Web-site, genre, type, topic, style, start, duration, language, title, name, hyperlink including content reference, etc.

5 Said first selection criterion can then be helpful for the user and to the device from which content is requested, i.e. helpful to find and subsequently transmit content having the first selection criterion, i.e. said selection criterion in general may reflect content interest(s) of a specific user, the user can therefore avoid to surf through many available channels if the device is a TV, or surf through many Web sites if his device is an Internet PC or accessed via a server from a client PC in order to find his content. The user can apply this
10 step instead.

Said selection criterion can therefore be understood as the users own profile of interests.

As discussed, the request is transmitted to the server or to another, i.e. the second device; in general, requests are put to the network (of devices or peers) as a whole
15 including the server(s) as stated in this step. Although, due to the nature of the network, the requests will not reach all peers in the network, they should reach at least one server, e.g. via a Kazaa like super-node that is or knows a server. In the beginning the server will only have the content available and participate in transferring the content to the requester, here said first device. If a certain number of peers have downloaded the content, the server may stop
20 offering it because it will be available from elsewhere, i.e. from said number of peers. This is in fact dealt with by means of steps 200, 300 and 400.

In step 200, the first content satisfying said first selection criterion may be transferred to said first device from the server. This is only in the case when said server previously has acknowledged said first device as a legal recipient of said first content, e.g.
25 through an eventually registration, and when said first content is available only on said server.

Subsequently, the server will note that said first device now has the requested first content available for the other devices. This implies that if the same request (for content) arrives again to the server, the first device will then be the direct content supplier instead of
30 the server. The latter – in fact redirecting of content - is dealt with in step 300.

Alternatively, instead of step 200, in step 300, said first request is redirected to a third device. Said third device is known to the server as a device in fact still having the requested first content. Subsequently, said first content satisfying said first selection criterion is transferred to said first device from the device re-directed to, i.e. from said third device.

The server will currently check that said content in fact still is on the third device, in case the user of the third device removes or removed the particular content, the server will subsequently find out. In that case, the server must provide content it self or redirect the request (for content) to another fourth device (in place of said third device). In other words, the server currently checks that content is in fact still available on said third device, and that said third device is on line, if not, the request is redirected to another, i.e. to said fourth device, etc.

Alternatively, instead of step 200 or step 300, in step 400, first content satisfying said first selection criterion is transferred to said first device from the second device. This is only the case when said first content is in fact available on said second device; in that case the server is subsequently informed (by said second device) that said first content has actually been transferred to said first device. The reason for doing the latter is to enable the server to charge said first device for receiving content, in fact requested by it self. Conversely – as in next step - to enable the server to reward said second device for transferring (and sharing) content.

It is assumed that when any device (second or third) provides or supplies content, the content, in all cases, initially came from the server or is at least approved from the server to legally be available from the other device(s) (second or third) for an eventually subsequent transfer. At later occasions, one of the other devices (second or third) devices can provide content (originally legally approved by the server, etc) to even more devices. Further, after reception of content on the first device, this also can play the role of ‘content provider’, i.e. acting in the same manner as said second and third devices; in fact when more devices have received the same content (satisfying the same criterion) any of these - of course – play the role of ‘content provider’ in competition with other devices having the same content, this lowers the waiting time for a requesting device and provides for an improvement in sharing of content among devices, this in turn also offloads the server.

Generally, in step 200, 300 and 400, the server, the third device and the second device, respectively transferred content to said first device.

In step 500, said second or third device, which in fact transferred content to said first device, is then rewarded. However, it may be the case that the server transferred content itself; in this case none, i.e. neither the second nor the third device are rewarded. However, in the general case, the second or third device is rewarded; conversely, the more rare case, i.e. the server transferred content, it will not reward itself, but it may note the transfer primarily for statistical purposes.

In all cases, regardless from where (server, second or third device) content was transferred, i.e. in step 600, said first device is charged for reception of said first content. The charge may be dependent of a subscription fee or subscription agreements in general or on a per transfer basis (download); it may be dependent on a file length, value or duration, and / or combinations thereof. This is possible since content may be transferred embedded in or by means of said file.

Optionally, said method comprises the following two steps, which deals with the opposite situation, i.e. the server receives content:

In step 700, a second content satisfying a second selection criterion and the second selection criterion are uploaded to the server from a fourth device. The server should then subsequently ensure that said second content is free for virus, has the right quality level (sufficient high sampling rate, low noise, stereo, aliasing, etc) and, most importantly, is legal, for the latter the owner of the server may have agreed contracts (e.g. through licence, partly or in whole, an exclusive right, etc) with the original creator, owner or supplier of content to ensure that it can be legally distributed afterwards as discussed in the steps above. The second selection criterion is uploaded with the corresponding second content in order to make said second content searchable again, when requested as discussed in step 100. The second selection criterion will be of the same nature and structure as that of said first selection criterion.

In step 800, the fourth device is rewarded. The reward is given to the fourth device in return for uploading said second content (with the second selection criterion) to the server. The reward may be given in form of credits, rebate, discounts, etc. The reward can then be used by said fourth device, if it later obtains a third content, etc.

Generally when the device is denoted first, second, etc device, it is to be understood that any device can perform the mentioned tasks, i.e. even though a first device, only as disclosed in the above steps requests content, it - as well as the other devices - may perform any task as reflected in the steps above.

Rewards, credits, rebate, discounts, the task of charging are generally dealt with by the server, i.e. the server keeps a balance of in and outgoing payments for each device up and downloading content.

As discussed above, for or each item of content the device has to pay a small fee. When a device is charged, a subsequent payment can be done on a transaction basis or included in telecommunication fees. The latter can be in the form of an elevated rate (price/minute) for the transfer or included in a periodic bill. Subscription is also an option.

Part of the fee is direct payment to the content provider, which may be represented by said server. Part of the fee is used to award a discount to the device offering the content. I.e. users of devices can recap part of that fee by sharing desirable content.

5 For each piece of content that is downloaded from a device (to the server or to another device), the device is rewarded with credits. These credits can be in the form of rebates on the purchase of new songs, on telecommunication fees or on downloading content from other devices. The credits can be proportional to the amount of data transferred, e.g. the size of the file, or proportional to the value of the song.

10 The content shared by devices is verified by the server. Devices offering non-compliant, e.g. sub standard content can be excluded from the exchange based on the identification of the mobile phone identification, i.e. not satisfying the criterion in step 200 of acknowledgement.

15 In general, according to the present invention, a service for sanctioned P2P transfer between devices is set up. Peers or devices who want to share content are registered at the server and the content they offer may indexed, e.g. the Napster model.

20 The server may offer a comprehensive collection of content. This can be done using an intuitive interface for the selection of content. The offering of content can be enhanced by supporting information. If certain content is not offered by any peer (device), e.g. very new content, the server may offer the content. The latter is a temporary measure till (enough) peers (devices) offer the content. This amounts to a transition model. Initially most content may be hosted by the server but few peers will use the redirection service. If the amount of connected peers in the network grows the demand on the redirection service will increase but at the same time the amount of content provided at the server side can decrease. Hence, if the popularity (and therefore the use) of the system increases the server will not
25 have to be scaled up.

The transfer rate of content shared by peers is not guaranteed. This enables the definition of a lazy transfer mode to offer unused bandwidth at reduced price. If the premium service of voice communication uses more of the networks bandwidth, the bandwidth available to P2P transfers is reduced.

30 A computer readable medium may be magnetic tape, optical disc, digital versatile disk (DVD), compact disc (CD record-able or CD write-able), mini-disc, hard disk, floppy disk, smart card, PCMCIA card, etc.

In the claims, any reference signs placed between parentheses shall not be constructed as limiting the claim. The word "comprising" does not exclude the presence of

elements or steps other than those listed in a claim. The word "a" or "an" preceding an element does not exclude the presence of a plurality of such elements.

The invention can be implemented by means of hardware comprising several distinct elements, and by means of a suitably programmed computer. In the device claim
5 enumerating several means, several of these means can be embodied by one and the same item of hardware. The mere fact that certain measures are recited in mutually different dependent claims does not indicate that a combination of these measures cannot be used to advantage.

CLAIMS:

1. A method of peer to peer transfer of content, said method comprising the steps of:
- receiving and transmitting (100), from a first device (11), a first request with a first selection criterion for a first content to a server (18) or to a second device (12);
- 5 transferring (200) the first content satisfying said first selection criterion to said first device from the server, when said server previously has acknowledged said first device as a legal recipient of said first content and in case said first content is available only on said server, and noting that said first device subsequently has the requested first content available for other devices (14 , 15 , 16, 17); or
- 10 re-directing (300) said first request to a third device (13) on which the server knows that the requested first content is still available and transferring said first content satisfying said first selection criterion to said first device from the third device; or
- transferring (400) the first content satisfying said first selection criterion to said first device from the second device, when said first content is available on said second
- 15 device, and informing the server that said first content has been transferred to said first device from said second device; and
- rewarding (500) the one of said second or third device from which said first content was transferred to said first device, when content was transferred from one of these; and
- 20 charging (600) said first device for reception of said first content.
2. A method according to claim 1, said method further comprising the steps of:
- uploading (700) a second content satisfying a second selection criterion and the second selection criterion to the server from a fourth device; and
- 25 rewarding (800) the fourth device for uploading the second content and the second criterion to the server.

3. A method according to claim 1 or 2, characterized in that said content is copy protected content, such as numerical information, picture, video, sound and combinations thereof.
- 5 4. A method according to any one of claims 1 through 3, characterized in that said content comprises one or more selected from the group:
- a DVD picture and sound signal;
 - a CD sound signal;
 - a given digital audio format (e.g. MP3, WMA, Real Audio, WAV, etc);
 - 10 a given digital movie format (e.g. DivX, DVD/MPEG2, Avi, wmf, MOV, Real Video, etc);
 - a given picture format (e.g. JPEG, GIF, BMP, TIFF, etc); and/or
 - any such format that is capable of causing the device to emit a picture and/or sound signal, e.g. G72x, aiff, real.
- 15 5. A method according to any one of claims 1 through 4, characterized in that any of said devices is a video cassette recorder (VCR), a personal digital assistant (PDA), a mobile phone, a television, a radio, a DVD player, a CD player, an information panel, a web tablet, a smart remote, a peer or a personal computer.
- 20 6. A device comprising:
- means for receiving and transmitting a first request with a first selection criterion for a first content to a server (18) or to a second device (12);
 - means for receiving a redirected said first request (13) on which the server
 - 25 knows that the requested first content is still available on said device;
 - means for transferring the first content satisfying said first selection criterion to a first device, when said first content is available on said device, and means for informing the server that said first content has been transferred to said first device;
 - means for being rewarded for transfer of content; and
 - 30 means for being charged for reception of content.
7. A device according to claim 6 further comprising:
- means for uploading a second content satisfying a second selection criterion and the second selection criterion to the server; and

means for being rewarded for the upload of the second content and the second selection criterion to the server.

8. A server comprising:

5 means for receiving a first request with a first selection criterion for a first content;

means for transferring the first content satisfying said first selection criterion to a first device, when said server previously has acknowledged said first device as a legal recipient of said first content and in case said first content is available only on said server, and means for noting that said first device subsequently has the requested first content available for other devices (14 , 15 , 16, 17); and / or

means for re-directing said first request to a third device (13) on which the server knows that the requested first content is still available; and / or

15 means for being informed that said first content has been transferred to said first device from said third device;

means for rewarding the one of said second or third device from which said first content was transferred to said first device, when content was transferred from one of these; and

20 means for charging said first device for reception of said first content.

9. A server according to claim 8 further comprising:

means for being uploaded with a second content satisfying a second selection criterion and means for being uploaded with the second selection criterion from a fourth device; and

25 means for rewarding the fourth device for uploading the second content and the second criterion.

10. A system comprising:

30 means for receiving and transmitting, from a first device (11), a first request with a first selection criterion for a first content to a server (18) or to a second device (12);

means for transferring the first content satisfying said first selection criterion to said first device from the server, when said server previously has acknowledged said first device as a legal recipient of said first content and in case said first content is available only

on said server, and noting that said first device subsequently has the requested first content available for other devices (14 , 15 , 16, 17);

- 5 means for re-directing said first request to a third device (13) on which the server knows that the requested first content is still available and transferring said first content satisfying said first selection criterion to said first device from the third device;
- means for transferring the first content satisfying said first selection criterion to said first device from the second device, when said first content is available on said second device, and informing the server that said first content has been transferred to said first device from said second device;
- 10 means for rewarding the one of said second or third device from which said first content was transferred to said first device, when content was transferred from one of these; and
- means for charging said first device for reception of said first content.

- 15 11. A system according to claim 10 further comprising:
- means for uploading a second content satisfying a second selection criterion and the second selection criterion to the server from a fourth device; and
- means for rewarding the fourth device for uploading the second content and the second criterion to the server.

- 20 12. A computer system for performing the method according to any one of claims 1 through 5.

- 25 13. A computer program product comprising program code means stored on a computer readable medium for performing the method of any one of claims 1 through 5 when the computer program is run on a computer.




FIG. 1

2/2




FIG. 2

Electronic Acknowledgement Receipt

EFS ID:	35386159
Application Number:	16278107
International Application Number:	
Confirmation Number:	4936
Title of Invention:	SYSTEM PROVIDING FASTER AND MORE EFFICIENT DATA COMMUNICATION
First Named Inventor/Applicant Name:	Derry Shribman
Customer Number:	131926
Filer:	Yehuda Binder/Dorit Binder
Filer Authorized By:	Yehuda Binder
Attorney Docket Number:	HOLA-005-US10
Receipt Date:	11-MAR-2019
Filing Date:	17-FEB-2019
Time Stamp:	16:49:37
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no
------------------------	----

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Preliminary Amendment	Preliminary-Amendment.pdf	15951 <small>d5358969e33296a0cb15b54ba4a5dd66b509fc40</small>	no	2

Warnings:

Information:					
2	Drawings-only black and white line drawings	Corrected_FIG7.pdf	52370	no	1
			0283926ae8105b215daf4b05d7c0285e6550266a		
Warnings:					
The page size in the PDF is too large. The pages should be 8.5 x 11 or A4. If this PDF is submitted, the pages will be resized upon entry into the Image File Wrapper and may affect subsequent processing					
Information:					
3	Information Disclosure Statement (IDS) Form (SB08)	IDS5.pdf	1035536	no	6
			20ec5cb2d09da6e90d47073974d8259eb8b3b0b1		
Warnings:					
Information:					
4	Foreign Reference	WO2004094980.pdf	879969	no	19
			192e165015d1398f91d038c47d179e3486f6c135		
Warnings:					
Information:					
Total Files Size (in bytes):			1983826		
<p>This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.</p> <p><u>New Applications Under 35 U.S.C. 111</u> If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.</p> <p><u>National Stage of an International Application under 35 U.S.C. 371</u> If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.</p> <p><u>New International Application Filed with the USPTO as a Receiving Office</u> If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.</p>					

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

ATTY.'S DOCKET: HOLA-005-US10

In re Application of:) Confirmation No. 4936
Derry SHRIBMAN) Art Unit:
Appln. No.: 16/278,107) Examiner:
Filed: February 17, 2019) Washington, D.C.
For: SYSTEM PROVIDING FASTER)
AND MORE EFFICIENT DATA)
COMMUNICATION) March 11, 2019

PRELIMINARY AMENDMENT:

Honorable Commissioner for Patents
U.S. Patent and Trademark Office
Randolph Building, Mail Stop Amendments
401 Dulany Street
Alexandria, VA 22314

Sir:

Amendments to the drawings begin on page 2 of this
paper.

Appln. No. 16/278,107
Filed February 17, 2019

Amendments to the drawings

Submitted herewith is Figure 7 wherein the line quality has been improved.

No new matter was added.

Respectfully submitted,

By / Yehuda Binder /
 Yehuda Binder
 Registration No. 73,612

Tel: +972-54-4444577
Fax: +972-9-7442619
e-mail: yehuda@maypatents.com

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number	16278107
	Filing Date	2019-02-17
	First Named Inventor	Derry Shribman
	Art Unit	
	Examiner Name	
	Attorney Docket Number	HOLA-005-US10

U.S.PATENTS						Remove
Examiner Initial*	Cite No	Patent Number	Kind Code ¹	Issue Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear
	1	6895011	B1	2005-05-17	Harold Aaron Lassers	

If you wish to add additional U.S. Patent citation information please click the Add button.

Add

U.S.PATENT APPLICATION PUBLICATIONS						Remove
Examiner Initial*	Cite No	Publication Number	Kind Code ¹	Publication Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear
	1	20090232003	A1	2009-09-17	Jean-Philippe Vasseur	
	2	20140189802	A1	2014-07-03	Gregory MONTGOMERY	
	3	20030229785	A1	2003-12-11	Michael J. Daseke	
	4	20050027782	A1	2005-02-03	Rajkumar Jalan	
	5	20030229718	A1	2003-12-11	Theron Tock	

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number	16278107
Filing Date	2019-02-17
First Named Inventor	Derry Shribman
Art Unit	
Examiner Name	
Attorney Docket Number	HOLA-005-US10

6	20090292816	A1	2009-11-26	Craig S. Etchegoyen
7	20120166582	A1	2012-06-28	Yehuda BINDER

If you wish to add additional U.S. Published Application citation information please click the Add button.

FOREIGN PATENT DOCUMENTS

Examiner Initial*	Cite No	Foreign Document Number ³	Country Code ² i	Kind Code ⁴	Publication Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear	T ⁵
	1							

If you wish to add additional Foreign Patent Document citation information please click the Add button.

NON-PATENT LITERATURE DOCUMENTS

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T ⁵
	1	Reed et al, "Anonymous Connections and Onion Routing", Naval Research Laboratory, 03/1998 https://www.onion-router.net/Publications/JSAC-1998.pdf (Year: 1998)	

If you wish to add additional non-patent literature document citation information please click the Add button.

EXAMINER SIGNATURE

Examiner Signature	<input type="text"/>	Date Considered	<input type="text"/>
--------------------	----------------------	-----------------	----------------------

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

¹ See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. ² Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). ³ For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. ⁴ Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. ⁵ Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number	16278107
	Filing Date	2019-02-17
	First Named Inventor	Derry Shribman
	Art Unit	
	Examiner Name	
	Attorney Docket Number	HOLA-005-US10

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.

The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Yehuda Binder/	Date (YYYY-MM-DD)	2019-04-23
Name/Print	Yehuda Binder	Registration Number	73,612

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. **DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Anonymous Connections and Onion Routing

Michael G. Reed, Paul F. Syverson, and David M. Goldschlag *
Naval Research Laboratory

Abstract

Onion Routing is an infrastructure for private communication over a public network. It provides anonymous connections that are strongly resistant to both eavesdropping and traffic analysis. Onion routing's anonymous connections are bidirectional and near real-time, and can be used anywhere a socket connection can be used. Any identifying information must be in the data stream carried over an anonymous connection. An onion is a data structure that is treated as the destination address by onion routers; thus, it is used to establish an anonymous connection. Onions themselves appear differently to each onion router as well as to network observers. The same goes for data carried over the connections they establish. Proxy aware applications, such as web browsing and e-mail, require no modification to use onion routing, and do so through a series of proxies. A prototype onion routing network is running between our lab and other sites. This paper describes anonymous connections and their implementation using onion routing. This paper also describes several application proxies for onion routing, as well as configurations of onion routing networks.

1 Introduction

Is Internet communication private? Most security concerns focus on preventing eavesdropping [18], i.e., outsiders listening in on electronic conversations. But encrypted messages can still be tracked, revealing who is talking to whom. This tracking is called traffic analysis and may reveal sensitive information. For example, the existence of inter-company collaboration may be confidential. Similarly, e-mail users may not wish to

*Address: (For Reed and Syverson) Naval Research Laboratory, Center For High Assurance Computer Systems, Washington, D.C. 20375-5337, USA, phone: +1 202.767.2389, fax: +1 202.404.7942, e-mail: {last_name}@itd.nrl.navy.mil. (For Goldschlag) Divx, 570 Herndon Parkway, Herndon, VA 20170, USA, phone: +1 703-708-4028, fax: +1 703-708-4088, e-mail: david.goldschlag@divx.com

reveal who they are communicating with to the rest of the world. In certain cases anonymity may be desirable also: anonymous e-cash is not very anonymous if delivered with a return address. Web based shopping or browsing of public databases should not require revealing one's identity.

This paper describes how a freely available system, *onion routing*, can be used to protect a variety of Internet services against both eavesdropping and traffic analysis attacks, from both the network and outside observers. This paper includes a specification sufficient to guide both re-implementations and new applications of onion routing. We also discuss configurations of onion routing networks and applications of onion routing, including Virtual Private Networks (VPN), Web browsing, e-mail, remote login, and electronic cash.¹

A purpose of traffic analysis is to reveal who is talking to whom. The *anonymous connections* described here are designed to be resistant to traffic analysis, i.e., to make it difficult for observers to learn identifying information from the connection (e.g., by reading packet headers, tracking encrypted payloads, etc.). Any identifying information must be passed as data through the anonymous connections. Our implementation of anonymous connections, onion routing, provides protection against eavesdropping as a side effect. Onion routing provides bidirectional and near real-time communication similar to TCP/IP socket connections or ATM AAL5 [6]. The anonymous connections can substitute for sockets in a wide variety of unmodified Internet applications by means of proxies. Data may also be passed through a privacy filter before being sent over an anonymous connection. This removes identifying information from the data stream, to make communication anonymous too.

Although onion routing may be used for anonymous communication, it differs from anonymous remailers [7, 16] in two ways: Communication is real-time and bidirectional, and the anonymous connections are application independent. Onion routing's anonymous

¹Preliminary versions of portions of this paper have appeared in [28, 14, 24].

connections can support anonymous mail as well as other applications. For example, onion routing may be used for anonymous Web browsing. A user may wish to browse public Web sites without revealing his identity to those Web sites. That requires removing information that identifies him from his requests to Web servers and removing information from the connection itself that may identify him. Hence, anonymous Web browsing uses anonymized communication over anonymous connections. The Anonymizer [1] only anonymizes the data stream, not the connection itself. So it does not prevent traffic analysis attacks like tracking data as it moves through the network.

This paper is organized in the following way: Section 2 presents an overview of onion routing. Section 3 presents empirical data about our prototype. Section 4 defines our threat model. Section 5 describes onion routing and the application specific proxies in more detail. Section 6 describes the implementation choices that were made for security reasons. Section 7 describes how onion routing may be used in a wide variety of Internet applications. Section 8 contrasts onion routing with related work, and section 9 presents concluding remarks.

2 Onion Routing Overview

In onion routing, instead of making socket connections directly to a responding machine, initiating applications make connections through a sequence of machines called *onion routers*. The *onion routing network* allows the connection between the *initiator* and *responder* to remain anonymous. Anonymous connections hide who is connected to whom, and for what purpose, from both outside eavesdroppers and compromised onion routers. If the initiator also wants to remain anonymous to the responder, then all identifying information must be removed from the data stream before being sent over the anonymous connection.

Onion routers in the network are connected by long-standing (permanent) socket connections. Anonymous connections through the network are multiplexed over the longstanding connections. For any anonymous connection, the sequence of onion routers in a route is strictly defined at connection setup. However, each onion router can only identify the previous and next hops along a route. Data passed along the anonymous connection appears different at each onion router, so data cannot be tracked en route, and compromised onion routers cannot cooperate by correlating the data stream each sees. We will also see that they cannot make use of replayed onions or replayed data.

2.1 Operational Overview

The onion routing network is accessed via a series of *proxies*. An initiating application makes a socket connection to an *application proxy*. This proxy massages connection message format (and later data) to a generic form that can be passed through the onion routing network. It then connects to an *onion proxy*, which defines a route through the onion routing network by constructing a layered data structure called an *onion*. The onion is passed to the *entry funnel*, which occupies one of the longstanding connections to an onion router and multiplexes connections to the onion routing network at that onion router. That onion router will be the one for whom the outermost layer of the onion is intended. Each layer of the onion defines the next hop in a route. An onion router that receives an onion peels off its layer, identifies the next hop, and sends the embedded onion to that onion router. The last onion router forwards data to an *exit funnel*, whose job is to pass data between the onion routing network and the responder.

In addition to carrying next hop information, each onion layer contains key seed material from which keys are generated for crypting² data sent forward or backward along the anonymous connection. (We define *forward* to be the direction in which the onion travels and *backward* as the opposite direction.)

Once the anonymous connection is established, it can carry data. Before sending data over an anonymous connection, the onion proxy adds a layer of encryption for each onion router in the route. As data moves through the anonymous connection, each onion router removes one layer of encryption, so it arrives at the responder as plaintext. This layering occurs in the reverse order for data moving back to the initiator. So data that has passed backward through the anonymous connection must be repeatedly post-crypted to obtain the plaintext.

By layering cryptographic operations in this way, we gain an advantage over link encryption. As data moves through the network it appears different to each onion router. Therefore, an anonymous connection is as strong as its strongest link, and even one honest node is enough to maintain the privacy of the route. In link encrypted systems, compromised nodes can cooperate to uncover route information.

Onion routers keep track of received onions until they expire. Replayed or expired onions are not forwarded, so they cannot be used to uncover route information, either by outsiders or compromised onion

²We define the verb *crypt* to mean the application of a cryptographic operation, be it encryption or decryption.

routers. Note that clock skew between onion routers can only cause an onion router to reject a fresh onion or to keep track of processed onions longer than necessary. Also, since data is encrypted using stream ciphers, replayed data will look different each time it passes through a properly operating onion router.

Although we call this system onion routing, the routing that occurs here does so at the application layer of the protocol stack and not at the IP layer. More specifically, we rely upon IP routing to route data passed through the longstanding socket connections. An anonymous connection is comprised of portions of several linked longstanding multiplexed socket connections. Therefore, although the series of onion routers in an anonymous connection is fixed for the lifetime of that anonymous connection, the route that data actually travels between individual onion routers is determined by the underlying IP network. Thus, onion routing may be compared to loose source routing.

Onion routing depends upon connection based services that deliver data uncorrupted and in-order. This simplifies the specification of the system. TCP socket connections, which are layered on top of a connectionless service like IP, provide these guarantees. Similarly, onion routing could easily be layered on top of other connection based services, like ATM AAL5.

Our current prototype of onion routing considers the network topology to be static and does not have mechanisms to automatically distribute or update public keys or network topology. These issues, though important, are not the key parts of onion routing and will be addressed in a later prototype.

2.2 Configurations

As mentioned above neighboring onion routers are neighbors in virtue of having longstanding socket connections between them, and the network as a whole is accessed from the outside through a series of proxies. By adjusting where those proxies reside it is possible to vary which elements of the system are trusted by users and in what way. (For some configurations it may be efficient to combine proxies that reside in the same place, thus they may be only conceptually distinct.)

2.2.1 Firewall Configuration

In the *firewall configuration*, an onion router sits on the firewall of a sensitive site. This onion router serves as an interface between machines behind the firewall and the external network. Connections from machines behind the firewall to the onion router are protected by other means (e.g., physical security). To complicate

tracking of traffic originating or terminating within the sensitive site, this onion router should also route data between other onion routers. This configuration might represent the system interface from a typical corporate or government site. Here the application proxies (together with any privacy filters), and the onion proxies would typically live at the firewall as well. (Typically, there might only be one onion proxy.)

There are three important features of this basic configuration:

- Connections between machines behind onion routers are protected against both eavesdropping and traffic analysis. Since the data stream never appears in the clear on the public network, this data may carry identifying information, but communication is still private. (This feature is used in section 7.1.)
- The onion router at the originating protected site knows both the source and destination of a connection. This protects the anonymity of connections from observers outside the firewall but also simplifies enforcement of and monitoring for compliance with corporate or governmental usage policy.
- The use of anonymous connections between two sensitive sites that both control onion routers effectively hides their communication from outsiders. However, if the responder is not in a sensitive site (e.g., the responder is some arbitrary Web server) the data stream from the sensitive initiator must also be anonymized. If the connection between the exit funnel and the responding server is unencrypted, the data stream might otherwise identify the initiator. For example, an attacker could simply listen in on the connections to a Web server and identify initiators of any connection to it.

2.2.2 Remote Proxy Configuration

What happens if an initiator does not control an onion router? If the initiator can make encrypted connections to some remote onion router, then he can function as if he is in the firewall configuration just described, except that both observers and the network can tell when he makes connections to the onion router. However, if the initiator trusts the onion router to build onions, his association with the anonymous connection from that onion router to the responder is hidden from observers and the network. In a similar way, an encrypted connection from an exit funnel to a responder hides the

association of the responder with the anonymous connection .

Therefore, if an initiator makes an anonymous connection to some responder, and layers end-to-end encryption over that anonymous connection, the initiator and responder can identify themselves to one another, yet hide their communication from the rest of the world. So we can build virtual private networks without protected sites.

Notice, however, that the initiator trusts the remote onion router to conceal that the initiator wants to communicate with the responder, and to build an anonymous connection through other onion routers. The next section describes how to shift some of this trust from the first onion router to the initiator.

2.2.3 The Customer-ISP Configuration

Suppose, for example, an Internet Services Provider (ISP) runs a funnel that accepts connections from onion proxies running on subscribers' machines. In this configuration, users generate onions specifying a path through the ISP to the destination. Although the ISP would know who initiates the connection, the ISP would not know with whom the customer is communicating, nor would it be able to see data content. So the customer need not trust the ISP to maintain her privacy. Furthermore, the ISP becomes a *common carrier*, who carries data for its customers. This may relieve the ISP of responsibility both for whom users are communicating with and the content of those conversations. The ISP may or may not be running an onion router as well. If he is running an onion router, then it is more difficult to identify connections that terminate with his customers; however, he is serving as a routing point for other traffic. On the other hand, if he simply runs a funnel to an onion router elsewhere, it will be possible to identify connections terminating with him, but his overall traffic load will be less. Which of these would be the case for a given ISP would probably depend on a variety of service, cost, and pricing considerations. Note that in this configuration the entry funnel must have an established longstanding connection to an onion router just like any neighboring onion router. (Cf. section 5.6 for a description of how these are established.) But, in most other cases, where the funnel resides on the same machine as the onion router, establishing an encrypted longstanding connection should not be necessary since the funnel can be directly incorporated into the onion router.

3 Empirical Data

We invite readers to experiment with our prototype of onion routing by using it to anonymously surf the Web, send anonymous e-mail, and do remote logins. For instructions please see <http://www.itd.nrl.navy.mil/ITD/5540/projects/onion-routing/>.

One should be aware that accessing a remote onion router does not completely preserve anonymity, because the connection between a remote machine and the first onion router is not protected. If that connection were protected, one would be in the remote proxy configuration, but there would still be no reason to trust the remote onion router. If one had a secured connection to an onion router one trusted, our onion router could be used as one of several intermediate routers to further complicate traffic analysis.

We have recently set up a thirteen node distributed network of government, academic, and private sites. However, at press time we have not yet gathered performance data for this network. The data we present are for a network running on a single machine. In our experimental onion routing network, five onion routers run on a single Sun Ultra 2 2170. This machine has two 167 MHz processors, and 256MB of memory. Anonymous connections are routed through a random sequence of five onion routers. Connection setup time should be comparable to a more distributed topology. Data latency, however, is more difficult to judge. Clearly, data will travel faster over socket connections between onion routers on the same machine than over socket connections between different machines. However, on a single machine the removal or addition of layers of encryption is not pipelined, so data latency may be worse.

Onion routing's overhead is mainly due to public key cryptography and is incurred while setting up an anonymous connection. On our Ultra 2 running a fast implementation of RSA [2], a single public key decryption of a 1024 bit plaintext block using a 1024 bit private key and a 1024 bit modulus takes 90 milliseconds. Encryption is much faster, because the public keys are only 16 bits long. (This is why RSA signature verification is cheaper than signing). So, the public key cryptographic overhead for routes spanning five onion routers is just under 0.5 seconds. This overhead can be further reduced, either with specialized hardware, or even simply on different hardware (a 200 MHz Pentium would be almost twice as fast).

In practice, our connection setup overhead does not appear to add intolerably to the overhead of typical socket connections. Still, it can be further reduced.

There is no reason that the same anonymous connection could not be used to carry the traffic for several ‘real’ socket connections, either sequentially or multiplexed. In fact, the specification for HTTP 1.1 defines pipelined connections to amortize the cost of socket setup, and pipelined connections would also transparently amortize the increased cost of anonymous connection setup. We are currently updating our Web proxy to be HTTP 1.1 compliant.

4 Threat Model

This section outlines our threat model. It does not intend to quantify the cost of attacks, but to define possible attacks. Future work will quantify the threat. First some vocabulary. A session is the data carried over a single anonymous connection. Data is carried in fixed length cells. Since these cells are multiply encrypted and change as they move through an anonymous connection, tracking cells is equivalent to tracking markers that indicate when cells begin. In a marker attack, the attacker identifies the set of outbound connections that some distinguished marker may have been forwarded upon. By intersecting these sets for a series of distinguished markers belonging to the same session, an attacker may determine, or at least narrow, the set of possible next hops. In a timing attack, the attacker records a timing signature for a session that correlates data rate over time. A session may have a very similar timing signature wherever it is measured over a route, so cooperating attackers may determine if they carry a particular session.

We assume that the network is subject to both passive and active attacks. Traffic may be monitored and modified by both external observers and internal network elements, including compromised onion routers. Attackers may cooperate and share information and inferences. We assume roving attackers that can monitor part, but not all, of the network at a time.

Our goal is to prevent traffic analysis, not traffic confirmation. If an attacker wants to confirm that two endpoints often communicate, and he observes that they each connect to an anonymous connection at roughly the same time, more often than is statistically expected, it is reasonable to infer that the endpoints are indeed communicating. Notice that this attack is infeasible if endpoints live in protected networks behind trusted onion routers on firewalls.

If the onion routing infrastructure is uniformly busy, then passive external attacks are ineffective. Specifically, neither the marker nor timing attacks are feasible, since external observers cannot assign markers to sessions. Active attacks are possible since reducing the

load on the system makes the network easier to analyze (and makes the system not uniformly busy).

Passive internal attacks require at least two compromised onion routers. Since onion routers can assign markers to a session, both the marker and timing attacks are possible. Specifically, timing signatures can be broadcast, and other compromised onion routers can attempt to find connections with matching timing signatures.

Another attack that is only feasible as an internal attack is the volume attack. Compromised onion routers can keep track of the number of cells that have passed over any given anonymous connection. They can then simply broadcast totals to other compromised onion routers. Cell totals that are close to the same amount at the same time at different onion routers are likely to belong to the same anonymous connection.³

Active internal attacks amplify these risks, since individual onion routers can selectively limit traffic on particular connections. An onion router could, for example, force a particular timing signature on a connection, and advertise that signature.

5 Onion Routing Specifics

5.1 Onion Routing Proxies

A proxy is a transparent service between two applications that would usually make a direct socket connection to each other but cannot. For example, a firewall might prevent direct socket connections between internal and external machines. A proxy running on the firewall may enable such connections. Proxy aware applications are becoming quite common.

Our goal has been to design an architecture for private communication that would interface with *unmodified* applications, so we chose to use proxies as the interface between applications and onion routing’s anonymous connections. For applications that are designed to be proxy aware, (e.g., WWW browsers), we simply design appropriate interface proxies. Surprisingly, for certain applications that are not proxy aware (e.g., RLOGIN), we have also been able to design interface proxies.

Because it is necessary to bridge between applications and the onion routing network, proxies must understand both application protocols and onion routing protocols. Therefore, we modularize the design into components: the application proxy, the onion proxy, and the entry funnel. The application proxy bridges between a socket connection from an application and

³Thanks to Gene Tsudik for noting this attack and for helpful discussions.

a socket connection to the onion proxy. It is the obligation of the application proxy to massage the data stream so the onion proxy, the entry funnel and the exit funnel can be application independent. Specifically, the application proxy must prepend to the data stream a *standard structure* that identifies the ultimate destination by either hostname/port or IP address/port. Additionally, it must process a one byte return code from the exit funnel and either continue if no error is reported or report the onion routing error code in some application specific meaningful way. The application proxy may also contain an optional privacy filter for sanitizing the data stream.

Upon receiving a new request, the onion proxy builds an onion defining the route of an anonymous connection. (It may use the destination address in the prepended structure to help define the route.) It then passes the onion to the funnel, and repeatedly precrypts the standard structure. Finally, it passes the precrypted standard structure through the anonymous connection to the exit funnel, thus specifying the ultimate destination. From this point on, the onion proxy blindly relays data back and forth between the application proxy and the onion routing network (and thus the exit funnel at the other end of the anonymous connection). Of course, it must apply the appropriate keystreams to incoming and outgoing data when blindly relaying data.

The entry funnel multiplexes connections from onion proxies to the onion routing network. For the services we have considered to date, a nearly generic exit funnel is adequate. Its function is to demultiplex connections from the last onion router to the outside. When it reads a data stream from the terminating onion router the first datum received will be the standard structure specifying the ultimate destination. The exit funnel makes a socket connection to that IP address/port, reports a one byte status message back to the onion routing network (and thus back to the onion proxy which in turn forwards it back to the application proxy), and subsequently moves data between the onion routing network and the new socket. (For certain services, like RLOGIN, the exit funnel also infers that the new socket must originate from a trusted port.) Entry and exit funnels are not application specific but must understand the onion routing protocol, which defines how multiplexed connections are handled.

As an example, consider the application proxy for HTTP. The user configures his browser to use the onion routing proxy. His browser may send the proxy a request like
GET http://www.domino.com/showcase/ HTTP/1.0
followed by optional fields.

The application proxy is listening for new requests. Once it obtains the GET request, it creates the standard structure and sends it (along a new socket connection) to the onion proxy, to inform the onion proxy of the service and destination of the anonymous connection. The application proxy then modifies the GET request to GET /showcase/ HTTP/1.0 and sends it directly (through the anonymous connection) to the HTTP server www.domino.com, followed by the optional fields. Notice that the server name and http:// are eliminated from the GET request because the connection is made directly to the HTTP server.

The application proxy essentially makes a connection to www.domino.com, and issues a request as if it were a client. Once this request is transmitted to the server, all proxies blindly forward data in both directions between the client and the server until the socket is broken by either side.

For the anonymizing onion routing HTTP proxy, the application proxy proceeds as outlined above with one change: it is now necessary to sanitize the optional fields that follow the GET command because they may contain identity information. Furthermore, the data stream during a connection must be monitored, to sanitize additional headers that might occur during the connection. For our current anonymizing HTTP proxy, operations that store cookies on the user's browser (to track a user, for example) are removed. This reduces function, so applications that depend upon cookies (like online shopping baskets) may not work properly.

5.2 Implementation


This section presents the interface specification between the components in an onion routing system. To provide some structure to this specification, we will discuss components in the order that data would move from an initiating client to a responding server.

There are four phases in an onion routing system: network setup, which establishes the longstanding connections between onion routers; connection setup, which establishes anonymous connections through the onion router network; data movement over an anonymous connection; and the destruction and cleanup of anonymous connections. We will commingle the discussion of these below.

5.3 Application Proxy

The interface between an application and the application proxy is application specific. The interface between the application proxy and the onion proxy is defined as follows. For each new proxy request, the

application proxy first determines if it will handle or deny the request. If rejected, it reports an application specific error message and then closes the socket and waits for the next request. If accepted, it creates a socket to the onion proxy's well known port. The application proxy then sends a standard structure to the onion proxy of the form:



Version is currently defined to be 1. *Protocol* is either 1 for RLOGIN, 2 for HTTP, or 3 for SMTP. *Retry Count* specifies how many times the exit funnel should attempt to retry connecting to the ultimate destination. Finally, the *Addr Format* field specifies the form of the ultimate destination address: 1 for a NULL terminated ASCII string with the hostname or IP address (in ASCII form) immediately followed by another NULL terminated ASCII string with the destination port number, and all others currently undefined. The ultimate destination address is sent after this standard structure, and the application proxy waits for a one byte error code before sending data.

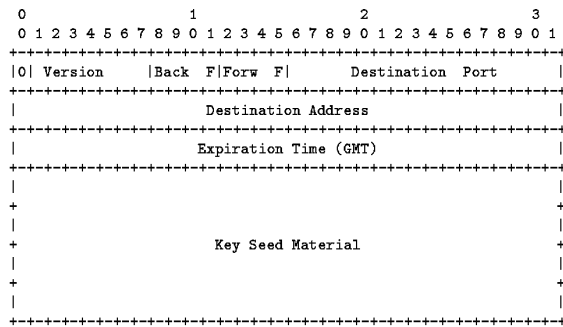
5.4 Onion Proxy

Upon receiving the standard structure, the onion proxy can decide whether to accept or reject the request based on the protocol, destination host, destination port, or the identity of the application proxy. If rejected, it sends an appropriate error code back to the application proxy, closes the socket, and waits for the next request. If accepted, it proceeds to build the onion and connects to the entry funnel of the first onion router, through the network, and to the exit funnel of the last. It next sends the standard structure to the exit funnel over the anonymous connection, and then passes all future data to and from the application proxy and anonymous connection. The repeated pre and post cryptptions and packaging of the standard structure and subsequent data is discussed later in section 5.6.

5.5 Onions

To build the anonymous connection to the exit funnel, the onion proxy creates an onion. An onion is a multi-layered data structure that encapsulates the route of the anonymous connection starting from the onion router for that exit funnel and working backward to the onion router at the entry funnel.

Each layer has the following structure:



As we will see below, the first bit must be zero for RSA public key cryptography to succeed. Following the zero bit is the *Version Number* of the onion routing system, currently defined to be 1.

The *Back F* field denotes the cryptographic function to be applied to data moving in the backward direction (defined as data moving in the direction opposite that which the onion traveled, usually toward the initiator's end of the anonymous socket connection) using *key₂* defined below. The *Forw F* field denotes the cryptographic function to be applied to data moving in the forward direction (defined as data moving in the same direction as that which the onion traveled, usually toward the responder's end of the anonymous socket connection) using *key₃* defined below. Currently defined cryptographic functions are: 0 for Identity (no encryption), 1 for DES OFB (output feedback mode) (56 bit key), and 2 for RC4 (128 bit key). The *Destination Address* and *Destination Port* indicate the next onion router in network order and are both 0 for the exit funnel. The *Expiration Time* is given in network order in seconds relative to 00:00:00 UTC January 1, 1970 (i.e., standard UNIX time(2) format) and specifies how long the onion router at this hop in the anonymous connection must track the onion against replays before it expires. *Key Seed Material* is 128 bits long and is hashed three times with SHA to produce three cryptographic keys (*key₁*, *key₂*, and *key₃*) of 128-bits each (the first eight bytes of each SHA output are used for DES and the first 16 bytes for RC4 keys).⁴

Since we use RSA public key cryptography with a modulus size of 1024-bits, the plaintext block size is 1024 bits and must be strictly less than the modulus numerically. To avoid problems, we force this relation by putting the most-significant bit first and setting it to 0 (the leading 0 above). Furthermore, the innermost layer of the onion is padded on the end with an additional 100 bytes prior to RSA encryption being

⁴Details on the cryptographic operations used in this paper can be found in [20, 26].

performed.

In version 1, an onion has five layers, but routes can be shorter. An onion is formed iteratively, innermost layer first. At each iteration, the first 128 bytes of the onion are encrypted with the public key of the onion router that is intended to decrypt that layer. The remainder of the onion is encrypted, using DES OFB with an IV (initialization vector) of 0 and *key*₁ (derived from *Key Seed Material* in that layer as defined above).⁵

Before discussing how onions and data are sent between onion routers, we will define onion router interconnection.

5.6 Onion Router Interconnection

During onion network setup (not to be confused with anonymous connection setup), longstanding connections between neighboring onion routers are established and keyed. The network topology is predefined and each onion router knows its neighbors.

To remain connected to each of its neighbors, onion routers must both listen for connections from neighbors and attempt to initiate connections to neighbors. To avoid deadlock and collision issues between pairs of neighbors, an onion router listens for connections from neighbors with “higher” IP/port addresses and initiates connections to neighbors with “lower” IP/port addresses. “Higher” and “Lower” are defined with respect to network byte ordering. (This was an expedient way to break symmetry. Ultimately we will want a more flexible solution. For example, when an onion router goes down, it should contact its neighbors upon coming back up. Requiring the neighbors to try to contact the down router until it responds is less efficient. This is not difficult to implement and we will do so in the future.)


The protocol has two phases: connection setup and keying. The initiating onion router opens a socket to a well known port of its neighboring onion router, and sends its IP address and well known port (the port is included to allow multiple onion routers to run on a single machine) in network order to identify itself. The keying phase ensues, using STS [9] which will generate two DES 56-bit keys. The link encryption over the longstanding connections is done by DES OFB with IVs of 0 and these two keys (one for data in each direction).

Once keyed, communication between onion routers is packaged into fixed sized *cells*, which allows for

⁵We use DES to encrypt the onion, and for link encryption between onion routers, because it has no licensing fees and can be used as a pseudorandom number generator. However, we would be happy to use a stronger pseudorandom number generator.

the multiplexing of both anonymous connections and control information over the longstanding connections. (Cell size was chosen to be compliant with ATM.) In version 1 of the onion routing system, there are four types of cells: PADDING (0), CREATE (1), DATA (2), and DESTROY (3).

Cells have the following structure:



The *ACI* (anonymous connection identifier) and *Command* fields are always encrypted using the link encryption between neighboring nodes. Additionally, the *Length* and *Payload* fields are encrypted using the link encryption between neighboring nodes if the command is either PADDING (0) or DESTROY (3). For CREATE (1) commands, the length is link encrypted, but the payload is already encrypted because it carries the onion. For DATA (2) commands, the length and entire payload are encrypted using the anonymous connection’s forward or backward cryptographic operations.

Each anonymous connection is assigned an ACI at each onion router, which labels an anonymous connection when it is multiplexed over the longstanding connection to the next onion router. ACIs must be unique on their longstanding connection but need not be globally unique.

To move an onion through the system, an onion router peels off the outermost layer, identifying the next hop. It checks the freshness (not expired and not replayed) of the onion, computes the necessary cryptographic keys, initializes the forward and backward cryptographic engines, chooses a new ACI for the next hop in the new connection, and then builds a data structure associated with that connection which maps incoming to outgoing ACIs and the cryptographic engines associated with forward and backward data. Since neighboring onion routers choose ACIs for each other on the thick pipe that they share, each is assigned half of the naming space. The neighboring onion router with a “higher address” chooses ACIs in the top half of that space, while its neighbor with the lower address chooses ACIs from the bottom half of that space. After the outermost layer of onion is peeled off, the rest of the onion is padded randomly to its original length, placed into CREATE cells, and then sent out in order to the appropriate neighbor. The payload of the last cell is padded with random bits to fill the cell if necessary (to

avoid traceability).

Data moves through an anonymous connection in DATA cells. At each onion router both the length and payload fields of a cell are crypted using the appropriate cryptographic engine. The new cell is sent out to the appropriate neighbor. The onion proxy must repeatedly crypt data to either add the appropriate layers of crypton on outgoing data, or remove layers of crypton from incoming data. When constructing a DATA cell from a plaintext data stream, the cell is (partially) filled, its true length is set, and all 45 bytes of the length and payload fields are repeatedly crypted using the stream ciphers defined by the onion. Therefore, when the cell arrives at the exit funnel, the length field reflects the length of the actual data carried in the payload.

If a connection is broken, a DESTROY command is sent to clean up state information. The ACI field of the DESTROY command carries the ACI of the broken connection. The length and payload must be random. Upon receipt of a DESTROY command, it is the responsibility of an onion router to forward the DESTROY appropriately and to acknowledge receipt by sending another DESTROY command back to the previous sender. After sending a DESTROY command about a particular ACI, an onion router may not send any more cells along that anonymous connection. Once an acknowledgment DESTROY message is received, an onion routing node considers the anonymous connection destroyed and the ACI can be used as a label for a new anonymous connection.

The PADDING command is used to inject data into a longstanding socket to further confuse traffic analysis. PADDING cells are discarded upon receipt.

Each onion router also reorders cells moving through it. All cells that arrive at an onion router within a fixed interval of time on any connection are mixed pseudo-randomly, except that the order of cells in each anonymous connection is preserved.

5.7 Exit Funnel

When a routing node receives an onion with *Destination Address* and *Destination Port* of 0, it knows it is the terminal onion router for the connection and passes the connection not to another onion router but to its own exit funnel. The funnel proceeds to read the standard structure that will be the first data across the anonymous socket connection, establishes a connection to the ultimate destination as indicated, and returns the status code. After this, it will blindly forward data between the anonymous connection and the connection to the responder's machine.

6 Implementation Vulnerabilities

An implementation of a secure design can be insecure. In this section, we describe several implementation decisions that were made for security considerations.

Onions are packaged in a sequence of cells that must be processed together. This onion processing involves a public key decryption operation which is relatively expensive. Therefore, it is possible to imagine an implementation that clears outgoing queues while an onion is being processed, and then outputs the onion. Therefore, any period of inactivity on the out-bound queues is likely to be followed by a sequence of onion cells being output on a single queue. Such an implementation makes tracking easier and should be avoided.

After processing at each onion router, onions are padded at the end to compensate for the removed layer. This padding must be random, since onions are not link encrypted between onion routers. Similarly, the length and payload of a DESTROY command must be new random content at each onion router; otherwise, compromised onion routers could track that payload.

In a multi-threaded implementation, there is a significant lure to rely upon apparent randomness in scheduling to reorder events. If reordering is important to the secure operation of the system, deliberate reordering is crucial, since low level system randomness may in fact be predictable.

There are two vulnerabilities for which we do not have good solutions. If part of the onion routing network is taken down, traffic analysis may be simplified. Also, if a longstanding connection between two onion routers is broken, it will result in many DESTROY messages, one for each anonymous connection that was routed through that longstanding connection. Therefore, a compromised onion router may infer from near simultaneous DESTROY messages that the associated anonymous connections had some common route. Delaying DESTROY messages hurts performance, since we require that a DESTROY message propagate to the endpoints to take down the connection that is visible to the user. Carrying the DESTROY message through the anonymous connection and garbage collecting dormant anonymous connections later would be ideal, but we do not know how to efficiently insert control information into a raw data channel, especially considering our layered encryption. One possibility is for the onion router on the initiator side of a break to send some large predetermined number of one bits back to the initiator followed by a message that the connection is destroyed. The onion proxy could then check for such a signal after it strips off each layer of each

packet, and notify the application proxy if it receives the signal. The initiator can contact the responder out of band, presumably through another anonymous connection, authenticate itself by some means as the initiator of the broken connection, and notify the responder of the break. Onion routers can either be notified directly by the onion proxy after some random delay or possibly garbage collect least recently used ACIs. We will continue to explore the feasibility of this and other possibilities.⁶

7 Applications

We first describe how to use anonymous connection in VPNs, anonymous chatting services, and anonymous cash. We then describe onion routing proxies for three Internet services: Web browsing, e-mail, and remote logins. These three onion routing proxies have been implemented. Anonymizing versions of these proxies that remove the identifying information that may be present in the headers of these services' data streams have been implemented as well.

7.1 Virtual Private Networks

If two sites wanted to collaborate, they could establish one or more long term tunnels that would multiplex many socket connections, or even raw IP packets, over a single anonymous connection. This would effectively hide who is collaborating with whom and what they are working on, without requiring the construction of an individual anonymous connection for each connection made. Such long term anonymous connections between enclaves provide the analog of a leased line over a public network. Note that the protection provided a VPN by onion routing is broader than that provided by encrypting firewalls. Basic encrypting firewalls encrypt payloads only. Thus, they protect confidentiality, but do nothing to protect against traffic analysis. IPSEC will protect traffic for individual connections by encapsulating packets in encrypted packets from the firewall, but this will not protect against institutional level traffic analysis. Communication between two such firewalls will still indicate a collaboration between the sites behind them. Constant padding may be added, but this is very expensive. And, unless many unrelated sites agree to do it, it still does not hide the existence of the VPN established between those sites that are so padding.

⁶Thanks to Gene Tsudik for some of the fundamental elements of this proposal.

7.2 Anonymous Chatting

Anonymous connections can be used in a service similar to IRC, where many parties meet to *chat* at some central server. The chat server may mate several anonymous connections carrying matching tokens. Each party defines the part of the connection leading back to itself, so no party has to trust the other to maintain its privacy. If the communicating parties layer end-to-end encryption over the mated anonymous connections, they also prevent the central server from listening in on the conversation.

7.3 Anonymous Cash

Certain forms of e-cash are designed to be anonymous and untraceable, unless they are double spent or otherwise misused. However, if a customer cannot contact a vendor without identifying himself, the anonymity of e-cash is undermined. For transactions where both payment and product can be conveyed electronically, anonymous connections can be used to hide the identities of the parties from one another [27].

How can the customer be prevented from taking his purchase without paying for it (e.g., by closing the connection early) or the vendor be prevented from taking the customer's e-cash without completing the transaction? This is a hard problem [12, 4]. In the case of a well known vendor, a practical solution is to require customers to pay first. The vendor is unlikely to deliberately cheat its customers since it may be caught in an audit.

7.4 Remote Login

We proxy remote login requests by taking advantage of the option `-l username to rlogin`. The usual `rlogin` command is of the form:

```
rlogin -l username server
```

To use `rlogin` through an onion routing proxy, one would type

```
rlogin -l username@server proxy
```

where *proxy* refers to the onion routing proxy to be used and both *username* and *server* are the same as specified above. A normal `rlogin` request is transmitted from a privileged port on the client to the well known port for `rlogin` (513) on the server as:

```
\0 username on client \0 username on server \0 terminal type \0
```

where *username on client* is the username of the individual invoking the command on the client machine, *username on server* is either the `-l` field (if specified) or

the username of the individual invoking the command on the client machine (if no `-l` is specified), and the *terminal type* is a standard termcap/linespeed specification. The server responds with a single zero byte if it will accept the connection or breaks the socket connection if an error has occurred or the connection is rejected. Our normal rlogin proxy therefore receives the initial request:

```
\0 username on client \0 username@server \0 terminal type \0
```

The proxy creates an anonymous connection to the RLOGIN port on the *server* machine and proceeds to send it a massaged request of the form:

```
\0 username \0 username \0 terminal type \0
```

Once this request is transmitted to the server, the proxy blindly forwards data in both directions between the client and server until the socket is broken by either side.

Notice that the onion router does not send the *server* the client's username on the client, so communication is anonymous, unless the data-stream subsequently reveals more information.

7.5 Web Browsing

Proxying HTTP requests follows the IETF HTTP V1.0 Specification [3]. An HTTP request from a client through an HTTP proxy is of the form:

```
GET http://www.server.com/file.html HTTP/1.0
```

followed by optional fields. Notice that an HTTP request from a client to a server is of the form:

```
GET file.html HTTP/1.0
```

also followed by optional fields. The server name and protocol scheme are missing, because the connection is made directly to the server.

As an example, a complete request from Netscape Navigator to an onion router HTTP proxy may look like this:

```
GET http://www.server.com/file.html HTTP/1.0
Referer: http://www.server.com/index.html
Proxy-Connection: Keep-Alive
User-Agent: Mozilla/3.0 (X11; I; SunOS 5.4 sun4m)
Host: www.server.com
Accept: image/gif, image/x-xbitmap, image/jpeg
```

The proxy must create an anonymous connection to `www.server.com`, and issue a request as if it were a

client. Therefore, the request must be massaged to remove the server name and scheme, and transmitted to `www.server.com` over the anonymous connection. Once this request is transmitted to the server, the proxy blindly forwards data in both directions between the client and server until the socket is broken by either side.

For privacy filtering of HTTP, the proxy proceeds as outlined above with one change. It is now necessary to sanitize the optional fields that follow the GET command because they may contain identity information. Furthermore, the data stream during a connection must be monitored, to sanitize additional headers that might occur during the connection.

The *Anonymizer* [1] also provides anonymous Web browsing. Users can connect to servers through the Anonymizer and it strips off identifying headers. This is essentially what our filtering HTTP proxy does. But packets can still be tracked and monitored. The Anonymizer could be used as a front end to the onion routing network to provide effective protection against traffic analysis. We discuss this further in section 8.

7.6 Electronic Mail

Electronic mail is proxied by utilizing the `user%host@proxy` form of e-mail address instead of the normal `user@host` form. This form should work with most current and older mail systems. Under this form, the client contacts the proxy server's well known SMTP port (25). Instead of the normal mail daemon listening to that port, the proxy listens and interprets what it receives following a strict state machine: wait for a valid HELO command, wait for a valid MAIL From: command, and then wait for a valid RCPT To: command. Each command argument is temporarily buffered. Once the RCPT To: command has been received, the proxy proceeds to create an anonymous connection to the destination server and relays the HELO and MAIL From: commands exactly as received. The RCPT To: command is massaged and forwarded. Any subsequent RCPT To: commands are rejected. Once the DATA request is transmitted to the server, the proxy forwards data in both directions from the client and server. An example of e-mail from `joe@sender.com` on the machine `sender.com` to `mary@recipient.com` via the `onion.com` onion router is given below. Joe types `mail mary%recipient.com@onion.com`. First the communications from the client on `sender.com` to the onion router SMTP proxy on `onion.com` is given, followed by the communications from the exit funnel to `recipient.com`:

220 `onion.com` SMTP Onion Routing Network.


```
HELO sender.com
250-onion.com -- Connection from
250 sender.com (2.0.0.1).
MAIL From: joe@sender.com
250 Sender is joe@sender.com.
RCPT To: mary%recipient.com@onion.com
```

The proxy massages the RCPT To: line to make the address `mary%recipient.com` and makes an anonymous connection to `recipient.com`. It then replays the massaged protocol to `recipient.com`:

```
220-recipient.com Sendmail 4.1/SMI-4.1 ready
220 at Wed, 28 Aug 96 15:15:00 EDT
HELO Onion.Routing.Network
250-recipient.com Hello Onion.Routing.Network
250 [2.0.0.5], pleased to meet you
MAIL From: joe@sender.com
250 joe@sender.com... Sender ok
RCPT To: mary%recipient.com
250 mary%recipient.com... Recipient ok
DATA
354 Enter mail, end with "." on a line by itself
```

At this point, the proxy forwards data in both directions, until a line containing only a period is sent from the sender to the recipient:

```
This is a note
.
```

The proxy forwards the line containing only a period to the recipient, and forwards the recipient's response to the sender. At that point, the proxy sends QUIT to the recipient, reads the response and closes the connection to the recipient. The proxy then waits for a command from the sender; if that command is QUIT, the proxy sends a response and closes its connection to the sender:

```
250 Mail accepted
QUIT
221 onion.com Service closing transmission channel.
```

If the command is not QUIT, then it is MAIL, and the protocol repeats. Anything else prompts an error response, and the proxy waits for the next correct command.

For the privacy filtered proxying of electronic mail, the proxy proceeds as outlined above with a few changes. It is now necessary to sanitize both the MAIL From: command and the header portion of the actual message body. Sanitization of the MAIL From: command is trivial with a simple substitution of `anonymous` for `joe@sender.com`. For the header sanitization, we have taken the conservative approach of deleting all headers, but this may be modified in the future to only remove identifying information and leave the remaining header information intact.

8 Comparisons with Related Work

Chaum [5] defines a layered object that routes data through intermediate nodes, called *mixes*. These intermediate nodes may reorder, delay, and pad traffic to complicate traffic analysis. In mixes, the assumption is that a single perfect mix adequately complicates traffic analysis, but a sequence of multiple mixes is typically used because real mixes are not ideal. Because of this, mix applications can use mixes in fixed order, and often do. Onion routers differ from mixes in at least two ways: onion routers are more limited in the extent to which they delay traffic at each node because of the real-time expectations that the applications demand of socket connections. Also, in a typical onion routing configuration, onion routers are also entry points to the onion routing network, and traffic entering or exiting at those nodes may not be visible. This makes it hard to track packets, because they may drop out of the network at any node, and new packets may be introduced at each node. While onion routing cannot delay traffic to the extent that mixes can, traffic between onion routers is multiplexed over a single channel and is link encrypted with a stream cipher. This makes it hard to parse the stream.

Anonymous remailers like Penet [17] strip headers from received mail and forward it to the intended recipient. They may also replace the sender's address with some alias, permitting replies. These sorts of remailers store sensitive state: the mapping between the alias and the true return address. Also, mail forwarded through a chain of remailers may be tracked because it appears the same to each remailer.

Mix based remailers like [7, 16] use mixes to provide anonymous e-mail services. Essentially, the mail message is carried in the innermost layer of the onion data structure. Another onion type structure, used for a return address, can be contained in the message. This makes the return path self contained, and the remailer essentially stateless. Onion routing shares many structures with Babel [16] but it uses them to build (possibly long lived) application independent connections. This makes anonymous connections accessible to a wide variety of applications. For application to e-mail it has both advantages and disadvantages. Onion routing's service makes an anonymous connection directly to the recipient's sendmail daemon. A disadvantage is that, since the connection is made in real-time, there is less freedom in mixing, which therefore might not be done as well. An advantage is that the anonymous connection is separated from the application, so anonymous e-mail systems are considerably simplified because the application specific part does not have to move data

through the network. Furthermore, because the onion routing network can carry many types of data, it has the potential to be more heavily utilized than a network that is devoted only to e-mail. Heavy utilization is the key to anonymity.

In [10], a structure similar to an onion is used to forward individual IP packets through a network. By maintaining tracking information at each router, ICMP error messages can be moved back along the hidden route. Essentially, a connection is built for each packet in a connectionless service. Although a followup paper [11] suggests that performance will be good, especially with hardware based public key cryptography, our experience suggests that both the cryptographic overhead of building onions and the tracking of onions against replay is not efficiently done on a packet-by-packet basis. However, it is easy to imagine an onion routing proxy that collects IP packets and forwards them over some anonymous connection. In this way, communication is anonymous at the IP layer, but connections need not be built for each IP packet. This anonymous IP communication may be more robust than our current architecture: it could survive a broken anonymous connection, since IP does not expect reliable delivery.

In [22], mixes are used to provide untraceable communication in an ISDN network. Here is a summary of that paper. In a phone system, each telephone line is assigned to a particular local switch (i.e., local exchange), and switches are interconnected by a (long distance) network. Anonymous calls in ISDN rely upon an anonymous connection between the caller and the long distance network. These connections are made anonymous by routing calls through a predefined series of mixes within each switch. The long distance endpoints of the connection are then mated to complete the call. (Notice that observers can tell which local switches are connected.) Also, since each phone line has a control circuit connection to the switch, the switch can broadcast messages to each line using these control circuits. So, within a switch a truly anonymous connection can be established: A phone line makes an anonymous connection to some mix. That mix broadcasts a token identifying itself and the connection. A recipient of that token can make another anonymous connection to the specified mix, which mates the two connections to complete the call.

Our goal of anonymous connections over the Internet differs from anonymous remailers and anonymous ISDN. The data is different, with real-time constraints more severe than mail, but somewhat looser than voice. Both HTTP and ISDN connections are bidirectional, but, unlike ISDN, HTTP connections are likely to be small requests followed by short bursts of returned

data. Most importantly, the network topology of the Internet is more akin to the network topology of the long distance network between switches, where capacity is a shared resource. In anonymous ISDN, the mixes hide communication within the local switch, but connections between switches are not hidden. This implies that all calls between two businesses, each large enough to use an entire switch, reveal which businesses are communicating. In onion routing, mixing is dispersed throughout the Internet, which improves hiding.

Pipe-net [8] is a proposal similar to onion routing. It has not been implemented, however. Pipe-net's threat model is more paranoid than onion routing's: it attempts to resist active attacks by global observers. For example, Pipe-net's connections carry constant traffic (to resist timing signature attacks) and disruptions to any connection are propagated throughout the network.

The Anonymizer is a Web proxy that filters the HTTP data stream to remove a user's identifying information, essentially as our filtering HTTP proxy does. For example, the Anonymizer will "strip out all references to your e-mail address, computer type, and previous page visited before forwarding your request" [1]. This makes Web browsing private in the absence of any eavesdropping or traffic analysis. The Anonymizer is vulnerable in three ways: First, it must be trusted. Second, traffic between a browser and the Anonymizer is sent in the clear, so that traffic identifies the true destination of a query, and includes the identifying information that the Anonymizer would filter. Third, even if traffic between the browser and the Anonymizer were encrypted, passive external observers could mount the volume attack mentioned in section 4. The Anonymizer, however, is now readily available to everyone on the Web.

NetAngels [21] is similar to the Anonymizer, except that it builds personal profiles of its subscribers and targets advertisements to match the profile. However, the profile is not released to the advertiser and is deleted when a subscription is canceled. Subscribers must trust NetAngels, and connections to the service are subject to the same attacks as the Anonymizer.

LPWA [19, 13] (formerly known as Janus) is a "proxy server that generates consistent untraceable aliases for you that enable you to browse the Web, register at web sites and open accounts, and be 'recognized' upon returning to your accounts, all while still *preserving your privacy*." Like the previous two, the LPWA proxy is at a server that is remote from the user application. It is thus subject to the same trust and vulnerability limitations.

It is possible, however, to shift trusted elements to

the user's machine (or to a machine on the boundary between his trusted LAN and the Internet). Shifting trust in this way can improve the security of other privacy services like the Anonymizer, NetAngels, and LPWA. Currently, those are centralized to provide an intermediary that masks the true source of a connection. If anonymous connections are used to hide the source address instead, the other functions of these services may run as a local proxy on the user's desktop. Security is improved because privacy filtering and other services are done on a trusted machine and because communication is resistant to traffic analysis. Also, there is no central point of failure.

Another approach to anonymous Web connections is Crowds [25]. Crowds is essentially a distributed and chained Anonymizer, with encrypted links between crowd members. Web traffic is forwarded to a crowd member, who flips a weighted coin and, depending on the result, forwards it either to some other crowd member or to the destination. This makes communication resistant to local observers.

9 Conclusion

This paper describes anonymous connections, their realization in onion routing, and some of their applications. Anonymous connections are resistant to both eavesdropping and traffic analysis. They separate the anonymity of the connection from the anonymity of communication over that connection. For example, two parties controlling onion routers can identify themselves to each other without revealing the existence of a connection between them. This paper demonstrates the versatility of anonymous connections by exploring their use in a variety of Internet applications. These applications include standard Internet services like Web browsing, remote login, and electronic mail. Anonymous connections can also be used to support virtual private networks with connections that are resistant to traffic analysis and that can carry connectionless traffic.

Anonymous connections may be used as a new primitive that enables novel applications in addition to facilitating secure versions of existing services [24]. Besides exploring other novel applications, future work includes a system redesign to improve throughput and an implementation of *reply onions* [15, 23]. Reply onions are basically reply addresses that enable connections to be established back to an anonymous party. We will be implementing other mechanisms for responding to anonymous connections as well. We are also beginning a detailed analysis of onion routing to enable a quantitative assessment of resistance to traffic analysis.

The onion routing network supporting anonymous connections can be configured in several ways, including a firewall configuration and a customer-ISP configuration, which moves privacy to the user's computer and may relieve the carrier of responsibility for the user's connections.

Onion routing moves the anonymous communications infrastructure below the application level, properly separating communication and applications. Since the efficacy of mixes depends upon sufficient network traffic, allowing different applications to share the same communications infrastructure increases the ability of the network to resist traffic analysis.

Acknowledgments

We have had helpful comments from and discussion with people too numerous to mention. We note especially the help of Birgit Pfitzmann, Gene Tsudik, and James Washington. We also thank the anonymous referees, the Levien family for hosting the onion dinner, and the Isaac Newton Institute for hosting one of the authors while some of this work was done. The fast UltraSparc implementation of RSA was done by Tolga Acar and Çetin Kaya Koç. This work was supported by ONR and DARPA.

References

- [1] The Anonymizer. <http://www.anonymizer.com>
- [2] T. Acar, B. S. Kaliski, Jr., and Ç. Koç. "Analyzing and Comparing Montgomery Multiplication Algorithms", *IEEE Micro*, 16(3):26-33, June 1996.
- [3] T. Berners-Lee, R. Fielding, and H. Frystyk. *Hypertext Transfer Protocol - HTTP/1.0*, <ftp://ds.internic.net/rfc/rfc1945.txt>
- [4] L. J. Camp, M. Harkavey, B. Yee, J. D. Tygar, "Anonymous Atomic Transactions", *Second USENIX Workshop on Electronic Commerce*, 1996.
- [5] D. Chaum. "Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms", *Communications of the ACM*, v. 24, n. 2, Feb. 1981, pp. 84-88.
- [6] D. E. Comer. *Internetworking with TCP/IP, Volume 1: Principles, Protocols, and Architecture*, Prentice-Hall, Engelwood Cliffs, New Jersey, 1995.

- [7] L. Cottrell. *Mixmaster and Remailer Attacks*, <http://obscura.obscura.com/~loki/remailer/remailer-essay.html>
- [8] W. Dai. Pipe-net, February 1995, post to the cypherpunks mailing list.
- [9] Whitfield Diffie, Paul C. van Oorschot, and Michael J. Wiener. "Authentication and Authenticated Key Exchanges". *Designs, Codes, and Cryptography*, 2:107–125, 1992.
- [10] A. Fasbender, D. Kesdogan, O. Kubitz. "Variable and Scalable Security: Protection of Location Information in Mobile IP", *46th IEEE Vehicular Technology Society Conference*, Atlanta, March 1996.
- [11] A. Fasbender, D. Kesdogan, O. Kubitz. "Analysis of Security and Privacy in Mobile IP", *4th International Conference on Telecommunication Systems Modeling and Analysis*, Nashville, March 1996.
- [12] M. Franklin and M. Reiter, "Fair Exchange with a Semi-Trusted Third Party", *Fourth ACM Conference on Computer and Communications Security*, Zurich, April 1997.
- [13] E. Gabber, P. Gibbons, Y. Matias, and A. Mayer. "How to Make Personalized Web Browsing Simple, Secure, and Anonymous", *Financial Cryptography '97*, February 1997, final proceedings to appear.
- [14] D. Goldschlag, M. Reed, and P. Syverson. "Privacy on the Internet", *INET '97*, Kuala Lumpur, June 1997.
- [15] D. Goldschlag, M. Reed, P. Syverson. "Hiding Routing Information", in *Information Hiding*, R. Anderson, ed., LNCS vol. 1174, Springer-Verlag, 1996, pp. 137–150.
- [16] C. Gülcü and G. Tsudik. "Mixing Email with Babel", *1996 Symposium on Network and Distributed System Security*, San Diego, February 1996.
- [17] J. Helsingius. www.penet.fi.
- [18] Internet Engineering Task Force. <http://www.ietf.org/>
- [19] <http://lpwa.com:8000/>
- [20] A. Menezes, P. van Oorschot, and S. Vanstone. *Handbook of Applied Cryptography*, CRC Press, 1997.
- [21] <http://www.netangels.com>
- [22] A. Pfitzmann, B. Pfitzmann, and M. Waidner. "ISDN-Mixes: Untraceable Communication with Very Small Bandwidth Overhead", *GI/ITG Conference: Communication in Distributed Systems*, Mannheim Feb, 1991, Informatik-Fachberichte 267, Springer-Verlag, Heidelberg 1991, pp. 451–463.
- [23] M. G. Reed, P. F. Syverson, and D. M. Goldschlag. "Proxies for Anonymous Routing", *Proc. 12th Annual Computer Security Applications Conference*, San Diego, CA, IEEE CS Press, December, 1996, pp. 95–104.
- [24] M. Reed, P. Syverson, and D. Goldschlag. "Protocols using Anonymous Connections: Mobile Applications", *1997 Security Protocols Workshop*, Paris, April 1997, final proceedings to appear.
- [25] M. Reiter and A. Rubin. *Crowds: Anonymity for Web Transactions (preliminary announcement)*, DIMACS Technical Reports 97-15, April 1997.
- [26] B. Schneier. *Applied Cryptography: Protocols, Algorithms and Source Code in C*, John Wiley and Sons, 1994.
- [27] D. Simon, "Anonymous Communication and Anonymous Cash", in *Advances in Cryptology-CRYPTO'96*, N. Kobitz, ed., LNCS vol. 1109, Springer-Verlag, 1996, pp. 61–73.
- [28] P. Syverson, D. Goldschlag, and M. Reed. "Anonymous Connections and Onion Routing", *Proceedings of the 1997 IEEE Symposium on Security and Privacy*, Oakland, CA, IEEE CS Press, May 1997, pp. 44–54.

Electronic Acknowledgement Receipt

EFS ID:	35851686
Application Number:	16278107
International Application Number:	
Confirmation Number:	4936
Title of Invention:	SYSTEM PROVIDING FASTER AND MORE EFFICIENT DATA COMMUNICATION
First Named Inventor/Applicant Name:	Derry Shribman
Customer Number:	131926
Filer:	Yehuda Binder/Dorit Binder
Filer Authorized By:	Yehuda Binder
Attorney Docket Number:	HOLA-005-US10
Receipt Date:	28-APR-2019
Filing Date:	17-FEB-2019
Time Stamp:	04:47:58
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no
------------------------	----

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Information Disclosure Statement (IDS) Form (SB08)	IDS6.pdf	1034794 <small>a0847d276775efa4d3ae43f5ed680e8b3aed0fb4</small>	no	4

Warnings:

Information:					
2	Non Patent Literature	003-Anonymous-Connections.pdf	1365617 3ebf1dc9eafeef33782fd6e2f09d8169dc346320	no	15
Warnings:					
Information:					
Total Files Size (in bytes):				2400411	
<p>This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.</p> <p><u>New Applications Under 35 U.S.C. 111</u> If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.</p> <p><u>National Stage of an International Application under 35 U.S.C. 371</u> If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.</p> <p><u>New International Application Filed with the USPTO as a Receiving Office</u> If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.</p>					



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

Table with 4 columns: APPLICATION NUMBER (16/278,107), FILING OR 371(C) DATE (02/17/2019), FIRST NAMED APPLICANT (Derry Shribman), ATTY. DOCKET NO./TITLE (HOLA-005-US10)

CONFIRMATION NO. 4936

PUBLICATION NOTICE

131926
May Patents Ltd. c/o Dorit Shem-Tov
P.O.B 7230
Ramat-Gan, 5217102
ISRAEL



Title:SYSTEM PROVIDING FASTER AND MORE EFFICIENT DATA COMMUNICATION

Publication No.US-2019-0182360-A1

Publication Date:06/13/2019

NOTICE OF PUBLICATION OF APPLICATION

The above-identified application will be electronically published as a patent application publication pursuant to 37 CFR 1.211, et seq. The patent application publication number and publication date are set forth above.

The publication may be accessed through the USPTO's publically available Searchable Databases via the Internet at www.uspto.gov. The direct link to access the publication is currently http://www.uspto.gov/patft/.

The publication process established by the Office does not provide for mailing a copy of the publication to applicant. A copy of the publication may be obtained from the Office upon payment of the appropriate fee set forth in 37 CFR 1.19(a)(1). Orders for copies of patent application publications are handled by the USPTO's Public Records Division. The Public Records Division can be reached by telephone at (571) 272-3150 or (800) 972-6382, by facsimile at (571) 273-3250, by mail addressed to the United States Patent and Trademark Office, Public Records Division, Alexandria, VA 22313-1450 or via the Internet.

In addition, information on the status of the application, including the mailing date of Office actions and the dates of receipt of correspondence filed in the Office, may also be accessed via the Internet through the Patent Electronic Business Center at www.uspto.gov using the public side of the Patent Application Information and Retrieval (PAIR) system. The direct link to access this status information is currently https://portal.uspto.gov/pair/PublicPair. Prior to publication, such status information is confidential and may only be obtained by applicant using the private side of PAIR.

Further assistance in electronically accessing the publication, or about PAIR, is available by calling the Patent Electronic Business Center at 1-866-217-9197.

Office of Data Management, Application Assistance Unit (571) 272-4000, or (571) 272-4200, or 1-888-786-0101

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number	16278107
	Filing Date	2019-02-17
	First Named Inventor	Derry Shribman
	Art Unit	
	Examiner Name	
	Attorney Docket Number	HOLA-005-US10

U.S.PATENTS						Remove
Examiner Initial*	Cite No	Patent Number	Kind Code ¹	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	9015335	B1	2015-04-21	Samuel S. Gigliotti	
	2	7788378	B2	2010-08-31	Ravi T. Rao	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS						Remove
Examiner Initial*	Cite No	Publication Number	Kind Code ¹	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20130304796	A1	2013-11-14	Steven J. Jackowski	
	2	20120164980	A1	2012-06-28	Vinh Van Phan	
	3	20010054020	A1	2001-12-20	Brian E. Barth	
	4	20160105530	A1	2016-04-14	Derry Shribman	

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number	16278107
Filing Date	2019-02-17
First Named Inventor	Derry Shribman
Art Unit	
Examiner Name	
Attorney Docket Number	HOLA-005-US10

5	20070050522	A1	2007-03-01	Adam J. Grove
6	20090216887	A1	2009-08-27	Andreas Hertle
7	20130080575	A1	2013-03-28	Matthew Browning Prince

If you wish to add additional U.S. Published Application citation information please click the Add button.

FOREIGN PATENT DOCUMENTS

Examiner Initial*	Cite No	Foreign Document Number ³	Country Code ²	Kind Code ⁴	Publication Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear	T ⁵
	1	2922275	EP	B1	2016-03-23	Axis AB		

If you wish to add additional Foreign Patent Document citation information please click the Add button.

NON-PATENT LITERATURE DOCUMENTS

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T ⁵
	1	"Keep Alive" - Imperva, 2019 https://www.imperva.com/learn/performance/keep-alive (2019) (3 pages)	
	2	Third party observation filed on June 21, 2019 in PCT Application No. PCT/IL2018/050910 (7 pages)	
	3	ETF named: IPv6 Tunnel Broker, April 1999 - First uploaded document submitted with third party observation dated June 21, 2019 (13 pages)	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number	16278107
	Filing Date	2019-02-17
	First Named Inventor	Derry Shribman
	Art Unit	
	Examiner Name	
	Attorney Docket Number	HOLA-005-US10

4		RFC 3053 (January 2001) named: IPv6 Tunnel Broker - Secod uploaded document submitted with third party observation dated June 21, 2019 (13 pages)
---	--	---

If you wish to add additional non-patent literature document citation information please click the Add button

EXAMINER SIGNATURE

Examiner Signature		Date Considered	
--------------------	--	-----------------	--

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

¹ See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. ² Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). ³ For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. ⁴ Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. ⁵ Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number	16278107
	Filing Date	2019-02-17
	First Named Inventor	Derry Shribman
	Art Unit	
	Examiner Name	
	Attorney Docket Number	HOLA-005-US10

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.

The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Yehuda Binder/	Date (YYYY-MM-DD)	2019-07-01
Name/Print	Yehuda Binder	Registration Number	73612

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. **DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Home > [Learning Center](#) > [Performance](#) > [Keep Alive](#)

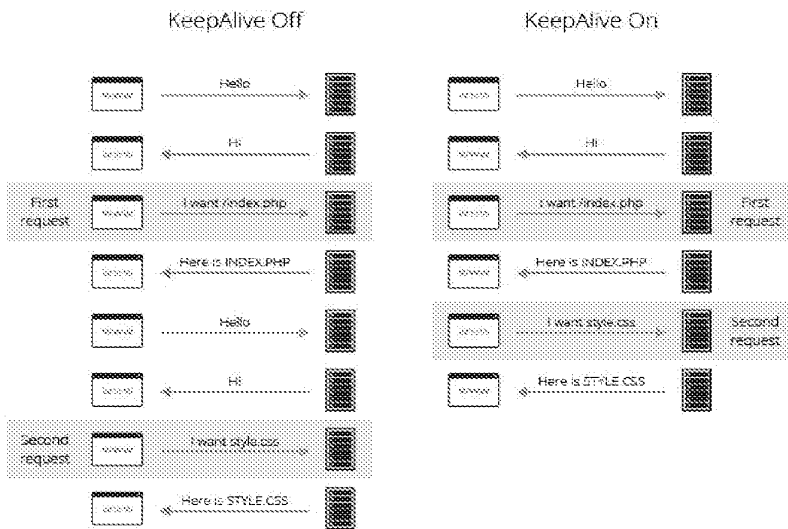
Keep Alive

5k Views [Performance](#) [CDN Guide](#) [Connection Optimization](#)

What is HTTP Keep-Alive

HTTP keep-alive, a.k.a., HTTP persistent connection, is an instruction that allows a single TCP connection to remain open for multiple HTTP requests/responses.

By default, HTTP connections close after each request. When someone visits your site, their browser needs to create new connections to request each of the files that make up your web pages (e.g. images, Javascript, and CSS stylesheets), a process that can lead to high [page load times](#).



Enabling the keep-alive header allows you to serve all web page resources over a single connection. Keep-alive also reduces both CPU and memory usage on your server.

Enabling the Keep-Alive Header

In the event that keep-alive is not enabled on your server, it can be turned on by adding the following code to your .htaccess file:

Learning Objectives

- Understand the concept of HT keep-alive
- Learn how to enable the keep-header
- Appreciate the benefits of connection keep-alive
- Learn about CDNs and keep-connections

What do you want to learn?

Related Topics

- > [Reverse Proxy](#)
- > [What is a CDN](#)
- > [The Essential CDN Guide](#)
- > [Cache Control](#)
- > [CDN Caching](#)
- > [CDN Infrastructure](#)
- > [Round Trip Time \(RTT\)](#)
- > [CDN and SSL/TLS](#)

By using this site you agree to the use of cookies for analytics, personalized content, ads, and as described in our [Cookie Policy](#). [Learn More](#)

```

imperva
<!-- Module mod_header.c -->
Header set Connection Keep-Alive
</Module>

```

Within the 'Connection keep-alive' header, the following two directives can affect its functionality.

1. **MaxKeepAliveRequests** – This directive sets the maximum number of requests for every keep-alive connection. When determining this figure, it's important to take into account the number of files on your website that a user might want to access.
2. **KeepAliveTimeout** – This directive sets the time that a server should wait for user requests before a new TCP connection needs to be established. This figure should be set according to how frequently your website is visited, i.e., sites with high traffic volumes will want to have a large timeout value to limit the number of TCP connection requests.

The Benefits of Connection Keep Alive

The HTTP keep-alive header maintains a connection between a client and your server, reducing the time needed to serve files. A persistent connection also reduces the number of TCP and SSL/TLS connection requests, leading to a drop in [round trip time \(RTT\)](#).

Establishing a TCP connection first requires a three-way handshake – a mutual exchange of SYN and ACK packets between a client and server before data can be transmitted. Using the keep-alive header means not having to constantly perform this process. This results in:

- ⌘ **Network resource conservation** – It's less taxing on network resources to use a single connection per client.
- ⌘ **Reduced network congestion** – Reducing the number of TCP connections between your servers and clients can lead to a drop in network congestion.
- ⌘ **Decreased latency** – Reducing the number of three-way handshakes can lead to improved site latency. This is especially true with [SSL/TLS connections](#), which require additional round-trips to encrypt and verify connections.

CDNs and Keep-Alive Connections

Keep-alive connections allow CDNs to reduce your site's RTT while still providing SSL/TLS security benefits.

The Imperva CDN uses keep-alive to maintain an open connection with your origin-server in between user sessions, for a few minutes at a time—as long as your site is visited while the connection is open, your CDN doesn't need to engage in any new SSL/TLS negotiations.

This saves a considerable amount of overhead that would have been used to initiate a new connection request with your origin for every new user request. Instead, each CDN proxy server is able to leverage its open connection to download resources for many users at once.

Cloud Security
FlexProtect Plans
Simplifying our Portfolio

[Products](#) [Support](#)

Case Studies
[Partners](#) [Resources](#) [About us](#)
Learning Center
Industry Solutions

Cloud Application Security
Documentation Portal
API Integration

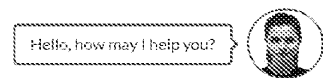
Events
Partners
Careers

[Login](#) [Request](#)

or [Contact Us](#)

[in](#) [f](#) [t](#)

[d](#) [v](#) [o](#)



By using this site you agree to the use of cookies for analytics, personalized content, ads, and as described in our [Cookie Policy](#). [Learn More](#)



(11) **EP 2 922 275 B1**

(12) **EUROPEAN PATENT SPECIFICATION**

(45) Date of publication and mention
of the grant of the patent:
23.03.2016 Bulletin 2016/12

(51) Int Cl.:
H04L 29/08 (2006.01)

(21) Application number: **14161503.9**

(22) Date of filing: **25.03.2014**

(54) **TUNNEL BROKER IN A SERVICE ORIENTED ARCHITECTURE**

TUNNEL-BROKER IN EINER DIENSTORIENTIERTEN ARCHITEKTUR

COURTIER DE TUNNEL DANS UNE ARCHITECTURE ORIENTÉE SERVICES

(84) Designated Contracting States:
**AL AT BE BG CH CY CZ DE DK EE ES FI FR GB
GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO
PL PT RO RS SE SI SK SM TR**

(30) Priority: **18.03.2014 US 201414218601**

(43) Date of publication of application:
23.09.2015 Bulletin 2015/39

(73) Proprietor: **Axis AB
223 69 Lund (SE)**

(72) Inventors:
• **Edlund, Björn
241 61 Löberöd (SE)**

- **Ståhl, Joachim
224 72 Lund (SE)**
- **Roubert, Joakim
224 75 Lund (SE)**
- **Ranbro, Mikael
241 35 Eslöv (SE)**
- **Olsson, Staffan
244 60 Furulund (SE)**
- **Hartzell, Ted
234 34 Lomma (SE)**

(74) Representative: **Awapatent AB
P.O. Box 1066
251 10 Helsingborg (SE)**

(56) References cited:
US-A1- 2002 087 707 US-B1- 7 779 086

EP 2 922 275 B1

Note: Within nine months of the publication of the mention of the grant of the European patent in the European Patent Bulletin, any person may give notice to the European Patent Office of opposition to that patent, in accordance with the Implementing Regulations. Notice of opposition shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

Description

[0001] This disclosure generally relates to a tunnel broker in a service oriented architecture system.

BACKGROUND INFORMATION

[0002] A network of devices may communicate over a network and may form part of a system that provides an array of various services. Different devices may provide different services at different times and the system may need to keep track of which services are available at which devices. When a device is added, removed, or modified, for example, the configuration of the system changes. In a system with a large number of devices, this may result in frequent need for reconfiguration, which consumes system resources. Thus, keeping track of available services at different devices may be a challenging task.

[0003] US 7,779,086 B1 discloses that a client calls a service on a server via a client handle H1 and may repeatedly call the service on the service via the client handle H1. A client, or a Cluster Name Server (CNS), may identify a node cluster associated with the client upon receipt of a request from the client. From the node cluster, the CNS may identify the appropriate server that provides the requested service.

[0004] US 2002/087707 A1 discloses an extended Domain Name Server (eDNS) system that maps a domain name to a group of servers. A first server is selected and a client device connects to the first server. If the first server becomes unavailable, the first server sends a handoff request to a second server as well as a redirection notice to the client device. The client device then establishes a connection to the second server.

SUMMARY

[0005] The invention is defined by a method according to claim 1. Further embodiments are set out in the dependent claims 2-7.

[0006] A possible advantage of this method is that a client device using or accessing a service has to only provide specifications for the service, without needing to discover or provide a network address for the service in order to establish a communication tunnel. Another possible advantage of this method is that if the service fails at a first device, the second end of the tunnel is moved to a second device that hosts the service without interruption to the client device and without the client device discovering that the second end of the tunnel has been moved. Another possible advantage is that the client device may not be aware that the client device is accessing the service through a communication tunnel. These advantages may result in a technical effect of reducing processing time and conserving system resources. The method may include fewer advantages, different advantages, or additional advantages to the ones described

above.

[0007] Additionally, the method may include sending a search query that specifies the requested service property to a service registry, wherein the service registry includes a list of services available in one or more nodes of the system; receiving search results from the service registry, wherein the search results include a list of one or more nodes having the requested service property; and selecting the first node in the system that hosts the first service instance having the requested service property may include selecting the first node from the list of one or more nodes having the requested service property. Sending a search query to a service registry, and receiving search results that include a list of one or more nodes having the requested service property, may provide an advantage of identifying nodes that hosts services that satisfy the requirements of the requested service property.

[0008] Additionally, the method may include determining that the communication tunnel should be updated; and selecting the second node in the system that hosts the second instance having the requested service property may be based on determining that the communication tunnel should be updated. Determining that the communication tunnel should be updated, and selecting the second node based on determining that the communication tunnel should be updated, may provide an advantage of responding to changing conditions, thereby continuing to provide to the client access to the service, having the requested service property, via the communication tunnel.

[0009] Additionally, determining that the communication tunnel should be updated may include re-sending the search query to the service registry at particular intervals; and receiving updated search results from the service registry, wherein the updated search results include an indication that the first node no longer hosts the first service instance having the requested service property or that the first service instance no longer has the requested service property. Re-sending the search query to the service registry at particular intervals, and receiving updated search results from the service registry, may provide an advantage of keeping track of which services are available at which nodes.

[0010] Additionally, determining that the communication tunnel should be updated may include at least one of receiving an indication from the service registry that the first node no longer hosts the first service instance having the requested service property or that the first service instance no longer has the requested service property; receiving a message from the client device that the first node is unreachable; or receiving a message from the first node that the first node is unable to provide the first service to the client device. Receiving an indication from a service registry that a node no longer hosts a service instance having a requested service property, or that the service instance no longer has the requested service property, may provide an advantage of determin-

ing when the second end of the communication tunnel should be moved to another node.

[0011] Additionally, determining that the communication tunnel should be updated may include determining that another node in the system is a better match for the requested service property; and selecting the second node in the system that hosts the second instance having the requested service property may be based on determining that another node in the system is a better match for the requested service property. Determining that another node in the system is a better match for the requested service property, and selecting the second node based on determining that another node in the system is a better match for the requested property, may provide an advantage of managing the communication tunnel to provide the best possible match to the requested service property for the client device.

[0012] Additionally, determining that another node in the system is a better match for the requested service property may include re-sending the search query to the service registry at particular intervals; and receiving updated search results from the service registry, wherein the updated search results include an indication that another node in the system is a better match for the requested service property. Re-sending the search query to the service registry at particular intervals, and receiving updated search results from the service registry, may provide an advantage of keeping track of which services are available at which nodes.

[0013] Additionally, determining that another node in the system is a better match for the requested service property may include receiving an indication from the service registry that the second node hosts the second service instance, wherein the second service instance is a better match for the requested service property. Receiving an indication from a service registry that another node in the system is a better match for the requested service property may provide an advantage of enabling the tunnel broker to determine that the second end of the communication tunnel should be moved to another node.

[0014] Additionally, the method may include determining one or more network connection metrics for a connection from the client device to particular nodes of the nodes included in the list of one or more nodes having the requested service property; and selecting the first node from the list of one or more nodes having the requested service property may be based on the determined one or more network connection metrics. Determining one or more network connection metrics for a connection from the client device to particular nodes may provide an advantage of enabling the tunnel broker to select a node that will provide a better connection to the client device via the communication tunnel.

[0015] Additionally, moving the second end of the communication tunnel from the first node to the second node may be done transparently with respect to the client device. Moving the second end of the communication tunnel transparently with respect to the client device may pro-

vide an advantage of enabling the client device to access the service without any perceived interruptions.

[0016] Additionally, the requested service property may include one or more of a particular service interface; a particular operating system associated with the service; a particular processing capacity associated with the service; a particular storage capacity associated with the service; a particular bandwidth associated with the service; a particular location associated with the service; a particular codec associated with the service; a particular domain associated with the service; or a particular security level associated with the service. Configuring the tunnel broker to select a node based on these properties may benefit the system by enabling a client device to request a service that includes at least one of the properties.

[0017] The invention is further defined by a computer device according to claim 8. Further embodiments are set out in the dependent claims 9-13.

[0018] A possible advantage of this computer device is that a client device using or accessing a service has to only provide specifications for the service, without needing to discover or provide a network address for the service in order to establish a communication tunnel. Another possible advantage of this computer device is that if the service fails at a first device, the tunnel broker may move the second end of the tunnel to a second device that hosts the service without interruption to the client device and without the client device discovering that the second end of the tunnel has been moved. Another possible advantage is that the client device may not be aware that the client device is accessing the service through a communication tunnel. These advantages may result in a technical effect of reducing processing time and conserving system resources. The computer device may include fewer advantages, different advantages, or additional advantages to the ones described above.

[0019] Additionally, the tunnel broker may be further configured to send a search query that specifies the requested service property to a service registry, wherein the service registry includes a list of services available in one or more nodes of the system; receive search results from the service registry, wherein the search results include a list of one or more nodes having the requested service property; and wherein, when selecting the first node in the system that hosts the first service instance having the requested service property, the tunnel broker may be further configured to select the first node from the list of one or more nodes having the requested service property. Sending a search query to a service registry, and receiving search results that include a list of one or more nodes having the requested service property, may provide an advantage of identifying nodes that hosts services that satisfy the requirements of the requested service property.

[0020] Additionally, the tunnel broker may be further configured to determine that the communication tunnel should be updated; and the tunnel broker may be con-

figured to select the second node in the system that hosts the second instance having the requested service property based on determining that the communication tunnel should be updated. Determining that the communication tunnel should be updated, and selecting the second node based on determining that the communication tunnel should be updated, may provide an advantage of enabling the tunnel broker to respond to changing conditions, thereby continuing to provide to the client access to the service, having the requested service property, via the communication tunnel.

[0021] Additionally, when determining that the communication tunnel should be updated, the tunnel broker may be further configured to re-send the search query to the service registry at particular intervals; and receive updated search results from the service registry, wherein the updated search results include an indication that the first node no longer hosts the first service instance having the requested service property or that the first service instance no longer has the requested service property. Re-sending the search query to the service registry at particular intervals, and receiving updated search results from the service registry, may provide an advantage of enabling the tunnel broker to keep track of which services are available at which nodes.

[0022] Additionally, when determining that the communication tunnel should be updated, the tunnel broker may be further configured to at least one of receive an indication from the service registry that the first node no longer hosts the first service instance having the requested service property or that the first service instance no longer has the requested service property; receive a message from the client device that the first node is unreachable; or receive a message from the first node that the first node is unable to provide the first service to the client device. Receiving an indication from a service registry that a node no longer hosts a service instance having a requested service property, or that the service instance no longer has the requested service property, may provide an advantage of enabling the tunnel broker to determine that the second end of the communication tunnel should be moved to another node.

[0023] Additionally, when determining that the communication tunnel should be updated, the tunnel broker may be further configured to determine that another node in the system is a better match for the requested service property; and the tunnel broker may be configured to select the second node in the system that hosts the second instance having the requested service property based on determining that another node in the system is a better match for the requested service property. Determining that another node in the system is a better match for the requested service property, and selecting the second node based on determining that another node in the system is a better match for the requested property, may provide an advantage of enabling the tunnel broker to manage the communication tunnel to provide the best possible match to the requested service property for the

client device.

[0024] Additionally, when determining that another node in the system is a better match for the requested service property, the tunnel broker may be further configured to receive an indication from the service registry that the second node hosts the second service instance, wherein the second service instance is a better match for the requested service property. Re-sending the search query to the service registry at particular intervals, and receiving updated search results from the service registry, may provide an advantage of enabling the tunnel broker to keep track of which services are available at which nodes.

[0025] Additionally, the tunnel broker may be configured to move the second end of the communication tunnel from the first node to the second node transparently with respect to the client device. Moving the second end of the communication tunnel transparently with respect to the client device may provide an advantage of enabling the client device to access the service without any perceived interruptions.

[0026] According to yet another aspect, a computer device may include logic configured to implement a tunnel broker configured to receive a request from a client device for a service in a system, the service having a requested service property; send a search query that specifies the requested service property to a service registry, wherein the service registry includes a list of services available in one or more nodes of the system; receive search results from the service registry, wherein the search results include a list of one or more nodes having the requested service property; select a first node in the system that hosts a first service instance having the requested service property from the list of one or more nodes having the requested service property; establish a communication tunnel between the client device and the selected first node, wherein the communication tunnel includes a first end at the client device and a second end at the first node; determine that the that the first node no longer hosts the first service instance having the requested service property or that the first node has become unreachable; select a second node in the system that hosts a second service instance having the requested service property, in response to determining that the first node no longer hosts the first service instance having the requested service property or that the first node has become unreachable; and move the second end of the communication tunnel from the first node to the second node transparently with respect to the client device.

[0027] A possible advantage of this computer device is that a client device using or accessing a service has to only provide specifications for the service, without needing to discover or provide a network address for the service in order to establish a communication tunnel. Another possible advantage of this computer device is that if the service fails at a first device, the tunnel broker may move the second end of the tunnel to a second device that hosts the service without interruption to the client

device and without the client device discovering that the second end of the tunnel has been moved. Another possible advantage is that the client device may not be aware that the client device is accessing the service through a communication tunnel. These advantages may result in a technical effect of reducing processing time and conserving system resources. The computer device may include fewer advantages, different advantages, or additional advantages to the ones described above.

BRIEF DESCRIPTION OF THE DRAWINGS

[0028]

Fig. 1 is a block diagram illustrating an exemplary environment according to one or more embodiments described below;

Fig. 2 is a block diagram illustrating exemplary components of a device of Fig. 1;

Fig. 3 is a block diagram illustrating exemplary functional layers of a device of Fig. 1;

Fig. 4A is a block diagram illustrating exemplary functional components of a service layer of Fig. 3;

Fig. 4B is a block diagram illustrating the functionality of the service registry of Fig. 4A;

Fig. 4C is a block diagram illustrating exemplary functional components of the service registry of Fig. 4A;

Fig. 4D is a block diagram of an exemplary property table for a particular service that may be stored by the service registry of Fig. 4A;

Fig. 5A is a block diagram illustrating functional components of an overlay network layer of Fig. 3;

Fig. 5B is a block diagram of a tree of an exemplary functional overlay network;

Fig. 6 is a block diagram illustrating functional components of a tunnel broker;

Fig. 7A is a block diagram illustrating components that may be stored in the service registry of Fig. 4A;

Fig. 7B is a block diagram illustrating components that may be stored in the tunnel database of Fig. 6;

Fig. 8 is a flowchart of a process for setting up and managing a communication tunnel according to an implementation described herein; and

Figs. 9A-9C are diagrams of exemplary scenarios of setting up and managing a communication tunnel according to an implementation described herein.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

[0029] The following detailed description refers to the accompanying drawings. The same reference numbers in different drawings identify the same or similar elements.

[0030] A system, such as a surveillance system, may include a large number of devices that can provide services. If a client device requests a particular service, a

tunneling protocol may be used to encapsulate traffic via a communication tunnel between the client device and a device providing the service. A tunnel can be set up manually by the client device or by a tunnel broker that selects the endpoints of the tunnel and establishes the tunnel between the endpoints. In order to establish the tunnel, the client device or the tunnel broker needs to know the topology of the network. For example, the client device may need to know the network address of the device providing the service and will request a tunnel to be set up between the network address of the client device and the network address of the device providing the service. Thus, the client device may first need to determine the network address of the device providing the service. Furthermore, if the device providing the service becomes unavailable, if the service becomes unavailable, or if the properties of the service change, the client device may need to locate another device in the system that provides the service, may need to determine the network address of the new device, and may need to request the tunnel broker to set up a new communication tunnel to the new device. In a system with a large number of devices with changing service capabilities, such a process may be slow and consume a large number of resources. These issues result in a technical problem in need of a solution.

[0031] Implementations described herein relate to a tunnel broker in a system based on a service oriented architecture (SOA). In system based on a SOA, functionality is discretized into services. A service is a self-contained cohesive unit of functionality. Communication with a service is performed through a service interface that has a defined message format. The communication process is independent of the implementation of the service. The service may provide end user functionality and the service interface may be designed to be understandable by business people. Furthermore, each service is independent of other services and the boundaries of the service are explicit. Thus, if one service crashes, other services will not be affected. Therefore, each service may run as a different process, for example.

[0032] Services provided by a node in the system are stored in a service registry. The service registry stores properties for each service, such as a service identifier, an operating system associated with the service, location coordinates of the node on which the service is running, processing capacity associated with the service, bandwidth capacity associated with the service, and/or other types of properties associated with the service. Not all nodes in the system may include a service registry. Thus, some service registries may store services available at other nodes in the system. Furthermore, service registries in the system may be topologically interconnected and a second service registry may be accessible through a first service registry. If the first service registry receives a search query and does not identify a match for the search query, the first service registry may forward the search query to the second service registry. Thus, to a client submitting a search query to locate a service, the

service registries in the system may appear as a single distributed service registry.

[0033] In order to solve the technical problem described above, the SOA system includes a tunnel broker. The tunnel broker is configured to establish and manage a communication tunnel from a client device to a service having a requested service property. The communication tunnel may have a first end at the client device and a second end at a device providing the service having the requested service property. The tunnel broker may change the device at the second end of the communication tunnel in a manner transparent to the client device. For example, instead of requesting a communication tunnel based on network addresses or device identifiers (e.g., a tunnel between network address 12.11.1.43:1233 and network address 143.223.123.1:22), a client may request a communication tunnel between a service host with a service property VIDEO_MONITORING_SERVICE and a service host with a service property CAMERA_FRONT_DOOR. The tunnel broker may identify a device with the service property CAMERA_FRONT_DOOR and may set up a communication tunnel to the client device (e.g., the VIDEO_MONITORING_SERVICE service host). The client device may not be aware of the actual network address or device identifier of the device providing the CAMERA_FRONT_DOOR service. Furthermore, if the device providing the CAMERA_FRONT_DOOR service becomes unavailable or stops hosting the service, the tunnel broker may identify another device in the system that host a service with the requested CAMERA_FRONT_DOOR service and may switch the second end of the communication tunnel to the other device transparently with respect to the client device.

[0034] The tunnel broker may receive a request from a client device for a service having a requested service property and may send a search query to a service registry for a list of system nodes that host a service with the requested service property. The service property may specify a particular service interface, a particular operating system, a particular processing capacity, a particular storage capacity, a particular bandwidth and/or bitrate, a particular location, a particular codec, a particular network domain, a particular security level, and/or another type of service property. The request may specify multiple requested service properties. Thus, the phrase "requested service property" may refer to multiple requested service properties. The service registry may return a list of one or more system nodes that host a service instance with the requested service property or properties and the tunnel broker may select one of the system nodes from the list. The system node may be selected based on the best match to the requested service property or properties and/or based on other factors, such as connection metrics. For example, if two nodes host a service instance with the requested property and one of the nodes has a higher quality connection to the client device, the tunnel broker may select the node with the higher quality con-

nection.

[0035] The tunnel broker may establish a communication tunnel with a first end at the client device and the second end at the selected system node providing the requested service. The tunnel may be set up using, for example, Tunnel Setup Protocol (TSP). After the communication tunnel is set up, the tunnel broker may determine that the communication tunnel should be updated. As an example, the tunnel broker may re-send the search query to the service registry at particular intervals and may receive updated search results from the service registry. The updated search results may indicate that the selected system node is not available, that the selected system node no longer hosts the service, that the properties of the hosted service no longer match the requirements of the requested service property, and/or that another system node hosts a service instance that better matches the requirements of the requested service property.

[0036] As another example, the tunnel broker may receive an update from the service registry without re-sending the search query. For example, the tunnel broker may set up a subscription for changes to service matching the search query and the service registry may send periodic updates to the tunnel broker. As yet another example, the tunnel broker may receive a message from the client device that the selected node is not reachable via the communication tunnel or may receive a message from the selected node that the selected node is no longer hosting the requested service.

[0037] In response to determining that the communication tunnel should be updated, the tunnel broker may select another node that hosts a service instance with the requested service property and may move the second end of the communication tunnel from the first selected node to the second selected node. The client device may continue to use the service without detecting that the tunnel broker has moved the second end of the communication tunnel.

[0038] In some implementations, a one-to-many communication tunnel may be established and maintained by the tunnel broker. For example, a client device may request a communication tunnel to a particular number of service instances having a requested service property. Thus, tunnel broker may set up a communication tunnel with a first end at the client device and with multiple second ends, each second end being connected to a particular service instance having the requested property.

[0039] Thus, the tunnel broker, as described herein, provides a possible advantage in that a client device using or accessing a service has to only provide specifications for the service, without needing to discover or provide a network address for the service in order to establish a communication tunnel. Another possible advantage of the tunnel broker is that if the service fails at a first device, the second end of the tunnel is moved to a second device that hosts the service without interruption to the client device and without the client device discov-

ering that the second end of the tunnel has been moved. Another possible advantage is that the client device may not even be aware that the client device is accessing the service through a communication tunnel.

[0040] Fig. 1 is a block diagram of an exemplary environment 100 in which the systems and/or methods described can be implemented. As shown in the embodiment of Fig. 1, environment 100 includes a network 110, sub-networks 120-A to 120-N (referred to collectively as "sub-networks 120" and individually as "sub-network 120"), devices 130-A-A to 130-N-K (referred to collectively as "devices 130" and individually as "device 130"), and administration device 150. Device 130-N-K refers to the Kth device 130 in sub-network 120-N. In this embodiment, the components in environment 100 form a service-oriented architecture (SOA) system service bus 140.

[0041] Network 110 enables sub-networks 120 and/or devices 130 to communicate with each other. Network 110 may include one or more circuit-switched networks and/or packet-switched networks. For example, in one embodiment, network 110 includes a local area network (LAN), a wide area network (WAN), a metropolitan area network (MAN), a Public Switched Telephone Network (PSTN), an ad hoc network, an intranet, the Internet, a fiber optic-based network, a wireless network, and/or a combination of these or other types of networks.

[0042] Sub-network 120 may include a LAN (e.g., a Layer 2 network) and/or a private network (e.g., a Layer 3 network). Sub-network 120 may interconnect one or more devices 130. For example, sub-network 120-A may interconnect devices 130-A-A to 130-A-J. Device 130 may include any device configured to communicate via SOA system service bus 140, for example.

[0043] Device 130 may include a server computer device, such as a Hypertext Preprocessor (PHP) server device, a C program server device, a Linux server device, a Windows server device, and/or another type of server device; a personal computer device, such as a desktop, laptop, tablet, a mobile communication device, and/or another type of personal computer device running Windows, Linux, Android, iOS, and/or another operating system; a monitoring device, such as a visible light camera, an infrared (IR) camera, a heat signature camera; a microphone; an alarm sensor, such as a motion sensor, a heat sensor, a pressure sensor, and/or another type of alarm sensor; a microcontroller computer device; and/or another type of computer device. While devices 130 are shown as connected to a sub-network 120, a particular device 130 may connect directly to network 110.

[0044] In one embodiment, SOA system service bus 140 is implemented between devices 130 on top of an existing network topology. SOA system service bus 140 may enable different types of devices 130, and/or devices 130 implemented using different platforms, to communicate using a service oriented architecture. SOA system service bus 140 may enable a first device 130 to request a particular service from any device 130 (e.g., itself or another device 130). Thus, a client (e.g., itself a "service"

or a "client service") hosted by first device 130 may call upon a service hosted by a second device 130 (e.g., when the service is not available in first device 130). A first service (e.g., in first device 130) that requests another service (e.g., in second device 130) is referred to as a "client" or a "client service" as having initiated the request. The first service may also provide services to other services in the network, for example.

[0045] In one embodiment, a service is accessed via a standardized service interface. Each type of service may be associated with a particular service interface (e.g., a different service interface). A client requesting a service may thus communicate with a service interface and the client may be agnostic with respect to the actual implementation of the service. In other words, implementations of services communicate with each other using protocols defined by the service interfaces so that each implementation does not have to be concerned with the others' implementations. A running service implementation, associated with a particular service interface, may be referred to as a service instance. A device 130 that includes a service host (e.g., a device that hosts a service) may keep track of available service instances with a service registry (e.g., a list or database of services). SOA system service bus 140 may enable communication between devices 130 to locate a requested service by searching service registries of service hosts in devices 130.

[0046] Administration device 150 may enable an administrator to configure or otherwise manage SOA system service bus 140. For example, administration device 150 may include a portable communication device (e.g., a mobile phone, a smart phone, a phablet device, a global positioning system (GPS) device, and/or another type of wireless device); a personal computer or workstation; a server device; a laptop, tablet, or another type of portable computer; and/or any type of device with communication capability.

[0047] Like network 110, sub-network 120 may include one or more circuit-switched networks and/or packet-switched networks. For example, sub-network 120 may include a LAN, a WAN, a MAN, a PSTN, an ad hoc network, an intranet, the Internet, a fiber optic-based network, a wireless network, and/or a combination of these or other types of networks.

[0048] Although Fig. 1 shows exemplary components of environment 100, in other implementations, environment 100 may include fewer components, different components, differently arranged components, or additional components than depicted in Fig. 1. Additionally or alternatively, any one device in environment 100 (or any group of devices) may perform functions described as performed by one or more other devices in environment 100.

[0049] Fig. 2 is a block diagram illustrating exemplary components of device 130. As shown in Fig. 2, device 130 may include a bus 210, a processor 220, a memory 230, an input device 240, an output device 250, and a communication interface 260.

[0050] Bus 210 may include a path that permits communication among the components of device 130. Processor 220 may include any type of single-core processor, multi-core processor, microprocessor, latch-based processor, and/or processing logic (or families of processors, microprocessors, and/or processing logics) that interprets and executes instructions. In other embodiments, processor 220 may include an application-specific integrated circuit (ASIC), a field-programmable gate array (FPGA), and/or another type of integrated circuit or processing logic.

[0051] Memory 230 may include any type of volatile and/or dynamic storage device that may store information and/or instructions, for execution by processor 220, and/or any type of non-volatile storage device that may store information for use by processor 220. For example, memory 230 may include a random access memory (RAM) or another type of dynamic storage device, a read-only memory (ROM) device or another type of static storage device, a content addressable memory (CAM), a magnetic and/or optical recording memory device and its corresponding drive (e.g., a hard disk drive, optical drive, etc.), and/or a removable form of memory, such as a flash memory.

[0052] Input device 240 may allow an operator to input information into device 130. Input device 240 may include, for example, a keyboard, a mouse, a pen, a microphone, a remote control, an audio capture device, an image and/or video capture device, a touch-screen display, and/or another type of input device. In one embodiment, device 130 may be managed remotely and may not include input device 240. In other words, device 130 may be "headless" and may not include a keyboard, for example.

[0053] Output device 250 may output information to an operator of device 130. Output device 250 may include a display, a printer, a speaker, and/or another type of output device. For example, device 130 may include a display, which may include a liquid-crystal display (LCD) for displaying content to the customer. In one embodiment, device 130 may be managed remotely and may not include output device 250. In other words, device 130 may be "headless" and may not include a display, for example.

[0054] Communication interface 260 may include a transceiver (e.g., a transmitter and/or a receiver) that enables device 130 to communicate with other devices and/or systems. Communications interface 260 may communicate via wireless communications (e.g., radio frequency, infrared, and/or visual optics, etc.), wired communications (e.g., conductive wire, twisted pair cable, coaxial cable, transmission line, fiber optic cable, and/or waveguide, etc.), or a combination of wireless and wired communications. Communication interface 260 may include a transmitter that converts baseband signals to radio frequency (RF) signals and/or a receiver that converts RF signals to baseband signals. Communication interface 260 may be coupled to an antenna for transmitting

and receiving signals.

[0055] Communication interface 260 may include a logical component that includes input and/or output ports, input and/or output systems, and/or other input and output components that facilitate the transmission of data to other devices. For example, communication interface 260 may include a network interface card (e.g., Ethernet card) for wired communications and/or a wireless network interface (e.g., a WiFi) card for wireless communications. Communication interface 260 may also include a universal serial bus (USB) port for communications over a cable, a Bluetooth™ wireless interface, a radio-frequency identification (RFID) interface, a near-field communications (NFC) wireless interface, and/or any other type of interface that converts data from one form to another form.

[0056] As described below, device 130 may perform certain operations relating to a tunnel broker configured to establish and manage a communication tunnel based on a requested service property. Device 130 may perform these operations in response to processor 220 executing software instructions contained in a computer-readable medium, such as memory 230. A computer-readable medium includes a non-transitory memory device. A memory device may be implemented within a single physical memory device or spread across multiple physical memory devices. The software instructions may be read into memory 230 from another computer-readable medium or from another device. The software instructions contained in memory 230 may cause processor 220 to perform processes described herein. Alternatively, hardware (e.g., fixed) circuitry may be used in place of, or in combination with, software instructions to implement processes described herein. Thus, implementations described herein are not limited to any specific combination of hardware circuitry and software.

[0057] Although Fig. 2 shows exemplary components of device 130, in other implementations, device 130 may include fewer components, different components, additional components, or differently arranged components than depicted in Fig. 2. Additionally or alternatively, one or more components of device 130 may perform one or more tasks described as performed by one or more other components of device 130. Administration device 150 may be configured similarly as device 130.

[0058] Fig. 3 is a block diagram illustrating exemplary communication layers of device 130. The functional components of device 130 may be implemented, for example, by processor 220 executing instructions from memory 230. Additionally or alternatively, the functional components of device 130 may be implemented via hardwired (e.g., fixed) circuitry of one or more ASICs. As shown in Fig. 3, device 130 may include a service layer 310, an overlay network layer 320, and a device layer 330.

[0059] Service layer 310, in one embodiment, enables clients to search for service instances of a particular service type and enables clients to send requests to particular service instances. A service may be accessed via a

standardized service interface that, in one embodiment, is agnostic to the actual implementation of the service. A service instance may be associated with explicit boundaries. In this embodiment, a particular process running on device 130, and/or particular data stored on device 130, either resides within the service instance or outside of the service instance without ambiguity. A service instance may be autonomous with respect to other service instances. For example, a particular service instance may be modified (e.g., code may be rewritten) without negatively impacting other service instances interacting with the particular service instance. A service may share a schema and/or a contract with other service instance (of the same type or of different type), but, in one embodiment, does not share the service implementation. A schema specifies the format and content of messages sent or received by the service interface. A contract specifies permissible sequences of messages sent or received by the service interface.

[0060] One or more services may be deployed together as a bundle. A bundle may correspond to service that functions as a deployment unit in the system. A node in the system that is able to deploy a particular bundle, corresponding to a grouping of one or more services, functions as a bundle host. A bundle repository service may store a collection of bundles in the system. Thus, when service manager select to deploy a service, the service manager may need to locate a bundle host that is able to deploy a bundle associated with the service. The service manager may contact the service registry to locate the bundle repository service. The service manager may then contact the bundle repository service to identify a bundle. The service manager may select a bundle and may then search the service registry to identify a suitable bundle host that may deploy the selected bundle. The service manager may then contact the bundle host and may instruct the bundle host to deploy the bundle associated with the service.

[0061] Overlay network layer 320, in one embodiment, implements an overlay network on top of an existing network topology. Overlay network layer 320 may be responsible for routing traffic through firewalls and/or dealing with network address translation (NAT) in the underlying network topology. In one embodiment, the overlay network topology (e.g., which may be different than the underlying network topology) includes nodes organized in a tree structure. The overlay network topology logically connects the nodes. In other embodiments, the overlay network topology may include a different type of structure (e.g., a mesh topology). Each service host in a device 130 may correspond to a node in the overlay network and may be assigned a node identifier (ID). As noted above, a device 130 may include multiple service hosts and/or multiple nodes. Device 130 may be described as including one host that corresponds to one node. The nodes may be connected via the network topology, such as a routing tree, and a node may send a message to another node via the routing tree. In one embodiment, a

node may send a message to another node via the underlying network topology without the message traversing the overlay network topology. Each node may store information (e.g., addresses of the underlying network, such as network 110) to reach its neighbors in the overlay network (as well as the underlying network). Overlay network layer 320 may correspond to a communication layer between the nodes and may use multiple network topologies to realize a particular function. For example, when searching service registries for a particular type of service, overlay network layer 320 may traverse edges of a tree of nodes while searching through service registries. In one embodiment, when sending a message from a first node to a second node, overlay network layer 320 may send the message directly from the first node to the second node, rather than by following edges of the tree. Overlay network layer 320 may provide node IDs to service layer 310 and service layer 310 may send messages to particular node IDs without needing to know the underlying network topology.

[0062] In one embodiment, device layer 330 performs device discovery during initial installation of SOA system service bus 140. Device layer 330 and/or overlay network layer 320 may also perform node discovery subsequent to initial installation, and/or may rediscover lost nodes that went offline and that rejoin the overlay network at a later time. In one embodiment, overlay network layer 320 manages a shared secret for the overlay network, such as a certificate, that enables the nodes to verify each other's identity. Overlay network layer 320 may form a topology (e.g., a routing tree or mesh) for the overlay network based on one or more metrics of nearness. However, a message from a first node to a second node need not traverse the routing tree and may instead be sent directly from the first node to the second node. In another embodiment, the message from the first node to the second node traverses the routing tree. Furthermore, overlay network layer 320 may send multicast messages based on multicast groups. Moreover, overlay network layer 320 may provide a quality of service (QoS) guarantee to service layer 310.

[0063] While network layer 320 generally deals with "nodes," device layer 330 generally deals with "devices." Device layer 330 corresponds to the lower levels of functionality of device 130, including functionality required to communicate using the underlying network topology (e.g., network 110 and/or sub-network 120). For example, in some implementations, device layer 330 may implement Layers 1 through 6 of the Open Systems Interconnection (OSI) model (e.g. the Physical layer, Data Link layer, Network layer, Transport layer, Session layer, and Presentation layer). Implementation of these layers may include routing Ethernet frames, routing Internet Protocol (IP) packets, session management, encrypting and decrypting packets, retransmitting lost packets, etc.

[0064] Although Fig. 3 shows exemplary functional components of device 130, in other implementations, device 130 may include fewer functional components, dif-

ferent functional components, differently arranged functional components, or additional functional components than depicted in Fig. 3. Additionally, any one of the components (or any group of components) of device 130 may perform functions described as performed by one or more other functional components of device 130.

[0065] Fig. 4A is a block diagram illustrating exemplary functional components of service layer 310. As shown in Fig. 4A, service layer 310 includes a service host 315. Service host 315 may include one or more services 410-A to 410-N (referred to collectively as "services 410" and individually as "service 410"), one or more clients 420-A to 420-K (referred to collectively as "clients 420" and individually as "client 420"), a message dispatcher 430, and a service registry 440.

[0066] Service 410 corresponds to a service instance associated with service host 315 of service layer 310 of device 130. In one embodiment, service 410 includes a service interface 412 and a service implementation 414. Service interface 412 may include a communication protocol, such as a standardized communication protocol. In one implementation, the communication protocol includes a unique name and version. Service interface 412 may be specified using a Simple Object Access Protocol (SOAP) interface specification, a JavaScript Object Notation (JSON) interface specification, and/or another type of interface specification. Service implementation 414 includes the implementation of service 410. Service implementation 414 processes requests received via service interface 412 and/or responds to service requests through service interface 412. Service interface 412 may convert responses received from service implementation 414 into a particular format compatible with the proper protocol, which client 420 uses to exchange messages with service 410.

[0067] In one embodiment, client 420 requests a service instance of a particular service type by sending a request to service registry 440. Once a service instance is identified and selected, client 420 may send a request to the identified and selected particular service instance via message dispatcher 430. As discussed above, clients 420 may also be services 410. The term "client" or "client service" identifies the service as one that is requesting another service.

[0068] Message dispatcher 430 receives incoming messages from client 420 and directs them to service 410 that is the intended recipient of the incoming message. Furthermore, message dispatcher 430 may receive messages from a service and send the message to a particular client 420. If the destination of the incoming message is not on the same device 130 as message dispatcher 430, then the message may be forwarded to the overlay network layer 320 for forwarding to the correct device 130. Services 410 and clients 420 may function as endpoints in the overlay network implemented by overlay network layer 320. Thus, in one embodiment, overlay network layer 320 may maintain a routing table based on the routing tree of the overlay network. The

routing table may include a list of next hop destinations for particular node IDs. Message dispatcher 430 may identify a next hop destination for the outgoing ID and may provide the message to overlay network layer 320 for delivery. Thus, in this embodiment, message dispatcher 430 implements a request-response messaging mechanism.

[0069] Service registry 440 maintains a list of deployed services 410 along with properties associated with the deployed services (e.g., instances of services). Exemplary components of service registry 440 are described in more detail below with reference to Fig. 4C. A service 410 may register with service registry 440 by providing service registry 440 with a description of the service (e.g., including properties of the service). Because clients 420 may also be services 410, clients 420 may also register with service registry 440.

[0070] Fig. 4B is a block diagram illustrating the functionality of service registry 440. As shown in Fig. 4B, service registry 440 may receive search queries from clients 420. A search query may specify a particular service type, one or more requested properties for the particular service type, a requested number of hits, and/or one or more other parameters. Service registry 440 may identify services 410 that satisfy the search query. If the number of requested hits is not satisfied by service registry 440, service registry 440 may forward a query to another service registry 440 (e.g., an adjacent service registry 440) in the overlay network. In one embodiment, service registry 440 does not select a particular service instance based on a search query. Rather, in this embodiment, service registry 440 returns the results of the query to client 420 and client 420, which originated the query, may select a particular service instance from the search results. In another embodiment, service registry 440 selects the particular service instance based on the search query from the results of the query.

[0071] Although Figs. 4A and 4B show exemplary functional components of service layer 310, in other implementations, service layer 310 may include fewer functional components, different functional components, differently arranged functional components, or additional functional components than depicted in Figs. 4A and 4B. Additionally, any one of the components (or any group of components) of service layer 310 may perform functions described as performed by one or more other functional components of service layer 310.

[0072] Fig. 4C is a block diagram illustrating exemplary functional components of service registry 440. As shown in Fig. 4C, service registry 440 may include a host service registry database (DB) 442, a query handler 444, and a service registry cache 446.

[0073] Host service registry DB 442 may maintain a list of services 410 hosted by service host 315 and/or properties of those services. An example of a service listed in host service registry DB 442 and properties of the service is provided below with respect to Fig. 4D. Host service registry DB 442 may be populated by serv-

ices 410 registering with service registry 440. Host service registry DB 442 may also expose an interface for adding or removing listed services and reading or writing properties of the services hosted by service host 315 and/or write service properties. In one embodiment for example, host service registry DB 442 may maintain a list of services 410 hosted by a service host 315 on a different device 130. The service host 315 on the different device may list its services in a service registry on another device using the exposed interface. Furthermore, host service registry DB 442 may expose a search query service interface accessible by other service registries. Thus, other service registries may use the search query service interface to determine whether host service registry DB 442 includes an entry that satisfies a particular query. In one embodiment, services listed in host service registry DB 442 may expire (e.g., be removed from DB 442 after a period of time if not refreshed) to help prevent DB 442 from storing outdated information.

[0074] Host service registry 442 may receive a subscription request from a service manager, may store the subscription request, and may forward the subscription request to all adjacent service registries.

[0075] Host service registry 442 may determine whether a service matches the subscription request and may send a subscription notification back to a service manager that originated the subscription request if a matching service is identified. Furthermore, host service registry 442 may determine whether an update to a stored service is associated with a subscription. If an update is associated with a subscription, host service registry 442 may send a subscription notification to the service manager (or another type of service) that originated the subscription request for the associated subscription.

[0076] Query handler 444 may handle queries received from client 420. In one embodiment, given a query, query handler 444 first searches the local host service registry DB 442, followed by service registry cache 446. Query handler 444 may issue a call to other service registries if the query has not been satisfied, for example. Service registry cache 446 may store data from remote service registries 440. Each service host 315 may maintain a local service registry 440 and services 410 that register with service host 315 are registered in the local service registry 440. A query from client 420 that cannot be satisfied by the local service registry 440 may be sent to one or more neighboring service hosts 315 to see if the neighboring service hosts 315 have service registries 440 that include services that satisfy the query. The remote service registry 440 may return results of the query back to the local service registry 440 and the results may be stored in service registry cache 446. In some implementations, parent nodes may cache data for their children nodes, while children nodes may not cache data for their parent nodes. In one embodiment, services listed in service registry cache 446 may expire (e.g., be removed from cache 446 after a period of time if not refreshed) to help prevent cache 446 from storing outdated

information.

[0077] Although Fig. 4C shows exemplary functional components of service registry 440, in other implementations, service registry 440 may include fewer functional components, different functional components, differently arranged functional components, or additional functional components than depicted in Fig. 4C. Additionally, any one of the components (or any group of components) of service registry 440 may perform functions described as performed by one or more other functional components of service registry 440.

[0078] Fig. 4D is a block diagram of an exemplary property table 460 for a particular service that may be stored by service registry 440. In one embodiment, an instance of a service (e.g., each instance) is associated with a property table, such as table 460. Host service registry database DB 442 may store a property table for each service registered with the corresponding service registry 440. In one embodiment, as described above, the information stored in any one service registry DB 442 may be different than information stored in other service registry databases. Exemplary property table 460 includes eight fields: ID field 462, interface field 464, service format field 468, transport protocol field 470, CPU ranking 472, disk space field 474, and RAM field 476.

[0079] Instance ID field 462 uniquely defines the instance of the particular service. The instance ID (possibly along with the node ID) may uniquely identify the service instance from any other services (of the same type or different type) in the network. In one embodiment, instance ID field 462 is an integer. In table 460, the instance ID is 6529 as an example.

[0080] Interface field 464 identifies the name of the interface of the service. In this case, the interface field 464 may also identify the type of service by the type of interface. For example, table 460 identifies the interface as "STORAGE SERVICE". Service format field 468 identifies the format used by the instance of the service. As an example, table 460 identifies the service format as "JSON". Transport protocol field 470 identifies the protocol used by the instance of the service. As an example, table 460 identifies the service format as "NODE PROTOCOL".

[0081] CPU ranking field 472 identifies the performance of the CPU associated with the service instance. In one embodiment, a scale is used (e.g., 1 to 100). Table 460 identifies the CPU ranking as 20/100 for the service in CPU ranking field 742. RAM field 476 identifies the amount of random-access memory available to the service. Table 460 identifies the available RAM as 2 GB in field 476.

[0082] Although Fig. 4D shows exemplary components of property table 460, in other implementations, property table 460 may include fewer components, different components, differently arranged components, or additional components than depicted in Fig. 4D.

[0083] Fig. 5A is a block diagram illustrating functional components of overlay network layer 320. As shown in

Fig. 5A, overlay network layer 320 may include a node manager 510, a communication manager 520, and a multicast manager 530.

[0084] Node manager 510 may provide node information, such as a node ID, to other nodes in the overlay network. Furthermore, node manager 510 may maintain a list of nodes in the overlay network. Node manager 510 may perform node discovery to identify new nodes added to the overlay network and/or to rediscover lost nodes that have re-joined the overlay network. Node manager 510 may also determine the topology of the network, as described below (e.g., which nodes are nearest other nodes).

[0085] Communication manager 520 may enable nodes to communicate with each other. Communication manager 520 may implement a mechanism to traverse the tree of the overlay network. Tree traversal may be performed in connection with search queries of service registries or when a direct communication method to another node is not available. Furthermore, communication manager 520 may implement a direct communication method that may enable particular nodes of the overlay network to communicate directly without having to traverse the tree of the overlay network.

[0086] Multicast manager 530 may implement a multicast mechanism. The multicast mechanism may be used to send a message to the members of a multicast group (e.g., all the members). Furthermore, the multicast mechanism may be used to implement a subscribe-notify messaging pattern. Thus, an event associated with a particular service instance may be used to trigger a message sent to the nodes that have subscribed to messages from the particular service instance. Multicast manager 530 may include an application layer multicast manager or a multicast manager from lower OSI layers.

[0087] Although Fig. 5A shows exemplary functional components of overlay network layer 320, in other implementations, overlay network layer 320 may include fewer functional components, different functional components, differently arranged functional components, or additional functional components than depicted in Fig. 5A. Additionally, any one of the components (or any group of components) of overlay network layer 320 may perform functions described as performed by one or more other functional components of overlay network layer 320.

[0088] Fig. 5B is a block diagram of an exemplary topology of an overlay network 540. As shown in the example of Fig. 5B, overlay network 540 includes nodes N1 to N7. Nodes N1 and N2 are in multicast group 560-1. Node N1 includes service endpoints S1 and S3 and client endpoint C1. Node N3 is the parent node to nodes N1 and N2. Node N3 includes a service endpoint S7 and a client endpoint C3.

[0089] Nodes N6 and N7 are in multicast group 560-2 and node N7 includes client endpoint C2 and service endpoints S5 and S6. Node N5 is the parent node to nodes N6 and N7 and includes service endpoint S9.

Nodes N3 and N5 are in multicast group 560-3. Node N4 is the parent node to nodes N3 and N5 and is the root node of overlay network 540. Furthermore, node N4 is in multicast group 560-4 and includes service endpoint S8. Although parent nodes in the topology of network 540 have two child nodes, in other implementations, parent nodes may have more than two child nodes.

[0090] Assuming each service endpoint is associated with a service registry 440, a search query may traverse overlay functional network 540 as follows. Assume service endpoint S7 in node N3 executes a search query to identify a particular service included in service endpoint S1 and service endpoint S5 (i.e. for which S1 and S5 are a match). Service endpoint S7 may send the search query to its local service registry, which may result in no matches in the search query. The local service registry may then identify adjacent service registries in the overlay network, which may include a service registry in node N1 and a service registry in node N4 (node N2 may not include a service registry, since there are no service endpoints associated with node N2). The service registry in node N1 may return a hit identifying service endpoint S1. The service registry in node N4 may return no hits and may forward the search query to its adjacent service registries, which in this case include service registries in nodes N3 and N5. However, since the service registry in node N3 has already processed the search, the search query may only be sent to the service registry in node N5. The service registry at node N5 may come up with no hits and may forward the search query to a service registry at node N7. Node N7 may identify service endpoint S5 as a hit and may return the results of the search query to node N4 and node N4 may forward the search results to service endpoint S7 in node N3.

[0091] Service endpoint S7 may then select communicate with either service endpoint S1 at node N1 or service endpoint S5 at node N7. In some implementations, service endpoint S7 may send a message to service endpoint S5 via nodes N4 and N5. In other implementations, service endpoint S7 may send a message to service endpoint S5 by communicating directly with node N7.

[0092] As another example, service endpoint S7 may only require the first match to the search query. Nodes may forward search queries to other nodes in a priority order that prioritizes nodes that are further down the tree. Thus, node N3 would forward the search query to nodes N1 and N2, before sending the search query to node N4, since nodes N1 and N2 are further down the tree (i.e., are children of node N3), while node N4 is further up the tree (i.e., is a parent of node N3). Since node N1 identifies a match for the search query, and service endpoint S7 only requires one match, the search may terminate before the search query is sent to node N4.

[0093] Fig. 6 is a block diagram illustrating functional components of a tunnel broker 600. Tunnel broker 600 may be configured to establish and manage a communication tunnel based on a requested service property. As shown in Fig. 6, tunnel broker 600 may include a client

interface 610, a service registry interface 620, a node selector 630, a node interface 640, and may communicate with a tunnel DB 650.

[0094] Client interface 610 may communicate with client 420. Client interface 610 may receive a request from client 420 for a communication tunnel to a service having a requested service property and may send an acknowledgement to client 420 that the request has been received. Furthermore, client interface 610 may configure client 420 for a first end of a communication tunnel. For example, client interface 610 may configure a node (e.g., device 130) hosting client 420 to receive and/or send packets, or other types of data units, encapsulated with tunnel headers associated with a communication tunnel.

[0095] Service registry interface 620 may communicate with service registry 440. For example, service registry interface 620 may generate a search query based on a request received by client interface 610 and may send the search query to service registry 440. Furthermore, service registry interface 620 may send a subscription request to service registry 440 to receive updates relating to the search query. Service registry interface 620 may receive search results from service registry 440 and may provide the search results to node selector 630.

[0096] Node selector 630 may select a particular node, and a particular service instance hosted by the node, as a second end of a communication tunnel. Node selector 630 may select the particular node and/or service instance based on search results obtained by service registry interface 620. For example, node selector 630 may rank the search results and may select a node and/or service instance that best matches the requested service property or properties. Additionally or alternatively, node selector 630 may select a node based on additional criteria. For example, node selector 630 may determine one or more connection metrics for the nodes in the received search results. A connection metric may correspond to a measure of the quality of connection between client 420 and each particular node in the received search results. A connection metric may include a total available bandwidth; a percentage bandwidth capacity; a highest, lowest, or average bitrate; a highest available Quality of Service (QoS); whether or not the client and the particular node are in the same domain; whether network address translation (NAT) is required between the client and the particular node; whether a firewall exists between the client and the particular node; and/or based on other types of connection metrics.

[0097] Node interface 640 may communicate with particular nodes in the system. For example, node interface 640 may configure a node (e.g., device 130) for a second end of a communication tunnel. For example, node interface 640 may configure the node to receive and/or send packets, or other types of data units, encapsulated with tunnel headers associated with the communication tunnel.

[0098] Tunnel DB 650 may store information relating to communication tunnels managed by tunnel broker

600. Exemplary information that may be stored in tunnel DB 650 is described below with reference to Fig. 7B.

[0099] Although Fig. 6 shows exemplary functional components of tunnel broker 600, in other implementations, tunnel broker 600 may include fewer functional components, different functional components, differently arranged functional components, or additional functional components than depicted in Fig. 6. Additionally, any one of the components (or any group of components) of tunnel broker 600 may perform functions described as performed by one or more other functional components of tunnel broker 600.

[0100] Fig. 7A is a block diagram illustrating components that may be stored in service registry 440. As shown in Fig. 7A, service registry 440 may include one or more service entries 701. Each service entry 401 may store information relating to a particular service hosted by service host associated with service registry 440. Service entry 401 may include a service field 710, node field 712, properties field 714, deployment field 716, and subscription field 718.

[0101] Service field 710 may identify a particular service associated with the service entry. For example, service field 710 may identify a service interface associated with the particular service. Node Field 712 may identify a particular node (e.g., device 130) associated with the particular service. In some implementations, a first node may maintain a service registry for a second node and may identify services associated with the second node in the service registry. Properties field 714 may store information identifying properties associated with the particular service. For example, properties field 714 may include information identifying a location associated with the service, an operating system associated with the service, a processing load associated with the service, a bandwidth capacity associated with the service, a memory capacity associated with the service, a storage capacity associated with the service, a sub-network and/or network domain associated with the service, a security level associated with the service, a codec type associated with the service, and/or another type of property.

[0102] Deployment field 716 may include information identifying whether the service is deployed or whether the service is available for deployment. Subscription field 718 may include information identifying subscriptions associated with the service. A service may be associated with one or more subscriptions. The subscription information may, for example, identify a particular tunnel broker 600 (e.g., based on a node ID) that has subscribed to notifications about changes to the service. Thus, if the service is deployed, made unavailable, if a property of the service changes, and/or if another type of change is detected, a notification may be sent to tunnel broker 600.

[0103] Although Fig. 7A shows exemplary components of service registry 440, in other implementations, service registry 440 may include fewer components, different components, differently arranged components, or additional components than depicted in Fig. 7A.

[0104] Fig. 7B is a block diagram illustrating components that may be stored in the tunnel DB 660. Tunnel DB 650 may store one or more tunnel records 751. Each tunnel record 751 may store information relating to a particular communication tunnel managed tunnel broker 600. Tunnel record 751 may include a tunnel ID field 760, a service properties field 762, a client field 764, and one or more node fields 770.

[0105] Tunnel ID field 760 may store an identifier that uniquely identifies a particular communication tunnel. Furthermore, tunnel ID field 760 may store tunnel header information (e.g., routing labels) associated with the particular communication tunnel. Service properties field 762 may store information identifying one or more requested service properties, such as a particular service interface, a particular operating system, a particular processing capacity, a particular storage capacity, a particular bandwidth and/or bitrate, a particular location, a particular codec, a particular network domain, a particular security level, and/or another type of service property.

[0106] Client field 764 may store information relating to client 420 that made the request for the communication tunnel. For example, client ID field 764 may store a node ID associated with client 420 in the overlay network. Furthermore, client ID field 764 may store a network address associated with client 420 in the underlying network (e.g., network 110, sub-network 120, etc.).

[0107] Each node field 760 may store information relating to a particular node in the system that has been identified as hosting a service that matches the requirements specified in the service properties field 762 of tunnel record 751. Node field 760 may include a node ID field 772, a properties field 774, and a status field 776.

[0108] Node ID field 772 may store information identifying the particular node. For example, node ID field 764 may store a node ID associated with the particular node in the overlay network. Furthermore, node ID field 764 may store a network address associated with the particular node in the underlying network (e.g., network 110, sub-network 120, etc.).

[0109] Properties field 774 may store information relating to the properties of a service instance, hosted by the particular node, which matches the request associated with the communication tunnel. For example, properties field 774 may include information identifying the service instance, information identifying the service interface of the service instance, and/or one or more properties associated with the service instance. Information in properties field 774 may be updated at particular intervals. For example, tunnel broker 600 may receive updates from service registry 440 at particular intervals based on a subscription request submitted by tunnel-broker 600 to service registry 440.

[0110] Status field 776 may include status information associated with the particular node. For example, the status information may include a search result rank for the particular node, whether or not a second end of a communication tunnel is established for the particular node,

one or more connection metrics associated with a connection from the client to the particular node, and/or other types of status information associated with the particular node.

[0111] Although Fig. 7B shows exemplary components of capabilities DB 640, in other implementations, capabilities DB 640 may include fewer components, different components, differently arranged components, or additional components than depicted in Fig. 7B.

[0112] Fig. 8 is a flowchart of a process for setting up and managing a communication tunnel according to an implementation described herein. In one implementation, the process of Fig. 8 may be performed by tunnel broker 600 in device 130. In other implementations, some or all of the process of Fig. 8 may be performed by another device or a group of devices separate from and/or including tunnel broker 600.

[0113] The process of Fig. 8 may include receiving a request from a client for a service having a requested service property (block 810). For example, a device 130 may include a service acting as a client that requests a service having a particular service property. Client 420 may first contact service registry 440 to request the location of the nearest tunnel broker. Service registry 440 may return a node ID identifying tunnel broker 600 to client 420. Client 420 may subsequently send a request to tunnel broker 600 to establish a communication tunnel to a service with a requested service property. The requested service property may include one or more of a particular service interface, a particular operating system, a particular processing capacity, a particular storage capacity, a particular bandwidth and/or bitrate, a particular location, a particular codec, a particular network domain, a particular security level, and/or another type of service property.

[0114] A search query may be sent to a service registry (block 820) and search results may be received from the service registry (block 830). For example, tunnel broker 600 may receive the request from client 420, may generate a search query based on the requested service properties, and may send the generated search query to the nearest service registry 440. Service registry 440 may evaluate the search query and may return search results that includes a list of one or more nodes in the system that match the requested service property or properties. If service registry 440 returns an empty list, indicating that a service with the requested property is not available in the system, tunnel broker 600 may generate an alert and may send the alert to client 420 and/or to administration device 150.

[0115] In some implementations, service registry 440 may return a list of nodes and/or service instances that most closely matches the requirements, even though none of the nodes and/or service instances satisfy all the requirements. In such situations, tunnel broker 600 may select the node and/or service instance that best matches the requirements associated with the requested service property or properties.

[0116] In some implementations, the client may request a one-to-many communication tunnel. For example, the client may request a communication tunnel to a particular number of service instances having a requested property. As an example, a video monitoring service client may request a communication tunnel to five different video streams from cameras at a particular location. In such implementations, tunnel broker 600 may request a particular number of search results for the search query.

[0117] A first node that hosts a first service instance having the requested service property may be selected based on the received search results (block 840). For example, tunnel broker 600 may select a node, and/or service instance hosted by the node, which best matches the requested service property. In some implementations, tunnel broker 600 may select the first node based on additional criteria. For example, tunnel broker 600 may select the first node based on one or more connection metrics associated with a connection between each particular node in the search results and the client device. Examples of connection metrics include a total available bandwidth for the connection; a percentage bandwidth capacity for the connection; a highest, lowest, or average bitrate for the connection; a highest available Quality of Service (QoS) for the connection; whether or not the client and the particular node are in the same domain; whether network address translation (NAT) is required between the client and the particular node; whether a firewall exists between the client and the particular node; and/or based on other types of connection metrics.

[0118] In situations in which the client requests a one-to-many communication tunnel, tunnel broker 600 may select multiple nodes for multiple second end of the communication tunnel and may individually set up the second ends of the communication tunnel for each selected node.

[0119] A communication tunnel may be established having a first end at the client device and a second end at the selected first node (block 850). For example, tunnel broker 600 may use TSP, or another protocol, to set up a communication tunnel between client 420 and the selected first node. For example, tunnel broker 600 may generate one or more tunnel headers and/or labels, may configure device 130 hosting client 420 to encapsulate packets, or other types of data units, with the generated tunnel headers and/or labels, and may configure the selected first node to encapsulate packets, or other types of data units, with the generated tunnel headers and/or labels. The communication tunnel may include a Multi-Protocol Label Switching (MPLS) tunnel, a Generic Routing Encapsulation (GRE) tunnel, an Internet Protocol (IP) Security (IPSec) tunnel, a Virtual Local Area Network (VLAN) tunnel, a Virtual Private Network (VPN) tunnel, and/or another type of communication tunnel.

[0120] A determination may be made that the communication tunnel should be updated (block 860). As an example, tunnel broker 600 may re-send the search query to service registry 440 at particular intervals and may

receive updated search results from service registry 440. The updated search results may indicate that the first node is not available, that the first node no longer hosts the service (e.g., that the service instance is no longer deployed), that the properties of the hosted service no longer match the requirements of the requested service property, and/or that another system node hosts a service instance that better matches the requirements of the requested service property.

[0121] As another example, tunnel broker 600 may receive an update from service registry 440 without re-sending the search query. For example, tunnel broker 600 may set up a subscription for changes to service matching the search query and service registry 440 may send periodic updates to tunnel broker 600. As yet another example, tunnel broker 600 may receive a message from client 420 that the first node is not reachable via the established communication tunnel or may receive a message from the first node that the first node is no longer hosting the requested service.

[0122] The search query may be re-sent to the service registry (block 870) and updated search results may be received from the service registry (block 880). In situations in which the determination to update the communication tunnel was not made based on an update received from service registry 440, tunnel broker 600 may re-send the search query to service registry 440 and may receive updated search results from service registry 440.

[0123] A second node that hosts a second service instance having the requested service property may be selected based on the updated search results (block 890). For example, tunnel broker 600 may select a second node, and/or service instance hosted by the second node, which best matches the requested service property based on the updates search results, and/or based on additional criteria, such as determined connection metrics for the nodes included in the updated search results.

[0124] The second end of the communication tunnel may be moved from the first node to the second node (block 895). For example, tunnel broker 600 may use TSP, or another protocol, to end the communication tunnel between client 420 and the first node and to set up a communication tunnel between client 420 and the selected second node. For example, tunnel broker 600 may configure the first node to stop encapsulating packets, or other types of data units, with the generated tunnel headers and/or labels associated with the communication tunnel and may configure the selected second node to encapsulate packets, or other types of data units, with the generated tunnel headers and/or labels. Blocks 870, 880, 890, and 895 may be repeated whenever tunnel broker 600 determined that the communication tunnel should be updated.

[0125] Figs. 9A-9C are diagrams of exemplary scenarios of setting up and managing a communication tunnel according to an implementation described herein. Fig. 9A illustrates an overlay network 910 that includes nodes N1, N2, N3, N4, and N5. Node N3 includes a service S3

corresponding to a video monitoring service client 910. Node N4 includes a service S4 corresponding to a tunnel broker service 920. Node N1 includes a service instance S1A, corresponding to a camera service 930 and a service instance S1B, corresponding to a camera service 940. Node N2 includes a service instance S2, corresponding to a camera service 950 and Node N5 includes a service instance S5, corresponding to a camera service 960.

[0126] Video monitoring service client 910 may require a camera service providing a video stream from a particular location, at a particular resolution, and at a particular bitrate. Thus, video monitoring service client 910 may send a request for a communication tunnel to tunnel broker 920 at node N4 and may specify the required service properties in the request. Tunnel broker 920 may generate a search query based on the request and may submit the search query to a service registry at node N4. The service registry at node N4 may forward the search query to service registries at the other nodes and other nodes may continue to forward the search query until a required number of search results are obtained or until all service registries have processed the search query.

[0127] Assume camera services 930, 940 and 950 satisfy the location and resolution requirements, but none of the available camera services fully satisfy the bitrate requirement. Furthermore, assume tunnel broker 920 determines that the connection from node N3 to node N1 is higher in quality than the connection from node N3 to node N2. Thus, tunnel broker service 920 may select node N1. Furthermore, assume camera service 930 has a higher bitrate than camera service 940 and, therefore, tunnel broker service 920 selects camera service 930. Tunnel broker service 920 may then establish tunnel 970 between video monitoring service client 910 and camera service 930. Video monitoring service client 910 may now receive a streaming video signal from camera service 930 via tunnel 970.

[0128] Continuing to Fig. 9B, assume that the bitrate associated with camera service 940 improves as a result of camera service 940 freeing up processing resources. The service registry of node N1 may include a subscription from tunnel broker 920 to receive updated to service properties associated with camera service 930 and camera service 940. Thus, the service registry of node N1 may send a subscription update to tunnel broker 920. Tunnel broker 920 may determine that camera service 940 is now a better match for the communication tunnel and may switch the second end of communication tunnel 970 to camera service 940 to establish communication tunnel 980.

[0129] Continuing to Fig. 9C, assume that node N1 experiences a node failure 990. Video monitoring service client 910 may stop receiving the video stream from camera service 940 and may send a message to tunnel broker service 920, indicating that node N1 has become unavailable. In response, tunnel broker service 920 may select the next best available service based on the most

recent search results associated with the search query based on the communication tunnel request generated by video monitoring service client 910. Tunnel broker service 920 may select camera service 950 at node N2 and may move the second end of the tunnel from node N1 to camera service 950 at node N2 to generate tunnel 995. Thus, video monitor service client 910 may continue to receive a video stream from the specified location at the best available bitrate and resolution.

[0130] In the preceding specification, various preferred embodiments have been described with reference to the accompanying drawings. It will, however, be evident that various modifications and changes may be made thereto, and additional embodiments may be implemented, without departing from the broader scope of the invention as set forth in the claims that follow. The specification and drawings are accordingly to be regarded in an illustrative rather than restrictive sense.

[0131] For example, while series of blocks have been described with respect to Fig. 8, and an order of signal flows have been described with respect to Figs. 9A-9C, the order of the blocks and/or signal flows may be modified in other implementations. Further, non-dependent blocks and/or signal flows may be performed in parallel.

[0132] It will be apparent that systems and/or methods, as described above, may be implemented in many different forms of software, firmware, and hardware in the implementations illustrated in the figures. The actual software code or specialized control hardware used to implement these systems and methods is not limiting of the embodiments. Thus, the operation and behavior of the systems and methods were described without reference to the specific software code--it being understood that software and control hardware can be designed to implement the systems and methods based on the description herein.

[0133] Further, certain portions, described above, may be implemented as a component that performs one or more functions. A component, as used herein, may include hardware, such as a processor, an ASIC, or a FPGA, or a combination of hardware and software (e.g., a processor executing software). The word "exemplary" as used herein means "as an example for illustration."

[0134] It should be emphasized that the terms "comprises" / "comprising" when used in this specification are taken to specify the presence of stated features, integers, steps or components but does not preclude the presence or addition of one or more other features, integers, steps, components or groups thereof.

[0135] The term "logic," as used herein, may refer to a combination of one or more processors configured to execute instructions stored in one or more memory devices, may refer to hardwired circuitry, and/or may refer to a combination thereof. Furthermore, a logic may be included in a single device or may be distributed across multiple, and possibly remote, devices.

[0136] For the purposes of describing and defining the present invention, it is additionally noted that the term

"substantially" is utilized herein to represent the inherent degree of uncertainty that may be attributed to any quantitative comparison, value, measurement, or other representation.

The term "substantially" is also utilized herein to represent the degree by which a quantitative representation may vary from a stated reference without resulting in a change in the basic function of the subject matter at issue.

[0137] No element, act, or instruction used in the present application should be construed as critical or essential to the embodiments unless explicitly described as such. Also, as used herein, the article "a" is intended to include one or more items. Further, the phrase "based on" is intended to mean "based, at least in part, on" unless explicitly stated otherwise.

Claims

1. A method, performed by a computer device (920), comprising:

receiving, by the computer device (920), a request from a client device (910) for a service in a system, the service having a requested service property;

selecting, by the computer device, a first node (930) in the system that hosts a first service instance having the requested service property; establishing, by the computer device (920), a communication tunnel (970) between the client device (910) and the selected first node (930), wherein the communication tunnel (970) includes a first end at the client device (910) and a second end at the first node (930), wherein establishing the communication tunnel (970) includes:

generating a tunnel label; configuring the client device (910) to encapsulate data units with the generated tunnel label; and

configuring the selected first node (930) to encapsulate data units with the generated tunnel label;

selecting, by the computer device (920), a second node (940) in the system that hosts a second service instance having the requested service property; and

moving, by the computer device (920), the second end of the communication tunnel (970) from the first node (930) to the second node (940) transparently with respect to the client device (910), wherein moving the second end of the communication tunnel (970) from the first node (930) to the second node (940) includes:

configuring the selected first node (930) to stop encapsulating data units with the generated tunnel label; and

configuring the selected second node (940) to encapsulate data units with the generated tunnel label.

2. The method of claim 1, further comprising:

sending a search query that specifies the requested service property to a service registry (440), wherein the service registry includes a list of services available in one or more nodes (130) of the system;

receiving search results from the service registry (440), wherein the search results include a list of one or more nodes (130) having the requested service property; and

wherein selecting the first node (930) in the system that hosts the first service instance having the requested service property includes selecting the first node (930) from the list of one or more nodes (130) having the requested service property.

3. The method of claim 2, further comprising:

determining that the communication tunnel (970) should be updated; and wherein selecting the second node (940) in the system that hosts the second instance having the requested service property is based on determining that the communication tunnel (970) should be updated.

4. The method of claim 3, wherein determining that the communication tunnel (970) should be updated includes:

re-sending the search query to the service registry (440) at particular intervals; and receiving updated search results from the service registry (440), wherein the updated search results include an indication that the first node (930) no longer hosts the first service instance having the requested service property or that the first service instance no longer has the requested service property.

5. The method of claim 3, wherein determining that the communication tunnel (970) should be updated includes at least one of:

receiving an indication from the service registry (440) that the first node (930) no longer hosts the first service instance having the requested service property or that the first service instance

- no longer has the requested service property; receiving a message from the client device (910) that the first node (930) is unreachable; or receiving a message from the first node (930) that the first node (930) is unable to provide the first service to the client device (910). 5
6. The method of claim 3, wherein determining that the communication tunnel (970) should be updated includes: 10
- determining that another node in the system is a better match for the requested service property; and 15
- wherein selecting the second node (940) in the system that hosts the second instance having the requested service property is based on determining that another node in the system is a better match for the requested service property. 20
7. The method of claim 2, further comprising:
- determining one or more network connection metrics for a connection from the client device (910) to particular nodes of the nodes included in the list of one or more nodes having the requested service property; and 25
- wherein selecting the first node (930) from the list of one or more nodes having the requested service property is based on the determined one or more network connection metrics. 30
8. A computer device comprising:
- logic configured to implement a tunnel broker (600) configured to: 35
- receive a request from a client device (910) for a service in a system, the service having a requested service property; 40
- select a first node (930) in the system that hosts a first service instance having the requested service property;
- establish a communication tunnel (970) between the client device (910) and the selected first node (930), wherein the communication tunnel (970) includes a first end at the client device (910) and a second end at the first node (930), wherein, when establishing the communication tunnel (970), the logic is further configured to: 45
- generate a tunnel label;
- configure the client device (910) to encapsulate data units with the generated tunnel label; and 50
- configure the selected first node (930) to encapsulate data units with the generated tunnel label; 55
- select a second node (940) in the system that hosts a second service instance having the requested service property; and
- move the second end of the communication tunnel (970) from the first node (930) to the second node (940) transparently with respect to the client device (910), wherein, when moving the second end of the communication tunnel (970) from the first node (930) to the second node (940), the logic is further configured to:
- configure the selected first node (930) to stop encapsulating data units with the generated tunnel label; and
- configure the selected second node (940) to encapsulate data units with the generated tunnel label.
9. The computer device of claim 8, wherein the tunnel broker (600) is further configured to:
- send a search query that specifies the requested service property to a service registry, wherein the service registry includes a list of services available in one or more nodes of the system; receive search results from the service registry, wherein the search results include a list of one or more nodes having the requested service property; and
- wherein, when selecting the first node (930) in the system that hosts the first service instance having the requested service property, the tunnel broker (600) is further configured to:
- select the first node (930) from the list of one or more nodes having the requested service property.
10. The computer device of claim 9, wherein the tunnel broker (600) is further configured to:
- determine that the communication tunnel (970) should be updated; and
- wherein the tunnel broker (600) is configured to select the second node (940) in the system that hosts the second instance having the requested service property based on determining that the communication tunnel (970) should be updated.
11. The computer device of claim 10, wherein when determining that the communication tunnel (970) should be updated, the tunnel broker (600) is further

configured to:

re-send the search query to the service registry at particular intervals; and
 receive updated search results from the service registry, wherein the updated search results include an indication that the first node (930) no longer hosts the first service instance having the requested service property or that the first service instance no longer has the requested service property.

12. The computer device of claim 10, wherein when determining that the communication tunnel (970) should be updated, the tunnel broker (600) is further configured to at least one of:

receive an indication from the service registry that the first node (930) no longer hosts the first service instance having the requested service property or that the first service instance no longer has the requested service property;
 receive a message from the client device (910) that the first node (930) is unreachable; or
 receive a message from the first node (930) that the first node (930) is unable to provide the first service to the client device (910).

13. The computer device of claim 10, wherein, when determining that the communication tunnel (970) should be updated, the tunnel broker (600) is further configured to:

determine that another node in the system is a better match for the requested service property; and
 wherein the tunnel broker (600) is configured to select the second node (940) in the system that hosts the second instance having the requested service property based on determining that another node in the system is a better match for the requested service property.

Patentansprüche

1. Verfahren, ausgeführt von einem Computergerät (920), Folgendes umfassend:

Empfangen einer Anfrage von einem Clientgerät (910) nach einem Dienst in einem System durch das Computergerät (920), wobei der Dienst eine angefragte Diensteseigenschaft aufweist,
 Auswählen eines ersten Knotens (930) im System, der eine erste Dienstinstanz hostet, welche die angefragte Diensteseigenschaft aufweist, durch das Computergerät,

Einrichten eines Kommunikationstunnels (970) zwischen dem Clientgerät (910) und dem ausgewählten Knoten (930) durch das Computergerät (920), wobei der Kommunikationstunnel (970) ein erstes Ende am Clientgerät (910) und ein zweites Ende am ersten Knoten (930) beinhaltet, wobei das Einrichten des Kommunikationstunnels (970) Folgendes beinhaltet:

Erzeugen einer Tunnelkennzeichnung, Konfigurieren des Clientgeräts (910) dafür, Dateneinheiten mit der erzeugten Tunnelkennzeichnung zu verkapseln, und Konfigurieren des ausgewählten ersten Knotens (930) dafür, Dateneinheiten mit der erzeugten Tunnelkennzeichnung zu verkapseln,

Auswählen eines zweiten Knotens (940) im System, der eine zweite Dienstinstanz hostet, welche die angefragte Diensteseigenschaft aufweist, durch das Computergerät (920) und transparentes Bewegen des zweiten Endes des Kommunikationstunnels (970) vom ersten Knoten (930) zum zweiten Knoten (940) in Bezug auf das Clientgerät (910), wobei das Bewegen des zweiten Endes des Kommunikationstunnels (970) vom ersten Knoten (930) zum zweiten Knoten (940) Folgendes beinhaltet:

Konfigurieren des ausgewählten ersten Knotens (930) dafür, das Verkapseln von Dateneinheiten mit der erzeugten Tunnelkennzeichnung zu beenden, und

Konfigurieren des ausgewählten zweiten Knotens (940) dafür, Dateneinheiten mit der erzeugten Tunnelkennzeichnung zu verkapseln.

2. Verfahren nach Anspruch 1, ferner Folgendes umfassend:

Senden einer Suchabfrage, welche die angefragte Diensteseigenschaft spezifiziert, an ein Dienstregister (440), wobei das Dienstregister eine Liste (130) von Diensten beinhaltet, die in einem oder mehreren Knoten des Systems verfügbar sind,

Empfangen von Suchergebnissen vom Dienstregister (440), wobei die Suchergebnisse eine Liste von einem oder mehreren Knoten (130) beinhalten, welche die angefragte Diensteseigenschaft aufweisen, und

wobei das Auswählen des ersten Knotens (930) im System, der die erste Dienstinstanz hostet, welche die angefragte Diensteseigenschaft aufweist, das Auswählen des ersten Knotens (930) von der Liste eines oder mehrerer Knoten (130)

- beinhaltet, welche die angefragte Dienst-eigenschaft aufweisen.
3. Verfahren nach Anspruch 2, ferner Folgendes um-fassend: 5
- Bestimmen, dass der Kommunikationstunnel (970) aktualisiert werden soll, und wobei das Auswählen des zweiten Knotens (940) im System, der die zweite Dienstinstanz hostet, welche die angefragte Dienst-eigenschaft aufweist, auf der Bestimmung basiert, dass der Kommunikationstunnel (970) aktuali-siert werden soll. 10
4. Verfahren nach Anspruch 3, wobei die Bestimmung, dass der Kommunikationstunnel (970) aktualisiert werden soll, Folgendes beinhaltet: 15
- erneutes Senden der Suchabfrage an das Dienstregister (440) in bestimmten Intervallen und Empfangen aktualisierter Suchergebnisse vom Dienstregister (440), wobei die aktualisierten Suchergebnisse einen Hinweis beinhalten, dass der erste Knoten (930) die erste Dienstinstanz, welche die angefragte Dienst-eigenschaft aufweist, nicht mehr hostet oder dass die erste Dienstinstanz nicht mehr die angefragte Dienst-eigenschaft aufweist. 20
5. Verfahren nach Anspruch 3, wobei die Bestimmung, dass der Kommunikationstunnel (970) aktualisiert werden soll, mindestens eines der Folgenden bein-haltet: 25
- Empfangen eines Hinweises vom Dienstregis-ter (440), dass der erste Knoten (930) die erste Dienstinstanz, welche die angefragte Dienst-eigenschaft aufweist, nicht mehr hostet oder dass die erste Dienstinstanz nicht mehr die angefragte Dienst-eigenschaft aufweist, 30
- Empfangen einer Nachricht vom Clientgerät (910), dass der erste Knoten (930) nicht erreichbar ist, oder 40
- Empfangen einer Nachricht vom ersten Knoten (930), dass der erste Knoten (930) nicht in der Lage ist, dem Clientgerät (910) den ersten Dienst zur Verfügung zu stellen. 45
6. Verfahren nach Anspruch 3, wobei die Bestimmung, dass der Kommunikationstunnel (970) aktualisiert werden soll, Folgendes beinhaltet: 50
- Bestimmen, dass ein anderer Knoten im System eine bessere Übereinstimmung für die ange-fragte Dienst-eigenschaft ist, und wobei das Auswählen des zweiten Knotens (940) im System, der die zweite Instanz hostet, welche die angefragte Dienst-eigenschaft aufweist, auf der Bestimmung basiert, dass ein an-derer Knoten im System eine bessere Überein-stimmung für die angefragte Dienst-eigenschaft ist. 55
7. Verfahren nach Anspruch 2, ferner Folgendes um-fassend: 10
- Bestimmen einer oder mehrerer Netzwerkver-bindungsmetriken für eine Verbindung vom Cli-entgerät (910) zu bestimmten Knoten der Kno-ten, die in der Liste eines oder mehrerer Knoten enthalten sind, welche die angefragte Dienst-eigenschaft aufweisen, und wobei das Auswählen des ersten Knotens (930) aus der Liste eines oder mehrerer Knoten, wel-che die angefragte Dienst-eigenschaft aufwei-sen, auf der Bestimmung einer oder mehrerer Netzwerkverbindungsmetriken basiert. 15
8. Computergerät, Folgendes umfassend: 20
- eine Logik, die dafür konfiguriert ist, einen Tun-nelbroker (600) auszuführen, der für Folgendes konfiguriert ist: 25
- Empfangen einer Anfrage von einem Client-gerät (910) nach einem Dienst in einem System, wobei der Dienst eine angefragte Dienst-eigenschaft aufweist, Auswählen eines ersten Knotens (930) im System, der eine erste Dienstinstanz hos-tet, welche die angefragte Dienst-eigen-schaft aufweist, Einrichten eines Kommunikationstunnels (970) zwischen dem Clientgerät (910) und dem ausgewählten Knoten (930), wobei der Kommunikationstunnel (970) ein erstes En-de am Clientgerät (910) und ein zweites En-de am ersten Knoten (930) beinhaltet, wo-bei die Logik beim Einrichten des Kommu-nikationstunnels (970) für Folgendes konfi-guriert ist: 30
- Erzeugen einer Tunnelkennzeichnung, Konfigurieren des Clientgeräts (910) dafür, Dateneinheiten mit der erzeug-ten Tunnelkennzeichnung zu verkap-seln, und Konfigurieren des ausgewählten ers-ten Knotens (930) dafür, Dateneinhei-ten mit der erzeugten Tunnelkenn-zeichnung zu verkapiteln, 35
- Auswählen eines zweiten Knotens (940) im System, der eine zweite Dienstinstanz hos-

tet, welche die angefragte Diensteigenschaft aufweist, und transparentes Bewegen des zweiten Endes des Kommunikationstunnels (970) vom ersten Knoten (930) zum zweiten Knoten (940) in Bezug auf das Clientgerät (910), wobei die Logik beim Bewegen des zweiten Endes des Kommunikationstunnels (970) vom ersten Knoten (930) zum zweiten Knoten (940) für Folgendes konfiguriert ist:

Konfigurieren des ausgewählten ersten Knotens (930) dafür, das Verkapseln von Dateneinheiten mit der erzeugten Tunnelkennzeichnung zu beenden, und Konfigurieren des ausgewählten zweiten Knotens (940) dafür, Dateneinheiten mit der erzeugten Tunnelkennzeichnung zu verkapseln.

9. Computergerät nach Anspruch 8, wobei der Tunnelbroker (600) ferner für Folgendes konfiguriert ist:

Senden einer Suchabfrage, welche die angefragte Diensteigenschaft spezifiziert, an ein Dienstregister, wobei das Dienstregister eine Liste von Diensten beinhaltet, die in einem oder mehreren Knoten des Systems verfügbar sind, Empfangen von Suchergebnissen vom Dienstregister, wobei die Suchergebnisse eine Liste von einem oder mehreren Knoten beinhalten, welche die angefragte Diensteigenschaft aufweisen, und wobei der Tunnelbroker (600) beim Auswählen des ersten Knotens (930) im System, der die erste Dienstinstanz hostet, welche die angefragte Diensteigenschaft aufweist, ferner für Folgendes konfiguriert ist:

Auswählen des ersten Knotens (930) von der Liste eines oder mehrerer Knoten, welche die angefragte Diensteigenschaft aufweisen.

10. Computergerät nach Anspruch 9, wobei der Tunnelbroker (600) ferner für Folgendes konfiguriert ist:

Bestimmen, dass der Kommunikationstunnel (970) aktualisiert werden soll, und wobei der Tunnelbroker (600) dafür konfiguriert ist, basierend auf der Bestimmung, dass der Kommunikationstunnel (970) aktualisiert werden soll, den zweiten Knotens (940) im System auszuwählen, der die zweite Dienstinstanz hostet, welche die angefragte Diensteigenschaft aufweist.

11. Computergerät nach Anspruch 10, wobei der Tunnelbroker (600) bei der Bestimmung, dass der Kommunikationstunnel (970) aktualisiert werden soll, ferner für Folgendes konfiguriert ist:

erneutes Senden der Suchabfrage an das Dienstregister in bestimmten Intervallen und Empfangen aktualisierter Suchergebnisse vom Dienstregister, wobei die aktualisierten Suchergebnisse einen Hinweis beinhalten, dass der erste Knoten (930) die erste Dienstinstanz, welche die angefragte Diensteigenschaft aufweist, nicht mehr hostet oder dass die erste Dienstinstanz nicht mehr die angefragte Diensteigenschaft aufweist.

12. Computergerät nach Anspruch 10, wobei der Tunnelbroker (600) bei der Bestimmung, dass der Kommunikationstunnel (970) aktualisiert werden soll, ferner für mindestens eines des Folgenden konfiguriert ist:

Empfangen eines Hinweises vom Dienstregister, dass der erste Knoten (930) die erste Dienstinstanz, welche die angefragte Diensteigenschaft aufweist, nicht mehr hostet oder dass die erste Dienstinstanz nicht mehr die angefragte Diensteigenschaft aufweist, Empfangen einer Nachricht vom Clientgerät (910), dass der erste Knoten (930) nicht erreichbar ist, oder Empfangen einer Nachricht vom ersten Knoten (930), dass der erste Knoten (930) nicht in der Lage ist, dem Clientgerät (910) den ersten Dienst zur Verfügung zu stellen.

13. Computergerät nach Anspruch 10, wobei der Tunnelbroker (600) bei der Bestimmung, dass der Kommunikationstunnel (970) aktualisiert werden soll, ferner für Folgendes konfiguriert ist:

Bestimmen, dass ein anderer Knoten im System eine bessere Übereinstimmung für die angefragte Diensteigenschaft ist, und wobei der Tunnelbroker (600) dafür konfiguriert ist, basierend auf der Bestimmung, dass ein anderer Knoten im System eine bessere Übereinstimmung für die angefragte Diensteigenschaft ist, den zweiten Knoten (940) im System auszuwählen, der die zweite Instanz hostet, welche die angefragte Diensteigenschaft aufweist.

Revendications

1. Procédé, effectué par un dispositif informatique (920), comprenant :

la réception, par le dispositif informatique (920), d'une demande en provenance d'un dispositif client (910) pour un service dans un système, le service comprenant une propriété de service demandée ;

la sélection, par le dispositif informatique, d'un premier noeud (930) dans le système qui héberge une première instance de service comprenant la propriété de service demandée ;

l'établissement, par le dispositif informatique (920), d'un tunnel de communication (970) entre le dispositif client (910) et le premier noeud sélectionné (930), dans lequel le tunnel de communication (970) comprend une première extrémité au dispositif client (910) et une deuxième extrémité au premier noeud (930), dans lequel l'établissement du tunnel de communication (970) comprend :

la génération d'une étiquette de tunnel ;
la configuration du dispositif client (910) pour encapsuler des unités de données avec l'étiquette de tunnel générée ; et
la configuration du premier noeud sélectionné (930) pour encapsuler des unités de données avec l'étiquette de tunnel générée ;

la sélection, par le dispositif informatique (920), d'un deuxième noeud (940) dans le système qui héberge une deuxième instance de service comprenant la propriété de service demandée ; et

le déplacement, par le dispositif informatique (920), de la deuxième extrémité du tunnel de communication (970) du premier noeud (930) au deuxième noeud (940) de manière transparente par rapport au dispositif client (910), dans lequel le déplacement de la deuxième extrémité du tunnel de communication (970) du premier noeud (930) au deuxième noeud (940) comprend :

la configuration du premier noeud sélectionné (930) pour arrêter l'encapsulation d'unités de données avec l'étiquette de tunnel générée ; et
la configuration du deuxième noeud sélectionné (940) pour encapsuler des unités de données avec l'étiquette de tunnel générée.

2. Procédé selon la revendication 1, comprenant en outre :

l'envoi d'une interrogation de recherche qui spécifie la propriété de service demandée à un registre de services (440), dans lequel le registre de services comprend une liste de services disponibles dans un ou plusieurs noeuds (130) du

système ;

la réception de résultats de recherche provenant du registre de services (440), dans lequel les résultats de recherche comprennent une liste d'un ou plusieurs noeuds (130) comprenant la propriété de service demandée ; et

dans lequel la sélection du premier noeud (930) dans le système qui héberge la première instance de service comprenant la propriété de service demandée comprend la sélection du premier noeud (930) dans la liste d'un ou plusieurs noeuds (130) comprenant la propriété de service demandée.

3. Procédé selon la revendication 2, comprenant en outre :

la détermination que le tunnel de communication (970) doit être mis à jour ; et
dans lequel la sélection du deuxième noeud (940) dans le système qui héberge la deuxième instance comprenant la propriété de service demandée est basée sur la détermination que le tunnel de communication (970) doit être mis à jour.

4. Procédé selon la revendication 3, dans lequel la détermination que le tunnel de communication (970) doit être mis à jour comprend :

le renvoi de l'interrogation de recherche au registre de services (440) à des intervalles particuliers ; et

la réception de résultats de recherche mis à jour en provenance du registre de services (440), dans lequel les résultats de recherche mis à jour comprennent une indication que le premier noeud (930) n'héberge plus la première instance de service comprenant la propriété de service demandée ou que la première instance de service ne comprend plus la propriété de service demandée.

5. Procédé selon la revendication 3, dans lequel la détermination que le tunnel de communication (970) doit être mis à jour comprend au moins l'une de :

la réception d'une indication en provenance du registre de services (440) que le premier noeud (930) n'héberge plus la première instance de service comprenant la propriété de service demandée ou que la première instance de service ne comprend plus la propriété de service demandée ;

la réception d'un message en provenance du dispositif client (910) que le premier noeud (930) est injoignable ; ou

la réception d'un message en provenance du

- premier noeud (930) que le premier noeud (930) est incapable de fournir le premier service au dispositif client (910).
6. Procédé selon la revendication 3, dans lequel la détermination que le tunnel de communication (970) doit être mis à jour comprend :
- la détermination qu'un autre noeud dans le système est une meilleure correspondance pour la propriété de service demandée ; et dans lequel la sélection du deuxième noeud (940) dans le système qui héberge la deuxième instance comprenant la propriété de service demandée est basée sur la détermination qu'un autre noeud dans le système est une meilleure correspondance pour la propriété de service demandée.
7. Procédé selon la revendication 2, comprenant en outre :
- la détermination d'une ou plusieurs métriques de connexion de réseau pour une connexion du dispositif client (910) à des noeuds particuliers parmi les noeuds inclus dans la liste d'un ou plusieurs noeuds comprenant la propriété de service demandée ; et dans lequel la sélection du premier noeud (930) dans la liste d'un ou plusieurs noeuds comprenant la propriété de service demandée est basée sur l'une ou plusieurs métriques de connexion de réseau déterminées.
8. Dispositif informatique comprenant :
- une logique configurée pour mettre en oeuvre un courtier de tunnel (600) configuré pour effectuer :
- la réception d'une demande en provenance d'un dispositif client (910) pour un service dans un système, le service comprenant une propriété de service demandée ; la sélection d'un premier noeud (930) dans le système qui héberge une première instance de service comprenant la propriété de service demandée ; l'établissement d'un tunnel de communication (970) entre le dispositif client (910) et le premier noeud sélectionné (930),
- dans lequel le tunnel de communication (970) comprend une première extrémité au dispositif client (910) et une deuxième extrémité au premier noeud (930), dans lequel, lors de l'établissement du tunnel de communication (970), la logique est en outre configurée pour effectuer :
- la génération d'une étiquette de tunnel ; la configuration du dispositif client (910) pour encapsuler des unités de données avec l'étiquette de tunnel générée ; et la configuration du premier noeud sélectionné (930) pour encapsuler des unités de données avec l'étiquette de tunnel générée ; la sélection d'un deuxième noeud (940) dans le système qui héberge une deuxième instance de service comprenant la propriété de service demandée ; et le déplacement de la deuxième extrémité du tunnel de communication (970) du premier noeud (930) au deuxième noeud (940) de manière transparente par rapport au dispositif client (910), dans lequel, lors du déplacement de la deuxième extrémité du tunnel de communication (970) du premier noeud (930) au deuxième noeud (940), la logique est en outre configurée pour effectuer :
- la configuration du premier noeud sélectionné (930) pour arrêter l'encapsulation d'unités de données avec l'étiquette de tunnel générée ; et la configuration du deuxième noeud sélectionné (940) pour encapsuler des unités de données avec l'étiquette de tunnel générée.
9. Dispositif informatique selon la revendication 8, dans lequel le courtier de tunnel (600) est en outre configuré pour effectuer :
- l'envoi d'une interrogation de recherche qui spécifie la propriété de service demandée à un registre de services, dans lequel le registre de services comprend une liste de services disponibles dans un ou plusieurs noeuds du système ; la réception de résultats de recherche provenant du registre de services, dans lequel les résultats de recherche comprennent une liste d'un ou plusieurs noeuds comprenant la propriété de service demandée ; et dans lequel, lors de la sélection du premier noeud (930) dans le système qui héberge la première instance de service comprenant la propriété de service demandée, le courtier de tunnel (600) est en outre configuré pour effectuer :
- la sélection du premier noeud (930) dans la liste d'un ou plusieurs noeuds comprenant la propriété de service demandée.
10. Dispositif informatique selon la revendication 9, dans lequel le courtier de tunnel (600) est en outre configuré pour effectuer :

la détermination que le tunnel de communication (970) doit être mis à jour ; et dans lequel le courtier de tunnel (600) est configuré pour effectuer la sélection du deuxième noeud (940) dans le système qui héberge la deuxième instance comprenant la propriété de service demandée sur la base de la détermination que le tunnel de communication (970) doit être mis à jour.

5

10

11. Dispositif informatique selon la revendication 10, dans lequel, lors de la détermination que le tunnel de communication (970) doit être mis à jour, le courtier de tunnel (600) est en outre configuré pour effectuer :

15

le renvoi de l'interrogation de recherche au registre de services à des intervalles particuliers ; et

la réception de résultats de recherche mis à jour en provenance du registre de service, dans lequel les résultats de recherche mis à jour comprennent une indication que le premier noeud (930) n'héberge plus la première instance de service comprenant la propriété de service demandée ou que la première instance de service ne comprend plus la propriété de service demandée.

20

25

12. Dispositif informatique selon la revendication 10, dans lequel, lors de la détermination que le tunnel de communication (970) doit être mis à jour, le courtier de tunnel (600) est en outre configuré pour effectuer au moins l'une de :

30

35

la réception d'une indication en provenance du registre de services que le premier noeud (930) n'héberge plus la première instance de service comprenant la propriété de service demandée ou que la première instance de service ne comprend plus la propriété de service demandée ; la réception d'un message en provenance du dispositif client (910) que le premier noeud (930) est injoignable ; ou

40

la réception d'un message en provenance du premier noeud (930) que le premier noeud (930) est incapable de fournir le premier service au dispositif client (910).

45

13. Dispositif informatique selon la revendication 10, dans lequel, lors de la détermination que le tunnel de communication (970) doit être mis à jour, le courtier de tunnel (600) est en outre configuré pour effectuer :

50

55

la détermination qu'un autre noeud dans le système est une meilleure correspondance pour la propriété de service demandée ; et

dans lequel le courtier de tunnel (600) est configuré pour effectuer la sélection du deuxième noeud (940) dans le système qui héberge la deuxième instance comprenant la propriété de service demandée sur la base de la détermination qu'un autre noeud dans le système est une meilleure correspondance pour la propriété de service demandée.




FIG. 1




FIG. 2




FIG. 3




FIG. 4A




FIG. 4B





FIG. 4C

PROPERTY TABLE
460



INSTANCE ID <u>462</u>	6529
INTERFACE <u>464</u>	STORAGE SERVICE
SERVICE FORMAT <u>468</u>	JSON
TRANSPORT PROTOCOL <u>470</u>	NODE PROTOCOL
CPU RANKING <u>472</u>	20/100
DISK SPACE <u>474</u>	1 TB
RAM <u>476</u>	2 GB

FIG. 4D

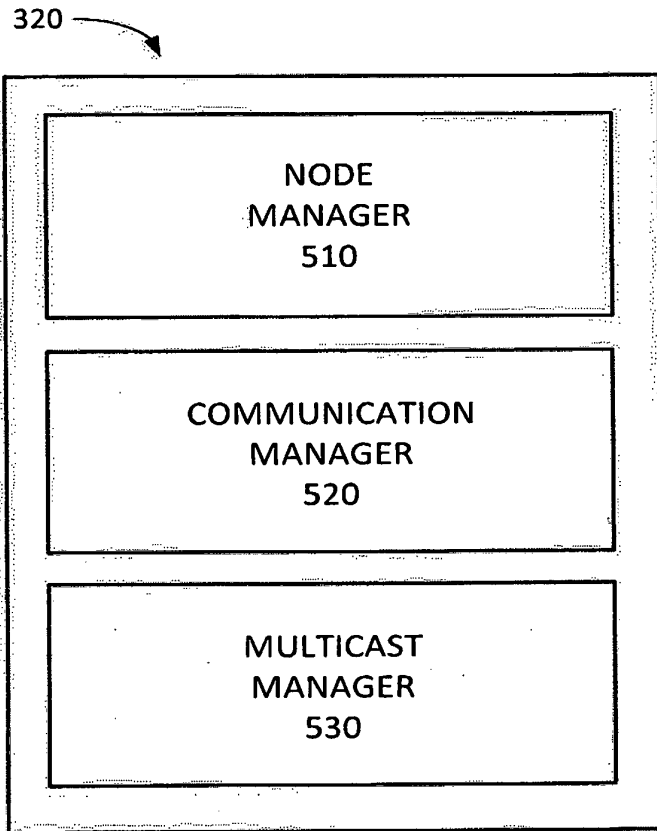


FIG. 5A




FIG. 5B




FIG. 6




FIG. 7A




FIG. 7B

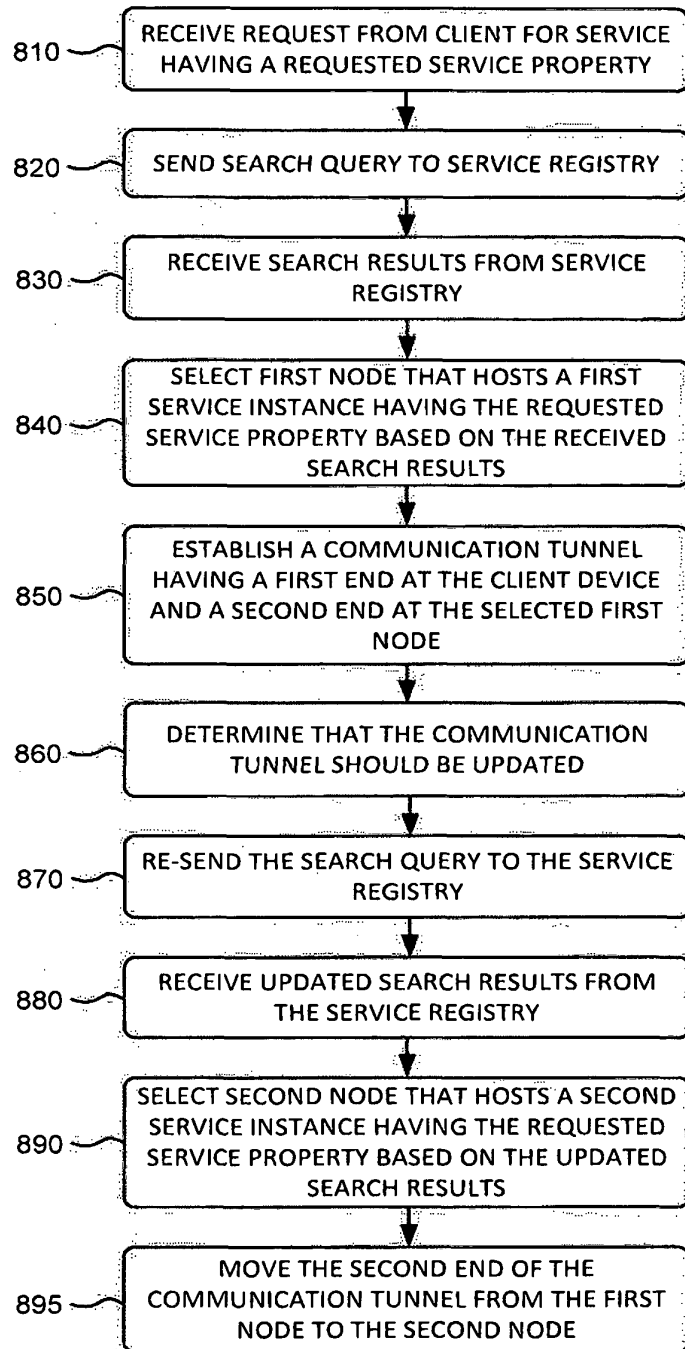


FIG. 8

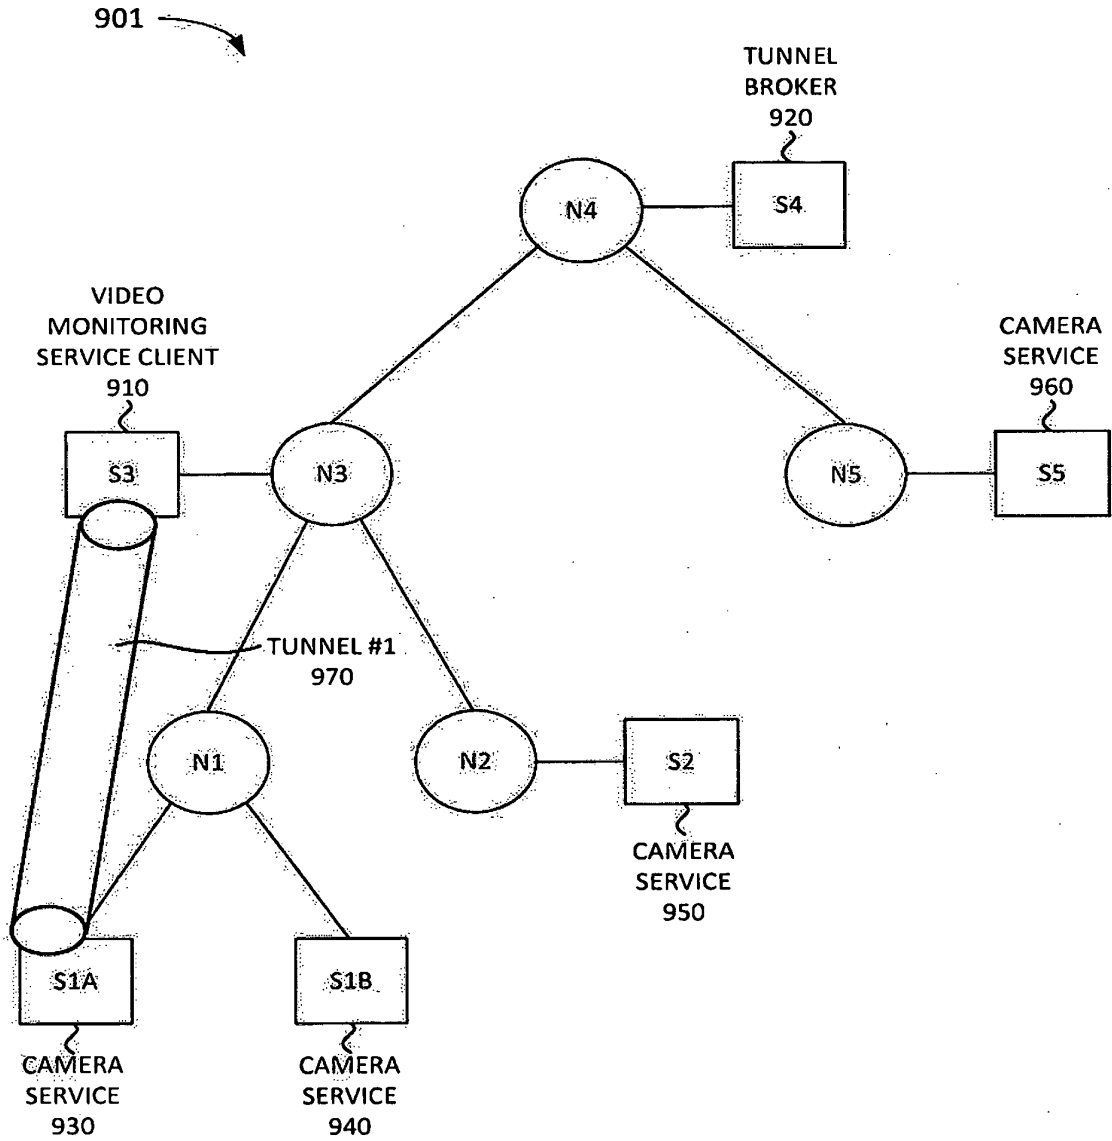


FIG. 9A




FIG. 9B




FIG. 9C

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- US 7779086 B1 [0003]
- US 2002087707 A1 [0004]

PATENT COOPERATION TREATY
PCT
THIRD PARTY OBSERVATION
(PCT Administrative Instructions Part 8)

Applicant's or agent's file reference HOLA-007-PCT	
International application number PCT/IL2018/050910	International filing date (day/month/year) 16 Aug 2018 (16/08/2018)
Applicant LUMINATI NETWORKS LTD.	
Third party observation submitted by Jurate BREIMELYTE	Observation submitted on behalf of Teso LT, UAB
Date of submission(day/month/year) 21 Jun 2019 (21/06/2019)	Language of observation English
<p>Basis and contents of observation</p> <p>1. The observation is made on the basis of the claims in the international application as filed.</p> <p>2. The observation comprises: References to documents: 5 Uploaded copies of documents: 5</p> <p>3. Further explanations: Uploaded copies of documents: 0</p>	

Citation # 1(Other) (# uploaded documents:1):

Identification of Document: the Internet Engineering Task Force draft documents IPv6 Tunnel Broker	Publication Date: 02 Apr 1999 (02/04/1999)
Link to document: https://datatracker.ietf.org/doc/html/draft-ietf-ngtrans-broker-00.txt	
DOI:	
Most relevant passages or drawings: Appendix section A.3 User, Tunnel and Tunnel Server management	Relevant to Claims: 1-15
Brief explanation of relevance: <p>We believe that the Claims 1-15 made in this patent application are identical and/or obvious consequences of certain prior art which we will further present. The main source of prior art we will use further is an identical setup described in the Internet Engineering Task Force draft documents "IPv6 Tunnel Broker" which had 6 versions officially released starting from the version 00 dated April 1999 and which later became RFC 3053 dated January 2001. As the draft language evolved certain items which were too obvious were removed from it and new ideas were added to it so we will use several versions of this document which we believe best illustrate the prior art relevant to specific claims of this patent application.</p> <p>Draft 0 describes the "tunnel server database" (Appendix section A.3 "User, Tunnel and Tunnel Server management"). This part of the prior art documents is relevant to the Claims 2-15. In Claim 2 as filed, the method of sending, receiving information between tunnel devices is disclosed. However, the same method was already described in the relevant prior art. Analyzing this claim and already having established that tunnel broker (hereinafter - TB) acts as first server and tunnel server (hereinafter referred as TS) acts as tunnel device it becomes obvious that the main idea of this claim is identical to the information contained in prior art document's Appendix section A.3:</p> <ul style="list-style-type: none"> • "The Tunnel Server database has one entry for each Tunnel Server; each entry has the following fields: ... IPv4 ... • The TB manages the service updating these databases". <p>As the TB and TS have been used for quite some time, the actual means how information about tunnel devices is obtained is irrelevant. It is obvious for a person skilled in the art, that sending and receiving a message with tunnel devices' data is a very straightforward way to get that information. The same section of the mentioned prior art source is relevant to Claims 3-15. The Appendix section A.3 "User, Tunnel and Tunnel Server management" indicates some of the possible fields which may be stored in tunnel server database about each tunnel server (same to Claim 3, 8, 10 of the present patent application). Claim 14 is obvious and does not constitute any practical addition to Claim 1. If the first and second messages would not contain a content identifier the tunnel device would not be able to identify the web server to which it needs to send an information request. Claims 4-7, 9, 11-13, 15 are obvious and dependent from the other, not new and obvious, claims. They do not create any novelty or inventive step that could get patent protection.</p>	

Citation # 2(Other) (# uploaded documents:1):

Identification of Document: the Internet Engineering Task Force draft documents IPv6 Tunnel Broker (version RFC 3053)	Publication Date: Jan 2001 (01/2001)
Link to document: https://datatracker.ietf.org/doc/html/rfc3053	
DOI:	
Most relevant passages or drawings: Section 2-4	Relevant to Claims: 1-15
Brief explanation of relevance: <p>Claim 1 should not be accepted because it is not new and is obvious compared to the relevant prior art. The prior art document Internet Engineering Task Force draft documents IPv6 Tunnel Broker (version RFC 3053) [hereinafter - RFC 3053) describes the general operation of "Tunnel broker" (hereinafter referred as TB) infrastructure.</p> <p>When comparing TB description to the elements in the prior art and Claim 1 of the present patent application, it is clear that both documents describe the same invention. In the prior art document, the TB acts as the first server (see section 2.1 "Tunnel Broker"), dual-stack node acts as a second server, tunnel server (hereinafter - TS) acts as a tunnel device (see section 2.2 "Tunnel server"). According to section 2.3 of the prior art document, the client can function as a standalone host or router. In the latter case, it is obvious for a person skilled in the art that client device according to this claim is any device connected to a standalone router.</p> <p>Analyzing all the steps of the claim in such setup the following conclusions can be drawn:</p> <ul style="list-style-type: none"> • It is obvious that the client device sends a request message to the second server with the identifier of the content which it receives. • The second server sends the first message to the first server which it receives (first part of section 2.3 "Using the Tunnel Broker: ...it should provide at least the following information...") • The first server selects a tunnel device from the list of tunnel devices and sends a second message using the selected IP address of the selected tunnel device which it receives (described in the second part of section 2.3 "Using the Tunnel Broker: The TB manages the client requests as follows") • The remaining steps are not described in the RFC in details but are obvious from figure 1 of the prior art source "Figure 1: The Tunnel Broker model" and the purpose of the tunnel itself (a reference in section 2 "Tunnel broker model"). Prior art document describes: "Tunnel brokers can be seen as virtual IPv6 ISPs, providing IPv6 connectivity to users already connected to the IPv4 Internet". Another relevant section is 4 "Use of the tunnel broker concept in other areas", which describes "The Tunnel Broker approach might be efficiently exploited also to automatically set-up and manage any other kind of tunnel)". <p>Claims 2-15 are dependent claims. They do not create any novelty or inventive step that could get patent protection. Therefore, at least the claims 1-15 should be rejected as not novel and obvious.</p>	

Citation # 3 (Patent/utility model) (# uploaded documents: 1):

Country code: US	Publication number: 20130080575	Document kind code: A1	
Patent Applicant/Patent Owner: Cloudflare Inc.	Title of invention: Distributing transmission of requests across multiple ip addresses of a proxy server in a cloud-base		
Link to document:			
Publication Date: 28 Mar 2013 (28/03/2013)	Filing Date: 27 Sep 2011 (27/09/2011)	Priority Date:	
Source of Abstract:	Accession number:	Publication Date of Abstract:	Retrieval Date of Abstract:

Most relevant passages or drawings: Paras. 15, 16-20, 26-28, 39, 46, 77, 90, 29-31, 73-75, 85-87, 98, 99, Fig. 1 and 5, Claims.	Relevant to Claims: 1, 2-6, 14
--	-----------------------------------

Brief explanation of relevance:

The claims of the current patent application are not new neither involve an inventive step. Identical method for content fetching is disclosed in the patent application US20130080575(A1) (hereinafter - '575). The publication discloses a method for fetching a first content over the Internet identified by a first content identifier by disclosing fetching "network resources" from origin servers, where the content may include "HTML pages, images, word processing documents, PDF files, movie files, music files, or other computer files." See, e.g., '575 publication, ¶¶ 26, 27; see also id. ¶¶ 15, 17 ("requested content"), 18 (same), 19, also Claim 1. The '575 publication discloses the content is identified in requests from client, for instance in HTTP requests sent to proxy server. See id. ¶¶ 45 (client request includes parameters such as "the type of requested content"), 46, 27, 39; Fig. 1. Identical content fetching method is disclosed in the current patent application claim 1.

The '575 publication discloses sending the first identifier to the first server by disclosing that client sends its IP address to DNS system when resolving a proxy server that will retrieve content from the original server. See, e.g., '575 publication, ¶¶ 15, 77 ("client device requests an IPv4 address for example.com (thus, the client device is an IPv4 enabled client)", 90 (same); Figs. 5 (step 510), 6 (step 610).

The '575 publication discloses sending a first request to the first server by disclosing that client sends a request to DNS system to provide a proxy server that will retrieve content from original server. See, e.g., '575 publication, ¶¶ 15, 77 ("client device 110A requests an IPv4 address for example.com"), 90 (same); Figs. 5 (step 510), 6 (step 610).

The '575 publication discloses receiving the second identifier from the first server by disclosing that client receives an IP address of server from DNS system in response to client's request for an address of a proxy. See, e.g., '575 publication, ¶¶ 15, 28 77 ("DNS system 140 returns to the client device 110A an IPv4 address that is mapped to the record for the requested domain. The IPv4 address is an address of the proxy server 120."), 90 (same); Figs. 5 (step 512), 6 (step 612).

The '575 publication discloses sending a second request to the second device using the second identifier, the second request including the first content identifier and the third identifier, by disclosing client sends a request to proxy server identified by DNS system. The request includes identification of the content type and the domain of origin server. See, e.g., '575 publication, ¶¶ 15, 28, 37, 45 ("a proxy server receives an incoming request from a client device and extracts a set of one or more parameters related to the request. For example, the set of parameters may include one or more of the following: the source IP address of the incoming request, the requested domain (e.g., as indicated in the Host header of an HTTP request) of the incoming request, the source port of the incoming request, cookie(s) of the incoming request, session identifier(s) of the incoming request (if different than the cookie(s)), the type of requested content (e.g., HTML, image, video, etc.), or any combination thereof.") (emphasis added), 78, 91 (same); Figs. 5 (step 514), 6 (step 614).

614). Prior art '575 [in Claim 1] discloses selecting a first one of a plurality of IP addresses of a same

protocol type of the proxy server for use as a source IP address for a second packet that carries an outgoing request, and transmitting a second packet that hosts the identified resource. A similar message are described in claim 14 comprising the content identifiers of the present application.

The '575 publication discloses the feature or receiving content, by disclosing that the selected proxy server sends the content retrieved from origin server in response to the requesting client. See, e.g., '575 publication, ¶¶ 16-20, 29-31, 73-75, 85-87, 98, 99; Figs. 5 (step 530), 6 (step 630).

And similarly to Claim 5 of the present application, '575 discloses [para. 78, 81, 88, 94, 100] that the IPv4 TCP connection may already be established between the proxy server and the origin server (e.g., if a single TCP connection is used for multiple requests, which may be sent from multiple, different, client devices). In Such cases, the operation may be skipped and the existing IPv4 TCP connection may be used. '575 teaches about using TCP connection: para. 78, 81, 88, 94, 100] "the client device and the proxy server establish a TCP connection. This TCP connection is referred herein as an IPv4 TCP connection since the client device initiates the TCP connection with the IPv4 address of the proxy server. After the TCP connection is established, at operation, the client device transmits an IPv4 packet that includes a resource request (e.g.an HTTP request) to the IPv4 address of the proxy server.

Citation # 4 (Patent/utility model) (# uploaded documents: 1):

Country code: US	Publication number: 7788378	Document kind code: B2	
Patent Applicant/Patent Owner: Microsoft Technology Licensing LLC		Title of invention: Apparatus and method for community relay node discovery	
Link to document:			
Publication Date: 26 Oct 2006 (26/10/2006)	Filing Date: 22 Apr 2005 (22/04/2005)	Priority Date:	
Source of Abstract:	Accession number:	Publication Date of Abstract:	Retrieval Date of Abstract:
Most relevant passages or drawings: claim 1, col 1, ln 45-60;		Relevant to Claims: 1, 5, 9	
<p>Brief explanation of relevance:</p> <p>US7788378B2 discloses [col 1, ln 45-60] a method of discovering a community relay node within a network community wherein the community relay node is operatively coupled to an access-protected client and adapted to facilitate communication between the access-protected client and a requesting client. It teaches about receiving a request message from a requesting client relating to a request for a community relay node, associating the request message with a serverless name resolution protocol name, selecting a community relay node from among a list of community relay nodes based on the serverless name resolution protocol name, wherein the list of community relay nodes comprises at least one internet protocol address associated with a community relay node, and returning an internet protocol address of the selected community relay node to the requesting client.</p> <p>US7788378B2 teaches in claim 1, about the request message sent and received by the client, selecting a community relay node, returning an internet protocol address to the requesting client. That is identical to Claim 1 in the present application. US7788378B2 teaches that by establishing a connection between the two clients, the community relay node may subsequently relay communications between the first and second clients. Similarly, already established communication is described in Claim 5 of the present application.</p>			

Citation # 5 (Patent/utility model) (# uploaded documents: 1):

Country code: US	Publication number: 20090216887	Document kind code: A1	
Patent Applicant/Patent Owner: Alcatel Lucent SAS		Title of invention: Method of establishing a connection	
Link to document:			
Publication Date: 27 Aug 2009 (27/08/2009)	Filing Date: 09 Dec 2008 (09/12/2008)	Priority Date: 13 Dec 2007 (13/12/2007)	
Source of Abstract:	Accession number:	Publication Date of Abstract:	Retrieval Date of Abstract:
Most relevant passages or drawings: Claims 5 and 7		Relevant to Claims:	
<p>Brief explanation of relevance:</p> <p>US20090216887A1 teaches in Claim 5, that when choosing a relay candidate establishing a connection between a first peer in a P2P network, one or more direct connections that are already established in the peer-to-peer network can be re-used for the relayed connection between the first peer and the second peer. A use of already established communication is described in Claim 11 of the current application.</p> <p>US20090216887A1 in Claim 7 teaches about method of establishing a connection, comprising the step of: sending, by the first peer, a relay discovered message comprising a peer identifier of the relay peer to the second peer;</p> <p>sending, by the second peer, a relayed connection request comprising connection information of the second peer to the relay peer;</p> <p>sending, by the relay peer, a relayed connection response comprising connection information of the relay peer to the second peer; and</p> <p>establishing a direct connection between the second peer and the relay peer if no direct connection between the second peer and the relay peer is already established.</p>			

ADVANCE E-MAIL

From the INTERNATIONAL BUREAU

PCT

COMMUNICATION IN CASES FOR WHICH
NO OTHER FORM IS APPLICABLE

To:

BINDER (SHEM TOV), Dorit
11 Shu'alei Shimshon St.
P.O.B. 7230
5217102 Ramat-Gan
ISRAËL

Date of mailing (<i>day/month/year</i>) 24 June 2019 (24.06.2019)	
Applicant's or agent's file reference HOLA-007-PCT	REPLY DUE see paragraph 1 below
International application No. PCT/IL2018/050910	International filing date (<i>day/month/year</i>) 16 August 2018 (16.08.2018)
Applicant LUMINATI NETWORKS LTD.	

- REPLY DUE within months/days from the above date of mailing
 NO REPLY DUE, however, see below
 IMPORTANT COMMUNICATION
 INFORMATION ONLY

2. COMMUNICATION:

Please find attached a copy of the non-patent literature documents uploaded with the third party observation dated 21 June 2019 (21.06.2019).

The International Bureau of WIPO 34, chemin des Colombettes 1211 Geneva 20, Switzerland	Authorized officer Biarge-Thurre Marina e-mail pct.team9@wipo.int Telephone No. +41 22 338 74 09
---	--

Internet Engineering Task Force
INTERNET DRAFT

Authors
Alain Durand (IMAG)
Paolo Fasano (CSELT)
Ivano Guardini (CSELT)
Domenico Lento (CSELT)

2 April 1999
Expires 1 October 1999

IPv6 Tunnel Broker
<[draft-ietf-ngtrans-broker-00.txt](#)>

Status of Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as ``work in progress.''

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Abstract

The IPv6 global Internet as of today is mostly build using tunnels over the existing IPv4 infrastructure. Those tunnels are difficult to configure and maintain in a large scale environment. The 6bone has proven that large sites and ISPs can do it, but this process is too complex for the isolated end user who already has an IPv4 connection and would like to enter the IPv6 world. The motivation for the development of the tunnel broker model is to help the early IPv6 adopters to hook up to the 6bone and to provide them stable, permanent IPv6 addresses and DNS names. The concept of the tunnel broker was first presented at Orlando's IETF in December 1998. Two implementations were demonstrated in Grenoble IPng & NGtrans interim meeting.

Durand Fasano Guardini Lento

Expires 1 October 1999

[Page 1]

Internet Draft

[draft-ietf-ngtrans-broker-00.txt](#)

1. Introduction

The growth of IPv6 networks started mainly using the transport facilities offered by the current Internet. This fact brought to the development of several techniques to manage IPv6 over IPv4 tunnels. At present most of the 6bone networks is built using manual tunneling over the Internet. The main drawback of this approach is the overwhelming management load for network administrators, who have to perform heavy configuration operations for each tunnel. Several attempts to reduce this management overhead have been proposed [1-3]. Nevertheless all of them present drawbacks that prevent from wide usage:

- [1] was introduced to use automatic tunnels with IPv4 compatible addresses. This approach does not solve the address exhaustion problem of IPv4. Also there is a great fear to include the complete IPv4 routing table in the IPv6 one and just making the routing table size problem worse by multiplying it by 5.
- [2] is the 6over4 mechanism. This is a site local mechanism to use IPv4 multicast as a layer 2 media. It does not solve the problem to connect an isolated user to the global IPv6 Internet.
- [3] is the 6to4 mechanism to embed IPv4 tunnel addresses into IPv6 prefixes to automatically discover tunnel endpoints. Some important technical issues such as source address selection and global routing are currently debated in the IETF. But the main difference are in the premises of the two approaches: 6to4 consider that isolated sites are to be dynamically connected in the absence of native IPv6 infrastructure and tunnel brokers consider the pre-existence of a large IPv6 global network.

This document presents an alternative approach based on the provision of Tunnel Brokers (TBs) to automatically manage tunnel requests coming from the users. This approach is expected to be useful to stimulate the growth of IPv6 interconnected hosts and to allow to early IPv6 network providers the provision of easy access to their IPv6 networks.

Section 2 provides an overall description of the Tunnel Broker Model; section 3 reports known limitations to the model; section 4 addresses security issues. A first implementation of the Tunnel Broker service is described in Appendix to this document.

2. Tunnel Broker Model

Tunnel brokers can be seen as virtual IPv6 ISP, providing IPv6 connectivity to users already connected to the IPv4 Internet. In the global IPv6 Internet it is expected that many tunnel brokers will be available and the user will just have to pick one. The list of the tunnel brokers should be referenced on a "well known" web page on <http://www.ipv6.org> to allow users to choose the "closest" one, the "cheapest" one, or any other one.

The tunnel broker model is based on a set of functional elements depicted in figure 1.




Figure 1: the Tunnel Broker model

2.1 Tunnel Broker

The TB is a place where users connect to register and activate tunnels. The TB manages tunnels creation, modification and deletion on behalf of the users. It shares the load of tunnel end-points on the network side among potentially several tunnel servers. It sends configuration orders to the relevant tunnel server when tunnels are to be created or modified. The TB also register the user in the DNS.

2.2 Tunnel server

A tunnel server is a dual stack (IPv4 & IPv6) router connected to the global Internet. Upon configuration order from the tunnel broker, it creates, modifies or deletes the half part of the tunnel toward the user. It can also maintain some statistics on the usage of the tunnels.

2.3 Using the Tunnel Broker

The client of the service is a dual-stack IPv6 node (host or router) connected to Internet. Approaching the TB, the client must provide the following information:

- the IPv4 address of the client side of the tunnel
- a nickname to be used for the registration in the DNS of the global IPv6 addresses assigned to both sides of the tunnel
- the client function (i.e. standalone host or router)

Besides, if the client machine is an IPv6 router willing to provide geographical connectivity to several IPv6 hosts, the client should be required to provide also some information about the amount of IPv6 addresses required. This allows the TB to allocate to the client an IPv6 subnet well fit to his needs instead of a single IPv6 address. Otherwise an IPv6 prefix of pre-defined length should be assigned to any client acting as an IPv6 router. The TB manages the client requests as follows:

- it first designates (e.g. according to some load sharing criteria defined by the network administrator) a Tunnel Server to be used as the actual tunnel end-point at the network side;
- it chooses the IPv6 prefix (/64 or /48) to be used;
- it fixes a lifetime for the tunnel;
- it configures the network side of the tunnel;
- it registers tunnel end-points addresses in the DNS;
- it prepares activation and de-activation scripts to be run on the client machine for easy configuration of the client side.

Then the TB sends back configuration information to the user, including tunnel parameters and DNS names. The lifetime of the IPv6 addresses are supposed to be relatively long and potentially longer than the lifetime of the IPv4 connection of the user. This will allow the user to get semipermanent IPv6 addresses and associated DNS names even though he is connected to the Internet via a dial-up link and get dynamically his IPv4 addresses by DHCP.

There are many technical alternatives to realize the interactions among the various entities in the tunnel broker model. The communication protocol used between the TB and the user could be based on SNMP, on an extension of DHCPv6, on an ad-hoc protocol or even on just some web forms filled up by the user. In a similar way, the communication protocol used between the TB and the tunnel servers is also implementation dependant. It could be some simple RSH commands, SNMP or an ad-hoc protocol specially designed or something else. Finally the Dynamic DNS Update protocol [4] should be used for automatic DNS update (i.e. to add or delete AAAA, A6 and PTR records from the DNS zone reserved for tunnel broker users) controlled by the TB. A simple alternative would be for the TB to use a small set of RSH commands to dynamically update the direct and inverse databases on the authoritative DNS server for the tunnel broker users zone (e.g. broker.isp-name.com).

2.4 Open issues

Real usage of the TB service may require to introduce accounting/billing functions.

3. Known limitations

This mechanism may not work if the user is using private IPv4 addresses behind a NAT box.

4. Security Considerations

The TB service raises several security issues. All interactions between the functional elements of the proposed architecture need to be secured, i.e.:

- the interaction between the client and TB;
- the interaction between the TB and the Tunnel Server;
- the interaction between the TB and the DNS.

Furthermore, if the client chooses to run the configuration scripts provided by TB, these scripts must be executed as root. The security techniques adopted for each of the required interaction is dependent on the implementation choices. For the client - TB interaction, the usage of http allows the exploitation of standard secure http features (SSL, S-HTTP). If e-mail exchanges are used standard mechanisms to secure e-mail can be used (PGP, PEM). For the interactions that use SNMP, the security issues are basically the same as those of securing SNMP. Otherwise if RSH commands are used standard IPsec mechanisms may apply. If the TB - DNS server interaction is a dynamic DNS update procedure, the security issues are the same discussed in [5] Finally TBs may face denial of service attack. They must implement some sort of protection against this.

5. Acknowledgments

Some of the ideas refining the tunnel broker model came from discussion with Perry Metzger and Marc Blanchet.

6. References

- [1] Gilligan, R., Nordmark, E., "Transition Mechanisms for IPv6 Hosts and Routers", RFC 1933, April 1996.
- [2] Carpenter, B., Jung, C., "Transmission of IPv6 over IPv4 Domains without Explicit Tunnels", draft-ietf-ipngwg-6over4-02.txt, January 1998.
- [3] Carpenter, B., Moore, K., "Connection of IPv6 Domains via IPv4 Clouds without Explicit Tunnels", draft-ietf-ngtrans-6to4-00.txt, January 1999.
- [4] Vixie, P., Editor, Thomson, T., Rekhter, Y., and J. Bound, "Dynamic Updates in the Domain Name System (DNS UPDATE)", RFC 2136, April 1997.
- [5] Eastlake, D., "Secure Domain Name System Dynamic Update", RFC 2137, April 1997.

7. Author's addresses

Alain Durand
IMAG
Direction des Moyens Informatiques
BP 53
38041 GRENOBLE CEDEX 9
France
Tel: +33 4 76635703
Mail: Alain.Durand@imag.fr

Durand Fasano Guardini Lento

Expires 1 October 1999

[Page 5]

Internet Draft

draft-ietf-ngtrans-broker-00.txt

Paolo Fasano S.p.A.
CSELT
Switching and Network Services - Networking
via G. Reiss Romoli, 274
10148 TORINO
Italy
Tel: +39 011 2285071
Mail: paolo.fasano@cse.lt.it

Ivano Guardini
CSELT S.p.A.
Switching and Network Services - Networking
via G. Reiss Romoli, 274
10148 TORINO
Italy
Tel: +39 011 2285424
Mail: ivano.guardini@cse.lt.it

Domenico Lento
CSELT S.p.A.
Switching and Network Services - Networking
via G. Reiss Romoli, 274
10148 TORINO
Italy
Tel: +39 011 2286993
Mail: domenico.lento@cse.lt.it

Durand Fasano Guardini Lento

Expires 1 October 1999

[Page 6]

Internet Draft

draft-ietf-ngtrans-broker-00.txt

Appendix Implementation Example

This appendix describes an early implementation of the TB service developed at CSELT, based on widely available communication tools. The basic communications between the clients and the TB run over http. The client uses a browser and can access a WWW Server providing the TB service interface. This interface offers two different hyperlinks, one for the new users and another for the registered users.

The new user has to provide some identification data (Name, Company and e-mail address) and a nickname to be used as:

- the username to login as registered user
- the name identifying the user in the DNS database

This information is submitted to the TB with a POST method. The TB starts the user configuration procedure and sends back an e-mail to the user providing a password for accessing the registered user pages and the name registered in the DNS database.

The registered user has the possibility to create a new tunnel, to view tunnel information, to change tunnel parameters and to remove an established tunnel (only one active tunnel per user is allowed). To create a new tunnel, the user has to provide some additional information:

- the IPv4 address of the user-side tunnel end-point (the TB pre-fill this field using http carried browser information);
- the O.S./IPv6 implementation used;
- if the user end-point of the tunnel will be on a host or router.

If the user requests to use a router as tunnel end-point a new form is pushed to the user asking:

- motivation;
- life-time.

Then the user submit this information to the TB and the tunnel configuration procedure takes place.

A registered user who has already set-up a tunnel can view a display of the following tunnel parameters:

- Server IPv4 Address
- Server IPv6 Address
- Server IPv6 Link Local Addr
- Client IPv4 Address
- Client IPv6 Address
- Client IPv6 Link Local Addr
- Expiration Date

The user can also modify the Client IPv4 Address if this is changed, can

extend the tunnel life-time a day before the Expiration date and can delete the tunnel anytime. The communication between the client and the TB may be secured using SSL (access to the TB using the https scheme).

A.1 User configuration procedure

When the TB receives a request of registration by a new user, it operates as follows:

- uses the nickname to build a name identifying that user in the DNS system;
- updates an internal user database;
- sends an e-mail back to the user.

A.2 Tunnel configuration procedure

Once a registered user asked for the creation of a tunnel providing all the required information the TB first checks if the user requested to terminate the tunnel on a router or on a host. If the user choice was a router the request is put in a pending state and managed administratively: the administrator of the TB has the possibility to accept or refuse the motivations and lifetime indicated by the user. If the user choice was a host the TB acts automatically as follows:

- i) verifies if resources are available to set-up a new tunnel (otherwise puts the user request in a pending state and go to step viii);
- ii) selects a Tunnel Server from the list of available Tunnel Servers on the basis of simple number-of-tunnels balancing criteria;
- iii) selects an IPv6 prefix to be used for assigning IPv6 addresses to the tunnel end-points;
- iv) sets an Expiration Date for the tunnel (default 7 days);
- v) configures the Tunnel Server;
- vi) updates the DNS server;
- vii) prepares activation and de-activation scripts for tunnel configuration on the user side;
- viii) pushes to the user browser a new page displaying the results of the tunnel request: if OK the new page displays tunnel parameters and hyperlinks to the activation and de-activation scripts.

The user who receives positive acknowledgment can then execute (downloading the scripts or not) the activation script to configure the user side of the tunnel. There is still the possibility for a user that do not want to run the configuration scripts or that has an IPv6 implementation not supported by the TB to set up his/her end-point of the tunnel manually. At the end of this procedure the user is IPv6 connected and identified by his/her own name in the DNS.

A similar procedure is performed when the user selected a router as tunnel end-point and the Administrator accepted the request.

A.3 User, Tunnel and Tunnel Server management

The TB maintains three databases, one for users, one for active tunnels and the last one for Tunnel Servers. The User database has one entry for each user of the service; each entry has the following fields:

- Username
- Password
- DNS entry
- Firstname
- Lastname
- Company
- Country
- E-Mail
- Has Tunnel (yes/not for active tunnel)
- Tunnel Count (number of tunnel creation performed by the user)

The Tunnel database has one entry for each active tunnel; each entry has the following fields:

- Identifier
- Owner
- User IPv4 Address
- Server IPv4 Address
- User Global IPv6 Address
- Server Global IPv6 Address
- User Link Local Address
- Server Link Local Address
- User OS Type
- Creation Date
- Expiration Date
- Standalone
- Manual

The Tunnel Server database has one entry for each Tunnel Server; each entry has the following fields:

- Identifier
- IPv4
- IPv6
- OS
- Use TBSP
- Used Standalone Tunnels
- Max Standalone Tunnels
- Used Router Tunnels
- Max Router Tunnels

The TB manages the service updating these databases. An Administrator Interface gives to the TB manager a full control (add, modify and remove any time) over users, tunnels and Tunnel Servers.

In order to access to the administrative web pages, the TB administrator has to log as Registered User using the administrative username and

Durand Fasano Guardini Lento

Expires 1 October 1999

[Page 9]

Internet Draft

draft-ietf-ngtrans-broker-00.txt

password. The page presented to the administrator contains hyperlinks to the following sections:

- Administrator Profile Change
- User Administration
- Tunnel Server Administration
- Tunnel Administration

The Administrator Profile Change lets the administrator to change his password.

The User Administration section, once selected, allows the administrator to interact with the User database in order to list the database content or delete an entry in the database. If the administration deletes an entry with an associated tunnel, the tunnel is released.

The Tunnel Server Administration section allows the administrator to manage the data contained in the Tunnel Server database. The page presented to the TB superuser contains hyperlinks to the following subsections:

- Tunnel Server List (the content of the Tunnel Server database is displayed with the relevant informations);
- Add Tunnel Server (this hyperlink allows the insertion of a new Tunnel Server; the administrator is asked for the Tunnel Server informations as described in the previous section);
- Modify Tunnel Server (this subsection is used by the administrator to change the information of a Tunnel Server, e.g. Max Standalone Tunnels);
- Delete Tunnel Server (causes the removal of the selected Tunnel Server entry from the Tunnel Server database; the tunnels managed by this Tunnel Server are released).

The Tunnel Administration section is used to perform tunnel management. The page presented to the administrator contains hyperlinks to the following subsections:

- Tunnel List (the content of the Tunnel database is displayed to the administrator)
- Manual Setup (allows the TB superuser to setup manually a tunnel)
- Release (causes the release of the selected tunnel)
- Change Parameters (allows the update of the data associated to a tunnel)
- Pending Router Request (displays the list of the user requests for a tunnel towards a router; two hyperlinks are associated to each entry allowing the administrator to accept or refuse the request).

A.4 Modularity

The Tunnel Broker implements a plugin-like mechanism for adding support for new Tunnel Servers or client operating systems without modifying the TB scripts or breaking the service. To achieve this result the scripts

has to follow a predefined template and are kept in a plugin directory checked at every request for a new tunnel. This implies that the list of supported Tunnel Servers and client OSs is built dynamically, based on the content of the plugin directory.

A.4.1 Script directory structure

The scripts for interacting with users, Tunnel Servers and DNS are stored in a plugin directory structured as following:

```
<TB plugin home directory>
|
+--- script
|
|   +--- dns
|   |
|   +--- server
|   |   |
|   |   +--- local
|   |   |   |
|   |   |   +--- act
|   |   |   |
|   |   |   +--- deact
|   |   |
|   |   +--- remote
|   |   |   |
|   |   |   +--- act
|   |   |   |
|   |   |   +--- deact
|   |
|   +--- client
|   |   |
|   |   +--- act
|   |   |
|   |   +--- deact
```

The scripts have to be inserted in the proper subdirectory accordingly with their functionality (eg. a Tunnel Server activation script for a remote Tunnel Server in inserted in directory <TB home>/script/server/remote/act).

This tree is scanned by the CGI program every time a user requests scripts for tunnel activation/deactivation and at the insertion of a new Tunnel Server for building the appropriate list of supported OSes.

A.4.2 Client scripts

Client scripts are used to help a TB user to configure his/her own host. In order to support a new client architecture, a TB administrator has to provide both activation and deactivation scripts for the selected configuration. These scripts must include the following keywords, that

Scripts follow a naming convention:

- activation and deactivation scripts must have the same name
- scripts filename has the structure <OS-StackType>.<extension> (eg. PERL scripts for FreeBSD hosts using INRIA IPv6 implementation could have as filename FreeBSD-INRIA.pl); the <OS-StackType> name is used as the name displayed in the user interface selection list.

will be replaced with proper values for the specific user request:

- `_ipv4client_` for the client IPv4 address;
- `_ipv4server_` for server IPv4 address;
- `_ipv6client_` for the client global IPv6 address;
- `_ipv6server_` for the server global IPv6 address;
- `_ipv6llclient_` for the client link local address;
- `_ipv6llserver_` for the server link local address.

Every time a TB user interacts with the TB web pages in order to download the activation/deactivation scripts, the CGI provides keywords substitution with the correct values stored in the TB database.

A.4.3 Server scripts

Server scripts are used to both setup and release an IPv6 over IPv4 tunnel on a tunnel server. In order to support a new Tunnel Server, a TB administrator has to provide both activation and deactivation script for the new platform. These scripts are invoked by the CGI program at every tunnel setup or release. The following parameters are passed to the script :

```
<tunnel type> (could assume the values 'standalone' or 'router')  
<client IPv4 address>  
<server IPv4 address>  
<client global IPv6 address>  
<server global IPv6 address>  
<client local link address>  
<server local link address>.
```

The executed script has to return the value 0 on success and -1 on failure.

A.4.4 DNS scripts

DNS scripts are used to interact with the DNS in order to update its resolution tables. All parameters specific to the DNS (IP address, software, file structure, etc.) and the interaction mode between the TB and the DNS are embedded within the DNS scripts and do not affect other TB scripts. The TB uses a script called 'dns_act' to add a new entry in the DNS database and a script named 'dns_deact' to remove a host entry

from the DNS tables. Both scripts are invoked by passing two parameters:

```
<host name>  
<global IPv6 address>.
```

The executed script has to return the value 0 on success and -1 on failure.

A.5 CSELT's Tunnel Broker location

The TB service is up and running at:
<https://carmen.csel.it/ipv6tb>

The software implementing the TB is freely available at:
<http://carmen.csel.it/ipv6/download>

Network Working Group
Request for Comments: 3053
Category: Informational

A. Durand
SUN Microsystems, Inc
P. Fasano
I. Guardini
CSELT S.p.A.
D. Lento
TIM
January 2001

IPv6 Tunnel Broker

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2001). All Rights Reserved.

Abstract

The IPv6 global Internet as of today uses a lot of tunnels over the existing IPv4 infrastructure. Those tunnels are difficult to configure and maintain in a large scale environment. The 6bone has proven that large sites and Internet Service Providers (ISPs) can do it, but this process is too complex for the isolated end user who already has an IPv4 connection and would like to enter the IPv6 world. The motivation for the development of the tunnel broker model is to help early IPv6 adopters to hook up to an existing IPv6 network (e.g., the 6bone) and to get stable, permanent IPv6 addresses and DNS names. The concept of the tunnel broker was first presented at Orlando's IETF in December 1998. Two implementations were demonstrated during the Grenoble IPng & NGtrans interim meeting in February 1999.

1. Introduction

The growth of IPv6 networks started mainly using the transport facilities offered by the current Internet. This led to the development of several techniques to manage IPv6 over IPv4 tunnels. At present most of the 6bone network is built using manually configured tunnels over the Internet. The main drawback of this approach is the overwhelming management load for network administrators, who have to perform extensive manual configuration for each tunnel. Several attempts to reduce this management overhead

Durand, et al.

Informational

[Page 1]

REC_3053

IPv6 Tunnel Broker

January 2001

have already been proposed and each of them presents interesting advantages but also solves different problems than the Tunnel Broker, or poses drawbacks not present in the Tunnel Broker:

- the use of automatic tunnels with IPv4 compatible addresses [1] is a simple mechanism to establish early IPv6 connectivity among isolated dual-stack hosts and/or routers. The problem with this approach is that it does not solve the address exhaustion problem of IPv4. Also there is a great fear to include the complete IPv4 routing table into the IPv6 world because this would worsen the routing table size problem multiplying it by 5;
- 6over4 [2] is a site local transition mechanism based on the use of IPv4 multicast as a virtual link layer. It does not solve the problem of connecting an isolated user to the global IPv6 Internet;
- 6to4 [3] has been designed to allow isolated IPv6 domains, attached to a wide area network with no native IPv6 support (e.g., the IPv4 Internet), to communicate with other such IPv6 domains with minimal manual configuration. The idea is to embed IPv4 tunnel addresses into the IPv6 prefixes so that any domain border router can automatically discover tunnel endpoints for outbound IPv6 traffic.

The Tunnel Broker idea is an alternative approach based on the provision of dedicated servers, called Tunnel Brokers, to automatically manage tunnel requests coming from the users. This approach is expected to be useful to stimulate the growth of IPv6 interconnected hosts and to allow early IPv6 network providers to provide easy access to their IPv6 networks.

The main difference between the Tunnel Broker and the 6to4 mechanisms is that they serve a different segment of the IPv6 community:

- the Tunnel Broker fits well for small isolated IPv6 sites, and especially isolated IPv6 hosts on the IPv4 Internet, that want to easily connect to an existing IPv6 network;
- the 6to4 approach has been designed to allow isolated IPv6 sites to easily connect together without having to wait for their IPv4 ISPs to deliver native IPv6 services. This is very well suited for extranet and virtual private networks. Using 6to4 relays, 6to4 sites can also reach sites on the IPv6 Internet.

In addition, the Tunnel Broker approach allows IPv6 ISPs to easily perform access control on the users enforcing their own policies on network resources utilization.

This document is intended to present a framework describing the guidelines for the provision of a Tunnel Broker service within the Internet. It does not specify any protocol but details the general architecture of the proposed approach. It also outlines a set of viable alternatives for implementing it. Section 2 provides an overall description of the Tunnel Broker model; Section 3 reports known limitations to the model; Section 4 briefly outlines other possible applications of the Tunnel Broker approach; Section 5 addresses security issues.

2. Tunnel Broker Model

Tunnel brokers can be seen as virtual IPv6 ISPs, providing IPv6 connectivity to users already connected to the IPv4 Internet. In the emerging IPv6 Internet it is expected that many tunnel brokers will be available so that the user will just have to pick one. The list of the tunnel brokers should be referenced on a "well known" web page (e.g. on <http://www.ipv6.org>) to allow users to choose the "closest" one, the "cheapest" one, or any other one.

The tunnel broker model is based on the set of functional elements depicted in figure 1.




Figure 1: the Tunnel Broker model

2.1 Tunnel Broker (TB)

The TB is the place where the user connects to register and activate tunnels. The TB manages tunnel creation, modification and deletion on behalf of the user.

For scalability reasons the tunnel broker can share the load of network side tunnel end-points among several tunnel servers. It sends configuration orders to the relevant tunnel server whenever a tunnel has to be created, modified or deleted. The TB may also register the user IPv6 address and name in the DNS.

A TB must be IPv4 addressable. It may also be IPv6 addressable, but this is not mandatory. Communications between the broker and the servers can take place either with IPv4 or IPv6.

2.2 Tunnel server (TS)

A TS is a dual-stack (IPv4 & IPv6) router connected to the global Internet. Upon receipt of a configuration order coming from the TB, it creates, modifies or deletes the server side of each tunnel. It may also maintain usage statistics for every active tunnel.

2.3 Using the Tunnel Broker

The client of the Tunnel Broker service is a dual-stack IPv6 node (host or router) connected to the IPv4 Internet. Approaching the TB, the client should be asked first of all to provide its identity and credentials so that proper user authentication, authorization and (optionally) accounting can be carried out (e.g., relying on existing AAA facilities such as RADIUS). This means that the client and the TB have to share a pre-configured or automatically established security association to be used to prevent unauthorized use of the service. With this respect the TB can be seen as an access-control server for IPv4 interconnected IPv6 users.

Once the client has been authorized to access the service, it should provide at least the following information:

- the IPv4 address of the client side of the tunnel;
- a name to be used for the registration in the DNS of the global IPv6 address assigned to the client side of the tunnel;
- the client function (i.e., standalone host or router).

Moreover, if the client machine is an IPv6 router willing to provide connectivity to several IPv6 hosts, the client should be asked also to provide some information about the amount of IPv6 addresses required. This allows the TB to allocate the client an IPv6 prefix that fits its needs instead of a single IPv6 address.

The TB manages the client requests as follows:

- it first designates (e.g., according to some load sharing criteria defined by the TB administrator) a Tunnel Server to be used as the actual tunnel end-point at the network side;
- it chooses the IPv6 prefix to be allocated to the client; the prefix length can be anything between 0 and 128, most common values being 48 (site prefix), 64 (subnet prefix) or 128 (host prefix);
- it fixes a lifetime for the tunnel;
- it automatically registers in the DNS the global IPv6 addresses assigned to the tunnel end-points;
- it configures the server side of the tunnel;

- it notifies the relevant configuration information to the client, including tunnel parameters and DNS names.

After the above configuration steps have been carried out (including the configuration of the client), the IPv6 over IPv4 tunnel between the client host/router and the selected TS is up and working, thus allowing the tunnel broker user to get access to the 6bone or any other IPv6 network the TS is connected to.

2.4 IPv6 address assignment

The IPv6 addresses assigned to both sides of each tunnel must be global IPv6 addresses belonging to the IPv6 addressing space managed by the TB.

The lifetime of these IPv6 addresses should be relatively long and potentially longer than the lifetime of the IPv4 connection of the user. This is to allow the client to get semipermanent IPv6 addresses and associated DNS names even though it is connected to the Internet via a dial-up link and gets dynamically assigned IPv4 addresses through DHCP.

2.5 Tunnel management

Active tunnels consume precious resources on the tunnel servers in terms of memory and processing time. For this reason it is advisable to keep the number of unused tunnels as small as possible deploying a well designed tunnel management mechanism.

Each IPv6 over IPv4 tunnel created by the TB should at least be assigned a lifetime and removed after its expiration unless an explicit lifetime extension request is submitted by the client.

Obviously this is not an optimal solution especially for users accessing the Internet through short-lived and dynamically addressed IPv4 connections (e.g., dial-up links). In this case a newly established tunnel is likely to be used just for a short time and then never again, in that every time the user reconnects he gets a new IPv4 address and is therefore obliged either to set-up a new tunnel or to update the configuration of the previous one. In such a situation a more effective tunnel management may be achieved by having the TS periodically deliver to the TB IPv6 traffic and reachability statistics for every active tunnel. In this way, the TB can enforce a tunnel deletion after a period of inactivity without waiting for the expiration of the related lifetime which can be relatively longer (e.g., several days).

Another solution may be to implement some kind of tunnel management protocol or keep-alive mechanism between the client and the TS (or between the client and the TB) so that each tunnel can be immediately released after the user disconnects (e.g., removing his tunnel end-point or tearing down his IPv4 connection to the Internet). The drawback of this policy mechanism is that it also requires a software upgrade on the client machine in order to add support for the ad-hoc keep-alive mechanism described above.

Moreover, keeping track of the tunnel configuration even after the user has disconnected from the IPv4 Internet may be worth the extra cost. In this way, in fact, when the user reconnects to the Internet, possibly using a different IPv4 address, he could just restart the tunnel by getting in touch with the TB again. The TB could then order a TS to re-create the tunnel using the new IPv4 address of the client but reusing the previously allocated IPv6 addresses. That way, the client could preserve a nearly permanent (static) IPv6 address even though its IPv4 address is dynamic. It could also preserve the associated DNS name.

2.6 Interactions between client, TB, TS and DNS

As previously stated, the definition of a specific set of protocols and procedures to be used for the communication among the various entities in the Tunnel Broker architecture is outside of the scope of the present framework document. Nevertheless, in the remainder of this section some viable technical alternatives to support client-TB, TB-TS and TB-DNS interactions are briefly described in order to help future implementation efforts or standardization initiatives.

The interaction between the TB and the user could be based on http. For example the user could provide the relevant configuration information (i.e., the IPv4 address of the client side of the tunnel, etc.) by just filling up some forms on a Web server running on the TB. As a result the server could respond with an html page stating that the server end-point of the tunnel is configured and displaying all the relevant tunnel information.

After that, the most trivial approach would be to leave the user to configure the client end-point of the tunnel on his own. However, it should be highly valuable to support a mechanism to automate this procedure as much as possible.

Several options may be envisaged to assist the Tunnel Broker user in the configuration of his dual-stack equipment. The simplest option is that the TB could just prepare personalized activation and de-activation scripts to be run off-line on the client machine to achieve easy set-up of the client side tunnel end-point. This

solution is clearly the easiest to implement and operate in that it does not require any software extension on the client machine. However, it raises several security concerns because it may be difficult for the user to verify that previously downloaded scripts do not perform illegal or dangerous operations once executed.

The above described security issues could be elegantly overcome by defining a new MIME (Multipurpose Internet Mail Extension) content-type (e.g., application/tunnel) [4,5] to be used by the TB to deliver the tunnel parameters to the client. In this case, there must be a dedicated agent running on the client to process this information and actually set-up the tunnel end-point on behalf of the user. This is a very attractive approach which is worth envisaging. In particular, the definition of the new content-type might be the subject of a future ad-hoc document.

Several options are available also to achieve proper interaction between the broker and the Tunnel Servers. For example a set of simple RSH commands over IPsec could be used for this purpose. Another alternative could be to use SNMP or to adopt any other network management solution.

Finally, the Dynamic DNS Update protocol [6] should be used for automatic DNS update (i.e., to add or delete AAAA, A6 and PTR records from the DNS zone reserved for Tunnel Broker users) controlled by the TB. A simple alternative would be for the TB to use a small set of RSH commands to dynamically update the direct and inverse databases on the authoritative DNS server for the Tunnel Broker users zone (e.g. broker.isp-name.com).

2.7 Open issues

Real usage of the TB service may require the introduction of accounting/billing functions.

3. Known limitations

This mechanism may not work if the user is using private IPv4 addresses behind a NAT box.

4. Use of the tunnel broker concept in other areas

The Tunnel Broker approach might be efficiently exploited also to automatically set-up and manage any other kind of tunnel, such as a multicast tunnel (e.g., used to interconnect multicast islands within the unicast Internet) or an IPsec tunnel.

Moreover, the idea of deploying a dedicated access-control server, like the TB, to securely authorize and assist users that want to gain access to an IPv6 network might prove useful also to enhance other transition mechanisms. For example it would be possible to exploit a similar approach within the 6to4 model to achieve easy relay discovery. This would make life easier for early 6to4 adopters but would also allow the ISPs to better control the usage of their 6to4 relay facilities (e.g., setting up appropriate usage policies).

5. Security Considerations

All the interactions between the functional elements of the proposed architecture need to be secured:

- the interaction between the client and TB;
- the interaction between the TB and the Tunnel Server;
- the interaction between the TB and the DNS.

The security techniques adopted for each of the required interactions is dependent on the implementation choices.

For the client-TB interaction, the usage of http allows the exploitation of widely adopted security features, such as SSL (Secure Socket Layer) [7], to encrypt data sent to and downloaded from the web server. This also makes it possible to rely on a simple "username" and "password" authentication procedure and on existing AAA facilities (e.g., RADIUS) to enforce access-control.

For the TB-TS interaction secure SNMP could be adopted [8,9,10]. If the dynamic DNS update procedure is used for the TB-DNS interaction, the security issues are the same as discussed in [11]. Otherwise, if a simpler approach based on RSH commands is used, standard IPsec mechanisms can be applied [12].

Furthermore, if the configuration of the client is achieved running scripts provided by the TB, these scripts must be executed with enough privileges to manage network interfaces, such as an administrator/root role. This can be dangerous and should be considered only for early implementations of the Tunnel Broker approach. Transferring tunnel configuration parameters in a MIME type over https is a more secure approach.

In addition a loss of confidentiality may occur whenever a dial-up user disconnects from the Internet without tearing down the tunnel previously established through the TB. In fact, the TS keeps tunneling the IPv6 traffic addressed to that user to his old IPv4

address regardless of the fact that in the meanwhile that IPv4 address could have been dynamically assigned to another subscriber of the same dial-up ISP. This problem could be solved by implementing on every tunnel the keep-alive mechanism outlined in section 2.5 thus allowing the TB to immediately stop IPv6 traffic forwarding towards disconnected users.

Finally TBs must implement protections against denial of service attacks which may occur whenever a malicious user exhausts all the resources available on the tunnels server by asking for a lot of tunnels to be established altogether. A possible protection against this attack could be achieved by administratively limiting the number of tunnels that a single user is allowed to set-up at the same time.

6. Acknowledgments

Some of the ideas refining the tunnel broker model came from discussion with Perry Metzger and Marc Blanchet.

7. References

- [1] Gilligan, R. and E. Nordmark, "Transition Mechanisms for IPv6 Hosts and Routers", RFC 1933, April 1996.
- [2] Carpenter, B. and C. Jung, "Transmission of IPv6 over IPv4 Domains without Explicit Tunnels", RFC 2529, March 1999.
- [3] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds without Explicit Tunnels", Work in Progress.
- [4] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies", RFC 2045, November 1996.
- [5] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part Two: Media Types", RFC 2046, November 1996.
- [6] Vixie, P., Editor, Thomson, T., Rekhter, Y. and J. Bound, "Dynamic Updates in the Domain Name System (DNS UPDATE)", RFC 2136, April 1997.
- [7] Guttman, E., Leong, L. and G. Malkin, "Users' Security Handbook", FYI 34, RFC 2504, February 1999.
- [8] Wijnen, B., Harrington, D. and R. Presuhn, "An Architecture for Describing SNMP Management Frameworks", RFC 2571, April 1999.

- [9] Blumenthal, U. and B. Wijnen, "User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)", RFC_2574, April 1999.
- [10] Wijnen, B., Presuhn, R. and K. McCloghrie, "View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)", RFC_2575, April 1999.
- [11] Eastlake, D., "Secure Domain Name System Dynamic Update", RFC_2137, April 1997.
- [12] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC_2401, November 1998.

Durand, et al.

Informational

[Page 11]

RFC_3053

IPv6 Tunnel Broker

January 2001

3. Authors' Addresses

Alain Durand
SUN Microsystems, Inc
901 San Antonio Road
MPK17-202
Palo Alto, CA 94303-4900
USA

Phone: +1 650 786 7503
EMail: Alain.Durand@sun.com

Paolo Fasano S.p.A.
CSELT S.p.A.
Switching and Network Services - Networking
via G. Reiss Romoli, 274
10148 TORINO
Italy

Phone: +39 011 2285071
EMail: paolo.fasano@cse.lt.it

Ivano Guardini
CSELT S.p.A.
Switching and Network Services - Networking
via G. Reiss Romoli, 274
10148 TORINO
Italy

Phone: +39 011 2285424
EMail: ivano.guardini@cse.lt.it

Domenico Lento
TIM
Business Unit Project Management
via Orsini, 9
90100 Palermo
Italy

Phone: +39 091 7583243
EMail: dlento@mail.tim.it

Durand, et al.

Informational

[Page 12]

RFC_3053

IPv6 Tunnel Broker

January 2001

9. Full Copyright Statement

Copyright (C) The Internet Society (2001). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

Electronic Acknowledgement Receipt

EFS ID:	36456460
Application Number:	16278107
International Application Number:	
Confirmation Number:	4936
Title of Invention:	SYSTEM PROVIDING FASTER AND MORE EFFICIENT DATA COMMUNICATION
First Named Inventor/Applicant Name:	Derry Shribman
Customer Number:	131926
Filer:	Yehuda Binder/Dorit Binder
Filer Authorized By:	Yehuda Binder
Attorney Docket Number:	HOLA-005-US10
Receipt Date:	01-JUL-2019
Filing Date:	17-FEB-2019
Time Stamp:	03:58:11
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no
------------------------	----

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Information Disclosure Statement (IDS) Form (SB08)	005-US10-IDS-007.pdf	1035162 <small>477d396711cc1256037feac747293c342b098dbc</small>	no	5

Warnings:

Information:					
2	Non Patent Literature	003-KEEP-ALIVE.pdf	219277	no	3
			2ce5a2fa91540953b35c35c3ce9da8d2661cd978		
Warnings:					
Information:					
3	Foreign Reference	007-EP2922275.pdf	553291	no	39
			c1e9bf90ec6f918527140c333cf5af68b4c1f431		
Warnings:					
Information:					
4	Other Reference-Patent/App/Search documents	007-WO2019043687-TPOBS.pdf	306816	no	7
			998c61380ac96f993acd531cca738a8cd99ddef		
Warnings:					
Information:					
5		007-IL2018050910-NP-references.pdf	1113908	yes	27
			3305ca6da44b589a7be18701a091930ce7173aac		
	Multipart Description/PDF files in .zip description				
	Document Description		Start	End	
	Non Patent Literature		1	14	
	Non Patent Literature		15	27	
Warnings:					
Information:					
Total Files Size (in bytes):			3228454		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number	16278107
	Filing Date	2019-02-17
	First Named Inventor	Derry Shribman
	Art Unit	2459
	Examiner Name	
	Attorney Docket Number	HOLA-005-US10

U.S.PATENTS						Remove
Examiner Initial*	Cite No	Patent Number	Kind Code ¹	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	9177157	B2	2015-11-03	Yehuda Binder	

If you wish to add additional U.S. Patent citation information please click the Add button.

Add

U.S.PATENT APPLICATION PUBLICATIONS						Remove
Examiner Initial*	Cite No	Publication Number	Kind Code ¹	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20070142036	A1	2007-06-21	Johan Wikman	
	2	20090248793	A1	2009-10-01	Sanny Jacobsson	
	3	20130080575	A1	2013-03-28	Matthew Browning Prince	
	4	20110066924	A1	2011-03-17	Gregory Dorso	
	5	20120246273	A1	2012-09-27	Claudson F. Bornstein	

If you wish to add additional U.S. Published Application citation information please click the Add button.

Add

FOREIGN PATENT DOCUMENTS						Remove
--------------------------	--	--	--	--	--	--------

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT**
(Not for submission under 37 CFR 1.99)

Application Number	16278107
Filing Date	2019-02-17
First Named Inventor	Derry Shribman
Art Unit	2459
Examiner Name	
Attorney Docket Number	HOLA-005-US10

Examiner Initial*	Cite No	Foreign Document Number ³	Country Code ²ⁱ	Kind Code ⁴	Publication Date	Name of Patentee or Applicant of cited Document	Pages, Columns, Lines where Relevant Passages or Relevant Figures Appear	T ⁵
	1	2004094980	WO	A2	2004-11-04	FONTIJN, Wilhelmus, F., J. et al		

If you wish to add additional Foreign Patent Document citation information please click the Add button

NON-PATENT LITERATURE DOCUMENTS

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T ⁵
	1	Third-party submission under 37 CFR 1.290 filed on July 23, 2019 and entered in U.S. Appl. No. 16/140,749	
	2	Third-party submission under 37 CFR 1.290 filed on July 23, 2019 and entered in U.S. Appl. No. 16/140,785	
	3	Third-party submission under 37 CFR 1.290 filed on July 23, 2019 and entered in U.S. Appl. No. 16/214,433	
	4	Third-party submission under 37 CFR 1.290 filed on July 23, 2019 and entered in U.S. Appl. No. 16/214,451	
	5	Third-party submission under 37 CFR 1.290 filed on July 23, 2019 and entered in U.S. Appl. No. 16/214,476	
	6	Third-party submission under 37 CFR 1.290 filed on July 23, 2019 and entered in U.S. Appl. No. 16/214,496	
	7	Third-party submission under 37 CFR 1.290 filed on July 23, 2019 and entered in U.S. Appl. No. 16/292,363	

**INFORMATION DISCLOSURE
STATEMENT BY APPLICANT
(Not for submission under 37 CFR 1.99)**

Application Number	16278107
Filing Date	2019-02-17
First Named Inventor	Derry Shribman
Art Unit	2459
Examiner Name	
Attorney Docket Number	HOLA-005-US10

8	Third-party submission under 37 CFR 1.290 filed on July 22, 2019 and entered in U.S. Appl. No. 16/292,364
9	Third-party submission under 37 CFR 1.290 filed on July 23, 2019 and entered in U.S. Appl. No. 16/292,374
10	Third-party submission under 37 CFR 1.290 filed on July 23, 2019 and entered in U.S. Appl. No. 16/292,382
11	Third-party submission under 37 CFR 1.290 filed on July 25, 2019 and entered in U.S. Appl. No. 16/365,250
12	Third-party submission under 37 CFR 1.290 filed on July 25, 2019 and entered in U.S. Appl. No. 16/365,315
13	"Slice Embedding Solutions for Distributed Service Architectures" - Esposito et al., Boston University, 02/12/2011 http://www.cs.bu.edu/techreports/pdf/2011-025-slice-embedding.pdf (Year 2011) (16 pages)

If you wish to add additional non-patent literature document citation information please click the Add button

EXAMINER SIGNATURE

Examiner Signature		Date Considered	
--------------------	--	-----------------	--

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

¹ See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. ² Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). ³ For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. ⁴ Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. ⁵ Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)	Application Number	16278107
	Filing Date	2019-02-17
	First Named Inventor	Derry Shribman
	Art Unit	2459
	Examiner Name	
	Attorney Docket Number	HOLA-005-US10

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.

The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Yehuda Binder/	Date (YYYY-MM-DD)	2019-08-06
Name/Print	Yehuda Binder	Registration Number	73612

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. **DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt

EFS ID:	36791161
Application Number:	16278107
International Application Number:	
Confirmation Number:	4936
Title of Invention:	SYSTEM PROVIDING FASTER AND MORE EFFICIENT DATA COMMUNICATION
First Named Inventor/Applicant Name:	Derry Shribman
Customer Number:	131926
Filer:	Yehuda Binder/Dorit Binder
Filer Authorized By:	Yehuda Binder
Attorney Docket Number:	HOLA-005-US10
Receipt Date:	06-AUG-2019
Filing Date:	17-FEB-2019
Time Stamp:	10:00:33
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no
------------------------	----

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Other Reference-Patent/App/Search documents	16140749-Third-party-submission.pdf	435335 <small>3db83767eb028fa7aea6b0cf59843d845bdcfa539</small>	no	5

Warnings:

Information:					
2	Other Reference-Patent/App/Search documents	16140785-Third-party-submission.pdf	397514	no	5
			52aeb6d4b21b368b9e4ef8b4816cd34ede858a4e		
Warnings:					
Information:					
3	Other Reference-Patent/App/Search documents	16214433-Third-party-submission.pdf	400551	no	5
			dc3cd5a8fc0ee68851c6564c91e4004d36db1637		
Warnings:					
Information:					
4	Other Reference-Patent/App/Search documents	16214451-Third-party-submission.pdf	413417	no	5
			b6ac7ba98b0e2f1544cce389a0a17911b010310c		
Warnings:					
Information:					
5	Other Reference-Patent/App/Search documents	16214476-Third-party-submission.pdf	394016	no	5
			bc41f6602e015118dd9e1be9dbdc8f995dafd430		
Warnings:					
Information:					
6	Other Reference-Patent/App/Search documents	16214496-Third-party-submission.pdf	399422	no	5
			d31d47fc339d1d6dac947988a604e1df5b82be81		
Warnings:					
Information:					
7	Other Reference-Patent/App/Search documents	16292363-Third-party-submission.pdf	410250	no	5
			571f452ac0b39b92f309979874d5e19926b2fd80		
Warnings:					
Information:					
8	Other Reference-Patent/App/Search documents	16292364-Third-party-submission.pdf	374919	no	4
			33866a9e474fe311f80c730587b9880701eab460		
Warnings:					
Information:					

9	Other Reference-Patent/App/Search documents	16292374-Third-party-submission.pdf	437532	no	5
			5a8765bc7033926c22556a0201301477d2a d55f3		
Warnings:					
Information:					
10	Other Reference-Patent/App/Search documents	16292382-Third-party-submission.pdf	435095	no	5
			f7598bf172edfdaa05161b1cfefdf1bbccdb 173		
Warnings:					
Information:					
11	Other Reference-Patent/App/Search documents	16365250-Third-party-submission.pdf	398304	no	5
			1d19eac24a9fc4065d7aa803a7665258029 5297e		
Warnings:					
Information:					
12	Other Reference-Patent/App/Search documents	16365315-Third-party-submission.pdf	428698	no	5
			e7925c770e9cca3229ce393a637ad8e77b1 3953c		
Warnings:					
Information:					
13	Non Patent Literature	Slice-Embedding.pdf	1687423	no	16
			5fb1575ca50559d1d278d60059293a3bba 15fa6		
Warnings:					
Information:					
14	Foreign Reference	WO2004094980.pdf	879969	no	19
			a7e39f6bf1ca149cf34cc4672a914090a3d9 223		
Warnings:					
Information:					
15	Information Disclosure Statement (IDS) Form (SB08)	Third-party-submission-005-US10.pdf	1035225	no	5
			b0a4792c6aa24a51d88796a14e1ec5bf715 470bf		
Warnings:					
Information:					
Total Files Size (in bytes):				8527670	

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290 CONCISE DESCRIPTION OF RELEVANCE		
Application Number	16140749	
U.S. PATENTS		
Cite No	Patent Number	Concise Description of Relevance
U.S. PATENT APPLICATION PUBLICATION		
Cite No	Publication Number	Concise Description of Relevance
1	20090248793	Jacobsson discloses a method of retrieving content from a network (e.g. the Internet (par19)). The content is identified by a content identifier (see e.g.,par.34-35). Also discloses that other metadata may be associated with content and content sources. Jacobsson discloses the selection of the content source from the list of content sources. The selection process can include a optimization/prioritization algorithm that considers the geographic location of the content source (par 5-6, 35-36).Also discloses multiple devices, including clients (peers) and servers that can be content sources (par 23-24,42). The various content sources are identified by public or private IP address (par 35). Jacobsson discloses that requests for content can be sent to any of the content sources (client devices and/or servers). Jacobsson discloses the first device requesting and receiving the identified content (or portions thereof) from the various content sources (web servers and/or client devices)par 53.

2	20070142036	<p>Wikman discloses a mobile intermediate device for providing content from a geographically proximate origin device to a geographically remote client device, the mobile intermediate device having a mobile intermediate device address to which content requests are communicated from the client device, and mobile intermediate device comprising an input and an output, the output arranged to provide content to a remote client device upon receipt of a content request at the input from the remote client device. [0020-0021]</p> <p>Wikman discloses a method of providing content using a mobile intermediate device, the mobile intermediate device for providing content from a geographically proximate origin device to a geographically remote client device, the mobile intermediate device having a mobile intermediate device address to which content requests are communicated from the client device. [0024]</p>
---	-------------	---

FOREIGN PATENT DOCUMENTS

CiteNo	Foreign Document Number	Concise Description of Relevance
1	2004094980	<p>Fontijn discloses a method and system of peer-to-peer content transfer over the internet(1:1-5; 2:1-8). The content is identified by specific selection criteria entered by user and transmitted to a server(4:10-19).</p> <p>It is commonly known in the art and disclosed in 4:14-18, that devices in P2P networks request content, and as such, it would be understood that the devices requested the content, and such requests would include an identification of the content and a destination identifier.</p> <p>Fontijn discloses a group consisting of a plurality of P2P devices (Fig. 1; 4:32-5:3). Devices on the network are identified by IP address (2:12-16).</p> <p>Fontijn discloses that a second device is included in the group and selected out of the devices in the group (Fig. 1; 4:32-5:3). The Server can select the device (4:21-22). Fig 1. shows a network of devices and a server, where the server will redirect the transfer of content to [the] device <.> which then provide the content (4:21-22).</p>

NON-PATENT PUBLICATIONS

Cite No	Reference	Concise Description of Relevance

--	--	--

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16140749

U.S. PATENTS

Cite No	Patent Number	Kind Code ¹	Issue Date (YYYY-MM-DD)	First Named Inventor

U.S. PATENT APPLICATION PUBLICATIONS

Cite No	Publication Number	Kind Code ¹	Publication Date (YYYY-MM-DD)	First Named Inventor
1	20070142036	A1	2007-06-21	Johan Wikman
2	20090248793	A1	2009-10-01	Sanny Jacobsson

FOREIGN PATENTS AND PUBLISHED FOREIGN PATENT APPLICATIONS

Cite No	Foreign Document Number ³	Country Code ²	Kind Code ¹	Publication Date (YYYY-MM-DD)	Applicant, Patentee or First Named Inventor	T ⁵
1	2004094980	WO	A2	2004-11-04	Wilhelmus F. J. Fontijn	<input type="checkbox"/>

NON-PATENT PUBLICATIONS (e.g., journal article, Office action)

Cite No	Author (if any), title of the publication, page(s) being submitted, publication date, publisher (where available), place of publication (where available).	T ⁵	E ⁶

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16140749

		<input type="checkbox"/>	<input type="checkbox"/>
--	--	--------------------------	--------------------------

STATEMENTS

The party making the submission is not an individual who has a duty to disclose information with respect to the above-identified application under 37 CFR 1.56.

This submission complies with the requirements of 35 U.S.C. 122(e) and 37 CFR 1.290.

The fee set forth in 37 CFR 1.290(f) has been submitted herewith.

The fee set forth in 37 CFR 1.290(f) is not required because this submission lists three or fewer total items and, to the knowledge of the person signing the statement after making reasonable inquiry, this submission is the first and the only submission under 35 U.S.C 122(e) filed in the above-identified application by the party making the submission or by a party in privity with the party.

This resubmission is being made responsive to a notification of non-compliance issued for an earlier filed third-party submission. The corrections in this resubmission are limited to addressing the non-compliance. As such, the party making this resubmission: (1) requests that the Office apply the previously-paid fee set forth in 37 CFR 1.290(f), or (2) states that no fee is required to accompany this resubmission as the undersigned is again making the fee exemption statement set forth in 37 CFR 1.290(g).

Signature	/Jurate Breimelyte/	
Name/Print	Jurate Breimelyte	Registration Number (if applicable)

Examiner Signature	Date Considered	
--------------------	-----------------	--

***EXAMINER:** Signature indicates all documents listed above have been considered, except for citations through which a line is drawn. Draw line through citation if not considered. Include a copy of this form with next communication to applicant. 1. If known, enter kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16. See MPEP 901.04(a). 2. Enter the country or patent office that issued the document, by two-letter code under WIPO standard ST.3. See MPEP 1851. 3. For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. 4. If known, enter the kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16. See MPEP 901.04(a). 5. Check mark indicates translation attached. 6. Check mark indicates evidence of publication attached.

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290 CONCISE DESCRIPTION OF RELEVANCE		
Application Number	16140785	
U.S. PATENTS		
Cite No	Patent Number	Concise Description of Relevance
U.S. PATENT APPLICATION PUBLICATION		
Cite No	Publication Number	Concise Description of Relevance
1	20090248793	In Jacobsson, para. 64 discloses the relationship of client and server by virtue of computer programs running on the respective computers and having a client-server relationship to each other. Claim 9 similarly to the present invention discloses a method for providing content in a network: receiving, at the device and from the first server device, source information that identifies a selected content and obtaining the content using the received source information. Claim 16 further discloses sending a confirmation from the device to the first server device, the confirmation including information about a result of the requests generated by the device. Jacobsson discloses the selection of the content source from the list of content sources. The selection process can include an optimization/prioritization algorithm that considers the geographic location of the content source (paras 5-6, 35-36). This information is sent from the device to the server.

2	20070142036	Wikman (Claim 5, para 21; 67) discloses that a mobile intermediate device is used as a reverse proxy. An ordinary forward proxy is an intermediate server that sits between the client and the origin server. In order to get content from the origin server, the client sends a request to the proxy naming the origin server as the target and the proxy then requests the content from the origin server and returns it to the client. The reverse proxy then decides where to send those requests, and returns the content as if it was itself the origin (para 68). The mobile intermediate device address may be globally or locally unique. (para 65) discloses that the intermediate server, is configured to act as a reverse proxy. The access request can be forwarded by using, for instance, any suitable near field communication technique.
---	-------------	--

FOREIGN PATENT DOCUMENTS

CiteNo	Foreign Document Number	Concise Description of Relevance

NON-PATENT PUBLICATIONS

Cite No	Reference	Concise Description of Relevance

--	--	--

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16140785

U.S. PATENTS

Cite No	Patent Number	Kind Code ¹	Issue Date (YYYY-MM-DD)	First Named Inventor

U.S. PATENT APPLICATION PUBLICATIONS

Cite No	Publication Number	Kind Code ¹	Publication Date (YYYY-MM-DD)	First Named Inventor
1	20070142036	A1	2007-06-21	Johan Wikman
2	20090248793	A1	2009-10-01	Sanny Jacobsson

FOREIGN PATENTS AND PUBLISHED FOREIGN PATENT APPLICATIONS

Cite No	Foreign Document Number ³	Country Code ²	Kind Code ¹	Publication Date (YYYY-MM-DD)	Applicant, Patentee or First Named Inventor	T ⁵
						<input type="checkbox"/>

NON-PATENT PUBLICATIONS (e.g., journal article, Office action)

Cite No	Author (if any), title of the publication, page(s) being submitted, publication date, publisher (where available), place of publication (where available).	T ⁵	E ⁶

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16140785

		<input type="checkbox"/>	<input type="checkbox"/>
--	--	--------------------------	--------------------------

STATEMENTS

The party making the submission is not an individual who has a duty to disclose information with respect to the above-identified application under 37 CFR 1.56.

This submission complies with the requirements of 35 U.S.C. 122(e) and 37 CFR 1.290.

The fee set forth in 37 CFR 1.290(f) has been submitted herewith.

The fee set forth in 37 CFR 1.290(f) is not required because this submission lists three or fewer total items and, to the knowledge of the person signing the statement after making reasonable inquiry, this submission is the first and the only submission under 35 U.S.C 122(e) filed in the above-identified application by the party making the submission or by a party in privity with the party.

This resubmission is being made responsive to a notification of non-compliance issued for an earlier filed third-party submission. The corrections in this resubmission are limited to addressing the non-compliance. As such, the party making this resubmission: (1) requests that the Office apply the previously-paid fee set forth in 37 CFR 1.290(f), or (2) states that no fee is required to accompany this resubmission as the undersigned is again making the fee exemption statement set forth in 37 CFR 1.290(g).

Signature	/Jurate Breimelyte/		
Name/Print	Jurate Breimelyte	Registration Number (if applicable)	

Examiner Signature		Date Considered	
--------------------	--	-----------------	--

***EXAMINER:** Signature indicates all documents listed above have been considered, except for citations through which a line is drawn. Draw line through citation if not considered. Include a copy of this form with next communication to applicant. 1. If known, enter kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16. See MPEP 901.04(a). 2. Enter the country or patent office that issued the document, by two-letter code under WIPO standard ST.3. See MPEP 1851. 3. For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. 4. If known, enter the kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16. See MPEP 901.04(a). 5. Check mark indicates translation attached. 6. Check mark indicates evidence of publication attached.

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290 CONCISE DESCRIPTION OF RELEVANCE		
Application Number	16214433	
U.S. PATENTS		
Cite No	Patent Number	Concise Description of Relevance
U.S. PATENT APPLICATION PUBLICATION		
Cite No	Publication Number	Concise Description of Relevance
1	20110066924	Dorso discloses that the user device selects at least one relay peer based on the availability and capabilities of the relay peers. An optimization scheme may be employed to make such a selection (para 20). First user device 105 then contacts each of the potential peers and assesses the potential availability and capabilities of the peers (para 31). It should be appreciated that the factors for the optimization scheme may include, but are not limited to the amount of available resources of the relay peer, the historic reliability of the relay peer, and the length of time the relay peer has been running the protocol (para 32). Claims 3 and 15 of Dorso discloses the computer-implemented method where factors relating to said peer nodes comprises an amount of available resources of said peer nodes, wherein said resources are CPU and bandwidth.

2	20130080575	<p>Prince discloses a cloud-based proxy service provisioned through DNS. The proxy service is available over the Internet and does not require customers to install new hardware or software in order to support the service (para 13) (like in the specification of the present application). The proxy server (which may be identified to the client by a DNS server (e.g., para77)) may perform one or more actions on an incoming request received from a client device: determine whether the request is malformed; determine the type and/or size of the requested content; determine whether the origin server is offline; and determine whether the requested content is available in cache (para17). Responsive to the proxy server receiving a request (e.g., HTTP request) from a client device on an IPv4 connection and determining that the request should be transmitted to a destination origin server using an IPv6 connection, the proxy server transmits that request using the IPv6 connection, or vice versa (para 31).</p>
---	-------------	--

FOREIGN PATENT DOCUMENTS

CiteNo	Foreign Document Number	Concise Description of Relevance

NON-PATENT PUBLICATIONS

Cite No	Reference	Concise Description of Relevance

--	--	--

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16214433

U.S. PATENTS

Cite No	Patent Number	Kind Code ¹	Issue Date (YYYY-MM-DD)	First Named Inventor

U.S. PATENT APPLICATION PUBLICATIONS

Cite No	Publication Number	Kind Code ¹	Publication Date (YYYY-MM-DD)	First Named Inventor
1	20130080575	A1	2013-03-28	Matthew Browning Prince
2	20110066924	A1	2011-03-17	Gregory Dorso

FOREIGN PATENTS AND PUBLISHED FOREIGN PATENT APPLICATIONS

Cite No	Foreign Document Number ³	Country Code ²	Kind Code ¹	Publication Date (YYYY-MM-DD)	Applicant, Patentee or First Named Inventor	T ⁵
						<input type="checkbox"/>

NON-PATENT PUBLICATIONS (e.g., journal article, Office action)

Cite No	Author (if any), title of the publication, page(s) being submitted, publication date, publisher (where available), place of publication (where available).	T ⁵	E ⁶

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16214433

		<input type="checkbox"/>	<input type="checkbox"/>
--	--	--------------------------	--------------------------

STATEMENTS

The party making the submission is not an individual who has a duty to disclose information with respect to the above-identified application under 37 CFR 1.56.

This submission complies with the requirements of 35 U.S.C. 122(e) and 37 CFR 1.290.

The fee set forth in 37 CFR 1.290(f) has been submitted herewith.

The fee set forth in 37 CFR 1.290(f) is not required because this submission lists three or fewer total items and, to the knowledge of the person signing the statement after making reasonable inquiry, this submission is the first and the only submission under 35 U.S.C 122(e) filed in the above-identified application by the party making the submission or by a party in privity with the party.

This resubmission is being made responsive to a notification of non-compliance issued for an earlier filed third-party submission. The corrections in this resubmission are limited to addressing the non-compliance. As such, the party making this resubmission: (1) requests that the Office apply the previously-paid fee set forth in 37 CFR 1.290(f), or (2) states that no fee is required to accompany this resubmission as the undersigned is again making the fee exemption statement set forth in 37 CFR 1.290(g).

Signature	/Jurate Breimelyte/	
-----------	---------------------	--

Name/Print	Jurate Breimelyte	Registration Number (if applicable)	
------------	-------------------	--	--

Examiner Signature		Date Considered	
--------------------	--	-----------------	--

***EXAMINER:** Signature indicates all documents listed above have been considered, except for citations through which a line is drawn. Draw line through citation if not considered. Include a copy of this form with next communication to applicant. 1. If known, enter kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16. See MPEP 901.04(a). 2. Enter the country or patent office that issued the document, by two-letter code under WIPO standard ST.3. See MPEP 1851. 3. For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. 4. If known, enter the kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16. See MPEP 901.04(a). 5. Check mark indicates translation attached. 6. Check mark indicates evidence of publication attached.

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290 CONCISE DESCRIPTION OF RELEVANCE		
Application Number	16214451	
U.S. PATENTS		
Cite No	Patent Number	Concise Description of Relevance
U.S. PATENT APPLICATION PUBLICATION		
Cite No	Publication Number	Concise Description of Relevance
1	20130080575	Prince discloses a method for fetching a first content over the Internet identified by a first content identifier by disclosing fetching 'network resources' from origin servers, where the content may include 'HTML pages, images, word processing documents, PDF files, movie files, music files, or other computer files' (paras 26, 27, 15, 17-19 ('requested content')). Prince discloses that the content is identified in requests from client, for instance in HTTP requests sent to proxy server (para 45 (client request includes parameters such as 'the type of requested content'), paras 46, 27, 39; Fig. 1). Prince discloses sending the first identifier to the first server by disclosing the client sends its IP address to DNS system when resolving a proxy server that will retrieve content from original server (paras 15, 77 ('client device 110A requests an IPv4 address for example.com (thus, the client device 110A is an IPv4 enabled client)'), par 90 (same); Figs. 5 (step 510), 6 (step 610)).

2	20110066924	Dorso discloses fetching content over the Internet from a first server via a group of multiple devices. A first user device (105, see Fig. 1) can send communications to, or receive communications from, a second user device (e.g., 140 in Fig. 1) and/or a third user device (e.g., 145 in Fig. 1) (each of which could be 'a first server') via a number of relay peers (120, 125, 130, 135 in Fig. 1) ('which could mean a group of multiple devices')(also paras 31, 33, 41). Like in the claimed invention, Dorso discloses that the second user device can be a handheld mobile device (paras 18, 23-24, Claim 28, 34). Discloses that each device in the group of multiple devices is identified in the Internet by an associated group device identifier (par30 ('tracker peer 110 provides potential relay peers to first user device 105. E.g., tracker peer may provide data identifying first relay peer 120, second relay peer 125, and third relay peer 130 as peers or nodes associated with P2P computer environment').
3	20070142036	Wikman discloses a method of providing content using a mobile intermediate device, the mobile intermediate device for providing content from a geographically proximate origin device to a remote client device, the mobile intermediate device having a mobile intermediate device address to which content requests are communicated from the client device (Claim 1; para 24). Claim 5 and 6 discloses that the mobile intermediate device is arranged to function as a reverse proxy server and such device address is globally unique. It also discloses that content to a remote device is delivered upon receiving a request from the client device and that processing unit ('first server') is an actor between the client device and origin devices ('second server')(Claim 1). Claim 3 discloses that mobile intermediate device comprises a web server and the mobile intermediate device address is the web server address, and the content request is addressed to the web server (equal to present invention Claim 22,26)

FOREIGN PATENT DOCUMENTS

CiteNo	Foreign Document Number	Concise Description of Relevance

NON-PATENT PUBLICATIONS		
Cite No	Reference	Concise Description of Relevance

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16214451

U.S. PATENTS

Cite No	Patent Number	Kind Code ¹	Issue Date (YYYY-MM-DD)	First Named Inventor

U.S. PATENT APPLICATION PUBLICATIONS

Cite No	Publication Number	Kind Code ¹	Publication Date (YYYY-MM-DD)	First Named Inventor
1	20070142036	A1	2007-06-21	Johan Wikman
2	20110066924	A1	2011-03-17	Gregory Dorso
3	20130080575	A1	2013-03-28	Matthew Browning Prince

FOREIGN PATENTS AND PUBLISHED FOREIGN PATENT APPLICATIONS

Cite No	Foreign Document Number ³	Country Code ²	Kind Code ¹	Publication Date (YYYY-MM-DD)	Applicant, Patentee or First Named Inventor	T ⁵
						<input type="checkbox"/>

NON-PATENT PUBLICATIONS (e.g., journal article, Office action)

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16214451

Cite No	Author (if any), title of the publication, page(s) being submitted, publication date, publisher (where available), place of publication (where available).	T ⁵	E ⁶
		<input type="checkbox"/>	<input type="checkbox"/>

STATEMENTS

The party making the submission is not an individual who has a duty to disclose information with respect to the above-identified application under 37 CFR 1.56.

This submission complies with the requirements of 35 U.S.C. 122(e) and 37 CFR 1.290.

The fee set forth in 37 CFR 1.290(f) has been submitted herewith.

The fee set forth in 37 CFR 1.290(f) is not required because this submission lists three or fewer total items and, to the knowledge of the person signing the statement after making reasonable inquiry, this submission is the first and the only submission under 35 U.S.C. 122(e) filed in the above-identified application by the party making the submission or by a party in privity with the party.

This resubmission is being made responsive to a notification of non-compliance issued for an earlier filed third-party submission. The corrections in this resubmission are limited to addressing the non-compliance. As such, the party making this resubmission: (1) requests that the Office apply the previously-paid fee set forth in 37 CFR 1.290(f), or (2) states that no fee is required to accompany this resubmission as the undersigned is again making the fee exemption statement set forth in 37 CFR 1.290(g).

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290 CONCISE DESCRIPTION OF RELEVANCE		
Application Number	16214476	
U.S. PATENTS		
Cite No	Patent Number	Concise Description of Relevance
U.S. PATENT APPLICATION PUBLICATION		
Cite No	Publication Number	Concise Description of Relevance
1	20110066924	Dorso discloses fetching content over the Internet via a group of devices. The 'first user device 105 may use a web browser application to practice the present technology,' and '[t]he purpose of tracker peer 110 is to receive a request from a user device such as first user device' (par. 29-30). It is implicit that the user device could use an IP address, host name, or similar identifier to communicate with the tracker peer. In one embodiment the user device may divide the communication into portions and each portion is sent via a different relay peer (par. 21-22). In one embodiment, a P2P computer environment is a distributed network architecture that is composed of peers that make a portion of resources available directly to their peers without intermediary servers (par. 25;36-37; Claim 25).

2	20090248793	Jacobsson discloses a method of retrieving content from a network (e.g. the Internet (par. 19)). The content is identified by a content identifier (see e.g.,paras 34-35). Jacobsson teaches the coordinating system (i.e., first server) providing the list of content sources to the client device (paras 36;39;52). The content source may be identified by public or private IP address (par 35). Jacobsson discloses the first device requesting and receiving the identified content (or portions thereof) from the various content sources (web servers and/or client devices) (par 53).
3	20070142036	Wikman discloses a method of providing content using a mobile intermediate device, the mobile intermediate device for providing content from a geographically proximate origin device to a geographically remote client device, the mobile intermediate device having a mobile intermediate device address to which content requests are communicated (Claim 1; par 24). Claim 5 and 6 discloses that the mobile intermediate device is arranged to function as a reverse proxy server and such device address is globally unique. It also discloses that content to a remote device is delivered upon receiving a request from the client device and that processing unit ('first server') is an actor between the client device and origin devices ('second server')(Claim 1). Claim 3 discloses that mobile intermediate device comprises a web server and the mobile intermediate device address is the web server address, and the content request is addressed to the web server(as also recited in present invention Claim 22,26)

FOREIGN PATENT DOCUMENTS

CiteNo	Foreign Document Number	Concise Description of Relevance

NON-PATENT PUBLICATIONS		
Cite No	Reference	Concise Description of Relevance

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16214476

U.S. PATENTS

Cite No	Patent Number	Kind Code ¹	Issue Date (YYYY-MM-DD)	First Named Inventor

U.S. PATENT APPLICATION PUBLICATIONS

Cite No	Publication Number	Kind Code ¹	Publication Date (YYYY-MM-DD)	First Named Inventor
1	20070142036	A1	2007-06-21	Johan Wikman
2	20090248793	A1	2009-10-01	Sanny Jacobsson
3	20110066924	A1	2011-03-17	Gregory Dorso

FOREIGN PATENTS AND PUBLISHED FOREIGN PATENT APPLICATIONS

Cite No	Foreign Document Number ³	Country Code ²	Kind Code ¹	Publication Date (YYYY-MM-DD)	Applicant, Patentee or First Named Inventor	T ⁵
						<input type="checkbox"/>

NON-PATENT PUBLICATIONS (e.g., journal article, Office action)

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16214476

Cite No	Author (if any), title of the publication, page(s) being submitted, publication date, publisher (where available), place of publication (where available).	T ⁵	E ⁶
		<input type="checkbox"/>	<input type="checkbox"/>

STATEMENTS
<p>The party making the submission is not an individual who has a duty to disclose information with respect to the above-identified application under 37 CFR 1.56.</p> <p>This submission complies with the requirements of 35 U.S.C. 122(e) and 37 CFR 1.290.</p> <p><input type="checkbox"/> The fee set forth in 37 CFR 1.290(f) has been submitted herewith.</p> <p><input checked="" type="checkbox"/> The fee set forth in 37 CFR 1.290(f) is not required because this submission lists three or fewer total items and, to the knowledge of the person signing the statement after making reasonable inquiry, this submission is the first and the only submission under 35 U.S.C. 122(e) filed in the above-identified application by the party making the submission or by a party in privity with the party.</p> <p><input type="checkbox"/> This resubmission is being made responsive to a notification of non-compliance issued for an earlier filed third-party submission. The corrections in this resubmission are limited to addressing the non-compliance. As such, the party making this resubmission: (1) requests that the Office apply the previously-paid fee set forth in 37 CFR 1.290(f), or (2) states that no fee is required to accompany this resubmission as the undersigned is again making the fee exemption statement set forth in 37 CFR 1.290(g).</p>

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290 CONCISE DESCRIPTION OF RELEVANCE		
Application Number	16214496	
U.S. PATENTS		
Cite No	Patent Number	Concise Description of Relevance
U.S. PATENT APPLICATION PUBLICATION		
Cite No	Publication Number	Concise Description of Relevance
1	20090248793	Jacobsson discloses multiple devices, including clients (peers) and servers that can be content sources (par23-24,42). The various content sources are identified by public or private IP address (par35).If portions of content are being requested from multiple content sources, the request, for example, for a client device to provide the embedded video from a certain web page on a server would inherently include identifying information about that server so that the client device could provide the proper video (par 36-38).Jacobsson teaches the coordinating system (i. e., first server) providing the list of content sources to the client device (par 36;39;52).Claims 17-18 disclose forwarding, in response to the identifier: i) source information to the device that identifies each selected content source; ii) object information and source information the object information describing the object and the multiple content portions, the source information identifying each selected content source

2	20110066924	Dorso discloses fetching content over the Internet via a group of devices. 'The first user device 105 may use a web browser application to practice the presented technology' and 'the purpose of the tracker peer 110 is to receive a request from a user device such as first user device' (par 29-30). It is implicit that the user device could use an IP address, host name, or similar identifier to communicate with the tracker peer. In one embodiment the user device may divide the communication into portions and each portion is sent via a different replay peer (par. 21-22). In one embodiment, a P2P computer environment is a distributed network architecture that is composed of peers that make a portion of resources available directly to their peers without intermediary servers (par 25; 36-37, Claim 25).
---	-------------	--

FOREIGN PATENT DOCUMENTS

CiteNo	Foreign Document Number	Concise Description of Relevance

NON-PATENT PUBLICATIONS

Cite No	Reference	Concise Description of Relevance

--	--	--

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16214496

U.S. PATENTS

Cite No	Patent Number	Kind Code ¹	Issue Date (YYYY-MM-DD)	First Named Inventor

U.S. PATENT APPLICATION PUBLICATIONS

Cite No	Publication Number	Kind Code ¹	Publication Date (YYYY-MM-DD)	First Named Inventor
1	20110066924	A1	2011-03-17	Gregory Dorso
2	20090248793	A1	2009-10-01	Sanny Jacobsson

FOREIGN PATENTS AND PUBLISHED FOREIGN PATENT APPLICATIONS

Cite No	Foreign Document Number ³	Country Code ²	Kind Code ¹	Publication Date (YYYY-MM-DD)	Applicant, Patentee or First Named Inventor	T ⁵
						<input type="checkbox"/>

NON-PATENT PUBLICATIONS (e.g., journal article, Office action)

Cite No	Author (if any), title of the publication, page(s) being submitted, publication date, publisher (where available), place of publication (where available).	T ⁵	E ⁶

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16214496

		<input type="checkbox"/>	<input type="checkbox"/>
--	--	--------------------------	--------------------------

STATEMENTS

The party making the submission is not an individual who has a duty to disclose information with respect to the above-identified application under 37 CFR 1.56.

This submission complies with the requirements of 35 U.S.C. 122(e) and 37 CFR 1.290.

The fee set forth in 37 CFR 1.290(f) has been submitted herewith.

The fee set forth in 37 CFR 1.290(f) is not required because this submission lists three or fewer total items and, to the knowledge of the person signing the statement after making reasonable inquiry, this submission is the first and the only submission under 35 U.S.C. 122(e) filed in the above-identified application by the party making the submission or by a party in privity with the party.

This resubmission is being made responsive to a notification of non-compliance issued for an earlier filed third-party submission. The corrections in this resubmission are limited to addressing the non-compliance. As such, the party making this resubmission: (1) requests that the Office apply the previously-paid fee set forth in 37 CFR 1.290(f), or (2) states that no fee is required to accompany this resubmission as the undersigned is again making the fee exemption statement set forth in 37 CFR 1.290(g).

Signature	/Jurate Breimelyte/		
Name/Print	Jurate Breimelyte	Registration Number (if applicable)	

Examiner Signature		Date Considered	
--------------------	--	-----------------	--

***EXAMINER:** Signature indicates all documents listed above have been considered, except for citations through which a line is drawn. Draw line through citation if not considered. Include a copy of this form with next communication to applicant. 1. If known, enter kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16. See MPEP 901.04(a). 2. Enter the country or patent office that issued the document, by two-letter code under WIPO standard ST.3. See MPEP 1851. 3. For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. 4. If known, enter the kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16. See MPEP 901.04(a). 5. Check mark indicates translation attached. 6. Check mark indicates evidence of publication attached.

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290 CONCISE DESCRIPTION OF RELEVANCE		
Application Number	16292363	
U.S. PATENTS		
Cite No	Patent Number	Concise Description of Relevance
U.S. PATENT APPLICATION PUBLICATION		
Cite No	Publication Number	Concise Description of Relevance
1	20110066924	Dorso discloses fetching content over the Internet from a first server via a group of multiple devices. In Dorso, a first user device (105, see Fig. 1 below) can send communications to, or receive communications from, a second user device (e.g., 140 in Fig. 1) and/or a third user device (e.g., 145 in Fig. 1) (each could be 'a first server') via a number of relay peers (120, 125, 130, 135 in Fig. 1) (equal to 'a group of client devices' in present application) (par 31, 33, 41). The 'first user device 105 may use a web browser application to practice the present technology' and '[t]he purpose of tracker peer 110 is to receive a request from a user device such as first user device' (Par 29-30). It is implicit that the user device could use an IP address, host name, or similar identifier to communicate with the tracker peer. In one embodiment the user device may divide the communication into portions and each portion is sent via a different relay peer (par 21-22).

2	20130080575	<p>Prince discloses a method for fetching a first content over the Internet identified by a first content identifier by disclosing fetching 'network resources' from origin servers 130, where the content may include 'HTML pages, images, word processing documents, PDF files, movie files, music files, or other computer files.' par. 26, 27; 15, 17-19.</p> <p>Prince discloses that the second device is included in the group and that the step of selecting the second device out of the devices in the group by disclosing multiple proxy servers 120 and the DNS System 140 selecting one of the proxy servers 120 based, for example, on the origin server 130 the client device 110 is requesting to resolve. See, e.g., par. 13, 15, 23.</p> <p>In the Prince, the disclosed method of fetching the content over the Internet as provided in Fig. 5 corresponds to, for example, the disclosure of Fig. 5b in the present patent application. Both disclose the same method of receiving the content through the additional proxy.</p>
3	20090248793	<p>Jacobsson discloses multiple devices, including clients (peers) and servers that can be content sources. par. 23-24,42. The various content sources are identified by public or private IP address (par. 35).</p> <p>Jacobsson discloses that requests for content can be sent to any of the content sources (client devices and/or servers)(par. 36-38).</p> <p>Jacobsson discloses the first device requesting and receiving the identified content (or portions thereof) from the various content sources (web servers and/or client devices) (par. 53).</p> <p>Claim 9 of Jacobsson and its dependent claims disclose a method for providing content in a network, the method comprising: forwarding an identifier associated with an object from a device to a first server device in a network; receiving, at the device and from the first server device, source information that identifies a selected content source; and obtaining the content using the received source information.</p>

FOREIGN PATENT DOCUMENTS

CiteNo	Foreign Document Number	Concise Description of Relevance

NON-PATENT PUBLICATIONS		
Cite No	Reference	Concise Description of Relevance

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16292363

U.S. PATENTS

Cite No	Patent Number	Kind Code ¹	Issue Date (YYYY-MM-DD)	First Named Inventor

U.S. PATENT APPLICATION PUBLICATIONS

Cite No	Publication Number	Kind Code ¹	Publication Date (YYYY-MM-DD)	First Named Inventor
1	20090248793	A1	2009-10-01	Sanny Jacobsson
2	20130080575	A1	2013-03-28	Matthew Browning Prince
3	20110066924	A1	2011-03-17	Gregory Dorso

FOREIGN PATENTS AND PUBLISHED FOREIGN PATENT APPLICATIONS

Cite No	Foreign Document Number ³	Country Code ²	Kind Code ¹	Publication Date (YYYY-MM-DD)	Applicant, Patentee or First Named Inventor	T ⁵
						<input type="checkbox"/>

NON-PATENT PUBLICATIONS (e.g., journal article, Office action)

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16292363

Cite No	Author (if any), title of the publication, page(s) being submitted, publication date, publisher (where available), place of publication (where available).	T ⁵	E ⁶
		<input type="checkbox"/>	<input type="checkbox"/>

STATEMENTS

The party making the submission is not an individual who has a duty to disclose information with respect to the above-identified application under 37 CFR 1.56.

This submission complies with the requirements of 35 U.S.C. 122(e) and 37 CFR 1.290.

The fee set forth in 37 CFR 1.290(f) has been submitted herewith.

The fee set forth in 37 CFR 1.290(f) is not required because this submission lists three or fewer total items and, to the knowledge of the person signing the statement after making reasonable inquiry, this submission is the first and the only submission under 35 U.S.C. 122(e) filed in the above-identified application by the party making the submission or by a party in privity with the party.

This resubmission is being made responsive to a notification of non-compliance issued for an earlier filed third-party submission. The corrections in this resubmission are limited to addressing the non-compliance. As such, the party making this resubmission: (1) requests that the Office apply the previously-paid fee set forth in 37 CFR 1.290(f), or (2) states that no fee is required to accompany this resubmission as the undersigned is again making the fee exemption statement set forth in 37 CFR 1.290(g).

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290 CONCISE DESCRIPTION OF RELEVANCE		
Application Number	16292364	
U.S. PATENTS		
Cite No	Patent Number	Concise Description of Relevance
U.S. PATENT APPLICATION PUBLICATION		
Cite No	Publication Number	Concise Description of Relevance
1	20090248793	<p>Identically to the present invention, the prior art describes a method for providing content in a network. It also discloses receiving the identifier from a device, selecting a content source in the network, forwarding, in response to the identifier, source information to the device (Claim 1). Para 0006 and Claim 3 of the prior art discloses that the optimization algorithm can be configured so that the at least one content source is selected according to at least one approach the content source, for example the content source being located in a geographic location in which the device is located. Prior art also discloses the use of a geographic location with which the content source is associated. Para 0036 and Claim 4 discloses that in some implementations of the invention, preference can be given to the selection of content sources associated with the same geographic location (e.g., country), network, and/or internet service provider as the client device.</p>

FOREIGN PATENT DOCUMENTS		
CiteNo	Foreign Document Number	Concise Description of Relevance
1	2004094980	Fontijn discloses a method and system of peer-to-peer content transfer over the internet. See 1:1-5; 2:1-8. The content is identified by specific selection criteria entered by user and transmitted to a server or device storing the content. See 4:10-19. Fontijn discloses a group consisting of a plurality of P2P devices 12, 13, 14, 15, 16, and 17. See Fig. 1; 4:32-5:3. As in the recited application, in Fontijn the devices on the network are identified by IP address. See 2:12-16. Fontijn further discloses selecting the second device out of the devices in the group. See 2:5-16. See Fig. 1; 4:9-29. Fontijn discloses that it is known that users can select a device, see, e.g., 2:12-16, and that the Server can select the device when it redirects the first device. See, e.g., 4:21-22. Fontijn discloses that the second device may be selected based on the past activity of having downloaded the content. See, e.g., 8:8-10; Fig 2 (Step 400).

NON-PATENT PUBLICATIONS		
Cite No	Reference	Concise Description of Relevance

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16292364

U.S. PATENTS						
Cite No	Patent Number	Kind Code ¹	Issue Date (YYYY-MM-DD)	First Named Inventor		
U.S. PATENT APPLICATION PUBLICATIONS						
Cite No	Publication Number	Kind Code ¹	Publication Date (YYYY-MM-DD)	First Named Inventor		
1	20090248793	A1	2009-10-01	Sanny Jacobsson		
FOREIGN PATENTS AND PUBLISHED FOREIGN PATENT APPLICATIONS						
Cite No	Foreign Document Number ³	Country Code ²	Kind Code ¹	Publication Date (YYYY-MM-DD)	Applicant, Patentee or First Named Inventor	T ⁵
1	2004094980	WO	A2	2004-11-04	Wilhelmus F J Fontijn	<input type="checkbox"/>
NON-PATENT PUBLICATIONS (e.g., journal article, Office action)						
Cite No	Author (if any), title of the publication, page(s) being submitted, publication date, publisher (where available), place of publication (where available).				T ⁵	E ⁶

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16292364

		<input type="checkbox"/>	<input type="checkbox"/>
--	--	--------------------------	--------------------------

STATEMENTS

The party making the submission is not an individual who has a duty to disclose information with respect to the above-identified application under 37 CFR 1.56.

This submission complies with the requirements of 35 U.S.C. 122(e) and 37 CFR 1.290.

The fee set forth in 37 CFR 1.290(f) has been submitted herewith.

The fee set forth in 37 CFR 1.290(f) is not required because this submission lists three or fewer total items and, to the knowledge of the person signing the statement after making reasonable inquiry, this submission is the first and the only submission under 35 U.S.C 122(e) filed in the above-identified application by the party making the submission or by a party in privity with the party.

This resubmission is being made responsive to a notification of non-compliance issued for an earlier filed third-party submission. The corrections in this resubmission are limited to addressing the non-compliance. As such, the party making this resubmission: (1) requests that the Office apply the previously-paid fee set forth in 37 CFR 1.290(f), or (2) states that no fee is required to accompany this resubmission as the undersigned is again making the fee exemption statement set forth in 37 CFR 1.290(g).

Signature	/Jurate Breimelyte/		
Name/Print	Jurate Breimelyte	Registration Number (if applicable)	

Examiner Signature		Date Considered	
--------------------	--	-----------------	--

***EXAMINER:** Signature indicates all documents listed above have been considered, except for citations through which a line is drawn. Draw line through citation if not considered. Include a copy of this form with next communication to applicant. 1. If known, enter kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16. See MPEP 901.04(a). 2. Enter the country or patent office that issued the document, by two-letter code under WIPO standard ST.3. See MPEP 1851. 3. For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. 4. If known, enter the kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16. See MPEP 901.04(a). 5. Check mark indicates translation attached. 6. Check mark indicates evidence of publication attached.

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290 CONCISE DESCRIPTION OF RELEVANCE		
Application Number		16292374
U.S. PATENTS		
Cite No	Patent Number	Concise Description of Relevance
U.S. PATENT APPLICATION PUBLICATION		
Cite No	Publication Number	Concise Description of Relevance
1	20090248793	Jacobsson discloses a method of retrieving content from a network (e.g. the Internet (par19)). The content is identified by a content identifier (see e.g., par34-35). It also discloses that other metadata may be associated with content and content sources. Jacobsson discloses the selection of the content source from the list of content sources. The selection process can include a optimization/prioritization algorithm that considers the geographic location of the content source (par 5-6,35-36). Also discloses multiple devices, including clients (peers) and servers that can be content sources (23-24,42). The various content sources are identified by public or private IP address (35). Jacobsson discloses that requests for content can be sent to any of the content sources (client devices and/or servers). Also discloses the first device requesting and receiving the identified content (or portions thereof) from the various content sources (web servers and/or client devices)(53).

2	20070142036	Wikman discloses a method for providing content to a remotely located electronic device, which may be connectable to the Internet, by accessing content on a device (e.g. Internet server). It discloses accessing content on an origin device using a client device via intermediate device, with the knowledge of how to access the intermediate device, but without knowing the specific access information of the origin device. In this way, a client device can access the content (website, other content) contained on an origin device via intermediate device (par 19. Para 25 explains providing user selected content from respective origin device using near field connectivity between the intermediate device and the respective geographically proximate origin device from which the user selected content is sourced; and providing the content from the origin device to the client via the intermediate device using the intermediate device as the content source. The method may be performed using the Internet.
---	-------------	--

FOREIGN PATENT DOCUMENTS

CiteNo	Foreign Document Number	Concise Description of Relevance
1	2004094980	Fontijn discloses the transfer of content in P2P system. The device (first device) receives the selection criterion which it then sends to another device (a second device) and to a server. The said first device cannot know whether the server or another P2P device, has the requested content available. If the server has the content satisfying the selection criterion, it will provide it to the requesting device. However, in order to distribute network usage more efficient - if the server knows that another peer (device) has the requested content, the server will redirect the transfer of content to this device (p. 4, lines 10-30). Figure 2 shows a method of P2P transfer of content. The content is transferred among device in the P2P network, as a starting point only the server can provide content; later on content can be distributed to various devices. Claim 4 discloses that content can be very different: a DVD picture and sound signal; a given digital movie format; a given picture format.

NON-PATENT PUBLICATIONS

Cite No	Reference	Concise Description of Relevance

--	--	--

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16292374

U.S. PATENTS						
Cite No	Patent Number	Kind Code ¹	Issue Date (YYYY-MM-DD)	First Named Inventor		
U.S. PATENT APPLICATION PUBLICATIONS						
Cite No	Publication Number	Kind Code ¹	Publication Date (YYYY-MM-DD)	First Named Inventor		
1	20070142036		2007-06-21	Johan Wikman		
2	20090248793	A1	2009-10-01	Sanny Jacobsson		
FOREIGN PATENTS AND PUBLISHED FOREIGN PATENT APPLICATIONS						
Cite No	Foreign Document Number ³	Country Code ²	Kind Code ¹	Publication Date (YYYY-MM-DD)	Applicant, Patentee or First Named Inventor	T ⁵
1	2004094980	WO	A2	2004-11-04	Wilhelmus F. J. Fontijn	<input type="checkbox"/>
NON-PATENT PUBLICATIONS (e.g., journal article, Office action)						
Cite No	Author (if any), title of the publication, page(s) being submitted, publication date, publisher (where available), place of publication (where available).				T ⁵	E ⁶

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16292374

		<input type="checkbox"/>	<input type="checkbox"/>
--	--	--------------------------	--------------------------

STATEMENTS

The party making the submission is not an individual who has a duty to disclose information with respect to the above-identified application under 37 CFR 1.56.

This submission complies with the requirements of 35 U.S.C. 122(e) and 37 CFR 1.290.

The fee set forth in 37 CFR 1.290(f) has been submitted herewith.

The fee set forth in 37 CFR 1.290(f) is not required because this submission lists three or fewer total items and, to the knowledge of the person signing the statement after making reasonable inquiry, this submission is the first and the only submission under 35 U.S.C. 122(e) filed in the above-identified application by the party making the submission or by a party in privity with the party.

This resubmission is being made responsive to a notification of non-compliance issued for an earlier filed third-party submission. The corrections in this resubmission are limited to addressing the non-compliance. As such, the party making this resubmission: (1) requests that the Office apply the previously-paid fee set forth in 37 CFR 1.290(f), or (2) states that no fee is required to accompany this resubmission as the undersigned is again making the fee exemption statement set forth in 37 CFR 1.290(g).

Signature	/Jurate Breimelyte/		
Name/Print	Jurate Breimelyte	Registration Number (if applicable)	

Examiner Signature		Date Considered	
--------------------	--	-----------------	--

***EXAMINER:** Signature indicates all documents listed above have been considered, except for citations through which a line is drawn. Draw line through citation if not considered. Include a copy of this form with next communication to applicant. 1. If known, enter kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16. See MPEP 901.04(a). 2. Enter the country or patent office that issued the document, by two-letter code under WIPO standard ST.3. See MPEP 1851. 3. For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. 4. If known, enter the kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16. See MPEP 901.04(a). 5. Check mark indicates translation attached. 6. Check mark indicates evidence of publication attached.

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290 CONCISE DESCRIPTION OF RELEVANCE		
Application Number	16292382	
U.S. PATENTS		
Cite No	Patent Number	Concise Description of Relevance
U.S. PATENT APPLICATION PUBLICATION		
Cite No	Publication Number	Concise Description of Relevance
1	20090248793	Jacobsson discloses that other metadata may be associated with content and content sources. The various content sources are identified by public or private IP address (par 35). Also discloses that requests for content can be sent to any of the content sources (client devices and/or servers)(par 53). It discloses that another client device can receive content from other client devices and/or web servers. For example, if the first client device receives the second content, that first client device can become a content source for yet further client devices. Par 18 and FIG. 1 discloses a computer system for delivering content over a network. In one embodiment, the system can integrate with an existing web infrastructure using an existing communications protocol such as HTTP to transfer digital content over the Internet. Many types of content can be delivered, including, audio content, image content, video content, and application program content.

2	20070142036	<p>FIG. 1 illustrates an environment in which a remote server can be accessed via an intermediate server. The electronic devices are located remote to the electronic client device 110. All electronic devices have communication capabilities for accessing the Internet. Electronic devices can act as Internet web servers, and thus themselves also manage content for use on the Internet(par 32). The electronic devices are capable of transmitting and receiving data by packet switching using a standardized IP and possibly also other protocols. The electronic devices each comprise mobile Internet web servers, each of which can be independently remotely accessed directly by using a URL, globally unique for each handset(par 35). The invention relates to a network system in which the handset can be used. The network system comprises at least an intermediate device, remote origin device in its proximity, and client device used to access the content on the remote device via intermediate device(par 74).</p>
---	-------------	---

FOREIGN PATENT DOCUMENTS

CiteNo	Foreign Document Number	Concise Description of Relevance
1	2004094980	<p>Fontijn discloses that when any device (e.g., 12, 13, 14) supplies content, that content came from the Server (8:16-19). It is commonly known in the art and disclosed in Fontijn, that devices in P2P networks request content, and as such, it would be understood that the devices of Fontijn requested the content, and such requests would include an identification of the content and a destination identifier for the request (4:14-18). Fontijn discloses that a second device (in the present invention - selected IP address) is selected from the group (Fig. 1 (devices 13-17); 4:32-5:3). Discloses that it is known that users can select a device (2:12-16), and that the Server can select the device (4:21-22). Claim 1 discloses receiving and transmitting, from a first device (11), a first request with a first selection criterion for a first content to a server (18); transferring (200) the first content satisfying said first selection criterion to said first device from the server.</p>

NON-PATENT PUBLICATIONS

Cite No	Reference	Concise Description of Relevance

--	--	--

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16292382

U.S. PATENTS

Cite No	Patent Number	Kind Code ¹	Issue Date (YYYY-MM-DD)	First Named Inventor

U.S. PATENT APPLICATION PUBLICATIONS

Cite No	Publication Number	Kind Code ¹	Publication Date (YYYY-MM-DD)	First Named Inventor
1	20070142036	A1	2007-06-21	Johan Wikman
2	20090248793	A1	2009-10-01	Sanny Jacobsson

FOREIGN PATENTS AND PUBLISHED FOREIGN PATENT APPLICATIONS

Cite No	Foreign Document Number ³	Country Code ²	Kind Code ¹	Publication Date (YYYY-MM-DD)	Applicant, Patentee or First Named Inventor	T ⁵
1	2004094980	WO	A2	2004-11-04	Wilhelmus F. J. Fontijn	<input type="checkbox"/>

NON-PATENT PUBLICATIONS (e.g., journal article, Office action)

Cite No	Author (if any), title of the publication, page(s) being submitted, publication date, publisher (where available), place of publication (where available).	T ⁵	E ⁶

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16292382

		<input type="checkbox"/>	<input type="checkbox"/>
--	--	--------------------------	--------------------------

STATEMENTS

The party making the submission is not an individual who has a duty to disclose information with respect to the above-identified application under 37 CFR 1.56.

This submission complies with the requirements of 35 U.S.C. 122(e) and 37 CFR 1.290.

The fee set forth in 37 CFR 1.290(f) has been submitted herewith.

The fee set forth in 37 CFR 1.290(f) is not required because this submission lists three or fewer total items and, to the knowledge of the person signing the statement after making reasonable inquiry, this submission is the first and the only submission under 35 U.S.C. 122(e) filed in the above-identified application by the party making the submission or by a party in privity with the party.

This resubmission is being made responsive to a notification of non-compliance issued for an earlier filed third-party submission. The corrections in this resubmission are limited to addressing the non-compliance. As such, the party making this resubmission: (1) requests that the Office apply the previously-paid fee set forth in 37 CFR 1.290(f), or (2) states that no fee is required to accompany this resubmission as the undersigned is again making the fee exemption statement set forth in 37 CFR 1.290(g).

Signature	/Jurate Breimelyte/		
Name/Print	Jurate Breimelyte	Registration Number (if applicable)	

Examiner Signature		Date Considered	
--------------------	--	-----------------	--

***EXAMINER:** Signature indicates all documents listed above have been considered, except for citations through which a line is drawn. Draw line through citation if not considered. Include a copy of this form with next communication to applicant. 1. If known, enter kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16. See MPEP 901.04(a). 2. Enter the country or patent office that issued the document, by two-letter code under WIPO standard ST.3. See MPEP 1851. 3. For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. 4. If known, enter the kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16. See MPEP 901.04(a). 5. Check mark indicates translation attached. 6. Check mark indicates evidence of publication attached.

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290 CONCISE DESCRIPTION OF RELEVANCE		
Application Number	16365250	
U.S. PATENTS		
Cite No	Patent Number	Concise Description of Relevance
U.S. PATENT APPLICATION PUBLICATION		
Cite No	Publication Number	Concise Description of Relevance
1	20090248793	In Jacobsson, para. 64 discloses the relationship of client and server by virtue of computer programs running on the respective computers and having a client-server relationship to each other. Claim 9 similarly to the present invention discloses a method for providing content in a network: receiving, at the device and from the first server device, source information that identifies a selected content and obtaining the content using the received source information. Claim 16 further discloses sending a confirmation from the device to the first server device, the confirmation including information about a result of the requests generated by the device. Jacobsson discloses the selection of the content source from the list of content sources (paras 5-6, 35-36). This information is sent from the device to the server. It is known in the art, that different appliances connected to the Internet can be used as an intermediate devices (proxies).

2	20130080575	Prince discloses a method for fetching a first content over the Internet identified by a first content identifier by disclosing fetching network resources from origin servers, where the content may include HTML pages, images, or any computer files (paras 26-27; 15, 17-19). Further discloses that the content is identified in requests from client, for instance in HTTP requests sent to proxy server. (Paras 45 (client request includes parameters such as the type of requested content), 46, 27, 39; Fig. 1). Discloses sending a first request to the first server by disclosing that client sends a request to DNS system to provide a proxy server that will retrieve content from original server. (Paras 15, 77 (client device requests an IPv4 address for example.com (thus, the client device 110A is an IPv4 enabled client), 90; Figs. 5 (step 510), 6 (step 610)).
3	20070142036	Wikman discloses a mobile intermediate device for providing content from a geographically proximate origin device to a geographically remote client device, the mobile intermediate device having a mobile intermediate device address to which content requests are communicated from the client device, and mobile intermediate device comprising an input and an output, the output arranged to provide content to a remote client device upon receipt of a content request at the input from the remote client device. The mobile intermediate device may be arranged to function as a reverse proxy server (0020-0021; 0067). Different appliances can be understood as mobile devices and when connected to the Internet can perform a function of an exit node.

FOREIGN PATENT DOCUMENTS

CiteNo	Foreign Document Number	Concise Description of Relevance

NON-PATENT PUBLICATIONS		
Cite No	Reference	Concise Description of Relevance

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16365250

U.S. PATENTS

Cite No	Patent Number	Kind Code ¹	Issue Date (YYYY-MM-DD)	First Named Inventor

U.S. PATENT APPLICATION PUBLICATIONS

Cite No	Publication Number	Kind Code ¹	Publication Date (YYYY-MM-DD)	First Named Inventor
1	20070142036	A1	2007-06-21	Johan Wikman
2	20130080575	A1	2013-03-28	Matthew Browning Prince
3	20090248793	A1	2009-10-01	Sanny Jacobsson

FOREIGN PATENTS AND PUBLISHED FOREIGN PATENT APPLICATIONS

Cite No	Foreign Document Number ³	Country Code ²	Kind Code ¹	Publication Date (YYYY-MM-DD)	Applicant, Patentee or First Named Inventor	T ⁵
						<input type="checkbox"/>

NON-PATENT PUBLICATIONS (e.g., journal article, Office action)

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16365250

Cite No	Author (if any), title of the publication, page(s) being submitted, publication date, publisher (where available), place of publication (where available).	T ⁵	E ⁶
		<input type="checkbox"/>	<input type="checkbox"/>

STATEMENTS

The party making the submission is not an individual who has a duty to disclose information with respect to the above-identified application under 37 CFR 1.56.

This submission complies with the requirements of 35 U.S.C. 122(e) and 37 CFR 1.290.

The fee set forth in 37 CFR 1.290(f) has been submitted herewith.

The fee set forth in 37 CFR 1.290(f) is not required because this submission lists three or fewer total items and, to the knowledge of the person signing the statement after making reasonable inquiry, this submission is the first and the only submission under 35 U.S.C. 122(e) filed in the above-identified application by the party making the submission or by a party in privity with the party.

This resubmission is being made responsive to a notification of non-compliance issued for an earlier filed third-party submission. The corrections in this resubmission are limited to addressing the non-compliance. As such, the party making this resubmission: (1) requests that the Office apply the previously-paid fee set forth in 37 CFR 1.290(f), or (2) states that no fee is required to accompany this resubmission as the undersigned is again making the fee exemption statement set forth in 37 CFR 1.290(g).

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290 CONCISE DESCRIPTION OF RELEVANCE		
Application Number		16365315
U.S. PATENTS		
Cite No	Patent Number	Concise Description of Relevance
U.S. PATENT APPLICATION PUBLICATION		
Cite No	Publication Number	Concise Description of Relevance
1	20090248793	<p>Identically to the present invention, the prior art describes a method for providing content in a network. It also discloses receiving the identifier from a device, selecting a content source in the network, forwarding, in response to the identifier, source information to the device (Claim 1). Par 6 and Claim 3 of the prior art discloses that the optimization algorithm can be configured so that at least one content source is selected according to at least one approach, for example the content source being located in a geographic location in which the device is located. Par 36 and Claim 4 discloses that in some implementations of the invention, preference can be given to the selection of content sources associated with the same geographic location (e.g., country), network, and/or internet service provider as the client device. The content can include at least one content type selected from: audio content, image content, Video content, application program content, and combinations (par6).</p>

2	20070142036	<p>Prior art discloses a method to access the content (e.g. website, or other content) contained on an origin device via the intermediate device by a client device (0019;0046-56). Wikman discloses a method for providing content to a remotely located electronic device, connectable to the Internet, by accessing content on a device (e.g. an Internet server) located near a mobile intermediate device by using the mobile intermediate device (another Internet server) as a 'through conduit'. Discloses a mobile intermediate device for providing content from a geographically proximate origin device to a geographically remote client device, the mobile intermediate device having a mobile intermediate device address to which content requests are communicated from the client device, and mobile intermediate device comprising an input and an output, the output arranged to provide content to a remote client device upon receipt of a content request at the input from the remote client device (0020-21).</p>
---	-------------	--

FOREIGN PATENT DOCUMENTS

CiteNo	Foreign Document Number	Concise Description of Relevance
1	2004094980	<p>The content is identified by specific selection criteria entered by user and transmitted to a server or device storing the content. See 4:10-19. Fontijn discloses a group consisting of a plurality of P2P devices 12, 13, 14, 15, 16, and 17. See Fig. 1; 4:32-5:3. As in the recited application, in Fontijn the devices on the network are identified by IP address. See 2:12-16. Fontijn further discloses selecting the second device out of the devices in the group. See 2:5-16. See Fig. 1; 4:9-29. Fontijn discloses that it is known that users can select a device, see, e.g., 2:12-16, and that the Server can select the device when it redirects the first device. See, e.g., 4:21-22.</p>

NON-PATENT PUBLICATIONS

Cite No	Reference	Concise Description of Relevance

--	--	--

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16365315

U.S. PATENTS						
Cite No	Patent Number	Kind Code ¹	Issue Date (YYYY-MM-DD)	First Named Inventor		
U.S. PATENT APPLICATION PUBLICATIONS						
Cite No	Publication Number	Kind Code ¹	Publication Date (YYYY-MM-DD)	First Named Inventor		
1	20070142036	A1	2007-06-21	Johan Wikman		
2	20090248793	A1	2009-10-01	Sanny Jacobsson		
FOREIGN PATENTS AND PUBLISHED FOREIGN PATENT APPLICATIONS						
Cite No	Foreign Document Number ³	Country Code ²	Kind Code ¹	Publication Date (YYYY-MM-DD)	Applicant, Patentee or First Named Inventor	T ⁵
1	2004094980	WO	A2	2004-11-04	Wilhelmus F. J. Fontijn	<input type="checkbox"/>
NON-PATENT PUBLICATIONS (e.g., journal article, Office action)						
Cite No	Author (if any), title of the publication, page(s) being submitted, publication date, publisher (where available), place of publication (where available).				T ⁵	E ⁶

THIRD-PARTY SUBMISSION UNDER 37 CFR 1.290	Application Number	16365315

		<input type="checkbox"/>	<input type="checkbox"/>

STATEMENTS

The party making the submission is not an individual who has a duty to disclose information with respect to the above-identified application under 37 CFR 1.56.

This submission complies with the requirements of 35 U.S.C. 122(e) and 37 CFR 1.290.

The fee set forth in 37 CFR 1.290(f) has been submitted herewith.

The fee set forth in 37 CFR 1.290(f) is not required because this submission lists three or fewer total items and, to the knowledge of the person signing the statement after making reasonable inquiry, this submission is the first and the only submission under 35 U.S.C 122(e) filed in the above-identified application by the party making the submission or by a party in privity with the party.

This resubmission is being made responsive to a notification of non-compliance issued for an earlier filed third-party submission. The corrections in this resubmission are limited to addressing the non-compliance. As such, the party making this resubmission: (1) requests that the Office apply the previously-paid fee set forth in 37 CFR 1.290(f), or (2) states that no fee is required to accompany this resubmission as the undersigned is again making the fee exemption statement set forth in 37 CFR 1.290(g).

Signature	/Jurate Breimelyte/
-----------	---------------------

Name/Print	Jurate Breimelyte	Registration Number (if applicable)	
------------	-------------------	--	--

Examiner Signature		Date Considered	
--------------------	--	-----------------	--

*EXAMINER: Signature indicates all documents listed above have been considered, except for citations through which a line is drawn. Draw line through citation if not considered. Include a copy of this form with next communication to applicant. 1. If known, enter kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16. See MPEP 901.04(a). 2. Enter the country or patent office that issued the document, by two-letter code under WIPO standard ST.3. See MPEP 1851. 3. For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. 4. If known, enter the kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16. See MPEP 901.04(a). 5. Check mark indicates translation attached. 6. Check mark indicates evidence of publication attached.

Slice Embedding Solutions for Distributed Service Architectures

Flavio Esposito Ibrahim Matta Vatche Ishakian
 flavio@cs.bu.edu matta@cs.bu.edu visahak@cs.bu.edu

Computer Science Department
 Boston University
 Boston, MA

Technical Report BUCS-TR-2011-025

Abstract—Network virtualization provides a novel approach to run multiple concurrent virtual networks over a common physical network infrastructure. From a research perspective, this enables the networking community to concurrently experiment with new Internet architectures and protocols. From a market perspective, on the other hand, this paradigm is appealing as it enables infrastructure service providers to experiment with new business models that range from leasing virtual slices of their infrastructure to host multiple concurrent network services.

In this paper, we present the slice embedding problem and recent developments in the area. A slice is a set of virtual instances spanning a set of physical resources. The embedding problem consists of three main tasks: (1) resource discovery, which involves monitoring the state of the physical resources, (2) virtual network mapping, which involves matching users' requests with the available resources, and (3) allocation, which involves assigning the resources that match the users' query.

We also outline how these three tasks are tightly connected, and how there exists a wide spectrum of solutions that either solve a particular task, or jointly solve multiple tasks along with the interactions among them. To dissect the space of solutions, we introduce three main classification criteria, namely, (1) the type of constraints imposed by the user, (2) the type of dynamics considered in the embedding process, and (3) the allocation strategy adopted. Finally, we conclude with a few interesting research directions.

I. INTRODUCTION

We all became familiar with the layered reference model of ISO OSI as well as the layered TCP/IP architecture [47]. In these models, a layer is said to provide a *service* to the layer immediately above it. For example, the transport layer provides services (logical end-to-end channels) to the application layer, and the internetworking layer provides services (packet delivery across individual networks) to the transport layer.

The notion of distributed service architecture extends this service paradigm to many other (large scale) distributed systems.

Aside from the Internet itself, including its future architecture design, *e.g.*, NetServ [73] or RINA [23], with the term *distributed service architecture* we refer to a large scale distributed system whose architecture is based on a service paradigm.

Some examples are datacenter-based systems [39], Cloud Computing [36] (including high performance computing systems such as cluster-on-demand services), where the rentable resources can scale both up and down as needed, Grid Computing [45], overlay networks (*e.g.*, content delivery networks [6],

[10]), large scale distributed testbed platforms (*e.g.*, PlanetLab [65], Emulab/Netbed [77], VINI [7], GENI [31]), or Service-oriented Architecture (SoA), where web applications are the result of the composition of services that need to be instantiated across a collection of distributed resources [80].

A common characteristic of all the above distributed systems is that they all provide a service to a set of users or, recursively, to another service. In this survey, we restrict our focus on a particular type of service: a slice. We define a slice to be a set of virtual instances spanning a set of physical resources.

The lifetime span of a slice ranges from few seconds (in the case of cluster-on-demand services) to several years (in case of a virtual network hosting a content distribution service similar to Akamai, or even a GENI experiment hosting a novel architecture looking for new adopters to opt-in [34]). Therefore, the methods to acquire, configure, manipulate and manage such slices could be different across different service architectures. In particular, the problem of discovering, mapping and allocating physical resources (slice embedding) has different time constraints in each service architecture.¹

In some distributed service architecture applications, *e.g.* virtual network testbed, the slice creation and embedding time is negligible relative to the running time of the service they are providing. In many other applications, *e.g.* financial modeling, anomaly analysis, or heavy image processing, the time to solution — instant between the user, application or service requests a slice and the time of task completion — is dominated by or highly dependent on the slice creation and embedding time.

Therefore, to be profitable, most of those service architectures require agility—the ability to allocate and deallocate any physical resource (node or link) to any service at any time². Those stringent requirements, combined with the imperfect design of today's data center networks [35] and with the lack of an ideal virtualization technology [78], have recently motivated research on resource allocation [13], [82], [51], [35], [4], [70].

In this paper, we define the slice embedding problem—a

¹By resources we mean processes, storage capacity, and physical links, as well as computational resources such as processors.

²We extend the definition of agility as “ability to assign any server to any service” given by Greenberg *et al.* [35] by including links and, other resources along with a deallocation phase.

subarea of the resource allocation for service architectures—in Section II, we give a taxonomy (Section III), and we survey some of the recent solutions for each of its tasks (Sections IV, V and VI). Then, with the help of optimization theory, we model the three phases of the slice embedding problem as well as its tasks’ interactions (Section VIII). We point out how all the proposed approaches—including the related facility location problems (Section VII)—have considered either cases where the time to solution is practically equivalent to the running time of a slice, *i.e.* they did not consider the slice creation and embedding time at all, or they did not model some of the slice embedding tasks. In Section IX we discuss some interesting open research directions and finally, in Section X we conclude our discussion.

II. BACKGROUND AND AREA DEFINITION

A. Network Virtualization

Network virtualization provides a novel approach to running multiple concurrent virtual networks over a common physical network infrastructure. A physical network supports virtualization if it allows the coexistence of multiple virtual networks. Each virtual network is a collection of virtual nodes and virtual links that connect a subset of the underlying physical network resources. The most important characteristic of such virtual networks is that they are customizable (*i.e.*, can concurrently run different protocols or architectures, each tailored to a particular service or application [75]).

The interest in this technology has recently grown significantly because it will help the research community in the testing of novel protocols and algorithms in pseudo-real network environments [65], [77], [7], [28], as well as experimenting with novel Internet architectures as envisioned in [3]. This paradigm is particularly appealing to providers as it enables new business models: operators may in fact benefit from diversifying their infrastructure by leasing virtual networks to a set of customers [30], or by sharing costs in deploying a common infrastructure [11].

A recent survey on network virtualization can be found in [18]. The authors compare with a broad perspective, approaches related to network virtualization, *e.g.* virtual private networks and overlay networks. The paper also discusses economic aspects of service providers, analyzes their design goals (such as manageability or scalability), and overviews recent projects that use this technology (*e.g.* Planetlab [65] and GENI [31]). We narrow our focus on a more specific subarea of network virtualization (*i.e.* slice embedding), introducing a new taxonomy inspired by optimization theory for the three phases of the slice embedding problem. We leave our utility functions and model constraints as general as possible, so they can be instantiated, refined or augmented based on policies that would lead to efficient slice embedding solutions.

B. The Slice Embedding Problem

In this paper, we focus on a particular aspect of network virtualization, namely, the slice embedding problem.

A slice is defined as a set of virtual instances spanning a set of physical resources of the network infrastructure. The

slice embedding problem comprises the following three steps: resource discovery, virtual network mapping, and allocation.

Resource discovery is the process of monitoring the state of the substrate (physical) resources using sensors and other measurement processes. The monitored states include processor loads, memory usage, network performance data, etc. We discuss the resource discovery problem in Section IV.

Virtual network mapping is the step that matches users’ requests with the available resources, and selects some subset of the resources that can potentially host the slice. Due to the combination of node and link constraints, this is by far the most complex step in the slice embedding problem. In fact this problem is NP-hard [19]. These constraints include intra-node (*e.g.*, desired physical location, processor speed, storage capacity, type of network connectivity), as well as inter-node constraints (*e.g.*, network topology). We define the virtual network mapping problem in Section V.

Allocation involves assigning the resources that match the user’s query to the appropriate slice. The allocation step can be a single shot process, or it can be repeated periodically to either reassign or to acquire additional resources for a slice that has already been embedded.

C. Interactions in the Slice Embedding Problem

Before presenting existing solutions to the tasks encompassing the slice embedding problem, it is important to highlight the existence of interactions among these tasks, the nature of these interactions, how they impact performance, as well as the open issues in addressing these interactions.

In Figure 1, a user is requesting a set of resources. The arrow (1) going from the “Requests” to the “Discovery” block, represents user queries that could potentially have multiple levels of expressiveness and a variety of constraints. The resource discoverer (2) returns a subset of the available resources (3) to the principle in charge of running the virtual network mapping algorithm (4). Subsequently, the slice embedding proceeds with the allocation task. A list of candidate mappings (5) are passed to the allocator (6), that decides which physical resources are going to be assigned to each user. The allocator then communicates the list of winners (7)—users that won the allocation—to the discoverer, so that future discovery operations can take into account resources that have already been allocated. It is important to note that the slice embedding problem is essentially a closed feedback system, where the three tasks are solved repeatedly—the solution in any given iteration affects the space of feasible solutions in the next iteration.

D. Solutions to the Slice Embedding Problem

Solutions in the current literature either solve a specific task of the slice embedding problem, or are hybrids of two tasks. Some solutions jointly consider resource discovery and network mapping [41], [1], others only focus on the mapping phase [81], [54], [21], or on the interaction between virtual network mapping and allocation [79], [52], while others consider solely the allocation step [5], [9], [49], [33], [20]. Moreover, there are solutions that assume the virtual network mapping

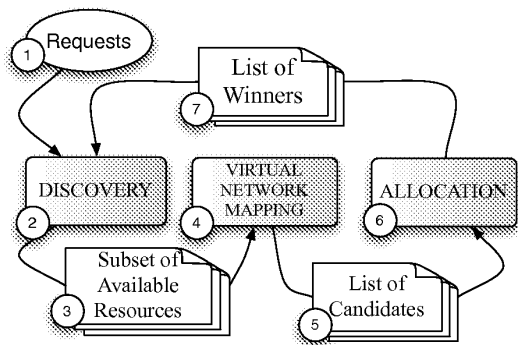


Fig. 1. Interactions and data exchanges in the slice embedding problem.

task is solved, and only consider the interaction between the resource discovery and allocation [68]. We do not discuss solutions that address the resource discovery task in isolation, since it is not different from classical resource discovery in the distributed system literature (see [60] for an excellent survey on the topic). In addition to considering one [81], [5] or more [62], [79] tasks, solutions also depend on whether their objective is to maximize users' or the providers' utility.

E. The novelty of the slice embedding problem

The slice embedding problem, or more specifically its constituent tasks, and network virtualization in general, may seem identical to problems in classical distributed systems. Network virtualization, however, is different in several ways, namely: (a) it enables novel business models, (b) it enables novel coexisting network approaches, and (c) it creates new embedding challenges that must be addressed.

Business models: network virtualization lays the foundations for new business models [22]. Network resources are now considered commodities to be leased on demand. The leaser could be an infrastructure or service provider, and the lessee could be another service provider, an enterprise, or a single user (e.g. a researcher in the case of virtual network testbed as in [31], [7], [38], [65], [28]). In those cases where the infrastructure is a public virtualizable network testbed (e.g. GENI [31]), the physical resources may not have any significant market value, since they are made available at almost no cost to research institutions.

Coexisting network approaches: the concept of multiple coexisting logical networks appeared in the networking literature several times in the past. The most closely related attempts are virtual private networks and overlay networks. A virtual private network (VPN) is a dedicated network connecting multiple sites using private and secured tunnels over a shared communication network. Most of the time, VPNs are used to connect geographically distributed sites of a single enterprise: each VPN site contains one or more customer edge devices attached to one or more provider edge routers [66].

An overlay network, on the other hand, is a logical network built on top of one or more existing physical networks. One substantial difference between overlays and network virtualization is that overlays in the existing Internet are typically implemented at the application layer, and therefore they may have limited applicability.

For example, they falter as a deployment path for radical architectural innovations in at least two ways: first, overlays have largely been in use as means to deploy narrow fixes to specific problems without any holistic view; second, most overlays have been designed in the application layer on top of the IP protocol, hence, they cannot go beyond the inherent limitations of the existing Internet [3].

In the case of VPNs, the virtualization level is limited to the physical network layer while in the case of overlays, virtualization is limited to the end hosts. Network virtualization introduces the ability to access, manage and control each layer of the current Internet architecture in the end hosts, as well as providing dedicated virtual networks.

Embedding challenges: although the research community has explored the embedding of VPNs in a shared provider topology, e.g., [26], usually VPNs have standard topologies, such as a full mesh. A virtual network in the slice embedding problem, however, may represent any topology. Moreover, resource constraints in a VPN or overlays are limited to either bandwidth requirements or node constraints, while in network virtualization, both link and node constraints may need to be present simultaneously. Thus, the slice embedding problem differs from the standard VPN embedding because it must deal with both node and link constraints for arbitrary topologies.

III. TAXONOMY

To dissect the space of existing solutions spanning the slice embedding tasks, as well as interactions among them, we consider three dimensions as shown in Figure 2: the *type of constraint*, the *type of dynamics*, and the *resource allocation approach*.

A. Constraint type

Users need to express their queries efficiently. Some constraints are on the nodes and/or links (e.g., minimum CPU requirement, average bandwidth, maximum allowed latency) while others consider inter-group [1] or geo-location constraints [17].

Based on this dimension, research work in this area assumes no constraints [81], considers constraints on nodes only [65], links only [55], [67], [37], or on both nodes and links [5], [79]. In addition, the order in which the constraints are satisfied is important as pointed out in [52]: satisfy the node constraints and then the link constraints [81], [79], or satisfy both constraints simultaneously [54], [52].

B. Dynamics

Each task in the slice embedding problem may differ in terms of its dynamics. In the resource discovery task, the

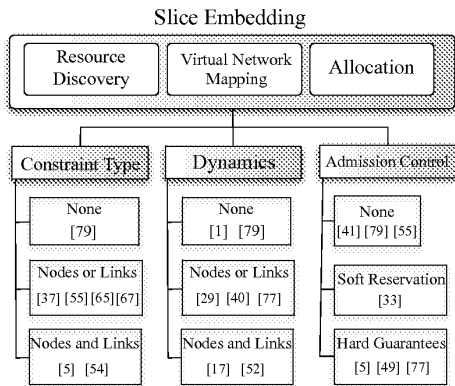


Fig. 2. Overview of the slice embedding taxonomy with classification of representative references.

status updates of each physical resource may be collected periodically [41], or on demand [1].

In the virtual network mapping task, virtual resources may be statically mapped to each physical resource [81], or they can move (*e.g.*, using path migrations [79] or by re-running the mapping algorithm [29]) to maximize some notion of utility [37]. Also, the mapping can focus only on one single phase at a time where each phase considers only nodes or links [81], [40], or simultaneously both nodes and links [52], [17].

Finally, the allocation task may be dynamic as well: users may be swapped in or out to achieve some Quality of Service (QoS) or Service Level Agreement (SLA) performance guarantees, or they can statically remain assigned to the same slice. An example of static assignment of a slice may be an infrastructure hosting a content distribution service similar to Akamai, whereas an example of dynamic reallocation could be a researcher's experiment being swapped out from/into the Emulab testbed [77].

C. Admission Control

As the substrate—physical infrastructure—resources are limited, some requests must be rejected or postponed to avoid violating the resource guarantees for existing virtual networks, or to maximize profit of the leased network resources. Some research work, however, does not consider any resource allocation [41], [54], [21], [81], [55], [52]. Others consider the resource allocation task, with [33] or without [49], [5], [79] guarantees to the user, *i.e.*, the resource allocation mechanism enforces admission to the users, or it only implements a tentative admission, respectively. An example of tentative admission is a system that issues tickets, without guarantee that those tickets can be exchanged with a resource later in time. The literature defines those tentative admission mechanisms that do not provide hard guarantees as *soft reservation* [33].

IV. RESOURCE DISCOVERY

Although researchers have developed, and in some cases deployed a number of resource discovery solutions for wide-

area distributed systems, the research in this area still has many open problems. Some of the existing distributed systems provide resource discovery through a centralized architecture, see, *e.g.*, Condor [53], Assign [67], or Network Sensitive Service Discovery (NSSD) [41]; others use a hierarchical architecture such as Ganglia [58], while XenosSearch [72], SWORD [62] and iPlane Nano [57] employ a decentralized architecture.

All of these systems allow users to find nodes that meet per-node constraints, except iPlane Nano that considers path metrics, while NSSD, SWORD, and Assign also consider network topologies. Unfortunately, none of these solutions analyze the resource discovery problem when the queried resources belong to multiple infrastructure or service providers. To obtain an efficient slice embedding, such cases would in fact require some level of cooperation (*e.g.*, by sharing some state), and such incentives to cooperate may be scarce.

As mentioned previously, we do not discuss solutions that address the resource discovery task in isolation, since it is not different from classical resource discovery in the distributed systems literature. Instead, we consider the resource discovery problem in combination with either the allocation or the network mapping task.

A. Discovery + allocation

We first discuss the interaction between discovery and allocation described in Network Sensitive Service Discovery (NSSD) [41]. The goal is to discover a service that meets a set of network properties specified by the user, and allocate it to the user.

This work emphasizes the importance of the interaction between discovery of network resources and their allocation to the users. The resource discovery task infers the network's performance metrics during its search and returns the best match with respect to some user criteria. In general, once a user's query is received, in existing systems either the provider (pure provider-side allocation) or the users (pure user-side allocation) execute the allocation task. If the allocation is done by the provider, users do not have to worry about anything after they submit a query, but may not know the quality of service they are going to get (in systems like PlanetLab for example, there are no service level agreements that the provider needs to meet). On the other hand, when the allocation is done by the user, each user needs to obtain a long list of candidates, as well as collect the status information of each candidate. Thus, the overhead of the discovery task is higher if users need to have the ability to choose the best set of resources. When the provider does the allocation instead, there may be no need to look at the complete set of resources as some heuristic (*e.g.* first fit) can be applied. Moreover, by showing the most available physical resources they own, providers could (indirectly) have to release information about their states, *e.g.*, information about which customer is hosted on a physical machine could be inferred [69].

To the best of our knowledge, NSSD is the first system that integrates the discovery and allocation tasks while enabling users to query static and dynamic network properties. Compared with pure provider-side allocation, NSSD allows users to

control the selection criteria by returning a list of candidates. Compared with pure user-side allocation, NSSD has lower overhead in the discovery task, as only a small number of candidates are returned. In this work, the resources to allocate are single servers, hence there is no virtual network mapping phase.

B. Discovery + virtual network mapping

We present SWORD [1], a system that considers the interaction between the resource discovery and the virtual network mapping tasks. SWORD is a resource discovery infrastructure for shared wide-area platforms such as PlanetLab [65]. We choose to describe SWORD as it is a well known network discovery system whose source code is available [74]. The system has been running on PlanetLab for several years. Some of the functionalities described in the original paper, however, are currently disabled. For example, the current implementation of SWORD runs in centralized mode, and inter-node and group requirements (*i.e.*, constraints on links and set of nodes, respectively), are not supported because no latency or bandwidth estimates are available.

Users wishing to find nodes for their application submit a resource request expressed as a topology of interconnected groups. A group is an equivalence class of nodes with the same per-node requirements (*e.g.*, free physical memory) and the same inter-node requirements (*e.g.*, inter-node latency) that is within each group. Supported topological constraints within and among groups include the required bandwidth and latency.

In addition to specifying absolute requirements, users can supply SWORD with per-attribute *penalty functions*, that map the value of an attribute (feature of a resource, such as load or delay) within the required range but outside an ideal range, to an abstract penalty value. This capability allows SWORD to rank the quality of the configurations that meet the applications' requirements, according to the relative importance of each attribute. Notice that these penalty values would be passed to the allocation together with the list of candidates.

Architecturally, SWORD consists of a distributed query processor and an *optimizer* which can be viewed as a virtual network mapper. The distributed query processor uses multi-attribute range search built on top of a peer-to-peer network to retrieve the names and attribute values of the nodes that meet the requirements specified in the user's query. SWORD's optimizer then attempts to find the lowest-penalty assignment of platform nodes (that were retrieved by the distributed query processor) to groups in the user's query—that is, the lowest-penalty embedding of the requested topology in the PlanetLab node topology, where the penalty of an embedding is defined as the sum of the per-node, inter-node, and inter-group penalties associated with that selection of nodes.

Due to the interaction between the distributed query processor (resource discovery task) and the optimizer (mapping task), SWORD is more than a pure resource discoverer. SWORD provides resource discovery, solves the network mapping task, but does not provide resource allocation. In particular, since PlanetLab does not currently support resource guarantees, a set of resources that SWORD returns to a user may no longer

meet the resource request at some future point in time. In light of this fact, SWORD supports a *continuous query* mechanism where a user's resource request is continually re-matched to the characteristics of the available resources, and in turn a new set of nodes are returned to the user. The user can then choose to migrate one or more instances of their application. This process is all part of the general feedback system outlined in Figure 1.

V. VIRTUAL NETWORK MAPPING

The virtual network mapping is the central phase of the slice embedding problem. In this section we define the problem of virtual network mapping, then we survey solutions that focus only on this phase, as well as solutions that cover interactions with the other two tasks of the slice embedding problem.

A. Problem definition

The virtual network mapping problem is defined as follows [52]:

Definition 1 (Network): A Network is defined as an undirected graph $G = (N, L, C)$ where N is a set of nodes, L is a set of links, and each node or link $e \in N \cup L$ is associated with a set of constraints $C(e) = \{C_1(e), \dots, C_m(e)\}$. A physical network will be denoted as $G^P = (N^P, L^P, C^P)$, while a virtual network will be denoted as $G^V = (N^V, L^V, C^V)$.

Definition 2 (Virtual Network Mapping): Given a virtual network $G^V = (N^V, L^V, C^V)$ and a physical network $G^P = (N^P, L^P, C^P)$, a virtual network mapping is a mapping of G^V to a subset of G^P , such that each virtual node is mapped onto exactly one physical node, and each virtual link is mapped onto a loop-free path p in the physical network. The mapping is called valid if all the constraints $C(e)$ of the virtual network are satisfied and do not violate the constraints of the physical network. More formally, the mapping is a function

$$M : G^V \rightarrow (N^P, \mathcal{P}) \quad (1)$$

where \mathcal{P} denotes the set of all loop-free paths in G^P . M is called a *valid mapping* if all constraints³ of G^V are satisfied, and for each $l^v = (s^V, t^V) \in L^V$, \exists a path $p : (s^P, \dots, t^P) \in \mathcal{P}$ where s^V is mapped to s^P and t^V is mapped to t^P .

Due to the combination of node and link constraints, the virtual network mapping problem is NP-hard. For example, assigning virtual nodes to the substrate (physical) network without violating link bandwidth constraints can be reduced to the multiway separator problem which is NP-hard [2].

To reduce the overall complexity, several heuristics were introduced, including backtracking algorithms [54], [52], simulated annealing as in Emulab [67], as well as heuristics that solve the node and link mapping independently.

³Examples of node constraints include CPU, memory, physical location, whereas link constraints may be delay, jitter, or bandwidth.

TABLE OF NOTATIONS		
Symbol	Page	Meaning
G	6	Undirected graph representing a general network
N	6	General set of nodes (or vertices) of a network
L	6	General set of links (or edges) of a network
C	6	General set of network constraints
$C^P (C^V)$	6	General set of physical (virtual) network constraints
$C(e) = \{C_1(e), \dots, C_m(e)\}$	6	Set of m constraints on the element e (node or link) of the network
$G^P (G^V)$	6	Undirected graph representing a physical (virtual) network
$N^P (N^V)$	6	Set of nodes or vertices of a physical (virtual) network
$L^P (L^V)$	6	Set of links or edges of a physical (virtual) network
\mathcal{P}	6	Set of loop-free physical paths in a physical network G^P
$l^v = (s^V, t^V)$	6	Virtual link starting from virtual node s^V , and ending in virtual node t^V
$p : (s^P, \dots, t^P) \in \mathcal{P}$	6	Physical path starting from physical node s^P , and ending in physical node t^P
M	6	Mapping function: $G^V \rightarrow (N^P, \mathcal{P})$
u'	7	Next physical node assigned in node mapping algorithm [81]
$S_{nmax} (S_{lmax})$	7	Maximum node (link) stress in G^P [81]
$S_N(v) (S_L(l))$	7	Current node (link) stress in G^P [81]
l	7	Index of physical links [81]
v	7	Index of physical nodes to map [81]
u	7	index of mapped physical nodes in node mapping algorithm [81]
$L(v)$	7	Set of links adjacent to physical node v [81]
$D(v, u)$	7	Distance between physical node v and u [81]
$\Pi(G^V)$	7	Revenue for allocating virtual network G^V [79]
CPU_r and bw_r	7	CPU and bandwidth required by the virtual network [79]
CPU_a and bw_a	7	CPU and bandwidth available on a physical network [79]
Ω	7	Price normalization factor [79]
$H(n^P)$	7	available resource on physical node n^P [81]
$R_N (R_L)$	8	Physical node (link) stress ratio [79]
$U^k(\cdot)$	8	Convex objective function run by virtual network k [37]
n_0	8	Number of virtual networks to simultaneously map [37]
$C^{(k)} = c_{ij}^{(k)}$	8	Binary matrix of capacity constraints for virtual network k using virtual path j on physical link l [37]
$y^{(k)}$	8	virtual link capacities for virtual network k [37]
$z^{(k)}$	8	Path rate vector for virtual network k [37]
$g^{(k)}$	8	General convex constraint for virtual network k [37]
\mathcal{D}	8	Matrix of physical link capacity
$w^{(k)}$	8	Weight assigned to virtual network k in the slice allocation phase
ω_{ij}	9	Weight (or utilization) imposed on resource j by user i ,
P_j	9	Price (in dollars) of the resource j [43]
U_j	9	Overall utilization of resource j [43]
\mathcal{R}_j	9	Physical CPU capacity of resource j in a <i>Colocation Game</i> [43]
$\mathcal{K}_j(i)$	9	Colocation cost for user i when mapped to resource j
a_{ij}	10	binary variable representing element i in the j^{th} set in a Set Packing Problem
w_j	10	Weight assigned to user requesting the set of resources —or objects— j in any allocation (Set Packing Problem)
y_j	10	Binary allocation variable for object j in a Set Packing Problem
$W(O)$	10	Set of users W (objects O) to be allocated in a Set Packing Problem
Q	10	Collection of subsets of objects in a Set Packing Problem
b_i	10	Number of copies for each object i in a Set Packing Problem
c_i	11	Cost of opening a facility at location i in a Facility Location Problem
d_{ij}	11	Cost of serving a user j from facility i
z_i	11	Binary variable showing whether or not the facility is selected at location i
x_{ij}	11	Binary variable that associates user j served by facility i in Facility Location Problem
x_i	11	Decision variable for location i , which is equal to one if the facility is selected
$f(\cdot), g(\cdot), h(\cdot)$	12	Utility functions for the discovery, virtual network mapping and allocation phase
$\gamma (\gamma_j)$	12	Number of virtual nodes (requested by user j)
$\psi (\psi_j)$	12	Number of virtual links (requested by user j)
$n_{ij}^V (n_{ij}^P)$	12	Decision variable on virtual (physical) node mappable (mapped) to user j
$l_{ij}^V (l_{ij}^P)$	12	Decision variable on virtual (physical loop-free path) link mappable (mapped) to user j
$\Theta_{ij} (\Phi_{kj})$	12	System's revenue when user j gets assigned to virtual node i (virtual link k .)
$C_i^n (C_k^l)$	12	Max virtual nodes (links) that can be simultaneously hosted on the physical node i (physical path k)

TABLE I
NOTATIONS USED IN THE PAPER.

B. Network mapping without constraints

The problem of static assignments of resources to a virtual network has been investigated in [81]. Since it is NP-hard, the authors proposed a heuristic to select physical nodes with lower *stress* (*i.e.*, with the lower number of virtual nodes already assigned to a given physical node), in an attempt to balance the load. The algorithm consists of two separate phases: node mapping and link mapping. The node mapping phase consists of an initialization step—cluster center localization—and an iterative subroutine—substrate node selection—that progressively selects the next physical node u' to which the next virtual node is mapped, *i.e.* the physical node with the least stress.

In particular, the center cluster is selected as follows:

$$u' = \arg \max_v \left\{ [S_{nmax} - S_N(v)] \sum_{l \in L(v)} [S_{lmax} - S_L(l)] \right\}$$

where S_{nmax} and S_{lmax} are the maximum node and link stress seen so far in the physical network, respectively. $S_N(v)$ is the stress on the physical node v , while $S_L(l)$ is the stress on the physical link l . $[S_{nmax} - S_N(v)]$ captures the availability of node v , while the availability on the links adjacent to v is captured by $\sum_{l \in L(v)} [S_{lmax} - S_L(l)]$.

The substrate node selection subroutine maps the remaining virtual nodes by minimizing a potential function proportional to both node and link stress on the physical network, *i.e.*:

$$u' = \arg \min_v \frac{\sum_{u \in V_A} D(v, u)}{S_{nmax} - S_N(v) + \epsilon}$$

where V_A is the set of already selected substrate nodes, v is an index over all physical nodes (so v could be the same as some u), ϵ is a small constant to avoid division by zero, and D is the distance between any two physical nodes v and u and it is defined as:

$$D(v, u) = \min_{p \in \mathcal{P}(u, v)} \sum_{l \in p} \frac{1}{S_{lmax} - S_L(l) + \epsilon}$$

where p is an element of all loop-free paths $\mathcal{P}(u, v)$ on the physical network that connects nodes u and v . The node mapping phase successfully terminates when all the virtual nodes are mapped.

The link mapping invokes a shortest path algorithm to find a minimum hop (loop-free) physical path connecting any pair of virtual nodes.

In the same paper, the authors modify this algorithm by subdividing the complete topology of a virtual network into smaller star topologies. These sub-topologies can more readily fit into regions of low stress in the physical network.

C. Network mapping with constraints

Many of the solutions to the virtual network mapping problem consider some constraints in the query specification. Lu and Turner [55] for example, introduce flow constraints in a mapping of a single virtual network. The NP-hard mapping problem is solved by greedily finding a backbone-star topology of physical nodes (if it exists, otherwise the slice cannot be

embedded), and the choice is refined iteratively by minimizing a notion of cost associated with the candidate topologies. The cost metric of a virtual link is proportional to the product of its capacity and its physical length. No guarantees on the convergence to an optimal topology mapping are provided, and only bandwidth constraints are imposed.

A novel outlook on the virtual network mapping problem for virtual network testbeds is considered in [21]. A topology and a set of (upper and lower bound) constraints on the physical resources are given, and a feasible mapping is sought. In order to reduce the search space of the NP-hard problem, a depth-first search with pruning as soon as a mapping becomes infeasible is used.

Another solution that considers embedding with constraints is presented in [52]. The authors propose a backtracking algorithm based on a subgraph isomorphism search method [48], that maps nodes and links simultaneously. The advantage of a single step node-link approach is that link constraints are taken into account at each step of the node mapping, therefore when a bad decision is detected, it can be adjusted by backtracking to the last valid mapping. With a two-stage approach instead, the remapping would have to be done for all the nodes, which is computationally expensive.

D. Network mapping + allocation

In all the solutions that focus only on the virtual network mapping task, only a single virtual network is considered (with or without constraints), and no resource allocation mechanism is provided. In case the mapping algorithm is designed for virtual network testbeds such as Emulab [77] or Planetlab [65], this may not be an issue except in rare cases, *e.g.*, during conference deadlines (see *e.g.*, Figure 1 in [5]). The lack of resource allocation is instead detrimental to an efficient slice embedding when the system aims to embed virtual networks (slices) that are profitable to the leasing infrastructure.

We discuss the case study of [79], that adds resource allocation to the virtual network mapping task, and hence introduces cooperation between the last two tasks of the slice embedding problem. The solution proposed in [79] is targeted specifically for infrastructure providers, as the physical resources considered—bandwidth and CPU—are assumed to be rentable. The authors define a revenue function R for each requested virtual network $G^V = (N^V, L^V)$ as:

$$\Pi(G^V) = \sum_{l^V \in L^V} bw_r(l^V) + \Omega \sum_{n^V \in N^V} CPU_r(n^V), \quad (2)$$

where $bw_r(l^V)$ and $CPU_r(n^V)$ are the bandwidth and the CPU requirements for the virtual link l^V and the virtual node n^V , respectively. L^V and N^V are the sets of requested virtual links and nodes, and Ω captures the price difference that the infrastructure provider may charge for CPU and bandwidth.

The algorithm is depicted in Figure 3: after collecting a set of requests, a greedy node mapping algorithm with the objective of maximizing the (long term) revenue R is run. In particular, the algorithm consists of the following three steps:

- 1) First the requests are sorted by revenue $\Pi(G^V)$ so that the most profitable mapping is sought with highest priority.




Fig. 3. Path splitting and migration mapping algorithm [79].

- 2) Then the physical nodes with insufficient available CPU capacity are discarded to reduce the complexity of the search.
- 3) Similarly to [81] (see Section V-B), a virtual node is mapped on the physical node n^P (if it exists) that maximizes the available resources H , where:

$$H(n^P) = CPU_a(n^P) \sum_{l^P \in L(n^P)} bw_a(l^P)$$

$CPU_a(n^P)$ and $bw_a(l^P)$ are the CPU and bandwidth available on the physical node n^P and link l^P , respectively, and $L(n^P)$ is the set of links adjacent to n^P .

After the node mapping, different link mapping algorithms are presented. First, the authors propose to use a *k-shortest path* algorithm [27]. The originality of this paper though, lies in the improvement of such a link assignment algorithm through two techniques: *path splitting* and *path migration*. In path splitting the virtual routers forward a fraction of the traffic through different physical paths to avoid congestion of critical physical links useful to host other virtual networks. Path migration instead is adopted to further improve the resource utilization as it consists of a periodic link mapping re-computation with a larger set of pre-mapped virtual networks, leaving unchanged both node mapping—virtual node cannot migrate on another physical node—and the path splitting ratios—fraction of the total virtual links requested to which at least two physical loop-free paths are assigned. After the link mapping algorithm, the slice requests that could not be embedded are queued for a re-allocation attempt, and they are definitively discarded if they fail a given number of attempts.

Inspired by [79] and by the PageRank algorithm [63], two topology-aware virtual network mapping and allocation algorithms (*Random Walk MaxMatch* and *Random Walk Breath First Search*) have been recently proposed [15]. The novelty, and common underlying idea of the two algorithms, is to use the same Markov chain model used in PageRank [63] to sort both physical and virtual nodes (instead of web pages), and map the most important virtual nodes to the most important physical nodes. A physical (virtual) node is highly ranked not only if it has available (required) CPU, and its adjacent links

have available (required) bandwidth (as in [79]), but also if its neighbors (recursively) have high rank.

After sorting both physical and virtual nodes, highly ranked virtual nodes are mapped to highly ranked physical nodes.

E. Dynamic approaches to network mapping and allocation

As mentioned in Section III-B, in the virtual network mapping task, virtual resources may be statically assigned to each physical resource, or they can be reassigned to maximize some notion of utility during the lifetime of a slice.

Many algorithms whose task is simply to discover feasible mappings are considered static, whether they use simulated annealing [67], genetic algorithms [77], or backtrack heuristics [54], [52]. A static resource assignment for multiple virtual networks though, especially when each virtual network needs to be customized to a particular application, can lead to lower performance and under utilization of the physical resources. Being aware of such inefficiencies, adaptive mechanisms to re-allocate physical resources, on demand or periodically, have been proposed.

Zan and Ammar [81] have proposed a dynamic version of their mapping algorithm, in which critical nodes and links in the physical network are periodically identified. To evaluate the current stress levels S_N and S_L for nodes and links, two metrics are defined: the node and link stress ratio (R_N and R_L). The former is the ratio between the maximum node stress and the average node stress across the whole physical network, while the latter is the ratio between the maximum link stress and the average link stress. Formally:

$$R_N = \frac{\max_{v \in N^P} S_N(v)}{|\sum_{v \in N^P} S_N(v)| / |N^P|}$$

$$R_L = \frac{\max_{l \in L^P} S_L(l)}{|\sum_{v \in L^P} S_L(l)| / |L^P|}$$

where N^P and L^P are the set of physical nodes and edges of the hosting infrastructure, respectively. R_N and R_L are periodically compared, and new requests are mapped optimizing the node stress if $R_N > R_L$, or the link stress if $R_N < R_L$. This process is iterated with the aim of minimizing the stress across the entire physical network.

Dynamic mapping approaches also include the solutions proposed in [55], since virtual links are iteratively reassigned, and in [79], due to the migration operations. Although without any considerations to the node constraints, also in [29] the authors consider a dynamic topology mapping for virtual networks.

A solution to the dynamic network mapping problem that uses optimization theory was presented in the *DaVinci* architecture—Dynamically Adaptive Virtual Networks for a Customized Internet [37]. A physical network with n_0 virtual mapped networks is considered. Each virtual network $k = 1, \dots, n_0$ runs a distributed protocol to maximize its own performance objective function $U^k(\cdot)$, assumed to be convex with respect to network parameters, efficiently utilizing the resources assigned to it. These objective functions, assumed to be known to a centralized authority, may vary with the

traffic class (*e.g.*, delay-sensitive traffic may wish to choose paths with low propagation-delay and keep the queues small to reduce queuing delay, while throughput-sensitive traffic may wish to maximize aggregate user utility, as a function of rate), and may depend on both virtual path rates $z^{(k)}$ and the bandwidth share $y^{(k)}$ of virtual network k over every physical link l .

The traffic-management protocols running in each virtual network are envisioned as the solution to the following optimization problem:

$$\begin{aligned} & \text{maximize} && U^{(k)}(z^{(k)}, y^{(k)}) \\ & \text{subject to} && C^{(k)} z^{(k)} \leq y^{(k)} \\ & && g^{(k)}(z^{(k)}) \leq 0 \\ & && z^{(k)} \geq 0 \end{aligned} \quad (3)$$

where $z^{(k)}$ are the variables (virtual path rates), $g^{(k)}(z^{(k)})$ are general convex constraints and $C^{(k)}$ defines the mapping of virtual paths over physical links. This means that there could be many flows on a single virtual network, *i.e.*, a virtual network k may host (allocate) multiple services. In particular, $c_{lj}^{(k)} = 1$ if virtual path j in virtual network k uses the physical link l and 0 otherwise.⁴

The dynamism of this approach lies in the periodic bandwidth reassignment among the n_0 hosted virtual networks. The physical network in fact runs another (convex) optimization problem, whose objective is to maximize the aggregate utility of all the virtual networks, subject to some convex constraints:

$$\begin{aligned} & \text{maximize} && \sum_k w^{(k)} U^{(k)}(z^{(k)}, y^{(k)}) \\ & \text{subject to} && C^{(k)} z^{(k)} \leq y^{(k)} \quad \forall k \\ & && \sum_k y^{(k)} \leq \mathcal{D} \\ & && g^{(k)}(z^{(k)}) \leq 0 \quad \forall k \\ & && z^{(k)} \geq 0 \quad \forall k \\ & \text{variables} && z^{(k)}, y^{(k)} \quad \forall k \end{aligned} \quad (4)$$

where $w^{(k)}$ is a weight (or priority) that a centralized authority in charge of embedding the slices assigns to each virtual network, and \mathcal{D} represents the physical capacities. Note how there are two levels of resource allocation in this model: each slice maximizes its utility by assigning capacity to each service hosted, and the physical network maximizes its utility by assigning resources to some slices.

As in [79], the DaVinci architecture allows (virtual) path splitting, causing packet reordering problems, and assumes the node mapping to be given. A more serious limitation is the assumption that physical links are aware of the performance objectives of all the virtual networks, which may not be possible in real world settings.

F. Distributed Virtual Network Mapping Solutions

All the previously discussed solutions assumed a centralized entity that would coordinate the mapping assignment. In other words, their solutions are limited to the intra-domain virtual network mapping. These solutions are well suited for

⁴As in [42], a system may in fact be hosted on a physical infrastructure by leasing a slice, and then provide other services by hosting (even recursively) other slices.

enterprises serving slices to their customers by using only their private resources. However, when a service must be provisioned using resources across multiple provider domains, the assumption of a complete knowledge of the substrate network becomes invalid, and another set of interesting research challenges arises.

It is well known that providers are not happy to share traffic matrices or topology information, useful for accomplishing an efficient distributed virtual network mapping. As a result, existing embedding algorithms that assume complete knowledge of the substrate network are not applicable in this scenario.

To the best of our knowledge, the first distributed virtual network mapping problem was devised by Houidi *et al.* [40]. The protocol assumes that all the requests are hub-spoke topologies, and runs concurrently three distributed algorithms at each substrate node: a *capacity-node-sorting* algorithm, a *shortest path tree* algorithm, and a *main mapping* algorithm. The first two are periodically executed to provide up to date information on node and link capacities to the main mapping.

For every element mapped, there has to be a trigger and a synchronization phase across all the nodes. The algorithm is composed of two phases: when all nodes are mapped, a shortest path algorithm is run to map the virtual links. The authors propose the use of an external signalling/control network to alleviate the problem of the heavy overhead.

In [17], the authors proposed a simultaneous node and link distributed class of mapping algorithms. In order to coordinate the node and the link mapping phases, the distributed mapping algorithm is run on the physical topology augmented with some additional logical elements (meta node and meta links) associated with the location of the physical resource.

In [16], the same authors describe a similar distributed (policy-based) inter-domain mapping protocol, based on geographic location of the physical network: PolyViNE. Each network provider keeps track of the location information of their own substrate nodes employing a hierarchical addressing scheme, and advertising availability and price information to its neighbors via a Location Awareness Protocol (LAP) — a hybrid gossiping - publish/subscribe protocol. Gossiping is used to disseminate information in a neighborhood of a network provider and pub/sub is employed so a provider could subscribe to other providers which are not in its neighborhood. PolyViNE also considers a reputation metric to cope with the lack of truthfulness in disseminating the information with the LAP protocol.

VI. ALLOCATION

Different strategies have been proposed when allocating physical resources to independent parties. Some solutions prefer practicality to efficiency, and adopt best effort approaches, (*see, e.g.*, PlanetLab [65]), while others let the (selfish) users decide the allocation outcome with a game [43], [42]. When instead it is the system that enforces the allocation, it can do it with [33] or without [5] providing guarantees. In the remainder of this section we focus first on the game theoretic solutions to resource allocation, and then on the latter case, describing first a set of solutions dealing with market-based mechanisms [5],

[49], [9], and then a reservation-based approach [33]. All those solutions focus solely on the standalone allocation task of the slice embedding problem.

A. Game-theory based allocation

Londoño *et al.* [43] defined a general pure-strategies colocation game which allows users to decide on the allocation of their requests. In their setting, customer interactions is driven by the rational behavior of users, who are free to relocate and choose whatever is best for their own interests. Under their model, a slice consists of a single node in a graph that needs to be assigned to a single resource. They define a cost function $\mathcal{K}_j(i)$ for user i when mapped to resource j as

$$\mathcal{K}_j(i) = P_j \frac{\omega_{ij}}{U_j} \quad (5)$$

where ω_{ij} is the weight (or utilization) imposed on resource j by user i , P_j is the price (in dollars) of the resource j , U_j is the overall utilization of resource j , which must satisfy its capacity constraint

$$U_j = \sum_{i \in J} \omega_i \leq \mathcal{R}_j \quad (6)$$

where J is the set of users mapped on resource j , and \mathcal{R}_j is the physical CPU capacity of resource j .

They define a rational “move” of user i from resource a to resource b if $\mathcal{R}_b(i) < \mathcal{R}_a(i)$. The game terminates when no user has a move that minimizes her cost. Note how the utility of a user (player) is higher if she can move to a more “loaded” resource, as she will share the cost with the other players hosted on the same resource.

The model has two interesting properties. First, the interaction among customers competing for resources leads to a Nash Equilibrium (NE), *i.e.* a state where no customer in the system has incentive to relocate. Second, it has been shown that the Price of Anarchy—the ratio between the overall cost of all customers under the worst-case NE and that cost under a socially optimal solution—is bounded by 3/2 and by 2 for homogeneous and heterogeneous resources, respectively. The authors also provide a generalized version of this game (General Colocation Game), in which resources to be allocated are graphs representing the set of virtual resources and underlying relationships that are necessary to support a specific user application or task. In this general case however, the equilibrium results no longer hold as the existence of a NE is not always guaranteed.

The work by Chen and Roughgarden [14] also introduces a game theoretical approach to link allocation in the form of source-destination flows on a shared network. Each flow has a weight and the cost of the link is split in proportion to the ratio between the weight of a flow and the total weights of all the flows sharing the physical link.

As shown, even recently by Chowdhury [17], in a centralized solution, the virtual network mapping problem can be thought of as a flow allocation problem where the virtual network is a flow to be allocated on a physical network.

These two game theoretic approaches may serve as inspiring example for new allocation strategies involving different

selfish principles for virtual service provisioning / competition. A system may in fact let the users play a game in which the set of strategies represent the set of different virtual networks to collocate with, in order to share the infrastructure provider costs.

B. Market-based allocation

When demand exceeds supply and not all needs can be met, virtualization systems’ goals can no longer be related to maximizing utilization, but different policies to guide resource allocation decisions have to be designed. A natural policy is to seek efficiency, namely, to allocate resources to the set of users that bring to the system the highest utility. To such an extent, the research community has frequently proposed market-based mechanisms to allocate resources among competing interests while maximizing the overall utility of the users. A subclass of solutions dealing with this type of allocation is represented by auction-based systems. An auction is the process of buying and selling goods or services by offering them up for bid, taking bids, and then selling them to the highest bidder.

Few examples where auctions have been adopted in virtualization-oriented systems are Bellagio [5], Tycoon [49] and Mirage [9]. They use a combinatorial auction mechanism with the goal of maximizing a social utility (the sum of the utilities for the users who get the resources allocated).

A *Combinatorial Auction Problem (CAP)* is equivalent to a *Set Packing Problem (SPP)*, a well studied integer program: given a set O of elements and a collection Q of subsets of these elements, with non-negative weights, SPP is the problem of finding the largest weight collection of subsets that are pairwise disjoint. This problem can be formulated as an integer program as follows: we let $\mathbf{y}_j = 1$ if the j^{th} set in W with weight w_j is selected and $\mathbf{y}_j = 0$, otherwise. Then we let $a_{ij} = 1$ if the j^{th} set in W contains element $i \in O$ and zero otherwise. If we assume also that there are b_i copies of the same element i , then we have:

$$\begin{aligned} & \text{maximize} && \sum_{j \in W} w_j \mathbf{y}_j \\ & \text{subject to} && \sum_{j \in W} a_{ij} \mathbf{y}_j \leq b_i \quad \forall i \in O \\ & && \mathbf{y}_j = \{0, 1\} \quad \forall j \in Q \end{aligned} \quad (7)$$

SPP is equivalent to a CAP if we think of the \mathbf{y}_j s as the users to be possibly allocated and requesting a subset of resources in O , and w_j as the values of their bids. Note that solving a set packing problem is NP-Hard [25]. This means that optimal algorithms to determine the winner in an auction are also NP-Hard. To deal with this complexity, many heuristics have been proposed. In [5] for example, the authors rely on a thresholding auction mechanism called SHARE [20], which uses a first-fit packing heuristic.

Another example of a system that handles the allocation for multiple users with an auction is Tycoon [49]. In Tycoon, users place bids on the different resources they need. The fraction of resource allocated to one user is her proportional share of the total bids in the system. For this reason, Tycoon’s allocation mechanism can also be considered best-effort: there are no guarantees that users will receive the desired fraction of the resources. The bidding process is continuous in the sense that

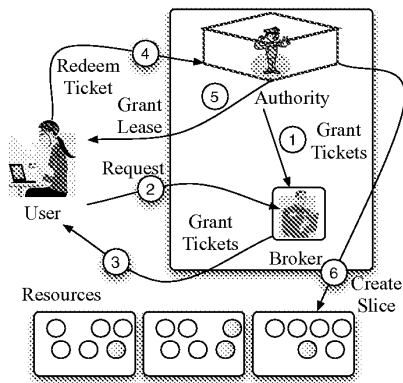


Fig. 4. Architecture and allocation phases in SHARP [33].

any user may modify or withdraw their bid at any point in time, and the allocation for all the users can be adjusted according to the new bid-to-total ratio.

As pointed out in [4], although market-based allocation systems can improve user satisfaction on large-scale federated infrastructures, and may lead to a social optimal resource allocation, there are few issues that should be taken into account when designing such mechanisms. In fact, the system may be exploited by users in many ways. Current auction-based resource allocation systems often employ very simple mechanisms, and there are known problems that may impact efficiency or fairness (see [4], Section 6). We report three of them here:

- *underbidding*: users know that the overall demand is low and they can drive the prices down.
- *iterative bidding*: often one shot auctions are not enough to reach optimal resource allocation but the iterations may not end by the time the allocations are needed.
- *auction sandwich attack*: occurs when users bid for resources in several time intervals. This attack gives the opportunity to deprive other users of resources they need, lowering the overall system utility.

C. Reservation-based allocation

As the last piece of this section on allocation approaches, we discuss a reservation-based system, SHARP [33] whose architecture is depicted in Figure 4. The system introduces a level of indirection between the user and the centralized *authority* responsible for authentication and for building the slice: the *broker or agent*. The authority issues a number of *tickets* to a number of brokers (usually many brokers responsible for a subset of resources are connected). Users then ask and eventually get tickets, and later in time, they redeem their tickets to the authority that does the final slice assignment (Figure 4).

This approach has many interesting properties but it may lead to undesirable effects. For example, coexisting brokers are allowed to split the resources: whoever has more requests should be responsible for a bigger fraction of them. This




Fig. 5. Different values of Oversubscription Degree tune allocation guarantees [33].

sharing of responsibilities may bring fragmentation problems as resources become divided into many small pieces over time. Fragmentation of the resources is a weakness, as the resources become effectively unusable being divided into pieces that are too small to satisfy the current demands.

One of the most relevant contributions of SHARP in the context of the slice embedding problem, is the rule of the *Oversubscription Degree (OD)*. The *OD* is defined as the ratio between the number of issued tickets and the number of available resources. When *OD* is greater than one, *i.e.*, there are more tickets than actual available resources, the user has a probability less than one to be allocated even though she owns a ticket. When instead *OD* is less or equal than one, users with tickets have guaranteed allocation (Figure 5).

Note how the level of guarantees changes with *OD*. In particular, when the number of tickets issued by the authority increases, the level of guarantees decreases. The authors say that the allocation policy tends to a first come first serve for *OD* that tends to infinity. In other words, if there are infinite tickets, there is no reservation at all, and simply the first requests will be allocated. The oversubscription degree is not only useful to control the level of guarantees (by issuing less tickets than available resources the damage from resource loss if an agent fails or becomes unreachable is limited), but it can be used also to improve resource utilization by means of statistical multiplexing the available resources.

VII. FACILITY LOCATION PROBLEMS

In this section we discuss a set of problems similar to slice embedding: the facility location problems. Facility location is a branch of operations research whose goal is to assign a number of facilities to a set of users, while minimizing a given cost function. An ample amount of literature exists on centralized [61], [76] or distributed [32], [50] solutions for this NP-hard problem [44].

The centralized facility location problem is defined as follows: suppose we are given n potential facility locations and a list of m users who need to be serviced from these locations. There is an initial fixed cost c_i of opening the facility at location i , while there is a cost d_{ij} of serving a user j from facility i . The goal is to select (open) a set of facility locations and to assign each user to one facility, while minimizing the cost.

In order to model this problem, we define a binary decision variable z_i for each location i , which is equal to one if the facility is selected, and 0 otherwise. In addition, we define a binary variable $x_{ij} = 1$ if user j is served by facility i , and 0

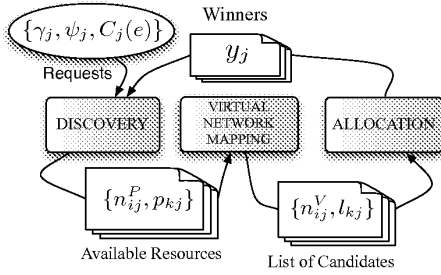


Fig. 6. Interactions and data exchanges in the slice embedding problem.

otherwise. The facility location problem is then formulated as follows:

$$\begin{aligned} & \text{minimize} && \sum_{i=1}^n c_i z_i + \sum_{j=1}^m \sum_{i=1}^n d_{ij} x_{ij} \\ & \text{subject to} && \sum_{i=1}^n x_{ij} = 1 \quad \forall j \\ & && x_{ij} \leq z_i \quad \forall i, \forall j \\ & && x_{ij}, z_i \in \{0, 1\} \quad \forall i, \forall j. \end{aligned} \quad (8)$$

The affine constraint $\sum_{i=1}^n x_{ij} = 1$ enforces a single facility to a user, while the constraint $x_{ij} \leq z_i$ ensures that if there is no facility at location i , *i.e.* $z_i = 0$, then user j cannot be served there, and we must have $x_{ij} = 0$.

The facility location and the slice embedding problems may look similar since both have the high level goal of assigning a set of resources to a set of users, and both solutions require knowledge of the resource availability to work efficiently. However, the two problems differ in many aspects: first, the facility location assignment algorithms usually assume no cooperation with the discovery protocol, while in the slice embedding problem the resource discovery is directly interacting with the other two phases, as we discuss in the next section. More importantly, the slice embedding problem assumes that resources are virtual instances of both nodes and edges of the physical infrastructure, as opposed to standalone facilities to be assigned to users. This detail leads to important differences in the assignment algorithms as explained in [79] and in [52]. Moreover, facility location problems assume that each and every user has to be assigned to only one physical resource (and the positive cost to the system of such assignment is minimized), while this assumption disappears in the slice embedding problem where, in general, there may not be the guarantee that every user is allocated.

VIII. ON MODELING THE SLICE EMBEDDING PROBLEM

In this section we use optimization theory to model the interactions between the three phases of the slice embedding problem. We first model each standalone phase — resource discovery, virtual network mapping, and allocation — and subsequently model the slice embedding problem as a whole by merging the three phases into a centralized optimization problem. Consider the ellipsoid in Figure 6, augmented from Figure 1 (we explain the rest of the notation throughout this section): user j requests a virtual network composed of $\gamma_j \in \mathbb{N}$ virtual nodes, $\psi_j \in \mathbb{N}$ virtual links and a vector of constraints $C_j(e) = \langle C_j(e_1), \dots, C_j(e_c) \rangle$ where e is a vector of

$c = \gamma_j + \psi_j$ elements — nodes and links — of the network.

Discovery: To model the resource discovery we introduce two binary variables, n_i^P and p_k that are equal to 1 if the i^{th} physical node and the k^{th} loop-free physical path, respectively, are available, and zero otherwise. An element is available if a discovery operation is able to find it, given a set of protocol parameters, *e.g.*, find all loop-free paths within a given deadline, or find as many available physical nodes as possible within a given number of hops.

If the system does not return at least γ physical nodes and ψ available loop-free physical paths among all the possible N nodes and P paths of the physical network G^P , then the user's request should be immediately discarded. Among all possible resources, the system may choose to return a set that maximizes a given notion of utility. Those utilities may have the role of selecting the resources that are closer — with respect to some notion of distance — to the given set of constraints $C(e)$. If we denote as $u_i \in \mathbb{R}$ and $\omega_k \in \mathbb{R}$ the utility of physical nodes and paths respectively, then the discovery phase of the slice embedding problem can be modeled as follows:

$$\begin{aligned} & \text{maximize} && f(n_i^P, p_k) = \sum_{i \in N} u_i n_i^P + \sum_{k \in P} \omega_k p_k \\ & \text{subject to} && \sum_{i \in N} n_i^P \geq \gamma \\ & && \sum_{k \in P} p_k \geq \psi \\ & && n_i^P, p_k \in \{0, 1\} \quad \forall i, \forall k \end{aligned} \quad (9)$$

After the discovery phase is completed, the vectors of available physical resources (n^P, p) are passed to the virtual network mapper.

Virtual Network Mapping: This phase takes as input all the available resources (subset of all the existing resources) $P' \subseteq P$ and $N' \subseteq N$, maps virtual nodes to physical nodes, virtual links to loop-free physical paths, and returns a list of candidates — virtual nodes and virtual links — to the allocator. To model this phase, we define two sets of binary variables $n_{ij}^V \forall i \in N'$, and $l_{kj} \forall k \in P', \forall j \in J$, where J is the set of users requesting a slice. $n_{ij}^V = 1$ if a virtual instance of node i could possibly be mapped to user j and zero otherwise, while $l_{kj} = 1$ if a virtual instance of the loop-free physical path k could possibly be mapped to user j , and zero otherwise. The virtual network mapping phase of the slice embedding problem can hence be modeled by the following optimization problem:

$$\begin{aligned} & \text{maximize} && g(n_{ij}^V, l_{kj}) = \sum_{j \in J} (\sum_{i \in N'} \Theta_{ij} n_{ij}^V + \sum_{k \in P'} \Phi_{kj} l_{kj}) \\ & \text{subject to} && \sum_{i \in N'} n_{ij}^V = \gamma_j \quad \forall j \in J \\ & && \sum_{k \in P'} l_{kj} = \psi_j \quad \forall j \in J \\ & && n_{ij}^V = n_{ij}^P \quad \forall i \in N' \quad \forall j \in J \\ & && l_{kj} \leq p_{kj} \quad \forall k \in P' \quad \forall j \in J \\ & && n_{ij}^V, n_{ij}^P, p_{kj}, l_{kj} \in \{0, 1\} \quad \forall i, \forall j, \forall k, \end{aligned} \quad (10)$$

where Θ_{ij} is the revenue that the system would get if user j gets assigned to virtual node i , and Φ_{kj} is the system's revenue if the user j gets the virtual link k . The first two constraints enforce that all the virtual resources requested by each user are mapped, the third constraint ensures that the one-to-one mapping between virtual and physical nodes is satisfied, and the fourth constraint ensures that at least one loop-free physical path is going to be assigned to each virtual link of

the requested slice.

Allocation: As soon as the virtual mapping candidates have been identified, a packing problem needs to be run, considering both user priorities and physical constraints. Enhancing the level of details from the standard set packing problem [71] to virtual nodes and links, we model the allocation phase of the slice embedding problem as follows:

$$\begin{aligned} & \text{maximize} && h(y_j) = \sum_{j \in J} w_j y_j \\ & \text{subject to} && \sum_{j \in J} n_{ij}^V y_j \leq C_i^n \quad \forall i \in N' \\ & && \sum_{j \in J} l_{kj} y_j \leq C_k^l \quad \forall k \in P' \\ & && y_j \in \{0, 1\} \quad \forall j \end{aligned} \quad (11)$$

where C_i^n and C_k^l are the number of virtual nodes and links respectively, that can be simultaneously hosted on the physical node i and physical path k , respectively, and y_j is a binary variable equal to 1 if user j has been allocated and zero otherwise. A weight w_j is assigned to each user j , and it depends on the allocation policy used (*e.g.* in first-come first-serve, $w_j = w \quad \forall j$, or in a priority based allocation w_j represents the importance of allocating user j 's request). As multiple resources are typically required for an individual slice, the slice embedding needs to invoke the appropriate resource allocation methods on individual resources, and it does so throughout this last phase. Each resource type may in fact have its own allocation policy (*e.g.*, either guaranteed or best-effort resource allocation models), and this phase only ensures that users will not be able to exceed physical limits or their authorized resource usage. For example, the system may assign a weight $w_j = 0$ to a user that has not yet been authorized, even though her virtual network could be physically mapped.

Slice Embedding: In order to clarify how the three phases of the slice embedding problem interact and how they may impact efficiency in network virtualization, we formulate a centralized optimization problem that considers the slice embedding problem as a whole. In particular, we model the three phases as follows:

$$\text{maximize} \quad \alpha \cdot f(n_{ij}^P, p_{kj}) + \beta \cdot g(n_{ij}^V, l_{kj}) + \delta \cdot h(y_j)$$

$$\text{subject to} \quad \sum_{i \in N} n_{ij}^P \geq \gamma_j \quad \forall j \quad (12a)$$

$$\sum_{k \in P} p_{kj} \geq \psi_j \quad \forall j \quad (12b)$$

$$\sum_i n_{ij}^V = \gamma_j \quad \forall j \quad (12c)$$

$$\sum_k l_{kj} = \psi_j \quad \forall j \quad (12d)$$

$$n_{ij}^V = n_{ij}^P \quad \forall i \quad \forall j \quad (12e)$$

$$l_{kj} \leq p_{kj} \quad \forall k \quad \forall j \quad (12f)$$

$$\sum_{j \in J} n_{ij}^V y_j \leq C_i^n \quad \forall i \quad (12g)$$

$$\sum_{j \in J} l_{kj} y_j \leq C_k^l \quad \forall k \quad (12h)$$

$$y_j \leq n_{ij}^V \quad \forall i \quad \forall j \quad (12i)$$

$$y_j \leq l_{kj} \quad \forall k \quad \forall j \quad (12j)$$

$$y_j, n_{ij}^P, p_{kj}, n_{ij}^V, l_{kj} \in \{0, 1\} \quad \forall i \quad \forall j \quad (12k)$$

where the first nine constraints (from (12a) to (12h)) are the same as in problems (9), (10) and (11), respectively, the two coupling constraints (12i) and (12j) guarantee that a user

is not allocated unless all the resources she queried can be mapped, and α , β and δ are normalization factors.

Note how constraints (12e), (12f) and constraints (12i) and (12j) bind the three phases of the slice embedding problem together. However, all the above constraints have never been simultaneously considered before in related literature. In [79] for example, the first two as well as the last two constraints are omitted (plus $\alpha = \delta = 0$), and a global knowledge of the resource availability is assumed. Other solutions that focus only on the virtual network mapping phase (for example [81]), omit even the capacity constraints (12g) and (12h).

From an optimization theory point of view, constraint omissions in general may result in sub-optimal solutions while constraint additions may lead to infeasible solutions. For example, the resource discovery constraints impact the other phases of the slice embedding, since a physical resource not found certainly cannot be mapped or allocated. Moreover, it is useless to run the virtual network mapping phase on resources that can never be allocated because they will exceed the physical capacity constraints. As a consequence, centralized or distributed solutions for the slice embedding problem as a whole seem to be a valuable research subarea of network virtualization.

IX. OPEN PROBLEMS

In this section we present some research challenges that are important to achieving efficient slice embedding. In general, due to its complexity, an efficient and largely scalable solution for the slice embedding problem that involves all the three tasks is still elusive.

A. Devising new heuristics and approximation algorithms

As described in Section V, the virtual network mapping is often split into node and link mappings to reduce the complexity. Note, however, that such assignments are not independent. In other words, solving them sequentially introduces sub-optimality. Researchers should therefore keep in mind that node assignments affect link assignments and vice-versa when devising heuristics for this particular task of the slice embedding problem.

Another interesting research direction is to devise heuristics for conflicting objectives. For example, it is not clear whether load balancing is the only way to improve system performance as done in [81]. One can think about optimizing other objectives such as bin packing on the physical resources to save power. Clearly these two optimization approaches are different and over the lifetime of a slice, one may need to optimize one more than the other. The *load profiling* technique presented in [59], seems to be a more generalized approach than bin packing and load balancing, where neither extreme is the objective, and the system attempts to match some target load distribution across the physical resources.

Although approximation algorithms have been discussed for similar problems (see for example [46] or in [12]), to the best of our knowledge, only in [16] they have been applied to the virtual network mapping task, thus leaving the modeling of the interaction with discovery and allocation open for further research.

B. Addressing scalability and cooperation among the slice embedding tasks

In all the solutions discussed, it is assumed that allocators have ubiquitous and updated information on the physical network. A resource allocator's ability to make effective and efficient use of the available resources, however, is governed by how much information is available to it at the time it needs to make a decision. Thus, its interaction with the resource discovery is key. An important factor in this interaction is how much data must be passed back and forth between the two components. While passing node information—how much resources are still available on each particular physical node—should be manageable, path information is $O(n^2)$ in the number of nodes, and hence will scale poorly.

Another open question is whether and how a system can achieve efficient allocation with partial information: although we are not the first to advocate that resource discovery and allocation in virtualization oriented architectures should work tightly together (Ricci *et al.* in [68] for example, claim that the Emulab testbed is being improved by keeping this design principle in mind), it is still not clear how much data should pass between the discoverer and the allocator, how often the two tasks need to communicate, and which subset of available resources should be advertised to the allocator.

C. Modeling interactions between the slice embedding tasks

Generally, when designing solutions that involve different tasks of the slice embedding problem, researchers may utilize (distributed) optimization techniques. It is in fact possible to view each phase of the slice embedding problem as a standalone optimization problem, where different principles try to optimize the different tasks of the slice embedding problem, passing around a limited amount of information, to obtain a globally optimal embedding solution. An efficiency-overhead trade-off analysis of the mechanisms that involve such message passing among the tasks encompassing the slice embedding problem could be helpful in designing novel virtualization-based systems. Such an analysis could also be generalized to the cooperation among any coexisting infrastructure services [30], with the help of (centralized or distributed) optimization theory [8], [24], control or even game theory, for those cases where the principles involved are selfish or do not have incentives to cooperate.

D. Dissecting distributed decomposition alternatives

A systematic understanding of the decomposability structures of the slice embedding problem may help obtain the most appropriate distributed algorithms, given the application. Decomposition theory provides tools to build analytic foundations for the design of modularized and distributed control of both physical and virtual networks.

For a given problem representation, there are often many choices of distributed algorithms, each leading to different outcome of the global optimality versus message passing tradeoff [56], [64]. Which alternative is the best depends on the specifics of the slice embedding application.

We believe that qualitative or quantitative comparisons across architectural decomposition alternatives of the slice embedding problem is an interesting research area. When designing novel (virtual) network architectures for specific applications, to understand where to place functionalities and how to interface them is an issue that could be more critical than the design of how to execute and implement the functionalities themselves.

E. Supporting multiple allocators

Since each allocator can only make scheduling decisions based on the jobs submitted to it, it seems challenging to make multiple allocators work together, and this opens an interesting research direction. Allocation solutions consider only the scheduling problem, but another interesting problem is what to do *after* the resources are allocated. Since an infrastructure should be able to host customized virtual networks, each with different goals and constraints, we believe that there is not a “right” type of resource allocator, but resource allocators of modern distributed service architectures should rather support different policies for different applications that they support; for example, some users should be able to be allocated in a first come first serve manner, others should have soft or hard reservation guarantees. An architecture that would support a range of allocation policies is still missing.

F. Protocol Design and Implementation

The recently proposed distributed service architectures (*e.g.* NetServ [73] or RINA [23]) are a promising petri dish for testing novel protocols and distributed applications. In the case of RINA for example, (recursive) slice embedding protocols could be designed and prototyped over virtualization-based platforms. In particular, (inspired by [37]), we believe that designing and implementing efficient protocols to guarantee a given Service Level Agreement among slices managed by the same, or by different providers, is an interesting research area. In the case of the RINA architecture [23], where “Distributed Inter-process communication Facilities (DIF)” — the building blocks of the architecture — can be thought of as slices, this would mean designing recursive protocols to enable service provisioning across multiple tier-level providers. In fact, a DIF, just as a slice, is a service building block that can be repeated and composed in layers to build wider scoped services that meet user requirements.

Moreover, as mentioned in Section VI-A, distributed protocols to capture competition and interactions among slice embedding providers could be devised, assuming cooperation among different principles providing the service, or by means of a marketplace that allows selfish behavior.

X. CONCLUSIONS

Network virtualization has been proposed as the technology that will allow growing and testing of novel Internet architectures and protocols, overcoming the weaknesses of the current Internet, as well as testing them in repeatable and reproducible network conditions. Moreover, taking cue from current trends