

- 35 -

time stamp in the refresh command is contemporaneous with the global TxTs at the server, a 40 second resistant timeout is exerted at the server before responding to the refresh request. This is implemented by the hold ring previously described. If during the period of the resistant timeout, the content of the HTML page at the server is modified by another viewer, the modified page is returned to the browser as the response to the refresh request. This is by the process of qualification and requalification previously described. If during the period of the resistant timeout, the content of the HTML page at the server does not change, the server will respond to the request with either: a true refresh of the same page by resubmitting the original, or near original data to the browser; or issue an HTTP Code 204 response indicating no change in page status or content.

In the manner described a plurality of IUAs 1 have a near contemporaneous shared view of a Web page. This is view is also kept up to date as soon as the Web page changes.

- 36 -

Claims:

1. A system for processing requests in a request-response client-server computer network in which a plurality of clients communicate with a server by issuing requests to read or write data, comprising:
- 5 - a request receiver for receiving a plurality of requests from the clients;
 - a request analyser for analysing each of the plurality of requests in turn passed from the request receiver;
 - 10 - a request qualifier for qualifying each received request as either an actionable request for which action should be taken and a response issued, or a deferrable request for which action should be delayed and a response deferred until an actionable state is reached;
 - 15 - a state indicator for indicating, for each deferrable request, when the actionable state is reached such that each such deferrable request is requalified as an actionable request; and
 - 20 - a response indicator for indicating for each actionable request to the server that each such actionable request should be actioned and a response provided.
- 25
2. A system according to claim 1, wherein the request analyser comprises a request queue store for storing requests in order of arrival of requests at the request receiver.
- 30
3. A system according to claim 2 or 3, wherein the request analyser comprises a request retriever for retrieving requests from the request queue store in order of arrival at the request receiver and for

- 37 -

passing the requests so retrieved to the request
qualifier.

- 5
4. A system according to claim 1, 2 or 3, wherein the
request analyser is arranged to determine whether the
request is to read or write data.
- 10
5. A system according to claim 4, further comprising a
controller for controlling operation of the request
analyser, request qualifier, state indicator and
response indicator, and for optionally suspending
operation of each, if the request analyser determines
that the ~~next~~ request is to write data, until the
data has been written.
- 15
6. A system according to any preceding claim, wherein
the request analyser is arranged to extract from each
request a time indicator indicative of when the last
transaction providing a response from the server to
the client occurred.
- 20
7. A system according to claim 6, wherein the time
indicator is a time stamp.
- 25
8. A system according to claim 6 or 7, wherein the
request qualifier is arranged to compare the
extracted time indicator with the time the requested
data was last modified.
- 30
9. A system according to claim 8, wherein the request
qualifier is arranged to qualify a request as an
actionable request if the time indicator indicates a
time earlier than the time the requested data was
last modified.

- 38 -

10. A system according to claim 8 or 9, wherein the request qualifier is arranged to qualify the request as deferrable if the time indicator indicates a time later than or equal to the time the requested data was last modified.
- 5
11. A system according to any of claims 6 to 10, wherein the state indicator comprises a held request store for storing deferrable requests in time indicator order.
- 10
12. A system according to claim 11, wherein the state indicator comprises a held request retriever for retrieving requests from the held request store in time indicator order, the qualifier being arranged to qualify by comparing the time the requested data was last modified with the time indicator, the state indicator arranged to indicate that an actionable state is reached if the time indicator indicates a time older than the time the requested data was last modified.
- 15
- 20
13. A system according to claim 12, wherein the held request retriever is arranged to periodically retrieve requests from the held request store.
- 25
14. A system according to claim 12 or 13, wherein the held request retriever is arranged to retrieve requests from the held request store in response to a notification from the server that data stored at the server has changed.
- 30
15. A method for processing requests in a request-response client-server computer network in which a plurality of clients communicate with a server by issuing requests to read or write data, comprising:
- 35

- 39 -

- receiving a plurality of requests from the clients at a receiver;
 - analysing each of the plurality of requests passed from the request receiver in a request analyser;
 - 5 - qualifying each request in a request qualifier as either an actionable request for which action should be taken and a response issued, or a deferrable request for which action should be delayed and a response deferred until an
 - 10 actionable state is reached;
 - indicating from a state indicator, for each deferrable request, when the actionable state is reached such that each such deferrable request is requalified as an actionable request; and
 - 15 - indicating from a response indicator for each actionable request to the server that each such actionable request should be actioned and a response provided.
- 20 16. A method according to claim 15, wherein the step of analysing further comprises storing requests in a request queue store in order of arrival of the requests at the request receiver.
- 25 17. A method according to claim 15, or 16, wherein the step of analysing further comprises retrieving requests from the request queue store in order of arrival at the request receiver and passing the requests so retrieved to the request qualifier.
- 30 18. A method according to any of claims 15, 16 or 17, wherein the step of analysing further comprises determining whether the request is to read or write data.
- 35

- 40 -

19. A method according to claim 18, further comprising
controlling operation of the request analyser,
request qualifier, state indicator and response
indicator with a controller, and optionally
suspending operation of each, if the request analyser
5 determines that the next request is to write data,
until the data has been written.
20. A method according to any of claims 15 to 19, wherein
the step of analysing each request comprises
10 extracting from each request a time indicator
indicative of when the last transaction providing a
response-from-the server to the client occurred.
21. A method according to claim 20, wherein the time
15 indicator is a time stamp.
22. A method according to claim 20 or 21 wherein, for a
request to read data, the request qualifier compares
the extracted time indicator with the time the
20 requested data was last modified.
23. A method according to claim 22, wherein the request
qualifier qualifies the request as an actionable
request if the time indicator indicates a time
25 earlier than the time the requested data was last
modified.
24. A method according to claim 22 or 23 wherein the
request qualifier qualifies the request as a
30 deferrable request if the time indicator indicates a
time later than or equal to the time the requested
data was last modified.

- 41 -

25. A method according to any of claims 20 to 24, further comprising storing deferrable requests in a held request store in time indicator order.
- 5 26. A method according to claim 25, further comprising retrieving requests from the held request store, comparing the time indicator of the requests with the time the requested data was last modified, wherein the actionable state is reached if the time indicator indicates a time earlier than the time the requested data was last modified.
- 10 27. A method according to claim 26, wherein the step of retrieving comprises periodically retrieving requests from the held request store.
- 15 28. A method according to claim 26 or 27, wherein the step of retrieving comprises retrieving requests from the held request store in response to a notification from the server that data stored at the server has changed.
- 20 29. A request-response client-server computer network comprising a plurality of clients and a server with which the client communicate by issuing requests to read or write data, comprising:
- 25 at each client:
- a client data store holding data provided from the server; and
 - a request provider for issuing requests to read or write data to the server;
- 30 at each server:
- a server data store holding a master copy of data provided to each client;
 - a request receiver for receiving a plurality of requests from the clients;
- 35

- 42 -

- a request analyser for analysing each of the plurality of requests in turn passed from the request receiver;
- a request qualifier for qualifying each received request as either an actionable request for which action should be taken and a response issued, or a deferrable request for which action should be delayed and a response deferred until an actionable state is reached;
- a state indicator for indicating, for each deferrable request, when the actionable state is reached such that each such deferrable request is requalified as an actionable request; and
- a response indicator for indicating for each actionable request that each such actionable request should be actioned by the server and a response provided from the server.

30. A network according to claim 29, wherein the request analyser comprises a request queue store for storing requests in order of arrival of requests at the request receiver.

31. A network according to claim 29 or 30, wherein the request analyser comprises a request retriever for retrieving requests from the request queue store in order of arrival at the request receiver and for passing the requests so retrieved to the request qualifier.

32. A network according to claim 29, 30 or 31, wherein the request analyser is arranged to determine whether the request is to read or write data.

- 43 -

- 5 33. A network according to claim 32, further comprising a controller for controlling operation of the request analyser, request qualifier, state indicator and response indicator, and for optionally suspending operation of each, if the request analyser determines that the next request is to write data, until the data has been written.
- 10 34. A network according to any of claims 29 to 33, wherein the request provider is arranged to issue requests to read data from the server data store, each request including a time indicator indicative of when the last transaction providing a response from the server to the client occurred.
- 15 35. A network according to any of claims 29 to 34, wherein the request analyser further comprises means for extracting from each request the time indicator.
- 20 36. A network according to claim 35, wherein the time indicator is a time stamp.
- 25 37. A network according to claim 35 or 36, wherein the request qualifier is arranged to compare the extracted time indicator with the time the requested data was last modified in the server data store.
- 30 38. A network according to claim 37, wherein the request qualifier is arranged to qualify a request as an actionable request if the time indicator indicates a time earlier than the time the requested data was last modified in the server data store.
- 35 39. A network according to claim 37 or 38, wherein the request qualifier is arranged to qualify the request as deferrable if the time indicator indicates a time

- 44 -

later than or equal to the time the requested data was last modified.

- 5 40. A network according to any of claims 35 to 39, wherein the state indicator comprises a held request store for storing deferrable requests in time indicator order.
- 10 41. A network according to claim 40, wherein the state indicator comprises held request retriever for periodically retrieving requests from the held request store, the qualifier being arranged to qualify by comparing the time the requested data was last modified with the time indicator, the state indicator arranged to indicate that an actionable state is reached if the time indicator indicates a time older than the time the requested data was last modified.
- 15 42. A network according to claim 34, wherein each client is arranged to repeatedly issue requests to read data from the server data store and to write the data received to the client data store.
- 20 43. A network according to claim 42, wherein the client data store and server data store hold structured data.
- 25 44. A network according to claim 43, wherein the structured data in the client data store forms a client database, and the time indicator is indicative of when the last transaction providing a response from the server to the client occurred.
- 30 45. A network according to claim 43 or 44, wherein the structured data in the server data store forms a
- 35

- 45 -

server database, and each portion of the server database includes a time field indicating the time at which the portion of data was last modified.

5 46. A network according to claim 42, wherein the client data store and server data store hold unstructured data.

10 47. A network according to claim 42, wherein the unstructured data is a page of HTML data.

15 48. A computer program product for controlling a computer in a request-response client-server computer network in which a plurality of clients communicate with a server by issuing requests to read or write data, comprising a recording medium readable by the computer; and means recorded on the recording medium for directing the computer to:

- 20 - receive a plurality of requests from the clients at a receiver;
- analyse each of the plurality of requests passed from the request receiver in a request analyser;
- 25 - qualify each request in a request qualifier as either an actionable request for which action should be taken and a response issued, or a deferrable request for which action should be delayed and a response deferred until an actionable state is reached;
- indicate from a state indicator, for each deferrable request, when the actionable state is reached such that each such deferrable request is requalified as an actionable request; and
- 30 - indicate from a response indicator for each actionable request to the server that each such actionable request should be actioned and a response provided.
- 35

- 46 -

49. A computer program product according to claim 48, wherein the step of analysing further comprises storing requests in a request queue store in order of arrival of the requests at the request receiver.
- 5
50. A computer program product according to claim 48, or 49, wherein the step of analysing further comprises retrieving requests from the request queue store in order of arrival at the request receiver and passing the requests so retrieved to the request qualifier.
- 10
51. A computer program product according to any of claims 48, 49 or 50, wherein the step of analysing further comprises determining whether the request is to read or write data.
- 15
52. A computer program product according to claim 51, the means further comprising a control subroutine for controlling operation of the request analyser, request qualifier, state indicator and response indicator with a controller, and suspending operation of each, if the request analyser determines that the next request is to write data, until the data has been written.
- 20
- 25
53. A computer program product according to any of claims 48 to 52, wherein the step of analysing each request comprises extracting from each request a time indicator indicative of when the requesting client last received a response from the server.
- 30
54. A computer program product according to claim 53, wherein the time indicator is a time stamp.

- 47 -

55. A computer program product according to claim 53 or 54 wherein, for a request to read data, the request qualifier compares the extracted time indicator with the time the requested data was last modified.
- 5 56. A computer program product according to claim 55, wherein the request qualifier qualifies the request as an actionable request if the time indicator indicates a time older than the time the requested data was last modified.
- 10 57. A computer program product according to claim 55 or 56 wherein the request qualifier qualifies the request as a deferrable request if the time indicator indicates a time younger than the time the requested data was last modified.
- 15 58. A computer program product according to any of claims 53 to 57, further comprising storing deferrable requests in a held request store in time indicator order.
- 20 59. A computer program product according to claim 58, further comprising periodically retrieving requests from the held request store, comparing the time indicator of the requests with the time the requested data was last modified, wherein the actionable state is reached if the time indicator indicates a time older than the time the requested data was last modified.
- 25 60. A system according to any of claims 1 to 14, wherein the requests and responses are HTTP requests and responses.
- 30

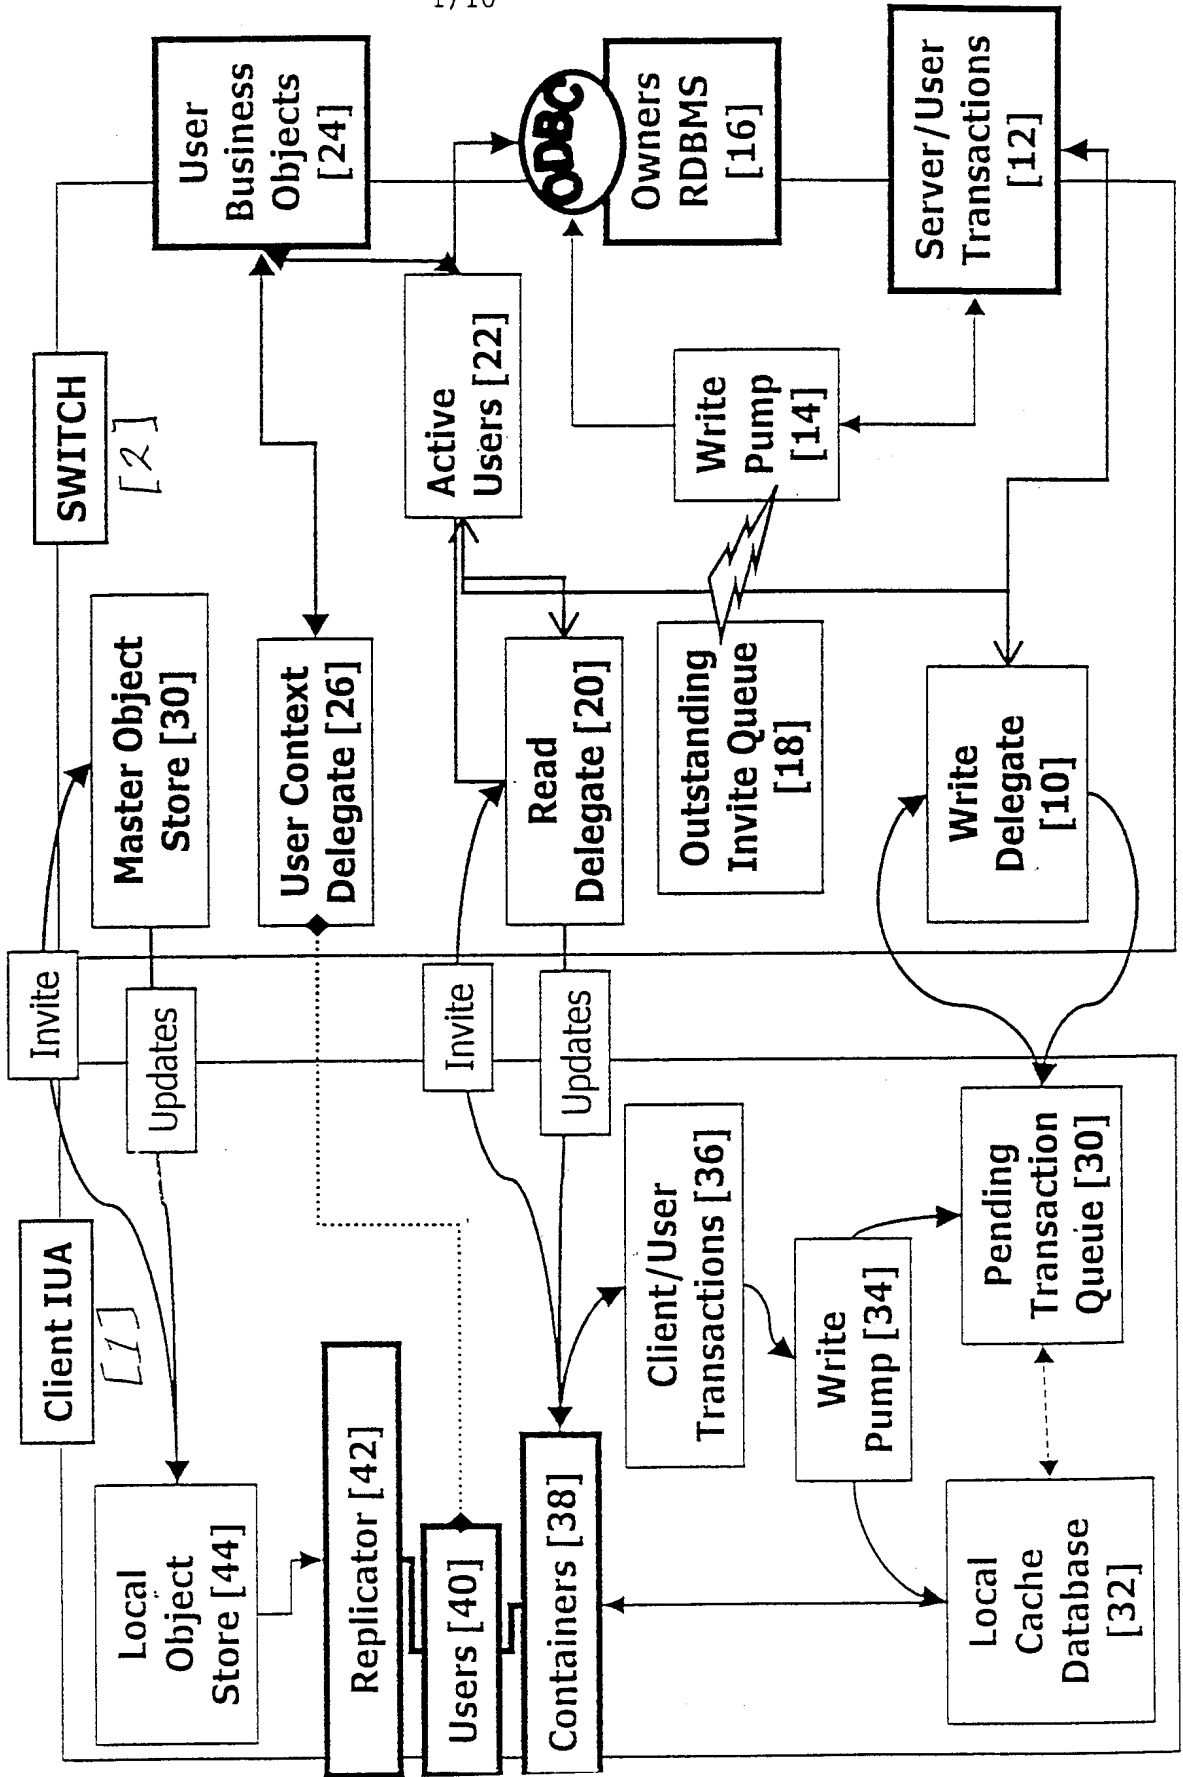
- 48 -

61. A method according to any of claims 15 to 28, wherein the requests and responses are HTTP requests and responses.

5 62. A network according to any of claims 29 to 47 wherein the requests and responses are HTTP requests and responses.

10 63. A computer program product according to any of claims 48 to 59 wherein the requests and responses are HTTP requests and responses.

FIGURE 1






FIGURE 1A




FIGURE 2

FIGURE 3




FIGURE 4




FIGURE 5




FIGURE 6




FIGURE 7





FIGURE 8






FIGURE 9





FIGURE 10

FIGURE 11





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.

INTERNATIONAL SEARCH REPORT

International Application No

PCT/GB 00/00487

A. CLASSIFICATION OF SUBJECT MATTER

IPC 7 G06F17/30 G06F17/60 G06F9/46

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 G06F H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

INSPEC, EPO-Internal, WPI Data, IBM-TDB

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>WURMAN P. R., WALSH W. E., WELLMAN M. P., O'MALLEY K. A.: "A CONTROL ARCHITECTURE FOR FLEXIBLE INTERNET AUCTION SERVERS" UNIVERSITY OF MICHIGAN, ARTIFICIAL INTELLIGENCE LABORATORY, 'Online! 6 February 1999 (1999-02-06), pages 1-12, XP002145391 Retrieved from the Internet: <URL:http://ai.eecs.umich.edu/people/wew/Papers/AB_Architecture2.ps.Z> 'retrieved on 2000-08-09! page 5; figure 2 page 7, line 1 - line 18 page 8, line 24 - line 27 page 9, line 8, paragraph 3.3.1 -page 10, last line</p> <p style="text-align: center;">---</p> <p style="text-align: center;">-/--</p>	<p>1-3, 6-10, 15-17, 20-24, 29-31, 34-39, 42-50, 53-57, 60-63</p>

Further documents are listed in the continuation of box C. Patent family members are listed in annex.

° Special categories of cited documents :

<p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"E" earlier document but published on or after the international filing date</p> <p>"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p>	<p>"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention</p> <p>"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone</p> <p>"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.</p> <p>"&" document member of the same patent family</p>
--	--

Date of the actual completion of the international search	Date of mailing of the international search report
22 August 2000	05/09/2000

Name and mailing address of the ISA	Authorized officer
<p>European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl, Fax: (+31-70) 340-3016</p>	<p>Archontopoulos, E</p>

INTERNATIONAL SEARCH REPORT

International Application No
PCT/GB 00/00487

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>WURMAN P. R., WELLMAN M. P., WALSH W. E.: "THE MICHIGAN INTERNET AUCTIONBOT: A CONFIGURABLE AUCTION SERVER FOR HUMAN AND SOFTWARE AGENTS" PROCEEDINGS OF THE SECOND INTERNATIONAL CONFERENCE ON AUTONOMOUS AGENTS (AGENTS-98), MINNEAPOLIS, MN, USA, 'Online! May 1998 (1998-05), pages 301-308, XP002145392 Retrieved from the Internet: <URL:ftp://ftp.eecs.umich.edu/people/wellman/agents98wurman.ps.Z> 'retrieved on 2000-08-09! page 2, column 2, line 16 - line 29 page 3, left-hand column, line 7 - line 20 page 4, left-hand column; figure 2 page 4, left-hand column, line 20 -right-hand column, line 5 page 5, right-hand column, line 38 - line 42 page 6, right-hand column, line 18 - line 33 page 7, column 1, line 1, paragraph 7 - line 3</p> <p style="text-align: center;">---</p>	<p>1-3, 6-10, 15-17, 20-24, 29-31, 34-39, 42-50, 53-57, 60-63</p>
A	<p>"EFFECTIVE LOCKING SCHEME FOR REAL-TIME APPLICATIONS" IBM TECHNICAL DISCLOSURE BULLETIN, US, IBM CORP., NEW YORK, vol. 36, no. 6B, 1 June 1993 (1993-06-01), pages 319-320, XP000377398 ISSN: 0018-8689 the whole document</p> <p style="text-align: center;">---</p>	<p>6-10, 20-24, 35-39, 53-57</p>
A	<p>US 5 649 099 A (NICHOLS DAVID A ET AL) 15 July 1997 (1997-07-15)</p> <p>column 15, line 54 -column 16, line 18</p> <p style="text-align: center;">-----</p>	<p>1-3, 6-10, 15-17, 20-24, 29-31, 35-39, 48-50, 53-57</p>

INTERNATIONAL SEARCH REPORT

information on patent family members

International Application No

PCT/GB 00/00487

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5649099	A	15-07-1997	NONE

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
31 May 2001 (31.05.2001)


PCT

(10) International Publication Number
WO 01/38984 A1

- (51) International Patent Classification⁷: **G06F 11/00,** (72) Inventors; and
H04N 7/14
- (21) International Application Number: PCT/US00/32551
- (22) International Filing Date: 29 November 2000 (29.11.2000)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
60/167,873 29 November 1999 (29.11.1999) US
60/170,388 13 December 1999 (13.12.1999) US
- (71) Applicants (for all designated States except US): **FUTURE TV TECHNOLOGIES, LTD.** [GB/GB]; Dublin (IE). **RZUCIDLO, Eugene, C.** [US/US]; Greenberg Traurig LLP, 885 Third Avenue, New York, NY 10022 (US).
- (72) Inventors; and
(75) Inventors/Applicants (for US only): **BROWN, Julian** [GB/GB]; 19 Ascham Road, Cambridge CB4 2BD (GB). **RAND, Ricky** [GB/GB]; Orchard House, 40 Barrington Road, Foxton, Cambridge CB2 6SJ (GB). **CLARK, Paul** [GB/GB]; The Old Mill, Mount Hawk, Truro, Cornwall TR4 8BL (GB).
- (74) Agents: **RZUCIDLO, Eugene, C.** et al.; Greenberg Traurig LLP, 885 Third Avenue, 21st Floor, New York, NY 10022-4834 (US).
- (81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, 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, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

[Continued on next page]

(54) Title: A SYSTEM AND METHOD FOR MAINTAINING FAULT TOLERANCE WHEN DELIVERING MEDIA ON DEMAND



(57) Abstract: A system and method is disclosed to deliver media on demand over a network. In particular, the present invention utilizes replication and redundancy to provide inherent fault tolerance to a subscriber of a media on demand system over a network. A plurality of copies of a subscriber selected media content are provided on a network. When a subscriber selects a particular selection of media content to use, a first copy of the selected media content is transmitted to the subscriber. This transmission is monitored for interruptions (200). If an interruption is detected (205), the point of interruption of a second copy of the subscriber selected media content on the network is determined (225) subsequently transmitted to the subscriber over the network, commencing from the point of interruption in the transmission of the first copy of the subscriber selected media content.



WO 01/38984 A1



(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

— *Before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments.*

Published:

— *With international search 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.

**A SYSTEM AND METHOD FOR MAINTAINING FAULT TOLERANCE
WHEN DELIVERING MEDIA ON DEMAND**

CROSS-REFERENCE TO RELATED APPLICATION

This application claims priority from U.S. provisional applications numbers 60/170,388 filed December 13, 1999 and 60/167,873 filed November 29, 1999, the disclosures of which are incorporated herein by reference.

FIELD OF THE INVENTION

This invention relates generally to a system and method to deliver media on demand over a network. In particular, the present invention utilizes replication, redundancy and/or rate pacing to provide inherent fault tolerance to a subscriber of a media on demand system over a network.

BACKGROUND OF THE INVENTION

Consumer demand for enhanced on-site digital entertainment, information and communication services are growing rapidly. These services, colloquially referred to as media on demand services, encompass many prior audio and video services, and may include: basic television service, AM/FM broadcast radio, premium television and video service, PPV programming, video-on-demand (VOD), audio-on-demand (AOD), interactive video and television (IVT), video games and other entertainment programs, educational information and programs, scientific and other database research, "home-shopping", infomercials, Internet access and the like.

Media on demand holds out the promise that almost every movie or musical composition ever made will be available to a subscriber of the media on demand service. Instead of driving to a video rental or music store to select a movie or recorded musical work to play, the

subscriber will be able to select any movie or musical composition stored in the systems media server library, and have that movie delivered to them over traditional distribution networks (i.e. cable systems, telephone networks, satellite systems, etc.). The use of various distribution networks to delivery media content to subscribers is well known in the art.

U.S. patent No. 5,905,522 to Lawler, U.S. patent No. 5,675,738 to Suzuki et al., U.S. patent No. 5,629,732 to Moskowitz et al., U.S. patent No. 5,606,359 to Youden et al., U.S. patent No. 5,790,174 to Richard, III et al.; and U.S. patent No. 5,550,577 to Verbiest et al., describe various media on demand services.

A fundamental aspect of a media on demand system is the provision of a large amount of data from a data storage system to many users almost simultaneously without significant processing time or delay. In addition, another aspect of a media on demand system is to provide inherent fault tolerance against delivery interruptions of the media content provided to subscribers.

SUMMARY OF THE INVENTION

The present invention, a system and method for maintaining fault tolerance when delivering media on demand through redundancy, replication and/or rate pacing is directed to system and method that satisfies the need to provide media on demand subscribers with reliable media replete with inherent fault tolerance.

In a preferred embodiment of the present invention, the instant method comprises the steps of: (i) providing a plurality of copies of media content on the network; (ii) providing a method for determining the location of a plurality of a particular selected media content on the network; (iii) providing a first copy of the particular

selected media content to a subscriber through a network transmission; (iv) providing a method for detecting an interruption of the transmission of the particular selected media content to the subscriber; (v) selecting a second copy of the particular selected media content; and (vi) transmitting the second copy of the particular selected media content to the subscriber over the network.

According to the present invention, the overall network is completely flexible in topology, and may consist of a number of differently sized arrays of clusters or individual clusters, connected together in any convenient manner, geographically distributed if required. In a preferred embodiment of the invention, each array of clusters is treated as an independent region in a hierarchical organization, and advertises itself as a unit to the higher-level network. See provisional patent application to Rand et al., Application No. 60/155,388, entitled "A System and Method for Large-Scale, Distributed, Personalized Media on Demand" for a detailed description of the network architecture.

Once a media subscriber selects an item of media content to receive, a request for the selected item of media content is made to a network manager. The selection may be made by a location-independent name.

The network manager determines at least one media server holding a copy of the selected media content out of a plurality of possible media servers using the location-independent name. Determining the location of the media content comprises sending a request to a first mapping server to look up a media server holding a copy of the selected item of media in a map held by the first mapping server. Each mapping server maintains a list of

all media content in its region, and the distance (hops) of the media content from the network manager. If known at the first mapping server, the request is responded to directly with the address of the first media server holding a copy of the selected media content. If not known at the first mapping server, the request is delegated to a second mapping server. This process is continued until a mapping server determines the location of the selected media on a media server and the distance (hops) of the media content from the network manager. The network manager then directs the closest media server maintaining the selected media content to deliver the media content to the subscriber.

If during the presentation of the selected media to the subscriber the media delivery stream is interrupted, the network manager is notified of this event and determines the point in time the media server ceased delivery of the media content stream. A request is then made to the network manager to find a different copy of the selected media.

Similar to above, the network manager determines the location of the selected media content by sending a request to a first mapping server to determine a media server holding a different copy of the selected media content. If known at the first mapping server, the request is responded to directly with the address of the next closest media server holding the different copy of the selected media content. If not known at the first mapping server, the request is delegated to a second mapping server. This process is continued until a mapping server determines the location of the different copy of the selected media content on a media server. The media server is then directed to deliver the closest different copy of the selected media content to the

requesting subscriber at the exact point in time of the media content stream that the original copy of the media content ceased delivery.

In one embodiment of the invention, the change from one media source to another media source appears seamless with the addition of a buffer storage system. In another embodiment of the invention, the buffer storage system employs rate pacing with a feedback loop to assist in replenishment of the buffer storage system.

BRIEF DESCRIPTION OF THE DRAWINGS

FIGURE 1 shows an illustrative block diagram depicting a minimal embodiment of a media on demand system in accordance with the invention.

FIGURE 2 shows a block diagram depicting the fault tolerance method according to one embodiment of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

A minimal embodiment of a media on demand system in accordance with one embodiment of the present invention is shown in Figure 1. This minimal embodiment is shown for the purpose of clarity. However, the system comprising the present invention is designed to scale well beyond the size depicted here. Figure 1 shows two clusters, C1 and C2 connected by packet switch network N1. There are common elements and differences in each network. Common elements include packet switches PS1 and PS2, and cluster or network managers CM1 and CM2.

Packet switches PS1 and PS2 are high-bandwidth packet switches such as are known in the art, and provide internal links within each cluster and also inter-cluster links via packet switch network N1.

Cluster manager unit CM1 is an industrialized computer hosting at least mapping server MAP1. Cluster

manager unit CM2 is an industrialized computer hosting at least mapping server MAP2.

Mapping server MAP1 obtain maps of content availability from media servers MS11 and MS12 connected to local packet switch PS1. Similarly, mapping server MAP2 obtain maps of content availability from media servers MS21 and MS22 connected to local packet switch PS2. Mapping servers MAP1 and MAP2 then broadcast advertisements of the content and services resident on their media servers to corresponding mapping services in other clusters in network region 1 through a flooding process such as is known in the art for routing topology advertisement. Conversely, on receipt of content and service advertisements from other clusters, mapping servers MAP1 and MAP2 generate a global map of the availability of content and services for use by the cluster managers CM1 and CM2 respectively.

Media servers MS11, MS12, MS21 and MS22 are units which store items of media content as files on hard disks or other memory devices, and which can be instructed to deliver the items of media content as a plurality of interleaved, substantially regular streams of data onto a network. Such units are well known in the prior art, and may be constructed from commonly available computer equipment or custom hardware as is required.

Cluster C1 also provides a cache media server CMS13 through which items of media content are directed the first time they are requested, and from which they can be served directly without using external network bandwidth on subsequent requests from subscribers on the same cluster.

In cluster C1, subscriber access network SAN1 may be one of a number of mechanisms for distribution of high-speed data to residential or business subscribers, well

known in the art, including but no limited to ADSL over traditional telephony and QAM/MCNS over traditional community antenna television networks.

The subscriber interface is set top box STB1. A minimal embodiment of STB1 comprises: a microprocessor, random access memory and non-volatile storage for boot software; a network interface connected to subscriber access network SAN1; an input device to input commands to STB1; and a media player output to play audio and/or video on a media player, such as a television or stereo.

According to the present invention, the overall network is completely flexible in topology, and may consist of a number of differently sized arrays of clusters or individual clusters, connected together in any convenient manner, geographically distributed if required. See provisional patent application to Rand et al., Application No. 60/155,388, entitled "A System and Method for Large-Scale, Distributed, Personalized Media on Demand" for a further description of the network.

Once a media subscriber selects an item of media content to receive, a request for the selected item of media content is made to a network cluster manager. The selection may be made by a location-independent name. The system and method of selecting and delivering media on demand to a subscriber is fully described in a provisional patent application to Rand et al., Application No. 60/155,388, entitled "A System and Method for Large-Scale, Distributed, Personalized Media on Demand" which is incorporated by reference herein.

Initial Delivery of Media

A subscriber provides input to set top box STB1 for a selection of media content. Set top box STB1 transmits this selection to cluster manager CM1 through subscriber access network SAN1 and packet switch PS1. Cluster

manager CM1 transmits this selection to mapping server MAP1 and receives a list of all available media responsive to the subscriber's selection in the region, and the distance (hops) from cluster manager CM1 to the media server maintaining the requested media. In a preferred embodiment of the invention, the list for multiple copies of the selected media is maintained in hierarchical order, from the closest media server to the furthest media server. In another preferred embodiment of the invention, all copies of the selected media equidistant from the cluster manager are arranged in a round-robin fashion.

If the requested media content is known at mapping server MAP1 (i.e. in region 1), the request is responded to directly with the address of the first media server holding a copy of the selected media content. If not known at mapping server MAP1 (i.e. not within region 1), cluster manager CM1 communicates the subscriber's selection to a cluster manager (not shown) outside of the region through packet switch work N1. This process is continued until a mapping server determines the location of the selected media on a media server and the distance (hops) of the media content from cluster manager CM1. The cluster manager for the mapping server determining the location of the subscriber's selected media content provides this information to cluster manager CM1 through packet switch network N1.

For the purpose of example, it is assumed that mapping server MAP1 determines the following locations for the subscriber's media request made through cluster manager CM1, and communicates this information to cluster manager CM1.

RANKING	LOCATION OF SUBSCRIBER'S	DISTANCE (HOPS) TO LOCATION
---------	--------------------------	-----------------------------

	REQUEST	
1	MS-11	1
2	MS-21	2
3	MS-22	3

The requesting cluster manager then directs the closest media server maintaining the selected media content to deliver the media content to the subscriber through the set top box. The transmission of the media content is originally at a higher rate than the normal subscriber user rate to fill a buffer located in the set top box. The transmission rate of the media content then slows as the buffer in the set top box fills. The request is made through the cluster manager for the cluster maintaining the selected media content.

The cluster manager then communicates the origination (location) of this media stream to the subscriber's set top box receiving the selected media content stream. Knowing the origination of the selected media stream, the set top box can then communicate directly with the media server for the purpose of controlling media delivery operations, such as fast-forward, rewind, stop and play, and provide dynamic rate pacing to facilitate fault tolerance in the delivery rate.

By way of the above referenced example, cluster manager CM1 communicates with media server MS-11 and requests the subscriber's selected media content to be delivered to set top box STB1 through subscriber access network SAN1. Cluster manager CM1 also communicates with set top box STB1 and provides the dispatch location (MS11) for the subscriber's selected media.

The transmission of the media is monitored and tracked as some multiple of a finite unit to determine

the amount of media transmitted. Any number of unit types may be used to monitor the media transmitted. In one embodiment of the invention, the aggregate transmission or subscriber receive time (in seconds, minutes, etc.) is used to monitor and track the selected media stream units. In another embodiment of the invention the frames transmitted or received (particularly for video applications) are used to monitor and track the selected media stream units. In still a further embodiment of the invention, the bytes of data transmitted or received are used to monitor and track the selected media stream units.

Any number of equipment forming the network can monitor the transmission of the requested media content. In a preferred embodiment of the invention, set top box STB1 monitors the transmission of the requested media content. In another embodiment of the invention, cluster manager CM1 monitors the transmission of the requested media content. In still a further embodiment of the invention, media server MS-11 monitors the transmission of the requested media content

In a preferred embodiment of the invention, set top box STB1 includes a buffer which stores a finite quantity of transmission units of the selected media content.

Fault Tolerance

A block diagram showing a method to maintain fault tolerance according to one embodiment of the invention is shown in figure 2. The illustrated method monitors a media stream being transmitted to the subscriber as shown in step 200. If set top box STB1 detects an interruption in the selected media content data stream originating from media server MS11 as shown in step 205, it automatically transmits a stop command to media server MS11, followed by a momentary rewind command and a play

command, as shown in step 210. If the set top box STB1 still detects an interruption in the selected media content data stream, as shown in step 215 it: (i) notifies cluster manager CM1 of the interruption, step 220; (ii) provides the interruption point, i.e. unit location in the transmission of the interruption, to cluster manager CM1, step 225; and (iii) requests another copy of the selected media content to cluster manager CM1, step 230. The unit location of the transmission interruption is provided as a multiple of the units used to monitor and track the transmission. In one embodiment of the invention, the units are a function of the selected media transmitted by the media server. In another embodiment of the invention, the units are functions of the selected media received by the set top box.

Cluster manager CM1 then communicates with mapping server MAP1 to determine the location of the next media server maintaining a copy of the selected media content. Once prompted by cluster manager CM1, mapping server MAP1 determines the next closest media server maintaining a copy of the selected media content. For the purpose of the above example, mapping server MAP1 determines that media server MS21 maintains the next closest copy of the selected media content. Cluster manager CM1 then communicates with cluster manager CM2 and requests that a copy of the selected media content be delivered from media server MS21 to set top box STB1 commencing at the unit location in the transmission that the interruption was detected. Cluster manager CM1 also communicates with set top box STB1 and provides the origination of the new selected media content. With the origination location, set top box STB1 has the ability to communicate directly

with media server MS21 to perform operational tasks such as fast-forward, rewind, play and stop.

In a preferred embodiment of the present invention, the steps of detecting the interruption by set top box STB1; notifying cluster manager CM1 of the interruption; determining the location of the second copy of the selected media content by mapping server MAP1; and transmitting the second copy of the selected media content from media server MS21 to the subscriber through set top box STB1, is completed prior to or immediately after the buffer in set top box STB1 is completely under-run.

In another embodiment of the invention, cluster manager CM1 directs cluster manager CM2 to cause a transmission of the selected media content from media server MS21 to cache media server CMS13. Cache media server CMS13 in turn transmits the selected media content to the subscriber via subscriber access network SAN1 and set top box STB1. The transmission between media server MS21 and cache media server CMS13 is made at a greater rate than the transmission from cache media server CMS13 and set top box STB1 to allow for buffering or caching in cache media server CMS13. In a preferred embodiment of the invention, cluster manager CM1 has discretion to instruct cache media server CMS13 to replicate and save the copy of the selected media content for future transmission to subscribers. Caching and replication may continue regardless of subsequent action by the subscriber.

Fault Tolerance Through Dynamic Rate Pacing

STB1 maintains a top pointer, which indicates where in the buffer storage system the next frame of data to be displayed begins. STB1 also maintains a bottom pointer, which indicates where in the buffer the next packet of

data received from the media server MS11 is to be appended. The top pointer is decremented by the frame size after data has been displayed on the monitor, while the bottom pointer will be decremented by the packet size after a packet has been received from the media server MS11. The adjustments of the pointer values will be performed modulo the buffer size. Thus, STB1 contains means to permit it to determine when the media server MS11 is sending data faster than STB1 is displaying it (overwriting the buffer), or when the media server MS11 is not sending data fast enough to refresh the display (underwriting the buffer). Because both the packet size and the frame size are fixed, STB1 can prevent an overwrite or underwrite condition, since it is able to signal the media server MS11 via the subscriber access network SAN1 to adjust the data delivery rate. STB1 calculates the amount by which the delivery rate is to be adjusted, and transmits that rate adjustment to the media server MS11.

When STB1 determines that the media server's MS11 present transmission rate will eventually overwrite or underwrite the buffer, STB1 signals to the media server MS11 to either slow or accelerate the packet transmission by some number of frames per second. Factors to be used in calculating whether and how much to speed or slow the transmission rate can include the transmission rate, the estimated time to an overrun or an underrun, the buffer size, and the rate at which data is displayed. The methods for calculating this adjustment are well known to those of ordinary skill in the art. An additional factor that can be included in this adjustment is how much data is to be maintained in the buffer storage system. For example, in one embodiment, in which the buffer has a capacity of 3 seconds of data, STB1 can maintain an

average of 1.5 seconds of data by signaling the media server MS11 to speed up transmission whenever the amount of data to be displayed falls below one second's worth, and to delay transmission whenever the amount of data exceeds two second's worth. In a preferred embodiment, these upper and lower size limits may be adjusted dynamically to allow for more precise control of the data transmission rate. STB1 communicates the rate adjustment to the media server MS11 via the previously mentioned communications channel over the SAN1.

In a representative system as previously described, when a stream is initially requested from the media server MS11 by STB1, it would normally be requested at a high rate of transmission, for example at 6 Mbits/sec, to load the buffer. In response to the request, the media server MS11 transmits the stream at the accelerated rate of 6 Mbits/sec in an effort to establish the buffer. In theory, to transmit a 6 Mbits/sec (6 megabits per second), the media server MS11 would transmit 1000 bytes of data every $10 \frac{2}{3}$ packet-starts (8,000 packets of 1000 bytes every second = 64 megabits per second total bandwidth). Since the system limits the media server MS11, however, to transmitting data only at the packet-start, the media server MS11 cannot transmit on $\frac{2}{3}$ of a packet start. Accordingly, the media server MS11 would transmit at that average rate, but in practice it could transmit on the 10th, the 21st, the 32nd, the 42nd, the 53rd, and then the 64th packet starts. It is apparent that in 64 packet starts, or 8 milliseconds, 6000 bytes, or 48,000 bits are transmitted. Transmitting 48,000 bits every 8 milliseconds is equivalent to the rate of 6 megabits per second. When STB1 has determined that the buffer is sufficiently full, it signals the media server MS11 to slow to 3 Mbits/sec. To maintain 3 Mbits/sec,

the media server MS11 would slow to transmitting a packet averaging every 21 1/3 packet-starts.

Due to clock drift and/or possibly other factors, even though the media server MS11 is transmitting packets every 21 1/3 packet-starts, the data may be arriving faster than it is being used at STB1. For example, if STB1 were displaying frames at the rate of 29.999 per second to the user due to an inaccuracy in its clock, the buffer would eventually overrun due to receiving one more frame every ten-thousand than it is displaying. Accordingly, when STB1 calculates that the data rate is too fast for the current conditions, it can request that the media server MS11 slow down, and, in fact, can specify the number of bits per second it desires to receive. The media server MS11 will respond to the request.

Although the present invention has been described in relation to particular preferred embodiment thereof, many variations and modification and other uses will become apparent to those skilled in the art. It is preferred, therefore, that the present invention be limited not by the specific disclosure herein, but only by the appended claims.

What is claimed is:

1. A method to provide inherent fault tolerance to a subscriber of media content over a network, said method comprising the steps of:

(a) providing a plurality of copies of the subscriber selected media content on the network;

(b) transmitting a first copy of the subscriber selected media content to the subscriber over the network;

(c) detecting an interruption of the transmission of the subscriber selected media content to the subscriber;

(d) determining the point of interruption in the transmission of the subscriber selected media content;

(e) determining the location of a second copy of the subscriber selected media content on the network; and

(f) transmitting the second copy of subscriber selected media content to the subscriber over the network, the transmission of the second copy of media content commencing from the point of interruption in the transmission of the first copy of the subscriber selected media content.

2. The method of claim 1 wherein the step of providing a plurality of copies of the subscriber selected media content further comprises the steps of:

(a) locating a first copy of the subscriber selected media content on the network; and

(b) transmitting the first copy of the subscriber selected media content to a caching storage device to create a second copy of the subscriber selected media content, the transmission to the caching storage device being at a transmission rate greater than a rate of transmission to the subscriber.

3. The method of claim 1 wherein the step of transmitting a first copy of the subscriber selected media content to the subscriber further comprises the steps of tracking the transmission of the media as some multiple of a finite unit to determine the amount of media transmitted.

4. The method of claim 1 wherein the step of transmitting a first copy of the subscriber selected media content to the subscriber further comprises the steps of:

(a) transmitting the first copy of the subscriber selected media content to a buffer, the transmission rate being modulated to maintain the level of subscriber selected media content in the buffer; and

(b) transmitting the first copy of the subscriber selected media content from the buffer to the subscriber.

5. The method of claim 1 wherein the step of detecting an interruption of the transmission of the subscriber selected media content further comprises the steps of:

(a) transmitting a stop command to stop the transmission of the first copy of the subscriber selected media content to the subscriber;

(b) transmitting a momentary rewind command to rewind a short portion of the first copy of the subscriber selected media content; and

(c) transmitting a play command to play institute the transmission of the first copy of the subscriber selected media content from the point of rewind.

6. The method of claim 1 wherein the step of determining the point of interruption in the transmission comprises the steps of:

(a) determining a unit location of the interruption in the transmission of the subscriber media content;

(b) notifying a media manager of the unit location of the interruption in the transmission of the subscriber media content.

7. The method of claim 1 wherein the step of determining the location of a second copy of the subscriber selected media content further comprises the steps of:

(a) generating a global map of available media content on the network responsive to the subscriber's selection;

(b) transmitting the global map of available media content on the network responsive to the subscriber's selection to a network manager;

(c) determining the distance from the location of the available media content responsive to the subscriber's selection to the subscriber.

8. A method to provide inherent fault tolerance to a subscriber of media content over a network, said method comprising the steps of:

(a) providing a first copy of the subscriber selected media content on the network;

(b) transmitting a first copy of the subscriber selected media content to a cache storage device to create a second copy of the subscriber selected media content;

(c) transmitting the second copy of the subscriber selected media content to the subscriber over the network, the transmission of the second copy of the media content to the subscriber being at a transmission rate less than the transmission of the first copy of the selected media content to the cache storage device;

(d) detecting an interruption of the transmission of the second copy of the subscriber selected media content to the subscriber;

(e) determining the point of interruption in the transmission of the second copy of the subscriber selected media content;

(f) transmitting the first copy of subscriber selected media content to the subscriber over the network, the transmission of the first copy of media content commencing from the point of interruption in the transmission of the second copy of the subscriber selected media content.

9. An apparatus for dynamically adjusting a data transmission rate between a data server and a data receiver in a network connecting at least one data server

with at least one data receiver, the server and receiver having separate control clocks, wherein the data receiver is transmitting the received data to a display device, said apparatus comprising:

(a) means for the data receiver to store data received from the data server prior to transmission to the display device;

(b) means for the data receiver to detect overwrite and underwrite conditions in the data storage means;

(c) means for the data receiver to calculate an adjustment to the data transmission rate when the detection means detects an overwrite or an underwrite; and

(d) means for communicating the data transmission rate adjustment to the data server.

10. The apparatus of claim 9, in which the data storage means comprises a digital memory buffer.

11. The apparatus of claim 10, in which the memory buffer is of a fixed size.

12. The apparatus of claim 9, in which the means for detecting overwrites or underwrites comprises a first pointer adapted to point to the beginning of the data and a second pointer adapted to point to the end of the data, wherein the relative positions of the first and second pointers are compared by said detecting means to determine whether the storage means is being overwritten or underwritten.

13. The apparatus of claim 9, in which the means for calculating an adjustment comprises a programmed micro-controller adapted to calculate a rate adjustment.

14. The apparatus of claim 9, in which the means for communicating the adjustment to the server comprises a communications channel connecting the server and the receiver.

15. The apparatus of claim 9, in which the means for communicating the adjustment to the server comprises a first communication channel connecting the receiver with a manager, and a second communications channel connecting the manager to the server.

16. An apparatus for dynamically adjusting a data transmission rate between a data server and a data receiver in a network connecting at least one data server with at least one data receiver, the server and receiver having separate control clocks, wherein the data receiver is transmitting the received data to a display device, said apparatus comprising:

(a) a digital memory buffer of fixed size for the data receiver to store data received from the data server prior to transmission to the display device, said data having a beginning and an end;

(b) a first pointer adapted to point to the beginning of the data in the memory buffer and a second pointer adapted to point to the end of the data in the memory buffer,

(c) a controller adapted to compare the relative position of the first and second pointers to detect

whether the memory buffer is being overwritten or underwritten;

(d) means for the data receiver to calculate an adjustment to the data transmission rate when an overwrite or underwrite condition has been detected; and

(e) a communications channel connecting the data server and the data receiver by which the data receiver can transmit the adjusted data transmission rate to the data server.

17. A method for a data receiver to dynamically adjust a data transmission rate from a data server in a network connecting at least one data server with at least one data receiver, the server and receiver having separate control clocks, the receiver having a data buffer of fixed size, said method comprising the steps of:

(a) storing data received from the data server in the buffer prior to transmitting the data to the display device;

(b) determining when the buffer is being overwritten or underwritten;

(c) calculating an adjustment to the data transmission rate; and

(d) communicating the adjustment to the data server.

18. The method of claim 17, wherein the data stored in the buffer has a beginning and an end, further comprising the step of positioning a first pointer to point to the beginning of said data and positioning a second pointer to point to the end of said data, such that the position

of the first pointer is compared to that of the second pointer and a determination is made as to whether the buffer is being overwritten or underwritten.

19. The method of claim 18, further comprising the steps of adjusting the position of the first pointer when data is transmitted from the buffer to the display device by the amount of data transmitted to the display device, and adjusting the position of the second pointer when data is received from the data server by the amount of data received from the server.

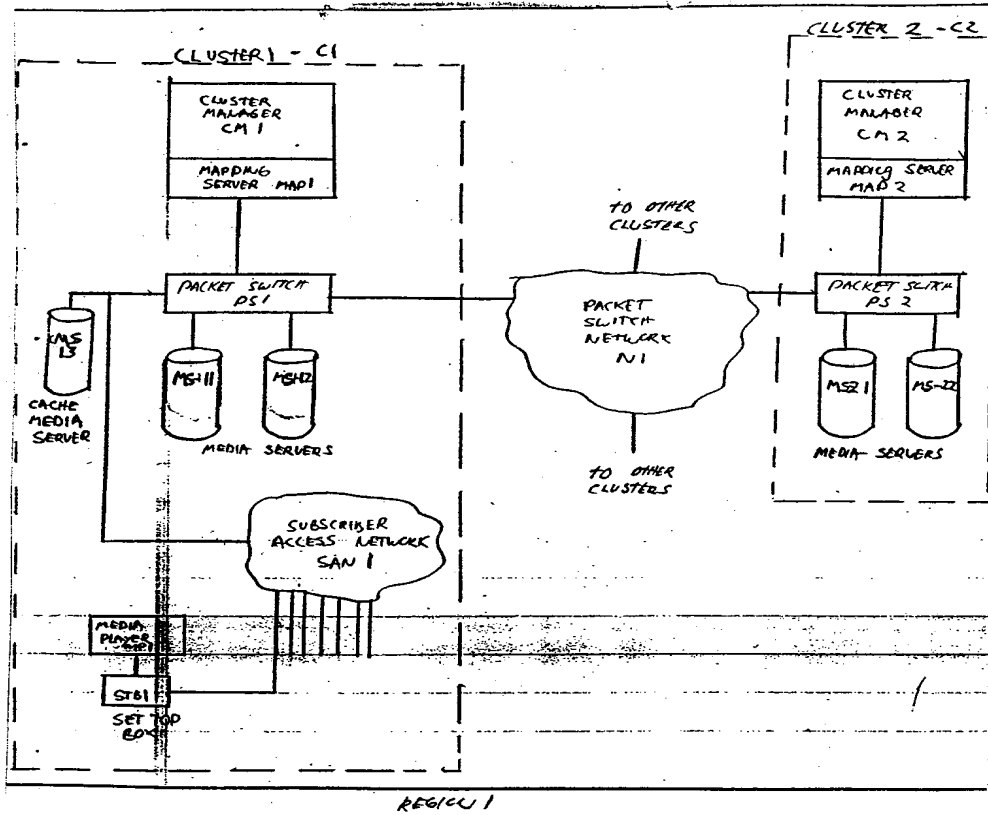
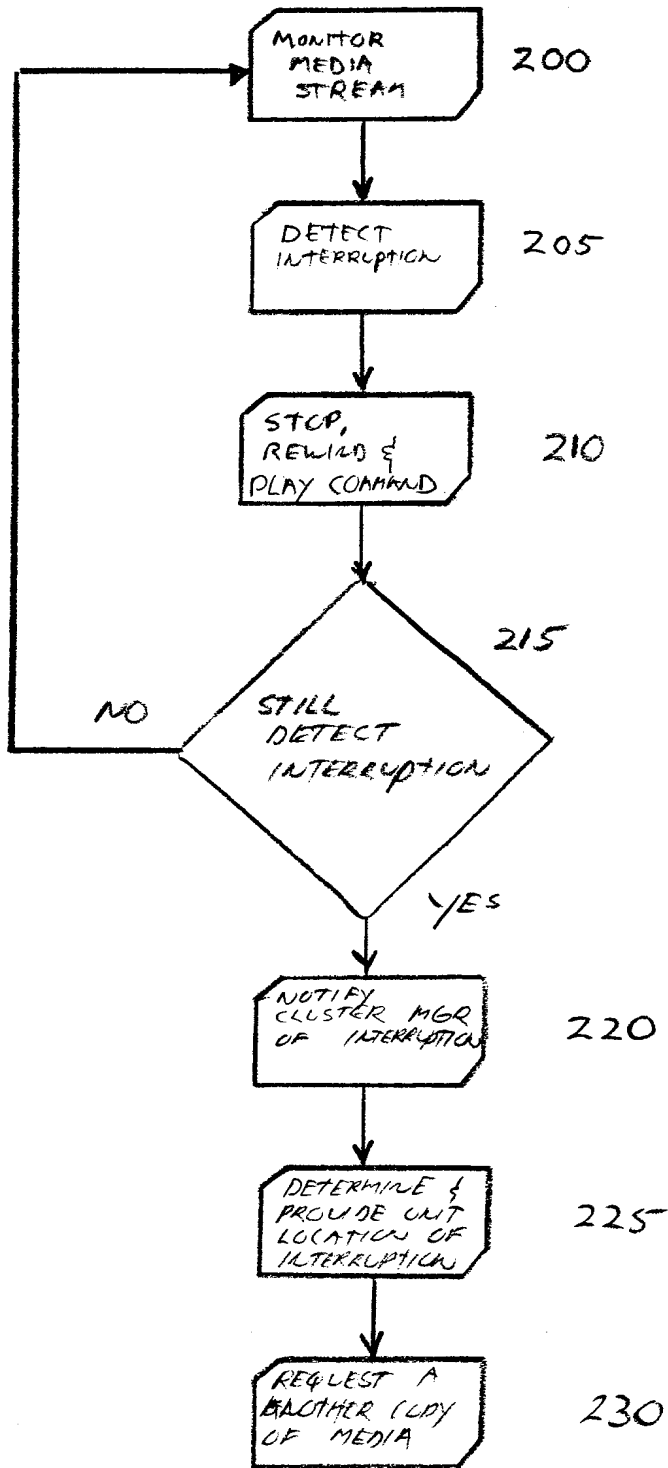


FIGURE 1

FIGURE 1

FIGURE 2



INTERNATIONAL SEARCH REPORT

International application No.
PCT/US00/32551

A. CLASSIFICATION OF SUBJECT MATTER
 IPC(7) : G06F 11/00, H04N 7/14
 US CL : 714/66, 7, 8, 10-12; 348/7, 12, 13; 709/240, 242; 710, 20, 21;
 According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED
 Minimum documentation searched (classification system followed by classification symbols)
 U.S. : 714/66, 7, 8, 10-12; 348/7, 12, 13; 709/240, 242; 710, 20, 21;

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
 WEST, JPAB, EPAB, TDBD, DWPI, PGPB

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 5,606,359 A (YOU DEN ET AL) 25 FEBURARY 1997, Entire document.	1-19
Y	US 5,699,503 A (BOLOS KY ET AL) 16 DECEMBER 1997, Entire document.	1-19
Y	US 5,815,146 A (YOU DEN ET AL) 29 SEPTEMBER 1998, Entire document.	1-19
Y,P	US 6,061,504 A (TZELNIC ET AL) 09 MAY 2000, Entire document.	1-19
Y	US 5,625,404 A (GRADY ET AL) 29 APRIL 1997, Entire document.	1-19

Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:	*T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
A document defining the general state of the art which is not considered to be of particular relevance	*X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
E earlier document published on or after the international filing date	*Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	* & * document member of the same patent family
O document referring to an oral disclosure, use, exhibition or other means	
P document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search 29 MARCH 2001	Date of mailing of the international search report 26 APR 2001
--	---

Name and mailing address of the ISA/US Commissioner of Patents and Trademarks Box PCT Washington, D.C. 20231 Facsimile No. (703) 305-3230	Authorized officer <i>Pagget Hand</i> NADEEM IQBAL Telephone No. (703) 308-5228
--	---

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
31 May 2001 (31.05.2001)

PCT

(10) International Publication Number
WO 01/38993 A1


- (51) International Patent Classification⁷: G06F 13/00
- (21) International Application Number: PCT/US00/32032
- (22) International Filing Date:
20 November 2000 (20.11.2000)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
09/450,855 29 November 1999 (29.11.1999) US
- (71) Applicant: BINARY BROADCASTING CORPORATION [US/US]; 542 Emerson Street, Palo Alto, CA 94301 (US).
- (72) Inventor: MELMON, Matthew; 20975 Valley Green Drive, #241, Cupertino, CA 95014 (US).
- (74) Agents: PISANO, Nicola, A. et al.; Fish & Neave, 1251 Avenue of the Americas, New York, NY 10020 (US).

- (81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, 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, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.
- (84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published:
— With international search 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.

(54) Title: SECURE DIGITAL MUSIC DISTRIBUTION



(57) Abstract: A portable secure digital audio player (24) and an Internet accessible service (22) for selecting music for download to the player are described. The service (22) provides access to a variety of MP3 encoded digital music. A subscriber uses the service to select a playlist, or to instruct the service to automatically generate a playlist. The service (22) may track the subscriber's musical preferences and copyright management information on the selected music. The player (24), which includes modem circuitry, connects to the service (22) either directly or through the Internet (23), identifies itself to the service using a unique identifier, and downloads the music in the playlist. Since the service may only download music to a secure player, the player may only accept music from the service, and the player cannot digitally transfer the music to any other device, the risk of piracy is reduced.



WO 01/38993 A1

SECURE DIGITAL MUSIC DISTRIBUTION

Field Of The Invention

The present invention relates generally to apparatus and methods for secure distribution of digital music. More specifically, the present invention provides a secure MP3 player and a service for providing digital music to the secure MP3 player.

Background Of The Invention

Recently, the distribution of digital music over the Internet has gained in popularity and visibility. This is due, in part, to the widespread use of the MP3 music compression format, which permits relatively high quality reproduction of music using music files that are small enough to be downloaded by many home users. A typical three minute song may be compressed into a file having a size of approximately 2.5 to 3 megabytes using MP3, while retaining sound quality comparable to the quality of music played from a compact disc. A home user connected to the Internet using a low cost modem capable of transferring 56 kilobits per second can download the song in a couple of minutes.

Numerous applications for playing digital music have been developed for use on personal

- 2 -

computers. Additionally, portable digital music playback devices have been developed that can download MP3 encoded music from a personal computer, and play the downloaded music. These portable music players

5 have numerous advantages over previously known portable music players, such as portable cassette players or portable CD players. For example, since portable digital music players typically store music in solid state memory, there are no mechanical parts that may

10 cause skipping or other audible glitches in the music when the device is used during jarring physical activities. Additionally, portable digital music players can be made very small, and are not confined to a particular form factor by the media that they play.

15 These advantages have made portable digital music players very popular, with several companies manufacturing a variety of models that play MP3 encoded music.

MP3 players are also being built for use in

20 automobiles, and as part of a home audio system. These digital music players often use a hard drive or a CD-ROM to store music. The primary advantage of such devices is their ability to store a large amount of music in a small space. For example, at approximately

25 1 megabyte per minute of CD-quality music stored in MP3 format, a single CD-ROM can store more than 10 hours of high quality music. Hard drives conforming to the size of a 3.5" drive bay on a personal computer, and having capacities exceeding 20 gigabytes are now available

30 inexpensively. Such a drive could store more than 300 hours of high quality MP3 encoded music.

The ability to distribute digital music over a network, such as the Internet, offers several

- 3 -

potential advantages to consumers and artists. Consumers could choose specific songs that they wish to purchase, pay for those songs, and download them over the Internet. There would be no need to travel to a music store, and the music could be delivered through a modem or other communication device within a few minutes. Since there is no manufacturing or packaging cost, and only minimal distribution cost involved in providing digital music, the prices charged for digital music distributed over the Internet could be lower than what is charged for music distributed by other means, while still maintaining a relatively high profit margin. Once downloaded, digital music, such as MP3 formatted music, provides the consumer a high degree of flexibility. The music may be stored in a searchable database, and played on any of a number of devices.

For artists, compressed digital music provides a good means of gaining exposure. For example, a new or unknown band may gain valuable exposure by making its music available in MP3 format for free download on the Internet. The music may be distributed through popular MP3 web sites, and may be made available on a web site dedicated to a particular style of music, or a particular band or artist.

The popularity of the MP3 digital music format may also be demonstrated by the number of MP3 files and web sites available on the Internet. Thousands (possibly tens of thousands) of MP3 files are available on hundreds of sites, with more files and sites being added daily. The term "mp3" is one of the most frequently searched terms on numerous Internet search engines and portal sites. Additionally, several companies, such as MP3.Com, Inc., of San Diego,

- 4 -

California, and EMusic.com Inc., of Redwood City, California, provide access to a wide variety of MP3 format music.

Unfortunately, a large portion of the MP3
5 encoded music available on the Internet is not being legally distributed. It is very easy for Internet users to make unauthorized copies of music off of CDs, to encode the music as an MP3 file, and to make the music available over the Internet, where unauthorized
10 copies may be made by anyone with Internet access. New web sites distributing pirated MP3 files spring up more rapidly than copyright holders are able to react to the presence of such sites.

Although some MP3 files, such as those
15 distributed by the companies mentioned above, are available legally, most of this music is by lesser known artists. Music by major artists is typically not legally available in MP3 format, or in any other format that may be distributed over the Internet. The
20 recording industry has been unwilling to make music available for Internet distribution, due in part to their concerns about piracy, content control, and marketing.

The largest concern of the recording industry
25 seems to be piracy. MP3 files provide no technological means of preventing perfect copies from being distributed illegally over the Internet. There is nothing in the MP3 format that prevents unauthorized copies from being played, that identifies the copyright
30 holders or the licensee, or that in any way insures that the copyright holders will be compensated for use of the music.

- 5 -

As a result of this lack of copyright management, the recording industry has actively opposed the wide scale adoption of MP3 as a format for distributing music over the Internet. Instead, the industry has partnered with technology companies to form the Secure Digital Music Initiative (SDMI), to develop a digital music format with acceptable copyright management. So far, there has been considerable disagreement among SDMI members, and although a proposed standard has been announced, there are not yet any digital music players that implement the SDMI format. Given the installed base of MP3 users, and the growing popularity of MP3, it is unclear whether the recording industry's efforts to impose its own music formats will be successful.

The recording industry is also concerned about issues of retaining control over music content, and over marketing issues. For example, although distribution of music over the Internet would permit songs to be sold individually, this does not match the typical way in which the recording industry distributes music. Typically, music is distributed as albums, which contain numerous songs, only one or two of which on any given album may be popular. Bundling numerous songs in albums also provides the industry with a degree of creative control over the musical content available to consumers.

Traditional recording industry methods of gaining exposure for songs and artists are also threatened by uncontrolled Internet distribution of digital music. For example, requiring that specific songs by specific artists be given air time as part of a broader licensing agreement with radio stations

- 6 -

provides the recording industry with an important means of promoting its product. Without an ability to influence the music that is "broadcast" over the Internet, the industry will lose this marketing tool.

5 The participation of the recording industry in the distribution of digital music may be vital to the future of both the recording industry, and the nascent digital music industry. Since the major record labels hold the rights to almost all popular music, and
10 have contracts with almost all popular artists that insure they will continue to hold these rights, it is unlikely that much popular music by major artists will be distributed in compressed digital format unless the concerns of the recording industry are addressed. At
15 the same time, consumers may be unwilling to pay for digital music players that place severe restrictions on the availability of digital music, and for digital music that offers no price, quality, or selection advantages over what is available on other media, such
20 as compact disks.

In addition to concerns about control and distribution, the music industry faces a problem due to the ever expanding choices being given to consumers. While highly motivated consumers are generally willing
25 to take the time to stay informed about the latest artists and bands, typical consumers do not have time to keep up with the ever expanding selection of music that is becoming available to them. The music industry faces the possibility that typical consumers will be
30 overwhelmed by the array of available choices, making purchasing decisions more difficult, and leading to decreased sales.

- 7 -

This problem is intensified by the ability of artists to make their recordings available directly to the public in digital form. Before the advent of digital music distribution, record labels and radio stations played a significant role in screening the music that would be made available. The ability to digitally distribute music, and to play music broadcasts over the Internet have decreased the ability of the recording and radio industries to keep the choices available to consumers at a manageable level.

Current methods of exposing typical consumers to the latest music, to assist their purchasing decisions, include print coverage in magazines and other sources, radio broadcasts, and live concerts. These methods are insufficient, because they tend to reach only motivated consumers, or are able to expose consumers to only a small portion of the available music.

For example, the variety and quantity of print media covering all aspects of the music world is vast, making it practically impossible for any but motivated consumers to develop a complete understanding from print sources. Similarly, only motivated consumers are willing to endure the expense and difficulty of obtaining tickets to live concerts. Although radio broadcasts are able to reach typical consumers, limited spectrum, concentration of broadcast ownership, and the constant push for higher ratings limit the content of radio broadcasts to only a small fraction of the content that must be promoted.

In view of the above, it would be desirable to provide a system through which digital music, encoded in a widely available format, may be securely

distributed to individual users through a communications network, such as the public telephone network or the Internet.

It would also be desirable to provide a portable digital music player capable of securely receiving digital music through a communications network, such as the public telephone network or the Internet.

It would further be desirable to provide a system for distributing digital music that provides payment to the owners of the distribution rights in the music, and provides the owners of the distribution rights with the ability to control the distribution of the music.

Additionally, it would be desirable to provide a system for distributing digital music that permits users to purchase licenses to download, store, and listen to digital music, to keep track of licensed music, and to download music into a portable playback device that permits the user to listen to the music.

It would further be desirable to provide a service to assist typical music consumers in making purchasing decisions by providing easy access to information about music, samples of the music of particular artists and groups, and automated music selections according to the consumer's preferences.

Summary Of The Invention

It is an object of the present invention to provide a system through which digital music, encoded in a widely available format, may be securely distributed to individual users through a

- 9 -

communications network, such as the public telephone network or the Internet.

It is also an object of the present invention to provide a portable digital music player capable of
5 securely receiving digital music through a communications network, such as the public telephone network or the Internet.

It is a further object of the present invention to provide a system for distributing digital
10 music that provides payment to the owners of the distribution rights in the music, and provides the owners of the distribution rights with the ability to control the distribution of the music.

Additionally, it is an object of the present
15 invention to provide a system for distributing digital music that permits users to purchase licenses to download, store, and listen to digital music, to keep track of licensed music, and to download music into a portable playback device that permits the user to
20 listen to the music.

It is a further object of the present invention to provide a service to assist typical music consumers in making purchasing decisions by providing
25 easy access to information about music, samples of the music of particular artists and groups, and automated music selections according to the consumer's preferences.

These and other objects of the present invention are met by providing an Internet accessible
30 service that permits a user to select music for download while keeping track of copyright management data on the selected music, and providing a portable secure digital audio player for playing the music.

- 10 -

Music from the service, typically encoded using the MP3 format, is downloaded directly into the secure digital audio player, and cannot be downloaded to other devices. The secure digital audio player is incapable
5 of downloading music from unauthorized sources, and incapable of transferring digital music to other devices once it is downloaded to the player.

Since music from the service may only be downloaded into the secure digital audio player, which
10 may not digitally transfer the music to other devices, there is a greatly reduced risk of piracy. Additionally, music may be offered on the service as single songs, or as packages which may be downloaded only as a unit, thereby providing an increased degree
15 of creative control and promotional ability to music distributors. Further control is available by restricting the number of times that a particular song may be played, the number of songs from the same album that may be selected for download, the time period over
20 which a song may be downloaded or played, or through other restrictions on the downloading and playing of digital music.

The service also offers access to a variety of articles from print media that provide information
25 on music, and that may assist a user in making purchasing decisions. Since the service may keep track of a user's downloads, it is able to determine a user's musical preferences, and may target the selection of articles displayed to the user's tastes. Additionally,
30 the service may serve a promotional purpose, by displaying featured artists, and providing access to playlists containing works by the featured artists.

- 11 -

The service may also permit a user to purchase CDs or other media containing music by the featured artist, or that is featured in the articles. Additionally, the service may provide users with an ability to purchase a license to download selected digital music into the secure digital audio player, and keep track of the licenses owned by each user of the service.

The secure digital audio player of the present invention has a unique identifier that is used to identify the player to the service, and modem circuitry that is used for establishing a direct connection between the secure digital audio player and the service. The secure digital audio player is programmed to establish a connection to a server associated with the service, identify itself to the server, and download digital music from the server. Once the music is downloaded, the secure digital audio player may be used to play the digital music.

To use the system of the present invention, a user (also referred to hereinbelow as a subscriber) uses a web browser to establish a connection with the service over the Internet. The user then uses the service to research music, and to select a playlist of music to be downloaded to the secure digital audio player.

Once the playlist has been selected, the secure digital audio player may establish a connection to the service, either directly or through the Internet, and download the music in the selected playlist into the secure digital audio player. The user may then listen to the music using the secure digital audio player as a portable music player.

- 12 -

Brief Description Of The Drawings

The above and other objects and advantages of the present invention will be apparent upon consideration of the following detailed description, taken in conjunction with the accompanying drawings, in which like reference characters refer to like parts throughout, and in which:

FIG. 1 shows an overview of the system of the present invention;

10 FIG. 2 is a diagram showing the contents of a media record stored in the digital media repository of the system of the present invention;

FIG. 3 is a diagram of the structure of a subscriber record stored in a subscriber database of the system of the present invention;

15 FIG. 4 shows display elements and relationships between display elements that may be displayed by an Internet accessible service in accordance with the principles of the present invention;

FIG. 5 is an example web browser window showing the Internet accessible service of the present invention;

25 FIG. 6 shows a portable secure digital audio player built in accordance with the principles of the present invention;

FIG. 7 shows an alternative embodiment of a portable secure digital audio player built in accordance with the principles of the present invention;

30 FIG. 8 shows the structure of the circuitry of a portable secure digital audio player built in

- 13 -

accordance with the principles of the present invention;

FIG. 9 shows the programmed routines in the program memory of a portable secure digital audio
5 player;

FIG. 10 shows a method in accordance with the present invention for downloading digital music from a server to a secure digital audio player;

FIG. 11 shows the structure of a radio
10 station playlist database for use with the present invention for automatically generating playlists; and

FIG. 12 is a flowchart of a method in accordance with the principles of the present invention for automatically generating playlists.

15 Detailed Description Of The Invention

FIG. 1 shows an overview of the system of the present invention. The system includes digital media repository 20, database 21, server 22, network 23, and secure digital audio player 24. A user uses Internet
20 access device 25 to access server 22 over network 23 to select music, and manage the playlists for his or her secure digital audio player 24. Secure digital audio player 24 may then receive music either over network 23, or through a connection to telephone
25 network 26.

Digital media repository 20 stores numerous digital music files, encoded in a well known, standard format, such as MP3. Additionally, digital media repository 20 may store articles from music magazines,
30 ratings and charts of top-selling music, and other data that may be useful in managing the system or in assisting users to make purchasing decisions. Digital

- 14 -

media repository 20 also may include copyright management information, such as who is the owner of the rights to the music, the price to be charged to users, and information on the number of times the music file
5 has been accessed. Alternatively, this copyright management information may be stored in database 21.

As described in detail hereinbelow, digital media repository 20 may also contain a variety of other information on music or other media files. This
10 information may include the name of a song, the performers, the length of the song, the album on which the song appears, its date of release, its record label, a genre (e.g. rock, jazz, alternative, country, etc.), ratings for music, and other information on a
15 music or media files that may be useful for managing the system, for generating playlists, or for assisting users in making purchasing decisions. As above, this information may be stored in database 21, instead of in digital media repository 20.

20 Digital media repository 20 may also include an area in which user music files may be stored. User files stored on the system should be free music files. Commercial music files may be identified by a number of techniques, such as encryption, digital "watermarking",
25 or other similar steganographic methods. The system may refuse to store commercial music files to which the user has not purchased an appropriate license.

Database 21 is used to track subscriber accounts on the system of the present invention. For
30 each registered subscriber of the system, a record in database 21 keeps track of personal information on the subscriber, a unique ID number associated with the subscriber's secure digital audio player (described in

- 15 -

detail hereinbelow), one or more user-selected
playlists, pointers to songs in digital media
repository 20 for which the subscriber is licensed, and
information on the subscriber's music subscriptions.
5 Additionally, database 21 may contain financial
information about the subscriber, such as a credit card
number to which new music purchases and subscriptions
are to be charged. Additionally, as described
hereinabove, database 21 may contain copyright
10 management information and other information on the
materials stored in digital media repository 20.

Database 21 may also store information on
radio stations, and the songs that various radio
stations are playing. As will be described in detail
15 hereinbelow, this information may be used to assist the
system in automatic generation of playlists for
download to secure digital audio player 24.

Server 22 provides a web-based interface to
database 21 and digital media repository 20. This web-
20 based interface may be employed by users, through
Internet access device 25, to select music, or by
secure digital audio player 24, to download music.

A user, using Internet access device 25,
which may, for example, comprise a personal computer or
25 set-top box, connects to server 22 through the
Internet. Server 22 then provides the user with access
to stored music, to playlists, to news about music and
the music industry, access to new music (either through
purchase, subscription, or free promotion), and to
30 information on featured artists or music. Server 22
may target this material to each user, displaying the
news stories, new songs, and featured artists and music
that are most likely to appeal to the user. The user

- 16 -

may use the web-based interface to server 22 to build a playlist of music or other audio content that he or she wants to download to secure digital audio player 24. Server 22 may also permit users to purchase CDs or
5 other traditional music distribution media either directly through server 22, or through music vendors who have made arrangements with the company that runs server 22.

Once the user has determined which music
10 should be downloaded to secure digital audio player 24, the user may disconnect from server 22, and may later reconnect to server 22 using secure digital audio player 24. As will be described in detail hereinbelow, secure digital audio player 24 preferably includes
15 modem circuitry, that allows it to connect to server 22 through telephone network 26. Alternatively, secure digital audio player 24 may connect to server 22 across the Internet.

In an alternative preferred embodiment, there
20 is no need to use Internet access device 25 to make a selection of music. Secure digital audio player 24 may be used in a "standalone" mode, to permit the user to select a playlist, or to request that server 22 generate a customized playlist for the user.
25 Optionally, users may also make selections using Internet access device 25, as described above. The selections that are available, and the nature of the user interface through which the user makes selections of music may depend on whether Internet access device
30 25 or secure digital audio player 24 is used to select music or other audio content from server 22.

In standalone mode, the user may connect to server 22 using secure digital audio player 24. Using

- 17 -

a user interface presented on a display of secure digital audio player 24, the user interacts with server 22 to select a playlist. Alternatively, if server 22 has sufficient data on the user's musical preferences, 5 the user may request that server 22 generate a playlist based on the user's preferences. Once a playlist has been selected or generated, server 22 may download the audio content of the playlist to secure digital audio player 24, as described hereinbelow.

10 As will be understood by one skilled in the relevant arts, use of secure digital audio player 24 to select music may decrease the selections available to the user, since the user's interactions will be limited by the user interface available on secure digital audio 15 player 24. Additionally, it may be more difficult to access news about the music industry, or other non-music content through secure digital audio player 24. This decreased ability to make selections and to access news and other non-music content is offset by 20 advantages in ease of use and mobility. Use of standalone mode permits music to be selected and downloaded without requiring use of Internet access device 25.

Once music has been selected, using Internet 25 access device 25 or secure digital audio player 24 in standalone mode, the music or other audio content must be downloaded from server 22 to secure digital audio player 24. Secure digital audio player 24 sends a unique identifier to server 22, which is used to access 30 the playlist that was compiled by the user. Server 22 then acquires the songs in the playlist from digital media repository 20, handles any required copyright management tasks, and downloads the songs to secure

- 18 -

digital audio player 24. Server 22 may optionally use the unique identifier to encrypt content before downloading content to secure digital audio player 24. This helps insure that the music being downloaded may
5 only be played on the secure digital audio player having the correct unique identifier, since the unique identifier must be used to decrypt the content. Thus, digital music content encrypted for one secure digital audio player could not be played on any other secure
10 digital audio player, or on other music playing devices.

In a preferred embodiment, before downloading music, secure digital audio player 24 may also send information such as identifiers for the songs that are
15 already stored in secure digital audio player 24, so that server 22 can download only songs in the playlist that are not already stored. Additionally, secure digital audio player 24 may send other information to server 22, such as the number of times that the user
20 has played each of the songs stored in secure digital audio player 24. This information may be useful for determining the user's preferences, and for copyright management.

It is important to note that secure digital
25 audio player 24 can only download music from server 22, or a similarly configured server. In particular, it is not possible to load music into secure digital audio player 24 without using a server that keeps copyright management information. Additionally, once music is
30 downloaded into secure digital audio player 24, it may not be transferred in digital form from secure digital audio player 24 to any other device. Further, server 22 will download music only to a properly identified

- 19 -

secure digital audio player, and may not download digital music to other devices.

If encryption based on the unique identifier is used, even if the content is intercepted during transmission, only the secure digital audio player having the proper unique identifier can decrypt and play the content. Thus, secure digital audio player 24 makes digital piracy of music much more difficult, since all music must be downloaded to secure digital audio player 24 from legitimate sources, and secure digital audio player 24 is incapable of transferring digital music.

Referring now to FIG. 2, the data stored in digital music repository 20 is described in greater detail. Digital media repository 20 preferably comprises numerous fields storing audio content, and a variety of information related to the audio content. Although the fields of digital music repository 20 are described individually, as part of a "row" in a flat database table, it will be understood by one skilled in the art that digital music repository 20 actually may comprise numerous interrelated tables in a relational database. One skilled in the relevant arts will further recognize that although the fields are described as containing data, they may instead contain pointers to data that is stored in separate tables.

As described below, each row in the table contained in digital media repository 20 comprises numerous fields. FIG. 2 shows these fields for one row of the table.

Audio content field 30 is a binary field that contains audio content, stored in MP3 format, or any other format that may digitally represent audio

- 20 -

content. Audio content field 30 is not required to contain the actual audio content for each entry in digital media repository 20, since there may be multiple references in the digital music repository to the same audio content. In such cases, audio content field 30 may simply point to a record in digital media repository 20 that contains the digital audio content.

Artist field 31 is a text field that contains the name of the individual or group that recorded the audio content. Each piece of audio content stored in digital media repository 20 should be associated with some artist or group.

Category field 32 contains a top-level category describing the audio content. Example of top-level categories include "music", "news", and "talk". In a preferred embodiment, category field 32 is a text field containing the name of the top-level category. Every piece of audio content in digital media repository 20 is categorized so that the system may provide users with an ability to choose audio content by category.

Genre field 33 contains a second-level category that describes the audio content. The available second-level categories will depend on the top-level category stored in category field 32. Examples of second-level categories for the "music" top-level category include "alternative", "rap", and "rock". Genre field 33 is preferably a text field containing the name of the second-level category, and is used by the system to permit selection of audio content by genre.

Imprint field 34 and promotion field 35 are used together to describe the record labels that

- 21 -

recorded, promoted, and distributed the audio content. Often a small record label will handle artist relationships and recording, while a major label, such as Columbia, Warner Brothers, or Universal, will handle promotion and distribution. Imprint field 34 contains the name of the "imprint" record label that recorded the audio content, while promotion field 35 contains the name of the "promotion" or "distribution" record label that promoted and/or distributed the audio content. Because the major labels sometimes handle recording, it is possible for a major label to appear in imprint field 34 for some audio content, and in promotion field 35 for other audio content. Since all audio content was recorded by some entity, imprint field 34 will typically contain data, while promotion field 35 will contain data less frequently. In a preferred embodiment, both imprint field 34 and promotion field 35 are text fields, and may be cross-checked against each other to protect against redundancy and corruption.

Album field 36 contains the name of the album on which the audio content appeared. For audio content having a top-level category of "music", album field 36 provides the name of an album. For audio content having a top-level category of "news" or "talk", album field 36 may contain the name of a show or collection from which the audio content was taken.

Single field 37 contains the name or title of the audio content. Each audio content item should have a name or title.

Magazine field 38 contains the names of one or more magazines that contain a reference to the album for the audio content. References to the album instead

- 22 -

of the particular single are used to provide an accurate picture of the overall media coverage and artist, and because magazine references to singles are less common than references to albums. Magazine field 5 38 is optional, and may not be applicable to some types of content, such as news or talk.

Issue field 39 contains the issue dates of the magazines in magazine field 38. If there are no magazines in magazine field 38, then issue field 39 10 will also be empty.

Cover field 40 identifies the artists appearing on the covers of the magazines in magazine field 38. If there are no magazines in magazine field 38, then cover field 40 will be empty. Additionally, 15 since not every magazine shows an artist on its cover, cover field 40 may be empty even if there is a corresponding entry in magazine field 38. Automatic playlist generators will favor music by artists that appear on the covers of leading magazines. Entries in 20 cover field 40 may be cross-checked against entries in artist field 31 to prevent redundancy and corruption.

Commentator field 41 contains the names of the authors of the media references in magazine field 38. If an author is not listed, commentator field 41 25 may be empty even if there is a corresponding reference in magazine field 38, or may contain a standard entry, such as "staff". Of course, if magazine field 38 is empty, commentator field 41 also will be empty.

Page field 42 is a numeric field that 30 contains the page numbers of the media references in magazine field 38. For some types of media references, page number may not be applicable, and if magazine field 38 is empty, page field 42 will also be empty.

- 23 -

Ranking field 43 contains normalized numerical rankings of an album from media references in magazine field 38, if the media references give reviews. The normalization process involves converting
5 the ranking system used in a media references into a number, ranging from one to five, with five being the best. Normalizing rankings from media references is typically performed by a human editorial staff. Since not all media references contain reviews, ranking field
10 43 may be empty. If data is present in ranking field 43, the information may be used to assist in automatically generating playlists.

Chart field 44 contains the names of the charts being used in the media references in magazine
15 field 38. Since not all media references have a chart, chart field 44 may be empty.

Position field 45 contains the positions of the audio content on the charts in chart field 44. Position field 45 may be used to assist in
20 automatically generating playlists. Chart field 45 may be empty, since charts may not always be available.

Downloads field 46 keeps track of the number of times the audio content has been downloaded to subscribers. This information may be useful in
25 copyright management, and in determining the popularity of songs with subscribers to the service. Popular songs may receive favorable treatment when the system automatically generates playlists.

Pricing field 47 contains information on one
30 or more pricing plans or restrictions that may apply to the audio content. A single piece of audio content may have multiple pricing plans. For example, a song may be available for free for a ten day promotional period,

- 24 -

or for a limited number of plays, and require a payment for further availability. It may be available at a reduced cost to particular market segments, and be available for free for three months to members of a particular subscription service. A song may have a high price for an unlimited download license, and a low price for a license for a limited time or limited number of plays. Some songs may not be available except in combination with other songs. Each of these options and restrictions would have a code entered in pricing field 47, along with the price to be charged for each option.

It will be understood by one skilled in the relevant arts that one or more of the above-described fields may be in separate tables within digital media repository 20. Additionally, aside from the audio content, the data stored in various of the fields may be stored in tables in database 21, instead of in digital media repository 20.

FIG. 3 is a diagram showing the structure of a record in database 21. The main purpose of database 21 is to keep track of subscriber accounts. Additionally, information gathered on the preferences and download habits of subscribers may be used to generate individually targeted playlists and marketing material. Generally, there will be a record in database 21 for each secure digital audio player manufactured. Database 21 may also store other information, such as portions of the data that are described above with reference to FIG. 2, or information on radio stations, as described hereinbelow.

- 25 -

It will be understood that although database 21 is described as a single flat table containing numerous records, database 21 may comprise numerous interrelated tables in a relational database.

5 Additionally, fields in a table in database 21 may contain data, or may contain pointers to other fields or tables in which data is stored.

Key field 50 is a primary search key for a subscriber. It is a unique identified that is
10 associated with each record in database 21.

Player ID field 51 contains a unique hardware identification code for a secure digital audio player device. This code is also contained in the secure digital audio player, and is used by the player to
15 identify itself to the system, and optionally to encrypt audio content.

Name field 52, address field 53, e-mail field 54, and phone field 55 contain, respectively, the name, address, e-mail address and telephone number of the
20 subscriber owning the secure digital audio player identified in player ID field. These fields are not required, and users who are concerned about privacy may refuse to fill in these fields. Failing to fill in these fields may prevent some special offers or other
25 features or promotional material from being available.

Player phone field 56 contains the phone number for the phone line that the secure digital audio player uses to connect to the system. It may be possible to automatically determine this information
30 when the secure digital audio player connects.

Credit field 57 contains credit card information for a subscriber. This field is not required, but numerous features of the system, such as

- 26 -

an ability to instantly purchase music for download, instantly purchase CDs, or join subscription plans may require that the subscriber provide credit card information.

5 Segment field 58 contains marketing segment information on the subscriber. Based on a subscriber's preferences, downloads, and other usage of the system, the subscriber may be placed in one or more market segments by the system. Segment field 58 may be used
10 by the system to target offers, advertisements, and other marketing material to the subscriber, as well as to determine which articles, artists, and music to feature for the user. Additionally, segment field 58
15 may be used by the system to generate targeted playlists for the subscriber.

 Playlists field 59 contains the user's playlists. Playlists typically comprise a query on the data stored in digital media repository 20. Alternatively, playlists may contain links to one or
20 more records in digital media repository 20, containing audio content that the subscriber is permitted to download into his or her secure digital audio player. In addition to accessing the playlists stored in
25 playlists field 59, a subscriber may have access to playlists or to individual songs through subscriptions, or through the system, which may make certain songs and playlists available to all users.

 Media field 60 contains links to each of the records in digital media repository 20 with audio
30 content that the subscriber has a right to download or include in a playlist. Some of the entries in media field 60 may be songs that the subscriber has uploaded to the system. Additionally, media field 60 may

- 27 -

contain information on the subscriber's rights with respect to each of the pieces of audio content, such as if the subscriber's right to download expires on a particular date, or after a predetermined number of
5 downloads. In addition to being able to access audio content included in media field 60, a user may be able to access audio content that is free to all users of the system, or that is accessed through a subscription.

Subscription field 61 contains a list of the
10 subscriptions held by the subscriber. Each subscription provides the subscriber with access to a predetermined set of audio content with predetermined usage conditions. The content available through a subscription may change over time, so that, for
15 example, a subscriber always has access to the top ten rap songs each month, or to all of the music produced by a favorite artist or group. Some subscriptions may require the a subscriber pay a fee (one-time or periodic), while other subscriptions may be available
20 for free. Some subscriptions may be subsidized by including advertising audio content.

It will be understood by one skilled in the relevant arts that one or more of the above-described fields may be in separate tables within database 21.
25 Additionally, other data on subscribers may be tracked, including information on each subscriber's favorite audio content in each category, the number of times that a subscriber downloads a particular type of audio content, a subscriber's favorite radio stations, and
30 the times of day when each subscriber typically downloads content into his or her secure digital audio player. Some of this additional information may require additional tables or fields in database 21,

- 28 -

while some may be obtained through use of online transaction analysis tools, which are typically provided by database vendors.

Referring now to FIG. 4, the elements that
5 make up the web-based interface to the system, that is used to select and purchase audio content, build playlists, enter personal information, read articles, and interact with the system are described. Access to these elements is provided through server 22, which
10 executes web server software, and permits a user to use Internet access device 25 to interact with the system. The elements shown in FIG. 4 are generated using web pages, scripts, and applets that communicate between Internet access device 25, server 22, database 21 and
15 digital media repository 20 to produce interactive views of playlists, articles, and other content. A subscriber may display or interact with one or more of the views of these elements at any given time.

Featured artist element 65 displays a window
20 identifying a recording artist or group that has recently received a relatively large amount of press coverage. Determining which artists have received prominent coverage may be determined by querying digital media repository 20 to determine which artists
25 appear in the most articles. Additionally, a subscriber's preferences and the genre and category of the music associated with the artist may be used to generate a featured artist element targeted to a particular subscriber.

30 The window displayed by featured artist element 65 shows information such as the name of the artist, pointers to recent articles about the artist, a picture of the artist, and a "create sampler" button,

- 29 -

that will send a query to digital media repository 20
that generates a playlist containing samples of the
artist's best, most recent, or most popular songs.
Additionally, the window displayed by featured artist
5 element 65 may contain links that permit a user to
purchase CDs by the featured artist.

Featured stories element 66 displays a window
containing links to one or more articles from the print
media that the service's editors considered to be
10 particularly noteworthy. Subscribers may customize
featured stories element 66 to determine the number of
featured stories that should be listed, and the general
subject matter (e.g., genre and category) of the
stories.

15 Quote element 67 displays a window containing
a quote of the day. Each day, the editorial staff will
select a noteworthy or amusing quote to display in
quote element 67.

Playlist selector element 68 displays a pop-
20 up list of the user's available playlists, and permits
the user to select a playlist that will be displayed in
the window of playlist element 69, and that will be
downloaded to secure digital audio player 24. Although
each subscriber may have numerous playlists, only one
25 playlist may be active at a time.

A subscriber may have numerous custom
playlists that may be selected using playlist selector
element 68. Additionally, there may be numerous freely
accessible playlists, promotional playlists, playlists
30 associated with featured artists, or playlists
associated with subscriptions that may be accessed
through playlist selector element 68. Playlist
selector 68 may also be used to request that the system

- 30 -

automatically generate a customized playlist for the user, as described hereinbelow.

Playlist selector element 68 also may provide subscribers with an ability to purchase access to
5 entire playlists of songs, or "electronic albums." By selling an entire playlist of songs at once, distributors are able to exert a degree of creative control over the distribution of their audio content. Optionally, these playlists may be uneditable,
10 providing distributors with even greater control over the manner in which their audio content is packaged and distributed.

Playlist element 69 displays a window that shows the contents of the currently active playlist
15 (i.e., the playlist selected in playlist selector element 68). For each item in a playlist the title of the audio content, and the artist or group that performed the audio content are displayed. Optionally, other information may be displayed by playlist element
20 69, such as the album from which the audio content was taken, the genre and category, the imprint label, or the promotion label. Additionally, playlist element 69 may include links that permit a user to purchase the CD or single on which items in the playlist appear. If a
25 subscriber "clicks" in the window of playlist element 69, playlist editor element 70 will become active.

Playlist editor element 70 provides a window that permits a user to edit the active playlist. Audio content may be rearranged, added, or removed from the
30 playlist. When audio content is added to the playlist, audio content selector element 71 may be displayed.

Audio content selector element 71 displays a window in which the subscriber may select audio content

- 31 -

for inclusion in a playlist. The window includes a list of the various pieces of audio content to which the user has access, either due to purchase, subscription, upload, promotional release, or free
5 release of the audio content to subscribers. The subscriber may select audio content from the list, or may acquire new audio content.

To acquire new audio content, the subscriber may select from the audio content available in digital
10 media repository 20. Audio content may be selected according to genre and category, artist or group, name, album, similarity of sound to another piece of audio content, or any other searchable criteria. For searches based on similarity of sound, known content
15 based retrieval techniques may be used.

If the content requires payment of a license fee for downloading, the price will be displayed, and the subscriber will have an option to complete the purchase using a credit card or other payment method.
20 If the audio content is available for free, or is part of a subscription plan in which the subscriber participates, then the subscriber may move the audio content directly into his or her list of available audio content, or into a playlist. The user may also
25 be given an option to purchase the CD or single on which the selected audio content appears.

In addition to selling individual pieces of audio content, distributors may make audio content available in groups. For example, numerous songs may
30 be made available only together, and with restrictions on the ability to download the songs individually, and the ability to rearrange their order. In this manner, additional creative control may be maintained.

- 32 -

Similar restriction mechanisms may be used to permit subscribers to download only a preset number of songs from any one album. For example, a distributor may allow a subscriber to download three tracks from a new album for free, and charge for any additional
5 downloads. This permits distributors to promote their albums, while decreasing the risk of lost sales.

Optionally, audio content may also be made available through use of audio upload element 72.
10 Audio upload element 72 permits subscribers to upload audio content from their Internet access devices, and store the audio content in digital media repository 20. Since audio content may only be downloaded into secure digital audio player 24 through server 22, this
15 provides subscribers with a way to make their own audio content available for download through server 22.

Typically, due to copyright concerns, subscribers may not be permitted to share audio content uploaded through audio upload element 72, so other
20 subscribers will not have access to the uploaded audio content. Additionally, audio upload element 72 may check uploaded content to attempt to make certain that it is not pirated. Depending on the encoding of the music, this may be accomplished through steganographic
25 techniques, such as digital watermarking, or as part of the data associated with audio content.

Uploaded audio content may occupy a considerable amount of storage space in digital media repository 20. Subscribers may be given a fixed amount
30 of storage space for uploaded audio content, and may be able to purchase the use of additional storage space.

As seen in FIG. 5, at any given time, one or more of the windows associated with the elements

- 33 -

described hereinabove with reference to FIG. 4 may be displayed on a screen of Internet access device 25. In the example shown in FIG. 5, windows associated with featured artist element 65, featured stories element 5 66, quote element 67, playlist selector element 68, and playlist element 69 are displayed in window 75 on a screen of Internet access device 25. Toolbar 76 provides an area where additional elements and functions that may be performed by the user may be 10 selected, and banner ad area 77 provides space in which targeted commercial messages may be displayed. Preferably, these areas and windows are displayed in the context of a web browser running on Internet access device 25, and the subscriber may customize the windows 15 that are displayed.

It will be understood by one skilled in the art that many of the windows and elements described with reference to FIGS. 4 and 5 provide information that may be useful to assist users in selecting music, 20 but are not required. In a minimal system, for example, a user could simply select from a number of preset playlists, or request that the system generate a playlist based on the user's preferences. In such a system, only playlist selector element 68, and, 25 optionally, playlist element 69 need be used.

Such a simplified selection process could be accessed through secure digital audio player 24 in standalone mode, rather than requiring Internet access device 25 for music selections. Additionally, by 30 limiting the user to being able to select only from preset playlists or system-generated playlists, the copyright issues involved in distributing the music may be minimized, and the creative control and marketing

- 34 -

concerns of music publishers may be better addressed. Copyright issues involved in distributing news stories, quotes, images of performers, and other non-music content may also be avoided by limiting the selection
5 of elements and windows that are available to users.

Referring now to FIG. 6, the portable secure digital audio player of the present invention is shown. It will be understood by one skilled in the art that the design shown in FIG. 6 is primarily intended to
10 illustrate the functional aspects of secure digital audio player 24. Changes to the ornamental design, the placement and number of controls, the shape and size of the screen, and other modifications may be made without departing from the invention. In a preferred
15 embodiment, numerous ornamental packages embodying the functionality of secure digital audio player 24 may be made available, to suit differences in personal taste.

Secure digital audio player 24 preferably comprises housing 80, display 82, navigation buttons
20 84, playback buttons 86, volume control 87, headphone jack 88, telephone jack 90, telephone through jack 92, and power supply connector 94. Optionally, secure digital audio player 24 may comprise a slot (not shown) into which additional memory for storing audio content
25 may be inserted.

Housing 80, which preferably comprises a molded plastic material, contains the other components that comprise secure digital audio player 24. Housing 80 may optionally include a panel (not shown) to permit
30 access to a battery compartment (not shown). Housing 80 may be tinted and molded to provide numerous ornamental design options for secure digital audio player 24.

- 35 -

Display 82 comprises a screen on which secure digital audio player 24 may display status information. In a preferred embodiment, display 82 comprises an LCD display having a resolution of approximately 200 pixels in a horizontal direction, and 100 pixels in a vertical direction. Display 82 preferably shows five rows from the currently loaded playlist, each row containing information on one item of audio content stored in secure digital audio player 24. The central (third) row is highlighted, and shows the currently selected audio content, while the first and second rows show the previous two audio content items, and the fourth and fifth rows show the next two audio content items. Each row shows a "track number" for the row, which indicates the row's position in the playlist, the name of the artist or group, the name of the song or other audio content, and a checkbox indicating whether the row is selected or not. Audio content in selected rows is played, while audio content in rows that are not selected is skipped. Since the information listed in a row may not fit within the space provided for the row, rows may be horizontally scrollable.

In addition to showing rows of information on audio content, display 82 may show various status indicators, such as the current volume level, the battery level, and dialing status indicators when secure digital audio player 24 is connecting to server 22. Display 82 also may be used to display playlist options, and permit a subscriber to select a playlist when secure digital audio player 24 is connected to server 22 in standalone mode.

Navigation buttons 84 comprise up arrow button 84a, right arrow button 84b, down arrow button

- 36 -

84c, left arrow button 84d, and selector button 84e. Up arrow button 84a and down arrow button 84c are preferably used to scroll the list of audio content shown on display 82 up and down. Right arrow button 84b and left arrow button 84d are preferably used to scroll the content of the central row on display 82 right and left. Selector button 84e is preferably used to toggle the selection status of the audio content item in the central row of display 82.

10 It will be understood that navigation buttons 84 may have purposes other than those described above, depending on the mode of operation of secure digital audio player 24. For example, when secure digital audio player 24 is connected to server 22 in standalone mode, navigation buttons 84 may be used to select a playlist to be downloaded from a menu of available playlists. For example, up arrow button 84a and down arrow button 84c may be used to scroll through options on a menu, while right arrow button 84b and left arrow button 84d may be used to move through levels of a menu hierarchy. By using these buttons, the user may scroll through selections representing categories of audio content, genres within a category, and preset playlists within a genre. Selector button 84e may be used to chose a final selection. If the final selection is a category, rather than a specific playlist, then server 22 will automatically generate a playlist of audio content within that category. Similarly, if the final selection is a genre, then server 22 will automatically generate a playlist (typically 40 songs) containing music within that genre. Generally, navigation buttons 84 may be used for any navigation and selection tasks.

- 37 -

Playback buttons 86 preferably comprise skip back button 86a, search back button 86b, pause button 86c, stop button 86d, play button 86e, search ahead button 86f, and skip ahead button 86g. Play button 86e
5 causes audio playback to begin at the currently selected (i.e., highlighted in the center row of display 82) audio content. Alternatively, play button 86e may start playing at the beginning of the playlist. Secure digital audio player 24 will continue with the
10 next audio content item in the playlist when the current audio content item ends.

Stop button 86d halts audio playback, and sets the current position of the audio playback to the beginning of the playlist. Pause button 86c also halts
15 audio playback, but does not alter the position of audio playback, so that audio playback may continue from the point at which playback was paused.

Search ahead button 86f and search back button 86b cause audio playback to occur at a
20 predetermined accelerated pace, in a forward or a backward direction, respectively. Skip ahead button 86g causes playback to skip to the beginning of the next audio content item in the playlist, and continue from there. Skip back button 86a causes playback to
25 skip immediately back to the start of the current audio content item, and proceed from there. Pressing skip back button 86a a second time shortly after skipping to the beginning of the current audio content may cause playback to skip back to the beginning of the previous
30 audio content item in the current playlist.

Volume control 87 comprises a wheel, knob, buttons, or other control mechanism that permits a user

- 38 -

of secure digital audio player 24 to adjust the volume of playback either up or down.

Headphone jack 88 permits a user to plug in a pair of headphones, so that he or she may listen to the audio content being played by secure digital audio player 24. Headphone jack 88 preferably comprises a 3.5mm female stereo mini jack.

Telephone jack 90 and telephone through jack 92 are used to connect secure digital audio player 24 to a public telephone system so that secure digital audio player 24 may connect to server 22 to download digital music. Telephone jack 90 connects a telephone line to modem circuitry inside of secure digital audio player 24, permitting the modem circuitry to establish a connection to server 22, and to transfer data between server 22 and secure digital audio player 24.

Telephone through jack 92 is a passthrough for the telephone line, and permits additional telephone devices to be connected through secure digital audio player 24. Telephone jack 90 and telephone through jack 92 preferably comprise standard RJ11 modular telephone jacks.

In an alternative embodiment, telephone jack 90 and telephone through jack 92 may be replaced with a network connector (not shown), and the modem circuitry within secure digital audio player 24 may be replaced with network interface circuitry. This permits secure digital audio player 24 to be connected to a cable modem, a DSL connection, or other high-speed network, through which secure digital audio player 24 may communicate with server 22. In a second alternative embodiment, telephone jack 90 and telephone through jack 92 may be removed, and secure digital audio player

- 39 -

24 may connect to a network using wireless communication technology.

Power supply connector 94 is used to recharge secure digital audio player 24. In a preferred
5 embodiment, secure digital audio player 24 includes an internal rechargeable battery (not shown), which may be recharged by plugging a connector at one end of a power cable (not shown) into power supply connector 94, and plugging the other end of the power cable into a
10 standard electrical socket.

Referring now to FIG. 7, an alternative embodiment of secure digital audio player 24 is shown, in which power supply connector 94, telephone jack 92, and telephone through jack 94 are located on cradle
15 100. Secure digital audio player 24 includes player contacts 102, that connect to cradle contacts 104 on cradle 100 when secure digital audio player 24 is placed in cradle 100. When secure digital audio player 24 is placed in cradle 100, it may be recharged, and
20 may establish communications with server 22. As discussed hereinabove, telephone jack 92 and telephone through jack 94 may be replaced with a network connector, permitting secure digital audio player 24 to communicate with server 22 through cradle 100 using a
25 high-speed network connection.

FIG. 8 shows a block diagram of the circuitry of secure digital audio player 24. Circuitry 110 is disposed within housing 80, and comprises processor 112, program memory 113, data memory 114, audio memory
30 115, communication circuitry 116, DAC circuitry 118, and analog audio circuitry 120.

Processor 112 controls the operation of secure digital audio player 24, by executing software

- 40 -

that decodes and plays audio content, accepts commands from the various buttons and controls, displays information on display 82, controls communication circuitry 116 to establish communications with server 5 22, and downloads music from server 22 into audio memory 115 using the protocols described hereinbelow. Processor 112 preferably comprises a microcontroller that integrates numerous functions of circuitry 110, such as the Cirrus Logic EP7209, available from Cirrus 10 Logic, of Fremont, California. The EP7209, for example includes on-chip ROM and SRAM that may be used as program memory 113 and/or data memory 114, as well as circuitry for controlling a liquid crystal display, interfacing with flash memory (i.e., audio memory 115), 15 and circuitry for interfacing with an audio digital-to-analog converter (DAC).

Program memory 113 stores the software that is executed by processor 112, and preferably comprises a non-volatile memory. Program memory 113 is 20 preferably rewritable, so that the software that controls secure digital audio player 24 may be updated, for instance, to handle new audio compression formats. One skilled in the art will understand that program memory 113 may be entirely or partially integrated with 25 processor 112.

Data memory 114 stores data used by the software executed by processor 112. Various of the functions performed by processor 112 may need temporary data storage space. This temporary space is provided 30 in data memory 114, which may comprise RAM, SRAM, or other types of memory. Data memory 114 may be partially or completely integrated with processor 112.

- 41 -

Audio memory 115 stores the audio content that is to be played by secure digital audio player 24. Audio memory 115 is preferably large, since the size of audio memory 115 determines the amount of audio content that may be stored on secure digital audio player 24. For example, for MP3 audio content, 64 megabytes will store approximately an hour of audio content. Audio memory 115 need only be fast enough to provide a continuous stream of data to an audio decoding program executing on processor 112, which may permit large, but relatively slow storage devices, such as miniature hard drives, to be used as all or part of audio memory 115. Typically, audio memory 115 will comprise a flash memory device, or other large non-volatile storage device.

It will be understood by one skilled in the relevant arts that all or part of audio memory 115 may be located on a removable device, such as a flash memory card. This permits a subscriber to increase the amount of memory available to secure digital audio player 24 by using larger capacity removable memory, or to select from among one or more removable memory devices to determine which audio content should be played.

Communication circuitry 116 is used to communicate data between secure digital audio player 24 and server 22. In a preferred embodiment, communication circuitry 116 comprises standard modem circuitry capable of transferring data over a telephone line at approximately 56 kilobits per second. Alternatively, communication circuitry 116 may comprise standard network interface circuitry, capable of handling high-speed communications between secure

- 42 -

digital audio player 24 and server 22. In another alternative embodiment, communication circuitry 116 may comprise wireless communication circuitry, permitting secure digital audio player 24 to receive digital audio content through a wireless connection. As explained hereinabove, communication circuitry 116 may be located within housing 80, or may be part of cradle 100.

DAC circuitry 118 comprises a standard audio DAC, that converts digital signals provided by processor 112 into analog signals. The analog signals produced by DAC circuitry 118 may then be amplified and filtered by analog audio circuitry 120. It will be understood by one skilled in the relevant arts that depending on the processor used, all or part of DAC circuitry 118 and analog audio circuitry 120 may be integrated into processor 112. Alternatively, analog audio circuitry 120 and DAC circuitry may be combined into a single part, depending on the choice of components that are used to construct circuitry 110.

Referring now to FIG. 9, an overview of the software executed by processor 112 is shown. Software 120 comprises file management routines 122, audio decoding and playback routines 124, user interface routines 126, communications routines 128, and control routines 129.

File management routines 122 manage the contents of audio memory 115. Each piece of audio content, including data such as the artist or group, and the name of the content, whether the content is enabled for playback, and any restrictions that apply to playback of the content, is stored as an "audio content file" in audio memory 115. File management routines 122 will typically include routines for adding

- 43 -

audio content files, deleting audio content files, reading audio content files, retrieving a directory of available audio content files, retrieving information associated with audio content files (i.e., the performer, name, and enabled status), and for finding the audio content file that occupies a given position in a playlist. Routines that perform these types of functions are well known, and will vary according to the type of memory that is used for audio memory 115, and the file format that is selected for the audio content files.

Audio decoding and playback routines 124 handle decoding compressed or encoded audio content, and playing back the audio content on secure digital audio player 24. In a preferred embodiment that uses MP3 encoded audio content, audio decoding and playback routines comprise a standard MP3 decoding routine. If the content is encrypted using the player's unique identifier, the decoding routines will also need to use the player's unique identifier to decrypt the audio content.

Additionally, audio decoding and playback routines 124 may comprise routines for manipulating the digital audio data that is derived from the decoding process. These routines may include routines for digitally boosting the bass, treble, or any given frequency range, adding digital audio effects to simulate particular playback environments, digital volume adjustment, and digital compression of the dynamic range of the audio. Routines to perform these functions are all well known, and will vary according to the type of effects that are desired.

- 44 -

Audio decoding and playback routines 124 also may include routines for handling copyright management. These routines may be used to apply playback restrictions associated with a particular piece of audio content. For example, if a particular piece of audio content were only permitted to be played over a ten day evaluation period, or could only be played five times, restriction routines in audio decoding and playback routines 124 would apply these restrictions before decoding or playing the audio content.

It will be understood that audio decoding and playback routines 124 may comprise numerous decoding routines, permitting secure digital audio player 24 to play music that has been encoded or compressed in a variety of formats. Additionally, audio decoding and playback routines 124 may comprise additional routines for copyright management, such as routines that will search digital audio for a "digital watermark" or other steganographic techniques. Audio decoding and playback routines may also keep totals on the number of times that an audio content item is played, for later transfer to server 22.

User interface routines 126 handle all of the interactions between secure digital audio player 24 and a user. User interface routines 126 includes routines that display information on display 82, receive input from navigation buttons 84 and playback buttons 86, and that control the operation of secure digital audio player based on the inputs received from the user. User interface routines 126 also may include routines for displaying and navigating through menus when secure digital audio player 24 is in standalone mode. User interface routines for handling these tasks are well

- 45 -

known, and vary according to the information to be displayed, and the buttons and other inputs to be handled.

Communication routines 128 handle the tasks
5 of establishing a connection between secure digital audio player 24 and server 22. Communication routines 128 include routines for establishing a connection, identifying the secure digital audio player, and sending and receiving data.

10 The method used to establish a connection depends on the type of physical connection that is used. For example, if secure digital audio player 24 includes modem circuitry, and connects to server 22 across a public telephone network, then establishing a
15 connection would typically require using the modem circuitry to dial a telephone number, establishing a modem connection, and then using a known protocol, such as point-to-point protocol (PPP) to transfer data. Alternatively, if secure digital audio player 24
20 includes a network interface, then standard network protocols may be used to establish a connection and transfer data.

Control routines 129 handle overall control of the system. These routines accept user input using
25 user interface routines 126, and determine which actions are to be taken based on the user input. For example, when play button 86d is pressed by the user, control routines 129 may use file management routines 122 to retrieve an audio content item, and invoke audio
30 decoding and playback routines 124 to decode and play the audio content item.

Similarly, control routines 129 may accept requests from server 22 when secure digital audio

- 46 -

player 24 is communicating with server 22 in standalone mode. For example, server 22 may request that a playlist menu be displayed, so that the user may select a playlist for download. In this case, control
5 routines 129 would use communication routines 128 to retrieve a menu to be displayed from server 22, user interface routines 126 to display the menu and obtain a selection from the user, and communication routines 128 to send the selection back to server 22.

10 Control routines 129 may also include routines for handling updates to the software of secure digital audio player 24. When new audio compression standards become available, for example, update routines may be used to download new software capable
15 of decoding and playing the new formats from server 22. The update routines would typically store the new software in nonvolatile program memory 113.

The exact workings of control routines 129 depend on the functions performed by secure digital
20 audio player 24, and are generally well known, or easily implemented by one skilled in the art.

Referring to FIG. 10, an overview of some of the functions performed by communication routines 128, and by routines executing on server 22 that handle
25 downloading music to secure digital audio players is described. At step 201, secure digital audio player 24 establishes a connection with server 22, and sends server 22 its unique identifier. This unique identifier helps insure that server 22 only sends audio
30 content to authorized secure digital audio players. As described hereinabove, the unique identifier may also be used to encrypt the audio content. Since secure digital audio player 24 may not transfer digital audio

- 47 -

content to any other device, and server 22 may only transfer digital audio content to an authorized secure digital audio player having the proper unique identifier, there is a greatly reduced risk of digital audio content being played or redistributed without authorization.

Server 22 receives the unique identifier at step 301, and looks up the unique identifier in database 21. Server 22 retrieves information on the subscriber, and on the selected playlist from database 21.

At step 302, server 22 requests a "cookie" (i.e., a package of data having a predetermined format and predetermined name) from secure digital audio player 24. Secure digital audio player 24 receives this request at step 202, and sends a "cookie" containing data on the audio content currently stored in secure digital audio player 24. The "cookie" also may contain other information, such as which audio content is selected, and the number of times that particular pieces of audio content have been played, for billing and copyright management purposes.

At step 303, server 22 receives the cookie, and compares the audio content currently stored in secure digital audio player 24 with the playlist to be downloaded. Any audio content that is already present on secure digital audio player 24 need not be downloaded again.

At step 304, server 22 retrieves the audio content for the selected playlist from digital media repository 20, and downloads the required audio content to secure digital audio player 24. Secure digital audio player 24 receives the audio content at step 203,

- 48 -

and uses file management routines 122 to store the audio content in audio memory 115.

The process of downloading audio content to secure digital audio player 24 may be interrupted, due, 5 for example, to conditions on the telephone line. If such interruptions occur, secure digital audio player 24 may reestablish a connection with server 22, and continue downloading audio content.

This entire process of connecting to server 10 22, sending the unique identifier, and receiving audio content may occur at a predetermined time without user intervention. Secure digital audio player 24 could, for example, be set up to connect to server 22 and download audio content in the middle of the night.

15 Alternatively, by using secure digital audio player 24 in standalone mode, the process of connecting and downloading could be interactive, permitting a subscriber to select a playlist to be downloaded while secure digital audio player 24 is connected to server 20 22. Although the user may not have as many options using this method as he or she would have by interacting with server 22 through Internet access device 25, this provides a way to select and download audio content even without access to a web browser or 25 other Internet access device.

Referring now to FIG. 11, a record from an additional radio playlist database that may be included in database 21 or digital media repository 20 is shown. As described in detail hereinbelow, this radio playlist 30 database may be used to assist in automatic generation of playlists suited to a user's musical preferences.

Radio station record 135 includes name field 136, format field 137, call letters field 138,

- 49 -

frequency field 139, market field 140, market size
field 141, city field 142, state field 143, telephone
field 144, and playlists field 145. The purpose of
radio station record 135 is to keep track of the
5 playlists of radio stations in major markets, to assist
in automatic generation of playlists.

Name field 136, and call letters field 138
are text fields containing the name of a radio station,
and its call letters, respectively. In some cases, the
10 name and call letters of a radio station may be the
same. Frequency field 139 may also be used to identify
a station, and contains the radio frequency at which
the station broadcasts.

Format field 137 contains information on the
15 format of the material broadcast by the radio station.
Generally, the format will be a broad description of
the type or genre of the content. For example, a
station may have a "pop" format, a "rock" format, a
"country" format, or a "talk" format.

20 Market field 140 contains information on the
market in which the radio station broadcasts. This may
be the name of the city, urban area, or region in which
the station broadcasts. Market size field 141 contains
the approximate number of listeners in the market in
25 which the station broadcasts.

City field 142, state field 143, and
telephone field 144 contain, respectively, the city in
which a station is located, the state in which a
station is located, and the telephone number of the
30 station. Telephone field 144 may serve as a proxy for
the location of the station, since telephone numbers
typically contain area codes and other coded prefix
numbers indicative of location.

- 50 -

Playlists field 145 contains information on the songs that the radio station has played recently. The playlists fields typically contain a date for each playlist, and the artists and singles that were played.

5 This radio station playlist information is used to help determine a user's musical preferences, and to assist in automatically generating a playlist for secure digital audio player 24. When a user first connects with server 22, he or she enters various
10 personal information, such as a name, address, phone number, and e-mail address as part of the registration process. Once server 22 knows the user's address, radio station information can be accessed, to display a list of radio stations that are in the user's market.
15 The user provides server 22 with information on his or her musical preferences by indicating which of the radio stations are his or her favorites.

Once the system knows which radio stations a user prefers, playlists containing music by the same
20 artists and in the same genres may be automatically generated. For copyright management reasons, it may be desirable to offer access only to such automatically generated playlists, and to pre-built promotional playlists, rather than permitting a user to create his
25 or her own playlists.

Referring to FIG. 12, a method for using information about a user's preferred radio station to automatically generate a playlist for the user is described. Such playlists will typically be generated
30 by server 22, but also may be generated by any other computers or servers that have access to database 21 and digital media repository 20.

- 51 -

In step 401, the user provides information on his or her favorite radio station. Typically, the user will be given a choice of several radio stations that broadcast in the user's area. This step is typically
5 performed when the user first registers with the service, with the option of changing the favorite radio station later.

At step 402, which is typically performed when the system needs to generate a playlist, the
10 system queries the radio playlist database to retrieve recent radio playlists for the user's favorite radio station. As described hereinabove, these radio playlists contain the artists and singles that have been played on the radio station.

15 In step 403, the system determines which artists are the top five artists in the recent radio playlists retrieved in step 402, based on how many times each artist appears in the recent radio playlists. The top five artists are "selected", and
20 will have songs included in the generated playlist.

In step 404, for each artist appearing in the radio playlists retrieved in step 402, the system will determine the genre for the artist by querying the digital media repository. For artists who perform in
25 multiple genres, the singles performed by that artist in the radio playlists will be used to determine which genre should be used.

For each of the genres determined at step 304, a number of slots in the generated playlist are
30 allocated in step 405. Typically, a generated playlist will contain approximately 40 songs. Five of these are used for songs by the five artists selected in step 303. The remaining approximately 35 songs are split

- 52 -

between the genres of music that were determined in step 304 in proportion to the amount of music in that genre in the radio playlists. Thus, if 40% of the artists in the radio playlists retrieved in step 402
5 fit in a particular genre, then approximately 40% (i.e. approximately 14) of the 35 songs to be allocated would be selected from that genre.

At step 406, for each genre in the generated playlist, the system queries digital media repository
10 20 to determine who the top n artists in the genre, where n is the number of songs in the generated playlist that are allocated to that genre. These artists are "selected", along with the five top artists from the radio playlist that were selected in step 403.
15 Thus, if there are seven songs in the generated playlist allocated to a genre, the system will attempt to pick the top seven artists in that genre, and select those artists. If it is impossible to find enough artists in a given genre, the top artists in the genre
20 may be repeated as necessary.

At step 407, for each selected artist, the system queries digital media repository 20 to select a single by that artist to fill one of the spaces in the generated playlist. The single selected preferably
25 meets a number of criteria. First, the single should not be a single that already appears in the generated playlist. The single should also preferably not be the same single that appears in the radio station playlists. Finally, the single should preferably have
30 a good ranking or review, be high on the charts, be mentioned or be from an album that is mentioned in media sources, or be from an album that received good reviews or ranking. If all these criteria cannot be

- 53 -

met, then the system may choose the single at random from among the works of the particular artist, or, if that is not possible, from the appropriate genre.

Once all of the songs in the generated
5 playlist are selected by the system, the generated playlist may be downloaded to secure digital audio player 24 as described hereinabove. Advantageously, as the playlists of the user's favorite radio station change, the mix of music that is downloaded to secure
10 digital audio player 24 will change in a similar manner, but without simply repeating the same material that the user can hear on the radio.

It will be apparent to one skilled in the relevant arts that there are many variations to the
15 above-described method that could be used to generate playlists based on a user's favorite radio stations. For example, the generated playlist could be based entirely on the top singles in the genres from the radio playlists, rather than on the artists.
20 Additionally, other factors, such as how often the user plays specific tracks from previous generated playlists, or artists and singles that are being heavily promoted by record companies, may be used to generate a playlist for secure digital audio player 24.

25 Although preferred illustrative embodiments of the present invention are described above, it will be evident to one skilled in the art that various changes and modifications may be made without departing from the invention. For example, the appearance and
30 placement of various features on secure digital audio player 24 or cradle 100 could be altered, or additional controls could be added. Additionally, server 22 may provide access to additional services and sources of

- 54 -

information on music, including digitally encoded music videos, or audio and video coverage of music-related topics. Further, it will be understood that the structure of database 21 or digital media repository 20
5 could be altered without substantially affecting the nature of the data that is stored, or the uses of the data by the system. It is intended in the appended claims to cover all such changes and modifications that fall within the true spirit and scope of the invention.

- 55 -

What Is Claimed Is:

1. A method of distributing digital audio content, the method comprising:

providing a server that makes the digital audio content available to a subscriber;

providing a player for use by the subscriber, the player comprising a unique identifier, a memory for storing digital audio content, communication circuitry, audio output circuitry, and a processor, the processor programmed to use the communication circuitry to download the digital audio content from the server and store the digital audio content in the memory, and play the digital audio content, wherein the player is incapable of transferring the digital audio content to other devices;

accepting a selection of digital audio content from the subscriber;

establishing communication between the player and the server;

sending the unique identifier to the server to identify the player;

downloading the selection of digital audio from the server to the player; and

recording copyright management information on the server for the selection of digital audio.

2. The method of claim 1, wherein accepting a selection of digital audio content from the subscriber comprises:

establishing Internet communication between a web browser and the server;

sending a subscriber identifier from the web browser to the server;

- 56 -

displaying available selections of audio content in the web browser;
using the web browser to choose a selection of audio content;
sending the selection of audio content to the server; and
using the subscriber identifier to associate the selection of audio content with the unique identifier.

3. The method of claim 2, further comprising displaying information related to music in the web browser to assist the subscriber in making purchasing decisions.

4. The method of claim 3, wherein displaying information related to music comprises determining preferences of the subscriber and displaying information based on the preferences of the subscriber.

5. The method of claim 4, wherein determining preferences of the subscriber comprises determining a category and genre of audio content that the subscriber most frequently downloads.

6. The method of claim 4, wherein determining preferences of the subscriber comprises determining items of audio content that are most frequently downloaded by the subscriber.

- 57 -

7. The method of claim 4, wherein determining preferences of the subscriber comprises determining items of audio content that are most frequently played by the subscriber on the player.

8. The method of claim 3, wherein displaying information related to music comprises displaying a featured artist.

9. The method of claim 3, wherein displaying information related to music comprises displaying a list of featured articles.

10. The method of claim 3, wherein displaying information related to music comprises displaying a quote of the day.

11. The method of claim 1, wherein providing a player to the user comprises providing a player in which the communication circuitry comprises modem circuitry, and wherein establishing communication between the player and the server comprises using the modem circuitry to establish communication over a public telephone network between the player and the server.

12. The method of claim 1, further comprising sending information from the player to the server, the information including a list of audio content stored in the memory of the player.

- 58 -

13. The method of claim 12, wherein downloading the selection of digital audio comprises downloading items in the selection of digital audio that are not already stored in the memory of the player.

14. The method of claim 12, wherein sending information from the player to the server further comprises sending copyright management information from the player to the server.

15. The method of claim 1, further comprising encrypting the selection of digital audio.

16. The method of claim 15, wherein encrypting the selection of digital audio comprises using the unique identifier to encrypt the digital audio.

17. The method of claim 1, wherein accepting a selection of digital audio from the subscriber comprises:

 sending one or more choices of digital audio content from the server to the player;

 displaying the one or more choices of digital audio content on a display of the player;

 providing a user interface on the player through which the subscriber makes the selection of digital audio from among the one or more choices of digital audio content; and

 sending the selection of digital audio to the server.

- 59 -

18. The method of claim 1, further comprising providing an automatic playlist choice that requests that a playlist be automatically generated, and automatically generating a playlist if the subscriber selects the automatic playlist choice.

19. The method of claim 18, wherein automatically generating a playlist comprises generating a playlist based on the musical preferences of the subscriber.

20. The method of claim 19, wherein generating a playlist based on the musical preferences of the subscriber further comprises:

determining a favorite radio station of the subscriber; and

generating a playlist based on songs that have recently been played on the favorite radio station of the subscriber.

21. The method of claim 20, wherein generating a playlist based on songs that have recently been played on the favorite radio station of the subscriber comprises:

determining an artist that performed each song;

determining a genre of music associated with each artist;

allocating a number of slots in the playlist to each genre;

compiling a list of top artists for each genre; and

- 60 -

choosing a single for inclusion in the playlist by each artist in the list of top artists.

22. The method of claim 21, wherein choosing a single for inclusion in the playlist comprises choosing a single that is high on the charts, has a good ranking or review, is mentioned in media sources, or that is taken from an album with a good ranking, review, or coverage in media sources.

23. The method of claim 21, wherein choosing a single for inclusion in the playlist comprises choosing a single at random.

24. A secure digital audio player comprising:

a processor;

a program memory that stores a plurality of routines for execution by the processor;

an audio memory that stores an audio content item that may be played by the secure digital audio player;

communication circuitry that establishes communication with a server and transfers data between the server and the secure digital audio player;

a unique identifier that identifies the secure digital audio player to the server;

a display; and

a user control that permits the user to enter a command,

wherein the plurality of routines comprise:

a communication routine that downloads the audio content item from the server;

- 61 -

a decoding and playing routine that decodes and plays the audio content item;

a file management routine that manages the audio content item stored in the audio memory;

a user interface routine that reads the user control to determine the command, and outputs information to the display; and

a control routine that controls operation of the secure digital audio player,

the secure digital audio player being incapable of downloading the audio content item from sources other than the server, and incapable of transferring the audio content item in digital form to any other device.

25. The secure digital audio player of claim 24, wherein the secure digital audio player is portable.

26. The secure digital audio player of claim 24, wherein the secure digital audio player comprises a portable unit and a cradle unit, and wherein the communication circuitry is disposed in the cradle unit.

27. The secure digital audio player of claim 24, wherein the communication circuitry comprises modem circuitry.

28. The secure digital audio player of claim 24, wherein the communication circuitry comprises network interface circuitry.

- 62 -

29. The secure digital audio player of claim 24, wherein the user control comprises a set of navigation buttons that are used to specify the audio content item to be played, and a set of playback buttons that are used to command the secure digital audio player to play the audio content item, stop playing the audio content item, pause playing the audio content item, skip to a next audio content item, or skip to a previous audio content item.

30. The secure digital audio player of claim 24, wherein the display shows the name of the audio content item, the artist who performed the audio content item, the position of the audio content item in a playlist, and whether the audio content item is selected.

31. The secure digital audio player of claim 30, wherein the display further shows the name, artist, position in the playlist, and selection status for a first previous audio content item, a second previous audio content item, a first subsequent audio content item, and a second subsequent audio content item.

32. The secure digital audio player of claim 24, wherein the audio content item is encoded using MP3 encoding, and wherein the decoding and playback routine comprises a routine that decodes MP3 encoded content.

- 63 -

33. The secure digital audio player of claim 24, wherein the audio content item is encrypted using the unique identifier, and wherein the decoding and playback routine comprises a routine that decrypts the audio content item using the unique identifier.

34. The secure digital audio player of claim 24, wherein the audio content item is associated with a usage restriction, and wherein the decoding and playback routine comprises a routine that enforces the usage restriction.

35. The secure digital audio player of claim 24, wherein the communication routine comprises:

a connection routine that establishes a connection between the secure digital audio player and the server;

an identifier routine that sends the unique identifier to the server; and

a download routine, that receives the audio content item from the server.

36. The secure digital audio player of claim 35, wherein the communication routine further comprises a contents routine that sends information about the contents of the audio memory to the server.

37. The secure digital audio player of claim 24, wherein the control routine comprises a standalone mode routine, that retrieves a menu of audio selections from the server, displays the menu of audio selections on the display, permits the user to make a selection

- 64 -

from the menu of audio selections, and sends the selection to the server.

38. The secure digital audio player of claim 24, wherein the control routine comprises an update routine, that permits contents of the program memory to be updated.

39. The secure digital audio player of claim 38, wherein the update routine updates the contents of the program memory with programmed routines sent to the secure digital audio player by the server.

40. A server for securely delivering digital audio content to a subscriber, the server comprising:

- a selection routine that accepts a selection of a playlist from the subscriber, the playlist specifying one or more audio content items;

- a connection to a digital media repository from which the server retrieves the one or more audio content items specified in the playlist;

- a connection to a subscriber database from which the server retrieves information about the subscriber;

- a communication channel, through which the server communicates with a secure digital audio player, the secure digital audio player identifying itself to the server through use of a unique identifier; and

- a programmed routine that receives the unique identifier from the secure digital audio player, uses the unique identifier to identify the playlist, and downloads the one or more audio content items specified in the playlist to the secure digital audio player,

- 65 -

wherein the server is further programmed to refuse to download digital audio content to any device other than the secure digital audio player.

41. The server of claim 40, wherein the server further comprises an Internet connection through which the server communicates with an Internet access device, and wherein the selection routine accepts the selection of the playlist through the Internet connection.

42. The server of claim 40, wherein the selection routine accepts the selection of the playlist from the secure digital audio player, through the communication channel.

43. The server of claim 42, wherein the server further comprises a menu routine, that sends a menu of playlists to the secure digital audio player through the communication channel.

44. The server of claim 40, wherein the digital media repository comprises a digital storage system storing a plurality of media records, each one of the plurality of media records comprising:

an audio content field storing a digital audio content item;

an artist field storing the artist or group that recorded the digital audio content item;

a classification field storing a value indicative of the type of content in the digital audio content item;

- 66 -

a title field storing name or title of the digital audio content item;

a media field storing information on other media related to the digital audio content item; and

a copyright management field storing copyright management information.

45. The server of claim 44, wherein the classification field comprises a category field, specifying a category for the digital audio content item, and a genre field, specifying a genre or subcategory for the digital audio content item.

46. The server of claim 44, wherein the media field comprises a magazine field, containing references to magazine articles on the digital audio content item.

47. The server of claim 44, wherein the media field comprises a chart field, containing references to ranking charts on which the digital audio content item appears, and positions of the digital audio content item in the ranking charts.

48. The server of claim 44 wherein the media field comprises a ranking field, containing ranking values for the digital audio content item derived from one or more media sources.

49. The server of claim 44, wherein the copyright management field contains information on the number of times that the digital audio content item has been downloaded.

- 67 -

50. The server of claim 44, wherein the copyright management field comprises a company field storing information on a company that distributes, promotes, or owns rights to the digital audio content item.

51. The server of claim 44, wherein each one of the plurality of media records further comprises a pricing field, containing pricing information on the digital audio content item.

52. The server of claim 51, wherein the pricing field stores information on one or more pricing plans associated with the digital audio content item, each of the one or more pricing plans including pricing information and zero or more restrictions on use.

53. The server of claim 40, wherein the subscriber database comprises a digital storage system storing a plurality of subscriber records, each one of the plurality of subscriber records comprising an identifier field storing the unique identifier of the secure digital audio player of the subscriber.

54. The server of claim 53, wherein each one of the plurality of subscriber records further comprises a personal information field containing personal information on the subscriber.

55. The server of claim 54, wherein the personal information field comprises:

a name field, storing a name of the subscriber;

- 68 -

an address field, storing a physical address of the subscriber;

an e-mail field, storing an e-mail address of the subscriber; and

a phone field, storing a phone number of the subscriber.

56. The server of claim 54, wherein the personal information field comprises a credit field, storing information on a credit card used by the subscriber.

57. The server of claim 53, wherein each one of the plurality of subscriber records further comprises a market segment field, that contains information on a market segment to which the subscriber belongs.

58. The server of claim 53, wherein each one of the plurality of subscriber records further comprises an access field, that identifies audio content items in the digital media repository to which the subscriber has access.

59. The server of claim 53, wherein each one of the plurality of subscriber records further comprises a playlists field, containing information on playlists that are available to the subscriber, each playlist identifying one or more audio content items in the digital media repository.

- 69 -

60. The server of claim 53, wherein each one of the plurality of subscriber records further comprises a subscriptions field, containing information on subscription services offered on the server to which the subscriber has access.

61. The server of claim 40, wherein the server further comprises a programmed routine that uses the Internet connection to cause a featured artist window to appear on an Internet access device used by the subscriber, the featured artist window containing the name of a featured artist, a link to an article related to the featured artist, and a link that causes the server to generate a playlist identifying one or more audio content items in the digital media repository performed by the featured artist.

62. The server of claim 40, wherein the server further comprises a programmed routine that uses the Internet connection to cause a featured stories window to appear on an Internet access device used by the subscriber, the featured stories window containing links to one or more featured stories.

63. The server of claim 40, wherein the server further comprises a playlist selection routine that uses the Internet connection to cause a playlist selector to be displayed on an Internet access device used by the subscriber, the playlist selector permitting the subscriber to designate a selected playlist that is sent to the server.

- 70 -

64. The server of claim 63, wherein the server further comprises a playlist display routine that uses the Internet connection to cause a playlist window to be displayed on the Internet access device used by the subscriber, the playlist window displaying information on one or more audio content items in the digital media repository that are in the selected playlist.

65. The server of claim 63, wherein the server further comprises a playlist edit routine that uses the Internet connection to cause a playlist editor window to be displayed on the Internet access device used by the subscriber, the playlist editor window permitting the subscriber to add an audio content item to the selected playlist, delete an audio content item from the selected playlist, and move an audio content item from a first position in the selected playlist to a second position in the selected playlist.

66. The server of claim 63, wherein the server further comprises a content selector routine that uses the Internet connection to cause an audio content selector window to be displayed on the Internet access device used by the subscriber, the audio content selector window permitting the user to select audio content for inclusion in the selected playlist.

67. The server of claim 66, wherein the audio content selector window further permits the user to purchase a license to download a selected audio content item.

- 71 -

68. The server of claim 66, wherein the audio content selector window further permits the user to request that the server search for audio content that matches a selected criterion.

69. The server of claim 68, wherein the selected criterion comprises similarity of sound to a selected audio content item.

70. The server of claim 63, wherein the server further comprises an upload routine that uses the Internet connection to cause an upload window to be displayed on the Internet access device used by the subscriber, the upload window permitting the user to upload an audio content item to the server, and instruct the server to store the audio content item in the digital media repository.

71. The server of claim 40, wherein the server further comprises a connection to a radio station database from which the server retrieves information on songs recently played on radio stations.

72. The server of claim 71, wherein the server further comprises a programmed routine that uses the radio station database to automatically generate the playlist for download to the secure digital audio player based on songs recently played on a favorite radio station of the subscriber.

- 72 -

73. A method of assisting in music purchasing decisions comprising:

providing an Internet accessible service that permits a user to specify musical preferences;

accepting a selection of a playlist for downloading, the playlist containing one or more digital music items;

providing a server from which the one or more digital music items in the selected playlist may be downloaded;

providing a digital audio player to the user, the digital audio player comprising a unique identifier, a memory for storing digital music, communication circuitry, audio output circuitry, and a processor, the processor programmed to use the communication circuitry to download the one or more digital music items in the playlist from the server and store the one or more digital music items in the memory, and play the one or more digital music items, wherein the player is incapable of transferring the one or more digital music items to other devices;

tracking and updating musical preferences of the user;

generating a promotional playlist conforming to the musical preferences of the user, and permitting the user to select the promotional playlist;

establishing communication between the digital audio player and the server;

sending the unique identifier to the server to identify the digital audio player; and

downloading the one or more digital music items in the selected playlist from the server to the player;

- 73 -

74. The method of claim 73, wherein providing an Internet accessible service further comprises providing a service that permits the user to purchase compact disks containing music.

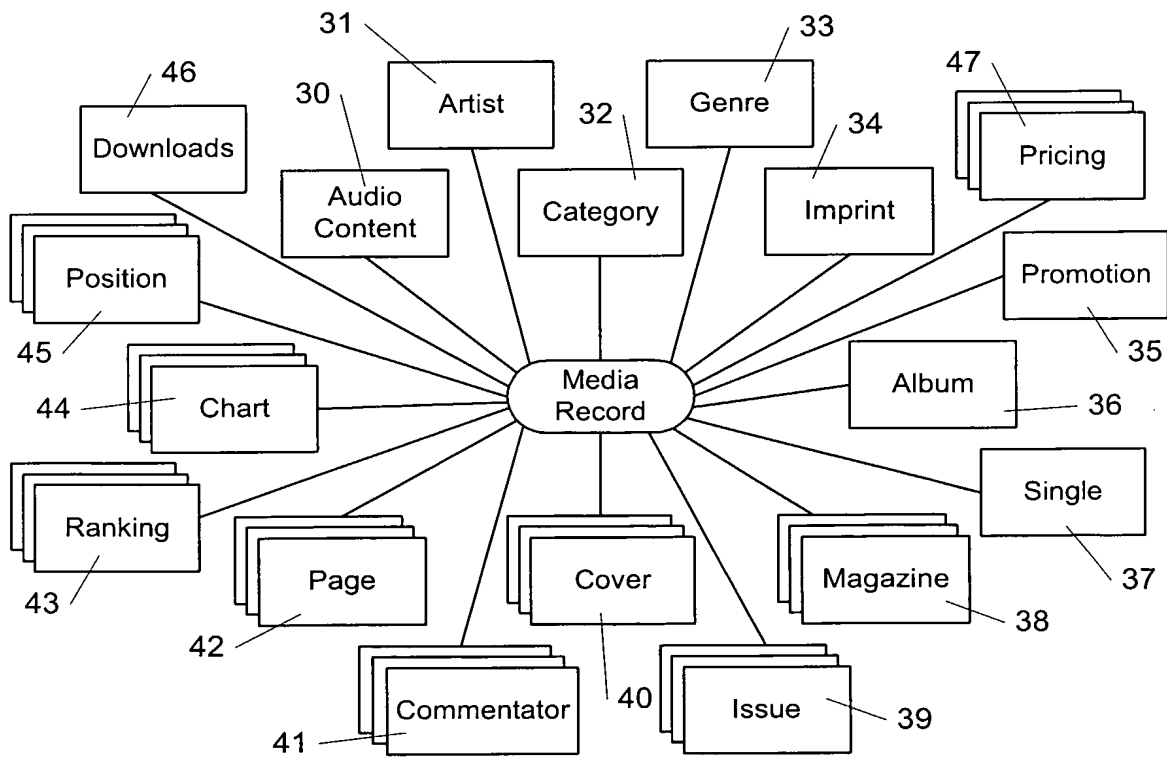



FIG. 2

2/7

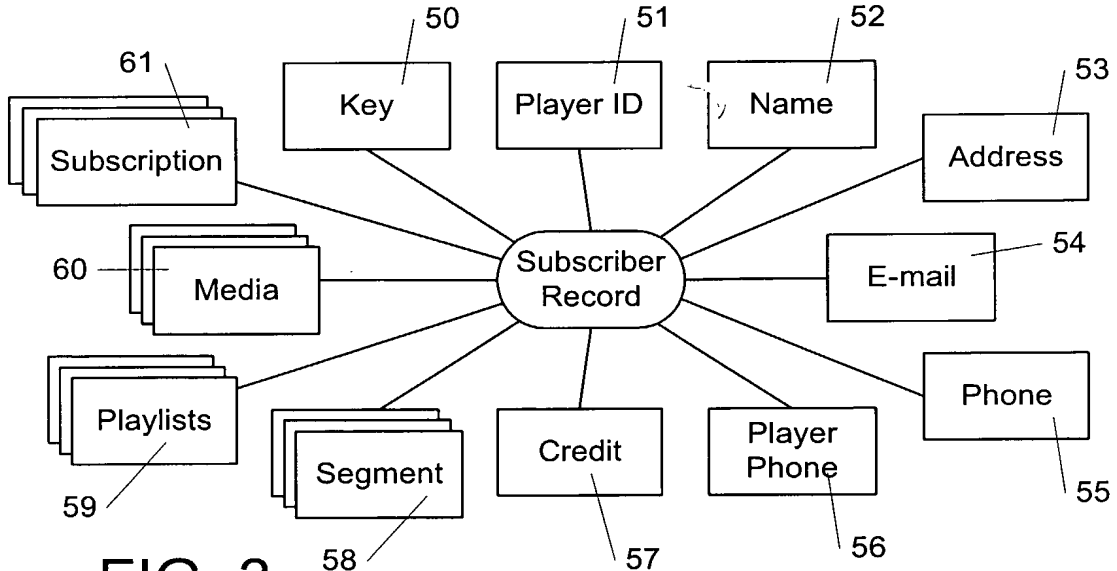


FIG. 3

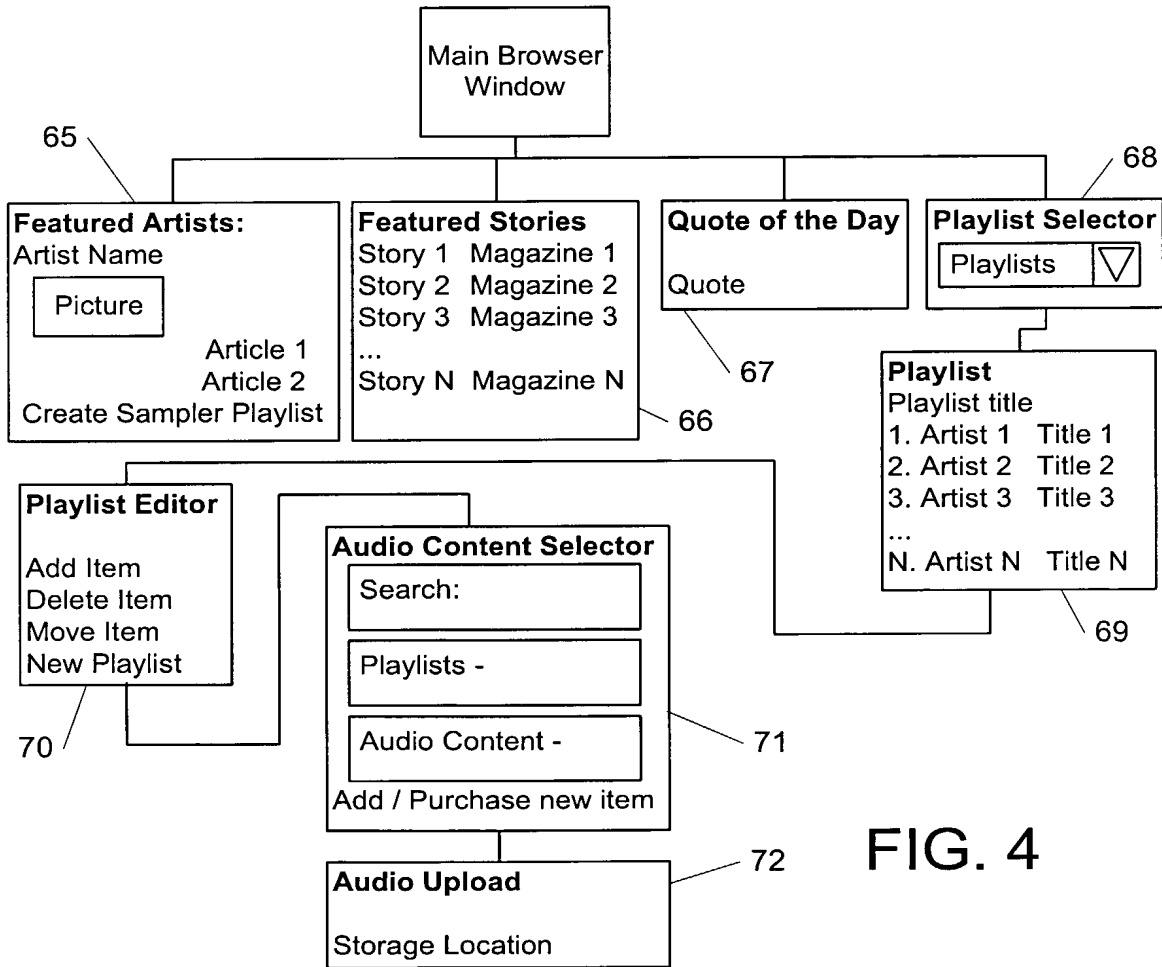


FIG. 4

65

77

68

75

69

66

67

76

Featured Artist

Misery "Manmade" Elliot

Picture

Rolling Stone, 816/817 July 9-22. 1999. p. 54
 Vibe. August 1999, p. 106
 The Source. 118 July 1999. p. 178
 create sampler →

Featured Stories

lost in the stars, The Wire, 104 June 99. p. 22
 SOM's First Step, Billboard, July 10, 1000. p. 1
 Richy Martin, Rolling Stone, 818 August 5, 1999. p. 48
 Geek Love, Hour Music Monthly, August 1999. p. 64
 the bittersweet symphony, Mojo, July 1999. p. 164
 The Industrial Disaster, Select, July 1999. p. 68
 Rodger Carter, Jazz Times, August 1999. p.60
 Approval please-And That, BBC Music, July 1999. p. 27

Quote of the Day

Brian would play Popeye the Sailor Man -
 Nobody could hear him anyway.
 Keith Richards. Mojo July 1999. p. 74

Playlist: Genne Sampler

Alternative, Country, Pop, Urban

Mary J. Slige	All that I Can Say
Jewel	Hands
Everhart	What it's like
Backstreet Boys	I Want it That Way
K-Ci & JoJo	Tell Me it's Real
Pearl Jam	Last Kiss
Kenny Rogders	The Greatest
Jennifer Lopez	If You Had My Love
Sarah McLechian	I will Remember You
Alabama	A Little More Time On You
Matchbox 20	Back 2 Good
Deathny's Child	Bits, Bits, Bits
Offspring	Why Don't You Get a Job?
Serwashed Ladies	Call and Answer
Dixie Chicks	Ready to Run
Sugar Ray	Someday
Lenny Kravitz	Fly Away
Sixpence Home the Richer	Kiss Me
Lit	My Own Worst Enemy
Berry White	Staying Power
Blaque	808
Lauryn Hill	Everything is Everything
Missy Elliott	All in My Girl
Lit Troy	Wanna Be a Bother
Ricky Martin	Livin'La Vida Loca
Enrique Iglosisa	Band
Eve B	Inside Out
Trina & Tamera	Joana
TLC	Unpesty
Smash Mouth	All Star
Nitty Gritty Dirt Band	Bang
Gemstones	So Ammodaus
Ted Beckman	She's So High
Sister Heart	All For You
Shenia Twain	That Don't Impress Me Much
Fearball	The Way
The McGrew	Something Like That
Madana	Beautiful Stranger
Alan Mach	Sit Here
Blessed Union of Souls	Hey Leonardo

BBCast.com

FIG. 5

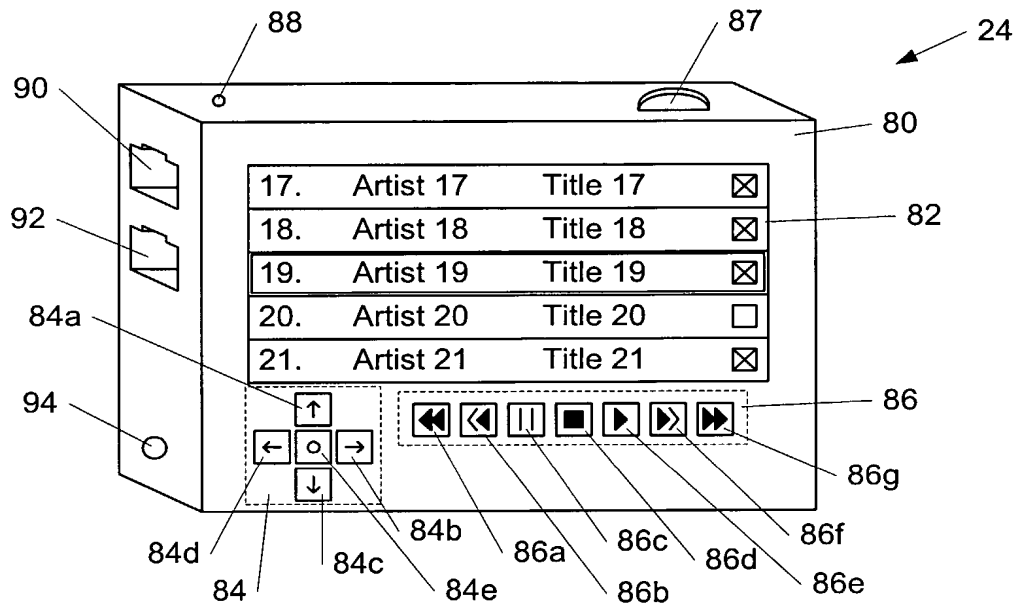


FIG. 6




FIG. 7




FIG. 8




FIG. 9

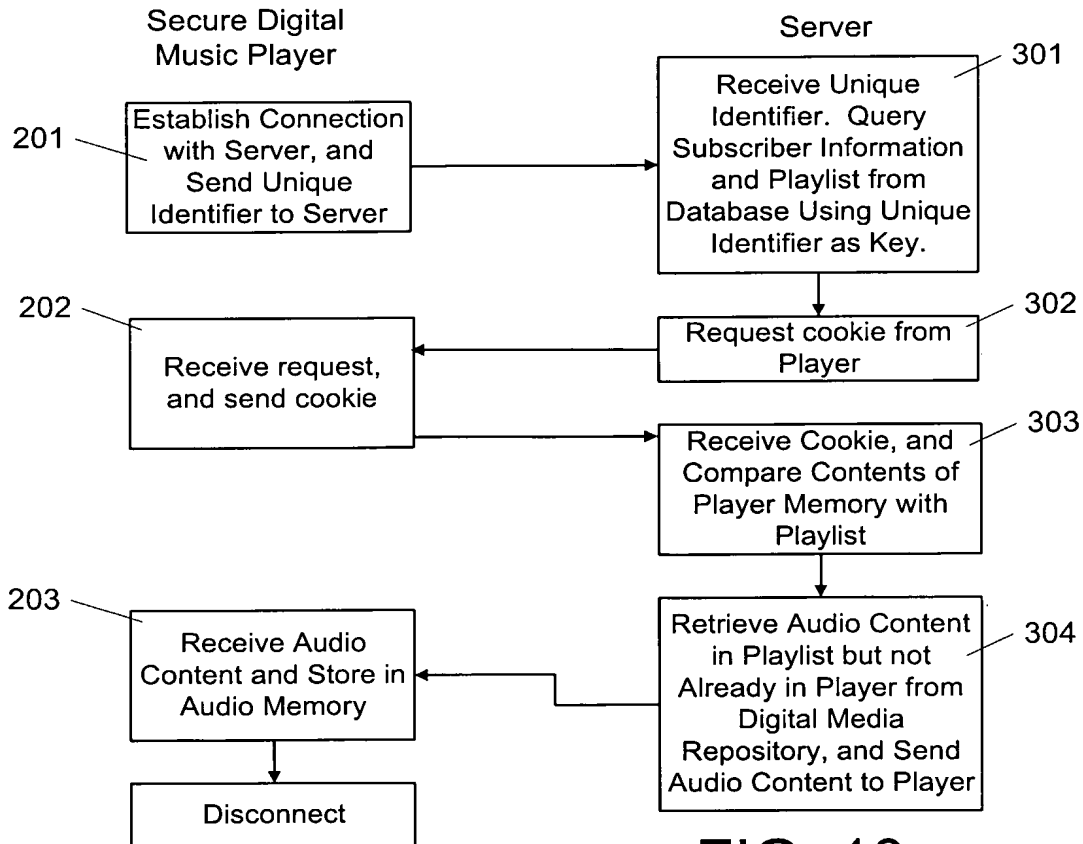


FIG. 10

7/7

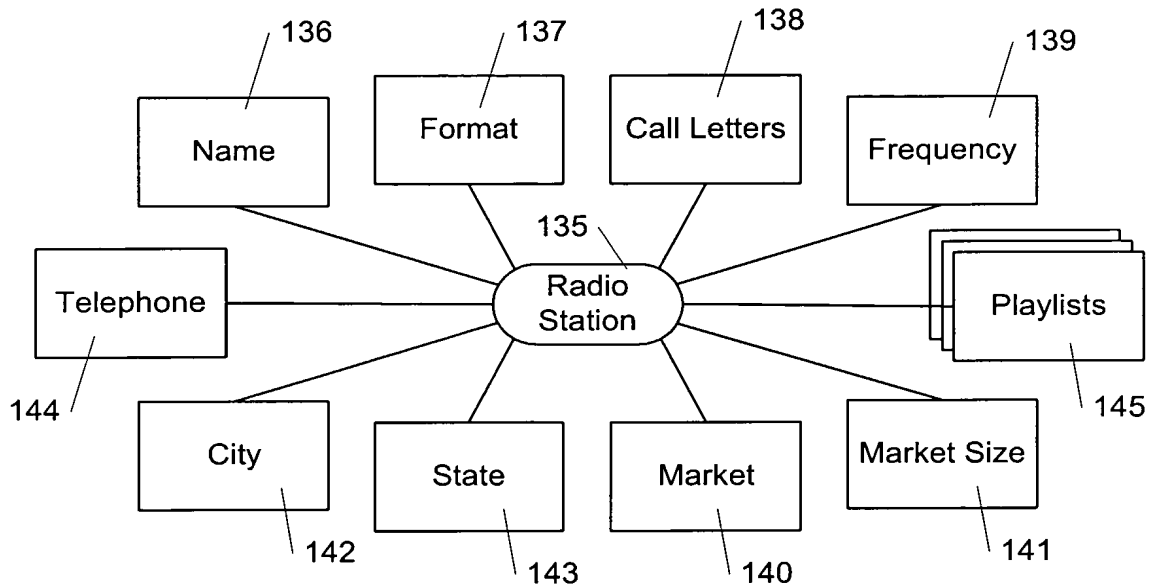


FIG. 11

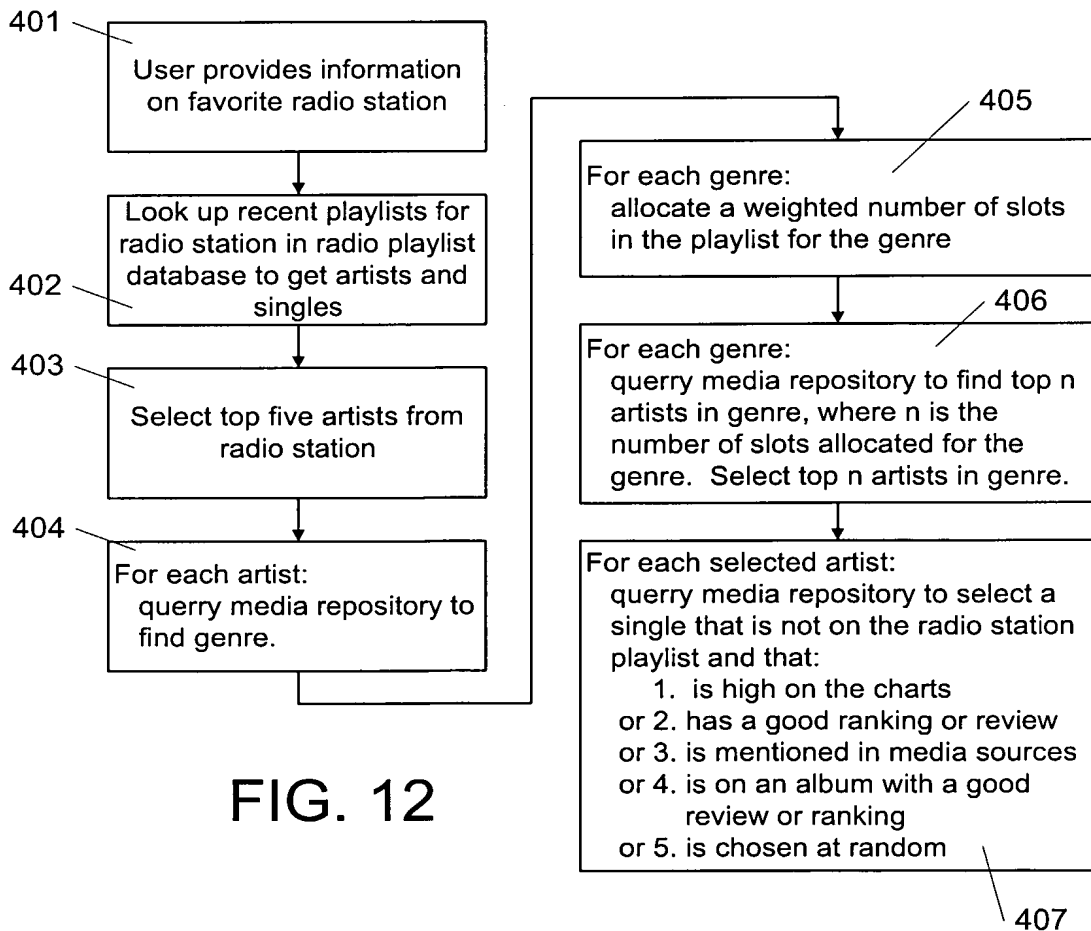


FIG. 12

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US00/32032

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) :G06F 13/00
US CL :709/219, 225, 227, 231
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 709/217, 219, 223, 224, 225, 227, 228, 231, 232, 313, 328, 329

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
EAST

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 5,900,608 A (IIDA) 04 May 1999, cols 1-6.	1-20, 24-74
Y	US 5,926,624 A (KATZ et al) 20 July 1999, cols 3-10.	1-20, 24-74
A	US 5,956,716 A (KENNER et al) 21 September 1999.	1-74
A	US 5,959,945 A (KLEIMAN) 28 September 1999.	1-74

Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"A" document defining the general state of the art which is not considered to be of particular relevance	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"E" earlier document published on or after the international filing date	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"&" document member of the same patent family
"O" document referring to an oral disclosure, use, exhibition or other means	
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search 04 JANUARY 2001	Date of mailing of the international search report 26 JAN 2001
--	--

Name and mailing address of the ISA/US Commissioner of Patents and Trademarks Box PCT Washington, D.C. 20231 Facsimile No. (703) 305-3230	Authorized officer VIET VU <i>James R. Matthews</i> Telephone No. (703) 305-9600
---	--

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
25 October 2001 (25.10.2001)

PCT

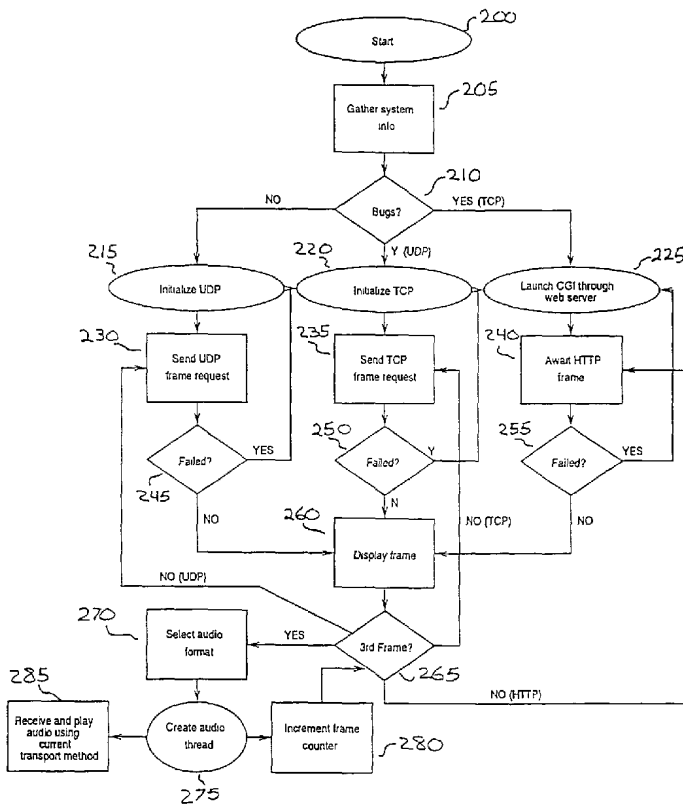
(10) International Publication Number
WO 01/80558 A2

- (51) International Patent Classification⁷: H04N 7/00
- (21) International Application Number: PCT/US01/40452
- (22) International Filing Date: 7 April 2001 (07.04.2001)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
09/548,989 14 April 2000 (14.04.2000) US
- (71) Applicant: SOLIDSTREAMING, INC. [US/US]; 32nd floor, 80 Pine Street, New York, NY 10005 (US).
- (72) Inventor: BASTONE, Daniel; Apt. 4A, 123 East 54th Street, New York, NY 10022 (US).
- (74) Agents: CHAU, Frank et al.; F. Chau & Associates, LLP, Suite 501, 1900 Hempstead Turnpike, East Meadow, NY 11554 (US).
- (81) Designated States (national): AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, 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, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZA, ZW.
- (84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: A SYSTEM AND METHOD FOR MULTIMEDIA STREAMING

Java Audio/Video Client



(57) Abstract: A system is provided for delivering streaming multimedia from a server to users via a communication network. Embedded clients at the users make requests for single datagrams. The server adapts to the variable bandwidths of the users and sends individual datagrams in response to the requesting user at the rate of available bandwidth of the requesting user.



WO 01/80558 A2



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.

A SYSTEM AND METHOD FOR MULTIMEDIA STREAMING**BACKGROUND OF THE INVENTION**1. Field of the Invention

5 The present invention relates to a device and method for delivering video and/or audio data. More particularly to a device and method for delivering video and audio data in real time (streaming) through a communication network.

10 2. Description of Related Art

 As high bandwidth medium such as DSL, cable, and T1 lines becomes more readily available for users of communication networks such as the Internet, content providers quickly move to take advantage of the increased
15 bandwidth to deliver livelier and snazzier content. Countless websites are capable of delivering live (streaming) video or multimedia in real time. Internet users having a high bandwidth connection medium and a desktop PC with a higher speed processor receive the
20 streaming video with little difficulty. Typically, software receivers such as the Real Player by Real Networks or Media Player by Microsoft are installed on the users' PCs. The receivers receive and decode the encoded video data

delivered by content provider websites and display the decoded data as streaming video to the user.

At the content providers end, the video is first compressed or encoded. A popular video/audio compression
5 format is MPEG, which is the format decoded by Microsoft's Media Player. MPEG (Moving Picture Experts Group) format is used in, for example, the Real Player and the Microsoft Media Player.

The MPEG technique for compressing digital video
10 includes use of Discrete Cosine Transform (DCT), Quantization, Huffman coding, and Motion Compensated Predictive coding, in which the differences in what has changed between an image and its proceeding image are calculated and only the differences are encoded. Predictive
15 coding requires interframe processing, i.e., data from neighboring frames are needed to successfully encode and decode an image; therefore, individual frames must be temporarily stored in a buffer and the image is encoded and decode using multiple frames. Buffering allows a server to
20 send data at a constant rate, regardless of the rate at which the client displays the data. A disadvantage of MPEG and the buffering technique is that a considerable amount of memory is needed for buffering. In portable devices, sufficient memory may not exist.

Another video compression technique, known as JPEG (Joint Photographic Expert Group), employs the MPEG coding process except predictive coding. Thus, JPEG compression does not rely on interframe processing, i.e., each frame is independently processed and has no processing relationship to another frame. Therefore, JPEG compression does not require frame buffering.

Various methods and protocols for delivering data are available depending on the bandwidth of the connecting medium. One example is the Transmission Control Protocol (TCP). The TCP typically functions in conjunction with the Internet Protocol (IP). The TCP provides reliable, stream-oriented connections that hide most of IP's shortcomings; i.e., the basic nature of IP cannot guarantee the data will be delivered correctly. The TCP/IP protocol suite gets its name because the TCP protocol is layered on top of the IP protocol. The TCP layer interfaces on one side to application processes and on the other side to the IP protocol.

TCP data is organized as a stream of bytes, much like a file. The datagram nature of the network is concealed. A mechanism (the Urgent Pointer) exists to let out-of-band data be specially flagged. Sequence numbers are used to coordinate which data has been transmitted and received. TCP will arrange for retransmission if it determines that

data has been lost. This method provides for reliable delivery. TCP will dynamically learn the delay characteristics of a network and adjust its operation to maximize the throughput without overloading the network, this gives TCP the quality of network adaptation. TCP manages data buffers, and coordinates traffic so its buffers will not overflow. Fast senders will be stopped periodically to keep up with slower receivers, resulting in flow control.

TCP operates in both directions (full duplex) and in an almost completely independent manner, akin to two independent byte streams traveling in opposite directions. No TCP mechanism exists to associate data in the forward and reverse byte streams. Only during connection start and close sequences can TCP exhibit asymmetric behavior, i.e., data transfer in the forward direction but not in the reverse, or vice versa.

Each endpoint of a TCP connection will have a buffer for storing data that is transmitted over the network before the application is ready to read the data. This lets network transfers take place while applications are busy with other processing, improving overall performance. To avoid overflowing the buffer, TCP sets a Window Size field in each packet it transmits. This field contains the amount of data that may be transmitted into the buffer. If this number falls to zero, the remote TCP can send no more data. It must

wait until buffer space becomes available and it receives a packet announcing a non-zero window size.

Sometimes, the buffer space is too small. This happens when the network's bandwidth-delay product exceeds the
5 buffer size. The simplest solution is to increase the buffer, but for extreme cases the protocol itself becomes the bottleneck (because it doesn't support a large enough Window Size). Under these conditions, the network is termed an LFN (Long Fat Network).

10 When a host transmits a TCP packet to its peer, it must wait a period of time for an acknowledgment. If the reply does not come within the expected period, the packet is assumed to have been lost and the data is re-transmitted. The time that the protocol will wait for a reply is a
15 variable. Over an Ethernet, no more than a few microseconds should be needed for a reply. If the traffic must flow over the wide-area Internet, a second or two might be reasonable during peak utilization times. If a communication device is on a satellite traveling toward Mars, minutes may be
20 required before a reply.

Round-Trip Time (RTT) estimates are an important performance parameters in a TCP exchange, especially when dealing with an indefinitely large transfer. All TCP implementations eventually drop packets and retransmit them,
25 no matter the quality of the link. If the RTT estimate is

too low, packets are re-transmitted unnecessarily; if too high, the connection can sit idle while the host waits to timeout.

The User Datagram Protocol (UDP) is used in higher
5 bandwidth communication links. UDP is a connectionless
protocol that, like TCP, runs on top of an IP network.
Unlike TCP/IP, UDP/IP provides very few error recovery
services, offering instead a direct way to send and receive
datagrams over an IP network. UDP is used primarily for
10 broadcasting messages over a network. UDP packets are
delivered like IP packets; connectionless datagrams that may
be discarded before reaching their targets. UDP is useful
when TCP would be too complex, too slow, or just
unnecessary.

15 UDP provides a few functions beyond that of IP. For
example, Port Numbers. UDP provides 16-bit port numbers to
let multiple processes use UDP services on the same host. A
UDP address is the combination of a 32-bit IP address and
the 16-bit port number. Unlike IP, UDP does checksum its
20 data, ensuring data integrity. A packet failing checksum is
simply discarded, with no further action taken.

A common gateway interface (CGI) is one way for a web
server to pass a web user's request to an application
program, which in turn passes data back to be forwarded to
25 the user. When the user requests a web page (for example, by

clicking on a hypertext link), the server retrieves the requested page, sending the page to the client. However, when a user submits a form on a web page, it usually needs to be processed by an application program. The web server typically passes the form information to a small application program (applet) that processes the data and may send back a confirmation message. This method for passing data between the server and the application is called the common gateway interface (CGI). The CGI is part of the web's Hypertext Transfer Protocol (HTTP).

For a client system having a lesser connection bandwidth, such as connection through a 56K modem, specific communication protocols can be established by an application program predownloaded or installed in the client's computer. A CGI can be used to pass a client's request to the application program. The CGI provides a consistent way for data to be passed from the user's request to the application program and back to the user. In other words, CGI operates in conjunction with clients and servers regardless of which operating system (OS) is being used by the parties, for example, Windows, Macintosh, UNIX and Linux, OS/390, or others. A CGI application may be written in a number of different languages.

An alternative to a CGI application is Microsoft's Active Server Page (ASP), in which a script embedded in a web page is executed at the server before the page is sent.

In a desktop client environment in which memory speed
5 and connection bandwidth are sufficient, the connection rate is stable and servers can 'push' the MPEG-type datagrams to the clients synchronously, i.e., frames are buffered at the server and the buffer content is dumped or transmitted to the clients at a substantially periodic rate. At the
10 desktop PC, the datagrams are also buffered because a single image depends on several datagrams.

As wireless applications and devices grow in popularity, the limitations of wireless applications and devices such as narrow bandwidth and limited memory and
15 processing capacity must be addressed. In particular, content providers who wish to deliver substantially the same streaming video contents to wireless users as desktop users must find a viable solution to overcome these limitations. Therefore, a need exists for a system and method for
20 delivery of streaming media to the client regardless of the client's available bandwidth.

SUMMARY OF THE INVENTION

These and other objects, features, and advantages of
25 the present invention will become apparent from the

following detailed description of illustrative embodiments thereof, which is to be used in connection with the accompanying drawings.

A method for streaming video data over a network in real time is provided. The method includes initializing a transport mode for the video data, sending a data request for a single frame of video data from a client to a server, retrieving the single frame from a memory at the server, and sending the video data to the client.

The step of initializing also includes listing available transport modes for the client, determine whether incompatibilities exist between the available transport modes and software, choosing a transport mode from the list, and initializing parameters of the transport mode at the client for client control of video steaming.

Initializing is performed by a client application which is capable of running in different operating systems. Alternatively, initializing is performed by a client embedded in a web page. The client embedded in the web page can be a common gateway interface and an active server page. The transport mode is chosen from the following, a UDP, a TCP, and an HTTP, however other modes are contemplated.

The steps of sending a data request for a single frame from the client to a server, retrieving the video data from

a shared memory, and sending the retrieved video data to the client, are repeated for each video frame.

Storing video includes capturing a thread from a specified source, and storing the captured thread in the server's shared area of memory.

According to another aspect of the present invention, a storage medium having a stored program which is executable by a processor for causing the processor to perform method steps for streaming video communication is also provided.

The method steps comprising requesting a packet representing a single datagram from a server over a communication network, receiving a requested packet, processing and displaying the requested packet, incrementing a datagram frame counter, requesting a next packet based on the frame counter value from the server, and asynchronously processing and displaying the next packet when received.

The communications between the processor and the server is preferably by Wireless Application Protocol (WAP). The server is accessed by said processor via HTTP. The packet representing a datagram is preferably JPEG encoded, and the step of asynchronously processing and displaying is independent of data from the step of processing and displaying the requested packet.

An apparatus is also provided for communicating streaming video data between a plurality of users and a

server connected by a communication network, comprising a stored program executable by a processor in said server for causing the server to receive requests for individual datagrams from the plurality of users and forwarding individual datagrams in response to each request to the user making the request at a rate based on available bandwidth of the user making the request.

BRIEF DESCRIPTION OF THE DRAWINGS

10 The preferred embodiments are described with reference to the drawings wherein:

 FIG. 1 is a diagram of a system for streaming data according to one embodiment of the invention;

 FIG. 2 is a flow diagram of a preferred embodiment of the invention for streaming audio/video data to a client;

 FIG. 3 is a flow diagram of a preferred embodiment of the invention for streaming video data;

 FIG. 4 is a flow diagram of a preferred embodiment of the invention for streaming audio data; and

20 FIG. 5 is a flow diagram of a method for streaming audio data over a CGI HTTP transport mode.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

Preferred embodiments of the present invention will be described below in more detail with reference to the accompanying drawings.

5 The present invention relates to a system and method for streaming video. The invention is implemented over a network of processors. The network can include, a local area network (LAN), a wide area network (WAN), an Intranet, an Internet, a wireless network, or the like. These networks
10 can be in any configuration including, for example, star, ring, and bus.

For purposes of the present invention a client will, by definition, be receiving and displaying data, while the server will be providing data to the client. A system and
15 method of the present invention enables the client to receive audio/video data in real time regardless of the bandwidth available to the client. This is achieved through the implementation of a dynamic bandwidth adaption method for managing the flow of data. As a result, devices such as,
20 PDAs, hand-held PCs, and various mobile devices with little bandwidth are able to receive streaming video in real time. It is readily appreciated by one ordinarily skilled in the art that the invention is not limited to these devices and is also suitable to desktop computers, servers, and the
25 like.

A system and method according to the present invention can be construed as a browser that allows a client to access the server's network to receive streamed data content. The browser is preferably WAP (Wireless Application Protocol) compatible, but can be deployed to any device including those which are not WAP compatible. The browser preferably adapts to different codecs which do not require interframe processing, for example, Motion JPEG and Wavelet, allowing the flexibility to port software to embedded devices having limited memory resources. Frame buffering at the server can be dispensed with because each frame can be independently processed and delivered.

As shown in FIG. 1, the server 110 having audio/video data is connected to a network 100, in this example a bus network. While FIG. 1 depicts a bus network 100, any other network capable of supporting a server and client is contemplated by the invention. Alternatively, both the server and client may be embodied in the same computer and operate without a network. Typically a server is a computer program that provides a service to another computer program, the client. A server can function as a client and a client as a server depending on whether services are being offered or requested. The server, for purposes of the invention, is a stored program including a script for downloading audio/video data and an applet. A client 120 or

130 is also connected to the network. The client 120 captures the data from the server 110. The capture can take place from a local capture card, a local looped file, a local file on-demand, or a remote IP address. Each will be explained below in more detail. It should be readily apparent to one ordinarily skilled in the art that the client is described for client 120 but the embodiment is applicable for multiple clients.

The client 120 is a script which may be stored in the memory of a Personal Digital Assistant (PDA), PC or embedded in a web page. The client is a application which is capable of running in different operating systems, for example, Windows CE, Palm OS, Nokia PDA's, Windows, Linux etc. Alternatively, the client is embedded in a web page, for example, a common gateway interface (CGI) or a Microsoft Active Server Page (ASP). A CGI may be written in different programming languages, for example, C, C++, Java, and Perl, though this is not a complete list of possible applicable languages. Execution of the client ensures cross-platform and cross-browser functionality. Referring to FIG. 2, upon startup 200, the client gathers information 205 regarding its own operating environment needed to choose available transport modes. The information is checked against known bugs or incompatibilities 210 with certain operating systems (OS) and browser combinations. An example of a bug includes,

the Java Interpreter used in Microsoft Internet Explorer v4.0 which has broken implementation of User Data Protocol (UDP). An appropriate transport mode is chosen according to the incompatibilities 210. The client will detect this and
5 default to suitable protocol, for example, Transmission Control Protocol (TCP) or Hypertext Transfer Protocol (HTTP) modes. The method of connection, e.g.UDP, TCP, HTTP, are referred to as a transports. The client initializes parameters which control the transport mode. This allows the
10 same client to operate all possible configurations, for example, audio, video, audio/video, on-demand, live, or other similar configurations.

Assuming that no incompatibilities have been detected, the client attempts a UDP connection 215 to the server to
15 retrieve a frame of video 230. In one embodiment, if the connection fails 245, a second attempt 220 is made to establish a UDP connection. After a specified number of attempts to establish a UDP connection, a TCP connection will be attempted. Note that in the present example, UDP is
20 preferable to TCP connections based on bandwidth, however other connections may provide superior bandwidth than UDP, in such a case the wider bandwidth connection would be attempted first and the UDP second. In an alternative embodiment, an attempt to establish a TCP connection is made
25 after the first failed UDP connection attempt.

Likewise, if the TCP connection fails 250 to retrieve a frame 235 then an attempt to establish a HTTP connection 225 is made. Alternatively, more than one attempt to establish a TCP connection will be made. Further attempts to establish a connection are made with progressively narrower transports until the transport mode with a narrowest acceptable bandwidth has been reached. For example, a HTTP connection. Attempts to establish this transport mode will be continue until a connection is established, 240/255. At this point it is assumed that the server is down or the network has failed. When the server comes up or the network improves, video will begin streaming once the connection has been established. In an alternative embodiment, the transport mode with the narrowest acceptable bandwidth will be attempted a specified number of times.

Once a successful video connection has been established and an audio format has been selected, an audio connection will be made and display the data will begin 260.

If at any point a connection fails 265, a connection will be attempted with the next bandwidth, for example from UDP to TCP. While in the alternative transport mode, the mode with a lower bandwidth, periodic attempts to establish a connection with the original transport will be made. No user intervention is needed during operation of the method.

Alternatively, if the connection is maintained, the datagram frame counter is incremented 280. A request is then made for the next packet based on the frame counter value from the server. The steaming video data will continue until
5 the client terminates the connection with the server.

If displaying a stream on-demand, the client is given the ability to fast forward, rewind, and jump to arbitrary parts of the stream. A clock is also displayed indicating the current location in the stream. The client initially
10 sends a query to the server requesting the length of the stream in frames. Proceeding with each frame, the server sends the frame number, allowing the client's clock to remain accurate regardless of the speed of the stream. When the user wishes to fast forward, rewind, or jump to a
15 specific location in the stream, it sends a request with the desired target frame number to the server. The server responds by seeking to the requested frame number and streaming from there. All other functions of rate adaption and protocol selection operate as normal.

20 An advantage to the present invention, related to the client's ability to move on the stream, is error resilience. That is, since each video frame is independent, if an error occurs in a frame the invention can either attempt to fix the error or drop the frame. When a frame is dropped, the
25 invention requests the next frame and preserves the

continuity of the video and/or audio display. This is especially advantageous in devices having limited processing power, and therefore, lacking the resources to fix errors.

Having described the role of the client above, the method will be described in reference to various types of capture. Other types of capture may also be implemented in the present invention.

The server may capture a thread 305 from a specified source, e.g., a local capture card. Example of a local capture card include the SunVideoPlus Osprey and the SunVideo (sun) capture board. In this instance, the server makes use of native Solaris threads to achieve rate-adaptive connections to the clients. A thread is a placeholder information associated with a single use of a program that can handle multiple concurrent users. From the program's point-of-view, a thread is the information needed to serve one individual user or a particular service request. If multiple users are using the program or concurrent requests from other programs occur, a thread is created and maintained for each of them. The thread allows a program to know which user is being served as the program alternately gets re-entered on behalf of different users. One way thread information is kept is by storing it in a special data area and putting the address of that data area in a register. The OS always saves the contents of the register when the

program is interrupted and restores it when it gives the program control again.

On startup, the server creates two threads which capture video 300 and audio 400 from the specified source. The server places the captured data 310 into a shared area of memory 315 which is accessible to all other threads within the server. Other threads are created which accept incoming requests for each transport type. These incoming requests for each client are handled as part of a non-blocking loop. Each transport type is handled differently.

For example, for UDP video 325 connections, the non-blocking loop accepts and services connections. Since UDP is "connectionless," there is no need to maintain a persistent connection to the client and each client connection is really a request for a single frame. The thread receives a single byte datagram from a client 340, immediately retrieving 355 and sending 370 back the frame currently stored in the shared memory segment 315. A datagram is a self-contained, independent entity of data carrying sufficient information to be routed from the source to the destination computer without reliance on earlier exchanges between this source and destination computer and the transporting network. This process continues indefinitely.

Since sending a UDP datagram is a non-blocking operation, the server need only have one thread. Thus, the

amount of CPU time and memory required from UDP is drastically less than that needed for other connections.

Another example is TCP connections 330. The server waits for a new connection to be established with a client 345. A non-blocking loop accepts incoming requests. The client corresponding to the requests are arranged in a connection pool list. Since TCP is a connected protocol, an independent thread is then created to service 360 each client. The service threads act like the UDP thread, awaiting 375 a single byte (request) from the client, the requested frame is retrieved 385 by the server from the shared memory 315. The retrieved frame is then sent 390 to the client.

The above two examples are rate adaptive solutions. By allowing the clients to control the flow of video frames rather than simply pushing each frame to each client as it is captured, data bottlenecks typically associated with streaming media over the Internet are eliminated.

As an illustration, two desktop PCs are connected to a server. The first by 100Mbps Ethernet, the second by a 28.8Kbps dialup Internet connection. The 100Mbps machine will receive video at the rate it is captured, up to 30fps. The rate of speed of the video is due in great part to the available bandwidth, the 100Mbps sends out frame requests to the server fast enough to allow for a 30 fps stream.

However, the 28.8Kbps machine requires more time to receive each frame and therefore sends out frame requests less frequently. Thus the rate of speed is about 3 to 4fps, and both machines will be at the same place in the stream. The illustration method of the present invention accomplishes this through client side requests for current frames. These requests may or may not be for the next sequential frame. Comparing the machines, the 100Mbps machine will show a higher quality video, while the 28.8Kbps machine will appear to be skipping frames in order to maintain a real time stream.

This approach has several advantages over conventional streaming methods. Current methods require re-encoding the video at several frame rates in order to support users with different bandwidth availability. The user is also required to select a rate beforehand that is appropriate for their connection. In contrast, the illustrative method according to the present invention needs only one rate of capture and encoding. A user no longer needs to know anything about their network bandwidth. The user will connect to the stream at a rate which adapts itself to the environment. Additionally, degraded network conditions will not create a bottleneck and will not cause the stream to fall behind real time, rather the frame rate will decrease and more frames will be skipped. When conditions improve, so will the frame

rate. Further, this approach will take advantage of high-bandwidth consumer access devices, such as xDSL and 38GHz wireless connections without any changes or updating the server of client.

5 Other transport modes are supported by the invention, for example, HTTP 335. HTTP is a one-way protocol and thus cannot perform true rate adaption. It is included for clients that do not have UDP, TCP, or another transport available to them. This may be due to a network firewall, or
10 a packet filter for example. Further, some Internet Service Providers (ISP) may not support connections on certain ports, or not support UDP at all. HTTP connections are achieved through the use of an external program which is activated by a web server 500, e.g. common gateway interface
15 (CGI). A main server creates one thread to accept and service connections from this external program via local UNIX domain sockets and UDP 505. Although UDP is used here, it is only for local redirection of frames from the main server to the CGI. The server waits for a CGI request 350.
20 The CGI requests frames at a steady rate from the main server 510. As the main server retrieves frames 365 from the shared memory 315, the frames are re-broadcast to the client 380 back through the web server 520 via HTTP. To avoid bottlenecks, from the audio push 525, this transport mode

transmits 380/515 at a default rate of 1fps in order to remain real time regardless of available bandwidth.

Another capture can be from a local file which is looped. This method obtains video and audio data from a single file that has been pre-encoded from a capture board and stored on the server. It operates exactly as above and uses the exact same thread structure. Data is read in from the specified file and placed into the shared memory segments at the rate it was encoded. When the end of the file is reached, it re-starts from the beginning.

Still another capture can be had from a local file on-demand. This method deals with more than one source of audio/video data being sent to all clients. Each client needs its own unique copy of the server which then acts as described above using the requested file. This is accomplished by creating another layer of threads which create entire server sets of threads for each client. While this creates a high amount of overhead, the data is being delivered in an compressed and encoded state, therefore it is not necessary to make deliveries in real time. In addition, the server returns the number of frames in the stream to the client as well as the frame number of each frame sent, therefore, the client's clock remains accurate. The server also receive requests from the client indicating that it wishes to jump forward of backward to a specific

frame, the server seeks to this location in the file and continues streaming.

Yet another capture is from a remote IP address, according to this method the server receives data from another server on a remote machine. The remote machine may be located anywhere on the network that is accessible from the local machine's network. The local server acts as a single client to the remote server, and pulls a stream from the remote server in the same rate-adaptive fashion as the Java client. The local server then places the data into its shared memory buffers, and re-broadcasts it as if it were being captured locally. All other functions of the local server operate as normal.

Because the local server is acting as a regular client, it does not matter how the remote server is captures data, it can be from any of the methods described above. It can even be capturing its data from another remote server, allowing for server chains to be created that, for example, re-broadcast a source feed on different networks.

Like video 300, audio data 400 may be transported by the methods described above: UDP 440, TCP 445, and HTTP 450. However, it is pushed data. As new audio data is captured 410, it is immediately sent out to all clients 430. This allows for a steady audio stream which will inherently be in sync with the video played in real time according to the

invention. The server also attempts to adapt the streamed audio to the clients available bandwidth. This is accomplished by the server making available multiple types of audio simultaneously, for example, raw uncompressed 410, Global System for Mobile communication (GSM) encoded audio 415, and G.728 encoded audio 420 (G728 is specified in ITU-T recommendation G.728, "Coding of speech at 16Kbit/s using low-delay code excited linear prediction"). Uncompressed audio 410 occupies 64k of bandwidth, G.728 420 occupies 16Kbps of bandwidth, and GSM 415 occupies 13Kbps of bandwidth. The client times the amount of time between two frames 265, for example, the second and third frames, and based on this time selects the appropriate available audio format 270. The server creates a thread 435 to accept incoming audio connections via each of the available transport modes: UDP 440, TCP 445, and HTTP 450. The server waits for a new connection 455 a, b, c to be established by a client. Upon connection, the client sends preferred audio format information to the server 460 a, b, c. The server then creates a service thread 275/465 a, b, c, for the client which delivers audio in the requested format. These service threads wait for a signal from the corresponding audio encoding thread 470a, b, c, that more audio has become available. The server retrieves audio data from a selected buffer 475a, b, c, then the selected buffer is sent to the

client 285/480a, b, c, and immediately wait for another signal 470a, b, c. HTTP connections will automatically default to the lowest bandwidth signal, since a degraded connection is assumed.

5 Advantageously, the system and method according to an illustrative embodiment of the invention functions in low bit rate environments, at about 9.6 Kbps, and suffers no degradation even during a transmission over a 14.4 Kbps network. Theoretically there is no upper bound as the
10 system utilizes an adaptive bandwidth method. That is, the system scales the data stream to the available bandwidth so that as bandwidth increases the stream rate will increase accordingly. The flow of data from the server to the client is intelligently managed so as to completely eliminate
15 network bottlenecks traditionally associated with streaming media. This is especially beneficial in wireless environments where available bandwidth may be limited or inconsistent. The end user is allowed to receive multimedia data without fore knowledge of the available bandwidth while
20 maintaining a real-time stream. This also frees the content provider from having to provide multiple media streams to accommodate differing connection speeds. In order to control the drain of the server's network capacity, the system and method provides for a cap. The cap is utilized by the server

to restrict usage of the stream to maintain the integrity of the server's network.

Having described preferred embodiments of a system and method for multimedia streaming, it is noted that
5 modifications and variations can be made by persons skilled in the art in light of the above teachings. It is therefore to be understood that changes may be made in the particular embodiments of the invention disclosed which are within the scope and spirit of the invention as outlined by the
10 appended claims. Having thus described the invention with the details and particularity required by the patent laws, what is claimed and desired protected by Letters Patent is set forth in the appended claims.

WHAT IS CLAIMED IS:

1. A method for streaming video data over a network in real time, comprising:

initializing a transport mode for the video data;

5 sending from a client to a server a data request for a single frame of video data;

retrieving the single frame from a memory at the server; and

sending the video data to the client.

10

2. The method for streaming video in Claim 1, wherein the steps of sending a data request for a single frame from the client to a server, retrieving the video data from a shared memory, and sending the retrieved video data to the client, are repeated for each video frame.

15

3. The method for streaming video as in claim 2, wherein said each video frame is processed for display independently from processing of another video frame.

20

4. The method for streaming video in Claim 1, wherein the step of initializing further comprises:

determining a list of available transport modes for the client;

determine incompatibilities between the available
transport modes and software;

choosing a transport mode from the list; and

initializing parameters of the transport mode at the
5 client for client control of video streaming.

5. The method for streaming video in Claim 4, wherein
the step of initializing is performed by a client
application which is capable of running in different
10 operating systems.

6. The method for streaming video in Claim 4, wherein
the step of initializing is performed by a client embedded
in a web page.

15 7. The method for streaming video in Claim 6, wherein
the client embedded in the web page is chosen from one of a
common gateway interface and an active server page.

20 8. The method for streaming video in Claim 1, wherein
the transport mode is chosen from the group consisting of a
UDP, a TCP, and a HTTP.

25

9. The method for streaming video in Claim 1, wherein storing video further comprises:

capturing a thread from a specified source; and

5 storing the captured thread in the server's shared area of memory.

10. A storage medium having a stored program which is executable by a processor for causing the processor to perform method steps for streaming video communication, the method steps comprising:

10 requesting a packet representing a single datagram from a server over a communication network;

receiving a requested packet;

processing and displaying said requested packet;

15 incrementing a datagram frame counter;

requesting a next packet based on the frame counter value from the server; and

asynchronously processing and displaying said next packet when received.

20

11. The method of claim 10, wherein communications between said processor and said server is by Wireless Application Protocol (WAP).

12. The method of claim 10, wherein said server is accessed by said processor via HTTP.

13. The method of claim 10, wherein said packet
5 representing a datagram is JPEG encoded.

14. The method of claim 10, wherein said step of asynchronously processing and displaying is independent of data from said step of processing and displaying said
10 requested packet.

15. An apparatus for communicating streaming video data between a plurality of users and a server connected by a communication network, comprising:

15 a stored program executable by a processor in said server for causing the server to: receive requests for individual datagrams from the plurality of users; and

forwarding individual datagrams in response to each request to the user making the request at a rate based on
20 available bandwidth of the user making the request.

16. The apparatus according to claim 15, wherein said plurality of users are wireless mobile devices.

17. The apparatus according to claim 15, wherein said server is accessible via HTTP.

5 18. The apparatus according to claim 15, wherein the datagrams are JPEG encoded.

1/5

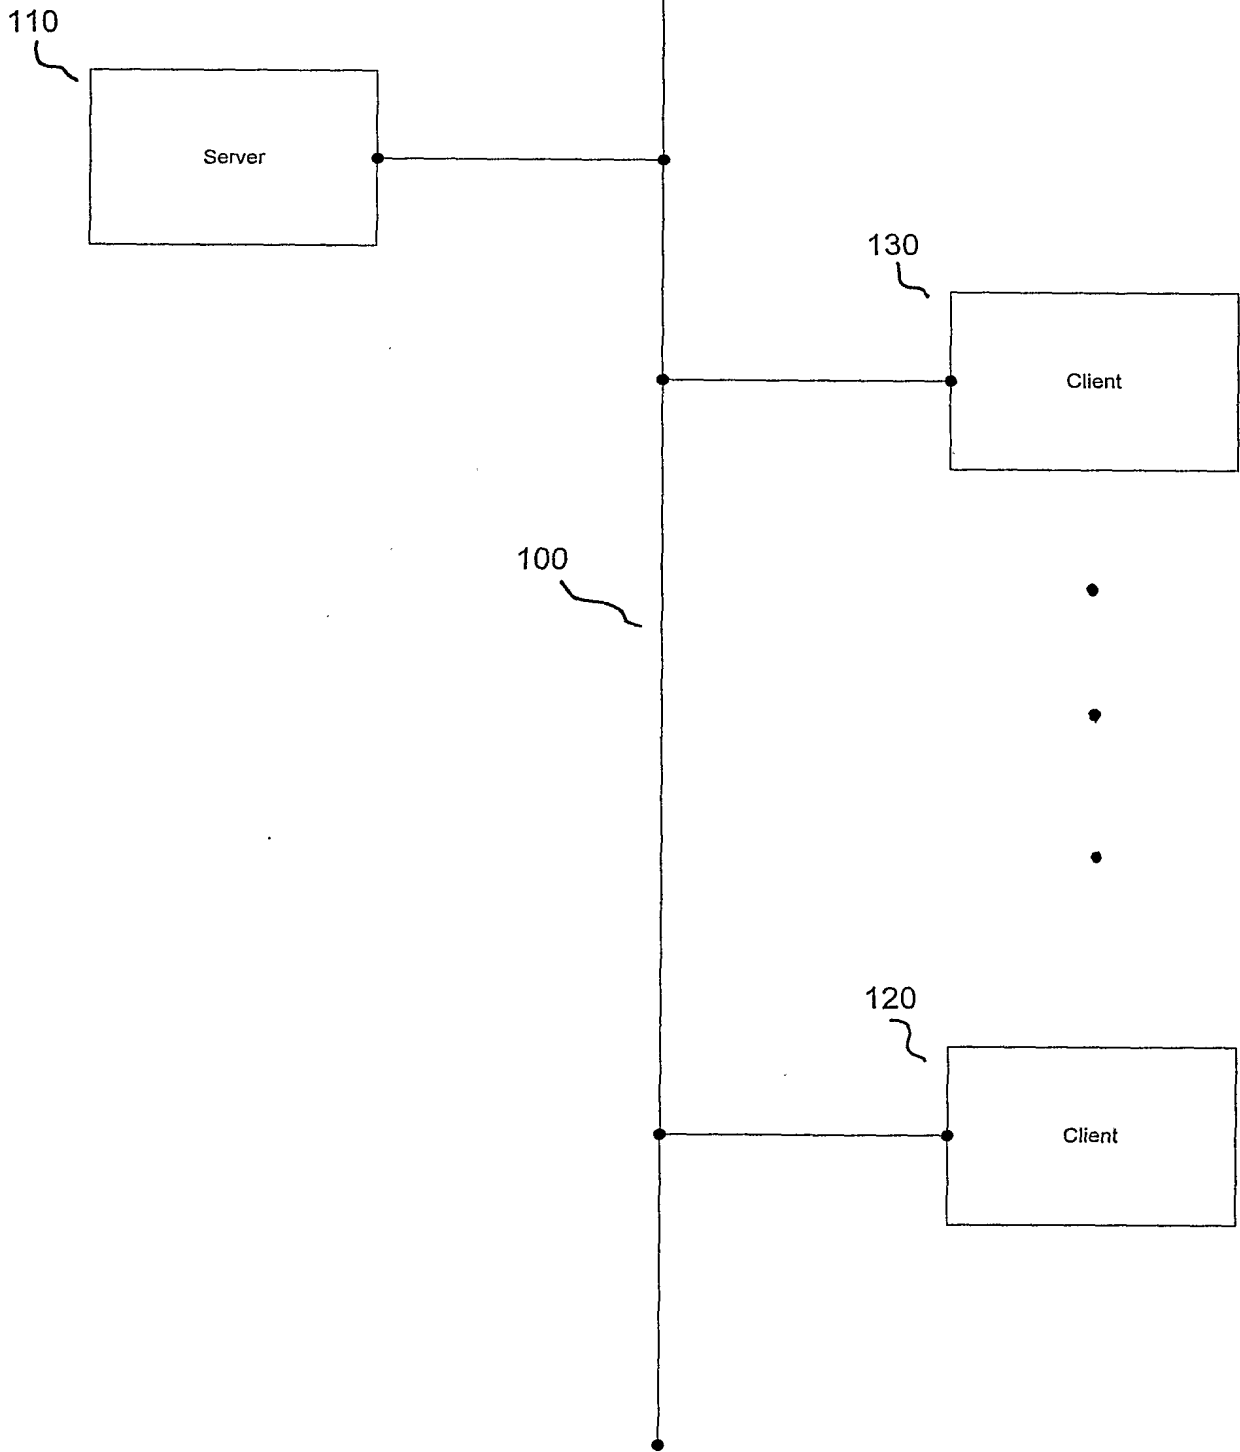


FIG. 1

Java Audio/Video Client




FIG. 2

Multithreaded Audio/Video Server

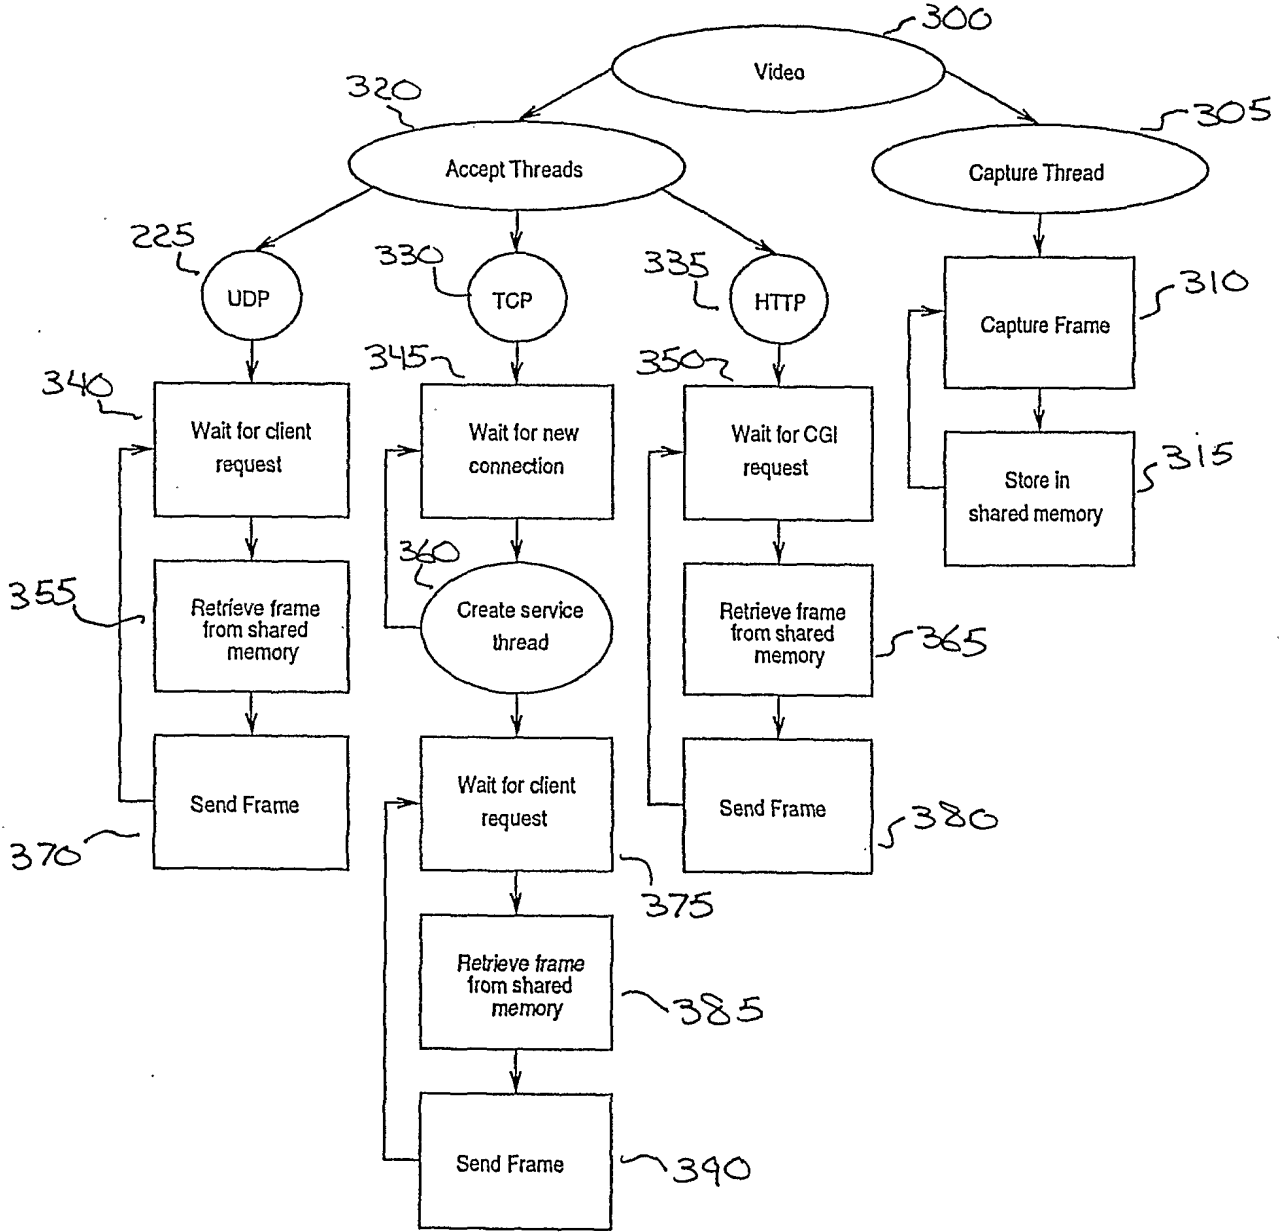


FIG. 3

Multithreaded Audio/Video Server

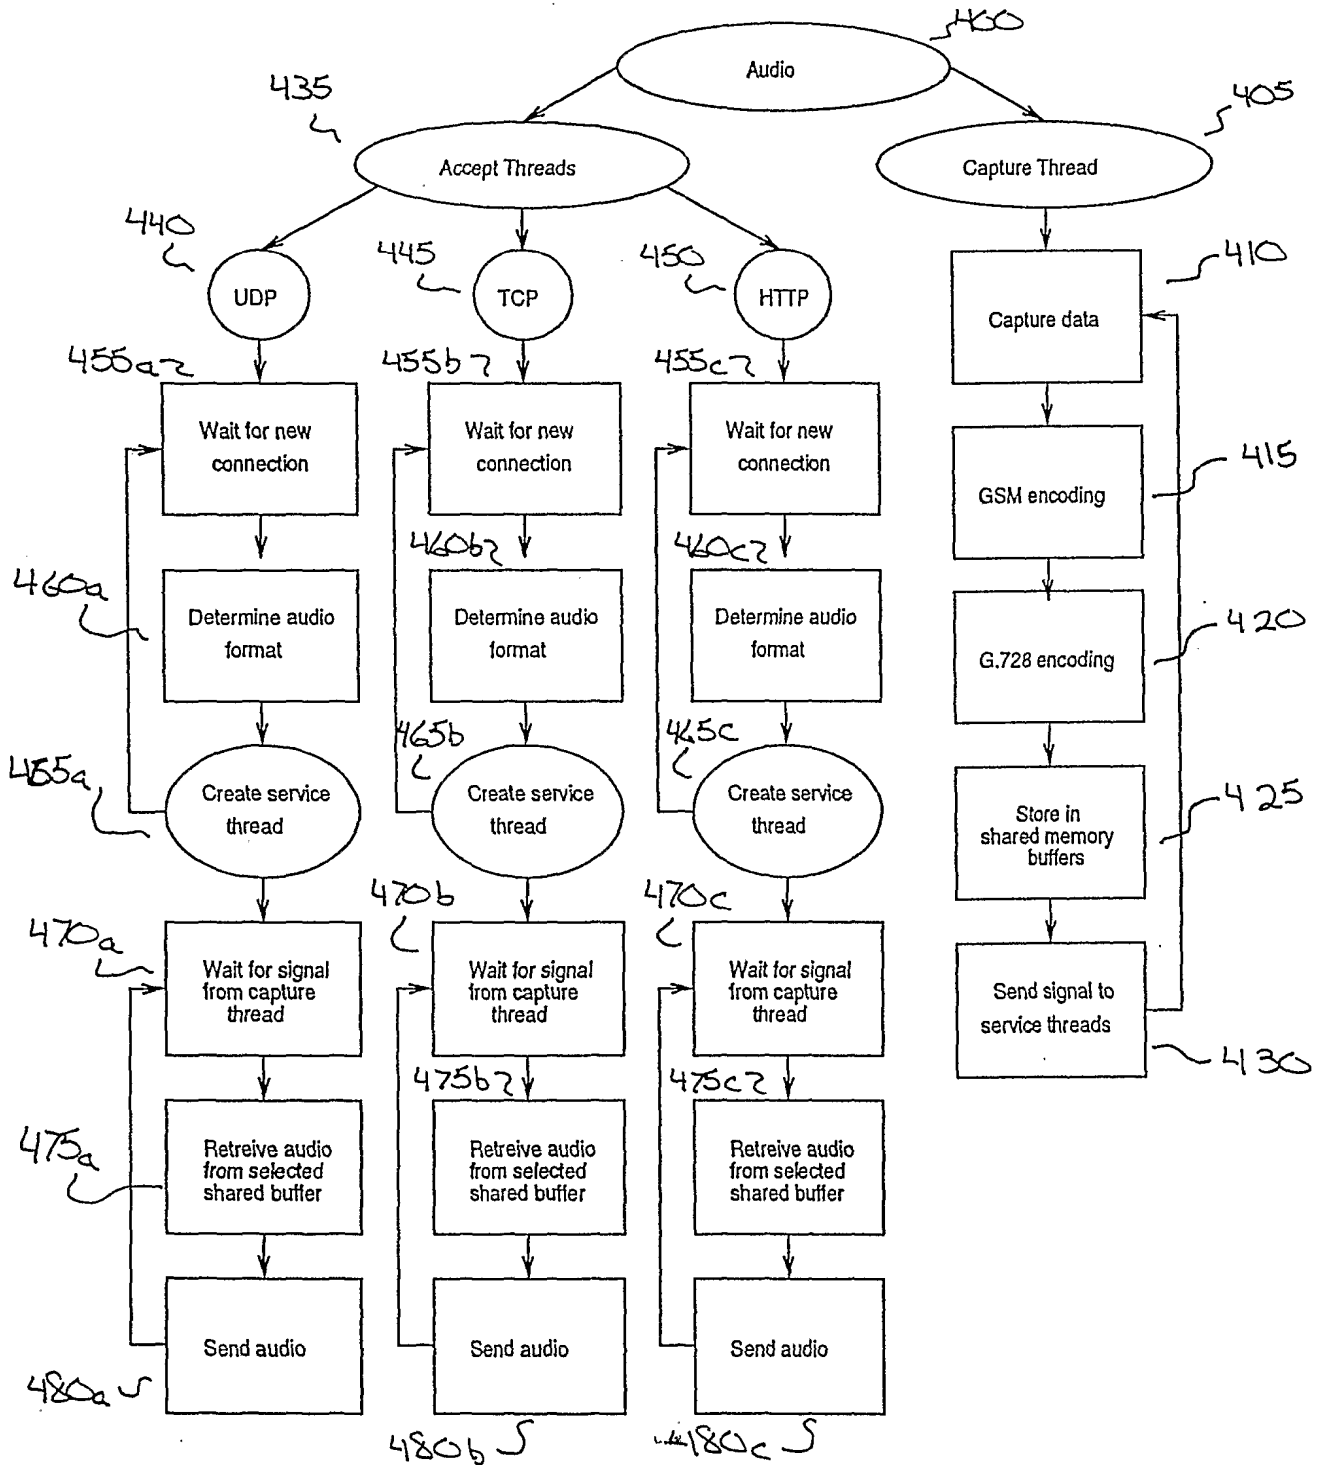


FIG. 4

CGI HTTP Transport Module

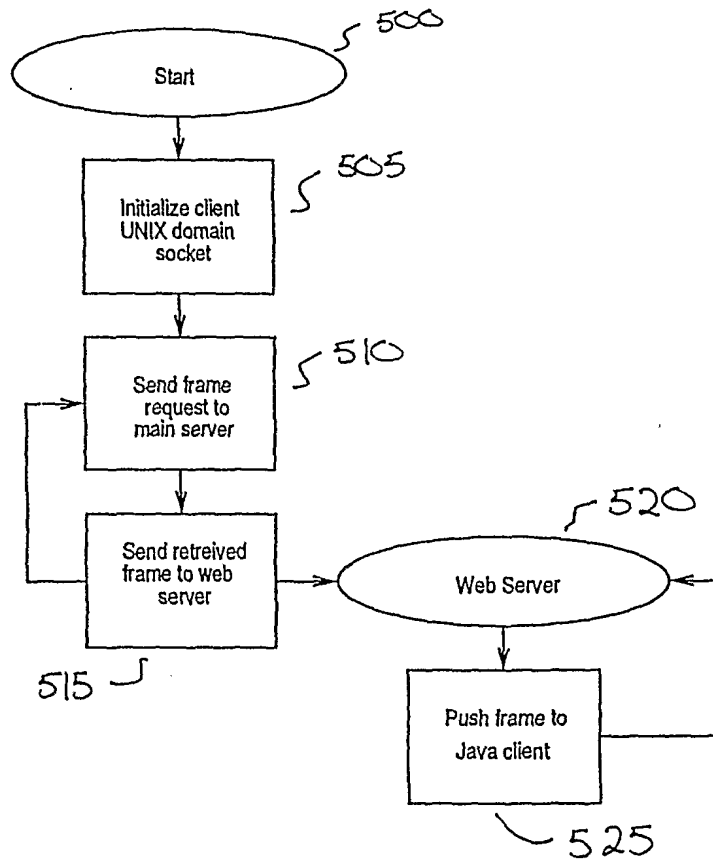


FIG. 5

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
25 October 2001 (25.10.2001)

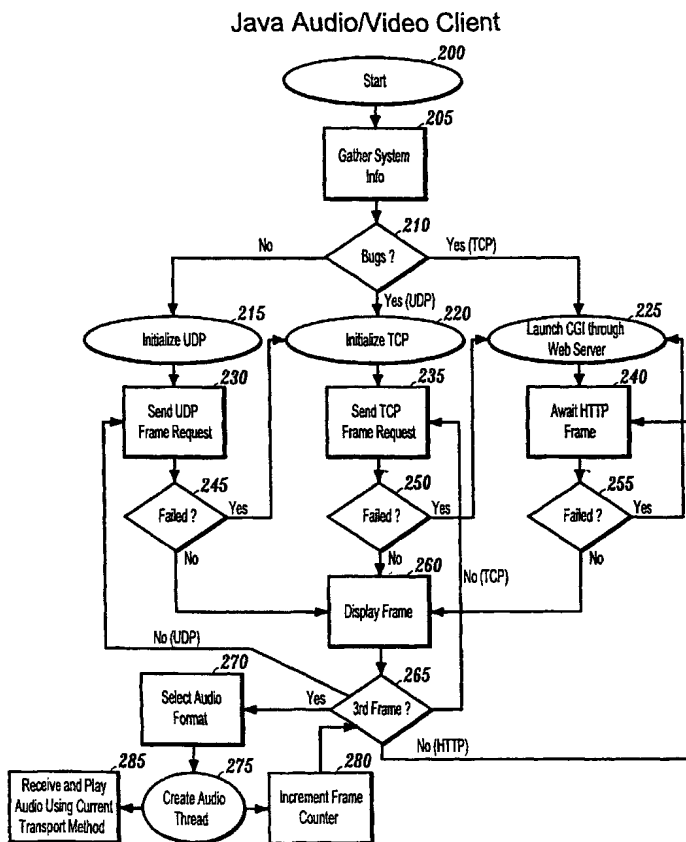
PCT

(10) International Publication Number
WO 01/80558 A3

- (51) International Patent Classification⁷: **H04N 7/24**
- (21) International Application Number: PCT/US01/40452
- (22) International Filing Date: 7 April 2001 (07.04.2001)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
09/548,989 14 April 2000 (14.04.2000) US
- (71) Applicant: **SOLIDSTREAMING, INC.** [US/US]; 32nd floor, 80 Pine Street, New York, NY 10005 (US).
- (72) Inventor: **BASTONE, Daniel**; Apt. 4A, 123 East 54th Street, New York, NY 10022 (US).
- (74) Agents: **CHAU, Frank et al.**; F. Chau & Associates, LLP, Suite 501, 1900 Hempstead Turnpike, East Meadow, NY 11554 (US).
- (81) Designated States (*national*): AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, 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, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZA, ZW.
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: A SYSTEM AND METHOD FOR MULTIMEDIA STREAMING



(57) Abstract: A system is provided for delivering streaming multimedia from a server to users via a communication network. Embedded clients at the users make requests for single datagrams. The server adapts to the variable bandwidths of the users and sends individual datagrams in response to the requesting user at the rate of available bandwidth of the requesting user.



WO 01/80558 A3



Published:

— *with international search 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.

(88) Date of publication of the international search report:

7 March 2002

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 01/40452

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04N7/24

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 H04N H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, PAJ, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 97 22201 A (XIE DONG ;CAMPBELL ROY H (US); CHEN ZHIGANG (US); TAN SEE MONG (US) 19 June 1997 (1997-06-19) abstract page 19, line 10 - line 22 page 21, line 8 - line 10 page 24, line 6 - line 10 page 24, line 16 -page 25, line 5 page 26, line 6 - line 16 ---	1-3,8, 15-18
X	US 5 929 850 A (HAASS JON C ET AL) 27 July 1999 (1999-07-27) abstract column 3, line 24 - line 43 column 6, line 9 - line 17 column 12, line 47 -column 13, line 30; figure 10 --- -/--	1-3

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

° Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *&* document member of the same patent family

Date of the actual completion of the international search

2 October 2001

Date of mailing of the international search report

16/10/2001

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
 NL - 2280 HV Rijswijk
 Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
 Fax: (+31-70) 340-3016

Authorized officer

Beaudet, J-P

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 01/40452

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5 999 979 A (RAVI HEMANTH SRINIVAS ET AL) 7 December 1999 (1999-12-07) the whole document ---	1-18
A	ERLANDSON C ET AL: "WAP - THE WIRELESS APPLICATION PROTOCOL" ERICSSON REVIEW, ERICSSON. STOCKHOLM, SE, no. 4, 1998, pages 150-153, XP000792053 ISSN: 0014-0171 the whole document ---	1-18
A	KALVA H ET AL: "IMPLEMENTING MULTIPLEXING, STREAMING, AND SERVER INTERACTION FOR MPEG-4" IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY, IEEE INC. NEW YORK, US, vol. 9, no. 8, December 1999 (1999-12), pages 1299-1312, XP000873453 ISSN: 1051-8215 paragraphs '0002!, '0005! -----	1-18

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 01/40452

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 9722201	A	19-06-1997	EP 0867003 A2 30-09-1998
			JP 2000515692 T 21-11-2000
			WO 9722201 A2 19-06-1997
US 5929850	A	27-07-1999	AU 3648197 A 21-01-1998
			EP 0909512 A1 21-04-1999
			JP 2001508602 T 26-06-2001
			WO 9800975 A1 08-01-1998
US 5999979	A	07-12-1999	EP 0956686 A1 17-11-1999
			JP 2000509592 T 25-07-2000
			WO 9834385 A1 06-08-1998
			EP 0956702 A1 17-11-1999
			WO 9834405 A1 06-08-1998
			US 6014706 A 11-01-2000
			US 6230172 B1 08-05-2001

CORRECTED VERSION

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
25 October 2001 (25.10.2001)

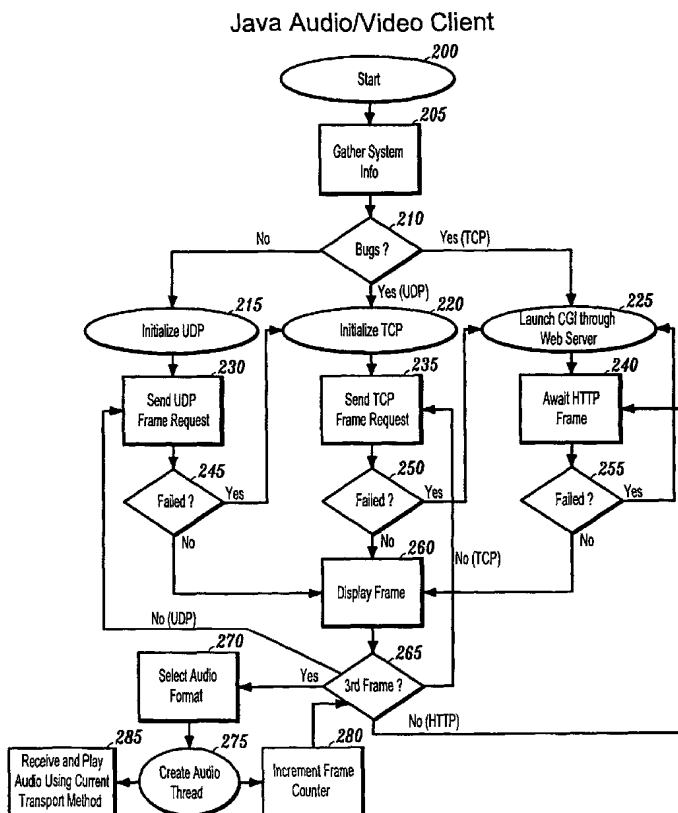
PCT

(10) International Publication Number
WO 01/080558 A3

- (51) International Patent Classification⁷: H04N 7/24
- (72) Inventor: BASTONE, Daniel; Apt. 4A, 123 East 54th Street, New York, NY 10022 (US).
- (21) International Application Number: PCT/US01/40452
- (74) Agents: CHAU, Frank et al.; F. Chau & Associates, LLP, Suite 501, 1900 Hempstead Turnpike, East Meadow, NY 11554 (US).
- (22) International Filing Date: 7 April 2001 (07.04.2001)
- (25) Filing Language: English
- (81) Designated States (national): AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, 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, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZA, ZW.
- (26) Publication Language: English
- (30) Priority Data: 09/548,989 14 April 2000 (14.04.2000) US
- (71) Applicant: SOLIDSTREAMING, INC. [US/US]; 32nd floor, 80 Pine Street, New York, NY 10005 (US).
- (84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian

[Continued on next page]

(54) Title: A SYSTEM AND METHOD FOR MULTIMEDIA STREAMING



(57) Abstract: A system is provided for delivering streaming multimedia from a server to users via a communication network. Embedded clients at the users make requests for single datagrams. The server adopts to the variable bandwidths of the users and sends individual datagrams in response to the requesting user at the rate of available bandwidth of the requesting user.



WO 01/080558 A3



patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report*

(88) Date of publication of the international search report:

7 March 2002

(48) Date of publication of this corrected version:

6 February 2003

(15) Information about Correction:

see PCT Gazette No. 06/2003 of 6 February 2003, Section II

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.

A SYSTEM AND METHOD FOR MULTIMEDIA STREAMING**BACKGROUND OF THE INVENTION**1. Field of the Invention

5 The present invention relates to a device and method for delivering video and/or audio data. More particularly to a device and method for delivering video and audio data in real time (streaming) through a communication network.

10 2. Description of Related Art

 As high bandwidth medium such as DSL, cable, and T1 lines becomes more readily available for users of communication networks such as the Internet, content providers quickly move to take advantage of the increased
15 bandwidth to deliver livelier and snazzier content. Countless websites are capable of delivering live (streaming) video or multimedia in real time. Internet users having a high bandwidth connection medium and a desktop PC with a higher speed processor receive the
20 streaming video with little difficulty. Typically, software receivers such as the Real Player by Real Networks or Media Player by Microsoft are installed on the users' PCs. The receivers receive and decode the encoded video data

delivered by content provider websites and display the decoded data as streaming video to the user.

At the content providers end, the video is first compressed or encoded. A popular video/audio compression format is MPEG, which is the format decoded by Microsoft's Media Player. MPEG (Moving Picture Experts Group) format is used in, for example, the Real Player and the Microsoft Media Player.

The MPEG technique for compressing digital video includes use of Discrete Cosine Transform (DCT), Quantization, Huffman coding, and Motion Compensated Predictive coding, in which the differences in what has changed between an image and its proceeding image are calculated and only the differences are encoded. Predictive coding requires interframe processing, i.e., data from neighboring frames are needed to successfully encode and decode an image; therefore, individual frames must be temporarily stored in a buffer and the image is encoded and decode using multiple frames. Buffering allows a server to send data at a constant rate, regardless of the rate at which the client displays the data. A disadvantage of MPEG and the buffering technique is that a considerable amount of memory is needed for buffering. In portable devices, sufficient memory may not exist.

Another video compression technique, known as JPEG (Joint Photographic Expert Group), employs the MPEG coding process except predictive coding. Thus, JPEG compression does not rely on interframe processing, i.e., each frame is independently processed and has no processing relationship to another frame. Therefore, JPEG compression does not require frame buffering.

Various methods and protocols for delivering data are available depending on the bandwidth of the connecting medium. One example is the Transmission Control Protocol (TCP). The TCP typically functions in conjunction with the Internet Protocol (IP). The TCP provides reliable, stream-oriented connections that hide most of IP's shortcomings; i.e., the basic nature of IP cannot guarantee the data will be delivered correctly. The TCP/IP protocol suite gets its name because the TCP protocol is layered on top of the IP protocol. The TCP layer interfaces on one side to application processes and on the other side to the IP protocol.

TCP data is organized as a stream of bytes, much like a file. The datagram nature of the network is concealed. A mechanism (the Urgent Pointer) exists to let out-of-band data be specially flagged. Sequence numbers are used to coordinate which data has been transmitted and received. TCP will arrange for retransmission if it determines that

data has been lost. This method provides for reliable delivery. TCP will dynamically learn the delay characteristics of a network and adjust its operation to maximize the throughput without overloading the network, this gives TCP the quality of network adaptation. TCP manages data buffers, and coordinates traffic so its buffers will not overflow. Fast senders will be stopped periodically to keep up with slower receivers, resulting in flow control.

TCP operates in both directions (full duplex) and in an almost completely independent manner, akin to two independent byte streams traveling in opposite directions. No TCP mechanism exists to associate data in the forward and reverse byte streams. Only during connection start and close sequences can TCP exhibit asymmetric behavior, i.e., data transfer in the forward direction but not in the reverse, or vice versa.

Each endpoint of a TCP connection will have a buffer for storing data that is transmitted over the network before the application is ready to read the data. This lets network transfers take place while applications are busy with other processing, improving overall performance. To avoid overflowing the buffer, TCP sets a Window Size field in each packet it transmits. This field contains the amount of data that may be transmitted into the buffer. If this number falls to zero, the remote TCP can send no more data. It must

wait until buffer space becomes available and it receives a packet announcing a non-zero window size.

Sometimes, the buffer space is too small. This happens when the network's bandwidth-delay product exceeds the
5 buffer size. The simplest solution is to increase the buffer, but for extreme cases the protocol itself becomes the bottleneck (because it doesn't support a large enough Window Size). Under these conditions, the network is termed an LFN (Long Fat Network).

10 When a host transmits a TCP packet to its peer, it must wait a period of time for an acknowledgment. If the reply does not come within the expected period, the packet is assumed to have been lost and the data is re-transmitted. The time that the protocol will wait for a reply is a
15 variable. Over an Ethernet, no more than a few microseconds should be needed for a reply. If the traffic must flow over the wide-area Internet, a second or two might be reasonable during peak utilization times. If a communication device is on a satellite traveling toward Mars, minutes may be
20 required before a reply.

Round-Trip Time (RTT) estimates are an important performance parameters in a TCP exchange, especially when dealing with an indefinitely large transfer. All TCP implementations eventually drop packets and retransmit them,
25 no matter the quality of the link. If the RTT estimate is

too low, packets are re-transmitted unnecessarily; if too high, the connection can sit idle while the host waits to timeout.

The User Datagram Protocol (UDP) is used in higher bandwidth communication links. UDP is a connectionless protocol that, like TCP, runs on top of an IP network. Unlike TCP/IP, UDP/IP provides very few error recovery services, offering instead a direct way to send and receive datagrams over an IP network. UDP is used primarily for broadcasting messages over a network. UDP packets are delivered like IP packets; connectionless datagrams that may be discarded before reaching their targets. UDP is useful when TCP would be too complex, too slow, or just unnecessary.

UDP provides a few functions beyond that of IP. For example, Port Numbers. UDP provides 16-bit port numbers to let multiple processes use UDP services on the same host. A UDP address is the combination of a 32-bit IP address and the 16-bit port number. Unlike IP, UDP does checksum its data, ensuring data integrity. A packet failing checksum is simply discarded, with no further action taken.

A common gateway interface (CGI) is one way for a web server to pass a web user's request to an application program, which in turn passes data back to be forwarded to the user. When the user requests a web page (for example, by

clicking on a hypertext link), the server retrieves the requested page, sending the page to the client. However, when a user submits a form on a web page, it usually needs to be processed by an application program. The web server typically passes the form information to a small application program (applet) that processes the data and may send back a confirmation message. This method for passing data between the server and the application is called the common gateway interface (CGI). The CGI is part of the web's Hypertext Transfer Protocol (HTTP).

For a client system having a lesser connection bandwidth, such as connection through a 56K modem, specific communication protocols can be established by an application program predownloaded or installed in the client's computer. A CGI can be used to pass a client's request to the application program. The CGI provides a consistent way for data to be passed from the user's request to the application program and back to the user. In other words, CGI operates in conjunction with clients and servers regardless of which operating system (OS) is being used by the parties, for example, Windows, Macintosh, UNIX and Linux, OS/390, or others. A CGI application may be written in a number of different languages.

An alternative to a CGI application is Microsoft's Active Server Page (ASP), in which a script embedded in a web page is executed at the server before the page is sent.

In a desktop client environment in which memory speed and connection bandwidth are sufficient, the connection rate is stable and servers can 'push' the MPEG-type datagrams to the clients synchronously, i.e., frames are buffered at the server and the buffer content is dumped or transmitted to the clients at a substantially periodic rate. At the desktop PC, the datagrams are also buffered because a single image depends on several datagrams.

As wireless applications and devices grow in popularity, the limitations of wireless applications and devices such as narrow bandwidth and limited memory and processing capacity must be addressed. In particular, content providers who wish to deliver substantially the same streaming video contents to wireless users as desktop users must find a viable solution to overcome these limitations. Therefore, a need exists for a system and method for delivery of streaming media to the client regardless of the client's available bandwidth.

SUMMARY OF THE INVENTION

These and other objects, features, and advantages of the present invention will become apparent from the

following detailed description of illustrative embodiments thereof, which is to be used in connection with the accompanying drawings.

A method for streaming video data over a network in real time is provided. The method includes initializing a transport mode for the video data, sending a data request for a single frame of video data from a client to a server, retrieving the single frame from a memory at the server, and sending the video data to the client.

The step of initializing also includes listing available transport modes for the client, determine whether incompatibilities exist between the available transport modes and software, choosing a transport mode from the list, and initializing parameters of the transport mode at the client for client control of video steaming.

Initializing is performed by a client application which is capable of running in different operating systems. Alternatively, initializing is performed by a client embedded in a web page. The client embedded in the web page can be a common gateway interface and an active server page. The transport mode is chosen from the following, a UDP, a TCP, and an HTTP, however other modes are contemplated.

The steps of sending a data request for a single frame from the client to a server, retrieving the video data from

a shared memory, and sending the retrieved video data to the client, are repeated for each video frame.

Storing video includes capturing a thread from a specified source, and storing the captured thread in the server's shared area of memory.

According to another aspect of the present invention, a storage medium having a stored program which is executable by a processor for causing the processor to perform method steps for streaming video communication is also provided.

The method steps comprising requesting a packet representing a single datagram from a server over a communication network, receiving a requested packet, processing and displaying the requested packet, incrementing a datagram frame counter, requesting a next packet based on the frame counter value from the server, and asynchronously processing and displaying the next packet when received.

The communications between the processor and the server is preferably by Wireless Application Protocol (WAP). The server is accessed by said processor via HTTP. The packet representing a datagram is preferably JPEG encoded, and the step of asynchronously processing and displaying is independent of data from the step of processing and displaying the requested packet.

An apparatus is also provided for communicating streaming video data between a plurality of users and a

server connected by a communication network, comprising a stored program executable by a processor in said server for causing the server to receive requests for individual datagrams from the plurality of users and forwarding individual datagrams in response to each request to the user making the request at a rate based on available bandwidth of the user making the request.

BRIEF DESCRIPTION OF THE DRAWINGS

10 The preferred embodiments are described with reference to the drawings wherein:

FIG. 1 is a diagram of a system for streaming data according to one embodiment of the invention;

15 FIG. 2 is a flow diagram of a preferred embodiment of the invention for streaming audio/video data to a client;

FIG. 3 is a flow diagram of a preferred embodiment of the invention for streaming video data;

FIG. 4 is a flow diagram of a preferred embodiment of the invention for streaming audio data; and

20 FIG. 5 is a flow diagram of a method for streaming audio data over a CGI HTTP transport mode.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

Preferred embodiments of the present invention will be described below in more detail with reference to the accompanying drawings.

5 The present invention relates to a system and method for streaming video. The invention is implemented over a network of processors. The network can include, a local area network (LAN), a wide area network (WAN), an Intranet, an Internet, a wireless network, or the like. These networks
10 can be in any configuration including, for example, star, ring, and bus.

For purposes of the present invention a client will, by definition, be receiving and displaying data, while the server will be providing data to the client. A system and
15 method of the present invention enables the client to receive audio/video data in real time regardless of the bandwidth available to the client. This is achieved through the implementation of a dynamic bandwidth adaption method for managing the flow of data. As a result, devices such as,
20 PDAs, hand-held PCs, and various mobile devices with little bandwidth are able to receive streaming video in real time. It is readily appreciated by one ordinarily skilled in the art that the invention is not limited to these devices and is also suitable to desktop computers, servers, and the
25 like.

A system and method according to the present invention can be construed as a browser that allows a client to access the server's network to receive streamed data content. The browser is preferably WAP (Wireless Application Protocol) compatible, but can be deployed to any device including those which are not WAP compatible. The browser preferably adapts to different codecs which do not require interframe processing, for example, Motion JPEG and Wavelet, allowing the flexibility to port software to embedded devices having limited memory resources. Frame buffering at the server can be dispensed with because each frame can be independently processed and delivered.

As shown in FIG. 1, the server 110 having audio/video data is connected to a network 100, in this example a bus network. While FIG. 1 depicts a bus network 100, any other network capable of supporting a server and client is contemplated by the invention. Alternatively, both the server and client may be embodied in the same computer and operate without a network. Typically a server is a computer program that provides a service to another computer program, the client. A server can function as a client and a client as a server depending on whether services are being offered or requested. The server, for purposes of the invention, is a stored program including a script for downloading audio/video data and an applet. A client 120 or

130 is also connected to the network. The client 120 captures the data from the server 110. The capture can take place from a local capture card, a local looped file, a local file on-demand, or a remote IP address. Each will be explained below in more detail. It should be readily apparent to one ordinarily skilled in the art that the client is described for client 120 but the embodiment is applicable for multiple clients.

The client 120 is a script which may be stored in the memory of a Personal Digital Assistant (PDA), PC or embedded in a web page. The client is a application which is capable of running in different operating systems, for example, Windows CE, Palm OS, Nokia PDA's, Windows, Linux etc. Alternatively, the client is embedded in a web page, for example, a common gateway interface (CGI) or a Microsoft Active Server Page (ASP). A CGI may be written in different programming languages, for example, C, C++, Java, and Perl, though this is not a complete list of possible applicable languages. Execution of the client ensures cross-platform and cross-browser functionality. Referring to FIG. 2, upon startup 200, the client gathers information 205 regarding its own operating environment needed to choose available transport modes. The information is checked against known bugs or incompatibilities 210 with certain operating systems (OS) and browser combinations. An example of a bug includes,

the Java Interpreter used in Microsoft Internet Explorer v4.0 which has broken implementation of User Data Protocol (UDP). An appropriate transport mode is chosen according to the incompatibilities 210. The client will detect this and default to suitable protocol, for example, Transmission Control Protocol (TCP) or Hypertext Transfer Protocol (HTTP) modes. The method of connection, e.g.UDP, TCP, HTTP, are referred to as a transports. The client initializes parameters which control the transport mode. This allows the same client to operate all possible configurations, for example, audio, video, audio/video, on-demand, live, or other similar configurations.

Assuming that no incompatibilities have been detected, the client attempts a UDP connection 215 to the server to retrieve a frame of video 230. In one embodiment, if the connection fails 245, a second attempt 220 is made to establish a UDP connection. After a specified number of attempts to establish a UDP connection, a TCP connection will be attempted. Note that in the present example, UDP is preferable to TCP connections based on bandwidth, however other connections may provide superior bandwidth than UDP, in such a case the wider bandwidth connection would be attempted first and the UDP second. In an alternative embodiment, an attempt to establish a TCP connection is made after the first failed UDP connection attempt.