U.S. Patent No. 7,124,205
(205 Patent)

IPR2017-1405 (Intel)
IPR2018-0336 (Dell)
IPR2017-1735 (Cavium)

*All citations herein are to the IPR2017-01405 case unless otherwise noted.

205 Patent: Instituted Grounds

- **Thia** in view of **Satran I** and **Satran II**
  - Claims 3, 9, 10, 16, 22, 27-33, 35, 36
  - Claims 31, 32-33 addressed in supplemental briefing

- Thia in view of **Satran I**, **Satran II**, and **Carmichael**
  - Claims 24-26

Ex. 1015 – Thia, Y.H., Woodside, C.M. Publication ("Thia")
Ex. 1053 – U.S. Patent No. 5,894,560 ("Carmichael")
Ex. 1056 – Satran, J. Publication ("Satran I")
Ex. 1057 – Satran, J. Publication ("Satran II")

205 Patent: Disputes

1. **Thia is enabling prior art**
2. Thia teaches the network interface device performing all network and transport layer processing
3. The combination of Thia, Satran I and Satran II discloses the challenged dependent claims
4. A POSA would have been motivated to combine Thia, Satran I and Satran II (as well as Carmichael)
5. Supplemental Briefing – claim 31 is indefinite or obvious in light of Thia, Satran I and Satran II
6. Motions to Amend 205 Patent should be denied

PO fails to identify why Thia is allegedly not enabling

- Patent Owner contends that Thia is an "inoperative device" and is therefore a non-enabling reference

Paper 32 (205 Corrected Response) at 22.

- Patent Owner's expert, Dr. Almeroth, essentially repeats (or copies verbatim) the opposition and does not provide any additional information or arguments

- A non-enabling reference can be prior art "for all that it teaches"

Id. (citing *Beckman Instruments v. LKB Produkter AB,* 892 F.2d 1547, 1551 (Fed. Cir. 1989)).

Dr. Lin: Thia is not a theoretical device

4.3 First Design: Design Steps

Figure 3 shows the steps followed in this study. There were three stages, a behavioural model, a structural or RTL model, and a gate level design. These gave us two kinds of feasibility check, that the logic we specified will execute the protocol within the environment we envisage, and that the design is technically feasible, for instance in a reasonable chip area.

Ex. 1015.008 (Thia).

SYNOPSYS was and still is one of the primary vendors of synthesis design tools used in the semiconductor industry to design semiconductor chips. A POSA would know that a gate-level design can be fabricated into a chip using well-known software tools and chip fabrication facilities. A POSA would have understood the teachings of Thia without the need for Thia to create a final chip.

Ex. 1223.005-.006 (Lin Reply Decl.) at ¶ 6.

- Thia discloses a design ready to be fabricated into a chip

PO fails to identify why Thia is not enabling

- The only alleged missing implementation details are not required by the 205 Patent claims or well-known to a POSA

> Accordingly, Thia does not contain any teaching or explanation of a dedicated NIA that implements or even tests TCP header bypass, let alone using flow keys and operation codes, and no explanation of how one would actually implement the theorized, generalized header prediction bypass stack in an actual NIC using the TCP protocol. Nor does Thia provide any teaching relating to the NIC performing packet data reassembly, such that certain header information is not passed to the host. These details are hardly trivial—they are exactly the prior art hurdles the '205 patent addressed and overcame. (Ex. 2026, ¶ 87)

Paper 32 (205 Corrected Response) at 27.

- Patent Owner's expert mirrors Patent Owner's Response

Thia is based on the well-known header prediction algorithm for TCP/IP

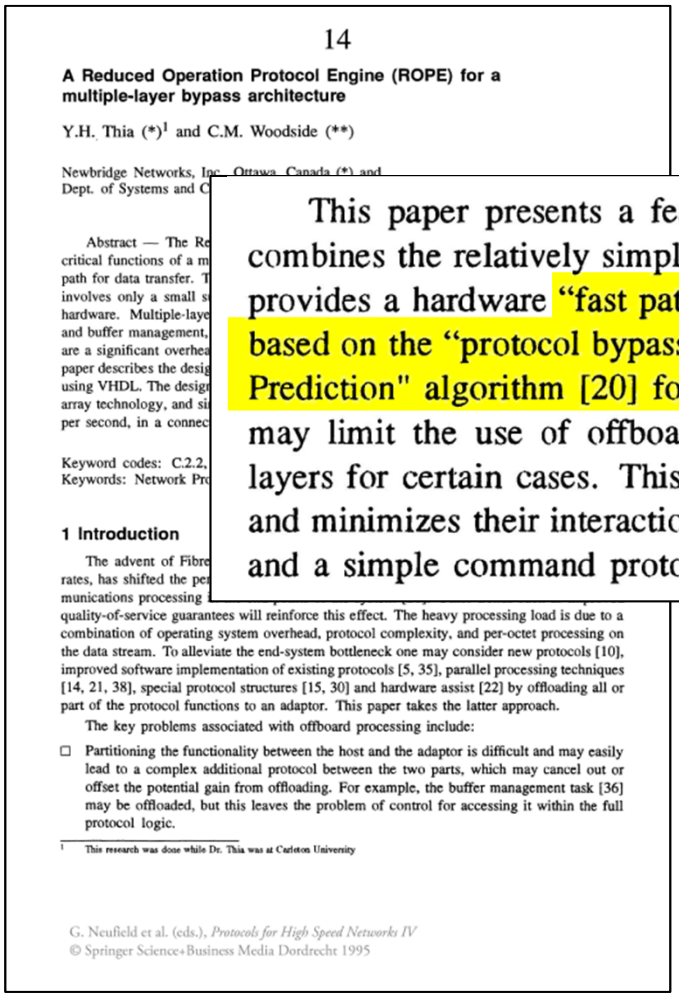

14

A Reduced Operation Protocol Engine (ROPE) for a multiple-layer bypass architecture

Y.H. Thia (*)[1] and C.M. Woodside (**)

Newbridge Networks, Inc., Ottawa, Canada (*) and
Dept. of Systems and C

Abstract — The Re
critical functions of a m
path for data transfer. T
involves only a small s
hardware. Multiple-laye
and buffer management,
are a significant overhea
paper describes the desig
using VHDL. The desig
array technology, and si
per second, in a connec

Keyword codes: C.2.2,
Keywords: Network Pr

1 Introduction

The advent of Fibre
rates, has shifted the pe
munications processing
quality-of-service guarantees will reinforce this effect. The heavy processing load is due to a combination of operating system overhead, protocol complexity, and per-octet processing on the data stream. To alleviate the end-system bottleneck one may consider new protocols [10], improved software implementation of existing protocols [5, 35], parallel processing techniques [14, 21, 38], special protocol structures [15, 30] and hardware assist [22] by offloading all or part of the protocol functions to an adaptor. This paper takes the latter approach.

The key problems associated with offboard processing include:

- Partitioning the functionality between the host and the adaptor is difficult and may easily lead to a complex additional protocol between the two parts, which may cancel out or offset the potential gain from offloading. For example, the buffer management task [36] may be offloaded, but this leaves the problem of control for accessing it within the full protocol logic.

[1] This research was done while Dr. Thia was at Carleton University

G. Neufield et al. (eds.), *Protocols for High Speed Networks IV*

This paper presents a feasibility study for a new approach to hardware assistance. It combines the relatively simple operations needed for data transfer across multiple layers and provides a hardware "fast path" for them, which will be efficient for bulk data transfer. It is based on the "protocol bypass concept" [37] which is a generalization of Jacobson's "Header Prediction" algorithm [20] for TCP/IP. Bypass solves the problems identified above, which may limit the use of offboard processing, by implementing an entire service through all layers for certain cases. This simplifies the interface between the host and the adaptor chip and minimizes their interaction, which is supported by an access test, some DMA processing and a simple command protocol. The chip design based on bypassing is called ROPE, for

Ex. 1015.002 (Thia); see also Paper 1 (205 Petition) at 21.

Dr. Lin: Thia is enabling to a POSA

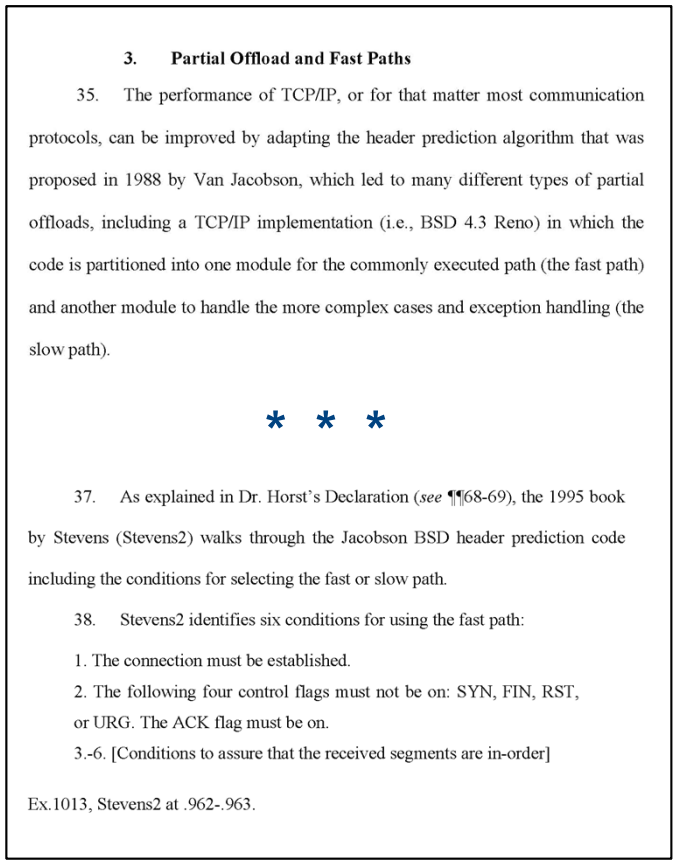

3. Partial Offload and Fast Paths

35. The performance of TCP/IP, or for that matter most communication protocols, can be improved by adapting the header prediction algorithm that was proposed in 1988 by Van Jacobson, which led to many different types of partial offloads, including a TCP/IP implementation (i.e., BSD 4.3 Reno) in which the code is partitioned into one module for the commonly executed path (the fast path) and another module to handle the more complex cases and exception handling (the slow path).

* * *

37. As explained in Dr. Horst's Declaration (*see* ¶¶68-69), the 1995 book by Stevens (Stevens2) walks through the Jacobson BSD header prediction code including the conditions for selecting the fast or slow path.

38. Stevens2 identifies six conditions for using the fast path:

1. The connection must be established.
2. The following four control flags must not be on: SYN, FIN, RST, or URG. The ACK flag must be on.
3.-6. [Conditions to assure that the received segments are in-order]

Ex.1013, Stevens2 at .962-.963.

See Ex. 1003.019, .021 (Lin Decl.) at ¶¶ 35-40;
see also Paper 1 (Petition) at 21, 23.

A POSA would have been able to understand and implement Thia's teachings, which is one of many implementations of Van Jacobson's header prediction

Ex. 1223.004-.005 (Lin Reply Decl.) at ¶ 5;
see also Ex. 1003.031, .070-.072 (Lin Decl.) at ¶ 54, A-12 – A-14.

205 Patent: Disputes

2. Thia teaches the network interface device performing all network and transport layer processing
   a. Thia teaches the network interface device performs all network layer processing
   b. Thia teaches the network interface device performs all transport layer processing

205 Patent: Claim 1

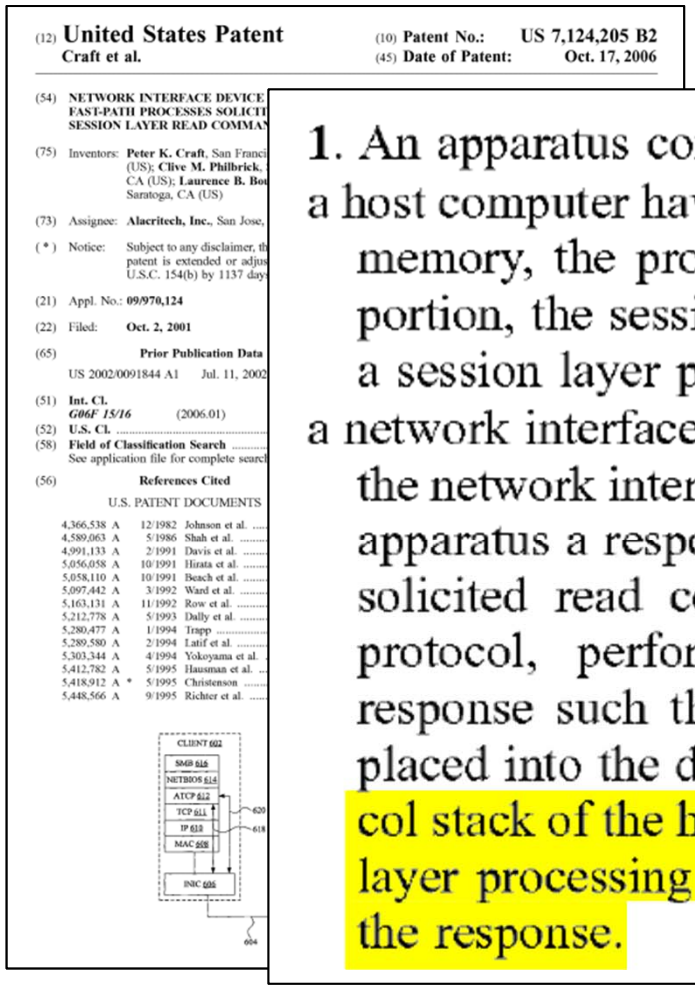

(12) **United States Patent**
Craft et al.

(10) **Patent No.: US 7,124,205 B2**
(45) **Date of Patent: Oct. 17, 2006**

(54) NETWORK INTERFACE DEVICE
FAST-PATH PROCESSES SOLICIT
SESSION LAYER READ COMMAN

(75) Inventors: **Peter K. Craft**, San Franci
(US); **Clive M. Philbrick**,
CA (US); **Laurence B. Bo**
Saratoga, CA (US)

(73) Assignee: **Alacritech, Inc.**, San Jose,

(*) Notice: Subject to any disclaimer, th
patent is extended or adjus
U.S.C. 154(b) by 1137 day

(21) Appl. No.: **09/970,124**

(22) Filed: **Oct. 2, 2001**

(65) **Prior Publication Data**

US 2002/0091844 A1 Jul. 11, 2002

(51) **Int. Cl.**
G06F 15/16 (2006.01)

(52) **U.S. Cl.** .....

(58) **Field of Classification Search** .....
See application file for complete search

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,366,538	A	12/1982	Johnson et al.
4,589,063	A	5/1986	Shah et al.
4,991,133	A	2/1991	Davis et al.
5,056,058	A	10/1991	Hirata et al.
5,058,110	A	10/1991	Beach et al.
5,097,442	A	3/1992	Ward et al.
5,163,131	A	11/1992	Row et al.
5,212,778	A	5/1993	Dally et al.
5,280,477	A	1/1994	Trapp
5,289,580	A	2/1994	Latif et al.
5,303,344	A	4/1994	Yokoyama et al.
5,412,782	A	5/1995	Hausman et al.
5,418,912	A *	5/1995	Christenson
5,448,566	A	9/1995	Richter et al.

1. An apparatus comprising:

a host computer having a protocol stack and a destination memory, the protocol stack including a session layer portion, the session layer portion being for processing a session layer protocol; and

a network interface device coupled to the host computer, the network interface device receiving from outside the apparatus a response to a solicited read command, the solicited read command being of the session layer protocol, performing fast-path processing on the response such that a data portion of the response is placed into the destination memory without the protocol stack of the host computer performing any network layer processing or any transport layer processing on the response.

Ex. 1001.052 (205 Patent), Claim 1; see also -.052, Claims 35 and 36.

205 Patent: Claim 22

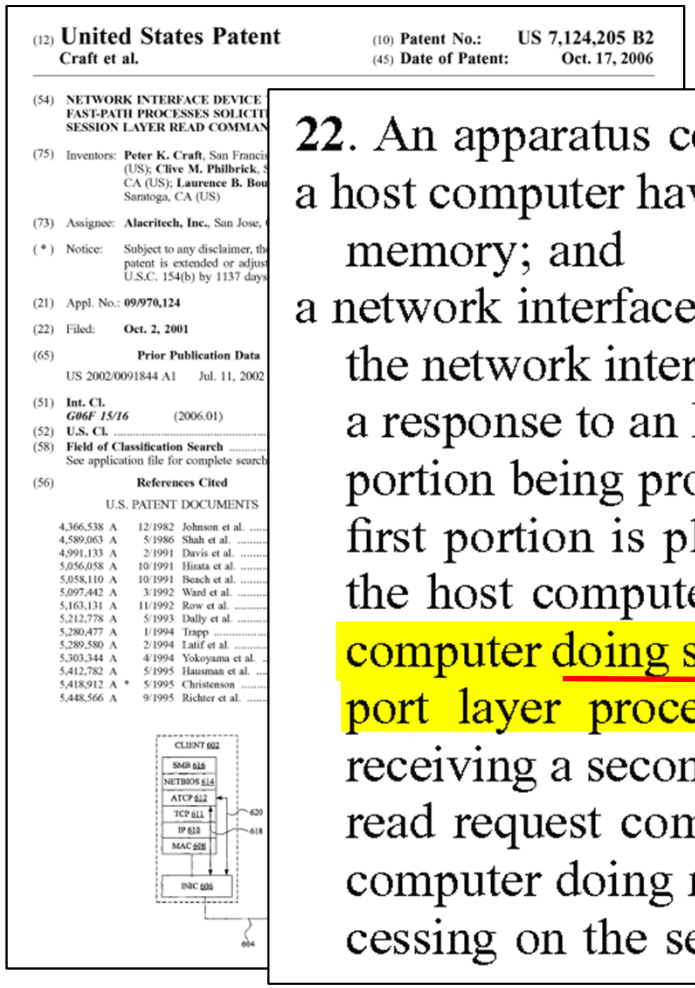

(12) **United States Patent**
Craft et al.

(10) **Patent No.: US 7,124,205 B2**
(45) **Date of Patent: Oct. 17, 2006**

(54) NETWORK INTERFACE DEVICE
FAST-PATH PROCESSES SOLICIT
SESSION LAYER READ COMMAN

(75) Inventors: **Peter K. Craft**, San Francis
(US); **Clive M. Philbrick**,
CA (US); **Laurence B. Bou**
Saratoga, CA (US)

(73) Assignee: **Alacritech, Inc.**, San Jose,

(*) Notice: Subject to any disclaimer, th
patent is extended or adjust
U.S.C. 154(b) by 1137 days

(21) Appl. No.: **09/970,124**

(22) Filed: **Oct. 2, 2001**

(65) **Prior Publication Data**

US 2002/0091844 A1 Jul. 11, 2002

(51) **Int. Cl.**
G06F 15/16 (2006.01)

(52) **U.S. Cl.** .....

(58) **Field of Classification Search** .....
See application file for complete search

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,366,538	A	12/1982	Johnson et al.	
4,589,063	A	5/1986	Shah et al.	
4,991,133	A	2/1991	Davis et al.	
5,056,058	A	10/1991	Hirata et al.	
5,058,110	A	10/1991	Beach et al.	
5,097,442	A	3/1992	Ward et al.	
5,163,131	A	11/1992	Row et al.	
5,212,778	A	5/1993	Dally et al.	
5,280,477	A	1/1994	Trapp	
5,289,580	A	2/1994	Latif et al.	
5,303,344	A	4/1994	Yokoyama et al.	
5,412,782	A	5/1995	Hausman et al.	
5,418,912	A *	5/1995	Christenson	
5,448,566	A	9/1995	Richter et al.	

22. An apparatus comprising:

a host computer having a protocol stack and a destination memory; and

a network interface device coupled to the host computer, the network interface device receiving a first portion of a response to an ISCSI read request command, the first portion being processed such that a data portion of the first portion is placed into the destination memory on the host computer with the protocol stack of the host computer doing substantially no network layer or transport layer processing, the network interface device receiving a second portion of the response to the ISCSI read request command, the protocol stack of the host computer doing network layer and transport layer processing on the second portion.

Ex. 1001.052 (205 Patent), Claim 22.

205 Patent: Claim 31

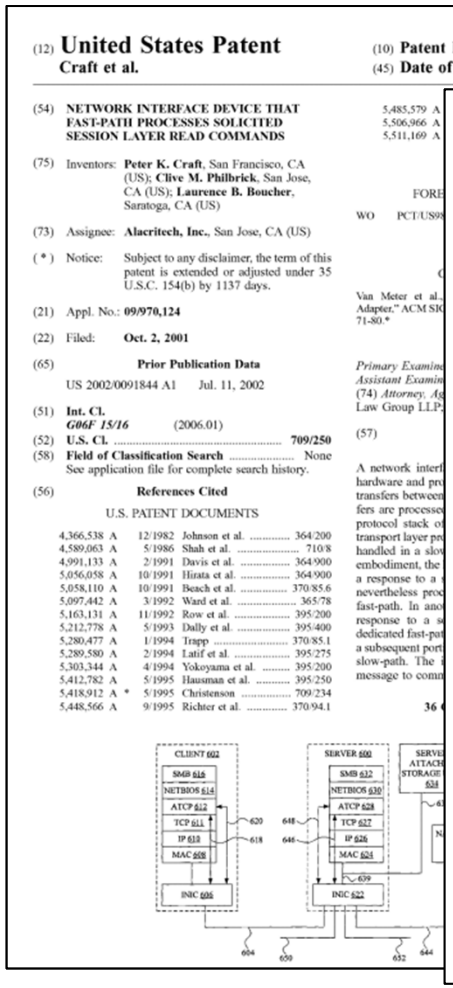

(12) **United States Patent**
Craft et al.

(10) **Patent No.: US 7,124,205 B2**
(45) **Date of Patent: Oct. 17, 2006**

(54) **NETWORK INTERFACE DEVICE THAT FAST-PATH PROCESSES SOLICITED SESSION LAYER READ COMMANDS**

(75) Inventors: **Peter K. Craft**, San Francisco, CA (US); **Clive M. Philbrick**, San Jose, CA (US); **Laurence B. Boucher**, Saratoga, CA (US)

(73) Assignee: **Alacritech, Inc.**, San Jose, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 1137 days.

(21) Appl. No.: **09/970,124**

(22) Filed: **Oct. 2, 2001**

(65) **Prior Publication Data**

US 2002/0091844 A1 Jul. 11, 2002

(51) **Int. Cl.**
G06F 15/16 (2006.01)

(52) **U.S. Cl.** **709/250**

(58) **Field of Classification Search** ..... None
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,366,538	A	12/1982	Johnson et al.	364/200
4,589,063	A	5/1986	Shah et al.	710/8
4,991,133	A	2/1991	Davis et al.	364/900
5,056,058	A	10/1991	Hirata et al.	364/900
5,058,110	A	10/1991	Beach et al.	370/85.6
5,097,442	A	3/1992	Ward et al.	365/78
5,163,131	A	11/1992	Row et al.	395/200
5,212,778	A	5/1993	Dally et al.	395/400
5,280,477	A	1/1994	Trapp	370/85.1
5,289,580	A	2/1994	Latif et al.	395/275
5,303,344	A	4/1994	Yokoyama et al.	395/200
5,412,782	A	5/1995	Hausman et al.	395/250
5,418,912	A *	5/1995	Christenson	709/234
5,448,566	A	9/1995	Richter et al.	370/94.1

31. An apparatus comprising:

a host computer having a protocol stack and a destination memory; and

means, coupled to the host computer, for receiving from outside the apparatus a response to an ISCSI read request command and for fast-path processing a portion of the response to the ISCSI read request command, the portion including data, the portion being fast-path processed such that the data is placed into the destination memory on the host computer without the protocol stack of the host computer doing significant network layer or significant transport layer processing, the means also being for receiving a subsequent portion of the response to the ISCSI read request command and for slow-path processing the subsequent portion such that the protocol stack of the host computer does network layer and transport layer processing on the subsequent portion.

Ex. 1001.052 (205 Patent), Claim 31.

205 Patent: Disputes

2. Thia teaches the network interface device performing all network and transport layer processing

   a. **Thia teaches the network interface device performs *all* network layer processing**

   b. Thia teaches the network interface device performs all transport layer processing

Thia: Bypass all host network layer processing in the data transfer phase

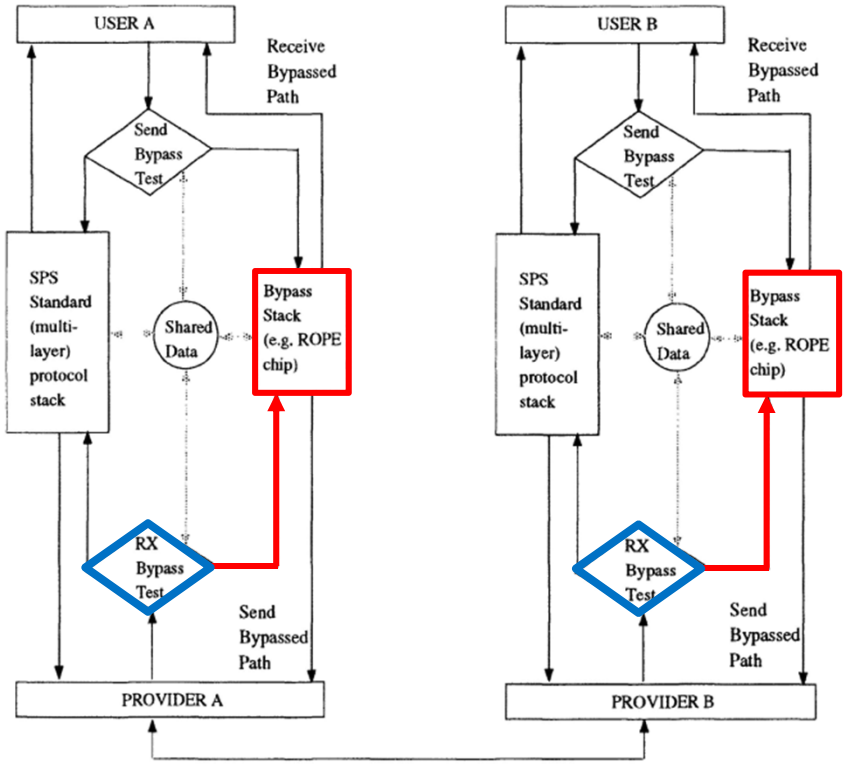

Figure 1 Bypass Architecture

Ex.1015.003 (Thia) at Fig. 1 (annotated);
see, e.g., Ex. 1003.071-072 (Lin Decl.) at A-13 – A-14;
see also Paper 1 (205 Petition) at 46-49.

"The bypass stack performs all the relevant protocol processing in the data transfer phase."

Ex. 1015.003 (Thia);
Paper 1 (205 Petition) at 31, 47, 67;
Ex. 1003.071 (Lin Decl.) at A-13;
see also Ex. 1223.006-.007 (Lin Reply Decl.) at ¶ 9.

Thia: Bypass multiple layers, including the network layer

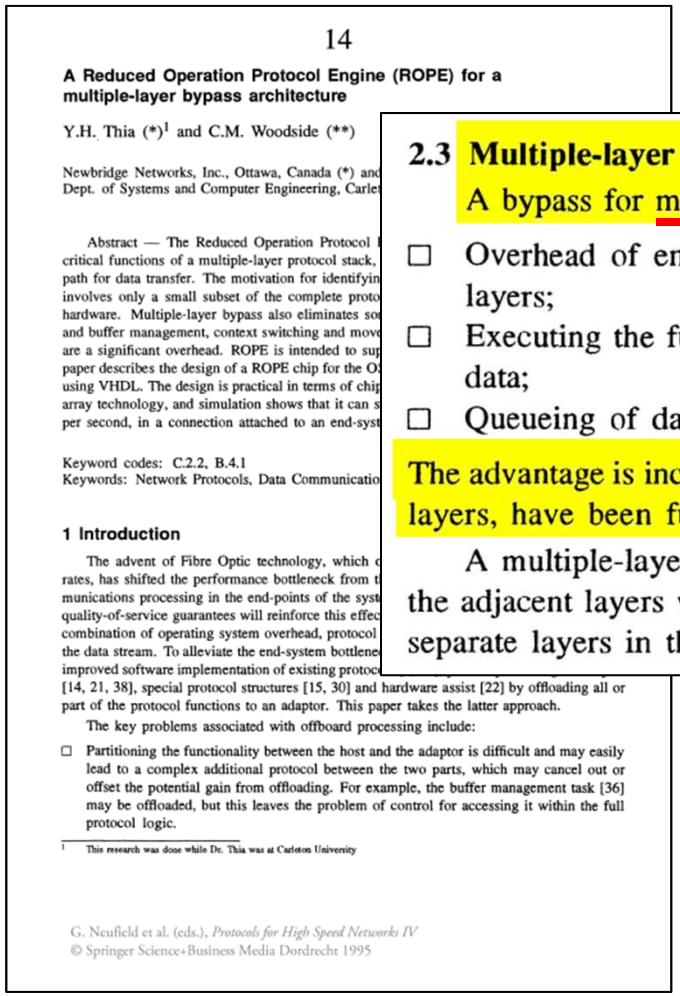

14

A Reduced Operation Protocol Engine (ROPE) for a multiple-layer bypass architecture

Y.H. Thia (*)[1] and C.M. Woodside (**)

Newbridge Networks, Inc., Ottawa, Canada (*) and
Dept. of Systems and Computer Engineering, Carle

Abstract — The Reduced Operation Protocol
critical functions of a multiple-layer protocol stack,
path for data transfer. The motivation for identifyin
involves only a small subset of the complete proto
hardware. Multiple-layer bypass also eliminates so
and buffer management, context switching and move
are a significant overhead. ROPE is intended to su
paper describes the design of a ROPE chip for the O
using VHDL. The design is practical in terms of chi
array technology, and simulation shows that it can s
per second, in a connection attached to an end-syst

Keyword codes: C.2.2, B.4.1
Keywords: Network Protocols, Data Communicatio

1 Introduction

The advent of Fibre Optic technology, which
rates, has shifted the performance bottleneck from t
munications processing in the end-points of the syst
quality-of-service guarantees will reinforce this effec
combination of operating system overhead, protocol
the data stream. To alleviate the end-system bottlene
improved software implementation of existing protoc
[14, 21, 38], special protocol structures [15, 30] and hardware assist [22] by offloading all or part of the protocol functions to an adaptor. This paper takes the latter approach.

The key problems associated with offboard processing include:

- Partitioning the functionality between the host and the adaptor is difficult and may easily lead to a complex additional protocol between the two parts, which may cancel out or offset the potential gain from offloading. For example, the buffer management task [36] may be offloaded, but this leaves the problem of control for accessing it within the full protocol logic.

[1] This research was done while Dr. Thia was at Carleton University

G. Neufield et al. (eds.), *Protocols for High Speed Networks IV*
© Springer Science+Business Media Dordrecht 1995

2.3 Multiple-layer bypass

A bypass for multiple layers instead of just one gives additional gains by avoiding:

- ☐ Overhead of encoding and decoding the interface control information passed between layers;
- ☐ Executing the full general protocol logic for the layers to decide how to manipulate the data;
- ☐ Queueing of data at layer boundaries.

The advantage is increased further in cases where some layers, like the network and application layers, have been further subdivided into sublayers.

A multiple-layer bypass path is a concatenation of processing procedures performed by the adjacent layers when they are simultaneously in the data transfer phase. Meanwhile, the separate layers in the SPS path handle the other phases.

Ex.1015.004 (Thia);
Paper 1 (205 Petition) at 47, 54, 59, 68, 84;
Paper 44 (205 Reply) at 9.

Thia: Bypass functions can be narrow or extended

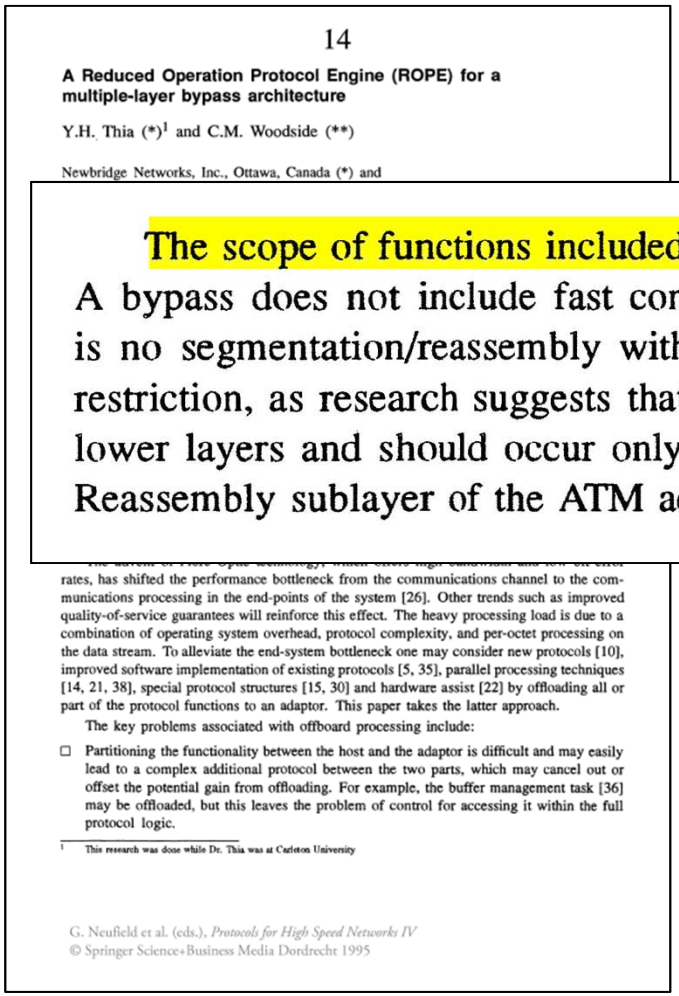

14

A Reduced Operation Protocol Engine (ROPE) for a multiple-layer bypass architecture

Y.H. Thia (*)[1] and C.M. Woodside (**)

Newbridge Networks, Inc., Ottawa, Canada (*) and

The scope of functions included in a bypass may be narrowly defined, or more extended. A bypass does not include fast connection setup but also does not interfere with it. There is no segmentation/reassembly within the bypass path, but we do not see this as a major restriction, as research suggests that fragmentation of PDUs should be restricted only to the lower layers and should occur only once in the protocol stack [23]. The Segmentation and Reassembly sublayer of the ATM adaptation layer is a good place for such functions [25].

rates, has shifted the performance bottleneck from the communications channel to the communications processing in the end-points of the system [26]. Other trends such as improved quality-of-service guarantees will reinforce this effect. The heavy processing load is due to a combination of operating system overhead, protocol complexity, and per-octet processing on the data stream. To alleviate the end-system bottleneck one may consider new protocols [10], improved software implementation of existing protocols [5, 35], parallel processing techniques [14, 21, 38], special protocol structures [15, 30] and hardware assist [22] by offloading all or part of the protocol functions to an adaptor. This paper takes the latter approach.

The key problems associated with offboard processing include:

- Partitioning the functionality between the host and the adaptor is difficult and may easily lead to a complex additional protocol between the two parts, which may cancel out or offset the potential gain from offloading. For example, the buffer management task [36] may be offloaded, but this leaves the problem of control for accessing it within the full protocol logic.

[1] This research was done while Dr. Thia was at Carleton University

G. Neufield et al. (eds.), *Protocols for High Speed Networks IV*

Ex. 1015.014 (Thia);
Paper 44 (Reply) at 9;
Ex. 1223.010-.011 (Lin Reply Decl.) at ¶ 15.

OSI model has multiple layers, which must be processed in order

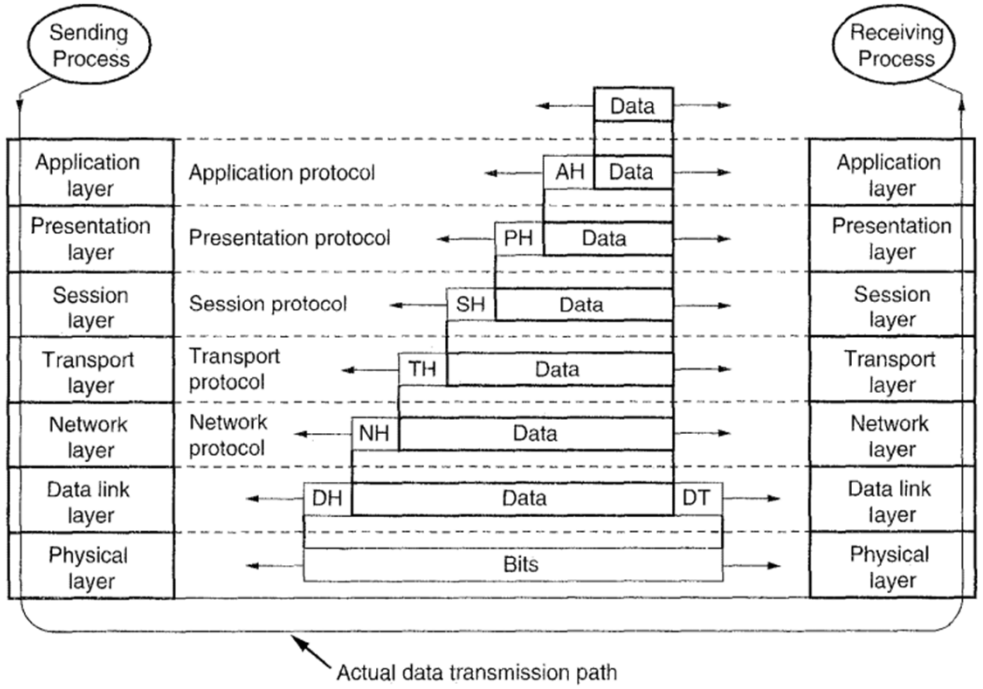

Fig. 1-17. An example of how the OSI model is used. Some of the headers may be null. (Source: H.C. Folts. Used with permission.)

Ex. 1006.052 (Tanenbaum96) at Fig. 1-17; Paper 1 (205 Petition) at 41.

- The network layer must be processed before the transport and session layers
- It is undisputed that Thia discloses processing the transport and session layers on the adapter

See e.g., Paper 1 (205 Petition) at 18-19;
Paper 44 (205 Reply) at 9-10;
Ex. 1223.006-.008 (Lin Reply Decl.) at ¶¶ 9-10;
Paper 31 (205 Response) at 2;
Ex. 2026.029 (Almeroth Decl.) at ¶ 68.

205 Patent: Disputes

2. Thia discloses performing all network and transport layer processing on the bypass path

   a. Thia's network interface device performs all network layer processing

   b. Thia's network interface device performs all transport layer processing

      i. The claims do not recite "reassembly"

      ii. Thia teaches placing data from in-order packets into host memory on the bypass path

      iii. "Segmentation/reassembly" discussed in Thia is below the transport layer

Thia's bypass transport layer includes "reassembly"

- Patent Owner admits that transport layer processing is performed on the bypass path, but argues that "reassembly" of incoming packets is missing from Thia:

> "Crucially, Thia does not disclose reassembling the incoming packets, which is a primary responsibility of the transport layer"

Paper 31 (205 Response) at 43.

The claims do not recite “reassembly”

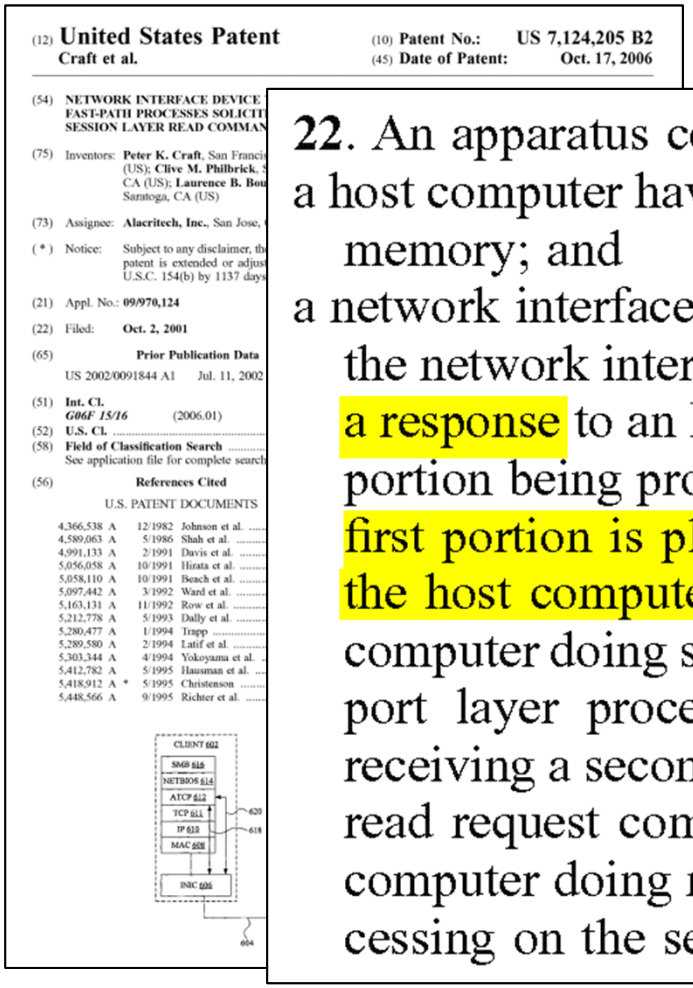

(12) **United States Patent**
Craft et al.

(10) **Patent No.: US 7,124,205 B2**
(45) **Date of Patent: Oct. 17, 2006**

(54) NETWORK INTERFACE DEVICE
FAST-PATH PROCESSES SOLICIT
SESSION LAYER READ COMMAN

(75) Inventors: **Peter K. Craft**, San Francis
(US); **Clive M. Philbrick**,
CA (US); **Laurence B. Bou**
Saratoga, CA (US)

(73) Assignee: **Alacritech, Inc.**, San Jose,

(*) Notice: Subject to any disclaimer, th
patent is extended or adjus
U.S.C. 154(b) by 1137 days

(21) Appl. No.: **09/970,124**

(22) Filed: **Oct. 2, 2001**

(65) **Prior Publication Data**

US 2002/0091844 A1 Jul. 11, 2002

(51) **Int. Cl.**
G06F 15/16 (2006.01)

(52) **U.S. Cl.** .....

(58) **Field of Classification Search** .....
See application file for complete search

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,366,538	A	12/1982	Johnson et al.	.....
4,589,063	A	5/1986	Shah et al.	.....
4,991,133	A	2/1991	Davis et al.	.....
5,056,058	A	10/1991	Hirata et al.	.....
5,058,110	A	10/1991	Beach et al.	.....
5,097,442	A	3/1992	Ward et al.	.....
5,163,131	A	11/1992	Row et al.	.....
5,212,778	A	5/1993	Dally et al.	.....
5,280,477	A	1/1994	Trapp	.....
5,289,580	A	2/1994	Latif et al.	.....
5,303,344	A	4/1994	Yokoyama et al.	.....
5,412,782	A	5/1995	Hausman et al.	.....
5,418,912	A *	5/1995	Christenson	.....
5,448,566	A	9/1995	Richter et al.	.....

22. An apparatus comprising:

a host computer having a protocol stack and a destination memory; and

a network interface device coupled to the host computer, the network interface device receiving a first portion of a response to an ISCSI read request command, the first portion being processed such that a data portion of the first portion is placed into the destination memory on the host computer with the protocol stack of the host computer doing substantially no network layer or transport layer processing, the network interface device receiving a second portion of the response to the ISCSI read request command, the protocol stack of the host computer doing network layer and transport layer processing on the second portion.

Ex. 1001.052 (205 Patent), Claim 22.

PO equates "reassembly" with placing data from each packet into host memory

> Other portions of PO's exemplary support make reassembly more apparent. For example, [0056] through [0058] describe an embodiment where "the data of the packet, without network or transport layer headers, is sent by direct memory access (DMA) unit 68 to the destination in file cache 80 or file cache 24 denoted by the CCB." In placing the data of each incoming packet in "file cache 24"—which is in host memory (Fig. 1)—the DMA unit on the network interface device thus reassembles the data. (Ex. 2305 ¶ 8). [0063] and [0064] similarly describe an

Paper 45 (Motion to Amend Reply) at 2;
see also Paper 70 (Supplemental Motion to Amend Reply) at 2.

205 Patent: Disputes

2. Thia discloses the performing all network and transport layer processing on the bypass path

   a. Thia's network interface device performs all network layer processing

   b. Thia's network interface device performs all transport layer processing

      i. Claims do not recite "reassembly"

      ii. Thia teaches placing data from in-order packets into host memory on the bypass path

      iii. "Segmentation/reassembly" discussed in Thia is below the transport layer

Thia: ROPE chip places data from each packet into host memory

A Reduced Operation Protocol Engine (ROPE) for a multiple-layer bypass architecture

Y.H. Thia (*)[1] and C.M. Woodside (**)

Newbridge Networks, Inc., Ottawa, Canada (*) and
Dept. of Systems and Computer Engineering, Carleton University, Ottaw

Abstract — The Reduced Operation Protocol Engine (ROPE) pres
critical functions of a multiple-layer protocol stack, based on the "bypas
path for data transfer. The motivation for identifying this separate proce
involves only a small subset of the complete protocol, which can then
hardware. Multiple-layer bypass also eliminates some inter-layer operat
and buffer management, context switching and movement of data across
are a significant overhead. ROPE is intended to support high-speed bul
paper describes the design of a ROPE chip for the OSI Session and Trans
using VHDL. The design is practical in terms of chip complexity and are
array technology, and simulation shows that it can support a data rate ap
per second, in a connection attached to an end-system.

Keyword codes: C.2.2, B.4.1
Keywords: Network Protocols, Data Communications Devices

1 Introduction

The advent of Fibre Optic technology, which offers high bandwidt
rates, has shifted the performance bottleneck from the communications
munications processing in the end-points of the system [26]. Other trend
quality-of-service guarantees will reinforce this effect. The heavy process
combination of operating system overhead, protocol complexity, and per
the data stream. To alleviate the end-system bottleneck one may consider
improved software implementation of existing protocols [5, 35], parallel pr
[14, 21, 38], special protocol structures [15, 30] and hardware assist [22]
part of the protocol functions to an adaptor. This paper takes the latter

The key problems associated with offboard processing include:

- Partitioning the functionality between the host and the adaptor is diffi
  lead to a complex additional protocol between the two parts, which
  offset the potential gain from offloading. For example, the buffer ma
  may be offloaded, but this leaves the problem of control for accessi
  protocol logic.

[1] This research was done while Dr. Thia was at Carleton University

G. Neufield et al. (eds.), *Protocols for High Speed Networks IV*

Layer	*Procedure*	***Bypass Chip***	*Host*	*Per-Octet (A)*	*Per-Packet (B)*	*Per-Group-Of-Packets Aggregated to Per-Packet for bulk data transfer (B)*	*Remarks*
Presentation	Encoding	X		X			
	Encryption	X		X			
	Compression	X		X			
	Context Alteration		X			X	
Session	Synchronization Management		X			X	
	Token management		X			X	
Transport (Class 4)	Checksum (Optional)	X		X			
	Timer Management	X			X	X	Depends on Implementation
	Generation of ACK packets (Flow Control)	X				X	
	Resequencing	X			X		
All 3 layers	Header Construction	X			X		
	Header Decode	X			X		
	Buffer Management	X			X		Minimized (Simple scheme)
	Context Switching	X			X	X	Moved away from host OS
	Data Copying	With multiple-layer bypass, data Copying within layers is elminated.			X		Use of dual-ported memory and DMA..

Table 1 Bypassable versus Non-bypassable functions

Ex. 1015.006 (Thia); see also Ex. 1223.006-.007 (Lin Reply Decl.) at ¶ 9; Ex. 1210.011-.014, .023-.025 (Lin Opp. Decl.) at ¶¶ 24-30, A-3 – A-5; Paper 39 (Opp. To Motion to Amend) at 11; Paper 1 (205 Petition) at 76.

Thia: Copy data portions of PDUs from the adaptor buffer to host memory

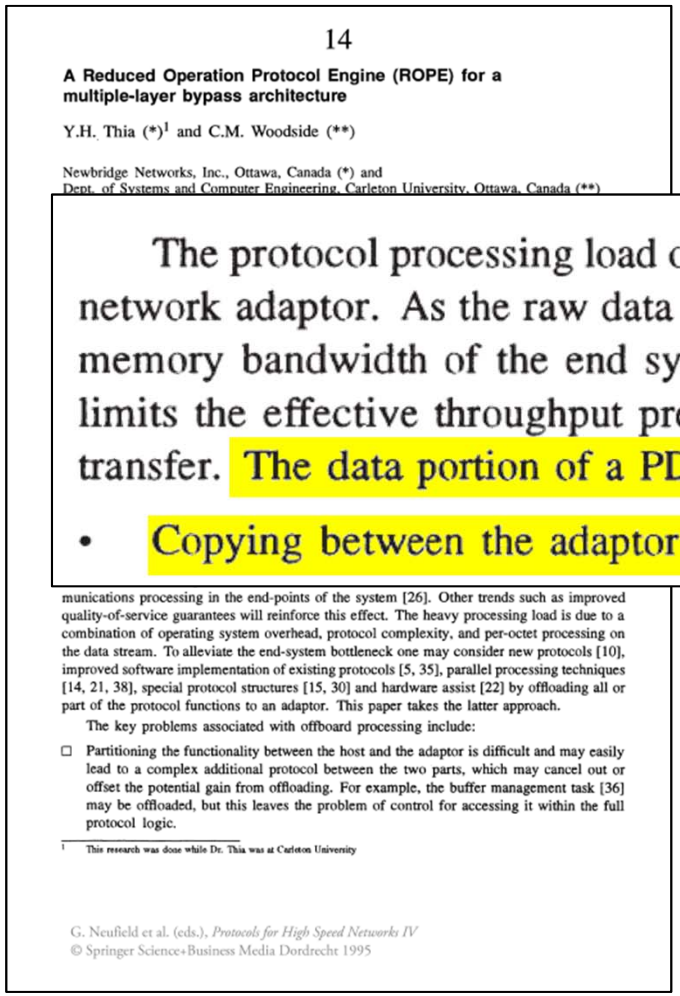

14

A Reduced Operation Protocol Engine (ROPE) for a multiple-layer bypass architecture

Y.H. Thia (*)[1] and C.M. Woodside (**)

Newbridge Networks, Inc., Ottawa, Canada (*) and
Dept. of Systems and Computer Engineering, Carleton University, Ottawa, Canada (**)

munications processing in the end-points of the system [26]. Other trends such as improved quality-of-service guarantees will reinforce this effect. The heavy processing load is due to a combination of operating system overhead, protocol complexity, and per-octet processing on the data stream. To alleviate the end-system bottleneck one may consider new protocols [10], improved software implementation of existing protocols [5, 35], parallel processing techniques [14, 21, 38], special protocol structures [15, 30] and hardware assist [22] by offloading all or part of the protocol functions to an adaptor. This paper takes the latter approach.

The key problems associated with offboard processing include:

- Partitioning the functionality between the host and the adaptor is difficult and may easily lead to a complex additional protocol between the two parts, which may cancel out or offset the potential gain from offloading. For example, the buffer management task [36] may be offloaded, but this leaves the problem of control for accessing it within the full protocol logic.

[1] This research was done while Dr. Thia was at Carleton University

G. Neufeld et al. (eds.), *Protocols for High Speed Networks IV*

The protocol processing load on an end system is typically shared between the host and the network adaptor. As the raw data bit rate supported by optical networks approaches the main memory bandwidth of the end system, the cost of moving data and of per-octet processing limits the effective throughput presented to the application process, especially for bulk data transfer. The data portion of a PDU may be physically moved for the following reasons:

- Copying between the adaptor buffer and the host system memory;

Ex. 1015.005 (Thia);
Paper 66 (Opp. To Motion to Amend) at 5;
Ex. 1262 (Lin Opp. to Supp. Motion to Amend) at ¶ 32;
Paper 39 (Opp. to Motion to Amend) at 12-13;
Ex. 1210.013-.014 (Lin Opp. to Motion to Amend) at ¶ 29.

Thia: Put incoming packets in the right order in the transport layer

4.5 Second Design, including major procedures for Transport Class 4 (Implemented)

This section describes extensions to the first design, which only supports Session BCS and TP2 functionality, to include some common TP4 functionality. Procedures for checksum, retransmission on timeout and resequencing were implemented. Extensions to the Session layer functionality and procedures for presentation layer conversion were not implemented, but are also discussed in section 6.

4.5.3 Retransmission and Resequencing

At the receiver end, out-of-sequence PDUs outside the flow-control window will be discarded. Otherwise, a PDU is buffered for resequencing. Duplicate TPDUs can be detected

Ex.1015.010 (Thia);
Paper 44 (205 Reply) at 13-14;
Paper 39 (205 Opp. To Motion to Amend) at 11-13;
Ex. 1210.011-.014 (Lin Opp. Decl.) at ¶¶ 24-30.

Thia: DMA data portions of PDUs to the host in the bypass path

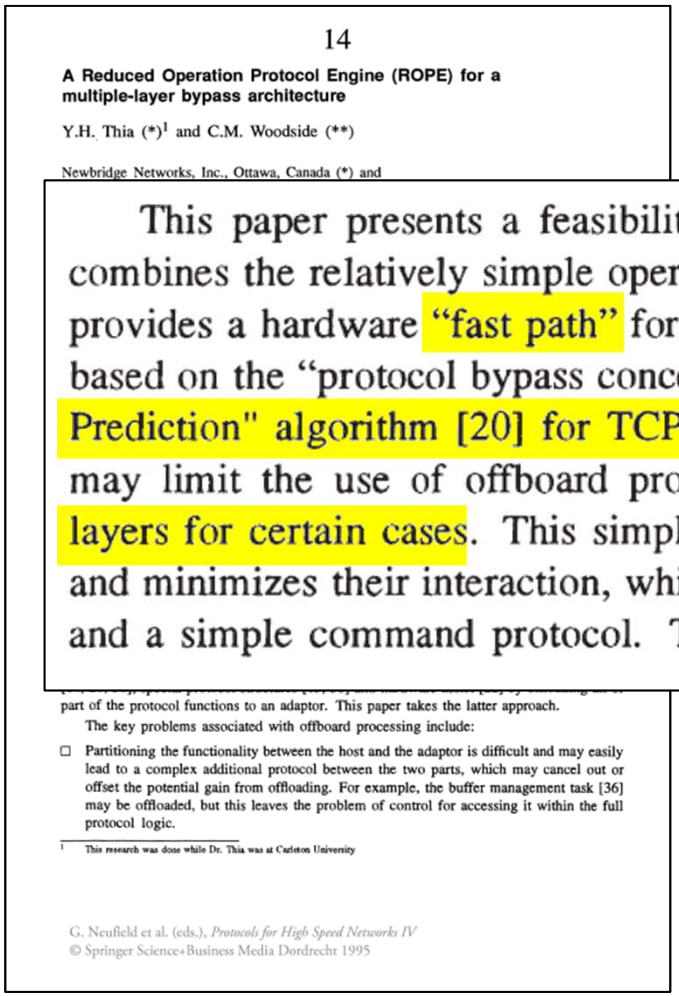

14

A Reduced Operation Protocol Engine (ROPE) for a multiple-layer bypass architecture

Y.H. Thia (*)[1] and C.M. Woodside (**)

Newbridge Networks, Inc., Ottawa, Canada (*) and

This paper presents a feasibility study for a new approach to hardware assistance. It combines the relatively simple operations needed for data transfer across multiple layers and provides a hardware "fast path" for them, which will be efficient for bulk data transfer. It is based on the "protocol bypass concept" [37] which is a generalization of Jacobson's "Header Prediction" algorithm [20] for TCP/IP. Bypass solves the problems identified above, which may limit the use of offboard processing, by implementing an entire service through all layers for certain cases. This simplifies the interface between the host and the adaptor chip and minimizes their interaction, which is supported by an access test, some DMA processing and a simple command protocol. The chip design based on bypassing is called ROPE, for

part of the protocol functions to an adaptor. This paper takes the latter approach.

The key problems associated with offboard processing include:

- Partitioning the functionality between the host and the adaptor is difficult and may easily lead to a complex additional protocol between the two parts, which may cancel out or offset the potential gain from offloading. For example, the buffer management task [36] may be offloaded, but this leaves the problem of control for accessing it within the full protocol logic.

[1] This research was done while Dr. Thia was at Carleton University

G. Neufeld et al. (eds.), *Protocols for High Speed Networks IV*

Ex. 1015.002 (Thia);
See also Paper 1 (205 Petition) at 21;
Ex. 1003.040, .065 (Lin Decl.) at ¶ 78, A-7.

205 Patent: Disputes

2. Thia discloses the performing all network and transport layer processing on the bypass path
   a. Thia's network interface device performs all network layer processing
   b. Thia's network interface device performs all transport layer processing
      i. Claims do not recite "reassembly"
      ii. Thia teaches placing data from in-order packets into host memory on the bypass path
      iii. The "segmentation/reassembly" discussed in Thia is below the transport layer

Thia's segmentation/reassembly for ATM is not transport layer reassembly

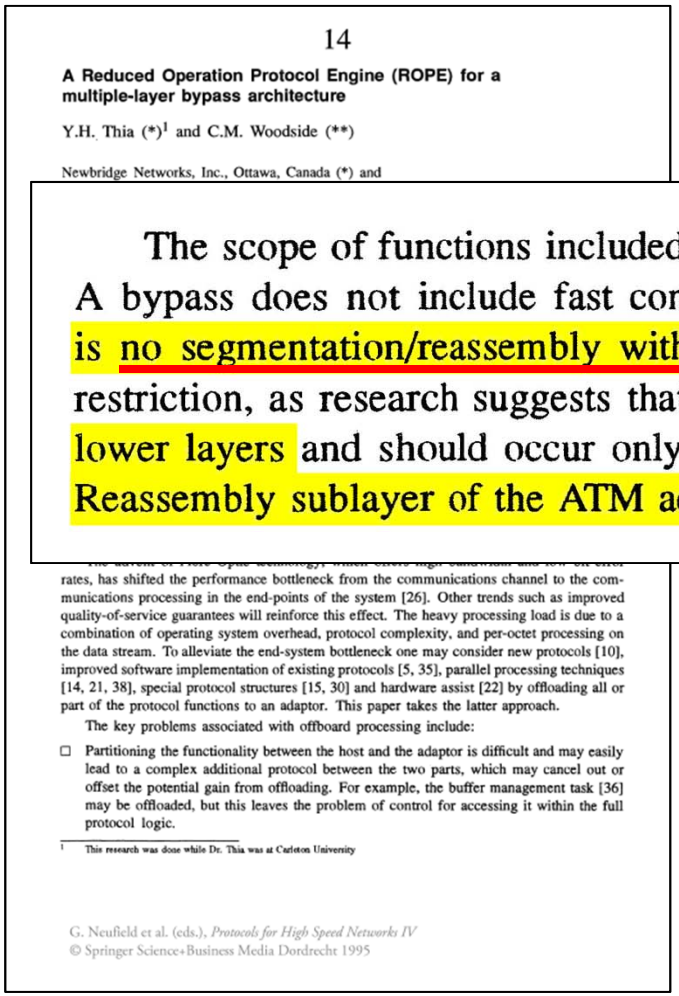

14

A Reduced Operation Protocol Engine (ROPE) for a multiple-layer bypass architecture

Y.H. Thia (*)[1] and C.M. Woodside (**)

Newbridge Networks, Inc., Ottawa, Canada (*) and

The scope of functions included in a bypass may be narrowly defined, or more extended. A bypass does not include fast connection setup but also does not interfere with it. There is no segmentation/reassembly within the bypass path, but we do not see this as a major restriction, as research suggests that fragmentation of PDUs should be restricted only to the lower layers and should occur only once in the protocol stack [23]. The Segmentation and Reassembly sublayer of the ATM adaptation layer is a good place for such functions [25].

rates, has shifted the performance bottleneck from the communications channel to the communications processing in the end-points of the system [26]. Other trends such as improved quality-of-service guarantees will reinforce this effect. The heavy processing load is due to a combination of operating system overhead, protocol complexity, and per-octet processing on the data stream. To alleviate the end-system bottleneck one may consider new protocols [10], improved software implementation of existing protocols [5, 35], parallel processing techniques [14, 21, 38], special protocol structures [15, 30] and hardware assist [22] by offloading all or part of the protocol functions to an adaptor. This paper takes the latter approach.

The key problems associated with offboard processing include:

- Partitioning the functionality between the host and the adaptor is difficult and may easily lead to a complex additional protocol between the two parts, which may cancel out or offset the potential gain from offloading. For example, the buffer management task [36] may be offloaded, but this leaves the problem of control for accessing it within the full protocol logic.

[1] This research was done while Dr. Thia was at Carleton University

G. Neufield et al. (eds.), *Protocols for High Speed Networks IV*
© Springer Science+Business Media Dordrecht 1995

Ex. 1015.014 (Thia).

Thia's "segmentation/reassembly" is fragmenting/re-assembling portions of packets at a layer below the transport layer.

See, e.g., Paper 44 (205 Reply) at 13;
Ex. 1223.009-.011 (Lin Reply Decl.) at ¶ 14;
Paper 39 (205 Opp. To Motion to Amend) at 9-11;
Ex. 1210.010-.011 (Lin Opp. Decl.) at ¶ 23.

Dr. Lin: Thia's segmentation/reassembly for ATM is not transport layer reassembly

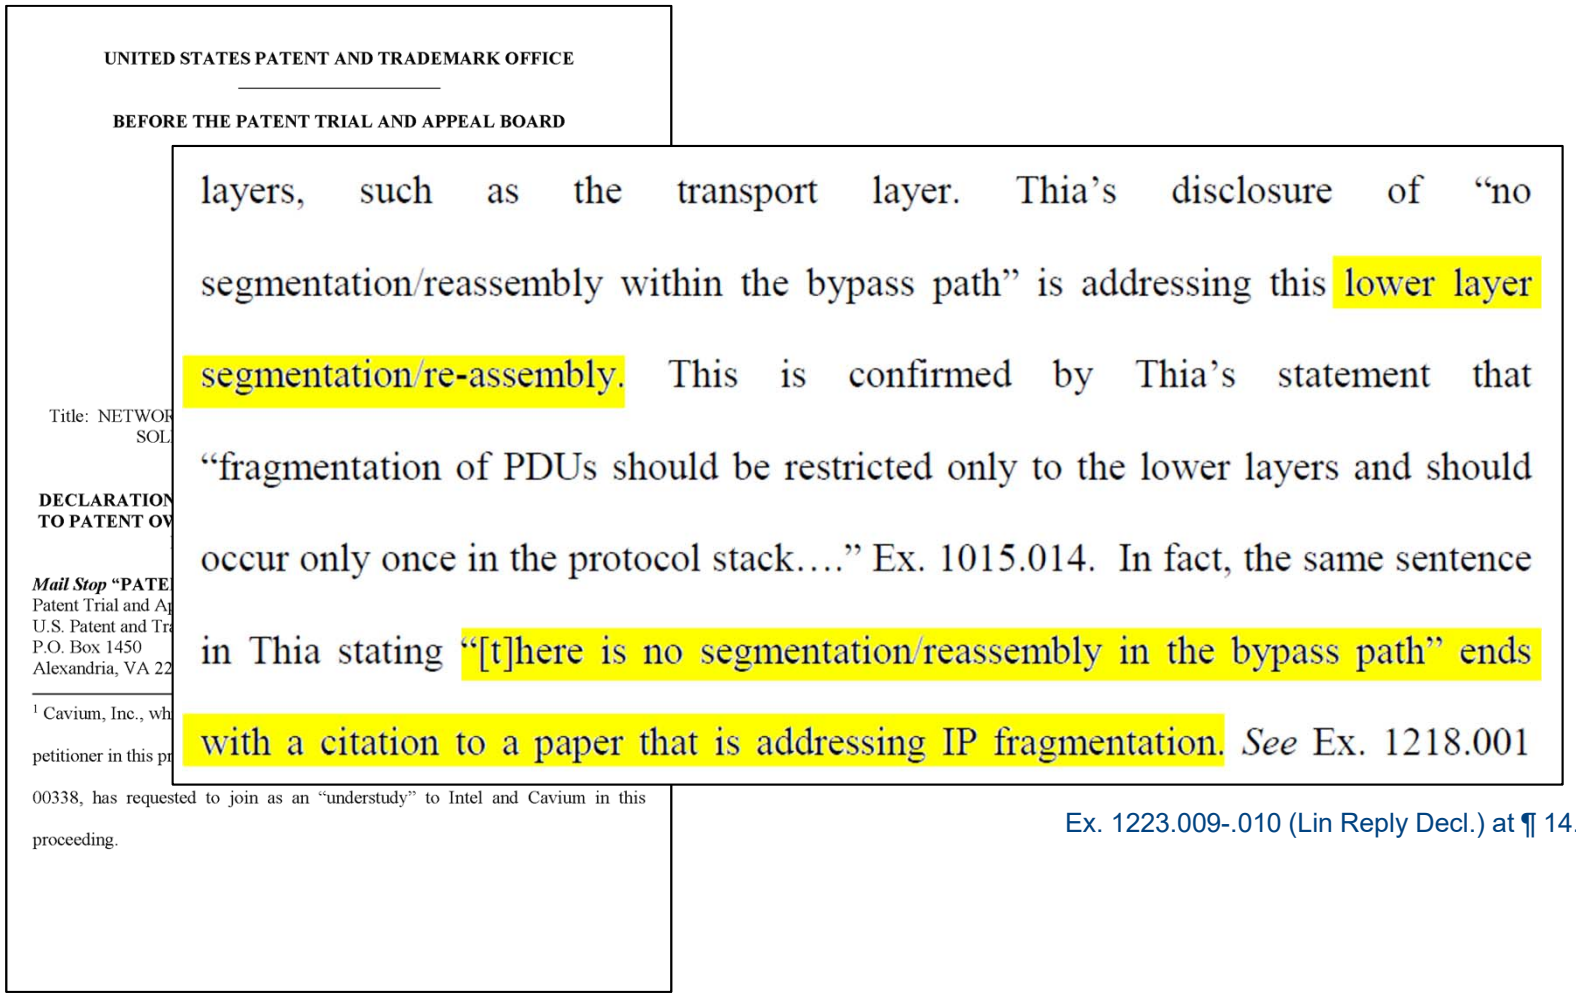

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

Title: NETWOR
SOL

DECLARATION
TO PATENT OV

Mail Stop "PATE
Patent Trial and Ap
U.S. Patent and Tr
P.O. Box 1450
Alexandria, VA 22

[1] Cavium, Inc., wh
petitioner in this pr
00338, has requested to join as an "understudy" to Intel and Cavium in this proceeding.

layers, such as the transport layer. Thia's disclosure of "no segmentation/reassembly within the bypass path" is addressing this lower layer segmentation/re-assembly. This is confirmed by Thia's statement that "fragmentation of PDUs should be restricted only to the lower layers and should occur only once in the protocol stack…." Ex. 1015.014. In fact, the same sentence in Thia stating "[t]here is no segmentation/reassembly in the bypass path" ends with a citation to a paper that is addressing IP fragmentation. *See* Ex. 1218.001

Ex. 1223.009-.010 (Lin Reply Decl.) at ¶ 14.

205 Patent: Fragmented packets are processed on “slow path” by host

- PO argues that “[0063] and [0064] similarly describe an embodiment in Figure 3 where ‘the data from the packet is sent 125 by DMA to the destination in the host file cache,’ again disclosing reassembly to a person of skill in the art.”

Paper 70 (Supplemental Motion to Amend Reply) at 2.

- But packets that are fragmented are diverted from the fast-path and processed conventionally by the host

> for the message. If the packet summary instead matches a CCB held in the INIC memory, the processor checks **114** for exception conditions which may include, e.g., fragmented or out of order packets and, if such an exception condition is found, flushes **116** the CCB and the packet to the host protocol stack **38** for protocol processing. For the case in

Ex. 2022 (205 Pub. App.) at ¶ [0063].

205 Patent: Disputes

1. Thia is enabling prior art
2. Thia discloses the network interface device performing all network and transport layer processing on the bypass path
3. **The combination of Thia, Satran I and Satran II discloses the challenged dependent limitations**
4. A POSA would have been motivated to combine Thia, Satran I and Satran II (as well as Carmichael)
5. Supplemental Briefing – claim 31 is indefinite or obvious in light of Thia, Satran I and Satran II
6. Motions to Amend 205 Patent should be denied

The challenged dependent limitations

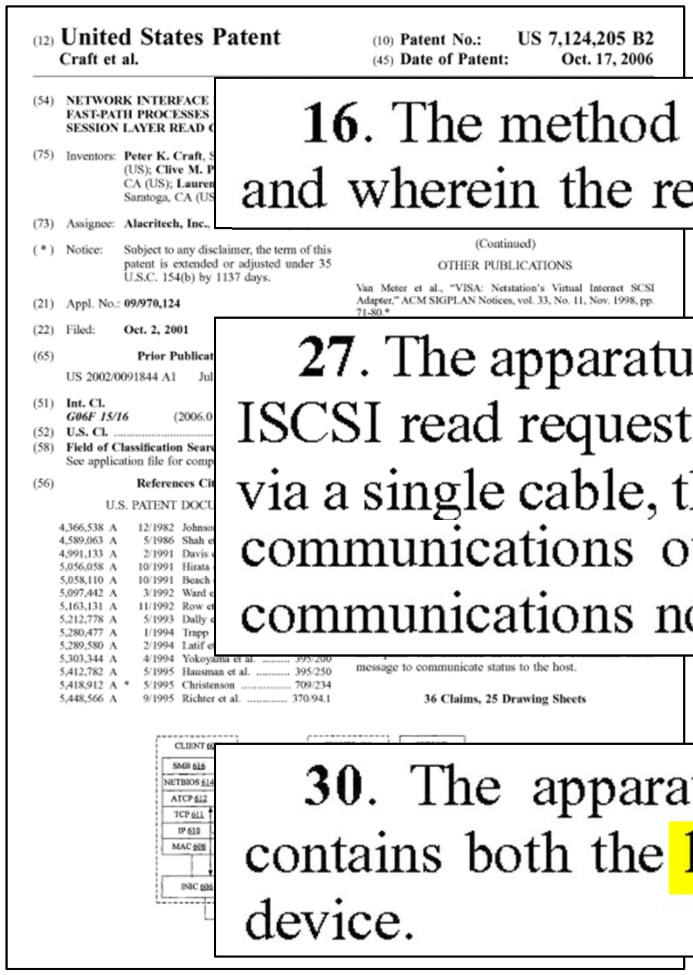
(12) **United States Patent**
Craft et al.

(10) **Patent No.: US 7,124,205 B2**
(45) **Date of Patent: Oct. 17, 2006**

(54) NETWORK INTERFACE
FAST-PATH PROCESSES
SESSION LAYER READ

(75) Inventors: **Peter K. Craft**, S
(US); **Clive M. P**
CA (US); **Lauren**
Saratoga, CA (US

(73) Assignee: **Alacritech, Inc.**,

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 1137 days.

(21) Appl. No.: **09/970,124**

(22) Filed: **Oct. 2, 2001**

(65) **Prior Publicat**
US 2002/0091844 A1 Jul

(51) **Int. Cl.**
G06F 15/16 (2006.0

(52) **U.S. Cl.** .....

(58) **Field of Classification Searc**
See application file for comp

(56) **References Cit**
U.S. PATENT DOCU

4,366,538 A 12/1982 Johnso
4,589,063 A 5/1986 Shah e
4,991,133 A 2/1991 Davis
5,056,058 A 10/1991 Hirata
5,058,110 A 10/1991 Beach
5,097,442 A 3/1992 Ward
5,163,131 A 11/1992 Row e
5,212,778 A 5/1993 Dally
5,280,477 A 1/1994 Trapp
5,289,580 A 2/1994 Latif
5,303,344 A 4/1994 Yokoyama et al. 395/200
5,412,782 A 5/1995 Hausman et al. 395/250
5,418,912 A * 5/1995 Christenson 709/234
5,448,566 A 9/1995 Richter et al. 370/94.1

(Continued)
OTHER PUBLICATIONS
Van Meter et al., "VISA: Netstation's Virtual Internet SCSI Adapter," ACM SIGPLAN Notices, vol. 33, No. 11, Nov. 1998, pp. 71-80.*

message to communicate status to the host.

36 Claims, 25 Drawing Sheets

CLIENT
SMB 616
NETBIOS 614
ATCP 612
TCP 611
IP 610
MAC 608
NIC 606

16. The method of claim **8**, wherein the bus is a PCI bus, and wherein the read request is a ISCSI read request.

Ex. 1001.051-.052 (205 Patent) at Claim 16.

27. The apparatus of claim **22**, wherein the response to the ISCSI read request command is received onto the computer via a single cable, the computer also receiving other network communications over the single cable, the other network communications not being ISCSI communications.

Ex. 1001.052 (205 Patent) at Claim 27.

30. The apparatus of claim **22**, wherein an enclosure contains both the host computer and the network interface device.

Ex. 1001.052 (205 Patent) at Claim 30.

The combination of Thia and Satran I and II discloses the challenged dependent claims

Claim 30 "enclosure"

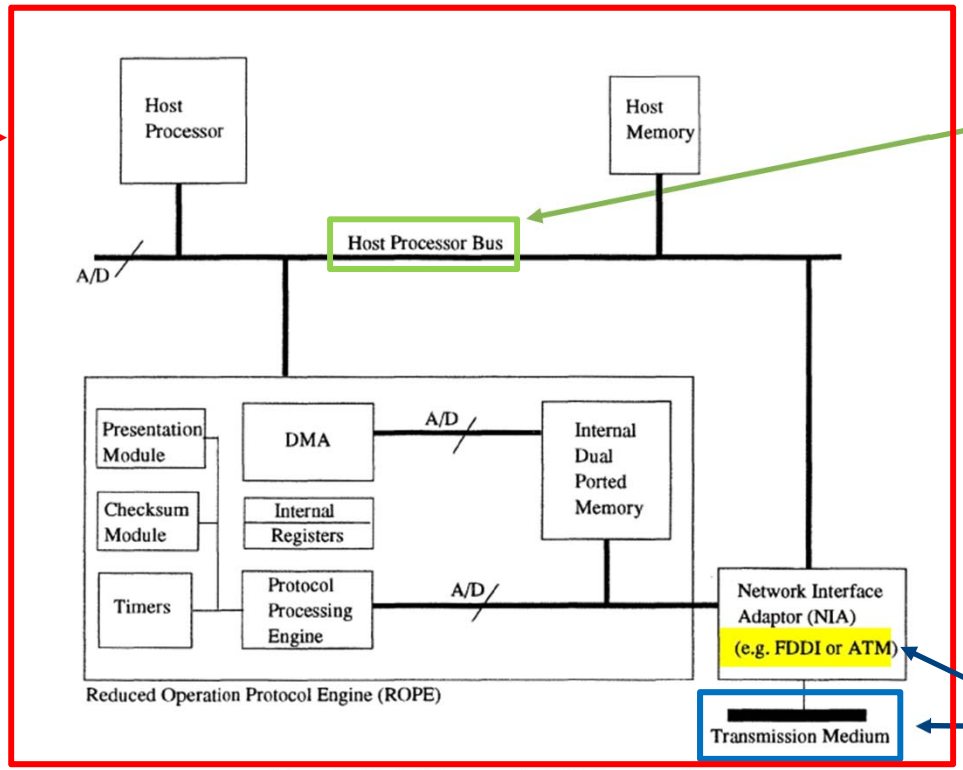

Figure 2 Block Diagram of VLSI bypass system

Claim 16 "PCI bus" (peripheral component interface)

Claim 27 "single cable"

Ex. 1015.007 (Thia) at Fig. 2 (annotated);
Paper 1 (205 Petition) at 62-63, 71-74, 78-80 (combined).

205 Patent: Disputes

4. A POSA would have been motivated to combine Thia, Satran I and Satran II (as well as Carmichael)
   a. **A POSA would have used Thia's bypass system with the iSCSI protocol of Satran I and Satran II**
   b. A POSA would have looked to both Satran I and Satran II
   c. The motivations to further include Carmichael are unrebutted by Patent Owner
   d. The Petition includes sufficient evidence regarding expectation of success

Thia's bypass would have been improved by Satran's iSCSI

A POSA would have been motivated to combine Thia and Satran. Ex.1003, Lin Decl. ¶¶92-99. Thia discloses a fast-path bypass system adapted for the OSI model. *See* Ex.1015, Thia at Abstract, at .008. Satran discloses a communication protocol that can be used with OSI. Ex.1003, Lin Decl. ¶97. Note that iSCSI is commonly known for TCP but it was known to a have applicability to the OSI model. *Id;* Ex.1001, 205 Patent at 38:47-51 ("The iSCSI layer 2412 is generally

Paper 1 (205 Petition) at 33;
see also Ex.1003.048 (Lin Decl.) at ¶¶ 94-97.

iSCSI operates at the session layer

- The 205 Patent acknowledges iSCSI operates at the session layer in the OSI protocol stack

> (labeled in FIG. **27** "ISCSI layer"). The ISCSI layer **2412** is generally considered to involve the application layer in the TCP/IP protocol stack model, whereas the ISCSI layer **2412** is generally considered to involve the session, presentation and application layers in the OSI protocol stack model.

Ex.1001.048 (205 Patent) at 38:47-51;
Paper 1 (205 Petition) at 43-44.

- Claim 3 requires that iSCSI operates at the session layer

> **3**. The apparatus of claim **1**, wherein the session layer protocol is ISCSI.

Ex.1001.051 (205 Patent) at Claim 3.

PO's criticism combining Thia and Satran is that Thia is a theoretical reference

> Petitioners' proffered "Motivations To Combine Thia and Satran" ***all*** erroneously assume that Thia discloses a real-world chip, and purport to offer motivations to combine that supposed real-world chip with either of the Satran references in a real-world environment. (Intel Pet. 34-36; Cavium Pet. 33-35.)
>
> Because Thia actually only discloses a feasibility study using a theoretical chip in a simulated environment, those proffered motivations fail. (Ex. 2026 at ¶ 114.)

Paper 31 (205 Response) at 50.

205 Patent: Disputes

4. A POSA would have been motivated to combine Thia, Satran I and Satran II (as well as Carmichael)
   a. A POSA would have used Thia's bypass system and ROPE chip with the iSCSI protocol and read requests of Satran I and Satran II
   b. <u>A POSA would have looked to both Satran I and Satran II</u>
   c. The motivations to further include Carmichael are unrebutted by Patent Owner
   d. The Petition includes sufficient evidence regarding expectation of success

A POSA would have looked to both Satran I and Satran II

- Petitioner addressed Satran I and Satran II together given their close relationship
  - Overlapping authors
  - Same functionality
  - Satran II is a revision of Satran I

Paper 1 (205 Petition) at 14-15, 32;
Ex. 1003.044 (Lin Decl.) at ¶ 87 and FN4.

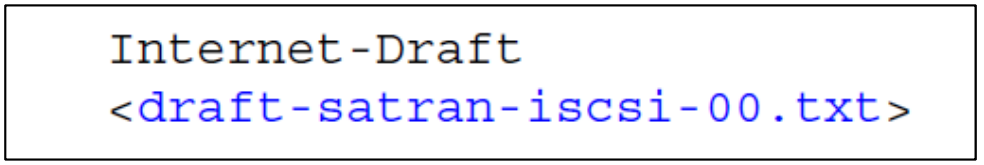

Ex.1056.001 (Satran I).

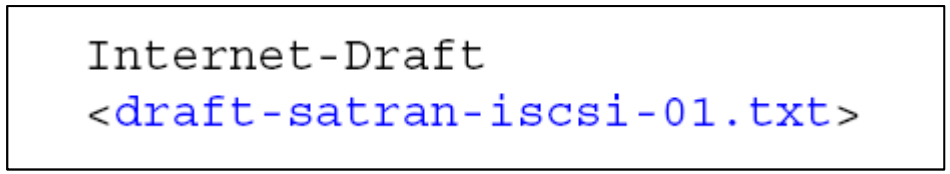

Ex.1057.001 (Satran II).

Satran I is an earlier version of Satran II

Intel's citations to either exhibit 1056 or 1057 for specific claim limitations. Neither references is referred to as a newer version of the other, nor is either reference a renamed version of the other (nor would they, as they both include an "expiration" date). In short, and by every objective indication, these are separate publications, authored by a different set of people, at different times, with different disclosures.

Paper 31 (205 Response) at 30.

future. SANs **418** and **420** may run a storage protocol such as SCSI over TCP/IP or SCSI Encapsulation Protocol. One such storage protocol is described by J. Satran et al. in the Internet-Draft of the Internet Engineering Task Force (IETF) entitled "iSCSI (Internet SCSI)," June 2000, which in an earlier Internet-Draft was entitled "SCSI/TCP (SCSI over TCP)," February 2000, both documents being incorporated by reference herein. Another such protocol, termed

Ex.1001.036 (205 Patent) at 13:28-35;
Paper 1 (205 Petition) at 44.
See also Paper 1 (205 Petition) at 32.

iSCSI (Internet SCSI)

Ex.1057.001 (Satran II).

SCSI/TCP (SCSI over TCP)

Ex.1056.001 (Satran I).

205 Patent: Disputes

4. A POSA would have been motivated to combine Thia, Satran I and Satran II (as well as Carmichael)
   a. A POSA would have used Thia's bypass system and ROPE chip with the iSCSI protocol and read requests of Satran I and Satran II
   b. A POSA would have looked to both Satran I and Satran II
   c. <u>The motivations to further include Carmichael are unrebutted by Patent Owner</u>
   d. The Petition includes sufficient evidence regarding expectation of success

A POSA would have been motivated to further include Carmichael

- Patent Owner does not address motivations to further include Carmichael

Paper 31 (205 Response) at 53-54.

- Evidence in Petition is unrebutted

Paper 44 (205 Reply) at 19.

205 Patent: Disputes

4. A POSA would have been motivated to combine Thia, Satran I and Satran II (as well as Carmichael)
   a. A POSA would have used Thia's bypass system and ROPE chip with the iSCSI protocol and read requests of Satran I and Satran II
   b. A POSA would have looked to both Satran I and Satran II
   c. The motivations to further include Carmichael are unrebutted by Patent Owner
   d. <u>The Petition includes sufficient evidence regarding reasonable expectation of success</u>

Dr. Lin: Reasonable expectation of success

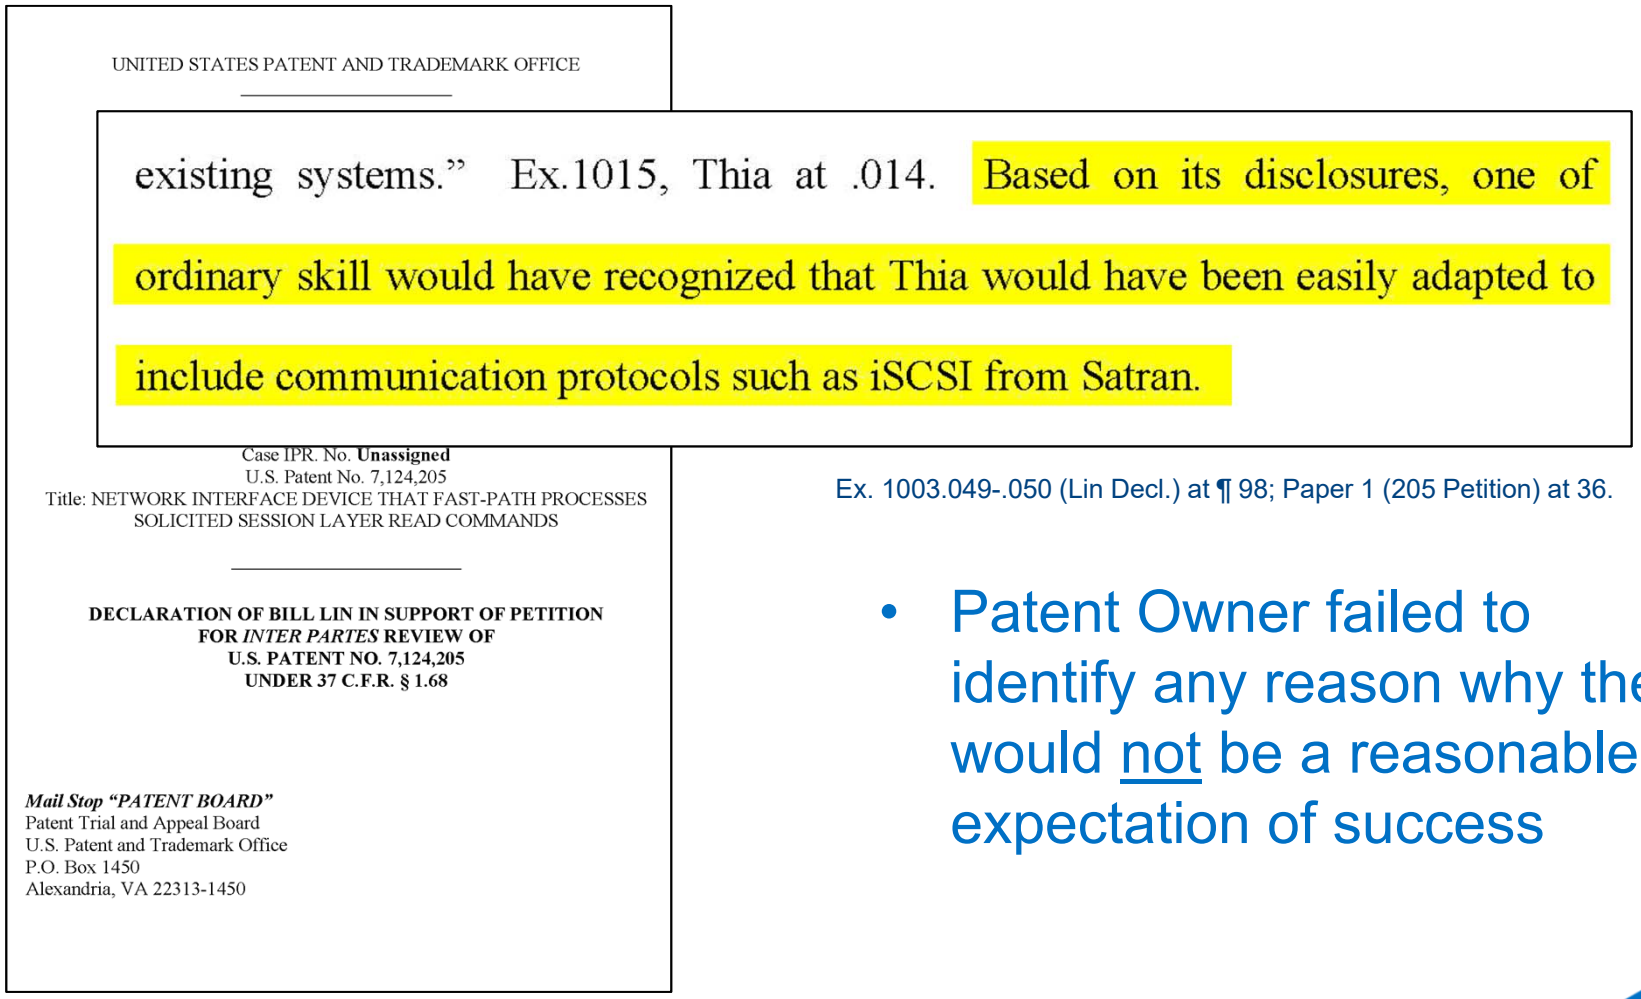

UNITED STATES PATENT AND TRADEMARK OFFICE

existing systems." Ex.1015, Thia at .014. Based on its disclosures, one of ordinary skill would have recognized that Thia would have been easily adapted to include communication protocols such as iSCSI from Satran.

Case IPR. No. **Unassigned**
U.S. Patent No. 7,124,205
Title: NETWORK INTERFACE DEVICE THAT FAST-PATH PROCESSES SOLICITED SESSION LAYER READ COMMANDS

DECLARATION OF BILL LIN IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW OF U.S. PATENT NO. 7,124,205 UNDER 37 C.F.R. § 1.68

Mail Stop "PATENT BOARD"
Patent Trial and Appeal Board
U.S. Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

Ex. 1003.049-.050 (Lin Decl.) at ¶ 98; Paper 1 (205 Petition) at 36.

- Patent Owner failed to identify any reason why there would not be a reasonable expectation of success

Dr. Lin: Combination not unduly complicated with a predictable result

UNITED STATES PATENT AND TRADEMARK OFFICE

Finally, a POSA would have recognized that combining Thia and Satran would not have been unduly complicated and would have had a predictable result – Thia providing fast-path capability to Satran communications. Ex.1003, Lin Decl. ¶98. Specifically, iSCSI has general applicability to the OSI model. *See, e.g.,*

DECLARATION OF BILL LIN IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW OF U.S. PATENT NO. 7,124,205 UNDER 37 C.F.R. § 1.68

Mail Stop "PATENT BOARD"
Patent Trial and Appeal Board
U.S. Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

Ex. 1003.049 (Lin Decl.) at ¶ 98; see also Paper 1 (205 Petition) at 36.

205 Patent: Disputes

5. Supplemental Briefing – claim 31 is indefinite or obvious in light of Thia, Satran I and Satran II
   a. The only new dispute is the corresponding structure for the means plus function elements in Claim 31

Claims 31-33 are invalid

- The Petition addressed claim 31 in the alternative
  - The claim is indefinite for lack of corresponding structure

    or

  - The claim is invalid based on the prior art grounds in the Petition (assuming the host is part of the “means”)

The petition addressed claim 31 in the alternative

Petitioner contends that this limitation is indefinite. Ex.1003, Lin Decl. ¶¶74-76. The claimed means must perform four functions: (1) receiving a response, (2) fast-path processing a portion of the response, (3) receiving a subsequent portion of the response, and (4) slow-path processing the subsequent portion. Importantly, the claimed means must also be coupled to the host computer. The 205 Patent discloses no structure coupled to the host that is capable of both fast-path and slow-path processing. *Id.* At best, the 205 Patent discloses a network adaptor that performs functions (1)-(3). *Id.; see, e.g.,* Ex.1001, 205 Patent at 8:61-64 and 39:48-49 (function 1, receiving), at 8:25-40 (function 2, fast-path processing), at 40:19-26 (function 3, receiving subsequent).

The 205 Patent, however, discloses that function 4 is performed by the host, not the network adaptor. *See, e.g., id.* 8:25-40, 36:17-21 (function 4, slow-path processing). Accordingly, the 205 Patent fails to disclose any structure, coupled to the host computer, that performs all four functions and is therefore indefinite.

Paper 1 (205 Petition) at 29.

Dispute is whether the last function requires the “means” to include the host

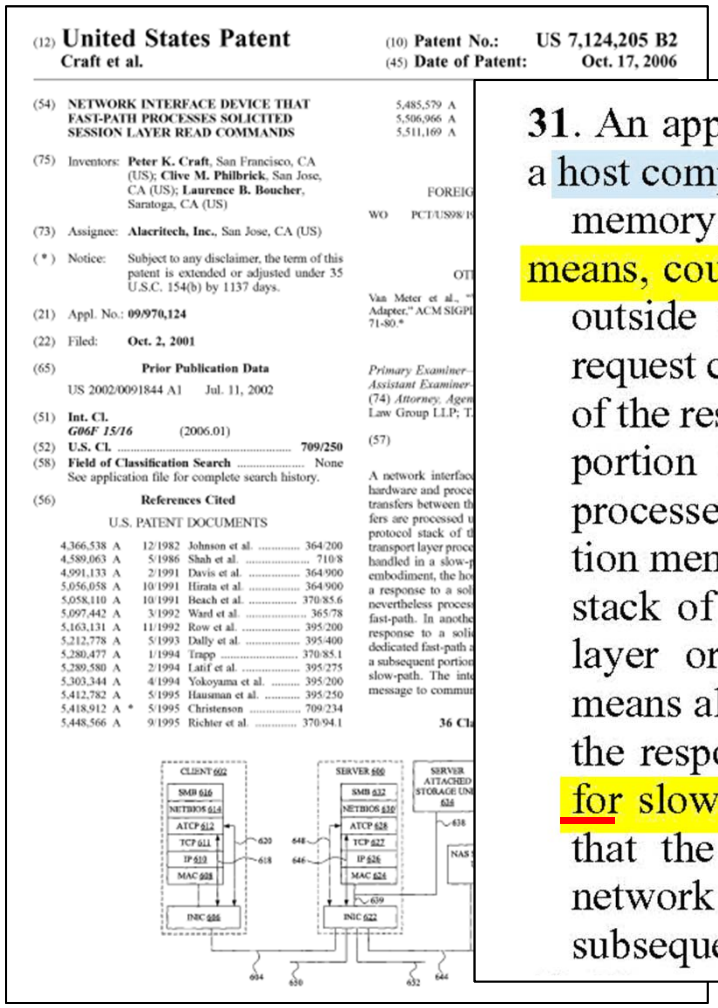

(12) **United States Patent**
Craft et al.

(10) **Patent No.: US 7,124,205 B2**
(45) **Date of Patent: Oct. 17, 2006**

(54) **NETWORK INTERFACE DEVICE THAT FAST-PATH PROCESSES SOLICITED SESSION LAYER READ COMMANDS**

(75) Inventors: **Peter K. Craft**, San Francisco, CA (US); **Clive M. Philbrick**, San Jose, CA (US); **Laurence B. Boucher**, Saratoga, CA (US)

(73) Assignee: **Alacritech, Inc.**, San Jose, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 1137 days.

(21) Appl. No.: **09/970,124**

(22) Filed: **Oct. 2, 2001**

(65) **Prior Publication Data**

US 2002/0091844 A1 Jul. 11, 2002

(51) **Int. Cl.**
G06F 15/16 (2006.01)

(52) **U.S. Cl.** **709/250**

(58) **Field of Classification Search** ..... None
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,366,538	A	12/1982	Johnson et al.	364/200
4,589,063	A	5/1986	Shah et al.	710/8
4,991,133	A	2/1991	Davis et al.	364/900
5,056,058	A	10/1991	Hirata et al.	364/900
5,058,110	A	10/1991	Beach et al.	370/85.6
5,097,442	A	3/1992	Ward et al.	365/78
5,163,131	A	11/1992	Row et al.	395/200
5,212,778	A	5/1993	Dally et al.	395/400
5,280,477	A	1/1994	Trapp	370/85.1
5,289,580	A	2/1994	Latif et al.	395/275
5,303,344	A	4/1994	Yokoyama et al.	395/200
5,412,782	A	5/1995	Hausman et al.	395/250
5,418,912	A *	5/1995	Christenson	709/234
5,448,566	A	9/1995	Richter et al.	370/94.1
5,485,579	A			
5,506,966	A			
5,511,169	A			

FOREIG

WO PCT/US98/1

OT

Van Meter et al., "
Adapter," ACM SIGP
71-80.*

Primary Examiner
Assistant Examiner
(74) *Attorney, Agen*
Law Group LLP; T

(57)

A network interfac
hardware and proce
transfers between th
fers are processed t
protocol stack of t
transport layer proce
handled in a slow-p
embodiment, the ho
a response to a sol
nevertheless proces
fast-path. In anothe
response to a soli
dedicated fast-path
a subsequent portio
slow-path. The int
message to commu

36 Cl

31. An apparatus comprising:

a host computer having a protocol stack and a destination memory; and

means, coupled to the host computer, for receiving from outside the apparatus a response to an ISCSI read request command and for fast-path processing a portion of the response to the ISCSI read request command, the portion including data, the portion being fast-path processed such that the data is placed into the destination memory on the host computer without the protocol stack of the host computer doing significant network layer or significant transport layer processing, the means also being for receiving a subsequent portion of the response to the ISCSI read request command and for slow-path processing the subsequent portion such that the protocol stack of the host computer does network layer and transport layer processing on the subsequent portion.

Ex.1001.052 (205 Patent) at Claim 31.

Function 4, “slow path processing,” is performed on the host

- No dispute that “slow path processing” is performed on the host

> Embodiments are described wherein an ISCSI read request command is sent from a network interface device to an ISCSI target. The network interface device does fast-path processing on an initial part of the response from the ISCSI target, but then switches to slow-path processing such that a subsequent part of the response from the ISCSI target is processed by the host protocol stack. In some embodiments,

Ex.1001.031 (205 Patent) at 4:40-46.

- Patent Owner never explains where the 205 specification links its proposed means (the network interface device) to function 4

PO's construction is at odds with the dependent limitations

- Patent Owner's construction – the INIC is the means – is directly contradicted by dependent claims 32 and 33:
  - If the means (INIC) performs "slow path processing *the subsequent portion* such that the protocol stack of the host does network layer and transport layer processing on the subsequent portion"
  - Then, Claims 32 and 33 are at odds because they expressly require that "the network layer and transport layer processing done *on the subsequent portion* by the means includes…"
  - The network layer and transport layer processing "*on the subsequent portion*" cannot be done by both the protocol stack of the host (claim 31) and the means (INIC) (claims 32 and 33)

Claims depending from claim 31

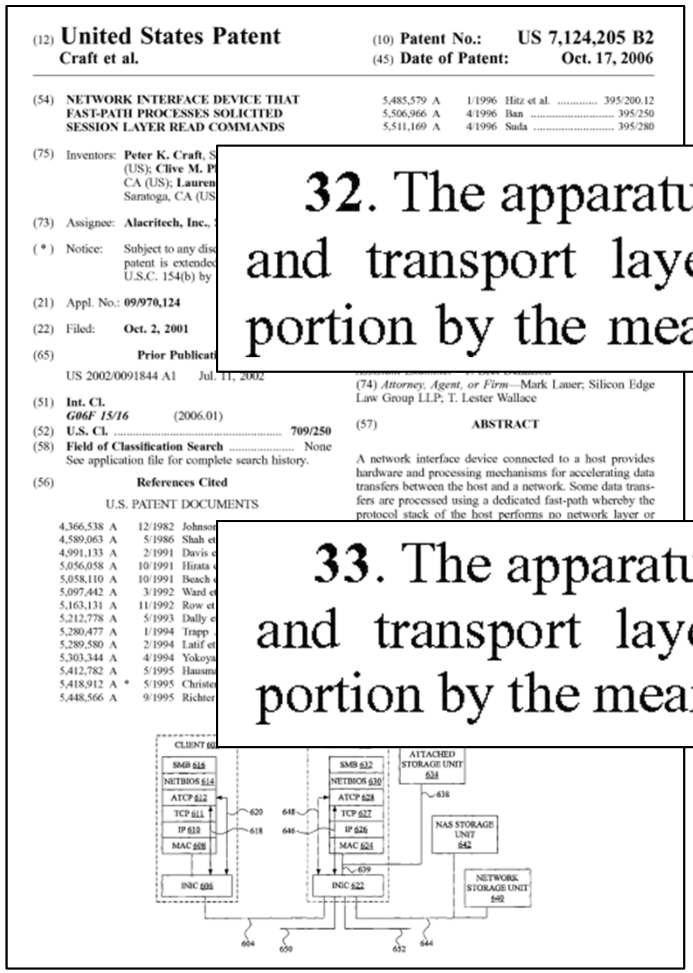

(12) **United States Patent**
Craft et al.

(10) **Patent No.: US 7,124,205 B2**
(45) **Date of Patent: Oct. 17, 2006**

(54) **NETWORK INTERFACE DEVICE THAT FAST-PATH PROCESSES SOLICITED SESSION LAYER READ COMMANDS**

(75) Inventors: **Peter K. Craft**, S (US); **Clive M. P** CA (US); **Lauren** Saratoga, CA (US

(73) Assignee: **Alacritech, Inc.**,

(*) Notice: Subject to any dis patent is extende U.S.C. 154(b) by

(21) Appl. No.: **09/970,124**

(22) Filed: **Oct. 2, 2001**

(65) **Prior Publicat**
US 2002/0091844 A1 Jul. 11, 2002

(51) **Int. Cl.**
G06F 15/16 (2006.01)

(52) **U.S. Cl.** **709/250**

(58) **Field of Classification Search** ..... None
See application file for complete search history.

(56) **References Cited**
U.S. PATENT DOCUMENTS

4,366,538	A	12/1982	Johnso	
4,589,063	A	5/1986	Shah et	
4,991,133	A	2/1991	Davis e	
5,056,058	A	10/1991	Hirata	
5,058,110	A	10/1991	Beach	
5,097,442	A	3/1992	Ward et	
5,163,131	A	11/1992	Row et	
5,212,778	A	5/1993	Dally e	
5,280,477	A	1/1994	Trapp	
5,289,580	A	2/1994	Latif et	
5,303,344	A	4/1994	Yokoya	
5,412,782	A	5/1995	Hausm	
5,418,912	A *	5/1995	Christe	
5,448,566	A	9/1995	Richter	

5,485,579	A	1/1996	Hitz et al.	395/200.12
5,506,966	A	4/1996	Ban	395/250
5,511,169	A	4/1996	Suda	395/280

(74) *Attorney, Agent, or Firm*—Mark Lauer; Silicon Edge Law Group LLP; T. Lester Wallace

(57) **ABSTRACT**

A network interface device connected to a host provides hardware and processing mechanisms for accelerating data transfers between the host and a network. Some data transfers are processed using a dedicated fast-path whereby the protocol stack of the host performs no network layer or

32. The apparatus of claim **31**, wherein the network layer and transport layer processing done on the subsequent portion by the means includes error condition handling.

Ex. 1001.052 (205 Patent), Claim 32.

33. The apparatus of claim **31**, wherein the network layer and transport layer processing done on the subsequent portion by the means includes exception condition handling.

Ex. 1001.052 (205 Patent), Claim 33.

Alternatively, Claims 31-33 are invalid if they are not found indefinite

- The Petition fully analyzed claim 31 under an interpretation where the last function (slow path processing) is performed by the host

> To the extent that the Office determines that this limitation is not means plus function or that the means includes the host computer, Thia in combination with Satran discloses this limitation. *Id.*

Paper 1 (205 Petition) at 86.

- Thia, Satran I and Satran II disclose fast path processing by the network interface device and slow path processing by the host in response to portions of an iSCSI read request

Paper 1 (205 Petition) at 81-87;
Ex. 1003.120-.131 (Lin Decl.) at A-62 – A-73.

205 Patent: Disputes

6. Motions to Amend 205 Patent should be denied

   a. **Patent Owner has not met its burden of production under 35 U.S.C. § 316(d) due to its failure to provide adequate written description support**

   b. The prior art discloses each limitation of the substitute claims

PO only provided string citations

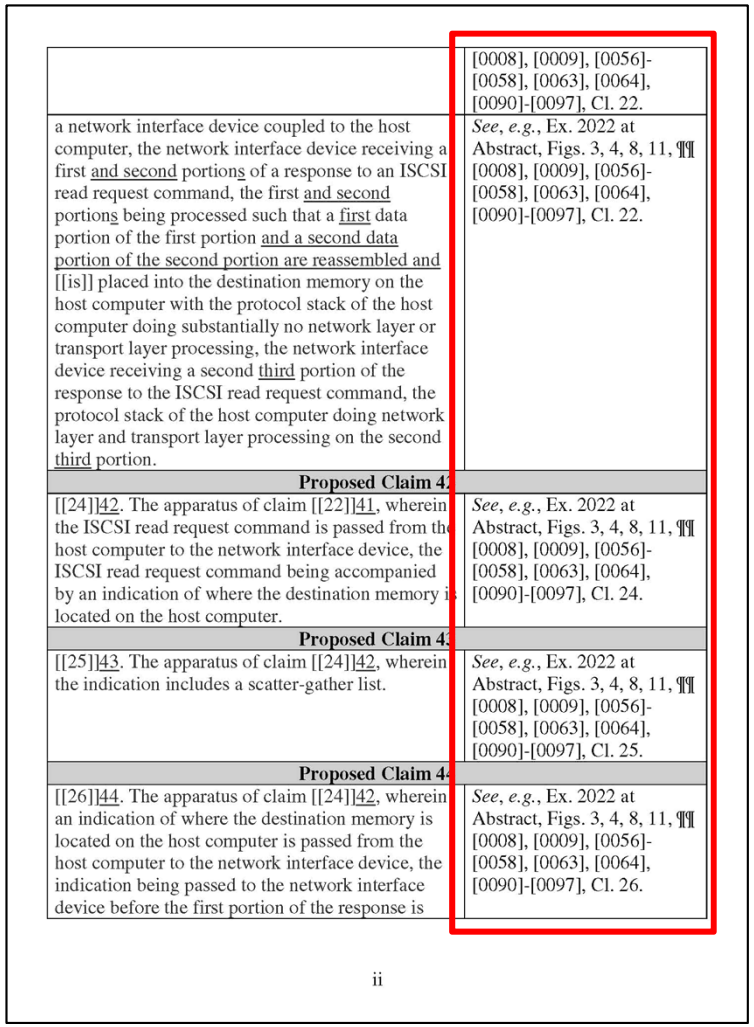

	[0008], [0009], [0056]-[0058], [0063], [0064], [0090]-[0097], Cl. 22.
a network interface device coupled to the host computer, the network interface device receiving a first and second portions of a response to an ISCSI read request command, the first and second portions being processed such that a first data portion of the first portion and a second data portion of the second portion are reassembled and [[is]] placed into the destination memory on the host computer with the protocol stack of the host computer doing substantially no network layer or transport layer processing, the network interface device receiving a second third portion of the response to the ISCSI read request command, the protocol stack of the host computer doing network layer and transport layer processing on the second third portion.	*See*, *e.g.*, Ex. 2022 at Abstract, Figs. 3, 4, 8, 11, ¶¶ [0008], [0009], [0056]-[0058], [0063], [0064], [0090]-[0097], Cl. 22.
Proposed Claim 4	
[[24]]42. The apparatus of claim [[22]]41, wherein the ISCSI read request command is passed from the host computer to the network interface device, the ISCSI read request command being accompanied by an indication of where the destination memory is located on the host computer.	*See*, *e.g.*, Ex. 2022 at Abstract, Figs. 3, 4, 8, 11, ¶¶ [0008], [0009], [0056]-[0058], [0063], [0064], [0090]-[0097], Cl. 24.
Proposed Claim 4	
[[25]]43. The apparatus of claim [[24]]42, wherein the indication includes a scatter-gather list.	*See*, *e.g.*, Ex. 2022 at Abstract, Figs. 3, 4, 8, 11, ¶¶ [0008], [0009], [0056]-[0058], [0063], [0064], [0090]-[0097], Cl. 25.
Proposed Claim 4	
[[26]]44. The apparatus of claim [[24]]42, wherein an indication of where the destination memory is located on the host computer is passed from the host computer to the network interface device, the indication being passed to the network interface device before the first portion of the response is	*See*, *e.g.*, Ex. 2022 at Abstract, Figs. 3, 4, 8, 11, ¶¶ [0008], [0009], [0056]-[0058], [0063], [0064], [0090]-[0097], Cl. 26.

ii

Paper 20, Appendix A (205 Motion to Amend) at ii.

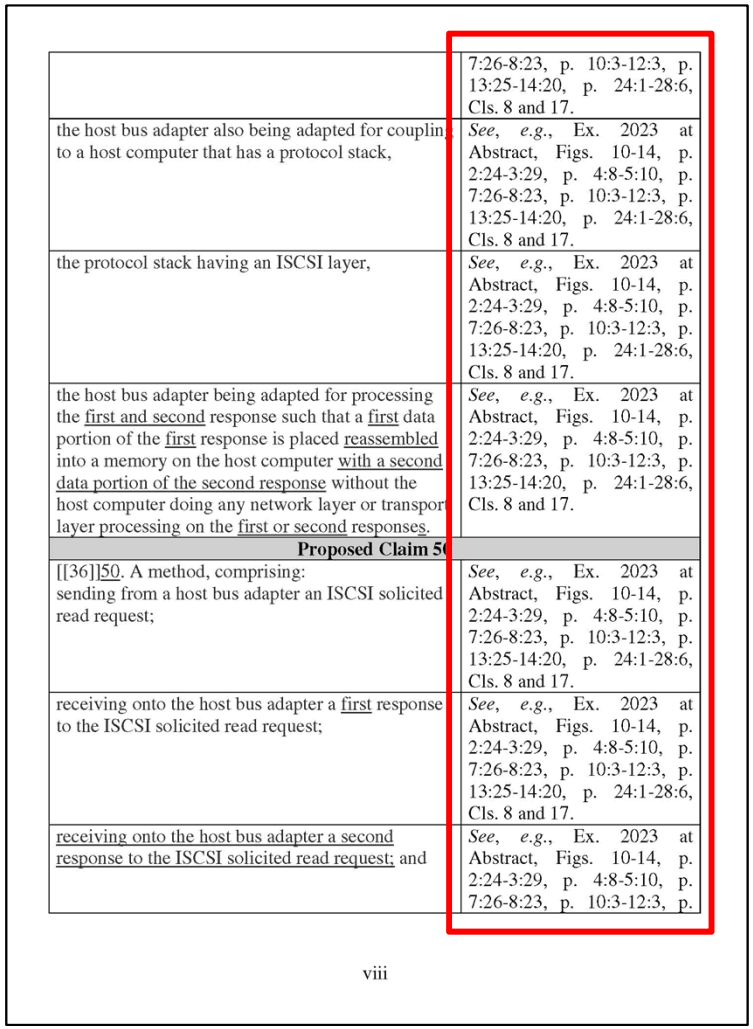

	7:26-8:23, p. 10:3-12:3, p. 13:25-14:20, p. 24:1-28:6, Cls. 8 and 17.
the host bus adapter also being adapted for coupling to a host computer that has a protocol stack,	*See*, *e.g.*, Ex. 2023 at Abstract, Figs. 10-14, p. 2:24-3:29, p. 4:8-5:10, p. 7:26-8:23, p. 10:3-12:3, p. 13:25-14:20, p. 24:1-28:6, Cls. 8 and 17.
the protocol stack having an ISCSI layer,	*See*, *e.g.*, Ex. 2023 at Abstract, Figs. 10-14, p. 2:24-3:29, p. 4:8-5:10, p. 7:26-8:23, p. 10:3-12:3, p. 13:25-14:20, p. 24:1-28:6, Cls. 8 and 17.
the host bus adapter being adapted for processing the first and second response such that a first data portion of the first response is placed reassembled into a memory on the host computer with a second data portion of the second response without the host computer doing any network layer or transport layer processing on the first or second responses.	*See*, *e.g.*, Ex. 2023 at Abstract, Figs. 10-14, p. 2:24-3:29, p. 4:8-5:10, p. 7:26-8:23, p. 10:3-12:3, p. 13:25-14:20, p. 24:1-28:6, Cls. 8 and 17.
Proposed Claim 5	
[[36]]50. A method, comprising: sending from a host bus adapter an ISCSI solicited read request;	*See*, *e.g.*, Ex. 2023 at Abstract, Figs. 10-14, p. 2:24-3:29, p. 4:8-5:10, p. 7:26-8:23, p. 10:3-12:3, p. 13:25-14:20, p. 24:1-28:6, Cls. 8 and 17.
receiving onto the host bus adapter a first response to the ISCSI solicited read request;	*See*, *e.g.*, Ex. 2023 at Abstract, Figs. 10-14, p. 2:24-3:29, p. 4:8-5:10, p. 7:26-8:23, p. 10:3-12:3, p. 13:25-14:20, p. 24:1-28:6, Cls. 8 and 17.
receiving onto the host bus adapter a second response to the ISCSI solicited read request; and	*See*, *e.g.*, Ex. 2023 at Abstract, Figs. 10-14, p. 2:24-3:29, p. 4:8-5:10, p. 7:26-8:23, p. 10:3-12:3, p.

viii

Paper 20, Appendix B (205 Motion to Amend) at viii.

PO failed to identify corresponding structure in the specification for each function

- PO did not identify corresponding structure for each function in claim 31 in its Response, supplemental Response, or Motion to amend

See Papers 31, 56, and 57.

- PO relies on the district court's construction and only provides string citations

> The district court's construction is supported by the ample disclosure of the INIC performing all the "means for" functions recited in claim 31. *See* Ex. 1010 at Figs. 1-2 and Figs. 5-12, 4:42–46, 17:13–18, 39:39–53, 39:57–62, 40:12–14.

Paper 56 (Supplemental Response) at 5-6.

205 Patent: Disputes

6. Motions to Amend 205 Patent should be denied

a. Patent Owner has not met its burden of production under 35 U.S.C. § 316(d) due to its failure to provide adequate written description support

b. <u>The prior art discloses each limitation of the substitute claims</u>

i. <u>"reassembled" / "assembles"</u>

PO's proposed amendments

- **wherein the fast-path processing reassembles the [data portion of the response/data of the packet] with a second [data portion of a second response/data portion of a second packet]** (claims 3, 9, 16)
- **wherein the fast-path processing assembles the data of the packet with a second data of a second packet** (claim 10)
- the first **and second** portion**s** being processed such that a **first** data portion of the first portion **and a second data portion of the second portion are reassembled and** placed into the destination memory (claim 22)
- processing the **first and second** response such that a **first** data portion of the first response is placed **reassembled** into a memory on the host computer **with a second data portion of the second response** (claims 35, 36)

Paper 20 (Motion to Amend) at Appendix C.

PO's proposed supplemental amendment

- means, coupled to the host computer, for receiving from outside the apparatus a response to an ISCSI read request command and for fast-path processing a **first and second** portion**s** of the response to the ISCSI read request command, the portion**s** including **first and second** data, the portion**s** being fast-path processed such that the **first and second** data **are** **reassembled** and placed into the destination memory on the host computer without the protocol stack of the host computer doing significant network layer or significant transport layer processing (claim 31)

Paper 57 (Second Motion to Amend) at Appendix C.

Thia discloses “reassembly” as claimed in PO’s contingent amendments

- As described in Petitioner’s reply to PO’s response and PO’s oppositions to PO’s motions to amend, Thia in combination with Satran I and Satran II and further in combination with Carmichael, disclose the “reassembly” limitations claimed in PO’s contingent amendments

Paper 44 (205 Reply) at 12-14;
Paper 39 (Opp. to Motion to Amend);
Paper 66 (Opp. to Supplemental Motion to Amend).

- See slides 306-314

U.S. Patent No. 9,055,104
(104 Patent)

IPR2017-01393 (Intel)
IPR2018-00374 (Dell)
IPR2017-01714 (Cavium)

*All citations herein are to the IPR2017-01405 case unless otherwise noted.

104 Patent: Instituted Grounds

- **Connery**
  - Claims 1, 6, 9, 12, and 15
  - Claim 22 in supplemental briefing

- **Connery** in view of **Boucher**
  - Claims 1, 6, 9, 12, and 15

Ex. 1043 - U.S. Patent No. 5,937,169 (“Connery”)
Ex. 1049 - WO00/13091 (“Boucher”)

104 Patent: Disputes

1. The prior art teaches all of the limitations of the examined claims

2. Supplemental briefing regarding claim 22

104 Patent: Disputes

1. The prior art teaches all of the limitations of the examined claims
   a. **“Prepending” (Claim 1)**
   b. “Sending . . . a Response to the Command” (Claims 1 and 12)
   c. “Prior to Receiving . . . an Acknowledgement” (Claims 1 and 12)
      i. No need to modify Connery
   d. “wherein receiving . . . a command to transmit data includes receiving . . . a pointer to the command” (Claim 9)

“Prepending” limitation

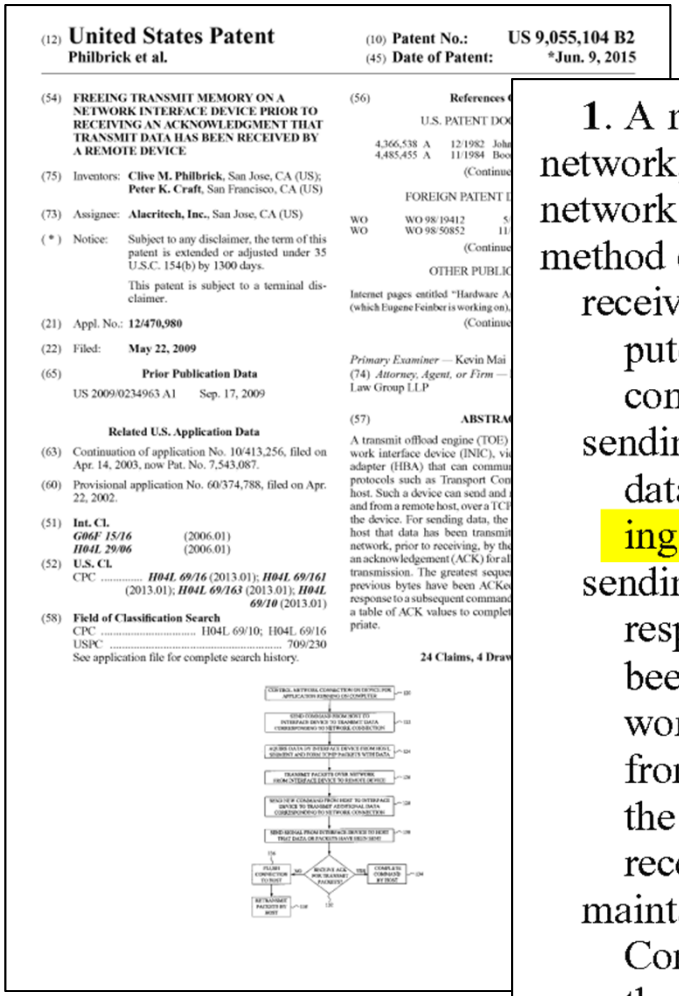

(12) **United States Patent**
Philbrick et al.

(10) **Patent No.: US 9,055,104 B2**
(45) **Date of Patent: *Jun. 9, 2015**

(54) **FREEING TRANSMIT MEMORY ON A NETWORK INTERFACE DEVICE PRIOR TO RECEIVING AN ACKNOWLEDGMENT THAT TRANSMIT DATA HAS BEEN RECEIVED BY A REMOTE DEVICE**

(75) Inventors: **Clive M. Philbrick**, San Jose, CA (US); **Peter K. Craft**, San Francisco, CA (US)

(73) Assignee: **Alacritech, Inc.**, San Jose, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 1300 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **12/470,980**

(22) Filed: **May 22, 2009**

(65) **Prior Publication Data**

US 2009/0234963 A1 Sep. 17, 2009

Related U.S. Application Data

(63) Continuation of application No. 10/413,256, filed on Apr. 14, 2003, now Pat. No. 7,543,087.

(60) Provisional application No. 60/374,788, filed on Apr. 22, 2002.

(51) **Int. Cl.**
G06F 15/16 (2006.01)
H04L 29/06 (2006.01)

(52) **U.S. Cl.**
CPC ***H04L 69/16*** (2013.01); ***H04L 69/161*** (2013.01); ***H04L 69/163*** (2013.01); ***H04L 69/10*** (2013.01)

(58) **Field of Classification Search**
CPC H04L 69/10; H04L 69/16
USPC 709/230
See application file for complete search history.

(56) **References**
U.S. PATENT DO
4,366,538 A 12/1982 John
4,485,455 A 11/1984 Boo
(Continue
FOREIGN PATENT
WO WO 98/19412 5
WO WO 98/50852 11
(Continue
OTHER PUBLIC
Internet pages entitled “Hardware A
(which Eugene Feinber is working on).
(Continue

Primary Examiner — Kevin Mai
(74) *Attorney, Agent, or Firm* —
Law Group LLP

(57) **ABSTRA**
A transmit offload engine (TOE)
work interface device (INIC), vi
adapter (HBA) that can commu
protocols such as Transport Con
host. Such a device can send and
and from a remote host, over a TCI
the device. For sending data, the
host that data has been transmit
network, prior to receiving, by th
an acknowledgement (ACK) for al
transmission. The greatest seque
previous bytes have been ACKe
response to a subsequent comman
a table of ACK values to complet
priate.

24 Claims, 4 Draw

1. A method for communication involving a computer, a network, and a network interface device of the computer, the network interface device being coupled to the network, the method comprising:

receiving, by the network interface device from the computer, a command to transmit application data from the computer to the network;

sending, by the network interface device to the network, data corresponding to the command, including prepending a transport layer header to at least some of the data;

sending, by the network interface device to the computer, a response to the command indicating that the data has been sent from the network interface device to the network, prior to receiving, by the network interface device from the network, an acknowledgement (ACK) that all the data corresponding to the command has been received; and

maintaining, by the network interface device, a Transport Control Protocol (TCP) connection that the command, the data and the ACK correspond to.

Ex. 1001.012 (104 Patent) at Claim 1.

Connery: The network interface device automatically segments data

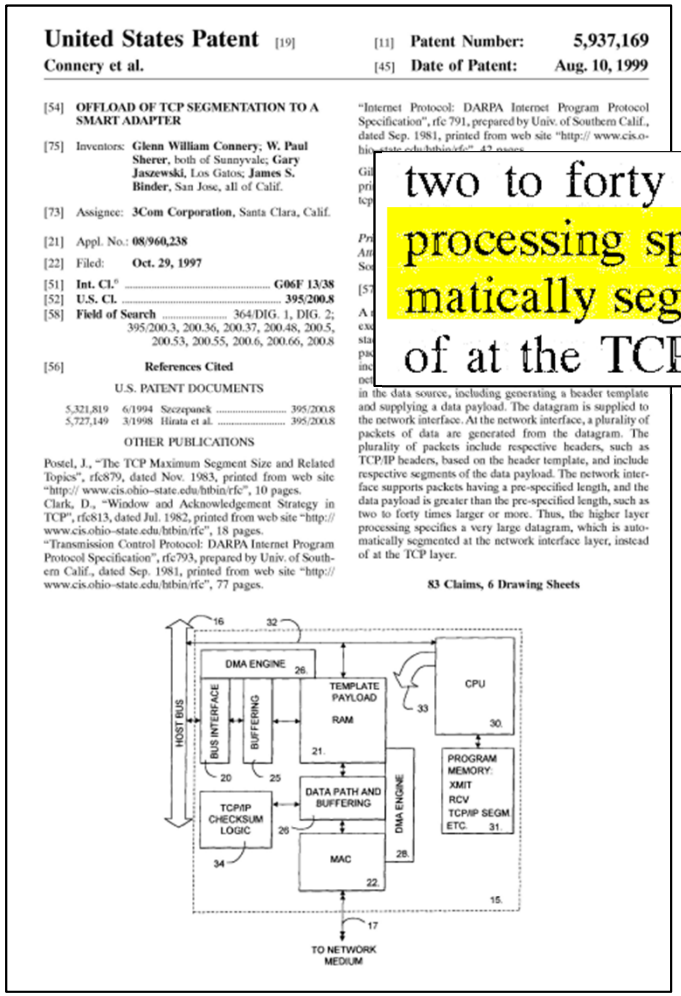

United States Patent [19]
Connery et al.

[11] **Patent Number: 5,937,169**
[45] **Date of Patent: Aug. 10, 1999**

[54] **OFFLOAD OF TCP SEGMENTATION TO A SMART ADAPTER**

[75] Inventors: **Glenn William Connery**, **W. Paul Sherer**, both of Sunnyvale; **Gary Jaszewski**, Los Gatos; **James S. Binder**, San Jose, all of Calif.

[73] Assignee: **3Com Corporation**, Santa Clara, Calif.

[21] Appl. No.: **08/960,238**

[22] Filed: **Oct. 29, 1997**

[51] **Int. Cl.**[6] **G06F 13/38**
[52] **U.S. Cl.** **395/200.8**
[58] **Field of Search** ..... 364/DIG. 1, DIG. 2; 395/200.3, 200.36, 200.37, 200.48, 200.5, 200.53, 200.55, 200.6, 200.66, 200.8

[56] **References Cited**

U.S. PATENT DOCUMENTS

5,321,819 6/1994 Szczepanek 395/200.8
5,727,149 3/1998 Hirata et al. 395/200.8

OTHER PUBLICATIONS

Postel, J., "The TCP Maximum Segment Size and Related Topics", rfc879, dated Nov. 1983, printed from web site "http:// www.cis.ohio-state.edu/htbin/rfc", 10 pages.
Clark, D., "Window and Acknowledgement Strategy in TCP", rfc813, dated Jul. 1982, printed from web site "http:// www.cis.ohio-state.edu/htbin/rfc", 18 pages.
"Transmission Control Protocol: DARPA Internet Program Protocol Specification", rfc793, prepared by Univ. of Southern Calif., dated Sep. 1981, printed from web site "http:// www.cis.ohio-state.edu/htbin/rfc", 77 pages.
"Internet Protocol: DARPA Internet Program Protocol Specification", rfc 791, prepared by Univ. of Southern Calif., dated Sep. 1981, printed from web site "http:// www.cis.o-

in the data source, including generating a header template and supplying a data payload. The datagram is supplied to the network interface. At the network interface, a plurality of packets of data are generated from the datagram. The plurality of packets include respective headers, such as TCP/IP headers, based on the header template, and include respective segments of the data payload. The network interface supports packets having a pre-specified length, and the data payload is greater than the pre-specified length, such as two to forty times larger or more. Thus, the higher layer processing specifies a very large datagram, which is automatically segmented at the network interface layer, instead of at the TCP layer.

83 Claims, 6 Drawing Sheets

two to forty times larger or more. Thus, the higher layer processing specifies a very large datagram, which is automatically segmented at the network interface layer, instead of at the TCP layer.

Ex. 1043.001 (Connery) at Abstract;
Ex. 1003.075 (Horst Decl.) at A-15;
Paper 1 (104 Petition) at 56-57.

Connery: “Prepending” the header

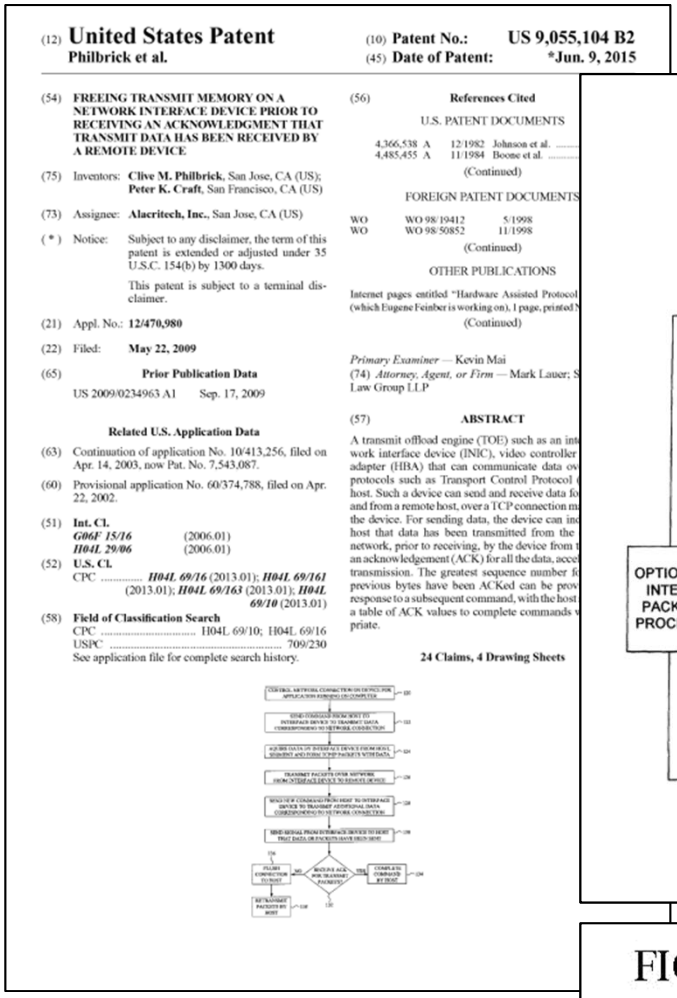

(12) **United States Patent**
Philbrick et al.

(10) **Patent No.: US 9,055,104 B2**
(45) **Date of Patent: *Jun. 9, 2015**

(54) **FREEING TRANSMIT MEMORY ON A NETWORK INTERFACE DEVICE PRIOR TO RECEIVING AN ACKNOWLEDGMENT THAT TRANSMIT DATA HAS BEEN RECEIVED BY A REMOTE DEVICE**

(75) Inventors: **Clive M. Philbrick**, San Jose, CA (US); **Peter K. Craft**, San Francisco, CA (US)

(73) Assignee: **Alacritech, Inc.**, San Jose, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 1300 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **12/470,980**

(22) Filed: **May 22, 2009**

(65) **Prior Publication Data**

US 2009/0234963 A1 Sep. 17, 2009

Related U.S. Application Data

(63) Continuation of application No. 10/413,256, filed on Apr. 14, 2003, now Pat. No. 7,543,087.

(60) Provisional application No. 60/374,788, filed on Apr. 22, 2002.

(51) **Int. Cl.**
G06F 15/16 (2006.01)
H04L 29/06 (2006.01)

(52) **U.S. Cl.**
CPC ***H04L 69/16*** (2013.01); ***H04L 69/161*** (2013.01); ***H04L 69/163*** (2013.01); ***H04L 69/10*** (2013.01)

(58) **Field of Classification Search**
CPC H04L 69/10; H04L 69/16
USPC 709/230
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,366,538 A 12/1982 Johnson et al. .....
4,485,455 A 11/1984 Boone et al. .....
(Continued)

FOREIGN PATENT DOCUMENTS

WO WO 98/19412 5/1998
WO WO 98/50852 11/1998
(Continued)

OTHER PUBLICATIONS

Internet pages entitled “Hardware Assisted Protocol (which Eugene Feinber is working on), 1 page, printed
(Continued)

Primary Examiner — Kevin Mai
(74) *Attorney, Agent, or Firm* — Mark Lauer; S
Law Group LLP

(57) **ABSTRACT**

A transmit offload engine (TOE) such as an int
work interface device (INIC), video controller
adapter (HBA) that can communicate data ov
protocols such as Transport Control Protocol
host. Such a device can send and receive data fo
and from a remote host, over a TCP connection m
the device. For sending data, the device can inc
host that data has been transmitted from the
network, prior to receiving, by the device from t
an acknowledgement (ACK) for all the data, accel
transmission. The greatest sequence number fo
previous bytes have been ACKed can be prov
response to a subsequent command, with the host
a table of ACK values to complete commands
priate.

24 Claims, 4 Drawing Sheets

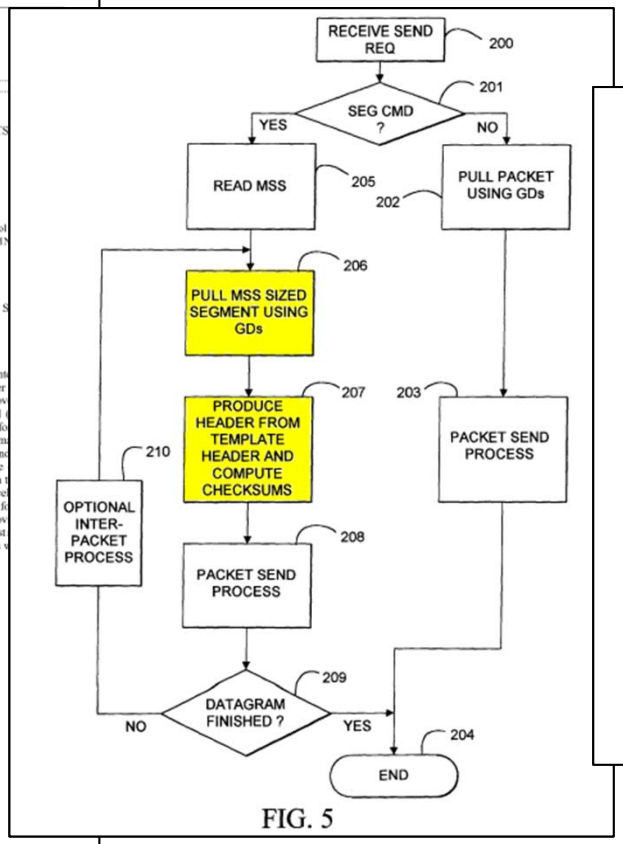

FIG. 5

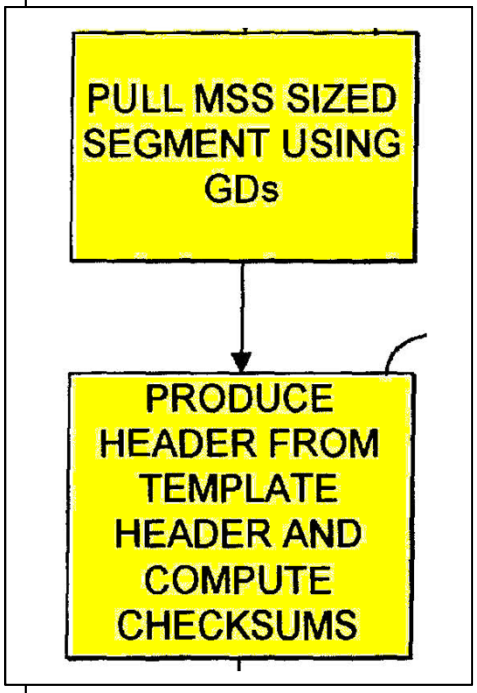

FIGS. **5–7** illustrate the processing executed in the smart network interface card according to the present invention.

Ex. 1043.014 (Connery) at 13:15-16, Fig. 5;
Ex. 1003.071-.074 (Horst Decl.) at A-12–A-13;
Paper 1 (104 Petition) at 53-54.

Dr. Horst: Both prepending header and appending data were obvious

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

INTEL CORPORATION

Finally, with respect to the claim term "*prepending*," it would have been obvious to a POSA that (1) the TCP header in Connery could have been prepended to the front of the data (as required by this claim limitation), or (2) the data could have been appended to the back of the header. Both are obvious. Specifically, given

Title: FREEING TRANSMIT MEMORY ON A NETWORK INTERFACE DEVICE PRIOR TO RECEIVING AN ACKNOWLEDGEMENT THAT TRANSMIT DATA HAS BEEN RECEIVED BY A REMOTE DEVICE

Declaration of Robert Horst, Ph.D. in Support of Petition for *Inter Partes* Review of U.S. Patent No. 9,055,104

Ex. 1003.074 (Horst Decl.) at A-14;
Paper 1 (104 Petition) at 55-56;
Paper 39 (104 Reply) at 7-8.

104 Patent: Disputes

1. The prior art teaches all of the limitations of the examined claims
   a. "Prepending" (Claim 1)
   b. <u>"Sending . . . a Response to the Command" (Claims 1 and 12)</u>
   c. "Prior to Receiving . . . an Acknowledgement" (Claims 1 and 12)
      i. No need to modify Connery
   d. "wherein receiving . . . a command to transmit data includes receiving . . . a pointer to the command" (Claim 9)

“Sending…a response to the command indicating data has been sent” limitation

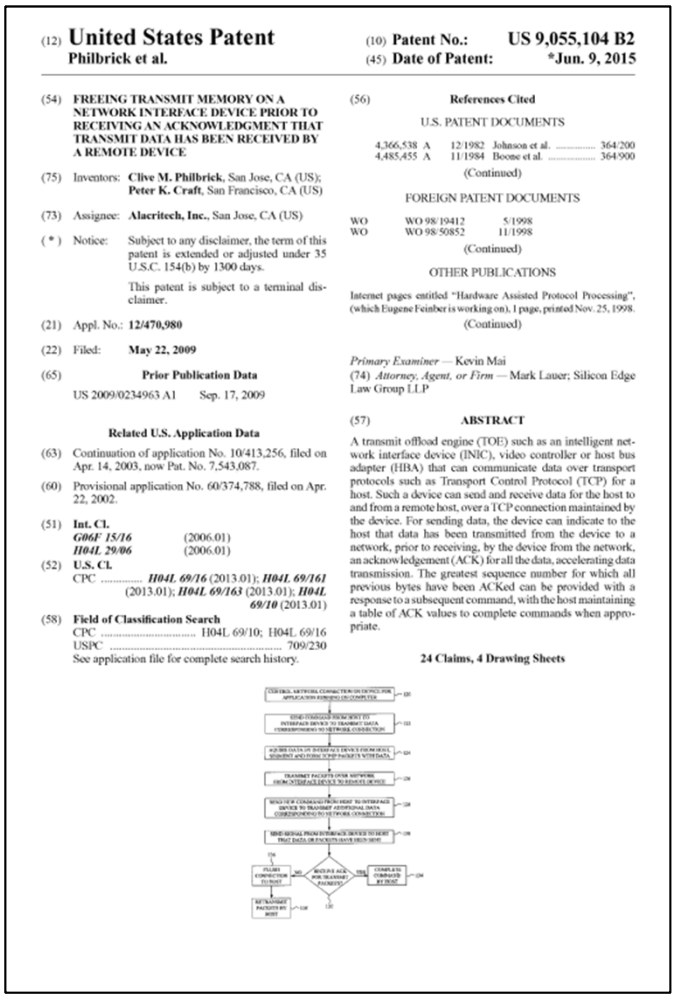

(12) **United States Patent**
Philbrick et al.

(10) **Patent No.: US 9,055,104 B2**
(45) **Date of Patent: *Jun. 9, 2015**

(54) **FREEING TRANSMIT MEMORY ON A NETWORK INTERFACE DEVICE PRIOR TO RECEIVING AN ACKNOWLEDGMENT THAT TRANSMIT DATA HAS BEEN RECEIVED BY A REMOTE DEVICE**

(75) Inventors: **Clive M. Philbrick**, San Jose, CA (US); **Peter K. Craft**, San Francisco, CA (US)

(73) Assignee: **Alacritech, Inc.**, San Jose, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 1300 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **12/470,980**

(22) Filed: **May 22, 2009**

(65) **Prior Publication Data**

US 2009/0234963 A1 Sep. 17, 2009

Related U.S. Application Data

(63) Continuation of application No. 10/413,256, filed on Apr. 14, 2003, now Pat. No. 7,543,087.

(60) Provisional application No. 60/374,788, filed on Apr. 22, 2002.

(51) **Int. Cl.**
G06F 15/16 (2006.01)
H04L 29/06 (2006.01)

(52) **U.S. Cl.**
CPC ***H04L 69/16*** (2013.01); ***H04L 69/161*** (2013.01); ***H04L 69/163*** (2013.01); ***H04L 69/10*** (2013.01)

(58) **Field of Classification Search**
CPC H04L 69/10; H04L 69/16
USPC 709/230
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,366,538 A 12/1982 Johnson et al. 364/200
4,485,455 A 11/1984 Boone et al. 364/900
(Continued)

FOREIGN PATENT DOCUMENTS

WO WO 98/19412 5/1998
WO WO 98/50852 11/1998
(Continued)

OTHER PUBLICATIONS

Internet pages entitled “Hardware Assisted Protocol Processing”, (which Eugene Feinber is working on), 1 page, printed Nov. 25, 1998.
(Continued)

Primary Examiner — Kevin Mai
(74) *Attorney, Agent, or Firm* — Mark Lauer; Silicon Edge Law Group LLP

(57) **ABSTRACT**

A transmit offload engine (TOE) such as an intelligent network interface device (INIC), video controller or host bus adapter (HBA) that can communicate data over transport protocols such as Transport Control Protocol (TCP) for a host. Such a device can send and receive data for the host to and from a remote host, over a TCP connection maintained by the device. For sending data, the device can indicate to the host that data has been transmitted from the device to a network, prior to receiving, by the device from the network, an acknowledgement (ACK) for all the data, accelerating data transmission. The greatest sequence number for which all previous bytes have been ACKed can be provided with a response to a subsequent command, with the host maintaining a table of ACK values to complete commands when appropriate.

24 Claims, 4 Drawing Sheets

1. A method for communication involving a computer, a network, and a network interface device of the computer, the network interface device being coupled to the network, the method comprising:

receiving, by the network interface device from the computer, a command to transmit application data from the computer to the network;

sending, by the network interface device to the network, data corresponding to the command, including prepending a transport layer header to at least some of the data;

sending, by the network interface device to the computer, a response to the command indicating that the data has been sent from the network interface device to the network, prior to receiving, by the network interface device from the network, an acknowledgement (ACK) that all the data corresponding to the command has been received; and

maintaining, by the network interface device, a Transport Control Protocol (TCP) connection that the command, the data and the ACK correspond to.

Ex. 1001.012 (104 Patent) at Claim 1.

*Aside from the addition of the underlined language, the highlighted language is identical for claim 12

Connery: A "large packet" is offloaded for processing by the smart adapter

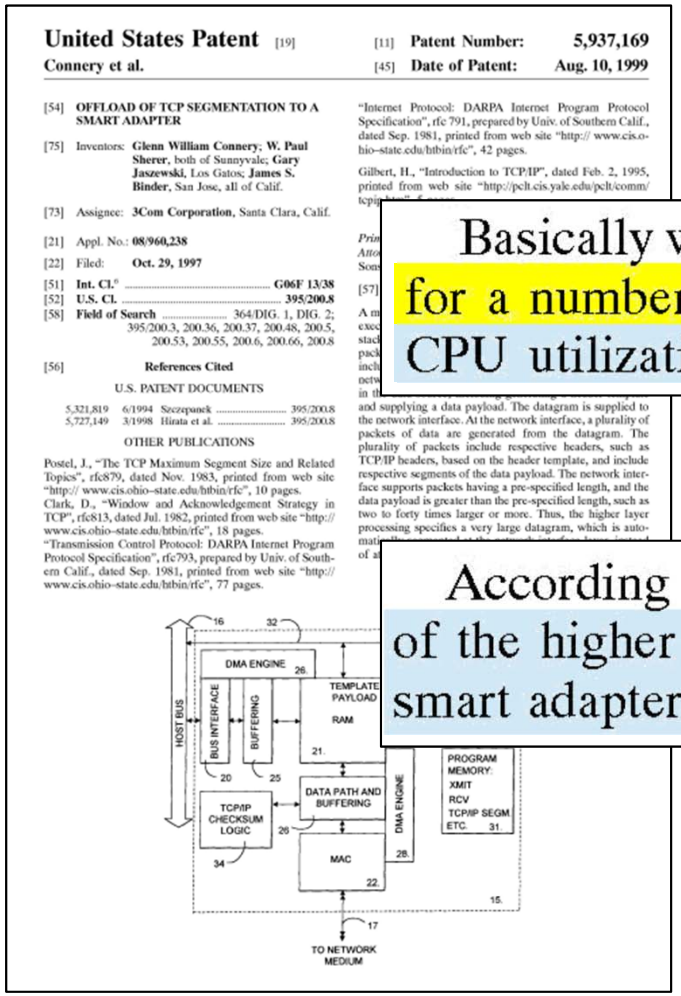

United States Patent [19]
Connery et al.

[11] **Patent Number: 5,937,169**
[45] **Date of Patent: Aug. 10, 1999**

[54] **OFFLOAD OF TCP SEGMENTATION TO A SMART ADAPTER**

[75] Inventors: **Glenn William Connery**, **W. Paul Sherer**, both of Sunnyvale; **Gary Jaszewski**, Los Gatos; **James S. Binder**, San Jose, all of Calif.

[73] Assignee: **3Com Corporation**, Santa Clara, Calif.

[21] Appl. No.: **08/960,238**

[22] Filed: **Oct. 29, 1997**

[51] **Int. Cl.**[6] **G06F 13/38**
[52] **U.S. Cl.** **395/200.8**
[58] **Field of Search** ..... 364/DIG. 1, DIG. 2; 395/200.3, 200.36, 200.37, 200.48, 200.5, 200.53, 200.55, 200.6, 200.66, 200.8

[56] **References Cited**

U.S. PATENT DOCUMENTS

5,321,819 6/1994 Szczepanek 395/200.8
5,727,149 3/1998 Hirata et al. 395/200.8

OTHER PUBLICATIONS

Postel, J., "The TCP Maximum Segment Size and Related Topics", rfc879, dated Nov. 1983, printed from web site "http:// www.cis.ohio-state.edu/htbin/rfc", 10 pages.
Clark, D., "Window and Acknowledgement Strategy in TCP", rfc813, dated Jul. 1982, printed from web site "http://www.cis.ohio-state.edu/htbin/rfc", 18 pages.
"Transmission Control Protocol: DARPA Internet Program Protocol Specification", rfc793, prepared by Univ. of Southern Calif., dated Sep. 1981, printed from web site "http://www.cis.ohio-state.edu/htbin/rfc", 77 pages.
"Internet Protocol: DARPA Internet Program Protocol Specification", rfc 791, prepared by Univ. of Southern Calif., dated Sep. 1981, printed from web site "http:// www.cis.ohio-state.edu/htbin/rfc", 42 pages.
Gilbert, H., "Introduction to TCP/IP", dated Feb. 2, 1995, printed from web site "http://pclt.cis.yale.edu/pclt/comm/

and supplying a data payload. The datagram is supplied to the network interface. At the network interface, a plurality of packets of data are generated from the datagram. The plurality of packets include respective headers, such as TCP/IP headers, based on the header template, and include respective segments of the data payload. The network interface supports packets having a pre-specified length, and the data payload is greater than the pre-specified length, such as two to forty times larger or more. Thus, the higher layer processing specifies a very large datagram, which is auto-

Basically we have substituted a single "large packet" send for a number of smaller sends. The goal is to reduce host CPU utilization and improve performance and scalability.

Ex. 1043.011 (Connery) at 7:47-49;
Ex. 1003.075 (Horst Decl.) at A-15;
Paper 1 (104 Petition) at 56.

According to the present invention, a significant portion of the higher layer transmit processing is offloaded onto a smart adapter. The present invention accomplishes this off-

Ex. 1043.008 (Connery) at 2:40-42;
Ex. 1003.074 (Horst Decl.) at A-15-16;
Paper 1 (104 Petition) at 57.

Connery: Segmentation offload reduces interrupts to one per “large packet”

United States Patent [19]
Connery et al.

[11] **Patent Number: 5,937,169**
[45] **Date of Patent: Aug. 10, 1999**

[54] **OFFLOAD OF TCP SEGMENTATION TO A SMART ADAPTER**

[75] Inventors: **Glenn William Connery**; **W. Paul Sherer**, both of Sunnyvale; **Gary Jaszewski**, Los Gatos; **James S. Binder**, San Jose, all of Calif.

[73] Assignee: **3Com Corporation**, Santa Clara, Calif.

“Internet Protocol: DARPA Internet Program Protocol Specification”, rfc 791, prepared by Univ. of Southern Calif., dated Sep. 1981, printed from web site “http:// www.cis.ohio-state.edu/htbin/rfc”, 42 pages.

Gilbert, H., “Introduction to TCP/IP”, dated Feb. 2, 1995, printed from web site “http://pclt.cis.yale.edu/pclt/comm/tcpip.htm”, 5 pages.

physicalization penalty for each packet. Also the number of interrupts on the host CPU is reduced to one per “large packet” rather than one per packet or one per some number of packets (algorithmic). Offloading TCP segmentation to the adapter will help reduce all of this quite significantly.

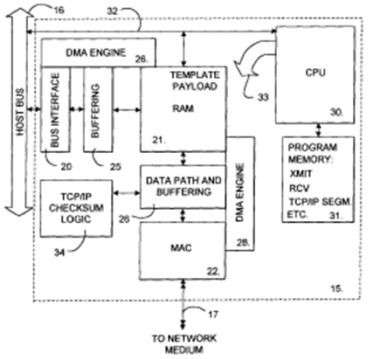

Ex. 1043.011 (Connery) at 7:60-64;
Ex. 1003.075-.076 (Horst Decl.) at A-15 – A-16;
Paper 1 (104 Petition) at 57.

Connery: Interrupts that are avoided can include transmit completion interrupts

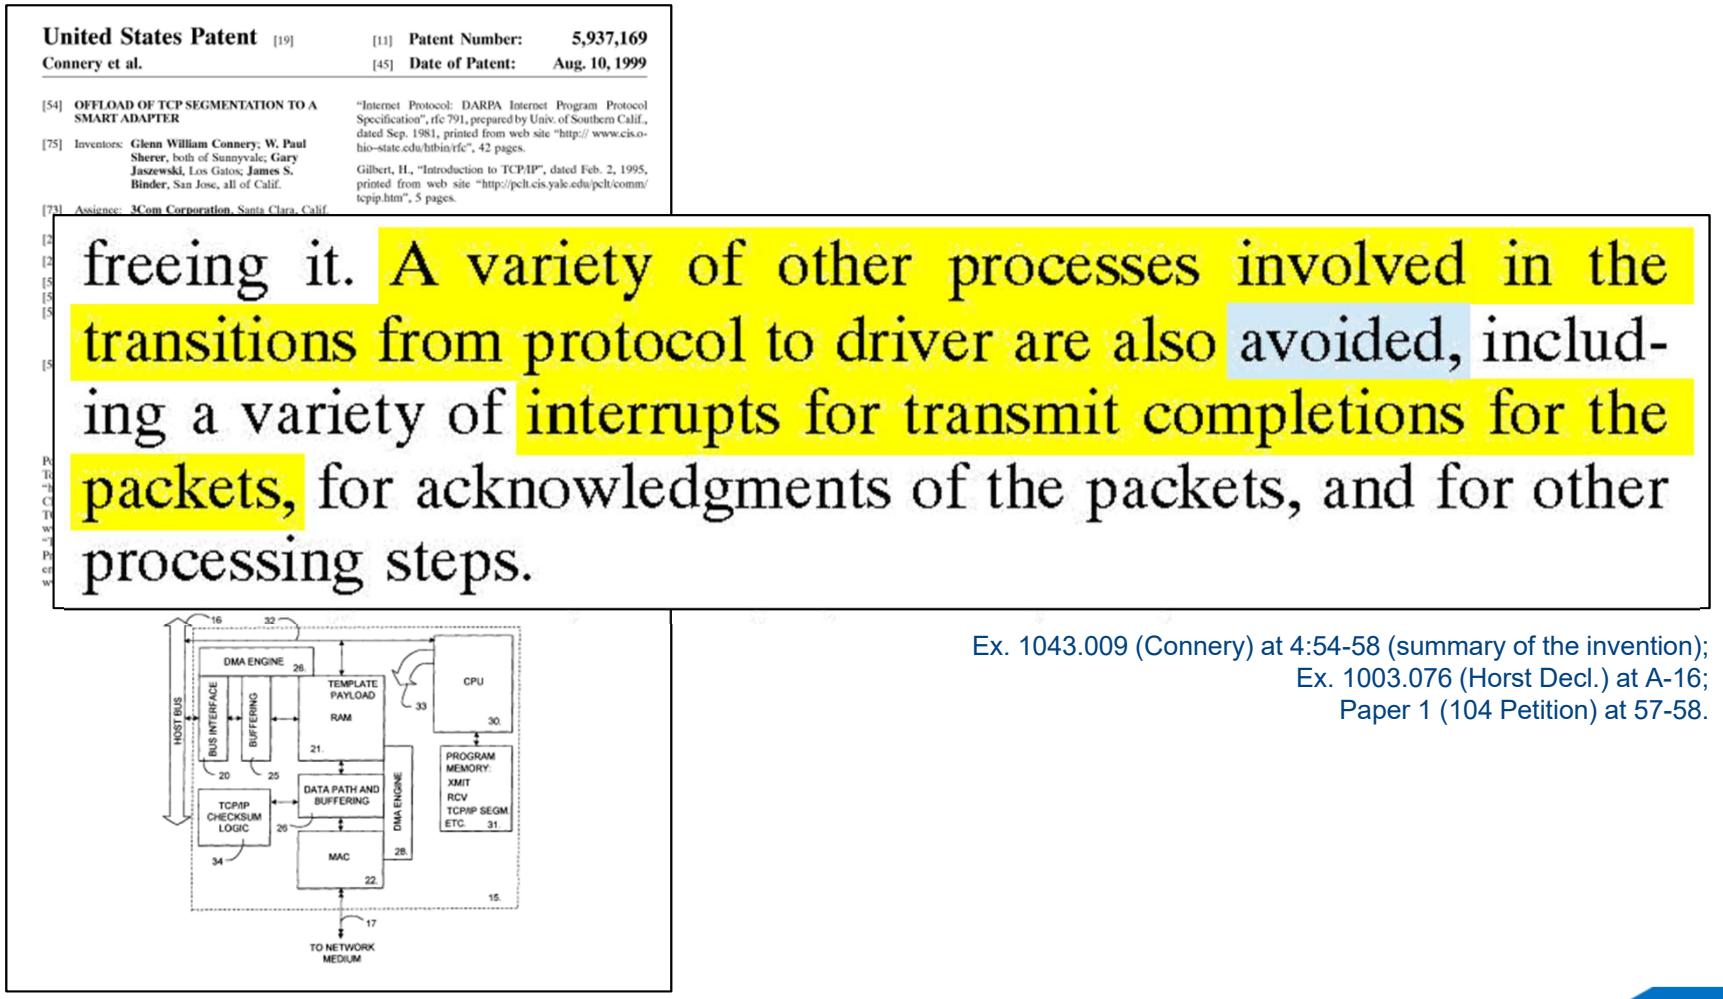

United States Patent [19]
Connery et al.

[11] Patent Number: 5,937,169
[45] Date of Patent: Aug. 10, 1999

[54] OFFLOAD OF TCP SEGMENTATION TO A SMART ADAPTER

[75] Inventors: Glenn William Connery; W. Paul Sherer, both of Sunnyvale; Gary Jaszewski, Los Gatos; James S. Binder, San Jose, all of Calif.

[73] Assignee: 3Com Corporation, Santa Clara, Calif.

"Internet Protocol: DARPA Internet Program Protocol Specification", rfc 791, prepared by Univ. of Southern Calif., dated Sep. 1981, printed from web site "http:// www.cis.ohio-state.edu/htbin/rfc", 42 pages.

Gilbert, H., "Introduction to TCP/IP", dated Feb. 2, 1995, printed from web site "http://pclt.cis.yale.edu/pclt/comm/tcpip.htm", 5 pages.

freeing it. A variety of other processes involved in the transitions from protocol to driver are also avoided, including a variety of interrupts for transmit completions for the packets, for acknowledgments of the packets, and for other processing steps.

Ex. 1043.009 (Connery) at 4:54-58 (summary of the invention);
Ex. 1003.076 (Horst Decl.) at A-16;
Paper 1 (104 Petition) at 57-58.

Dr. Horst: Connery's single interrupt is an "interrupt for transmit completion"

Accordingly, given Connery's goal of reducing host CPU utilization by substituting a single large packet send from the host for a number of smaller sends, and given Connery's disclosure that in its system "the number of interrupts on the host CPU is reduced to one per 'large packet' rather than one per packet or one per some number of packets" (all as explained above), a POSA would have understood that the "number of interrupts on the host CPU" that are "reduced to one per 'large packet'" is a send completion interrupt. This is the interrupt sent by the network interface device to the host computer indicating that data has been sent from the network interface device to the network. An interrupt like this is needed because otherwise the host would not know when the hardware is free to accept the next large packet to be sent. The transmit completion interrupt is sent in response to the host's original command to transmit data.

Title: F
DEVI
TRANSMIT DATA HAS BEEN RECEIVED BY A REMOTE DEVICE

Declaration of Robert Horst, Ph.D. in Support of Petition for *Inter Partes* Review of U.S. Patent No. 9,055,104

Ex. 1003.076-.077 (Horst Decl.) at A-16 – A-17;
Paper 1 (104 Petition) at 58-59.

Dr. Horst: “Transmit completion” means data was transmitted to the network

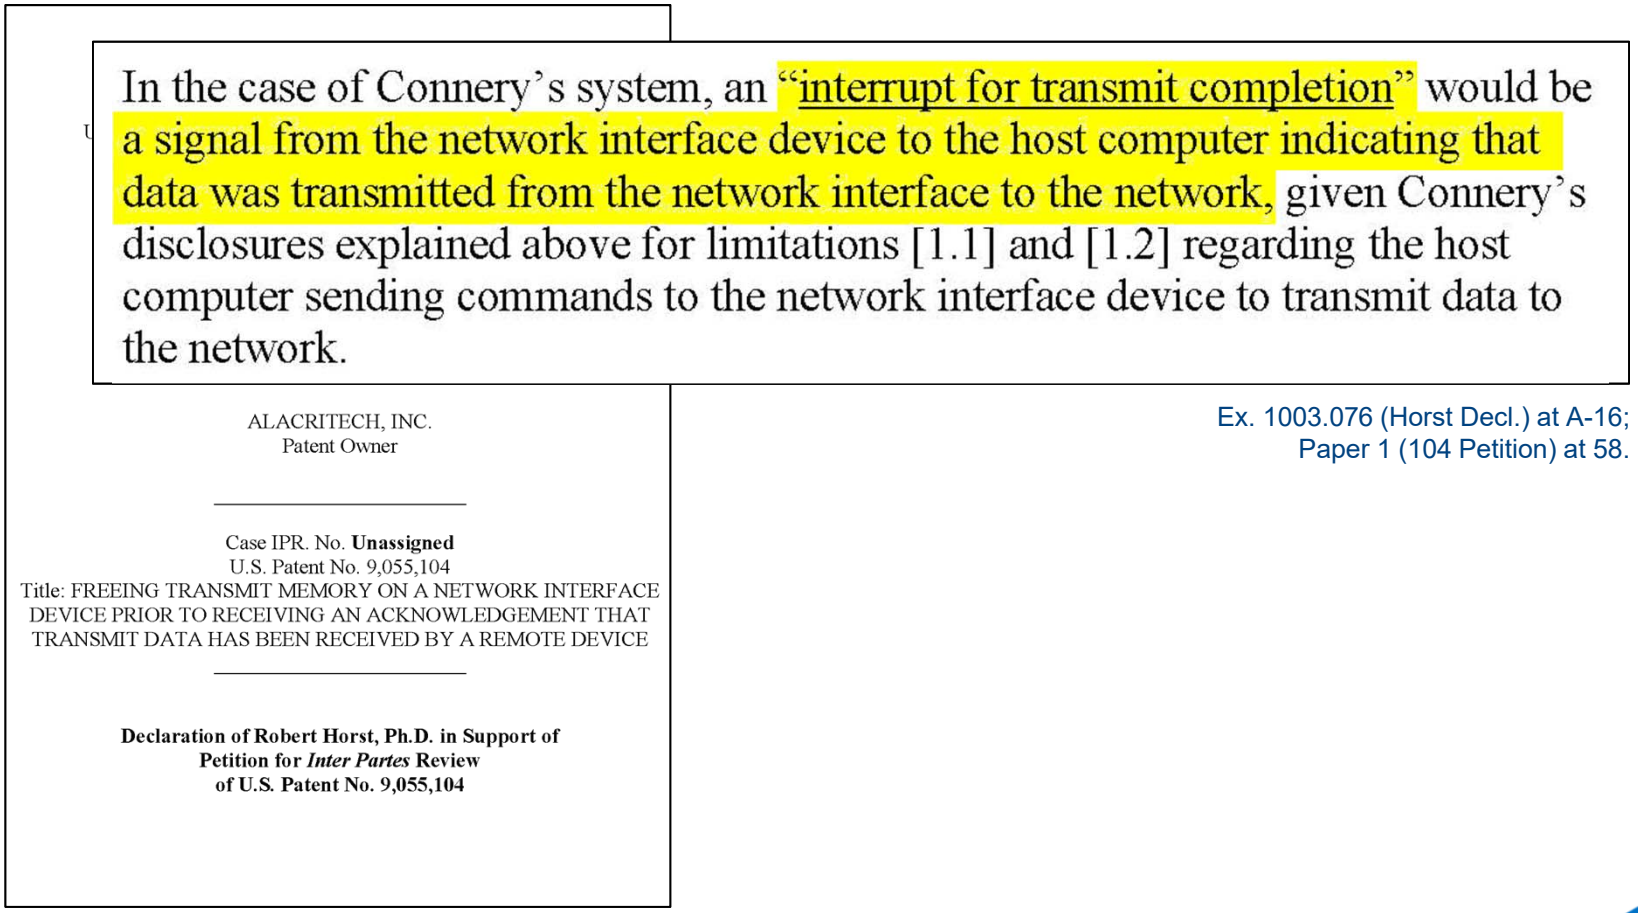

In the case of Connery’s system, an “interrupt for transmit completion” would be a signal from the network interface device to the host computer indicating that data was transmitted from the network interface to the network, given Connery’s disclosures explained above for limitations [1.1] and [1.2] regarding the host computer sending commands to the network interface device to transmit data to the network.

ALACRITECH, INC.
Patent Owner

Case IPR. No. **Unassigned**
U.S. Patent No. 9,055,104
Title: FREEING TRANSMIT MEMORY ON A NETWORK INTERFACE DEVICE PRIOR TO RECEIVING AN ACKNOWLEDGEMENT THAT TRANSMIT DATA HAS BEEN RECEIVED BY A REMOTE DEVICE

Declaration of Robert Horst, Ph.D. in Support of Petition for *Inter Partes* Review of U.S. Patent No. 9,055,104

Ex. 1003.076 (Horst Decl.) at A-16;
Paper 1 (104 Petition) at 58.

Dr. Horst: Alternatively, an “interrupt for transmit completion” was obvious choice

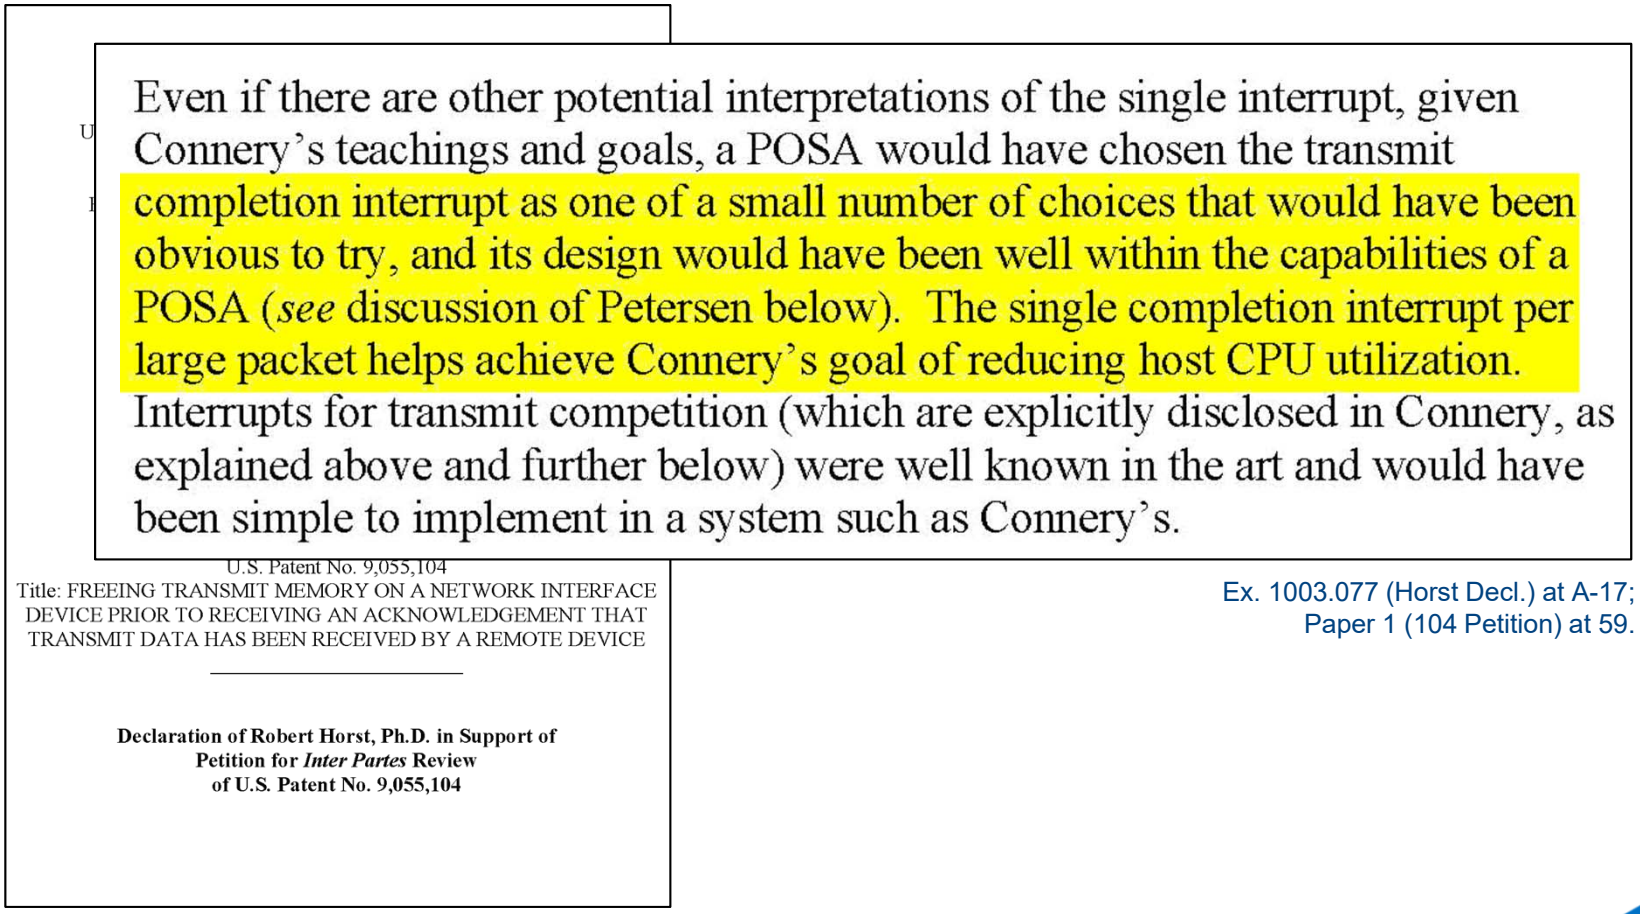

Even if there are other potential interpretations of the single interrupt, given Connery's teachings and goals, a POSA would have chosen the transmit completion interrupt as one of a small number of choices that would have been obvious to try, and its design would have been well within the capabilities of a POSA (*see* discussion of Petersen below). The single completion interrupt per large packet helps achieve Connery's goal of reducing host CPU utilization. Interrupts for transmit competition (which are explicitly disclosed in Connery, as explained above and further below) were well known in the art and would have been simple to implement in a system such as Connery's.

U.S. Patent No. 9,055,104
Title: FREEING TRANSMIT MEMORY ON A NETWORK INTERFACE DEVICE PRIOR TO RECEIVING AN ACKNOWLEDGEMENT THAT TRANSMIT DATA HAS BEEN RECEIVED BY A REMOTE DEVICE

Declaration of Robert Horst, Ph.D. in Support of Petition for *Inter Partes* Review of U.S. Patent No. 9,055,104

Ex. 1003.077 (Horst Decl.) at A-17;
Paper 1 (104 Petition) at 59.

Dr. Horst: Peterson corroborates that "transmit complete interrupt" was known

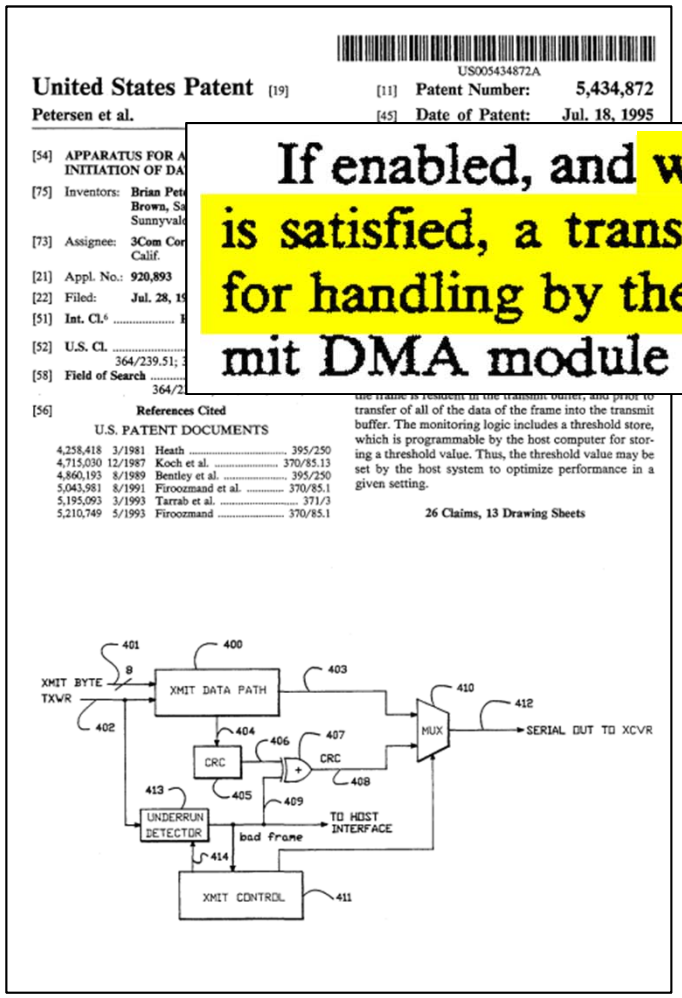

US005434872A

United States Patent [19]

Petersen et al.

[11] **Patent Number: 5,434,872**

[45] **Date of Patent: Jul. 18, 1995**

[54] APPARATUS FOR A
INITIATION OF DA

[75] Inventors: Brian Pet
Brown, Sa
Sunnyval

[73] Assignee: 3Com Cor
Calif.

[21] Appl. No.: 920,893

[22] Filed: Jul. 28, 19

[51] Int. Cl.[6] .....

[52] U.S. Cl. .....
364/239.51;

[58] Field of Search .....
364/2

[56] **References Cited**

U.S. PATENT DOCUMENTS

4,258,418	3/1981	Heath	395/250
4,715,030	12/1987	Koch et al.	370/85.13
4,860,193	8/1989	Bentley et al.	395/250
5,043,981	8/1991	Firoozmand et al.	370/85.1
5,195,093	3/1993	Tarrab et al.	371/3
5,210,749	5/1993	Firoozmand	370/85.1

the frame is resident in the transmit buffer, and prior to transfer of all of the data of the frame into the transmit buffer. The monitoring logic includes a threshold store, which is programmable by the host computer for storing a threshold value. Thus, the threshold value may be set by the host system to optimize performance in a given setting.

26 Claims, 13 Drawing Sheets

If enabled, and **when the transmit complete condition is satisfied, a transmit complete interrupt is generated for handling by the interrupt controller 60** in the transmit DMA module 67.

Ex. 1044.019 (Peterson) at 9:33-36;
Ex. 1003.077-.078 (Horst Decl.) at A-17 – A-18;
Paper 1 (104 Petition) at 59-60.

104 Patent: Disputes

1. The prior art teaches all of the limitations of the examined claims
   a. "Prepending" (Claim 1)
   b. "Sending . . . a Response to the Command" (Claims 1 and 12)
   c. <u>"Prior to Receiving . . . an Acknowledgement" (Claims 1 and 12)</u>
      i. <u>No need to modify Connery</u>
   d. "wherein receiving . . . a command to transmit data includes receiving . . . a pointer to the command" (Claim 9)

"…prior to receiving an acknowledgement" limitations

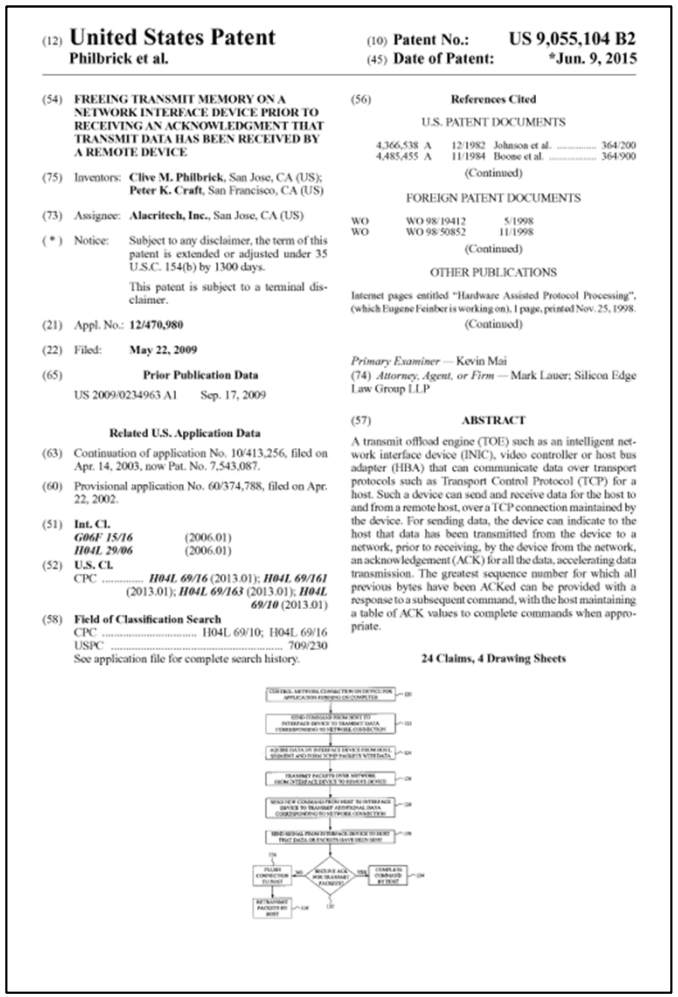

(12) **United States Patent**
Philbrick et al.

(10) **Patent No.: US 9,055,104 B2**
(45) **Date of Patent: *Jun. 9, 2015**

(54) **FREEING TRANSMIT MEMORY ON A NETWORK INTERFACE DEVICE PRIOR TO RECEIVING AN ACKNOWLEDGMENT THAT TRANSMIT DATA HAS BEEN RECEIVED BY A REMOTE DEVICE**

(75) Inventors: **Clive M. Philbrick**, San Jose, CA (US); **Peter K. Craft**, San Francisco, CA (US)

(73) Assignee: **Alacritech, Inc.**, San Jose, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 1300 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **12/470,980**

(22) Filed: **May 22, 2009**

(65) **Prior Publication Data**

US 2009/0234963 A1 Sep. 17, 2009

Related U.S. Application Data

(63) Continuation of application No. 10/413,256, filed on Apr. 14, 2003, now Pat. No. 7,543,087.

(60) Provisional application No. 60/374,788, filed on Apr. 22, 2002.

(51) **Int. Cl.**
G06F 15/16 (2006.01)
H04L 29/06 (2006.01)

(52) **U.S. Cl.**
CPC ***H04L 69/16*** (2013.01); ***H04L 69/161*** (2013.01); ***H04L 69/163*** (2013.01); ***H04L 69/10*** (2013.01)

(58) **Field of Classification Search**
CPC H04L 69/10; H04L 69/16
USPC 709/230
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,366,538 A 12/1982 Johnson et al. 364/200
4,485,455 A 11/1984 Boone et al. 364/900
(Continued)

FOREIGN PATENT DOCUMENTS

WO WO 98/19412 5/1998
WO WO 98/50852 11/1998
(Continued)

OTHER PUBLICATIONS

Internet pages entitled "Hardware Assisted Protocol Processing", (which Eugene Feinber is working on), 1 page, printed Nov. 25, 1998.
(Continued)

Primary Examiner — Kevin Mai
(74) *Attorney, Agent, or Firm* — Mark Lauer; Silicon Edge Law Group LLP

(57) **ABSTRACT**

A transmit offload engine (TOE) such as an intelligent network interface device (INIC), video controller or host bus adapter (HBA) that can communicate data over transport protocols such as Transport Control Protocol (TCP) for a host. Such a device can send and receive data for the host to and from a remote host, over a TCP connection maintained by the device. For sending data, the device can indicate to the host that data has been transmitted from the device to a network, prior to receiving, by the device from the network, an acknowledgement (ACK) for all the data, accelerating data transmission. The greatest sequence number for which all previous bytes have been ACKed can be provided with a response to a subsequent command, with the host maintaining a table of ACK values to complete commands when appropriate.

24 Claims, 4 Drawing Sheets

1. A method for communication involving a computer, a network, and a network interface device of the computer, the network interface device being coupled to the network, the method comprising:

receiving, by the network interface device from the computer, a command to transmit application data from the computer to the network;

sending, by the network interface device to the network, data corresponding to the command, including prepending a transport layer header to at least some of the data;

sending, by the network interface device to the computer, a response to the command indicating that the data has been sent from the network interface device to the network, prior to receiving, by the network interface device from the network, an acknowledgement (ACK) that all the data corresponding to the command has been received; and

maintaining, by the network interface device, a Transport Control Protocol (TCP) connection that the command, the data and the ACK correspond to.

Ex. 1001.012 (104 Patent) at Claim 1.

*Aside from the addition of the underlined language, the highlighted language is identical for claim 12

Unlike the transmit complete interrupt, an ACK is received from the destination

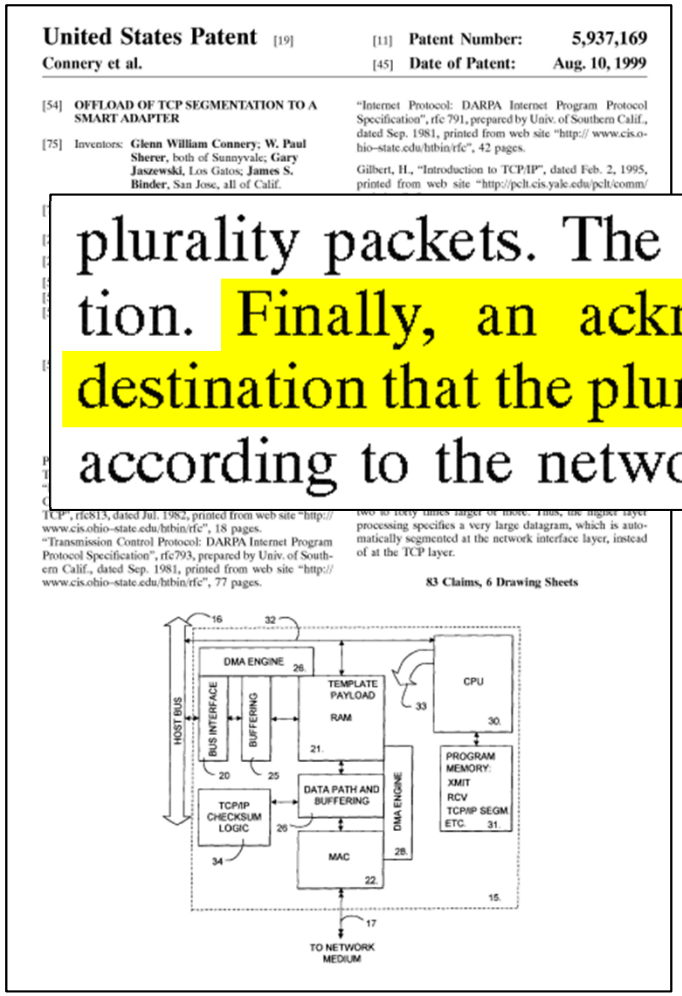

United States Patent [19]
Connery et al.

[11] Patent Number: 5,937,169
[45] Date of Patent: Aug. 10, 1999

[54] OFFLOAD OF TCP SEGMENTATION TO A SMART ADAPTER

[75] Inventors: Glenn William Connery; W. Paul Sherer, both of Sunnyvale; Gary Jaszewski, Los Gatos; James S. Binder, San Jose, all of Calif.

"Internet Protocol: DARPA Internet Program Protocol Specification", rfc 791, prepared by Univ. of Southern Calif., dated Sep. 1981, printed from web site "http:// www.cis.ohio-state.edu/htbin/rfc", 42 pages.

Gilbert, H., "Introduction to TCP/IP", dated Feb. 2, 1995, printed from web site "http://pclt.cis.yale.edu/pclt/comm/

TCP", rfc813, dated Jul. 1982, printed from web site "http:// www.cis.ohio-state.edu/htbin/rfc", 18 pages.
"Transmission Control Protocol: DARPA Internet Program Protocol Specification", rfc793, prepared by Univ. of Southern Calif., dated Sep. 1981, printed from web site "http:// www.cis.ohio-state.edu/htbin/rfc", 77 pages.

two to forty times larger or more. Thus, the higher layer processing specifies a very large datagram, which is automatically segmented at the network interface layer, instead of at the TCP layer.

83 Claims, 6 Drawing Sheets

plurality packets. The packets are then sent to the destination. Finally, an acknowledgment is received from the destination that the plurality of packets was successfully sent according to the network protocol.

Ex. 1043.009 (Connery) at 3:59-61 (Summary of the invention);
Ex. 1003.079-.080 (Horst Decl.) at A-19 – A-20;
Paper 1 (104 Petition) at 61-62.

Dr. Horst: Connery's transmit complete interrupt precedes receipt of ACKs due to latency

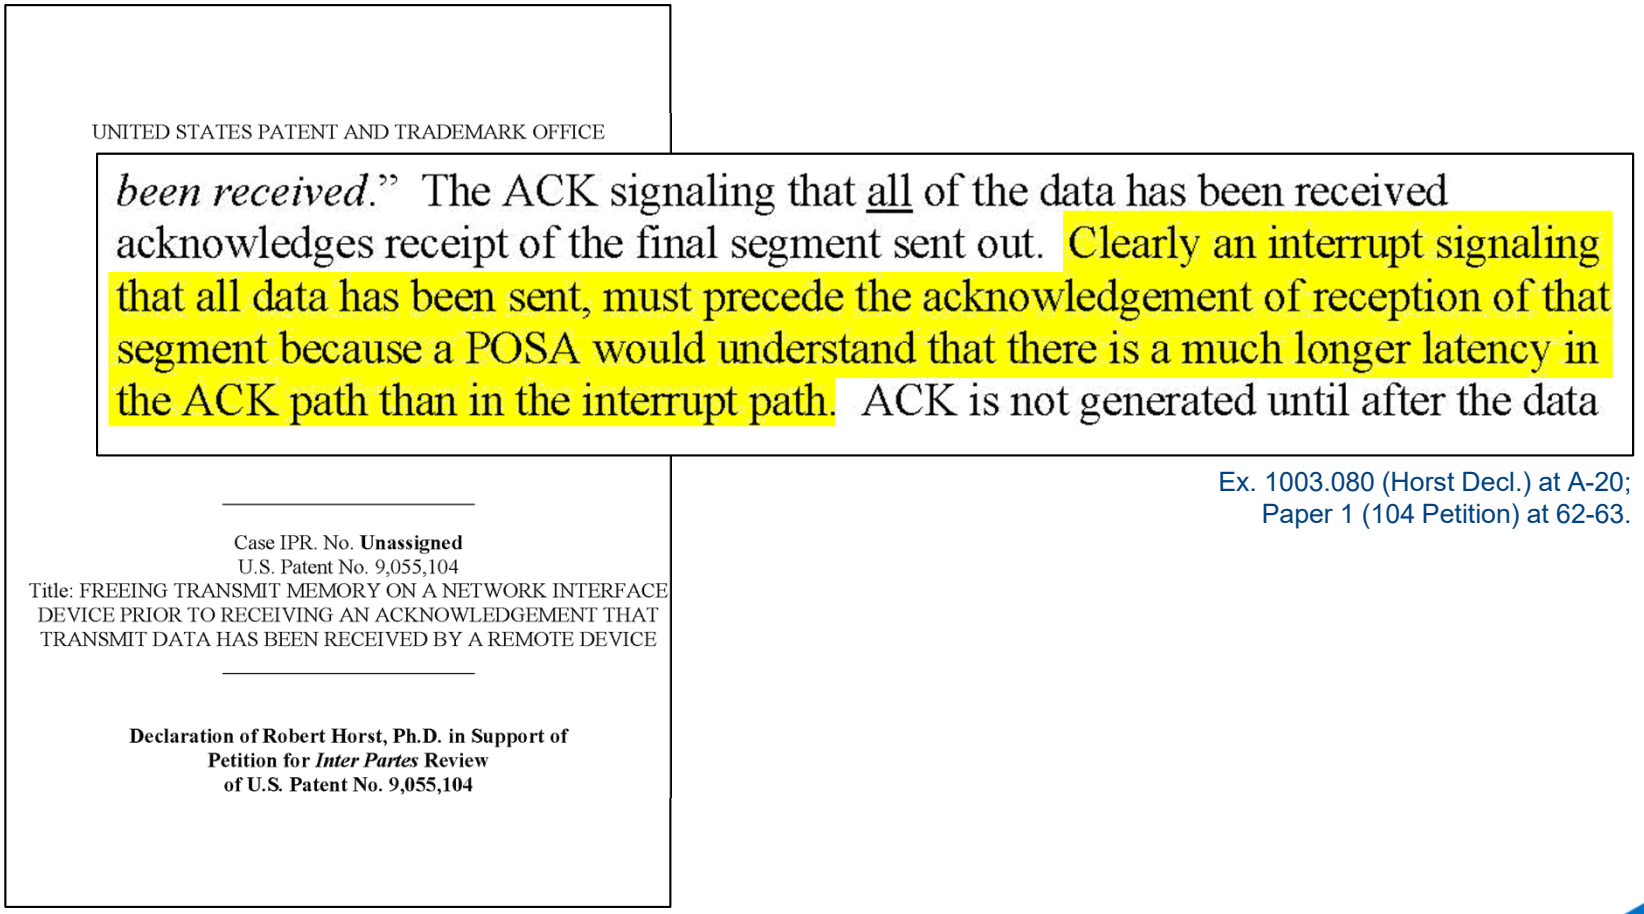

UNITED STATES PATENT AND TRADEMARK OFFICE

been received." The ACK signaling that all of the data has been received acknowledges receipt of the final segment sent out. Clearly an interrupt signaling that all data has been sent, must precede the acknowledgement of reception of that segment because a POSA would understand that there is a much longer latency in the ACK path than in the interrupt path. ACK is not generated until after the data

Case IPR. No. **Unassigned**
U.S. Patent No. 9,055,104
Title: FREEING TRANSMIT MEMORY ON A NETWORK INTERFACE DEVICE PRIOR TO RECEIVING AN ACKNOWLEDGEMENT THAT TRANSMIT DATA HAS BEEN RECEIVED BY A REMOTE DEVICE

Declaration of Robert Horst, Ph.D. in Support of Petition for *Inter Partes* Review of U.S. Patent No. 9,055,104

Ex. 1003.080 (Horst Decl.) at A-20;
Paper 1 (104 Petition) at 62-63.

Alacritech's expert, Dr. Min: Obvious that notification of data transmission occurs before ACK

Q. Is it because . . . the time that's required to give the host a notification that the transmission has been sent, that takes much less time than the round trip of the data and the ACK; is that right?

A. Yeah, that's right

Q. And you're saying that's obvious to a person of ordinary skill in the art?

A. Yes, of course.

Ex. 1077 (P. Min, March 21, 2017, Dep.) at 283:8-18 (objection omitted).
See also id. at 282:7-2;
Paper 1 (104 Petition) at 63, FN6;
Paper 39 (104 Reply) at 11-12.

No need to modify Connery's interrupts

- PO argues that there is no motivation to modify Connery's "interrupts on the host CPU to occur before the network interface receives an ACK that all the data has been received at the destination"

Paper 29 (Response) at 40-41.

  - However, there is no need to modify Connery, which teaches all claim limitations

Paper 39 (Reply) at 14.

Alternatively, a transmit complete interrupt would have been obvious to try

- A POSA would be motivated to use a transmit completion interrupt as one of a limited number of choices (e.g., Peterson discloses such interrupts) given Connery's goal of reducing CPU utilization
- A transmit complete interrupt has performance benefits

Ex. 1003.076-.078, .080-.081 (Horst Decl.) at A-16 - A-17, A-19 – A-20;
Paper 1 (Petition) at 58-59;
Paper 39 (Reply) at 14-15;
Ex. 1223.017 (Horst Reply Decl.) at 18;
Ex. 1077.282-.283 (Paul Min Dep., Mar. 21, 2017) at 282:7-25, 283:8-18.

PO argues there is no reasonable expectation of success if ACKs are not handled conventionally

actual application is beyond his or her skill."). As discussed above, the ACKs ensured reliability of transmission, so a modification removing that reliability mechanism from a network interface device would not have been intuitive or obvious to a POSITA. (Ex. 2026 at 108.) Further, there would not have been a reasonable expectation of success given the possible loss of reliability by not handling ACKs conventionally. (*Id.*) Since transmitting and processing ACKs was so deeply engrained in many standard protocols, such as TCP/IP, at the time, it would not have been obvious to a POSITA to ignore these messages. (*Id.*)

Paper 29 (Response) at 41.

Dr. Horst: ACKs *are* processed conventionally

- ACKs processed as required by TCP/IP

Ex. 1223.018-.019 (Horst Reply Decl.) at 17-18;
Paper 39 (104 Reply) at 15.

- Transmit complete interrupt is agnostic of ACKs, which are used for a different type of reliability:
  - Successful **receipt** of data at receiving system (ACK) vs successful **transmission** of data by the network interface (transmit complete)

Ex. 1223.018-.019 (Horst Reply Decl.) at 18-19;
Paper 39 (104 Reply) at 15.

104 Patent: Disputes

1. The prior art teaches all of the limitations of the examined claims
   a. "Prepending" (Claim 1)
   b. "Sending . . . a Response to the Command" (Claims 1 and 12)
   c. "Prior to Receiving . . . an Acknowledgement" (Claims 1 and 12)
      i. No need to modify Connery
   d. <u>"wherein receiving . . . a command to transmit data includes receiving . . . a pointer to the command" (Claim 9)</u>

“Receiving …a pointer to the command” limitation

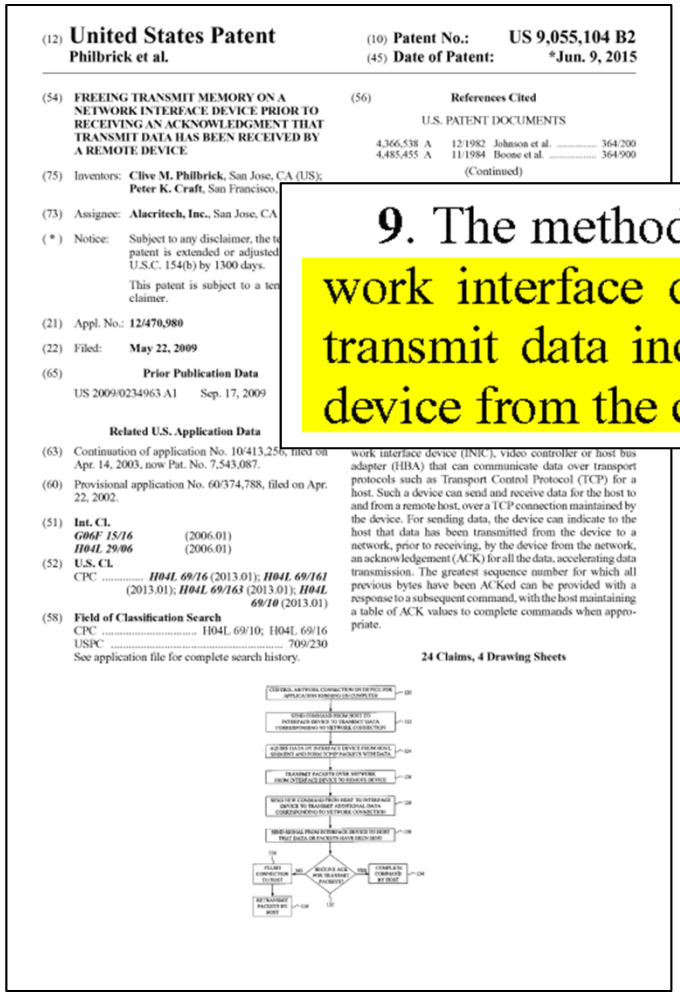

(12) **United States Patent**
Philbrick et al.

(10) **Patent No.: US 9,055,104 B2**
(45) **Date of Patent: *Jun. 9, 2015**

(54) **FREEING TRANSMIT MEMORY ON A NETWORK INTERFACE DEVICE PRIOR TO RECEIVING AN ACKNOWLEDGMENT THAT TRANSMIT DATA HAS BEEN RECEIVED BY A REMOTE DEVICE**

(75) Inventors: **Clive M. Philbrick**, San Jose, CA (US); **Peter K. Craft**, San Francisco,

(73) Assignee: **Alacritech, Inc.**, San Jose, CA

(*) Notice: Subject to any disclaimer, the t
patent is extended or adjusted
U.S.C. 154(b) by 1300 days.
This patent is subject to a ter
claimer.

(21) Appl. No.: **12/470,980**

(22) Filed: **May 22, 2009**

(65) **Prior Publication Data**
US 2009/0234963 A1 Sep. 17, 2009

Related U.S. Application Data

(63) Continuation of application No. 10/413,256, filed on Apr. 14, 2003, now Pat. No. 7,543,087.

(60) Provisional application No. 60/374,788, filed on Apr. 22, 2002.

(51) **Int. Cl.**
G06F 15/16 (2006.01)
H04L 29/06 (2006.01)

(52) **U.S. Cl.**
CPC ***H04L 69/16*** (2013.01); ***H04L 69/161*** (2013.01); ***H04L 69/163*** (2013.01); ***H04L 69/10*** (2013.01)

(58) **Field of Classification Search**
CPC H04L 69/10; H04L 69/16
USPC 709/230
See application file for complete search history.

(56) **References Cited**
U.S. PATENT DOCUMENTS
4,366,538 A 12/1982 Johnson et al. 364/200
4,485,455 A 11/1984 Boone et al. 364/900
(Continued)

work interface device (INIC), video controller or host bus adapter (HBA) that can communicate data over transport protocols such as Transport Control Protocol (TCP) for a host. Such a device can send and receive data for the host to and from a remote host, over a TCP connection maintained by the device. For sending data, the device can indicate to the host that data has been transmitted from the device to a network, prior to receiving, by the device from the network, an acknowledgement (ACK) for all the data, accelerating data transmission. The greatest sequence number for which all previous bytes have been ACKed can be provided with a response to a subsequent command, with the host maintaining a table of ACK values to complete commands when appropriate.

24 Claims, 4 Drawing Sheets

9. The method of claim **1**, wherein receiving, by the network interface device from the computer, a command to transmit data includes receiving, by the network interface device from the computer, a pointer to the command.

Ex. 1001.013 (104 Patent) at Claim 9.

Connery: A value for the transmit command may be passed by a pointer

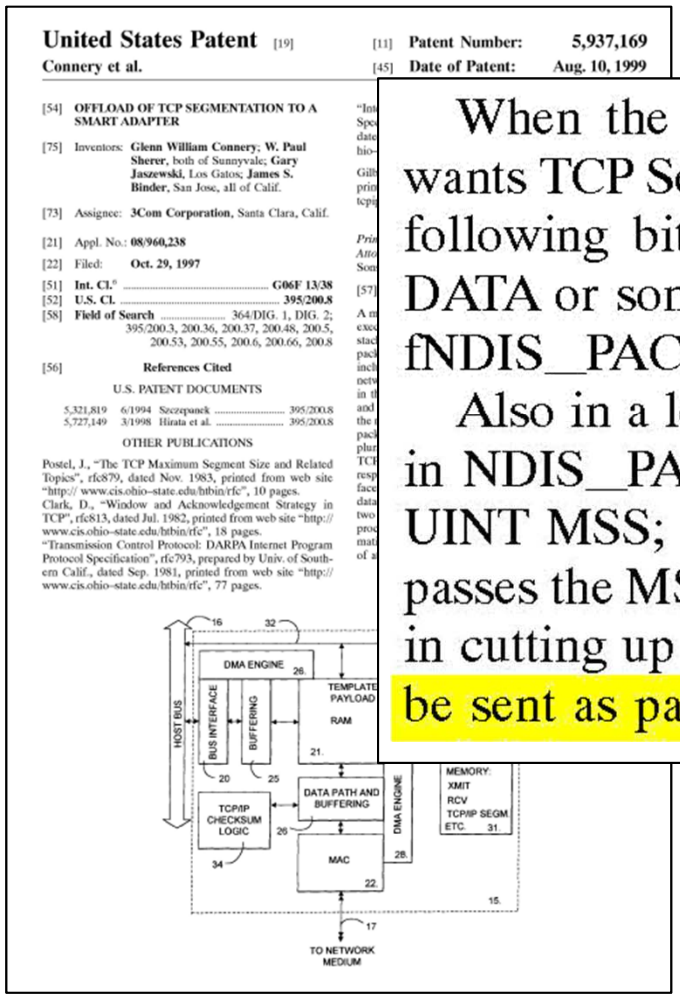

United States Patent [19]
Connery et al.

[11] Patent Number: 5,937,169
[45] Date of Patent: Aug. 10, 1999

[54] OFFLOAD OF TCP SEGMENTATION TO A SMART ADAPTER

[75] Inventors: Glenn William Connery; W. Paul Sherer, both of Sunnyvale; Gary Jaszewski, Los Gatos; James S. Binder, San Jose, all of Calif.

[73] Assignee: 3Com Corporation, Santa Clara, Calif.

[21] Appl. No.: 08/960,238

[22] Filed: Oct. 29, 1997

[51] Int. Cl.[6] G06F 13/38

[52] U.S. Cl. 395/200.8

[58] Field of Search ..... 364/DIG. 1, DIG. 2; 395/200.3, 200.36, 200.37, 200.48, 200.5, 200.53, 200.55, 200.6, 200.66, 200.8

[56] References Cited

U.S. PATENT DOCUMENTS

5,321,819 6/1994 Szczepanek 395/200.8
5,727,149 3/1998 Hirata et al. 395/200.8

OTHER PUBLICATIONS

Postel, J., "The TCP Maximum Segment Size and Related Topics", rfc879, dated Nov. 1983, printed from web site "http:// www.cis.ohio-state.edu/htbin/rfc", 10 pages.
Clark, D., "Window and Acknowledgement Strategy in TCP", rfc813, dated Jul. 1982, printed from web site "http:// www.cis.ohio-state.edu/htbin/rfc", 18 pages.
"Transmission Control Protocol: DARPA Internet Program Protocol Specification", rfc793, prepared by Univ. of Southern Calif., dated Sep. 1981, printed from web site "http:// www.cis.ohio-state.edu/htbin/rfc", 77 pages.

When the TCP/IP stack sends a datagram on which it wants TCP Segmentation to be done by the driver, it sets the following bit in Send Flags, or its equivalent in OOB_ DATA or some new location in NDIS_PACKET:
fNDIS_PACKET_TCPSEG

Also in a location in the OOB_DATA (or a new location in NDIS_PACKET) the value:
UINT MSS;
passes the MSS for this session down to the driver to be used in cutting up the payload into packets. This value may also be sent as part of a structure by passing a pointer instead.

Ex. 1043.012 (Connery) at 10:7-17;
Ex. 1003.094 (Horst Decl.) at A-34;
Paper 1 (104 Petition) at 75.

Dr. Horst: Passing a pointer was one of a limited number of ways of sending commands that was easily implemented

UNITED STATES PATENT AND TRADEMARK OFFICE

Connery's host computer sending the command itself to the network interface device, or instead (or additionally) sending a pointer to the command, would have been two of a limited number of well-known methods for the host computer to command the network interface device to transmit data, and a POSA could have easily implemented the pointer method with predicable results—namely, the proper transmission of the command.

Case IPR. No. **Unassigned**
U.S. Patent No. 9,055,104
Title: FREEING TRANSMIT MEMORY ON A NETWORK INTERFACE DEVICE PRIOR TO RECEIVING AN ACKNOWLEDGEMENT THAT TRANSMIT DATA HAS BEEN RECEIVED BY A REMOTE DEVICE

Declaration of Robert Horst, Ph.D. in Support of Petition for *Inter Partes* Review of U.S. Patent No. 9,055,104

Ex. 1003.095 (Horst Decl.) at A-35;
Paper 1 (104 Petition) at 75.

Dr. Horst: Matsunami corroborates passing a pointer to send a command was known to a POSA

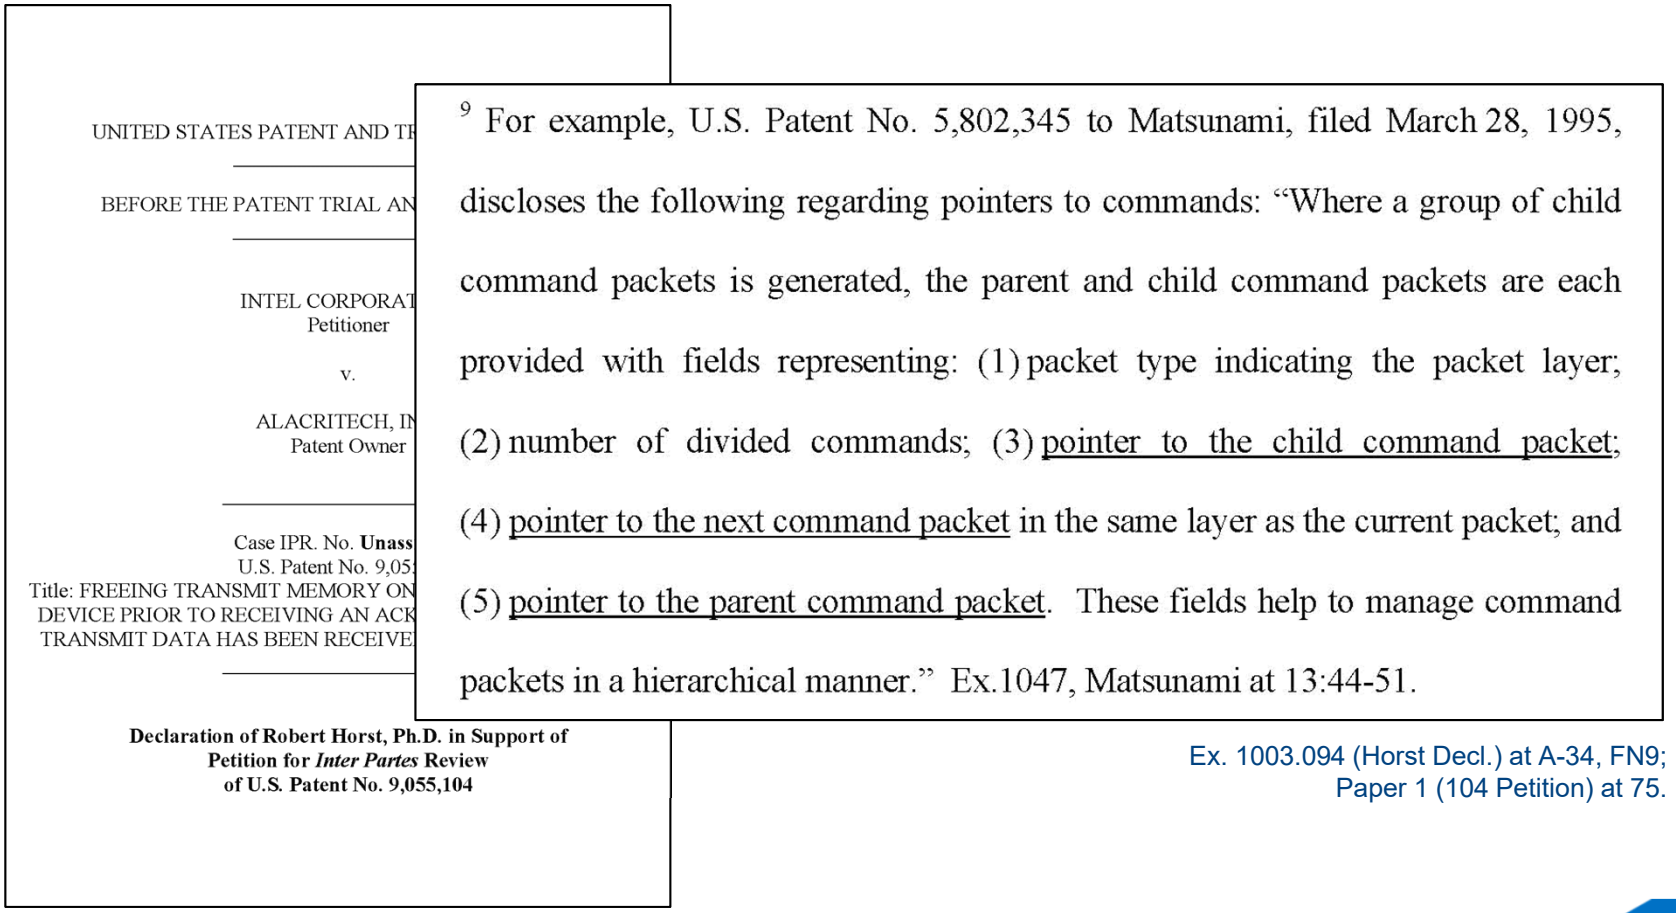

UNITED STATES PATENT AND TR

BEFORE THE PATENT TRIAL AN

INTEL CORPORAT
Petitioner

v.

ALACRITECH, IN
Patent Owner

Case IPR. No. **Unass**
U.S. Patent No. 9,05
Title: FREEING TRANSMIT MEMORY ON
DEVICE PRIOR TO RECEIVING AN ACK
TRANSMIT DATA HAS BEEN RECEIVE

Declaration of Robert Horst, Ph.D. in Support of Petition for *Inter Partes* Review of U.S. Patent No. 9,055,104

[9] For example, U.S. Patent No. 5,802,345 to Matsunami, filed March 28, 1995, discloses the following regarding pointers to commands: "Where a group of child command packets is generated, the parent and child command packets are each provided with fields representing: (1) packet type indicating the packet layer; (2) number of divided commands; (3) pointer to the child command packet; (4) pointer to the next command packet in the same layer as the current packet; and (5) pointer to the parent command packet. These fields help to manage command packets in a hierarchical manner." Ex.1047, Matsunami at 13:44-51.

Ex. 1003.094 (Horst Decl.) at A-34, FN9;
Paper 1 (104 Petition) at 75.

104 Patent: Disputes

2. Supplemental briefing regarding claim 22
   a. **Claim 22 is subject to §112(f)**
   b. The 104 Patent specification does not provide a corresponding structure for performing the various "means"
   c. Alternatively, if the claim is not subject to §112(f) or the disclosed "means" is a network interface, the grounds disclose all limitations of claim 22

Claim 22: “Means for” limitations

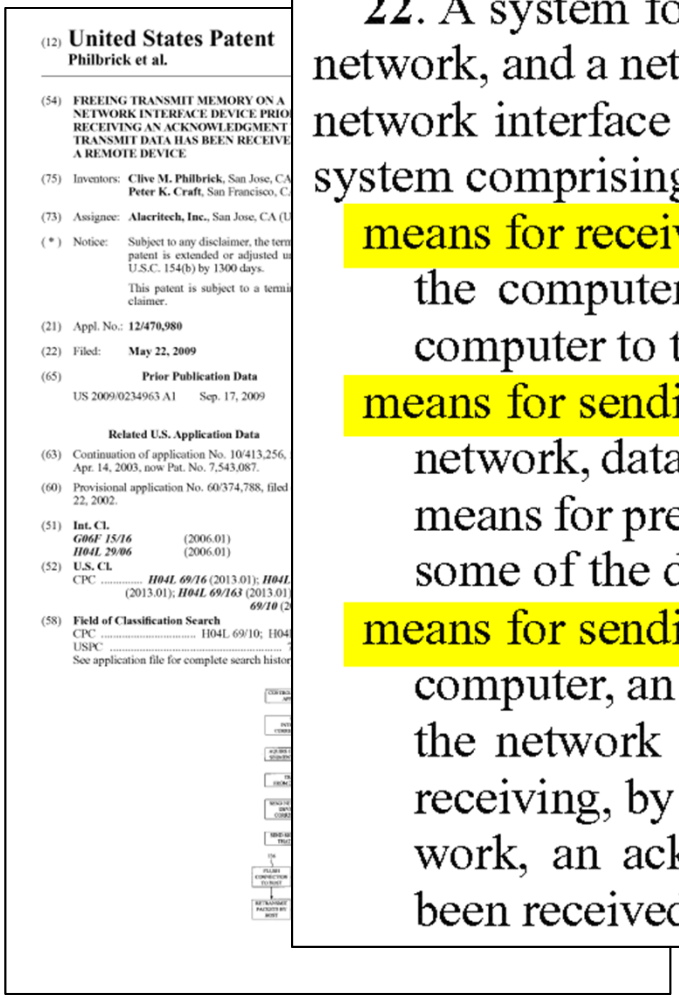

(12) **United States Patent**
Philbrick et al.

(54) FREEING TRANSMIT MEMORY ON A NETWORK INTERFACE DEVICE PRIO RECEIVING AN ACKNOWLEDGMENT TRANSMIT DATA HAS BEEN RECEIVE A REMOTE DEVICE

(75) Inventors: **Clive M. Philbrick**, San Jose, CA; **Peter K. Craft**, San Francisco, C

(73) Assignee: **Alacritech, Inc.**, San Jose, CA (U

(*) Notice: Subject to any disclaimer, the term patent is extended or adjusted u U.S.C. 154(b) by 1300 days.
This patent is subject to a termi claimer.

(21) Appl. No.: **12/470,980**

(22) Filed: **May 22, 2009**

(65) **Prior Publication Data**
US 2009/0234963 A1 Sep. 17, 2009

Related U.S. Application Data

(63) Continuation of application No. 10/413,256, Apr. 14, 2003, now Pat. No. 7,543,087.

(60) Provisional application No. 60/374,788, filed 22, 2002.

(51) **Int. Cl.**
G06F 15/16 (2006.01)
H04L 29/06 (2006.01)

(52) **U.S. Cl.**
CPC ***H04L 69/16*** (2013.01); ***H04L*** (2013.01); ***H04L 69/163*** (2013.01) ***69/10*** (2

(58) **Field of Classification Search**
CPC H04L 69/10; H04
USPC .....
See application file for complete search histor

22. A system for communication involving a computer, a network, and a network interface device of the computer, the network interface device being coupled to the network, the system comprising:

means for receiving, by the network interface device from the computer, a command to transmit data from the computer to the network;

means for sending, by the network interface device to the network, data corresponding to the command, including means for prepending a transport layer header to at least some of the data; and

means for sending, by the network interface device to the computer, an indication that the data has been sent from the network interface device to the network, prior to receiving, by the network interface device from the network, an acknowledgement (ACK) that the data has been received.

Ex. 1001.013 (104 Patent) at Claim 22.

Presumption that §112(f) applies when using "means for" + functional language

- "**means for receiving**, by the network interface device from the computer, **a command to transmit data from the computer to the network**;"
- "**means for sending**, by the network interface device to the network, **data corresponding to the command**,"
- "including **means for prepending a transport layer header to at least some of the data**; and"
- "**means for sending**, by the network interface device to the computer, **an indication that the data has been sent** from the network interface device to the network, prior to receiving, by the network interface device from the network, an acknowledgement (ACK) that the data has been received."

Ex. 1001 (104 Patent) at Claim 22.

The claimed “network interface device” is not sufficient corresponding structure

The specification discloses network interface devices that cannot perform the claimed functions.

Paper 48 (104 Supplemental Response) at 5;
Paper 55 (104 Supplemental Reply) at 2-3.

> The answer is to use two modes of operation: One in which the network frames are processed on the INIC through TCP and one in which the card operates like a typical dumb NIC. We call these two modes fast-path, and slow-path. In the slow-path case,

Paper 55 (104 Supplemental Reply) at 3-4.

> The transmit case works in much the same fashion. In slow-path mode the packets are given to the INIC with all of the headers attached. The INIC simply sends these packets out as if it were a dumb NIC. In fast-path mode, the host gives raw data to the INIC which it must carve into MSS sized segments, add headers to the data, perform checksums on the segment, and then send it out on the wire.

Paper 55 (104 Supplemental Reply) at 3-4;
Ex. 1031 (Alacritech 1997 Provisional Application) at 010.

PO did not overcome §112(f) presumption

- The District Court agreed that a network interface device is not sufficient structure

> Simply, while the various "means" of the claim may be part of or attached to the network interface device, they are not just the network interface device.

Paper 55 (104 Supplemental Reply) at 3;
Ex. 2030.042 (*Markman* Order) at 42.

104 Patent: Disputes

2. Supplemental briefing regarding claim 22
   a. Claim 22 is subject to §112(f)
   b. <u>The 104 Patent specification does not provide a corresponding structure for performing the various "means"</u>
   c. Alternatively, if the claim is not subject to §112(f) or the disclosed "means" is a network interface, the grounds disclose all limitations of claim 22

PO points to an “interface device” as the structure

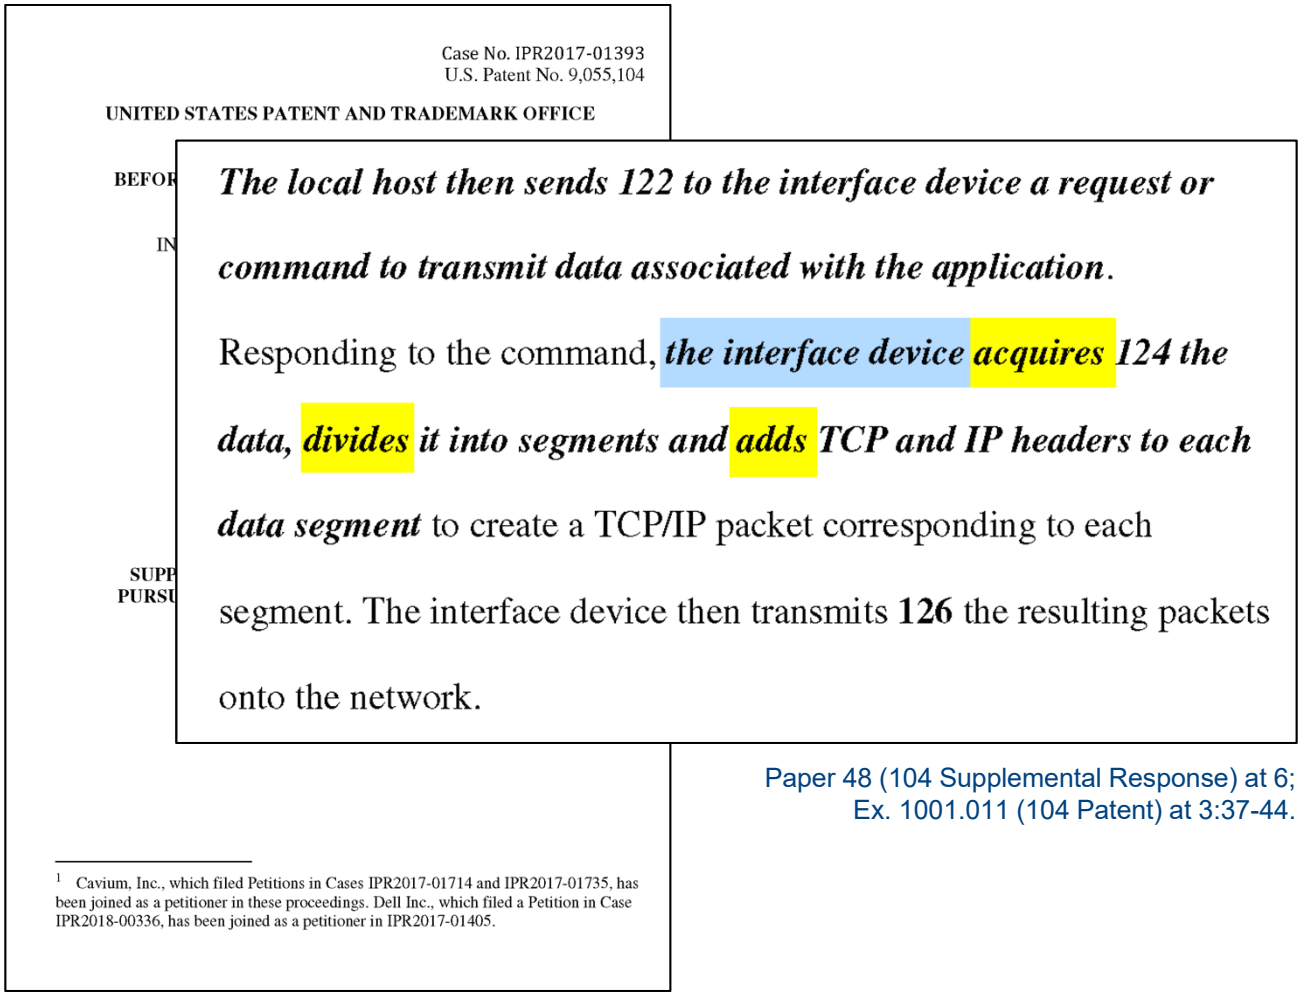

Case No. IPR2017-01393
U.S. Patent No. 9,055,104

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFOR

IN

SUPP
PURSU

The local host then sends 122 to the interface device a request or command to transmit data associated with the application.

Responding to the command, ***the interface device acquires 124 the data, divides it into segments and adds TCP and IP headers to each data segment*** to create a TCP/IP packet corresponding to each segment. The interface device then transmits **126** the resulting packets onto the network.

[1] Cavium, Inc., which filed Petitions in Cases IPR2017-01714 and IPR2017-01735, has been joined as a petitioner in these proceedings. Dell Inc., which filed a Petition in Case IPR2018-00336, has been joined as a petitioner in IPR2017-01405.

Paper 48 (104 Supplemental Response) at 6;
Ex. 1001.011 (104 Patent) at 3:37-44.

No structure on the interface device is identified for the functions

PO points to a “device” as the structure

Case No. IPR2017-01393
U.S. Patent No. 9,055,104

UNITED STATES PATENT AND TRADEMARK OFFICE

(‘104 Patent, 3:37-44; *see also id.* at 2:44-49 (“In one embodiment of the present disclosure this problem is solved by sending, from the device to the host, a signal that the data has been sent from the device to the network, prior to receiving, by the device from the network, an ACK that all the data has been received.”); 5:3-32,

Case IPR2017-01393[1]
U.S. Patent 9,055,104

SUPPLEMENTAL PATENT OWNER’S RESPONSE PURSUANT TO 35 U.S.C. § 313 AND 37 C.F.R. § 42.107

[1] Cavium, Inc., which filed Petitions in Cases IPR2017-01714 and IPR2017-01735, has been joined as a petitioner in these proceedings. Dell Inc., which filed a Petition in Case IPR2018-00336, has been joined as a petitioner in IPR2017-01405.

Paper 48 (104 Supplemental Response) at 6;
Ex. 1001.010 (104 Patent) at 2:44:49.

No structure in “the device” is identified for the functions

104 Patent does not identify any structure within the INIC

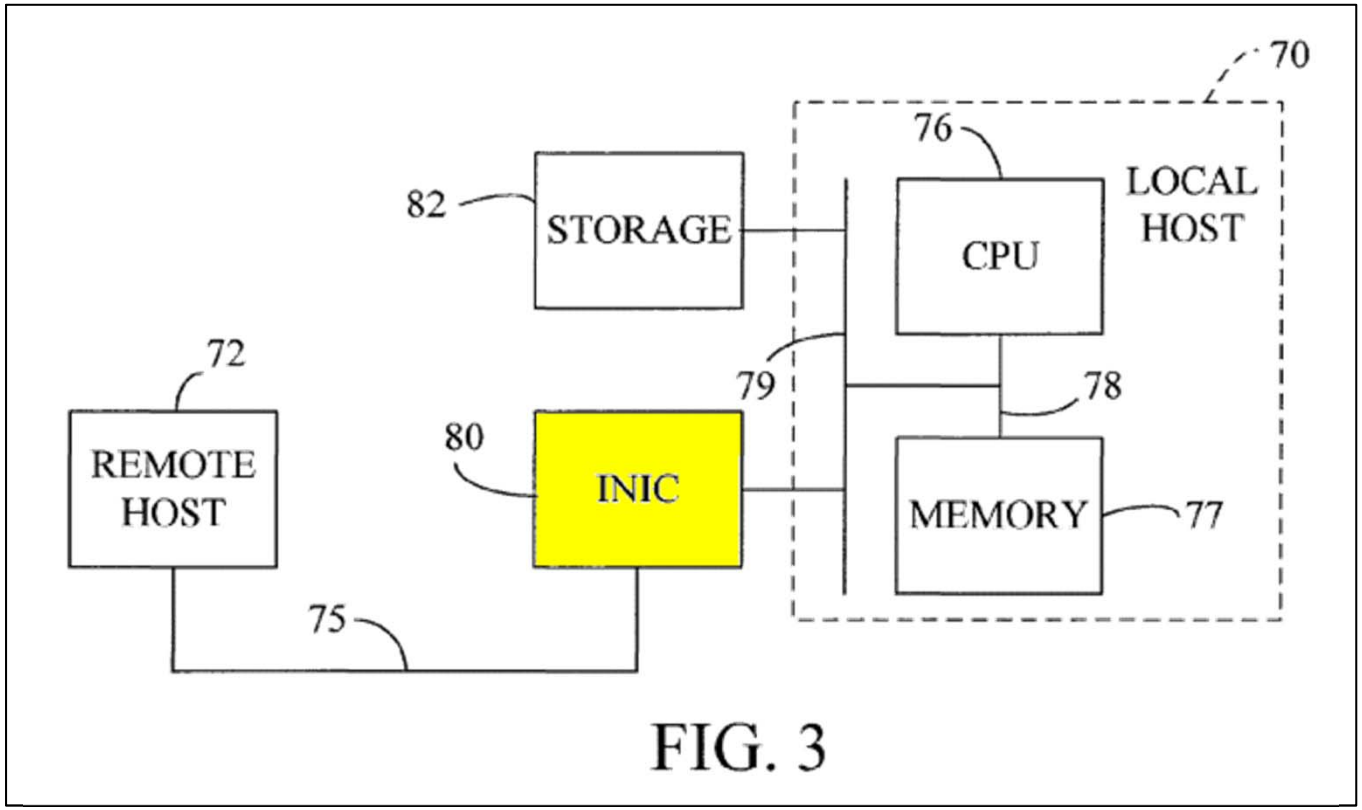

FIG. 3

Ex. 1001.007 (104 Patent) at Fig. 3.

104 Patent: Disputes

2. Supplemental briefing regarding claim 22
   a. Claim 22 is subject to §112(f)
   b. The 104 Patent specification does not provide a corresponding structure for performing the various "means"
   c. <u>Alternatively, if the claim is not subject to §112(f) or the disclosed "means" is a network interface, the grounds disclose all limitations of claim 22</u>

If claim 22 is not indefinite, then it is obvious for same reasons as claim 1

- Claim 22.p = claim 1.p (not challenged by PO)
- Claim 22.1 = claim 1.1 (not challenged by PO)
- Claim 22.2 = claim 1.2 (not challenged by PO)
- Claim 22.3 = claim 1.3 (addressed above)
- Claim 22.4 = claim 1.4, 1.5 (addressed above)

Ex. 1003.099-.100 (Horst Decl.) at A-39 – A-40;
Paper 1 (Petition) at 77-81;
Paper 55 (Supplemental Reply) at 4-6.

Evidence of Obviousness Far Outweighs Patent Owner's Alleged "Objective Evidence"

All citations refer to the docket for Case IPR2017-01391 unless otherwise noted.

Petitioner's arguments are the same for IPR2017-01392, -01393, -01405, -01406, -01409, -01410.

No evidence PO's products practice the claims

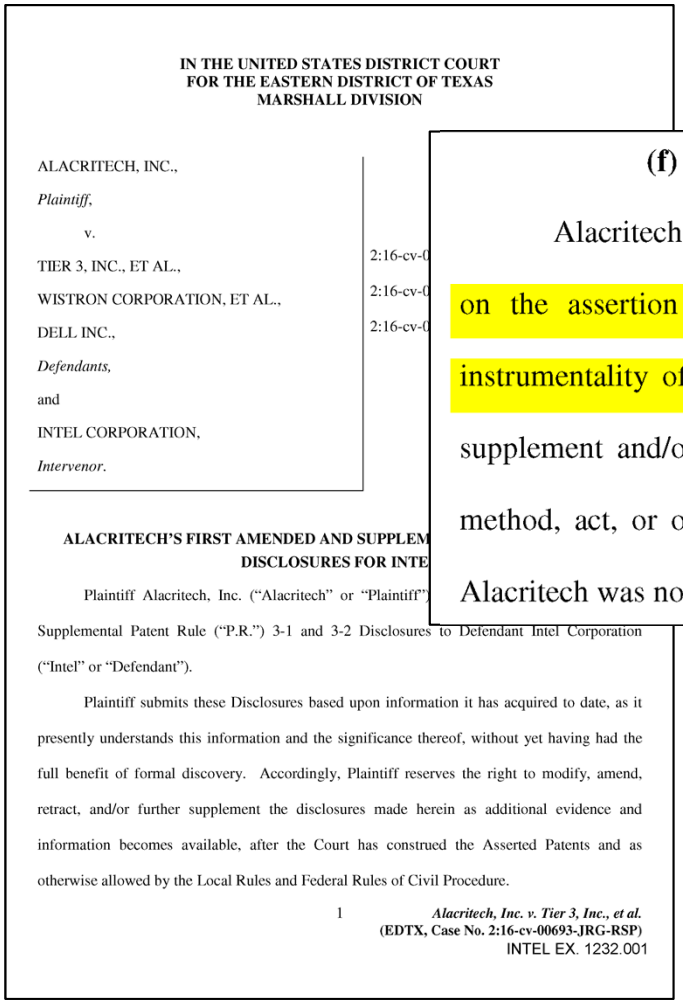

IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION

ALACRITECH, INC.,
Plaintiff,
v.
TIER 3, INC., ET AL.,
WISTRON CORPORATION, ET AL.,
DELL INC.,
Defendants,
and
INTEL CORPORATION,
Intervenor.

2:16-cv-0
2:16-cv-0
2:16-cv-0

ALACRITECH'S FIRST AMENDED AND SUPPLEM
DISCLOSURES FOR INTE

Plaintiff Alacritech, Inc. ("Alacritech" or "Plaintiff") Supplemental Patent Rule ("P.R.") 3-1 and 3-2 Disclosures to Defendant Intel Corporation ("Intel" or "Defendant").

Plaintiff submits these Disclosures based upon information it has acquired to date, as it presently understands this information and the significance thereof, without yet having had the full benefit of formal discovery. Accordingly, Plaintiff reserves the right to modify, amend, retract, and/or further supplement the disclosures made herein as additional evidence and information becomes available, after the Court has construed the Asserted Patents and as otherwise allowed by the Local Rules and Federal Rules of Civil Procedure.

1 ***Alacritech, Inc. v. Tier 3, Inc., et al.***
(EDTX, Case No. 2:16-cv-00693-JRG-RSP)
INTEL EX. 1232.001

(f) Alacritech Instrumentalities

Alacritech is still investigating this matter, however, at this time Alacritech is not relying on the assertion that its own apparatus, product, device, process, method, act, or other instrumentality of its own practices the claimed inventions. Alacritech reserves the right to supplement and/or amend this disclosure to identify any apparatus, product, device, process, method, act, or other instrumentality of its own that practices the Asserted Claims of which Alacritech was not aware at the time of these disclosures.

Paper 41 (036 Reply) at 20;
Ex. 1232.005 (Alacritech's First Amended and Supplemental Patent Local Rule 3-1 and 3-2 Disclosures).

PO's products were not commercially successful

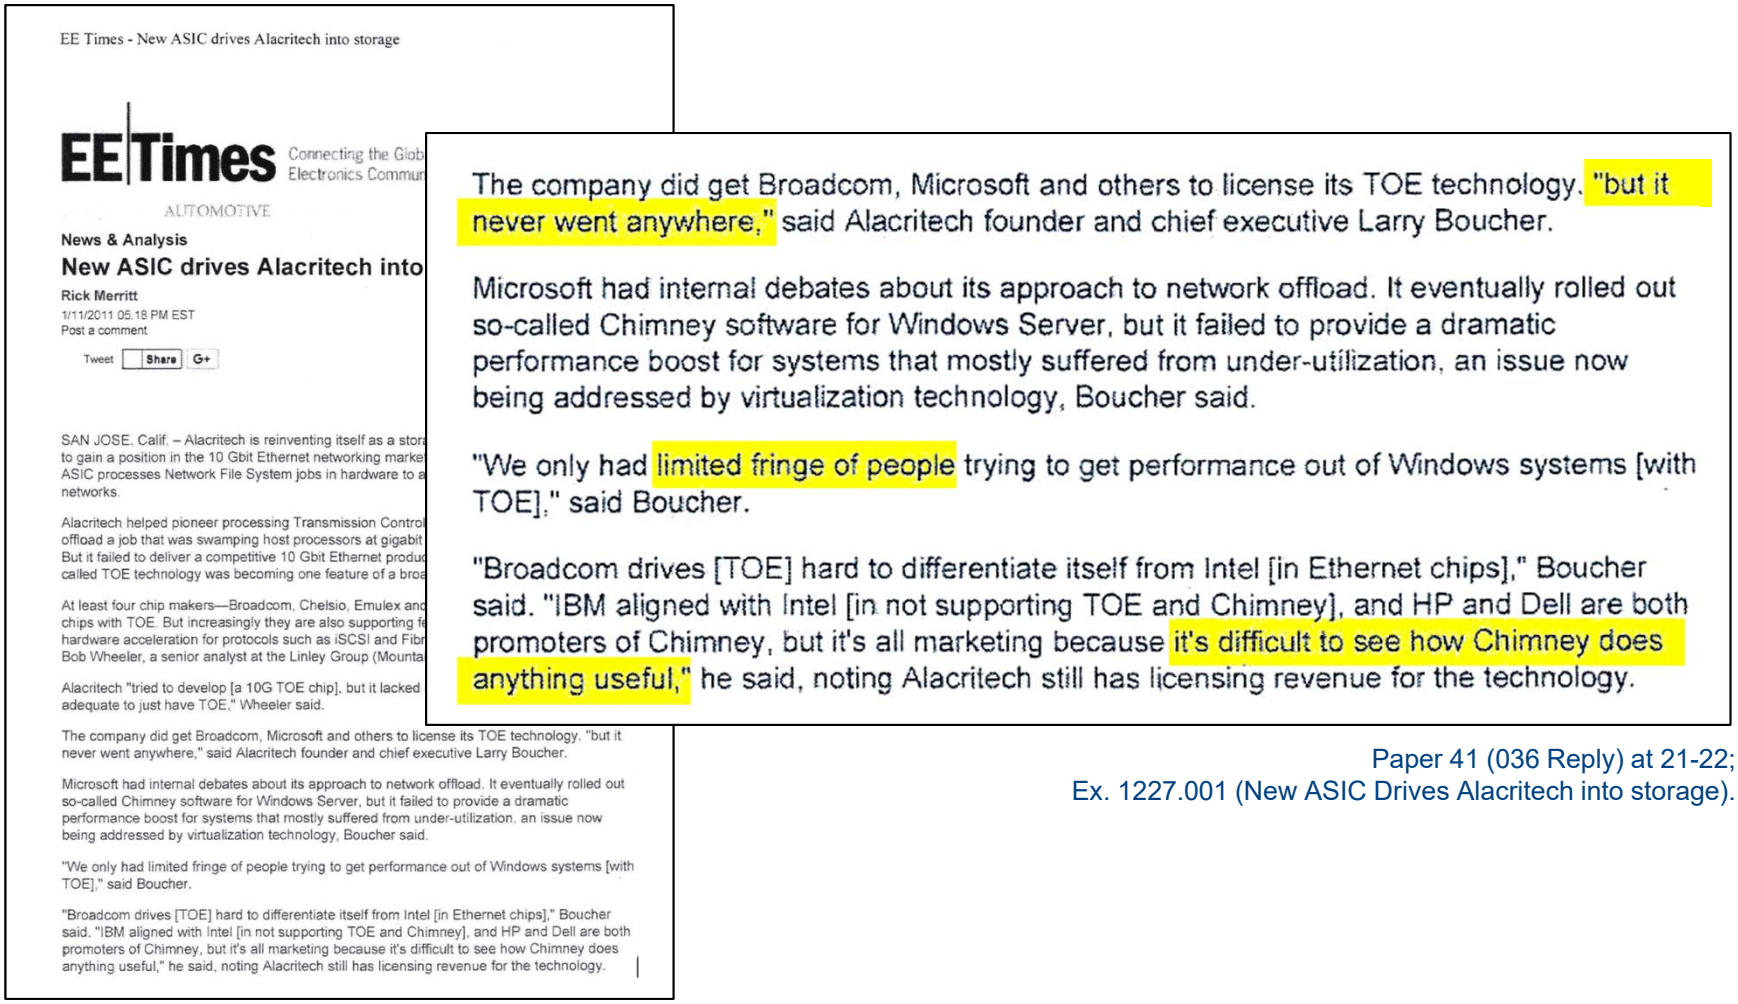

EE Times - New ASIC drives Alacritech into storage

EE Times Connecting the Glob… Electronics Commun…

AUTOMOTIVE

News & Analysis

New ASIC drives Alacritech into

Rick Merritt

1/11/2011 05:18 PM EST

Post a comment

Tweet Share G+

SAN JOSE, Calif. – Alacritech is reinventing itself as a stora… to gain a position in the 10 Gbit Ethernet networking marke… ASIC processes Network File System jobs in hardware to a… networks.

Alacritech helped pioneer processing Transmission Control… offload a job that was swamping host processors at gigabit… But it failed to deliver a competitive 10 Gbit Ethernet produc… called TOE technology was becoming one feature of a broa…

At least four chip makers—Broadcom, Chelsio, Emulex and… chips with TOE. But increasingly they are also supporting fe… hardware acceleration for protocols such as iSCSI and Fibr… Bob Wheeler, a senior analyst at the Linley Group (Mounta…

Alacritech "tried to develop [a 10G TOE chip], but it lacked … adequate to just have TOE," Wheeler said.

The company did get Broadcom, Microsoft and others to license its TOE technology, "but it never went anywhere," said Alacritech founder and chief executive Larry Boucher.

Microsoft had internal debates about its approach to network offload. It eventually rolled out so-called Chimney software for Windows Server, but it failed to provide a dramatic performance boost for systems that mostly suffered from under-utilization, an issue now being addressed by virtualization technology, Boucher said.

"We only had limited fringe of people trying to get performance out of Windows systems [with TOE]," said Boucher.

"Broadcom drives [TOE] hard to differentiate itself from Intel [in Ethernet chips]," Boucher said. "IBM aligned with Intel [in not supporting TOE and Chimney], and HP and Dell are both promoters of Chimney, but it's all marketing because it's difficult to see how Chimney does anything useful," he said, noting Alacritech still has licensing revenue for the technology.

The company did get Broadcom, Microsoft and others to license its TOE technology, "but it never went anywhere," said Alacritech founder and chief executive Larry Boucher.

Microsoft had internal debates about its approach to network offload. It eventually rolled out so-called Chimney software for Windows Server, but it failed to provide a dramatic performance boost for systems that mostly suffered from under-utilization, an issue now being addressed by virtualization technology, Boucher said.

"We only had limited fringe of people trying to get performance out of Windows systems [with TOE]," said Boucher.

"Broadcom drives [TOE] hard to differentiate itself from Intel [in Ethernet chips]," Boucher said. "IBM aligned with Intel [in not supporting TOE and Chimney], and HP and Dell are both promoters of Chimney, but it's all marketing because it's difficult to see how Chimney does anything useful," he said, noting Alacritech still has licensing revenue for the technology.

Paper 41 (036 Reply) at 21-22;
Ex. 1227.001 (New ASIC Drives Alacritech into storage).

“Conventional wisdom”: Use special purpose NICs for TCP/IP acceleration

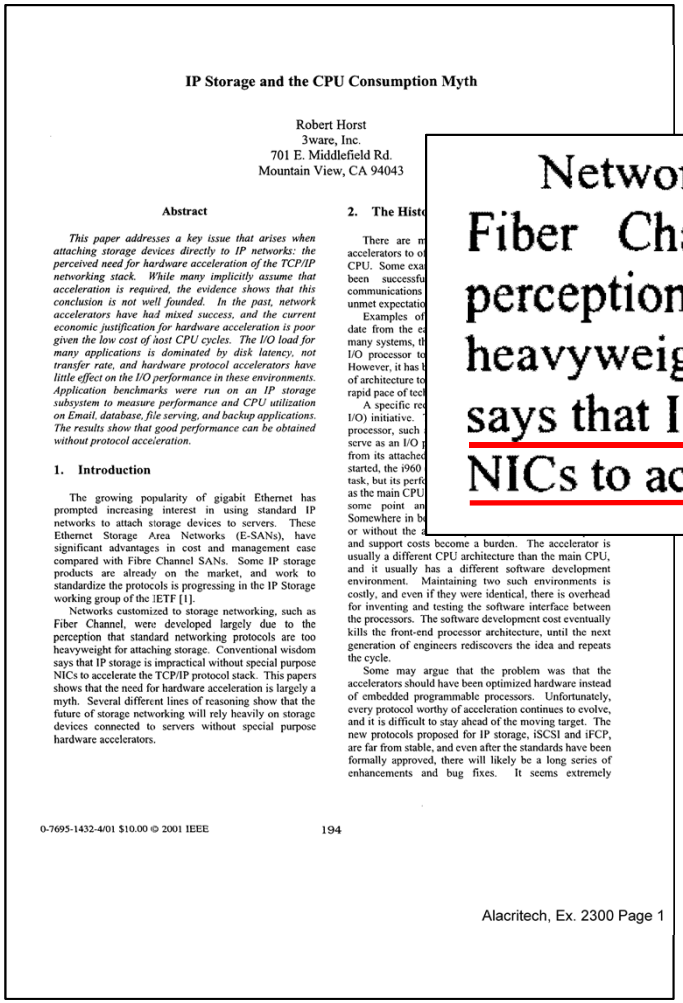

IP Storage and the CPU Consumption Myth

Robert Horst
3ware, Inc.
701 E. Middlefield Rd.
Mountain View, CA 94043

Abstract

This paper addresses a key issue that arises when attaching storage devices directly to IP networks: the perceived need for hardware acceleration of the TCP/IP networking stack. While many implicitly assume that acceleration is required, the evidence shows that this conclusion is not well founded. In the past, network accelerators have had mixed success, and the current economic justification for hardware acceleration is poor given the low cost of host CPU cycles. The I/O load for many applications is dominated by disk latency, not transfer rate, and hardware protocol accelerators have little effect on the I/O performance in these environments. Application benchmarks were run on an IP storage subsystem to measure performance and CPU utilization on Email, database, file serving, and backup applications. The results show that good performance can be obtained without protocol acceleration.

1. Introduction

The growing popularity of gigabit Ethernet has prompted increasing interest in using standard IP networks to attach storage devices to servers. These Ethernet Storage Area Networks (E-SANs), have significant advantages in cost and management ease compared with Fibre Channel SANs. Some IP storage products are already on the market, and work to standardize the protocols is progressing in the IP Storage working group of the IETF [1].

Networks customized to storage networking, such as Fiber Channel, were developed largely due to the perception that standard networking protocols are too heavyweight for attaching storage. Conventional wisdom says that IP storage is impractical without special purpose NICs to accelerate the TCP/IP protocol stack. This papers shows that the need for hardware acceleration is largely a myth. Several different lines of reasoning show that the future of storage networking will rely heavily on storage devices connected to servers without special purpose hardware accelerators.

2. The Histo

There are m
accelerators to o
CPU. Some exa
been successfu
communications
unmet expectatio

Examples of
date from the e
many systems, th
I/O processor te
However, it has l
of architecture to
rapid pace of tec

A specific re
I/O) initiative.
processor, such
serve as an I/O
from its attached
started, the i960
task, but its perfo
as the main CPU
some point an
Somewhere in b
or without the a
and support costs become a burden. The accelerator is usually a different CPU architecture than the main CPU, and it usually has a different software development environment. Maintaining two such environments is costly, and even if they were identical, there is overhead for inventing and testing the software interface between the processors. The software development cost eventually kills the front-end processor architecture, until the next generation of engineers rediscovers the idea and repeats the cycle.

Some may argue that the problem was that the accelerators should have been optimized hardware instead of embedded programmable processors. Unfortunately, every protocol worthy of acceleration continues to evolve, and it is difficult to stay ahead of the moving target. The new protocols proposed for IP storage, iSCSI and iFCP, are far from stable, and even after the standards have been formally approved, there will likely be a long series of enhancements and bug fixes. It seems extremely

0-7695-1432-4/01 $10.00 © 2001 IEEE

Alacritech, Ex. 2300 Page 1

> Networks customized to storage networking, such as Fiber Channel, were developed largely due to the perception that standard networking protocols are too heavyweight for attaching storage. Conventional wisdom says that IP storage is impractical without special purpose NICs to accelerate the TCP/IP protocol stack. This papers

Paper 41 (036 Reply) at 24;
Ex. 2300.001 (IP Storage and the CPU Consumption Myth).

Real Party in Interest is Correctly Named

*All citations refer to the docket for Case IPR2017-01391. Petitioner's arguments are the same for IPR2017-01392, -01393, -01405, -01406, -01409, -01410.

Board found RPI correctly named in institution decision

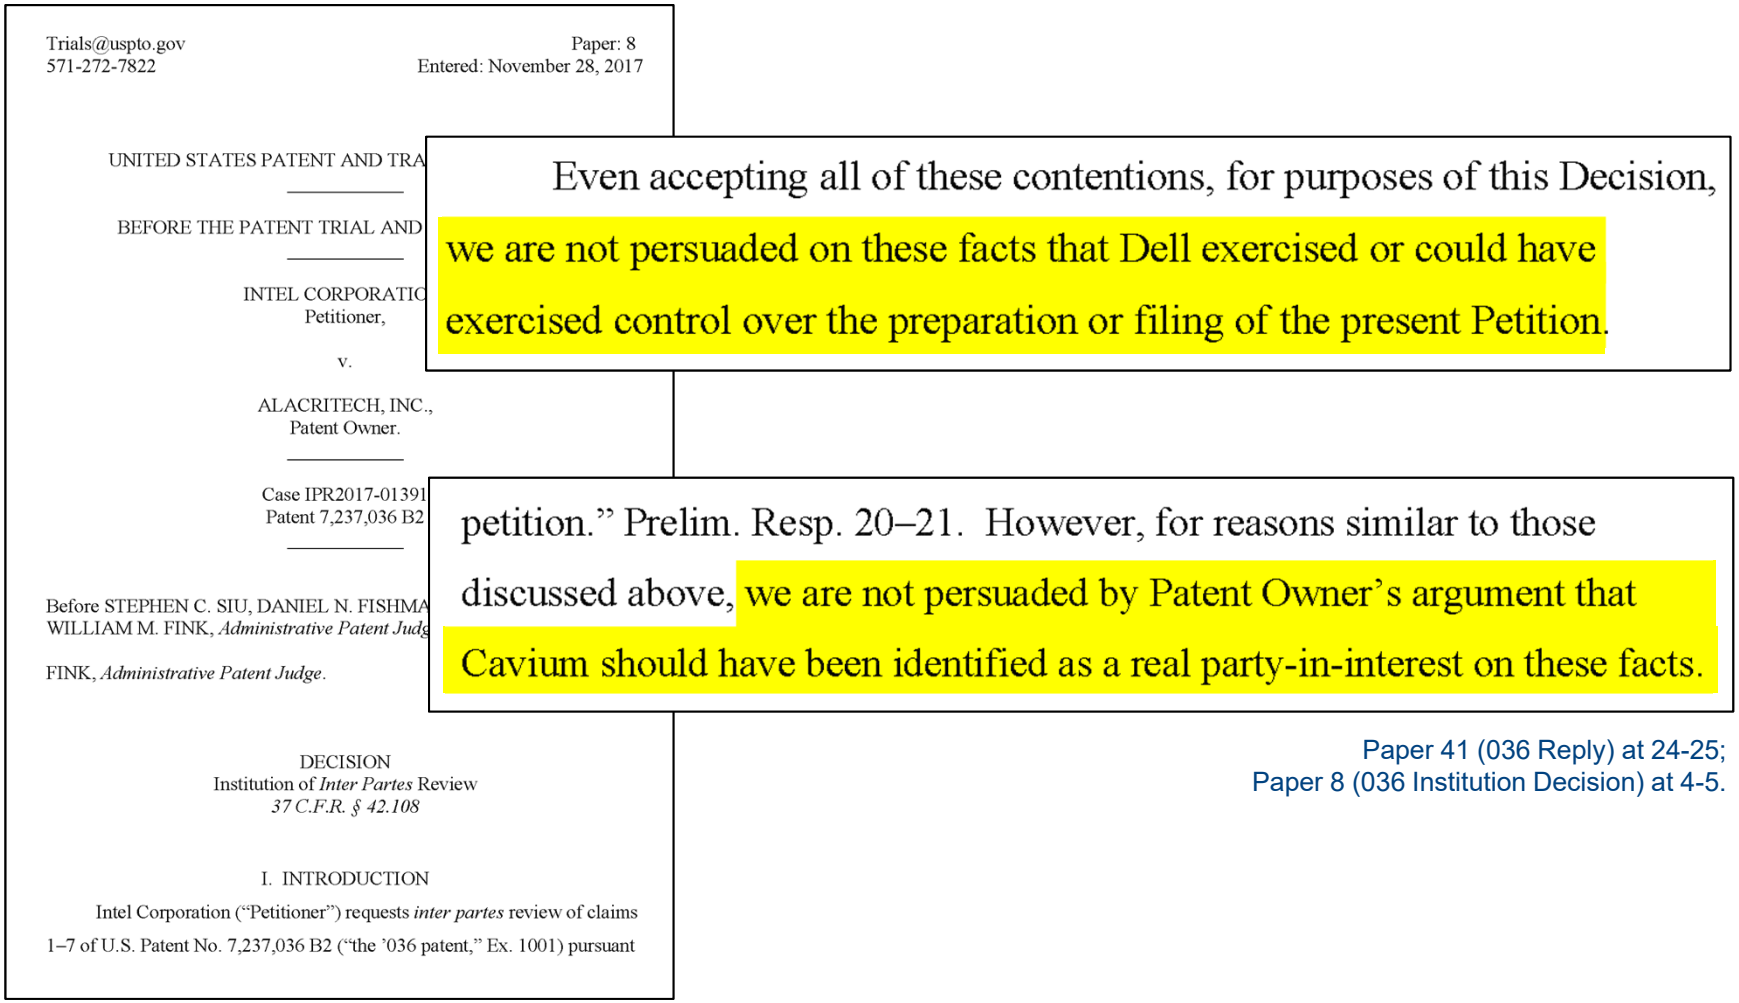

Trials@uspto.gov
571-272-7822

Paper: 8
Entered: November 28, 2017

UNITED STATES PATENT AND TRA

BEFORE THE PATENT TRIAL AND

INTEL CORPORATIC
Petitioner,

v.

ALACRITECH, INC.,
Patent Owner.

Case IPR2017-01391
Patent 7,237,036 B2

Before STEPHEN C. SIU, DANIEL N. FISHMA
WILLIAM M. FINK, *Administrative Patent Judg*

FINK, *Administrative Patent Judge*.

DECISION
Institution of *Inter Partes* Review
37 C.F.R. § 42.108

I. INTRODUCTION

Intel Corporation ("Petitioner") requests *inter partes* review of claims 1–7 of U.S. Patent No. 7,237,036 B2 ("the '036 patent," Ex. 1001) pursuant

Even accepting all of these contentions, for purposes of this Decision, we are not persuaded on these facts that Dell exercised or could have exercised control over the preparation or filing of the present Petition.

petition." Prelim. Resp. 20–21. However, for reasons similar to those discussed above, we are not persuaded by Patent Owner's argument that Cavium should have been identified as a real party-in-interest on these facts.

Paper 41 (036 Reply) at 24-25;
Paper 8 (036 Institution Decision) at 4-5.

PO relies only on speculation in arguing RPI is Incorrect

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND

INTEL CORP. and CAVIUM
Petitioners,

v.

ALACRITECH, INC.,
Patent Owner

Case IPR2017-01391[1]
U.S. Patent No. 7,237,03

CORRECTED PATENT OWNER'
PURSUANT TO 35 U.S.C. § 313 AND

[1] Cavium, Inc., filed a Petition in Case IPR2017-01718, and has been joined as a petitioner in this proceeding.

Each Petitioner identifies only one real party-in-interest: either Intel or Cavium. In doing so, the Petition at least fails to identify the other Petitioner (either Cavium or Intel), Dell Inc., Wistron, CenturyLink, and QLogic. Each of these parties are either a co-defendant or intervenor in Alacritech's patent infringement lawsuit over the challenged patent. Dell is both Intel's and Cavium's customer and indemnitee.[8] As Intel's and Cavium's Motions to Intervene filed in

Paper 41 (036 Reply) at 24-25;
Paper 30 (036 Response) at 57.

Alacritech accused Intel and Cavium of infringing the patents at issue

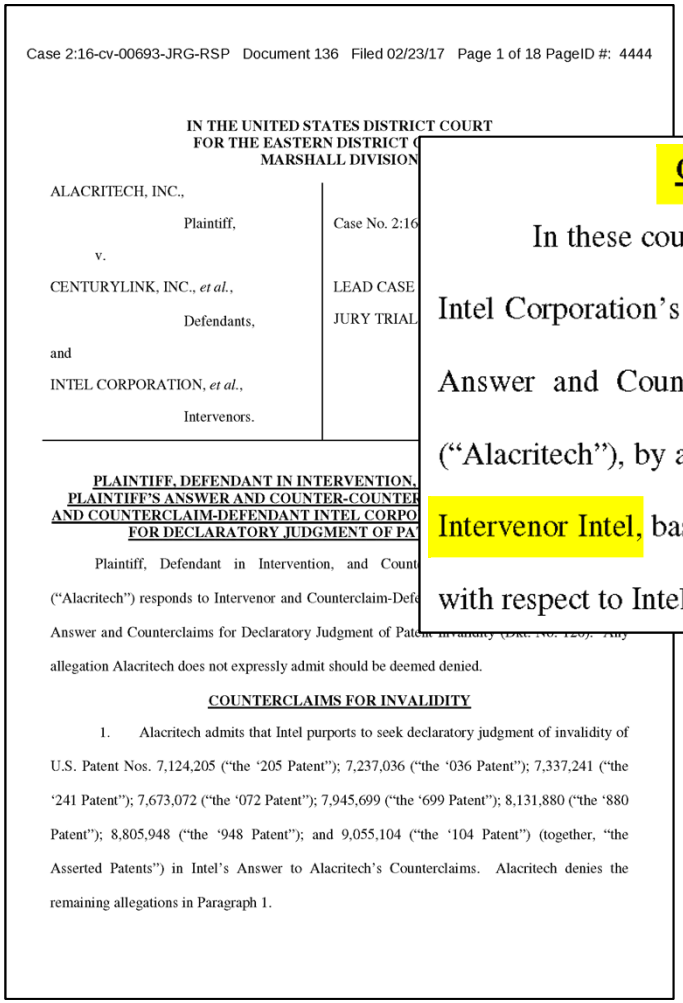

Case 2:16-cv-00693-JRG-RSP Document 136 Filed 02/23/17 Page 1 of 18 PageID #: 4444

IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT
MARSHALL DIVISION

ALACRITECH, INC.,
Plaintiff,
v.
CENTURYLINK, INC., *et al.*,
Defendants,
and
INTEL CORPORATION, *et al.*,
Intervenors.

Case No. 2:16
LEAD CASE
JURY TRIAL

PLAINTIFF, DEFENDANT IN INTERVENTION, PLAINTIFF'S ANSWER AND COUNTER-COUNTER AND COUNTERCLAIM-DEFENDANT INTEL CORPO FOR DECLARATORY JUDGMENT OF PA

Plaintiff, Defendant in Intervention, and Count ("Alacritech") responds to Intervenor and Counterclaim-Defe Answer and Counterclaims for Declaratory Judgment of Pate allegation Alacritech does not expressly admit should be deemed denied.

COUNTERCLAIMS FOR INVALIDITY

1. Alacritech admits that Intel purports to seek declaratory judgment of invalidity of U.S. Patent Nos. 7,124,205 ("the '205 Patent"); 7,237,036 ("the '036 Patent"); 7,337,241 ("the '241 Patent"); 7,673,072 ("the '072 Patent"); 7,945,699 ("the '699 Patent"); 8,131,880 ("the '880 Patent"); 8,805,948 ("the '948 Patent"); and 9,055,104 ("the '104 Patent") (together, "the Asserted Patents") in Intel's Answer to Alacritech's Counterclaims. Alacritech denies the remaining allegations in Paragraph 1.

COUNTERCLAIMS FOR PATENT INFRINGEMENT

In these counterclaims for patent infringement under 35 U.S.C. § 271 in response to both Intel Corporation's ("Intel") intervention in Case No. 2:16-cv-00692 to defend Dell, and Intel's Answer and Counterclaims (Dkt. 120) filed in the same matter, Plaintiff Alacritech, Inc. ("Alacritech"), by and through its undersigned counsel, complains and alleges as follows against Intervenor Intel, based on Alacritech's own personal knowledge and upon information and belief with respect to Intel's actions:

Paper 41 (036 Reply) at 25;
Ex. 1112.002 (Alacritech's Answer and Counterclaims against Intel);
See also Ex. 1233 (Alacritech's Answer and Counterclaims against Cavium).

Defendants and Cavium exercised no role in Intel's IPRs

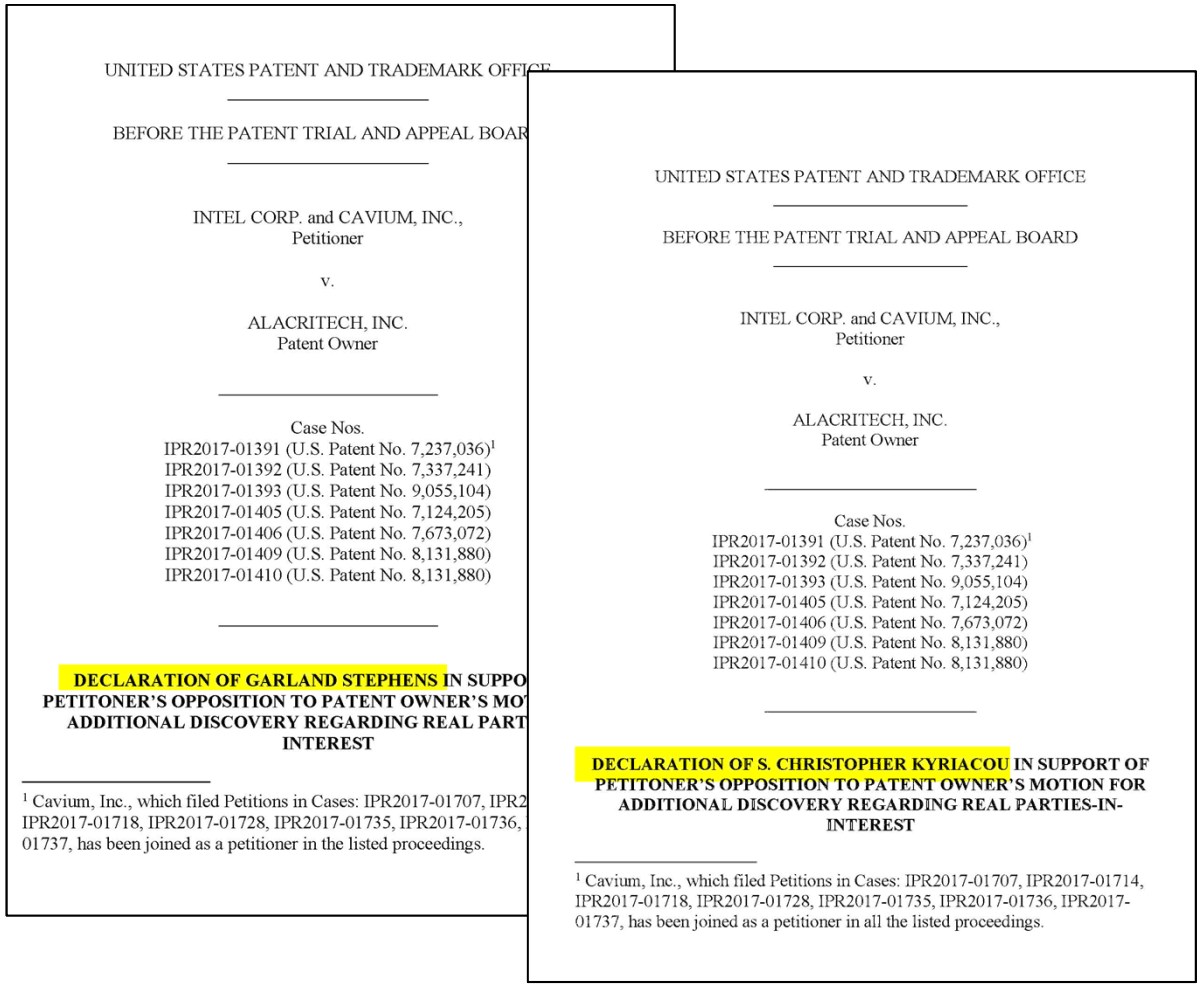

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOAR

INTEL CORP. and CAVIUM, INC.,
Petitioner

v.

ALACRITECH, INC.
Patent Owner

Case Nos.
IPR2017-01391 (U.S. Patent No. 7,237,036)[1]
IPR2017-01392 (U.S. Patent No. 7,337,241)
IPR2017-01393 (U.S. Patent No. 9,055,104)
IPR2017-01405 (U.S. Patent No. 7,124,205)
IPR2017-01406 (U.S. Patent No. 7,673,072)
IPR2017-01409 (U.S. Patent No. 8,131,880)
IPR2017-01410 (U.S. Patent No. 8,131,880)

DECLARATION OF GARLAND STEPHENS IN SUPPO
PETITONER'S OPPOSITION TO PATENT OWNER'S MO
ADDITIONAL DISCOVERY REGARDING REAL PART
INTEREST

[1] Cavium, Inc., which filed Petitions in Cases: IPR2017-01707, IPR2
IPR2017-01718, IPR2017-01728, IPR2017-01735, IPR2017-01736,
01737, has been joined as a petitioner in the listed proceedings.

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

INTEL CORP. and CAVIUM, INC.,
Petitioner

v.

ALACRITECH, INC.
Patent Owner

Case Nos.
IPR2017-01391 (U.S. Patent No. 7,237,036)[1]
IPR2017-01392 (U.S. Patent No. 7,337,241)
IPR2017-01393 (U.S. Patent No. 9,055,104)
IPR2017-01405 (U.S. Patent No. 7,124,205)
IPR2017-01406 (U.S. Patent No. 7,673,072)
IPR2017-01409 (U.S. Patent No. 8,131,880)
IPR2017-01410 (U.S. Patent No. 8,131,880)

DECLARATION OF S. CHRISTOPHER KYRIACOU IN SUPPORT OF PETITONER'S OPPOSITION TO PATENT OWNER'S MOTION FOR ADDITIONAL DISCOVERY REGARDING REAL PARTIES-IN-INTEREST

[1] Cavium, Inc., which filed Petitions in Cases: IPR2017-01707, IPR2017-01714, IPR2017-01718, IPR2017-01728, IPR2017-01735, IPR2017-01736, IPR2017-01737, has been joined as a petitioner in all the listed proceedings.

Paper 41 (036 Reply) at 24-25;
Ex. 1110 (Stephens Decl.);
Ex. 1111 (Kyriacou Decl.).

Tanenbaum96 Is Prior Art

All citations refer to the docket for Case IPR2017-01391 unless otherwise noted.

Petitioner's arguments are the same for IPR2017-01392, -01406, -01409, -01410.

Argument summary

1. **It is improper to raise public availability in a Motion to Exclude**

2. Patent Owner admits Tanenbaum96 was publicly available

3. Librarian declaration establishes Tanenbaum96 was publicly available to a POSA

Public availability of Tanenbaum96 raised for first time in Motion to Exclude

- Patent Owner did not raise public availability in Patent Owner's Response to Petitions

- Raised for the first time in Motions to Exclude

It is improper to raise public availability in a motion to exclude

"A motion to exclude must explain why the evidence is not admissible (e.g., relevance or hearsay) but may not be used to challenge the **sufficiency of the evidence** to prove a particular fact."

PTAB Trial Practice Guide, 77 Fed. Reg. 48756, 48767 § II(K) (Aug. 14, 2012) (emphasis added).
Update to PTAB Trial Practice Guide, §(K), p. 16 (August 2018).

"A motion to exclude is the **wrong vehicle to challenge public availability**, which is a substantive issue that goes to the sufficiency of the evidence, not to admissibility at issue here."

Arista Networks v. Cisco Sys., Case IPR2016-00303, Paper 53 at 9 (May 25, 2017) (emphasis added).
Paper 60 (Opp. to Motion to Exclude) at 3-4.

Argument summary

1. It is improper to raise public availability in a Motion to Exclude

2. **Patent Owner admits Tanenbaum96 was publicly available**

3. Librarian declaration establishes Tanenbaum96 was publicly available to a POSA

PO's expert taught Tanenbaum96 before alleged priority date Oct. 1997

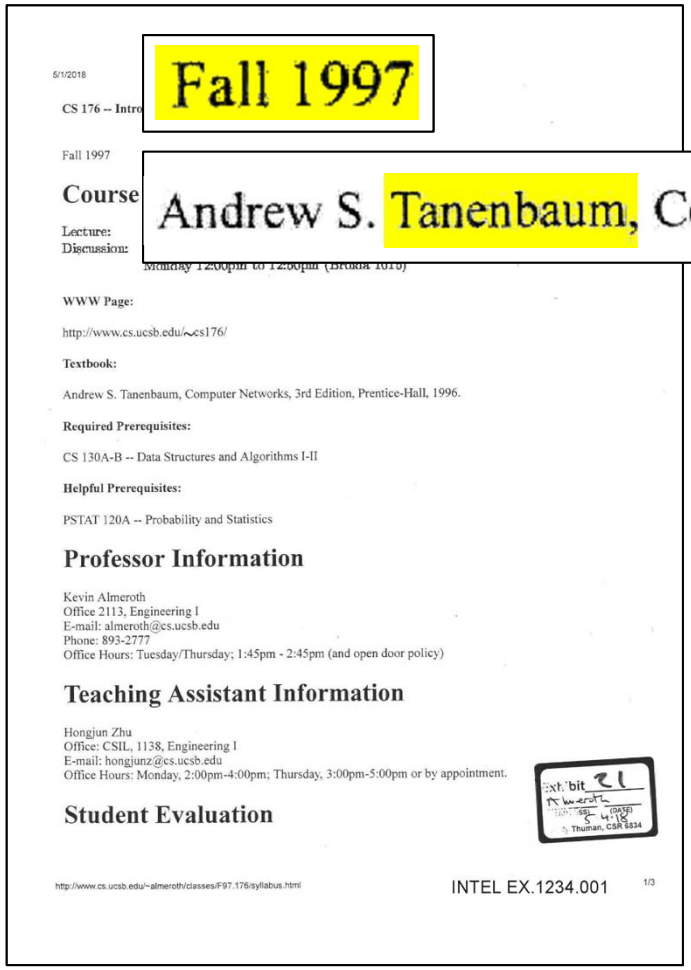

5/1/2018

CS 176 -- Intro

Fall 1997

Course

Lecture:
Discussion:

Monday 12:00pm to 12:00pm (Broida 1013)

WWW Page:

http://www.cs.ucsb.edu/~cs176/

Textbook:

Andrew S. Tanenbaum, Computer Networks, 3rd Edition, Prentice-Hall, 1996.

Required Prerequisites:

CS 130A-B -- Data Structures and Algorithms I-II

Helpful Prerequisites:

PSTAT 120A -- Probability and Statistics

Professor Information

Kevin Almeroth
Office 2113, Engineering I
E-mail: almeroth@cs.ucsb.edu
Phone: 893-2777
Office Hours: Tuesday/Thursday; 1:45pm - 2:45pm (and open door policy)

Teaching Assistant Information

Hongjun Zhu
Office: CSIL, 1138, Engineering I
E-mail: hongjunz@cs.ucsb.edu
Office Hours: Monday, 2:00pm-4:00pm; Thursday, 3:00pm-5:00pm or by appointment.

Student Evaluation

Exhibit 21
K. Almeroth
SGL 5-16-18 (DATE)
J. Thuman, CSR 6834

http://www.cs.ucsb.edu/~almeroth/classes/F97.176/syllabus.html

INTEL EX.1234.001

1/3

Fall 1997

Andrew S. Tanenbaum, Computer Networks, 3rd Edition, Prentice-Hall, 1996.

Q. Dr. Almeroth, do you recognize Exhibit 21?

A. It looks like the front page for the first course at UCSB that I taught.

Q. And the textbook was the Tanenbaum '96, right, that's the basis for several of the grounds that we've been talking about today and yesterday, right?

A. Yes. It was

Paper 60 (036 Opp. to Motion to Exclude) at 5-6;
Ex. 1234.001 (Almeroth Depo., Ex. 21);
Ex. 1225 (Almeroth Depo.) at 474:21-475:2.

PO's expert taught Tanenbaum96 before alleged priority date Oct. 1997

	Date	Lecture Topic	Reading	Assignment
1	Th 9/25	Course Overview	Chap 1, especially Sect 1.4	
2	T 9/30	Physical Layer Overview	skim Chap 2	
		Data Link Layer – Errors	Sect 3.1, 3.2	
		Data Link Layer – CRC	Sect 3.2	
3	Th 10/2	DLL Protocols	Sect 3.3,3.4	

Andrew S. Tanenbaum, Computer Networks, 3rd Edition, Prentice-Hall, 1996.

Required Prerequisites:

CS 130A-B -- Data Structures and Algorithms I-II

Helpful Prerequisites:

PSTAT 120A -- Probability and Statistics

Professor Informa

Kevin Almeroth
Office 2113, Engineering I
E-mail: almeroth@cs.ucsb.edu
Phone: 893-2777
Office Hours: Tuesday/Thursday; 1:4

Teaching Assistan

Hongjun Zhu
Office: CSIL, 1138, Engineering I
E-mail: hongjunz@cs.ucsb.edu
Office Hours: Monday, 2:00pm-4:00

Student Evaluatio

http://www.cs.ucsb.edu/~almeroth/classes/F97.176/

Paper 60 (036 Opp. to Motion to Exclude) at 6-7; Ex. 1234.003 (Almeroth Dep., Ex. 21).

Q. And then the third page of Exhibit 21. What is that?
A. That's the course outline.
Q. And it includes dates. Are those dates of individual classes?
A. They are.

Paper 41 (036 Reply) at 10 n.5; Ex. 1225 (Almeroth Depo.) at 475:3-8.

PO file histories admit Tanenbaum96 was publicly available

- In IDS for each patent, PO identified Tanenbaum96's date of publication as 1996 under 37 CFR § 1.98

265	Andrew S. Tanenbaum, Computer Networks, Third Edition, 1996, ISBN 0-13-349945-6.

Paper 64 (072 Opp. to Motion to Exclude) at 13; Ex. 1002.270 (072 File History).

N	Andrew S. Tanenbaum, "Computer Networks," Third Edition, 1996, ISBN 0-13-349945-6.

Paper 62 (241 Opp. to Motion to Exclude) at 13; Ex. 1002.270 (241 File History).

O	Andrew S. Tanenbaum, Computer Networks, Third Edition, 1996, ISBN 0-13-349945-6.

Paper 59 (880 1409 Opp. to Motion to Exclude) at 13; Ex. 1002.268 (880 File History).

Andrew S. Tanenbaum, "Computer Networks," Third Edition, 1996, ISBN 0-13-349945-6.

Paper 60 (036 Opp. to Motion to Exclude) at 13; Ex. 1001.003 (036 Patent).

PO patents describe Tanenbaum96 as a college-level textbook

(12) **United States Patent**
Boucher et al.

(10) **Patent No.: US 7,237,036 B2**
(45) **Date of Patent: Jun. 26, 2007**

(54) **FAST-PATH APPARATUS FOR RECEIVING DATA CORRESPONDING A TCP CONNECTION**

(75) Inventors: **Laurence B. Boucher**, Saratoga, CA (US); **Stephen E. J. Blightman**, San Jose, CA (US); **Peter K. Craft**, San Francisco, CA (US); **David A. Higgen**, Saratoga, CA (US); **Clive M. Philbrick**, San Jose, CA (US); **Daryl D. Starr**

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,366,538 A 12/1982 Johnson et al. 364/200

(Continued)

FOREIGN PATENT DOCUMENTS

WO WO/98/19412 5/1998

(Continued)

(51) **Int. Cl.**
G06F 13/38 (2006.01)
G06F 15/17 (2006.01)

(52) **U.S. Cl.** **709/245**; 709/236; 709/230; 370/474; 370/396; 370/469

(58) **Field of Classification Search** 709/245, 709/236, 230, 202; 370/474, 230, 396, 469; 707/2–4, 10; 712/19, 52

See application file for complete search history.

passed back to the host for message processing by the host. The INIC contains specialized hardware circuits that are much faster at their specific tasks than a general purpose CPU. A preferred embodiment includes a trio of pipelined processors with separate processors devoted to transmit, receive and management processing, with full duplex communication for four fast Ethernet nodes.

22 Claims, 89 Drawing Sheets

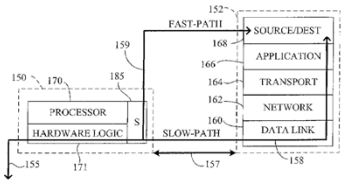

> The above description of layered protocol processing is simplified, as college-level textbooks devoted primarily to this subject are available, such as Computer Networks, Third Edition (1996) by Andrew S. Tanenbaum, which is incor-

Paper 60 (036 Opp. to Motion to Exclude) at 13;
Ex. 1001 (036 Patent) at 4:47-50.

Argument summary

1. It is improper to raise public availability in a Motion to Exclude

2. Patent Owner admits Tanenbaum96 was publicly available

3. **Librarian declaration establishes Tanenbaum96 was publicly available**

Majors declaration establishes Tanenbaum96 was publicly available

- Library acquired the book **from the publisher** by November 1, 1996

> 4. When a monograph is received and cataloged by the SCU Library, the date of cataloging is set and retained in the catalog record. The catalog date ("Cat Date") for Tanenbaum is November 1, 1996 (see Exhibit A). The volume would have been available for use within a few weeks of that date.

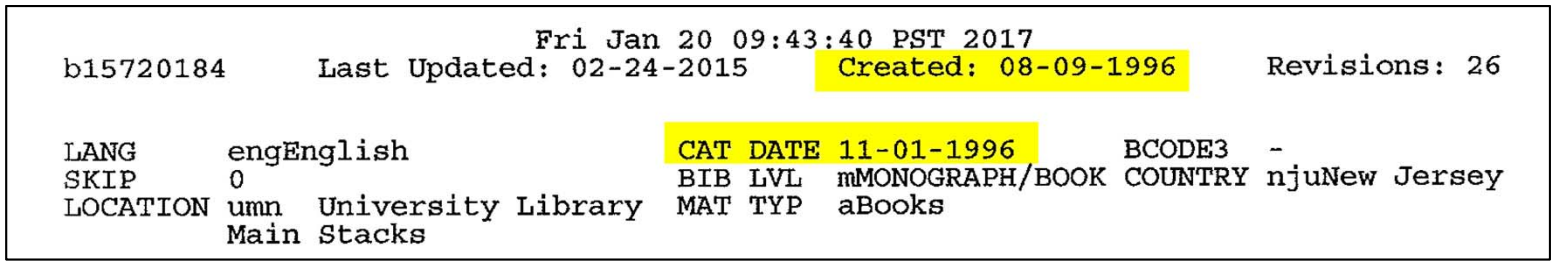

```
                                Fri Jan 20 09:43:40 PST 2017
b15720184      Last Updated: 02-24-2015      Created: 08-09-1996      Revisions: 26

LANG      engEnglish                    CAT DATE 11-01-1996      BCODE3   -
SKIP      0                             BIB LVL  mMONOGRAPH/BOOK COUNTRY njuNew Jersey
LOCATION  umn  University Library       MAT TYP  aBooks
          Main Stacks
```

Paper 60 (036 Opp. to Motion to Exclude) at 7-8, 12;
Ex. 1011.001, -.003 (Majors Declaration).

Majors declaration is admissible

- Made on declarant's personal knowledge of library procedures dating back to 1992
- Explains that Tanenbaum96 was received on August 9, 1996 and catalogued on November 1, 1996
- Attaches catalogue record and signed under the penalty of perjury

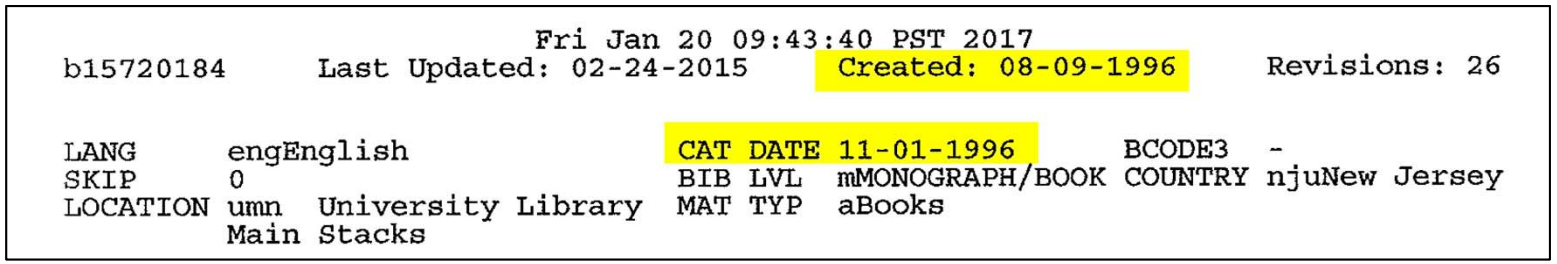

```
                              Fri Jan 20 09:43:40 PST 2017
b15720184      Last Updated: 02-24-2015      Created: 08-09-1996      Revisions: 26

LANG      engEnglish                    CAT DATE 11-01-1996      BCODE3  -
SKIP      0                             BIB LVL  mMONOGRAPH/BOOK COUNTRY njuNew Jersey
LOCATION  umn  University Library       MAT TYP  aBooks
          Main Stacks
```

Paper 60 (036 Opp. to Motion to Exclude) at 7-8;
Ex. 1011.003 (Majors Declaration)

Authenticity objection to catalog record is untimely and unfounded

- PO waived objection because it did not object to Majors Declaration or Catalog Record as being unauthentic

- Majors Declaration authenticates Catalog Record:

> 4. When a monograph is received and cataloged by the SCU Library, the date of cataloging is set and retained in the catalog record. The catalog date ("Cat Date") for Tanenbaum is November 1, 1996 (see Exhibit A). The volume would

Paper 60 (036 Opp. to Motion to Exclude) at 8-9;
Ex. 1011.001 (Majors Declaration).

Catalog record falls under hearsay exceptions

- Majors Declaration:

> 4. When a monograph is received and cataloged by the SCU Library, the date of cataloging is set and retained in the catalog record. The catalog date ("Cat Date") for Tanenbaum is November 1, 1996 (see Exhibit A). The volume would

- Establishes it falls under exceptions to hearsay
  - Rule 803(6) – Records of regularly conducted activity, as established by the Majors Declaration and description of policies and procedures of the library
  - Rule 807(a) – Provided by a University Library, offered to show public availability, more probative on public availability than other evidence that can be reasonably obtained, and is in the interests of justice

Paper 60 (036 Opp. to Motion to Exclude) at 8-9;
Ex. 1011.001 (Majors Declaration).