Ltd. 1995), pages 224-239.

I, Lisa Rowlison de Ortiz, declare:

1. I am the Head of Catalog & Metadata Services at University of California, Berkeley ("UC Berkeley") library. I am familiar with the UC Berkeley library system, including the library catalog and policies and procedures regarding the receipt, indexing, and availability of books and periodicals.

2. According to UC Berkeley Library policies and procedures, Library items are indexed in the library catalog and are made freely available to the faculty and student body of UC Berkeley as well as to the general public.

3. The UC Berkeley library holds a copy of a chapter by Y.H. Thia and C.M. Woodside, "A Reduced Operation Protocol Engine (ROPE) for a multiple-layer bypass architecture," published in the book Protocol for High Speed Networks IV, 1st Edition (TJ Press Ltd. 1995), pages 224-239. ("Thia").

4. When a monograph is received and cataloged by the UC Berkeley Library, the date of cataloging is set and retained in the catalog record. The catalog date ("Cat Date") for Thia is February 26, 1996 (*see* Exhibit A). Furthermore, after the volume is labeled and sent to its shelving location the date of receipt by this shelving location is stored in an internal note. This information shows that the volume was received by the Engineering Library on March 20, 1996. *Id.* The volume would have been available to the public within a few days of that date.

thereon. I declare that all statements made of my knowledge are true, and that all statements made on information and belief are believed to be true.

Executed on January 27, 2017, in Berkeley, California.

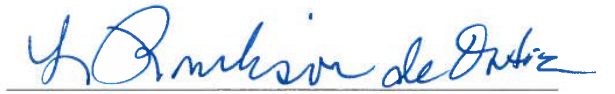

Lisa Rowlison de Ortiz

Protocols
High Spee
Networks

Edited by

Gerald Neufield and M
Department of Computer Science
University of British Columbia
Vancouver
Canada

Sponsored by IFIP WG6.1/WG6.4 in co
IEEE Com Soc

Published by Chapman & Hall on beha
International Federation for Information

CHAPMAN & HA
London · Glasgow · Wei
Melbourne · Madras

PREFACE

Welcome to the fourth IFIP workshop on protocols for high
workshop follows three very successful workshops held in 2
Stockholm (1993) respectively. We received a large number
contributions. This year, forty papers were received of w
papers and four were presented as poster papers. Although
the program committee decided to keep the number of
accommodate more discussion in keeping with the format of a

Many people have contributed to the success of this works
program committee who, with the additional reviewers, help
We are thankful to all the authors of the papers that were
organizations which have contributed financially to this
CICSR, UBC, MPR Teltech and Newbridge Networks.

Mabo Ito
Gerald Neufeld

Departments of Electrical Engineering and Computer Science
University of British Columbia

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 (**)

Abstract — The Reduced Operation Protocol Engine (ROPE) presented here offloads critical functions of a multiple-layer protocol stack, based on the "bypass concept" of a fast path for data transfer. The motivation for identifying this separate processing path is that it involves only a small subset of the complete protocol, which can then be implemented in hardware. Multiple-layer bypass also eliminates some inter-layer operations such as queue and buffer management, context switching and movement of data across layers, all of which are a significant overhead. ROPE is intended to support high-speed bulk data transfer. The paper describes the design of a ROPE chip for the OSI Session and Transport layer protocols, using VHDL. The design is practical in terms of chip complexity and area, using current gate array technology, and simulation shows that it can support a data rate approaching 1 gigabit 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 bandwidth and low bit error 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

- Non-protocol-specific processing is a large part of the ples include interrupt handling, context switching a in deeply layered protocol stacks.
- The choice of hardware for the adaptor depends or supports. In [2, 22] where the transport protocol la full protocol stack can be offloaded, general purpose because of the complexity of existing protocols, V data link layer has been disappointing so far. In [8 support TCP checksums. Also, some newer lightwei designed for VLSI implementation [1, 3].
- There is a tradeoff between performance, flexibility the frequently executed portion of the protocol ren significant advantage in providing hardware support tasks in the host software for flexibility.
- As host processing speed continues to outpace mem bandwidth approaches the processor memory bandv movement on the workstations down to the minimur

This paper presents a feasibility study for a new ap combines the relatively simple operations needed for data provides a hardware "fast path" for them, which will be based on the "protocol bypass concept" [37] which is a g Prediction" algorithm [20] for TCP/IP. Bypass solves th may limit the use of offboard processing, by implemer layers for certain cases. This simplifies the interface bet and minimizes their interaction, which is supported by an and a simple command protocol. The chip design based Reduced Operation Protocol Engine. The contribution of interface and the chip operation, and to report on a VI chip design. It appears to be feasible to support an end- approaching 1 Gbps.

The next section introduces the bypass concept, its Section 3 analyzes the key protocol processing overhea for a bypass VLSI implementation. Sections 4, 5 and 6 c chip using the industry standard hardware description lan Section 7.

2 The Bypass Concept

A bypass adds an additional path for certain operati original software. Conformance to the protocol is maintain through the normal "heavyweight" path. A bypass path can or for both together, and is compatible with other end-syste

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS
Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS
Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS
Sync your system to PACER to automate legal marketing.