Protocols for High-Speed Networks VI

PROTOCOLS FOR HIGH-SPEED NETWORKS VI

IFIP TC6 WG6.1 & WG6.4 /
IEEE ComSoc TC on Gigabit Networking
Sixth International Workshop on
Protocols for High-Speed Networks (PfHSN '99)
August 25-27, 1999, Salem, Massachusetts, USA

Edited by

Joseph D. Touch
University of Southern California,
Information Sciences Institute
USA

James P.G. Sterbenz
BBN Technologies, GTE
USA

KLUWER ACADEMIC PUBLISHERS
BOSTON / DORDRECHT / LONDON

Distributors for North, Central and South America:
Kluwer Academic Publishers
101 Philip Drive
Assinippi Park
Norwell, Massachusetts 02061 USA
Telephone (781) 871-6600
Fax (781) 871-6528
E-Mail <kluwer@wkap.com>

Distributors for all other countries:
Kluwer Academic Publishers Group
Distribution Centre
Post Office Box 322
3300 AH Dordrecht, THE NETHERLANDS
Telephone 31 78 6392 392
Fax 31 78 6546 474
E-Mail <services@wkap.nl>


Electronic Services <http://www.wkap.nl>

Library of Congress Cataloging-in-Publication Data

Printed on acid-free paper.

Printed in the United States of America.

PROTOCOLS FOR HIGH-SPEED NETWORKS VI

IFIP - The International Federation for Information Processing

IFIP was founded in 1960 under the auspices of UNESCO, following the First World Computer Congress held in Paris the previous year. An umbrella organization for societies working in information processing, IFIP's aim is two-fold: to support information processing within its member countries and to encourage technology transfer to developing nations. As its mission statement clearly states,

IFIP's mission is to be the leading, truly international, apolitical organization which encourages and assists in the development, exploitation and application of information technology for the benefit of all people.

IFIP is a non-profitmaking organization, run almost solely by 2500 volunteers. It operates through a number of technical committees, which organize events and publications. IFIP's events range from an international congress to local seminars, but the most important are:

- The IFIP World Computer Congress, held every second year;
- open conferences;
- working conferences.

The flagship event is the IFIP World Computer Congress, at which both invited and contributed papers are presented. Contributed papers are rigorously refereed and the rejection rate is high.

As with the Congress, participation in the open conferences is open to all and papers may be invited or submitted. Again, submitted papers are stringently refereed.

The working conferences are structured differently. They are usually run by a working group and attendance is small and by invitation only. Their purpose is to create an atmosphere conducive to innovation and development. Refereeing is less rigorous and papers are subjected to extensive group discussion.

Publications arising from IFIP events vary. The papers presented at the IFIP World Computer Congress and at open conferences are published as conference proceedings, while the results of the working conferences are often published as collections of selected and edited papers.

Any national society whose primary activity is in information may apply to become a full member of IFIP, although full membership is restricted to one society per country. Full members are entitled to vote at the annual General Assembly, National societies preferring a less committed involvement may apply for associate or corresponding membership. Associate members enjoy the same benefits as full members, but without voting rights. Corresponding members are not represented in IFIP bodies. Affiliated membership is open to non-national societies, and individual and honorary membership schemes are also offered.

COMMITTEE MEMBERS

WORKSHOP CO-CHAIRS:

James P. G. Sterbenz	*BBN Technologies, GTE, U.S.A.*
Joseph D. Touch	*USC / ISI, U.S.A.*

PROGRAM COMMITTEE:

Sujata Banerjee	*Univ. of Pittsburgh, U.S.A.*
Ernst Biersack	*Eurecom, France*
Gregor von Bochmann	*University of Ottawa, Canada*
Jon Crowcroft	*UCL, United Kingdom*
Christophe Diot	*Sprint Advanced Technology Lab, U.S.A.*
Julio Escobar	*SENACYT, Panama*
Per Gunningberg	*Uppsala University, Sweden*
Marjory Johnson	*RIACS/NASA Ames Research Center, U.S.A.*
Guy Leduc	*University of Liege, Belgium*
Bryan Lyles	*Sprint Advanced Technology Lab, U.S.A.*
Gerald Neufeld	*Siara Systems, U.S.A.*
Craig Partridge	*BBN Technologies, GTE, U.S.A.*
Harry Rudin	*IBM Zurich Research Lab, Switzerland*
Douglas C. Schmidt	*Washington University, U.S.A.*
Ahmed Tantawy	*IBM, U.S.A.*
Martina Zitterbart	*Tech. University Braunschweig, Germany*

PREFACE

This year marks the 10th anniversary of the IFIP International Workshop on Protocols for High-Speed Networks (PfHSN). It began in May 1989, on a hillside overlooking Lake Zurich in Switzerland, and arrives now in Salem Massachusetts 6,000 kilometers away and 10 years later, in its sixth incarnation, but still with a waterfront view (the Atlantic Ocean). In between, it has visited some picturesque views of other lakes and bays of the world: Palo Alto (1990 – San Francisco Bay), Stockholm (1993 – Baltic Sea), Vancouver (1994 – the Strait of Georgia and the Pacific Ocean), and Sophia-Antipolis / Nice (1996 – the Mediterranean Sea).

PfHSN is a workshop providing an international forum for the exchange of information on high-speed networks. It is a relatively small workshop, limited to 80 participants or less, to encourage lively discussion and the active participation of all attendees. A significant component of the workshop is interactive in nature, with a long history of significant time reserved for discussions. This was enhanced in 1996 by Christophe Diot and Wallid Dabbous with the institution of Working Sessions chaired by an "animator," who is a distinguished researcher focusing on topical issues of the day. These sessions are an audience participation event, and are one of the things that makes PfHSN a true "working conference."

This year, we received twice as many papers as we could accommodate, and have a record number of Working Sessions, the latter including charging and accounting, multicast, wireless, as well as a retrospective of this workshop. We have half as many countries represented as we have papers, making this a truly international year for PfHSN.

During the past ten years we have seen high-speed networking grow from a niche research area to commodity commercial products, such as Gigabit Ethernet, OC-192 SONET, and ATM. PfHSN has tracked the evolution of high-speed network research, from protocol concepts to their implementation, from applications to switching, from OSI to TCP and ATM. On the way, we have visited a diverse set of research areas, including parallelization, copy reduction, hardware protocols, congestion control, host–network interface design, video, quality of service, and layer integration.

Harry Rudin and Robin Williamson brought us that first workshop, and in their preface hoped that we would be able to discuss issues in designing and implementing high-level protocols, to bring the promise of high-speed networking to fruition. As we sit at our desktops, using 100 Mbps and Gbps commercial LANs connected by Gbps service provider connections, it is clear that their vision has become real. It is also clear that the challenges of

high-performance networking are ever more present, as the World-Wide Web continues to be nick-named the "World-Wide Wait."

The issue, raised ten years ago, was whether we need new protocols to support high-speed networks, or whether clever tuning would suffice. Ten years of research has not completely answered that question. We do know that networks will continue to get faster, that protocols will continue to provide richer capabilities, and that the trade-offs will shift, as technology evolves. So it goes with all computing.

Five years later, when we were in the midst of the ATM hype, James Sterbenz presented the keynote address which questioned *Protocols for High Speed Networks: Life after ATM?* ATM had taken experimental fast packet switching research in the mid 1980's, to a set of standards prematurely cast with serious technical flaws, to a bandwagon on which both researchers and vendors were jumping and thinking that ATM *was* high speed networking. Even before the ATM bubble burst, PfHSN remained focused on high-speed networking in general, rather than any particular technology in particular (such as ATM). In fact, of the PfHSN papers from the mid 90's which did concentrate on ATM, a substantial number were focused on the problems with ATM, both in isolation and as a second network layer under IP.

This workshop owes its existence and its longevity to Harry Rudin, who continued to shepherd it through IFIP as chair of WG 6.1, until passing the torch to Guy Leduc earlier this year. We thank them, as well as the past Workshop Co-Chairs Harry Rudin, Robin Williamson, Marjory Johnson, Per Gunningberg, Björn Pehrson, Stephen Pink, Mabo Ito, Gerald Neufeld, Wallid Dabbous, and Christophe Diot for the initiative and work to keep PfHSN going for ten years. We would also like to thank the authors, and all the participants, for making this workshop productive and worthwhile. We thank the Program Committee for taking the time to personally review papers and for their collaboration in creating this year's Working Sessions. We thank the authors for their interest, active participation, and assistance with the process of publishing their work herein. We thank Jeanine Yamazaki and Sungita Patel for their help in pulling PfHSN '99 all together.

Finally, some notes about the future: High-speed networking has become a pervasive aspect of network research, partly due to the success of this workshop. IFIP's Conference on High-Performance Networks (HPN), and IEEE's Gigabit Networking Workshop (GBN), to name only a few. With such success comes an opportunity: IFIP is considering a merger of PfHSN IFIP Conference on High-Performance Networking (HPN). Part of this merger will involve a name shift, some unification of the titles of both. But we hope it will retain part of the unique character of PfHSN, particularly the animated Working Session, and the open forum for discussion it provides. These sessions, like orchestral performances, are conducted by a

distinguished person, but are created only as the result of the participation and collaboration of the entire group. We hope these sessions, which we feel embody the best aspects of PfHSN, will continue, whatever the name of the forum.

In this light, the current volume contains a full contents and author index of the last ten years of PfHSN workshop. We hope that this will serve not only as a useful index to past PfHSNs, but also encourage people to look back at the evolution of high speed networking research over the last ten years and think about the future.

While the end-to-end and higher layer issues remain as elusive as ever for high speed networking, a new set of challenges at the lower layers are emerging. These not only motivate new research at the lover layers, but continue to drive research at the end-system, end-to-end protocols, middleware, and the application interface. Some of these new challenges include:

1. Wireless networks are of increasing importance in the arena of smart spaces and wearable networks. The scale and density of networks in which each individual may have 10 to 100 nodes is daunting (consider a stadium with 100 000 people and 10 000 000 networked devices in range). There is no reason to believe that the demand for bandwidth will not vastly exceed what is readily available in the radio spectrum.

2. Optical technology is finally maturing to the degree that the prospect of all-optical portions of the Internet are conceivable, and are worthy of research. There are significant issues remaining, not only in how to do all-optical networking, but how to do this in the context of the Internet as we know it.

3. We are now considering how to extend the Internet to Mars for the upcoming Mars space missions; it is only sensible to do this in the context of how to build an Interplanetary Internet. Not only do latencies on the order of 10 to 100 minutes cause us to rethink how we can do interactive applications because the usual tricks to effectively beat the speed of light break down, but current protocols simply break (for example timers and end-to-end state synchronization). While space bandwidths are currently low relative to fiber standards, the hint of free space laser communications causes us to consider the prospect of bandwidth-×-delay products of petabits or exabits in flight.

So, just as we said at the first PfHSN in 1989, and mid-decade in 1994, there is much interesting research to be done. We will continue to struggle with the issue Harry Rudin and Robin Williamson raised in 1989: whether new protocols are needed or clever re-engineering of current protocols will suffice, and it will likely continue to be a combination of the two.

Finally, two things have happened in the high-speed networking community: First, an increasing amount of networking research involves

performance as a fundamental aspect, and thus researchers who would have not originally identified themselves as part of the high-speed community are concerned with these issues on a daily basis. Second, many of us who were focused on high-speed networking *per se* have moved on to orthogonal disciplines, of which high-speed networking is a significant component. And while we suffer from the problem that "nothing is *just* high speed networking, and everything is high speed networking", there remains a community of people who are deeply interested in high-speed networking as a discipline. Hopefully, we will continue to be able to meet and work at forums such as PfHSN.

James P. G. Sterbenz

Joseph D. Touch

CONTENTS

PART 1

KEYNOTE AND WORKING SESSIONS

THE OPTICAL FUTURE OF HIGH-SPEED NETWORKS

C. Qiao
State University of New York, Buffallo

Abstract:

With the ever increasing optical transmission rate, now exceeding 1 Tb/s on a single fiber thanks to Wavelength Division Multiplexing (WDM) technologies, electronic routing/switching is quickly becoming a performance bottleneck in high-speed backbone networks. WDM optical networking, though rooting in the physical and link layers, is by no means yet another ordinary low layer technology as it may affect the designs of the upper electronic layer. Specifically, WDM optical networks can be configured to bypass intermediate electronic components by switching/routing data in the optical domain. This will not only reduce electronic processing and I/O loads, but also provide bit-rate and coding format transparency. However, this may also lead to unexpected "shortest" and "alternate" paths in the electronic layer. In addition, primitive optical logic, and especially the lack of optical memory (buffer) are major challenges in order to realize the vision of building a bandwidth-abundant infrastructure, which is ubiquitous and yet efficient, based on WDM optical networks. In this talk, I will describe these issues and along relevant optical switching paradigms, namely wavelength-routing (as a form of optical circuit switching), optical packet switching, and optical burst switching (OBS), and discuss how the next generation Optical Internet may support QoS and provide multicast services in the WDM optical layer.

Chunming Qiao earned his BS degree in Computer Science and Engineering from University of Science and Technology of China (USTC) in Hefei, P.R. China. He received the Andrew-Mellon Distinguished doctoral fellowship award from University of Pittsburgh and later earned his Ph.D. degree in Computer Science there in 1993. He was an Assistant Professor in the ECE Department, and now an Associate Professor in the newly created CSE Department at SUNY at Buffalo. His research interests cover the two converging areas of computers and communications and currently, he is conducting research on Wavelength Division Multiplexed (WDM) networks and internetworks (e.g. IP over WDM) which has been funded by NSF and Telcordia Technologies (formerly Bellcore). Dr. Qiao has published more than 50 papers in IEEE Trans. on Computers, Trans on Parallel and Distributed Systems, Trans. on Communications and Trans. on Networking as well as other journals and conference proceedings. He has served as a Co-Chair for the annual All-optical Networking conference since 1997,. a Program Vice Co-Chair for the 1998 International Conf. on Computer Communications and Networks (IC3N), panel/session organizer at the IEEE Globecom'99 and IEEE MILCOM'96, and program committee members and session chairs in over a dozen of other conferences and workshops. He has been featured in a couple of interviews shown in the Alcatel's Communications Week and SPIE's OE Reports. He is also an editor of the Journal on High-Speed Networks (JHSN) and the new Optical Networks Magazine, and a member of IEEE Computer Society and IEEE Communications Society.

PfHSN '99 WORKING SESSIONS

The working session is a key feature of Protocols for High-Speed Networks. The sessions for 1999 are described below.

PROTOCOL ISSUES FOR CHARGING, PRICING, AND ACCOUNTING

Animator: Bryan Lyles, Sprint Advanced Technology Lab, U.S.A.

MULTICAST ISSUES

Animator: Marjory Johnson, RIACS/NASA Ames Research Center, U.S.A.
Panelists: Radia Perlman, Sun Microsystems, U.S.A.
Kevin Almeroth, U.C. Santa Barbara, U.S.A.

10 YEARS OF PFHSN: WHERE HAVE WE BEEN, WHERE ARE WE GOING?

Animators: James P. G. Sterbenz, BBN Technologies, GTE, U.S.A.
Joe Touch, USC/ISI, U.S.A.

WIRELESS ISSUES

Animator: Julio Escobar, SENACYT, Panama