Internet Engineering Task Force
INTERNET-DRAFT
draft-sarolahti-tsvwg-crosslayer-00.ps
Expires: April 2007

P. Sarolahti
Nokia Research Center
S. Floyd
ICIR

15 October 2006

Cross-layer Indications for Transport Protocols

Status of this Memo

Abstract

This document discusses different types of explicit cross-layer signaling and notification mechanisms that have been proposed to enhance end-to-end transport performance. We analyze the different mechanisms in a framework based on what level of network support they require, and discuss the benefits and disadvantages different approaches have. The objective is to get a common understanding of the problem space, with pointers to past discussions on this topic, and to describe the possible next steps towards removing barriers from explicit cross-layer communication in future protocols.

Table of Contents

1. Introduction

In the past there have been different proposals on enhancing transport protocol (usually TCP) performance by means of explicit cross-layer signaling. The cross-layer signaling can be local notifications from the lower or upper protocol layers of the host device, or it can be explicit communication between the transport peers and the network between them. The mechanisms for cross-layer signaling inside a host implementation are largely dependent on the operating system architecture, and therefore not of interest for the IETF. However, the explicit signaling between the network and the end-hosts involves several considerations on the network behavior that we try to capture in this document.

Cross-layer signaling could be used, for example, for delivering hints to a TCP sender about the characteristics of the network path, to allow the sender to adjust its sending rate more efficiently than what would be possible using the traditional TCP probing mechanisms. While designing the possible uses of such signaling, a careful consideration needs to be made of what can be done within the limits of the congestion control principles [RFC2914, RFC2581], without endangering the network stability and fairness towards other flows. Often this determines whether end-hosts can negotiate directly without network support, or whether some or all of the routers along the network path need to support the signaling mechanism.

One of the guiding architectural principles of the Internet has been that the network should be stateless, with the transmission state and intelligence residing at the end hosts [Cla88].

Although today this principle has been ignored more than once by the different types of Network Address Translators (NATs) and stateful firewalls, it is an important consideration when evaluating the cross-layer signaling methods. While many of the signaling mechanisms discussed in this document conform to this principle, others do require some additional state in the network. Adding new bits of state in the network is not necessarily a bad thing, but the design should be such that loss of the state would not cause serious fate-sharing problems that might prevent the network's packet forwarding function from working.

This document casts an overview on different types of explicit cross-layer signaling, and discusses their possibilities and challenges. The objective of the final document is to get a common understanding of the problem space and to describe the possible next steps towards removing barriers from explicit cross-layer communication in future protocols.

1.1. Roles of different protocol layers

It is difficult to find a course book on computer networks that would not begin with a description of the different protocol layers, usually according to the ISO reference model. For example, [Hal96, Figure 1.11] summarizes the different protocol layers as follows:

* Physical layer (1): Mechanical and electrical network interface definitions.

* Link layer (2): Data link control (framing, data transparency, error control).

* Network layer (3): Network routing, addressing, call set-up, and clearing.

* Transport layer (4): End-to-end message transfer (connection management, error control, fragmentation, flow control).

* Session layer (5): Dialog and synchronization control for application entities.

* Presentation layer (6): Transfer syntax negotiation, data representation transformations.

* Application layer (7): File transfer, access and management, document and message interchange, job transfer and manipulation-

  Although ISO reference model layering is not explicitly visible in many of the IETF protocols, and some protocols might do tasks of more than one layer, it is possible to find places of different IETF protocols in this model. Usually the proposed cross-layer enhancements concern interactions between the link layer, network layer, and transport layer.

1.2. Conventions and Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Possible Benefits of Explicit Signaling

The past proposals to add new explicit signaling mechanisms have been motivated in the following ways.

* Mobility: One of the key design principles of the IP mobility protocols has been to isolate the mobility from the upper protocol layers. While this makes sense architecturally, it has been observed that hiding the mobility event from upper layer protocols can lead to suboptimal performance. A significant amount of research has been conducted to investigate and optimize the performance of the upper layers after a mobile hand-off occurs (e.g., [SKDK06]). It has been observed that with better awareness of mobility, benefits on transport protocol performance can be achieved.

* High delay-bandwidth networks: TCP slow-start is known to be inefficient when used over a very high-speed network path, or over a network path with large propagation delay. The TCP startup performance could be improved with explicit information about the current available capacity of the connection path [SAF06, KHR02].

* Non-congestion losses: a traditional research topic on TCP performance has been detection and response to packet reordering and to packet losses that are not caused by congestion. Possible performance benefits for being able to detect non-congestion losses have been evaluated in a number of documents [KSE+04].

* Quality of Service: In addition to DiffServ and RSVP QoS signaling, there have been other proposals for smaller enhancements regarding real-time transmission of packets. For example, Packet Lifetime Discard [GL03] assigns a lifetime for IP packets to allow routers to drop packets that have exceeded their lifetime, thereby saving network resources. An earlier work proposed adjusting the link-layer reliability mode in GSM networks based on the transport protocol in use [LR99], to allow more timely handling of UDP packets.

3. Classification of Explicit Signaling Mechanisms

We discuss a number of explicit signaling mechanisms by placing them into the following catgories. Each of the classes below has some common characteristics that determine how powerful the given type of signaling can be, and what kind of challenges can be expected.

A high-level classification is made to in-band and out-of-band signaling mechanisms. In-band signaling goes with the normal data traffic. The benefit is that in-band messages are known to share the same path as the data, and generally use less header overhead than a separate message. In addition, with in-band messages the sender is often able to know about a loss event via the transport protocol's normal acknowledgment mechanisms. The disadvantage of in-band signaling is that if some routers or middleboxes drop a packet because of an unknown option, the accompanying data also gets lost, causing a performance penalty for the data transfer.

* In-band signaling without network involvement.

  There have been several proposals to place information inside the transport header, for example using a TCP option. The information flows between the end hosts, but should not be read or modified by the network routers, because routers do not typically investigate transport headers. Furthermore, IPsec may prevent the routers from reading the option altogether.

* In-band signaling requiring support from some routers along the path.
  Typically mechanisms in this class use the IP header in one way or another; either through using the reserved (nowadays the DiffServ) bits in the IP header, or using an IP option. If the mechanism is such that it is useful without all routers supporting it, the deployment seems realistic: even long network paths can be supported and IP tunnels do not defeat the mechanism. ECN (Explicit Congestion Notification) [RFC3168] is one example of in-band signaling requiring support from some routers along a path.

* In-band signaling requiring support from all routers along the path.
  If all routers need to support an IP-level option, the deployment becomes challenging. It is difficult to get universal support outside the host's own network domain. Verifying that all routers have indeed processed an explicit notification is also difficult in a network environment that might consist of IP tunnels, different types of middleboxes and other types of special network equipment. Quick-Start [FAJS06] is an example of in-band signaling requiring support from all routers along a path.

* Out-of-band signaling without network involvement.
  ICMP is one of the oldest out-of-band signaling methods in the Internet. If an incoming packet causes an error condition, for example from using a protocol not supported at the receiver, the receiver generates an ICMP error message. Usually a portion of the transport header is included in the ICMP payload as additional information for the sender.

* Out-of-band signaling requiring support from some routers along the path.
  Network routers can also generate ICMP messages, for example, to implement Path MTU discovery by sending ICMP errors for packets exceeding link MTU. The functionality is similar to the case of an end-host generating an ICMP message. However, if IPsec encryption is employed, it may be impossible for the ICMP message to include portions of the transport header.

* Out-of-band signaling requiring support from all routers along the path.
  This class of mechanisms uses dedicated packets that are processed at the router. The Resource ReserVation Protocol (RSVP) [RFC2205] is one such mechanism. Typically a router alert IP option is needed to notify the router forwarding function that the payload of the IP packet contains data to be processed by the router.

4. Proposed Explicit Cross-layer Mechanisms

In this section ee discuss some of the explicit cross-layer signaling mechanisms that have been proposed in the past.

4.1. In-band signaling without network involvement

TCP's response to connectivity change indications such as mobility have been discussed in an Internet-Draft [SEE+06]. The draft describes a "connectivity-change indication" TCP option. The option could be used by a mobile receiver to indicate the other end that it has moved and the congestion control characteristics of a path should be re-evaluated. When connectivity is re-established after a short disruption, the option could also be used by the receiver that detects re-establishment of link connectivity to trigger a fast retransmission at the transmitting end. The TCP option would not require support from network, but it would need to be supported by both ends of the connection.

4.2. In-band signaling requiring support from some routers along the path

Explicit Congestion Notification (ECN) [RFC3168] uses a two-bit ECN field in the IP header to allow routers to indicate congestion in the network before they have to start dropping packets due to buffer overflow. ECN can be useful even if only a subset of routers implement it on the connection path. There were initial deployment problems with ECN because some routers in the network dropped packets with a non-zero ECN field, but we believe that today most of these routers have been fixed.

The use of the ECN field is taken further in an alternative protocol to use the field, called Re-ECN [BJSK06]. The protocol aims "to provide sufficient information in each IP datagram to be able to hold senders and whole networks accountable for the congestion they cause downstream, before they cause it."

Different forms of in-band signaling have also been proposed for dealing with corruption-based packet loss in wireless and satellite networks [KSE+04]. The paper on Explicit Transport Error Notification gives a taxonomy of different notification types, depending on the granularity of the notification, the direction of notification, the location of notification, and so on. The efficiency of cumulative error notification is investigated by simulation experiments. However, no specific packet format is proposed in the paper.

4.3. In-band signaling requiring support from all routers along the path

In Quick-Start [FAJS06], the sender uses an IP option to request permission from routers to send at a higher rate than the normal congestion control would allow. [FAJS06] specifies use of Quick-Start for TCP and discusses the challenges such a mechanism needs to address. Quick-Start router algorithms and their configuration are analyzed further in [SAF06], and [SKDK06] gives an initial analysis of Quick-Start in wireless environments with vertical hand-offs between different wireless link technologies.

IP tunnels impose a serious challenge for Quick-Start and other similar mechanisms, as discussed in [FAJS06]. We discuss IP tunnels further in Section 6.2 of this document. Another problem with deploying new IP options is that some routers or middle-boxes in the network silently drop packets containing known or unknown IP options [MAF04]. It has not been identified which network devices do this, and on which basis.

IP options has also been proposed for Path MTU Discovery [Wel03]. The design principles of IP-option-based Path MTU Discovery are roughly similar to Quick-Start, as are the challenges it faces.

Explicit Control Protocol (XCP) [KHR02] is a proposal for a full-fledged congestion control protocol involving the interaction of routers and the end-hosts. Although XCP can be considered to be more than just a cross-layer signaling mechanism, it also needs to consider the above-mentioned challenges.

Variable-structure congestion Control Protocol (VCP) is another proposed congestion control proposal using explicit feedback from routers. VCP leverages the ECN bits to let routers indicate their load information [XSSK05]. Based on the VCP bits, a TCP sender could apply either Multiplicative Increase, Additive Increase, or Multiplicative Decrease of the congestion window.

4.4. Out-of-band signaling without network involvement

While the Internet Control Message Protocol (ICMP) [RFC792, RFC2463] is not a cross-layer signaling mechanism but rather an IP-layer signaling mechanism, it is the Internet's basic mechanism for communicating error conditions and other notifications to the end hosts, and could be used as a bearer for cross-layer notifications. Many ICMP errors, such as "Protocol unreachable", originate from the intended destination of the IP packet. Part of the IP payload is included in the ICMP message to help the sender identify the transport protocol flow that caused the error, and to provide means of a basic (but rather weak) check of authenticity of the ICMP message. ICMP is an unreliable mechanism: if the ICMP message is lost in the network, neither the sender or the receiver gets a direct indication of the loss.

4.5. Out-of-band signaling requiring support from some routers along the path.

The ICMP messages can also originate from the network, as with a "Destination unreachable" message. The router can include part of the IP payload in the message, but if the original data packet is tunneled, for example, in an IPsec tunnel, the payload cannot be used to identify the transport protocol flow, and the sender cannot do much about checking the authenticity of the ICMP message.

4.6. Out-of-band signaling requiring support from all routers along the path.

The Resource ReSerVation Protocol (RSVP) uses separate out-of-band messages on top of IPv4 or IPv6 to make Quality-of-Service signaling [RFC2205]. The data sender sends a RSVP "Path" message to the data receiver that includes a Router Alert IP option telling the routers on the path to investigate the RSVP message contents closer. Each router adds its IP address to the message to enable routing of the Reservation (Resv) messages sent in the reverse direction to follow exactly the same network path. The Resv message does not use the Router Alert option, but is rather explicitly routed on a hop-by-hop basis between the network routers using the state established earlier. In addition to the Path and Resv messages, RSVP has a few other message types delivered on a hop-by-hop basis.

Recently the IETF has specified a NSIS (Next Steps in Signaling) framework to handle signaling in the Internet. The Generic Internet Signaling Transport (GIST) protocol has been specified to transport the application-specific signaling messages over the Internet [SH06]. GIST messages are transfered using TCP or UDP as the transport protocol, depending on whether a reliable connection-oriented service or a connectionless service is desired. The use of SCTP to carry GIST messages is also under investigation. GIST has some common characteristics with RSVP: it uses a Router Alert option to wake up the GIST-aware routers along the path, and for further signaling explicit hop-by-hop routing can be applied using the state established at routers.

5. Past IETF Activities

This section discusses the past history of the IETF in considering link-layer triggers and other issues of cross-layer communication.

The IAB has an internet-draft on "Architectural Implications of Link Indications" that summarizes current proposals, describes the architectural issues and provides examples of appropriate and inappropriate uses of link layer indications [Abo06]. The document also gives a history of the integration of link indications within the Internet architecture.

The "Performance Implications of Link Characteristics (pilc)" working group produced seven RFCs concerning different types of links and their effects on transport protocols. The PILC working group did not explicitly consider cross-layer interactions; however, the Performance Enhancing Proxies document [RFC3135] gives guidelines for designing proxies that could also be useful considerations for network devices with cross-layer functionality.

The Triggers for Transport BOF in November 2002 [TrigTranBof] discussed triggers such as "Link Up", "Link Down", and "Packet Discarded". The necessity of a focused and narrow problem statement was discussed, with a need to define the semantics and uses of triggers in a exact way. It was questioned whether different wireless link technologies would be able to reliably produce the required information for the trigger, and what kind of responses would be appropriate at the transport protocol. The consensus was that the "Link Up" trigger might be viable , but that a "Link Down" trigger would be more difficult to be implement in a way that would be useful to the transport protocol. The BOF did not result in the creation of a working group.

The Transport Service at the Intermediary BOF (intersec) [IntersecBof] in March 2003 proposed to work on an architecture that helps performance enhancing middleboxes interoperate better with end-to-end transport protocols, especially with end-to-end security. No working group was established.

BOFs on "Access Link Intermediaries Assisting Services (alias)" [AliasBof1, AliasBof2] were held in two consecutive IETF meetings in 2003, continuing from the trigtran and intersec BOFs. ALIAS extended the discussion to middle-boxes that explicitly signal their existence and capabilities to the transport end-points (and vice-versa). ALIAS included an extensive discussion of security issues, along with a discussion of whether the possible benefits of such intermediaries would be clear enough to make the work worthwhile. No working group was created.

Recently there was a BOF proposal for IETF-66 in July 2006 called "Transport-Enhancing Refinements to the Network Layer Interface (ternli)". This didn't get as far as the above mentioned related BOFs, because the BOF was canceled before the IETF meeting. Instead, a group of people were gathered in an ad-hoc meeting in Montreal discussing the problem space. TERNLI was motivated in part by the research conducted in the MOBOPTS IRTF group about the effects of mobility to transport protocols. While there was an agreement that an explicit signaling mechanism between the transport and network layers should not be limited to mobility, there was a discussion of how the responsibilities should be divided between the Transport and Internet Areas. It was discussed that the work in these two areas is rather disconnected, and it is not always known what related work is being done in the other area. It was agreed that it is worthwhile to continue the discussion on the related research at least informally on a mailing list established under the IETF servers. The Jabber log of the meeting can be found at [http://www.ietf.org/meetings/ietf-logs/tsvwg/2006-07-11.html].

6. Possible Problems with Cross-layer Mechanisms

As described in the summary above on past IETF activities, there is some interest and motivation in starting up new work on explicit cross-layer communication. This section attempts to summarize what is known about some of the open issues in this area.

Today the Internet contains a wide variety of different types of middleboxes, tunnels, and advanced packet handling technologies that could cause problems for protocols that assume a simple architecture of interconnected routers with simple packet forwarding algorithms. In addition, the layer-two technologies have become more complex and difficult to model and understand correctly. In this section we list some common challenges to cross-layer mechanisms.

6.1. Security Issues

A cross-layer signaling protocol needs protective measures that are strong enough to make attacks on the protocol difficult and reasonably unprofitable. At the same time, if an otherwise light-weight protocol has heavy-weight security mechanisms, the cost of the security procedures may outweigh the possible benefits of the protocol.

For in-band mechanisms that use reserved header bits or IP options, the receiver of the packet can be expected to check that the IP addresses and transport ports match the existing connection, and that the sequence numbers in the packet belong to the currently valid window. Therefore, blind attacks generated outside the packet transmission path have a reasonably low probability of succeeding. For example, most TCP connections survive comfortably in the Internet, although security of the basic TCP has been discovered to be insufficient in certain mission-critical long-term connections [SD06]. However, an attacker on a connection path that is able to read the transport and IP headers has a good chance of causing harm to a connection, particularly if the packet contains additional explicit information about the connection, for example in an IP option. IPsec can protect the transport header, but does not protect a mutable IP option that can be modified by routers along the path.

Out-of-band messages do not necessarily include the additional context from the transport protocol, so they can be an easier target for blind attackers. If a transport protocol context exists, for example when the message is triggered by a data packet, the sending router of the out-of-band signaling message can include the transport header with the message, as for example ICMP does, to authorize the message based on the "proof" that the message sender is located on the packet transmission path. However, this is unlikely to be useful when IPsec is used to encrypt the data traffic. To more securely authenticate the sender of a signaling message a more elaborate security framework is needed. In many cases the complexity of such a security framework causes the costs of the mechanism to defeat the possible benefits.

The routers on the connection path can also try to cheat a cross-layer signaling mechanism. A first-hop router that is located in the same administrative domain with the transport end-host may have an incentive to game the protocol to the end-host's benefit. For example, in the case of Explicit Congestion Notification, a router could try to erase the Congestion Experienced bit on the packet, or a Quick-Start-aware router could try to game a better transmission rate for the transport sender. ECN and Quick-Start both use random content in the header fields called Nonces to make it more difficult for routers and receivers to misuse the protocol. Nonces usually do not provide full protection against misuse, but rather make cheating difficult enough to be unprofitable.

[TBD: Check out RSVP for IPsec [RFC2207]. However, that document does not address the tunnel mode of IPsec.]

6.2. IP Tunnels

IP tunnels are a challenge for an explicit cross-layer protocol that requires the participation of routers, because the tunnel isolates the original IP header inside an outer header. In particular, IPsec tunnels are intended to protect the inner header, which makes it possible that exposing content from the inner IP header could be a violation of the underlying security policy -- even if it was technically possible for a tunnel ingress to do so. The Quick-Start specification includes a thorough discussion of problems with IP tunnels [FAJS06].

6.3. Non-conformant routers and middleboxes

[MAF04] observes that for 70% of the destinations tested, TCP SYN packets with unknown IP options were either lost in the network or ignored by the receiving web server. ([MAF04] was not able to determine further why these connections failed when unknown IP options were added to the TCP SYN packets.) The presence of routers or middleboxes that drop packets containing unknown IP options would be a major obstacle to any cross-layer mechanisms that depended on the use of IP options.

6.4. Processing efficiency

Packets with IP options are assumed to take the slow-path processing path in most routers, as opposed to the optimized fast-path. If the use of IP options or other mechanisms requiring router attention gained in popularity, the impact on the processing efficiency of routers would have to be considered. In the Quick-Start proposal, it is assumed that Quick-Start-capable routers would rate-limit the number of Quick-Start requests that are processed, to preserve router efficiency and to protect against possible attacks on the routers themselves.

6.5. Path vs. link indications

In some proposed cross-layer protocols, it has been unclear whether information from a single local link or portion of a path is sufficient to change the transmission rate of the sender. When a link indication causes a reduction of the transmission rate, there is usually no problem as long as reasonable measures are taken to protect from third-party blind attacks and to ensure that the link indication comes from a link that is still on the end-to-end path. However, it is generally not acceptable for a link indication to cause an increase in the transmission rate without first checking the status of the whole network path between the sender and receiver. Although the position of the IETF is probably solid on this principle, in several cases the last-hop link is "known" to be the bottleneck on the connection path, and in such cases there is a high likelihood that the bandwidth of the last-hop link determines the transmission rate between the sender and receiver. Considering the difficulties of globally deploying an end-to-end cross-layer scheme that requires feedback from all routers along a path, it is possible that in the future parties will find it tempting to use local link information in ways that are considered inappropriate by the IETF.

6.6. Other concerns

[TBD: traffic normalizers]

[TBD: Problem of limited TCP option space]

[TBD: layer 2 queues and operation logic]

7. Proposals for Future Actions

* How tunnels should support cross-layer mechanisms?

* How to fix or go around routers and middleboxes that drop packets with unknown options?

* How about fast-path processing at routers?

Normative References

[RFC793] J. Postel. Transmission Control Protocol. RFC 793, September 1981.

[RFC2119] S. Bradner. Key words for use in RFCs to Indicate Requirement Levels. BCP

14, RFC 2119, March 1997.

[RFC2460] S. Deering and R. Hinden. Internet Protocol, Version 6 (IPv6) Specification. RFC 2460, December 1998.

[RFC2581] M. Allman, V. Paxson, and W. Stevens. TCP Congestion Control. RFC 2581. April 1999.

[RFC2914] S. Floyd. Congestion Control Principles.

[RFC3168] K.K. Ramakrishnan, S. Floyd, and D. Black. The Addition of Explicit Congestion Notification (ECN) to IP. RFC 3168, Proposed Standard, September 2001.

Informative References

[AliasBof1] Access Link Intermediaries Assisting Services (alias) Bof minutes from IETF-57, Vienna, Austria, July 2003. Available at http://www3.ietf.org/proceedings/03jul/250.htm

[AliasBof2] Access Link Intermediaries Assisting Services (alias) Bof minutes from IETF-58, Minneapolis, MN, USA, November 2003. Available at http://www3.ietf.org/proceedings/03nov/248.htm

[RFC792] J. Postel. Internet Control Message Protocol. RFC 792, September 1981.

[RFC2205] R. Braden (ed.). Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification. RFC 2205, September 1997.

[RFC2207] R. Berger and T. O'Malley. RSVP Extensions for IPSEC Data Flows. RFC 2207, September 1997.

[RFC2463] A. Conta and S. Deering. Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification. RFC 2463, December 1998.

[RFC3135] J. Border, M. Kojo, J. Griner, G. Montenegro, and Z. Shelby. Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations. RFC 3135, June 2001.

[RFC3234] B. Carpenter and S. Brim, Middleboxes: Taxonomy and Issues, RFC 3234, February 2002.

[Abo06] B. Aboba (ed.). Architectural Implications of Link Indications, Internet-Draft "draft-iab-link-indications-05.txt", July 2006. Work in progress.

[BJSK06] B. Briscoe, A. Jacquet, A. Salvatori, and M. Koyabe. Re-ECN: Adding Accountability for Causing Congestion to TCP/IP. Internet-Draft "draft-briscoe-tsvwg-re-ecn-tcp-02", June 2006. Work in progress.

[Cla88] D. D. Clark. The Design Philosophy of the DARPA Internet Protocols. In Proceedings of ACM SIGCOMM '88, pages 106--114, Stanford, CA, USA.

[FAJS06] S. Floyd, M. Allman, A. Jain, and P. Sarolahti. Quick-Start for TCP and IP. Internet-Draft "draft-ietf-tsvwg-quickstart-05.txt", July 2006. Work in progress.

[GL03] A. Gurtov and R. Ludwig. Lifetime Packet Discard for Efficient Real-Time Transport over Cellular Links. ACM Mobile Computing and Communications Review, 7(4):32-45, October 2003

[Hal96] F. Halsall. Data Communications, Computer Networks and Open Systems, Fourth edition. Addison-Wesley, 1996.

[IntersecBof] Transport Service at the Intermediary BOF (intersec) minutes from IETF-56, San Francisco, CA, USA, March 2003. Available at http://www3.ietf.org/proceedings/03mar/248.htm

[KHR02] D. Katabi, M. Handley, and C. Rohrs. Congestion Control for High Bandwidth-Delay Product Networks. In Proceedings of ACM SIGCOMM 2002, Pittsburgh, PA, USA, August 2002.

[KSE+04] R. Krishnan, J. Sterbenz, W. Eddy, C. Partridge, and M. Allman. Explicit Transport Error Notification (ETEN) for Error-Prone Wireless and Satellite Networks. Computer Networks, 46(3), October 2004

[LR99] R. Ludwig and B. Rathonyi. Link Layer Enhancements for TCP/IP over GSM. In Proceedings of the Conference on Computer Communications (IEEE Infocom), New York, USA, March 1999.

[MAF04] A. Medina, M. Allman, and S. Floyd. Measuring Interactions Between Transport Protocols and Middleboxes. Internet Measurement Conference 2004, August 2004. URL "http://www.icir.org/tbit/".

[RW03] M. Rossi and M. Welzl. On the Impact of IP Option Processing, Preprint-Reihe des Fachbereichs Mathematik - Informatik, No. 15, Institute of Computer Science, University of Innsbruck, Austria, October 2003.

[RW04] M. Rossi and M. Welzl. On the Impact of IP Option Processing - Part 2, Preprint-Reihe des Fachbereichs Mathematik - Informatik, No. 26, Institute of Computer Science, University of Innsbruck, Austria, July 2004.

[SAF06] P. Sarolahti, M. Allman, and S. Floyd. Determining an Appropriate Sending Rate Over an Underutilized Network. To appear in Computer Networks Journal, Elsevier, August 2006.

[SEE+06] S. Schuetz, L. Eggert, W. Eddy, Y. Swami, and K. Le. TCP Response to Lower-Layer Connectivity-Change Indications. Internet-Draft "draft-schuetz-tcpm-tcp-rlci-00", May 2006. Work in progress.

[SH06] H. Schulzrinne and R. Hancock. GIST: General Internet Signaling Transport. Internet-Draft "draft-ietf-nsis-ntlp-10", July 2006. Work in progress.

[SKDK06] P. Sarolahti, J. Korhonen, L. Daniel, and M. Kojo. Using Quick-Start to

Improve TCP Performance with Vertical Hand-offs. To appear in IEEE Workshop on Wireless Local Networks (WLN) 2006, Tampa, FL, USA, November 2006.

[SD06] R. Stewart, M. Dalal (ed.). Improving TCP's Robustness to Blind In-Window Attacks. Internet-Draft "draft-ietf-tcpm-tcpsecure-05.txt", June 2006. Work in progress.

[TrigTranBof] Triggers for Transport (trigtran) Bof minutes from IETF-56, San Francisco, CA, USA, March 2003. Available at http://www3.ietf.org/proceedings/03mar/251.htm

[Wel03] M. Welzl, PMTU-Options: Path MTU Discovery Using Options. Expired Internet-Draft "draft-welzl-pmtud-options-01.txt", February 2003. URL "http://www.welzl.at/research/publications/".

[XSSK05] Y. Xia, L. Subramanian, I. Stoica, and S. Kalyanaraman. One More Bit Is Enough. In Proceedings of SIGCOMM 2005, August 2005.

[ZDPS01] Y. Zhang, N. Duffield, V. Paxson, and S. Shenker, On the Constancy of Internet Path Properties, Proc. ACM SIGCOMM Internet Measurement Workshop, November 2001.

AUTHORS' ADDRESSES

Pasi Sarolahti
Nokia Research Center
P.O. Box 407
FI-00045 NOKIA GROUP
Finland
Phone: +358 50 4876607
Email: pasi.sarolahti@nokia.com

Sally Floyd
Phone: +1 (510) 666-2989
ICIR (ICSI Center for Internet Research)
Email: floyd@icir.org
URL: http://www.icir.org/floyd/