

Improving TCP-friendliness for mHIP

Tatiana Polishchuk¹, Andrei Gurto¹²

¹ Helsinki Institute for Information Technology HIIT,
PO Box 19800, 00076 Aalto, Finland
WWW home page: <http://www.hiit.fi>

² Centre for Wireless Communications, Oulu, Finland

Abstract. Multihomed environments are getting increasingly common, especially for mobile users. mHIP was designed to provide secure multipath data transmission for the multihomed hosts and boost throughput of a single TCP connection by effectively distributing data over multiple available paths.

In this paper we develop a TCP-friendly congestion control scheme for mHIP secure multipath scheduling solution. We enable two-level control over aggressiveness of the multipath flows to prevent stealing bandwidth from the traditional transport connections in the shared bottleneck. We demonstrate how to achieve a desired level of friendliness at the expense of inessential performance degradation. A series of simulations verifies that the proposed congestion control for mHIP meets the criteria of TCP-compatibility, TCP-equivalence and TCP-equal share, preserving friendliness to UDP and another mHIP traffic.

Keywords: Internet, HIP, multipath routing, TCP-friendliness, goodput

1 Introduction

Multipath data transfer is a promising technique for enhancing reliability of Internet connections. New mobile devices and laptops are equipped with several network interfaces (e.g., WLAN, GPRS, 3G) and have multiple links to the Internet, which results in availability of multiple paths between a source and destination end host.

TCP [20] comprises a major share of the total Internet traffic. Among its other management tasks, TCP controls segment size, the rate at which data is exchanged, and network traffic congestion [21]. However, traditional TCP flow is constrained to use one path only per one connection between two communicating hosts. There are efforts within the networking community to overcome this limitation. Most of these efforts rely on the mechanisms which aggressively compete for network resources. Naive designs and implementations risk substantial unfairness to well-behaved TCP flows. Proper per-flow congestion control is required to limit aggressiveness of the proposed multipath solutions.

Other multipath communication methods, proposed to efficiently utilize multiple access links, unable to take advantage of all available multipath bandwidth

because they do not properly consider end-to-end delay of packet transmission. Out-of-order data arrivals at a receiver cause unpredictable underutilization of spare network capacity. Packet reordering and non-congestion packet loss can significantly degrade TCP performance.

TCP-friendliness has emerged as a measure of correctness in Internet congestion control. The notion of TCP-friendliness was introduced to restrict non-TCP flows from exceeding the bandwidth of a conforming TCP running under comparable conditions. Protocols commonly meet this requirement by using some form of AIMD (Additive Increase Multiplicative Decrease) congestion window management, or by computing a transmission rate based on equations derived from AIMD model.

In the prior work [8] we proposed a multipath HIP solution, which combines the advantages of HIP advanced security with the benefits of multipath routing. mHIP uses multiple parallel flows simultaneously in order to boost throughput of the TCP connection. The multipath scheduler takes into account rapidly changing parameters of the available paths, including the TCP queuing delay at a sender and the network delay, and sends each data packet through the path with the earliest estimated time of arrival. Simple congestion control measures were suggested to provide reliable multipath data delivery.

In this paper we study TCP-friendliness of HIP multipath design with respect to coexisting connections. The contributions of this work include the development of a two-level congestion control concept for a reliable multipath data transmission and methods of tuning aggressiveness of individual flows from the multipath bundle in order to provide a desirable level of TCP-friendliness while avoiding significant performance degradation.

The rest of the paper is organized as follows. Section 2 summarizes the related work. Preliminaries are presented in Section 3 and contain the review of multipath HIP simple congestion control and definitions of TCP-friendliness. Section 4 presents the step-by-step work which was done to enable TCP-friendly congestion control for mHIP. Conclusions and future work are given in Section 5.

2 Related Work

Despite the fact that multiple multipath solutions for multihomed hosts has recently emerged, multipath routing is not yet widely deployed in practice. Researchers study advantages of its implementation on different layers of the TCP/IP stack.

Transport layer solutions, such as SCTP [13], MPTCP [6], TCP-MH [15], can naturally obtain the most recent information on the quality of different paths and detect congestion situations in timely manner. For example, SCTP can perform measurements across several paths simultaneously, and then map flows on one or another path.

Network layer approaches ([2], [5]) are easy to deploy and totally transparent to applications and involve only minimal changes in contrary to the application and transport layer solutions which involve many changes in the infrastructure.

Wedge-layer approaches, implemented in HIP [7], LIN6 [11], MIP6 [14], conduct multiaddressing support in a functional layer between IP and transport. They have an advantage of being able to maintain multiaddressing information across transport associations. The transport activity between two endpoints may well be able to use multiaddressing immediately and with no further administrative overhead. Moreover, edge-based locator exchange protocols can be incorporated without necessitating modification to any host's IP or transport modules, which makes them the best choice to provide multihoming and multipath functionality for legacy Internet applications and transport protocols.

There is an effort in the community to create new methods which effectively and TCP-friendly utilize a spare network capacity. In [9] authors created a parallel multipath TCP solution, which controls data transmission over coordinated multiple TCP connections. They stressed the importance of TCP-friendliness for multipath schemes and suggested a way to find a balance between effectiveness and fairness. Their work provided a motivation to design a TCP-friendly congestion control over multipath flows inside one TCP connection.

When data packets are sent over several paths inside one connection they can experience different end-to-end delays and arrive out of order. In case of TCP traffic, packet reordering causes significant performance degradation. The authors of [16] surveyed and analyzed relevant techniques on coping with multipath TCP packet reordering. They conclude that there exists no one-fits-all solution to solve the problem of packet reordering for multipath TCP. Basing on the methods [3], [4], [17], [24] we suggest the improvement for multipath HIP which reduced the level of reordering on the receiver and significantly improved TCP-friendliness of our scheme.

According to the resource pooling principle [23] when several subflows of one connection share a bottleneck, their resource consumption adds up. Multipath connections with a large number of TCP-friendly subflows can compete unfairly against a smaller number of regular TCP connections. Each subflow is as aggressive as a single TCP, and a bundle of n TCP-friendly subflows will hence use an approximately n times greater share of the bottleneck resource than they should. TCP-fair multipath connection should displace no more TCP traffic than a traditional TCP stream would displace. A number of methods [9], [18], [10] were proposed to study and solve the TCP-fairness problem.

Although the current implementation of mHIP was not intended to achieve the TCP-fairness criterion, the two-level congestion control scheme proposed further in this paper will provide TCP-fairness of mHIP by default, and the preliminary experiments with competing mHIP flows inside one connection confirmed this assumption. mHIP multipath scheduling assumes the paths are bottleneck-disjoint. This automatically liberates us from the necessity to prove TCP-fairness of our solution since multiple flows of a single multipath HIP connection never share the same bottleneck link. When it is not possible to guarantee bottleneck independence of the paths a coupled congestion control for congestion management [18] was recently suggested by MPTCP working group. The complexity and effects of applying such measures are out of the scope of this paper.

3 Preliminaries

3.1 TCP-friendliness Definitions

TCP-friendliness is a generic term describing a scheme that aims to use no more bandwidth than TCP uses. In this paper we study mHIP congestion control in view of the criteria proposed in [22]:

A *TCP-compatible* flow, in the steady state, should use no more bandwidth than a TCP flow under comparable conditions, such as packet-loss rate and round-trip time (RTT). However, a TCP-compatible congestion control scheme is not preferred if it always offers far lower throughput than a TCP flow.

A *TCP-equivalent* scheme merely ensures the same throughput as TCP when they experience identical network conditions. Although a TCP-equivalent scheme consumes TCP-equivalent bandwidth when working by itself, it may not coexist well with TCP in the Internet.

TCP-equal share is a more realistic but more challenging criterion than TCP-equivalence and states that a flow should have the same throughput as TCP if competing with TCP for the same bottleneck. A TCP-equivalent flow may not be TCP-equal share, but the opposite is always true.

To be able to meet all three criteria a TCP-friendly scheme should use the same bandwidth as TCP in a steady-state region, while being aggressive enough to capture the available bandwidth and being responsive enough to protect itself from congestion, as the packet-loss condition changes in the paths in the transient state. *Aggressiveness* of a scheme describes how the scheme increases the throughput of a flow before encountering the next packet loss, while *responsiveness* describes how the scheme decreases the throughput of a flow when the packet-loss condition becomes severe.

In what follows we will examine the ability of our multipath solution to adhere to the proposed definitions of TCP-friendliness. We evaluate its performance using the factor of friendliness $FF(flow) = \frac{T(flow)}{T(TCP)}$ as the measure. Here $T(\cdot)$ denotes the average flow throughput in Mbps. $FF = 1$ indicates the solution satisfies the strongest TCP-equal share criterion, while solution resulting in $FF > 1$ is more aggressive than a typical TCP and the one with $FF < 1$ may be not TCP-compatible.

3.2 Review of Multipath HIP with Simple Congestion Control

In the prior research [8] HIP multipath scheduling showed a potential to aggregate about 99% of the sum of individual paths bandwidth. Simple congestion detection and avoidance are able to prevent the sending rate of the multipath traffic from significant degradation caused by congestion in the paths. Before we start evaluating mHIP congestion control scheme in the view of TCP-friendliness criteria, we recall how it operates.

1. *Connection establishment*

During the base exchange HIP obtains information about the number of available interfaces on both communicating hosts and the number of available paths with the initial parameters such as available bandwidth and propagation delay.

2. *Updating parameters of the paths*

mHIP uses HIP signaling packets for path probing. The frequency of heartbeats can vary depending on the particular setup.

3. *Sending data*

HIP multipath scheduler optimally splits data among the paths according to their capacities. The details of scheduling algorithm are provided in [8]. mHIP stores packet-to-path assignments at the sender and also in the ESP packet headers, which are used according the HIP standard [12]. SPI number, specified in the packet header corresponds to the path which is assigned to deliver this particular packet.

4. *Congestion control*

Marking and multipath congestion avoidance techniques provide a simple congestion control for mHIP. One packet per round-trip time is marked on the departure to each path. The expected delivery time of the marked packet is stored at the sender and then compared to its actual arrival time value on the receipt of the corresponding ACK. If the estimated delivery time and the actual arrival time of the marked packet are noticeably different, the scheduler considers the path to be congested.

Multipath congestion avoidance technique specifies two indicators of the path congestion:

- Case 1: standard TCP *dupack action*, when the sender is retransmitting the packet after the receipt of three duplicate acknowledgments from the receiver;
- Case 2: observed delivery time of the marked packet exceeds its corresponding expected delivery time by more than some preset value.

If any of the two indicators suggest congestion, the path is temporarily closed and the packets are redirected to the other available paths. mHIP sends regular probes to the congested path to detect when the path becomes again free for reliable data transmission.

5. *Assumptions and limitations*

Our approach corresponds to the class of disjoint multipath routing [19]. The paths are restricted to have independent bottlenecks. The scheduler resides at the sender side, no information from the receiver is available other than TCP acknowledgments (ACKs) received by the sender. At least one available path should not be congested at any given point of time.

4 Improving mHIP Step by Step

Next we examine mHIP congestion control in the view of TCP-friendliness criteria. We analyze the reasons why multipath flows not always fairly share available bandwidth with TCP and propose the methods to improve TCP-friendliness of our multipath solution.

4.1 Experimental Evaluation of mHIP with Simple Congestion Control

All simulations presented in this work were run using ns-2 network simulator [1]. A new protocol agent was implemented on the basis of TCP New Reno to deal with the multipath flow controlled by HIP. Existing TCP and UDP modules were also used to simulate external cross-traffic competing with HIP multipath flows for bottleneck bandwidth.

Consider a simulation model shown in Figure 1. A TCP traffic flow, controlled by multipath HIP, is sent from $n0$ to $n1$ over two available paths: $Path1 = n0 - n2 - n1$ and $Path2 = n0 - n3 - n1$ with the bandwidth of 8Mbps and 4Mbps respectively. Since multipath scheduler is distributing the traffic according to bandwidth-delay product of the paths, for simplicity the propagation delay is fixed to be the same for all the links and equals 30 ms. mHIP is calculating the end-to-end propagation delays in the paths, they can consist of any number of connected links and intermediate nodes. Node $n4$ is used for the path $n2 - n1 - n4$ construction, which accommodates a standard TCP New Reno flow, competing against one flow from the mHIP bundle for the bottleneck link $n2 - n1$. Drop-Tail scenario was used to manage the bottleneck link, its size is 1.5 times the bandwidth-delay product of the link. The packet size in each flow is 1250 bytes. The simulation runs for 20 seconds, which we believe is sufficient to reflect the difference between the proposed congestion control solutions. Appropriate $Rwin$ values were used at the receivers to allow maximum throughput of the flows.

We begin our first experiment with an empty network and then allocate multipath HIP subflows to the two end-to-end paths. At the same time we start sending a TCP traffic from $n2$ to $n4$, which will compete with mHIP flow in




Fig. 1. 2-path simulation model.

the bottleneck link $n2 - n1$. To simulate variable network conditions we also introduce cross-traffic to *Path2*. A 4Mbps UDP flow was scheduled between 5 and 11 seconds of the simulation run, triggering a congestion situation in *Path2*.

Figure 2 shows mHIP and TCP New Reno flow throughputs, averaged over 0.1 sec. As one can clearly conclude from the chart the flows do not share the bottleneck bandwidth fairly. mHIP (dotted curve) occupies more bandwidth, with the average of $T(mHIP1) = 3.98\text{Mbps}$ and TCP takes just $T(TCP) = 3.56\text{Mbps}$ resulting in the friendliness factor $FF = \frac{T(mHIP1)}{T(TCP)} = 1.11$.




Fig. 2. mHIP flow competes with TCP New Reno flow in *Path1*.

Lets try to understand the reason why mHIP starts starving the TCP flow during the particular time period. In the beginning of the simulation run mHIP and TCP flows share the bandwidth mostly fair. At some point after 5 seconds the marking technique reports a congestion situation, resulting from the competition with UDP cross-traffic in *Path2*. Let w be the number of packets at the sender, which corresponds to the *cwnd* value of the global TCP flow controlled by mHIP. The multipath scheduler sends w_1 packets to *Path1* and w_2 packets to *Path2* in the share correspondent to path characteristics with the total $w_1 + w_2 = w$. According to the congestion avoidance scheme *Path2* is closed and all the traffic from the congested *Path2* is rerouted to *Path1*, meaning that at this same time *Path1* receives not only its own share w_1 but also extra w_2 packets. In this region mHIP is dominating and stealing bandwidth from the competing TCP transport transmission in the bottleneck link $n2 - n1$. The proposed congestion control method is definitely more aggressive than AIMD policy of a typical TCP.

4.2 Designing TCP-friendly Congestion Control for mHIP

We want our mHIP connections to coexist with other traffic providing opportunities for all to progress satisfactory. To limit aggressiveness of the flow growth we propose the following two-level congestion control scheme - per-path AIMD plus TCP global stream congestion control on top of it, and introduce a sender-side buffer to provide better control on the packet sequence in congestion situations.




Fig. 3. Two-level multipath congestion control.

The proposed twofold congestion control scheme is illustrated in Figure 3. Global congestion controller coordinates the work of the individual per-path controllers and balances traffic load between the paths according to their available capacity. If *cwnd* capacity of the quickest path is exceeded, the path with the next minimum estimated arrival time is chosen.

An important property of the proposed scheme is that per-path controllers are connected so that the aggregated congestion window is a simple sum of per-flow congestion windows. Same rule applies to the threshold values. Connecting per-path congestion control parameters in such a way we guarantee the resulting multipath bundle behaves as a single TCP if all are sent to the same path.

Below we summarize the proposed updates to the mHIP multipath scheduling design presented in subsection 3.2. Parts 1,2 and 5 (connection establishment, path parameters updates and assumptions) remain unchanged, while there are some additions to the rest:

3. *Sending data*

After per-path congestion control limitations were introduced the scheduler takes in consideration the current sizes of per-path congestion windows. If *cwnd* capacity of the best path is exceeded, the path with the next minimum estimated arrival time is chosen. If there is no available capacity in any of the paths, the packet is placed to the sender-side buffer until new ACK arrives.

4. *Congestion control*

Marking is now removed from the congestion control scheme. Multipath congestion avoidance retains only one congestion indication, the standard TCP *dupack* event. Upon receipt of a preset number of *dupacks* (3 for standard TCP) the scheduler determines from which path the packet is missing and halves *cwnd* and *ssthresh* values of the corresponding path. This action reduces data intake in the congested path and automatically redirects traffic to the other paths which have available capacity. If there is no capacity in the paths, extra data goes to the sender-side buffer. Maximum capacity of the buffer is set to TCP receiver window size *Rwin*, making it capable to occupy the maximum flight-size number of packets in case of severe congestion situations.

4.3 Experimental Evaluation of mHIP with the Updated Congestion Control

To validate correctness of the proposed congestion control scheme we repeat the experiments with the simulation scenario described in Section 4.1. Again, one of the multipath HIP flows sent to *Path1* meets with the external TCP flow in the bottleneck link $n2 - n1$, while the other flow sent to *Path2* is interrupted by UDP cross-traffic in the link $n0 - n3$.

The resulting throughputs of the two flows competing in *Path1* are shown in Figure 4. mHIP average flow throughput is $T(mHIP) = 3.56\text{Mbps}$ and TCP takes about $T(TCP) = 3.98\text{Mbps}$ resulting in the fairness factor $FF = \frac{T(mHIP)}{T(TCP)} = 0.89$. Now we observe the opposite extreme: mHIP flow behaves too leniently and is not able to occupy available bandwidth effectively. In the following section we analyze the problem and propose a method to solve it.

4.4 Balancing between Aggressiveness and Responsiveness

Competition with the external traffic naturally influences effectiveness of multipath scheduling. Mistakes in the expected delivery time estimations result in the output sequence reordering at the receiver. TCP sender receives multiple *dupacks* in response to reordering, which mHIP scheduler treats as an indication of congestion. In response to the congestion mHIP scheduler halves congestion window of the corresponding path, reducing aggressiveness of the traffic flow. This precaution could be too strict in case when the missing sequence numbers are not lost but just slightly delayed in competition with the external flows.




Fig. 4. mHIP flow controlled by the proposed twofold multipath congestion control is suppressed by TCP.




Fig. 5. mHIP flow 1 friendly coexists with TCP New Reno flow.

To differentiate between the reordering signals and actual losses we propose the following modifications to mHIP congestion control scheme. First, we increase *dupthresh* value defining the number of *dupacks* which serve as an indication of congestion. This method is proposed in the related work [4], [24] as a cure from the mild packet reordering. Compared with the default *dupthresh* of three, the proposed techniques improves connection throughput by reducing the number of unnecessary retransmissions. But one should adjust *dupthresh* value carefully since making it too large slows down the reaction of the system to the actual losses and can significantly degrade the overall performance in the networks with high loss rates.

Additionally we introduce a new time variable ADDR (allowable delay due to reordering), which counts how much time has elapsed since the congestion situation in some path was reported. If the missing sequence number has arrived successfully during this allowable time period and the corresponding ACK arrives to the sender, *cwnd* and *ssthresh* of the path should be returned to the values prior to congestion notification. ADDR is chosen to be less than the shortest RTT among the paths used to deliver multipath flow. It will assure accurate differentiation between the packets delayed due to reordering and their duplicates retransmitted after the loss was reported. If the original packet arrives, the retransmitted one is naturally disregarded by the receiver.

4.5 Final Validation

Below we provide the final validation of effectiveness of the proposed modifications to mHIP congestion control. Again, we repeat the experiment described in Section 4.1 with the last version of mHIP with two-level congestion control scheme and all the proposed modifications applied.

Figure 5 illustrates significant improvement in TCP-friendliness of the mHIP flow when it competes against TCP for the bottleneck link bandwidth. Finally both mHIP and TCP flows are able to achieve comparable average throughputs of $T(mHIP1) = 3.80\text{Mbps}$ and $T(TCP) = 3.71\text{Mbps}$ with the friendliness factor $FF = \frac{T(mHIP1)}{T(TCP)} = 1.02$. The competition demonstrated high variation about the average during a short stabilization phase. This unfairness is rather moderate and can be tolerated as far as the flows quickly achieve stability and later coexist friendly.

4.6 UDP-friendliness

An interesting observation is that the second mHIP flow in *Path2* behaves also about friendly competing against the UDP cross-traffic which we used to simulate variable network conditions between 5 and 11 seconds. On this interval mHIP achieves the throughput of $T(mHIP2) = 4.20\text{Mbps}$. The solid curve in Figure 6 corresponds to the UDP cross-traffic flow with the average flow throughput $T(UDP) = 3.98\text{Mbps}$. The flows fight during negligible time period and then find stability to share the bottleneck about fairly with a moderate unfairness of $FF = \frac{T(mHIP2)}{T(UDP)} = 1.05$.




Fig. 6. mHIP flow 2 competes almost friendly with UDP cross-traffic.

4.7 TCP-compatibility and TCP-equivalence

According to the definitions TCP-compatible flow, in the steady state, should use no more bandwidth than a TCP flow under comparable conditions, while TCP-equivalent scheme ensures the same throughput as TCP when they experience identical network conditions. We send mHIP to the empty 2-path network with no cross-traffic to determine how effectively the protocol is able to use a spare network capacity in the steady state.

Figure 7 shows mHIP flow occupies no more available bandwidth than a TCP flow sent to the same path making it TCP-compatible. Moreover, mHIP achieves the same average flow throughput of 7.8Mbps as TCP in the steady state and thus meets the criteria of TCP-equivalence.

4.8 Friendliness to the Other mHIP

Another interesting question is whether mHIP competes friendly against the other mHIP connections. We run six multipath HIP connections in the simulated network scenario similar to the one we used in the experiments presented above, but now we have three parallel paths connecting the common source and destination with the following path bandwidths: 8 Mbps, 4 Mbps and 4 Mbps and the corresponding propagation delays: 60ms, 60ms and 20ms, which provide some diversity in the network parameters. Figure 8 demonstrates how the total network bandwidth is divided between the six multipath HIP bundles. The comparison shows a tolerable unfairness with the friendliness factor differing from 0.92 to 1.08. We conclude that multiple mHIP connections can coexist in the shared multipath network quite friendly to each other.




Fig. 7. Testing TCP-compatibility and equivalence of mHIP.




Fig. 8. Six mHIP connections share 3-path network about fair.

4.9 The Cost of Friendliness

We achieved the desired level of TCP-friendliness for our multipath HIP solution and would like to evaluate the cost in terms of performance degradation paid for this improvement.

We calculate the total throughput TT of the traffic flow controlled by multipath HIP. In the experiment where mHIP with simple congestion control policy demonstrated an excessive unfriendliness competing against TCP NewReno, $TT(mHIP) = 6.45\text{Mbps}$. After we applied a series of modifications to mHIP congestion control, similar experiment with the TCP-friendly mHIP resulted in $TT(mHIP) = 5.30\text{Mbps}$, which corresponds to $\sim 18\%$ performance reduction. A number of experiments with different network conditions confirmed the desired TCP-friendliness can be achieved at the cost of about 15-20% performance degradation.

5 Conclusions and Future Work

We showed a way how to tune aggressiveness of the multipath data transmission controlled by mHIP without losing its responsiveness in competition with cross-traffic. We designed a twofold congestion control scheme, and adjusted it to meet the TCP-friendliness definitions. Simulation results verify the improved congestion control algorithm meets TCP-compatibility, TCP-equivalence and TCP-equal share criteria under the proposed testing scenarios, and allows mHIP to coexist friendly with the other TCP, UDP and mHIP connections.

The work could be extended to provide a method to dynamically adjust mHIP congestion control variables and enable adaptivity to random congestion scenarios including extreme cases. We will continue examining mHIP friendliness in competition against different transport protocols other than TCP and compare the results against the alternative multipath proposals.

Acknowledgments

This work was supported in part by TEKES as part of the Future Internet program of the ICT cluster of the Finnish Strategic Centers for Science, Technology and Innovation.

References

1. Network simulator ns-2. <http://www.isi.edu/nsnam/ns/>, last checked 15/02/2010.
2. S. Barre and O. Bonaventure. Shim6 implementation report : Linshim6. Internet draft, draft-barre-shim6-impl-03.txt, September 2009.
3. S. Bhandarkar and A. L. N. Reddy. TCP-DCR: Making TCP robust to non-congestion events. In N. Mitrou, K. P. Kontovasilis, G. N. Rouskas, I. Iliadis, and L. F. Merakos, editors, *NETWORKING*, volume 3042 of *Lecture Notes in Computer Science*, pages 712–724. Springer, 2004.

4. E. Blanton and M. Allman. On making TCP more robust to packet reordering. *ACM Computer Communication Review*, 32:2002, 2002.
5. K. Chebrolu, B. Raman, and R. R. Rao. A network layer approach to enable TCP over multiple interfaces. *Wirel. Netw.*, 11(5):637–650, 2005.
6. A. Ford, C. Raiciu, S. Barre, and J. Iyengar. Architectural guidelines for multipath TCP development. Technical report, June 2010. Internet draft, draft-ietf-mptcp-architecture-01, work in progress.
7. A. Gurtov. *Host Identity Protocol (HIP): Towards the Secure Mobile Internet*. Wiley and Sons, 2008.
8. A. Gurtov and T. Polishchuk. Secure multipath transport for legacy Internet applications. In *Proc. of BROADNETS'09*, Madrid, Spain, Sept. 2009.
9. T. J. Hacker, B. D. Noble, and B. D. Athey. Improving throughput and maintaining fairness using parallel TCP. In *IEEE InfoCom*, 2004.
10. T. Ishida, K. Ueda, and T. Yakoh. Fairness and utilization in multipath network flow optimization. In *Proc. of 2006 IEEE International Conference on Industrial Informatics*, pages 1096–1101, 2006.
11. M. Ishiyama, M. Kunishi, and F. Teraoka. An analysis of mobility handling in LIN6. In *Proc. of International Symposium on Wireless Personal Multimedia Communications (WPMC'01)*, Aug. 2001.
12. P. Jokela, R. Moskowitz, and P. Nikander. Using the Encapsulating Security Payload (ESP) transport format with the Host Identity Protocol (HIP). IETF RFC 5202, Mar. 2008.
13. A. Jungmaier, E. Rescorla, and M. Tuexen. Transport layer security over Stream Control Transmission Protocol. RFC 3436, IETF, Dec. 2002.
14. J. Kempf, J. Arkko, and P. Nikander. Mobile IPv6 security. *Wirel. Pers. Commun.*, 29(3-4):389–414, 2004.
15. K.-H. Kim and K. G. Shin. Improving TCP performance over wireless networks with collaborative multi-homed mobile hosts. In *Proc. of the 3rd Int. conf. on Mobile systems, applications, and services (MobiSys'05)*, pages 107–120, June 2005.
16. K.-C. Leung, V. O. Li, and D. Yang. An overview of packet reordering in Transmission Control Protocol (tcp): Problems, solutions, and challenges. *IEEE Transactions on Parallel and Distributed Systems*, 18:522–535, 2007.
17. R. Ludwig and R. H. Katz. The Eifel algorithm: making TCP robust against spurious retransmissions. *SIGCOMM Comput. Commun. Rev.*, 30(1):30–36, 2000.
18. C. Raiciu. Coupled multipath-aware congestion control, March 2010. Work in progress.
19. S. Ramasubramanian, H. Krishnamoorthy, and M. Krunz. Disjoint multipath routing using colored trees. *Comput. Netw.*, 51(8):2163–2180, 2007.
20. W. R. Stevens. *TCP/IP illustrated (vol. 3): TCP for transactions, HTTP, NNTP, and the Unix domain protocols*. Addison Wesley Longman Publishing Co., Inc., Redwood City, CA, USA, 1996.
21. W. R. Stevens. TCP slow start, congestion avoidance, fast retransmit, and fast recovery algorithms. RFC 2001, IETF, Jan. 1997.
22. S.-C. Tsao and N. Chiao. Taxonomy and evaluation of TCP-friendly congestion-control schemes on fairness, aggressiveness, and responsiveness. *Network, IEEE*, 21(6):6–15, Nov. 2007.
23. D. Wischik, M. Handley, and M. B. Braun. The resource pooling principle. *SIGCOMM Comput. Commun. Rev.*, 38(5):47–52, 2008.
24. M. Zhang, B. Karp, S. Floyd, and L. Peterson. RR-TCP: A reordering-robust TCP with DSACK. In *Proc. of IEEE ICNP*, pages 95–106, 2003.