Multilink Transfer over Heterogeneous Networks

Dominik Kaspar* (Ph.D. Student), Audun Fosselie Hansen†* (Ph.D.), Carsten Griwodz‡* (Prof.)
*Simula Research Laboratory, Norway †Telenor R&I, Norway ‡University of Oslo, Norway
email:{kaspar, audunh, griff}@simula.no

I. Introduction

In a world with numerous heterogeneous wireless technologies (e.g. WLAN, HSPA, WiMAX, Bluetooth, etc.), devices such as cell phones and laptops are more and more often equipped with multiple communication interfaces. Motivated by this trend, major wireless providers, such as Telenor [1], are looking for solutions that can fully utilize multiple technologies when present.

While the term *multilink* refers to one host attached to many access networks, *multihoming* describes the method of maintaining multiple IP addresses on a single device, and the term *multipath* implies using different multihop end-to-end paths for routing. Although these three problems are closely related, our work focuses on multilink transfer and simultaneously utilizing all available interfaces. In theory [2], exploiting multiple links has several potential benefits: increased throughput by link aggregation, added fault tolerance by sending redundant data over independent links, and increased connectivity by combining several coverage areas.

II. Related Work and Preliminary Results

For several years related contributions have been proposed on many layers of the protocol stack, ranging from network level multihoming (e.g. [3]), over multipath transport for homogeneous and independent wired paths (e.g. [4]), to methods for application layer striping (e.g. [5]). A good overview of recent solutions to transport layer striping is given in [8], most of them focusing on modifications to TCP and SCTP (e.g. [9]). A lot of research has also focused on handover in case of failures on the primary path, for instance shim6 [6], a current IETF standardization effort to provide host-based and upper layer transparent multihoming. Vertical handover between wireless technologies has gained increased attention, and contributions to concurrent transfer over wireless networks are becoming more frequent (e.g. [7]).

The majority of existing solutions to network striping assume substantial protocol and end-host modifications, which may hinder general deployment. The fact that most of them have only been tested in simulations, often based on very simple assumptions about heterogeneity (e.g. [7], [10]), enforces our skepticism of deployability. As an initial study of link heterogeneity, we have set up an experiment to compare the characteristics of WLAN and HSDPA connections from a single location, using a laptop equipped with both technologies (shown with bold arrows in Fig. 1a). The client laptop is located 8 IP hops away from the server over the WLAN links, which is comparable to the 10 IP hops over the HSDPA link.

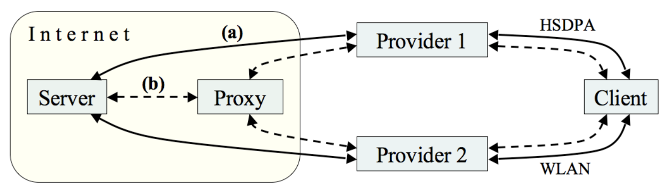

Fig. 1. A client device with connections to a server over two access networks, with (a) direct connections and (b) a proxy solution.

With this setup, ICMP echo messages were sent over a period of $24h$ (spaced $1s$ apart) from the client to the server over an HSDPA and two WLAN links. The comparison of loss, round-trip times and their standard deviation are shown in Table I. Throughput results are also shown (with the standard deviation in parenthesis), they were conducted from the same location by downloading a 42 MB large file 10 times from the server.

TABLE I
Round-trip time measurements over HSDPA and WLAN

	loss	min.	max.	mean	stdev.	TCP throughput
WLAN1	2.7%	2.2ms	5010ms	10.1ms	61.2ms	112 (10.4) KB/s
WLAN2	1.5%	3.2ms	1807ms	11.2ms	59.0ms	178 (54.5) KB/s
HSDPA	0.8%	67.5ms	1350ms	212ms	106ms	290 (13.7) KB/s

Figure 2 illustrates how the distribution of round-trip times varies significantly between WLAN and HSDPA (in this scenario, the first wireless hops dominate the total delay). The large differences in average and minimum RTTs contradict frequently made assumptions about the characteristics of heterogeneous links, which are often modeled too evenly and not very realistically (e.g. $30ms$ delay for both WLAN and UMTS [7]).

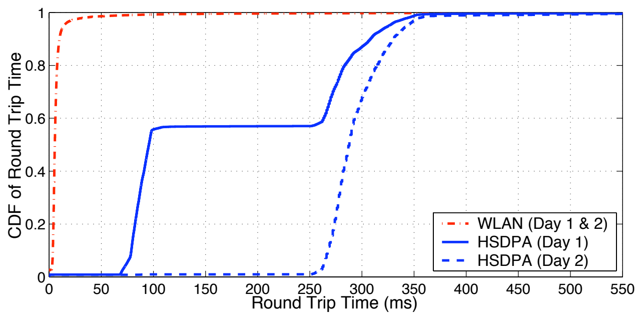

Fig. 2. CDF of 86400 RTTs between client and server on two test days (the two WLAN links are not distinguishable at this scale).

At present, no complete and readily deployable multilink solution exists due to implementation challenges and substantial network heterogeneity, such as dealing with transport layer packet reordering and optimally scheduling packets (or sessions) to available links that vary significantly in their characteristics. In fact, transport and congestion control over a single wireless channel is by itself considered a challenging problem [11].

We are convinced that a complete and most appropriate solution should consider several layers of the network stack and not claim one as the optimal target. Pure application striping will gain little from a multilink approach without support from the transport layer, the throughput may even decrease. The IP layer can solve parts of the multihoming implementation challenges; however without transport layer modifications, the aggregated bandwidth may not reach the application level. Transport layer modifications can increase the general throughput; however, to utilize the full potential, application-specific modifications are also necessary.

Thus, a major challenge is the implementation of an efficient cross-layer architecture for multilink support on end-hosts. Our goal is to approach this in a holistic way: to propose deployable modifications, support heterogeneity, and support cross-layer optimizations. In addition, we want to combine simulations and real-world experiments to obtain realistic evaluations and prove deployability.

III. Our Approach

It is our plan to develop a general architecture for supporting end-users to exploit multiple interfaces simultaneously. We believe that a cross-layer framework is required to optimally enable simultaneous data transfer over multiple interfaces with heterogeneous properties. Our ultimate goal is to develop a solution that can also handle the extreme dynamic heterogeneity expected in cognitive networks.

We are tackling these challenges with a bottom-up approach, enhancing the performance on the network and transport layers before designing application-specific optimizations. As a first step we will develop a multilink solution based on IP-in-IP tunneling, while evaluating different architectures and solutions on the transport level. From an ISP point of view, we envision the users to access any type of third-party services and servers on the Internet, also those out of the ISP's control. In contrast to sender-side methods (e.g. [8]), we therefore target a solution that does not require any modifications on the server. For minimizing changes needed at the clients, we pursue a proxy solution for connection splitting and packet modification. An important research topic will be to assign and design the tasks given to the proxy. The proxy may perform data assignment (e.g. path scheduling), reliable and efficient data delivery (e.g. controlling retransmissions), path monitoring, tunneling, and energy-optimization. Scalability is a key design requirement.

Currently, we are developing our testbed as described in Fig. 1b (dashed paths). The different solutions will be evaluated with this testbed over HSDPA and WLAN. The solutions will then serve as input to the MANGO project [12], where Telenor is one of the contributors. Application-specific middleware will then be able to take advantage of the developed transport layer solutions. Additional information on the state of this project are available at http://simula.no/research/networks/simtel.

References

[1] Telenor. [Online]. Available: www.telenor.com

[2] L. Golubchik et al., "Multi-path continuous media streaming: what are the benefits?" *Perform. Eval.*, 2002.

[3] D. S. Phatak, T. Goff, and J. Plusquellic, "IP-in-IP tunneling to enable the simultaneous use of multiple IP interfaces for network level connection striping," *Computer Networks*, 2003.

[4] J. R. Iyengar, P. D. Amer, and R. Stewart, "Concurrent multipath transfer using SCTP multihoming over independent end-to-end paths," *IEEE/ACM Trans. Netw.*, vol. 14, no. 5, 2006.

[5] B. Wang, W. Wei, Z. Guo, and D. Towsley, "Multipath live streaming via TCP: scheme, performance and benefits," in *ACM CoNEXT*. New York, NY, USA: ACM, 2007, pp. 1–12.

[6] E. Nordmark and M. Bagnulo, "Shim6: Level 3 multihoming shim protocol for IPv6," IETF Internet-Draft, 2008.

[7] C.-M. Huang and C.-H. Tsai, "WiMP-SCTP: Multi-path transmission using stream control transmission protocol (SCTP) in wireless networks," in *AINA Workshops (1)*, 2007, pp. 209–214.

[8] I. Gapanova, "Concurrent multipath transfer for heterogeneous networks," Master's thesis, TU Braunschweig, 2008.

[9] F. Perotto, C. Casetti, and G. Galante, "SCTP-based transport protocols for concurrent multipath transfer," *WCNC*, 2007.

[10] J. Liao, J. Wang, and X. Zhu, "cmpSCTP: An extension of SCTP to support concurrent multi-path transfer," *ICC*, 2008.

[11] B. Zhou, C. P. Fu, and V. Li, "TFRC Veno: An enhancement of TCP friendly rate control over wired/wireless networks," *ICNP 2007*, pp. 216–225, Oct. 2007.

[12] "Modular architecture and economic modeling of multilink networks (MANGO)," [Online] http://tinyurl.com/6yelam.