

Invited Paper

Potential benefit of flow-based routing in multihomed environments

Vlad Manilici*, Andreas Wundsam, Anja Feldmann and Pablo Vidales

Deutsche Telekom Laboratories, TU Berlin, Berlin, Germany

SUMMARY

The data rates provisioned by broadband Internet access connections continue to fall short of the requirements posed by emerging applications. However, the potential of statistical multiplexing of the last mile broadband connections remains unexploited even as the average utilisation of these connections remains low. Despite recent work in this area, two key questions remain unanswered: (a) what is the attainable benefit of broadband access sharing? and (b) how much of this benefit is realisable given real-world constraints? In this work we quantify the attainable benefit of a multihomed broadband access environment by proposing and evaluating several flow-based access sharing policies using a custom flow-based simulator. We then analyse how much of the performance benefit is lost due to real-world constraints by migrating from simulations to a test-lab environment employing a wireless network. Our results show that in today's broadband Internet access scenarios, a significant reduction in download times (up to a factor of 3) is achievable.
Copyright © 2009 John Wiley & Sons, Ltd.

1. INTRODUCTION

Driven by the wide adoption of wireless-enabled broadband access connections and the fact that time-based tariffs have all but disappeared, a new, opportunistic, low-mobility communication paradigm [1, 2] has emerged over the past few years, in which residential users *share* their broadband access links over wireless 802.11 connections with neighbours and other nomadic urban users. Since the average utilisation of a residential Internet connection still remains low [3], statistical multiplexing of user traffic across multiple broadband access connections over a shared wireless medium can be used to accommodate nomadic users, increase the peak data rates, and satisfy high instantaneous traffic demand of bandwidth-intensive applications.

Consider the scenario shown in Figure 1. It depicts a community of users accessing the Internet via DSL links provisioned by different ISPs that share an ad hoc wireless LAN. User 1 may experience congestion on her downlink connection due to serving a large download from mobile user Ext. 1. However, if User 3's Internet connection has

spare bandwidth, User 1 may be able to redirect her traffic to User 3's connection via multihoming.

Realities of current network infrastructure design and operation, such as inability of end users to reserve public IP addresses and lack of ISP support for packet striping, limit the extent to which statistical multiplexing can be leveraged for bandwidth aggregation. To overcome these limitations, several multihoming mechanisms, e.g. [4, 5], have been proposed to route one user's flows via a wireless shared medium through another user's broadband connection. Complementary to these efforts, our work investigates the potential rather than a mechanism for multihoming and attempts to answer a more fundamental question: what is the *attainable* benefit of flow-based re-routing in multihomed environments?

To answer this question, we employ hybrid simulations and real testbed experiments in various settings, using both synthetic and measurement-based traces for representing user traffic. We have developed a flow-level simulator, *FlowSim*, which simulates various flow routing algorithms based on actual flow traces from real networks as input,

* Correspondence to: Vlad Manilici, TU Berlin Deutsche Telekom Laboratories, Sekr. TEL 16, Ernst-Reuter-Platz 7, 10587 Berlin, Germany. E-mail: vlad@net.t-labs.tu-berlin.de.




Figure 1. Community network: shared broadband access.

and deployed a testbed using commodity wireless home gateways and DSL links. Synthetic Web traces generated using Surge [6] and residential access traces from the Munich Scientific Network (MWN) serve as traffic input. Our results show that for bandwidth intensive flows the average download time can be reduced by up to a factor of 3. This emphasises the substantial potential benefit of flow multiplexing among broadband connections in community networks. We expect further improvements by considering different delay distributions via different providers. But this is beyond the scope of this paper. To summarize, this work makes the following contributions:

- A flow-based simulator and an experimental setup for evaluating multihoming: reproducible results are obtained and the difference between ideal and progressively realistic experiments is quantified.
- Several multihoming strategies, ranging from idealistic ones, in which full knowledge of links and flows is assumed, to a practical and realisable one that is suitable for implementation on a router.
- An evaluation of the benefits and potential of the various flow routing strategies in multihomed environments using synthetic and real traffic loads.

The remainder of the paper is structured as follows: Section 2 discusses related work. In Section 3 we introduce our experimental approach and the simulator, and describe our testbed. Simulator and testbed results using measurement-based and Synthetic Web traces as input are discussed in Section 4. The paper concludes in Section 5.

2. RELATED WORK

The areas of local community networking and connectivity sharing have been well explored and some enterprise

solutions [7–9] have even been deployed in the real world. Settings where neighbours share a DSL connection are common, although not contractually allowed by most ISPs [10]. Bundling connectivity by utilizing multiple links in parallel is often used in access networks as well as backbones, where it is commonly known by the term ‘trunking’ or ‘channeling’. The mechanism assumes that the links are terminated at the same devices.

Commercial efforts similar to flow routing and targeted to centrally administered enterprises use the name *smart routing*. Smart routing is marketed by a number of companies, including Internap [11], Radware [12], Angran [13], Mushroom Networks [14] and Viprinet [15]. The proposed algorithms, independent of the general Internet routing, focus on cost in addition to network performance. Goldberg et al. [16] analyse the performance of smart routing and its potential both for interfering with BGP and for self-interference. The authors of Reference [17] show that smart routing can bring economical benefits even to the ISPs.

Akella et al. [7] find that multihoming has the potential of improving throughput performance by up to 25% compared to using only the best ISP connection. In addition, Akella et al. [18] report on an upper bound for the possible performance improvements with multihoming. It is roughly 40%. They also show that a careful choice of providers is necessary.

Closely related to flow management, MAR [4] provides a framework for leveraging the diversity of wireless access providers and channels to provide improved wireless data performance for the commuters via a wireless multihomed device that can be deployed in moving vehicles. However, their packet scheduling architecture requires the presence of a proxy server as a concentrator and support from the ISP side.

The MultiNet project [19] proposes a multihomed single-card device by means of virtualisation. In MultiNet the wireless card is continuously switched across multiple networks. In contrast to physically accessing multiple networks, our work brings forward the idea of routing flows across the network in order to achieve better performance.

A layer-4 protocol that supports multihoming intrinsically is SCTP, the *Stream Control Transmission Protocol* introduced by Stewart [20] and documented by RFC 4960 [21]. SCTP was primarily designed for Public Switched Telephone Network (PSTN) signaling messages over IP and still has to see wide deployment into popular TCP/IP protocol stack implementations. The address management takes place during the setup of the association and cannot be changed later.

pTCP [22] is an extension of TCP that enables bandwidth aggregation for multihomed devices. The protocol has been evaluated through simulations. Its mechanism is based on packet buffering and appears to us to require a large management overhead.

Habib et al. [23] propose the resurrection of the session layer for striping data from a single connection over several links. Implementing the proposal, however, requires extensive changes at the OS or the application level, which is also a requirement for SCTP or the IPv6 shim6 layer [24].

Recent work by Thompson et al. [5] evaluates a framework for end-host multihoming with flow routing based on RTT measurements and prediction of the flow sizes. However, the evaluation is limited to a proof of concept system consisting of two nodes and one specific flow routing algorithm. Another example of flow management is the work presented by Tao et al. [25]. They study the feasibility of improving performance by increasing path diversity through flow-based path switching. This work was evaluated using a wired experimental setup but no evaluation in wireless environments has been reported.

Papadopouli and Schulzrinne [26] propose the integration of channels with largely different QoS characteristics for serving streams with adaptable bandwidth. Their prototype operates as a multicast application with variable bandwidth. Standard applications such as WWW and email are unable to take advantage of the proposed system. Lad et al. [27] propose a high-level architecture for sharing connectivity in a coalition peering without discussing its realisation or presenting a performance evaluation.

Implementations of flow-routing so far require specific modifications to the deployed hard- or software base. *OpenFlow* [28], an emerging technology, promises to facilitate these adaptations and simplify implementation and deployment, by enabling an external entity, the controller, to control the forwarding table of Ethernet switches on a per-flow basis.

Complementary to the above approaches, the work reported in this paper offers an evaluation of multihoming in wireless environments. There are a plethora of papers that have reported on the difficulties posed by wireless environments and their impact on urban networks, see [2] or [29] to mention a few.

3. METHODOLOGY

To evaluate the attainable benefit of flow-based routing in multihomed environments, we combine simulations with

real testbed experiments. For our evaluation we derive the workload from user traces from the border router of the MWN to capture real user behaviour. However, since it is difficult to capture the reaction of the applications by only replaying traces we also rely on synthetic workloads generated by the Web traffic generator Surge [6]. This allows us to quantify potential performance improvements as experienced by the users.

3.1. Flow routing strategies

To obtain a baseline regarding the potential benefits of flow routing we investigate the following set of basic flow routing strategies.

Direct: the current practice for residential broadband connections: every flow is routed on the direct link, i.e. the DSL connection of the user originating the flow.

FatPipe: the ‘ideal case’; it bundles all DSL links into one fat pipe with bandwidth equal to the sum of the capacities of the individual connections. The only restriction is that no flow may exceed the bandwidth of the originating DSL connection. This strategy cannot be implemented in a real network and is only supported by the simulator.

MinLargeFlows: a ‘first algorithm’; it tries to minimise the number of bulky flows that share any of the DSL links. Therefore, it assigns each new flow to the link which, at the time the flow arrives, carries the lowest number of flows that have already transmitted some number of bytes, say 8 KB.

Moving flows from one link to another is difficult as it requires either support by all end-systems or transparent movement of the TCP state. Therefore, we do not allow in any algorithm re-routing a flow once it has been assigned to a certain link.

3.2. Flow routing testbed: FlowRoute

A testbed for conducting realistic experiments has been deployed at our site, emulating the setup in Figure 1. Four nodes have been deployed on one floor in our building, coexisting with other wireless networks, as in real residential scenarios. The distribution of the testbed nodes ensures wireless connectivity among them. Each node consists of a wireless router with IEEE 802.11a/b/g interfaces and a client machine which is also equipped with a wireless interface. The routers are directly connected to the Internet via 2 Mbps DSL lines and to the corresponding client machines using an Ethernet interface. The client

based flow routing system `FlowRoute` allows us to evaluate the benefits of flow routing in real-world testbeds. Hence, effects not easily captured in a simulator, such as the impact of different shared interconnection mediums, wired versus wireless Ethernet, etc., can be studied.

For some of the experiments we need the flexibility to tune the bandwidth of the access lines to the Internet. This functionality is provided by using the NistNet network emulator [30]. The NistNet emulator is installed on a separate machine that is accessible by all testbed nodes via the management network.

The software prototype that enables flow routing on the client machines, `FlowRoute` [31], consists of three main components: a preloaded shared library, called `libConnect`, which intercepts all client access to the Berkeley socket layer and delegates routing decisions. `Routed`, the unique central decision making module, has a global view on the link loads and makes the actual routing decisions based on the flow routing strategies previously described. The third component on the client machines, called `Proxy`, is responsible for re-directing the flows between a client and its destination and for informing the routed about the load of all shared DSL connections.

3.3. Simulator: *Flowsim*

We have developed a simulator, called `FlowSim`, to evaluate in a scalable manner the potential of flow routing strategies using flow-level traces as input, captured from real networks or testbed experiments. Popular packet level simulators such as `ns-2` or `SSFNet` are not suitable for our purpose, given the large number of flows considered here. To validate the simulator we reproduce the testbed experiments within `FlowSim` and compare the results.

We can use `FlowSim` to simulate network setups such as those shown in Figure 1, with various number of users, different capacities and delays for uplink and downlink directions and flow routing algorithm parameters. The interconnection capacity between DSL subscribers is not supposed to be a bottleneck and thus modelled as infinite. As such, the simulator does not distinguish if the community is interconnected via a wireless or wired network.

`FlowSim`[†] consists of two main components: a router and a scheduler. The *router* identifies interactive and bandwidth-limited flows. Bandwidth-limited flows are further classified by the direction (referred to as dominant

direction) in which they exceed a threshold into downlink-limited, uplink-limited and both-limited flows. Once a flow has been classified, it is routed using one of the flow routing strategies.

The role of the *scheduler* is to distribute the available bandwidth among competing flows, while approximating the dynamics of the TCP protocol. In order to achieve reasonable performance, it uses a fluid TCP model operating on discrete time slots, rather than on a per event basis.

The simulation results do not show a significant difference when using smaller time slots than 1/10 s, the current setting. The distribution of the available up and downstream bandwidth is done in a fair manner between all flows that use one link. Each bandwidth-limited flow exercises a TCP-style slow start before it is able to attain its fair share of the bandwidth. While connect and close delays are modelled, packet losses are ignored. Comparisons with results obtained from the testbed for the same set of traces used as input show that the approximations are reasonable. For each bandwidth-limited flow its fair share of the bandwidth is computed based on its dominant direction(s). The bandwidth share for the non-dominant direction is then set in proportion to the transmitted bytes. The results are then scaled to the available bandwidth.

As an initial validation of the simulator we, in Figure 2, compare the link utilisation's from a simulation run with that of an experiment run in the testbed. Each subplot shows the normalised bandwidth usage across time for both the simulator as well as the testbed run in different grey shades. With the exception of a few outliers and some lags, the curves match closely. This indicates that the simplifications within `FlowSim` are reasonable.




Figure 2. Link utilisation in terms of normalised bandwidth usage for both a simulator and a testbed run for one of the trace-based experiments and direct routing.

[†] For details and/or the downloadable software see: <http://www.net.t-labs.tu-berlin.de/~vlad/FlowRoutingSoftware>.




Figure 3. CCDF of flow durations with and without routing for the Surge Web workload and the MinLargeFlows strategy.

4. RESULTS

To evaluate the potential benefit of flow routing we explore its behaviour both in the simulator as well as in the testbed under different workloads. We start with a Web workload generated by Surge. Then we switch to a trace-based evaluation relying on traces from the MWN.

4.1. Synthetic Web workload

We use Surge [6] to generate a synthetic workload that resembles Web traffic. We updated its configuration parameters to reflect characteristics of today's Web (e.g. the median and mean HTTP object size of the current Web workload mix, including Web 2.0 and P2P, as observed in the MWN traces). We use the popular Apache2 software as our Web server. To impose a reasonable load we use four Surge instances per household (host). This Web workload results in an average utilisation of 0.39 Mbps per DSL link.

4.1.1. Testbed results. Figure 3 displays the *complementary cumulative distribution function* (CCDF) of the flow durations on a logarithmic X -axis for an experiment with NistNet and a wired interconnection network for the flow routing strategy MinLargeFlows. One can clearly see the benefit of flow routing for longer flows, as durations with flow routing are shorter than without. The corresponding logarithmic *probability density function* (PDF) is shown in Figure 4, highlighting that the flowrouting system imposes some overhead for short flows.

To quantify the achievable benefit, we compute the ratio of the durations, more precisely, we compute $duration(direct)/duration(routed)$. Larger values denote better performance under flow routing. Considering all flows, the mean benefit is 1.11 with a median of 0.90. At first, this




Figure 4. PDF of logarithms of the flow durations with and without routing for the Surge Web workload and the MinLargeFlows strategy.

might seem a bit disappointing. However, there are many short flows and they in this scenario all suffer from the overhead of our prototype flow routing implementation. But when we only consider flows with a duration larger than 0.5 s the mean improvement increased to 2.47 (median 1.83). This implies that there is a significant benefit for bulky flows.

We next switch from NistNet to using the actual 2 Mbps DSL lines, with the wired interconnection network. We observe that the results improve slightly in comparison to the emulated network. For all flows, the mean benefit is 1.20 (median 0.95). Flows that last longer than 0.5 s experience an improvement of 2.49 (median 2.03). The differences are explainable by small inaccuracies in NistNet.

Finally, we use the well connected wireless network in 802.11a mode as our interconnection network within the community. Somewhat surprisingly, the overall results improve slightly when compared to using a wired network. The mean (median) improvements for longer flows are 2.61 (2.11) and for all flows 1.27 (1.0). This shows that the wireless network is not the bottleneck. We reroute only 75% of the flows which imposes a load of less than 5 Mbps on the wireless network. Overall, we see that significant improvements are possible by taking advantage of multihoming.

4.1.2. Simulator results. We investigate the upper bound of the routing performance by using the simulator on the traces gathered from the Web workload experiments. This allows us to estimate the potential benefits and compare the performance prediction of the simulator against the performance improvements observed in the testbed. The simulator displays some deviations in flow duration, which are in part due to the simulator's assumption that bandwidth is shared fairly between flows. This assumption is known to work well on average but not on smaller time scales. As such, it is not surprising that the size of the deviations

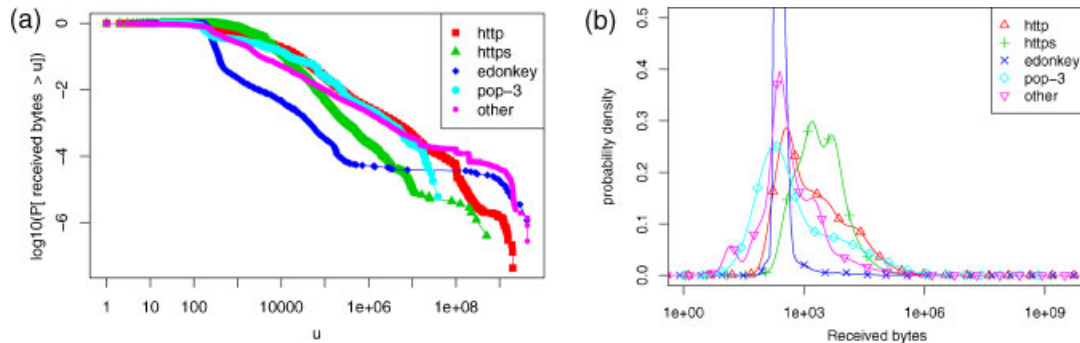


Figure 5. Trace analysis: received bytes per application. (a) CCDF; (b) PDF.

decrease as flow durations increase. In total, the ratio of the durations for bandwidth-limited flows in the simulation versus the testbed experiment has a mean of 1.07 and a median of 0.97. This indicates that, while for individual flows the predicted performance does not agree with the performance seen in the experiment, the overall results of the simulation match the experiment quite nicely.

We then compare the performance as predicted by the simulator with the performance of the actual experiment using the `MinLargeFlows` routing policy. The simulator reports a mean (median) improvement for flows longer than 0.5 s of 2.59 (2.09). In the Synthetic Web workload experiment with the same strategy, the inter-arrival times between Web connections are shorter than with the `Direct` policy, and the load is also higher. These predictions match the actually observed benefits of 2.47 very well. Interestingly, the predicted mean improvement by the simulator for `FatPipe` is only slightly better than for `MinLargeFlows`: with a mean improvement of 2.79 versus 2.24. This confirms that the achievable benefit for this scenario is indeed limited by the traffic properties of this specific Web workload.

4.2. Trace-based experiments

Using traces of real traffic to drive simulations as well as testbed experiments allows repeatability of results under realistic loads.

4.2.1. Workload. We use connection-level summaries of traffic traces captured at the border router of the MWN. The MWN provides 10 Gbps Internet upstream capacity to roughly 50 000 hosts at two major universities including student dormitories, with the daily traffic amounting to 3–6 TB. Ample bandwidth is available to all users of this network via high-speed local area networks. Thus, Internet

usage is not impaired by any bandwidth limitations of the access links.

In a typical 24-h workday trace from 24 April 2007 we identify approximately 37 million flows, out of which 21.1 million have incomplete connection information records, a typical component of today's traffic mix, such as SYN attempts (56%), rejected connections (27%), in progress connections (8%), and missing information (9%). For our experiments, we consider the remaining 15.9 million flows, through which 641 GB (182 GB) were transferred from (to) the Internet. These volumes correspond to an average bandwidth utilisation of 60 Mbps (17 Mbps) downstream (upstream).

To better understand the characteristics of the traffic, we classify the flows according to the application that most likely generated them, as identified by the port number. About 73.50% of the flows are HTTP, 7.83% HTTPS, 2.74% eDonkey, 0.52% POP-3 and 15.41% are other traffic. Figure 5(a) shows the CCDF (complimentary cumulative distribution function) of the received bytes for different applications on a log–log scale. We observe that the flow sizes of the applications are consistent with the heavy-tailed distributions that have previously been reported, e.g. [32]. The behaviour of eDonkey may seem strange at first, but this is due to its two traffic components—the control channel and data transfer connection. Part of the byte contributions of other P2P file sharing protocols are captured by the 'Other' class; hence, it is not surprising that the tails of the two curves coincide. The other P2P traffic is contained in the 'HTTP' class. The mean number of bytes over all connections is 43 409 and the median 795. Similar observations hold for the bytes sent per flow.

To investigate the actual flow size distributions, we plot the PDF of the logarithm of the transfer sizes on a loga-




Figure 6. Scatter plots of flow durations with and without routing for the MWN traces and the `MinLargeFlows` strategy. In addition, a contour plot is superimposed. (a) Full MWN trace; (b) 1/3MWN trace.

rhythmic X -axis, in Figure 5(b). HTTP/HTTPS exhibits the expected distribution with a small spike that corresponds to the typical header size. HTTPS has a larger mean but a smaller median. POP3 transfers are smaller, while eDonkey is on the extreme end with many short connections due to control traffic and several larger ones due to data transfers. The ‘Other’ category is a mix of unidentified flow types, which also seems to contain a significant amount of P2P traffic.

From this large collection of flow data we selected a subset of flows, referred to as *MWN* flow trace, that originated from IP addresses that are assigned to student residences, to ensure that we consider only genuine residential traffic. Students are provided Internet access only via 28 NAT gateways. The traffic via those NAT gateways imposes a load of about 0.34 Mbps on the upstream and 5.74 Mbps on the downstream. We assigned these 28 NAT gateways to our DSL links randomly. Using such a dense trace allows us to investigate times with peak traffic while keeping the durations of our experiments relatively short (30 min).

To quantify the performance of our flow re-routing against different traffic loads, we run experiments with the `MinLargeFlows` strategy on the testbed using the full MWN trace in addition to a random subselection of 1/3 of the total number of flows (reduced set).

4.2.2. Performance analysis. Given that the short flows suffer from the overhead imposed by the specific implementation chosen for our prototype we again concentrate on bulky flows—those that last longer than 0.5 s. Figure 6(a) and (b) show scatter plots of flow durations. Each point represents a flow, with the duration of the direct routing policy on the X -axis and the duration of the re-routed policy on the Y -axis. The overlaid contour lines plot the density of

Table 1. Summary of experimental results with regards to the achievable benefit for bulky flows (>0.5 sec.) under different workloads and experimental settings for both simulator and testbed.

Experiment	Policy	Mean	Median
Surge (Ethernet, NistNet)	<code>MinLargeFlows</code>	2.47	1.83
Surge (Ethernet, DSL)	<code>MinLargeFlows</code>	2.49	2.03
Surge (Wlan, NistNet)	<code>MinLargeFlows</code>	2.61	2.11
Surge (Simulator)	<code>MinLargeFlows</code>	2.59	2.09
Surge (Simulator)	<code>FatPipe</code>	2.79	2.24
1/3 MWN (Ethernet, NistNet.)	<code>MinLargeFlows</code>	2.04	1.02
MWN (Ethernet, NistNet)	<code>MinLargeFlows</code>	3.02	1.98
MWN (Simulator)	<code>MinLargeFlows</code>	2.47	1.99
MWN (Simulator)	<code>FatPipe</code>	3.73	2.40

measurement points in one region. More points below the $x = y$ diagonal imply more flows that benefit from routing. The density peak for the experiment with the full MWN trace is significantly below the diagonal, whereas the peak for the reduced MWN trace is closer to the diagonal.

The experiment shows improvements with respect to the flow durations: for the former case a mean factor of 3.02 (median 1.98) and in the latter case by a factor of 2.04 (median 1.02). When considering all flows including the short ones, the mean/median values are for the full MWN trace 2.21/1.12 and 1.17/0.94 for the reduced set. See Table 1 for a summary.

We plot the logarithmic density (PDF) of the flow durations ratio: $duration(direct)/duration(routed)$ in Figure 7. We note that the shift towards right (values larger than 1) is much more pronounced for the full trace experiment. This means that flow routing performs better in regions close to the link congestion. Figure 8 shows the CCDF of flow




Figure 7. Flow duration densities with different workloads.




Figure 8. CCDF of flow durations with and without routing for the MWN trace-based workload.

durations with and without flow routing using the `MinLargeFlows` strategy. We observe that the flows exhibit significantly shorter durations when flow routing is used.

The experimental results in the testbed for the standard `MinLargeFlows` routing algorithm compare well with the simulation. The median ratio of flow durations is 1.98 in the testbed and 1.99 in the simulator. Mean ratio of 3.02 is better in the testbed than 2.47 achieved in the simulator. This is partly due to the larger jitters that we observe in the testbed. The `FatPipe` policy achieves flow duration improvements by a mean factor of 3.73 (median 2.40) in the simulator.

Next, we explore the effect of the shared interconnection medium, wireless or wired. For this purpose we rely on the testbed. Average and median flow durations vary by less than 15% between a wireless shared medium with good connectivity and a wired one. The DSL emulation through `NistNet` achieves less than 3% average and median variation in flow lengths. Figure 9 compares the density distribution




Figure 9. Density of flow durations on different network media for the MWN experiment with routing.

of the flow durations when using DSL lines or `NistNet` for the Internet connection and Ethernet or Wireless for the testbed interconnections.

5. CONCLUSION

Broadband access connection sharing using wireless 802.11 technologies is emerging as an alternative approach towards nomadic broadband connectivity. Flow-based routing in conjunction with multihoming plays a key role in enabling such access sharing scenarios, by increasing the multiplexing of broadband connections. For this reason, it is important to evaluate the potential benefits of flow-based routing in multihomed environments.

In this work, we evaluate several flow re-routing strategies using synthetic workloads and captured user traces as traffic input. Our encouraging results show improvements of up to a factor of 3 in the download times of bandwidth-limited flows, with very small overhead. Bandwidth-demanding flows and an unbalanced distribution of load among the users of the sharing community particularly benefit from the flow-based routing approach. More importantly, full knowledge of flow characteristics is not needed, as indicated by the performance of `MinLargeFlows`.

ACKNOWLEDGMENT

The authors thank the Extended Hotspots team, including Petros Zerfos, Jatinder Singh and Robert Sombrutzki, for many fruitful discussions and help in setting up the experiments. This work was supported in part by a grant from Deutsche Telekom.

REFERENCES

1. Sastry N, Crowcroft J, Sollins K. Architecting citywide ubiquitous Wi-Fi access. In *ACM HotNets*, 2007.
2. Solarzski M, Vidales P, Schneider O, Zerfos P, Singh JP. An experimental evaluation of urban networking using IEEE 802.11 technology. In *Proceedings of IEEE OpComm*, 2006.
3. Odlyzko AM. Data networks are lightly utilized, and will stay that way. *Review of Network Economics*, 2003.
4. Rodriguez P, Chakravorty R, Chesterfield J, Pratt I, Banerjee S. Mar: a commuter router infrastructure for the mobile internet. In *MobiSys*, 2004.
5. Thompson N, He G, Luo H. Flow scheduling for end-host multihoming. In *IEEE Infocom*, 2006.
6. Barford P, Crovella M. Generating representative web workloads for network and server performance evaluation. In *SIGMETRICS/PERFORMANCE*, 1998.
7. Akella A, Seshan S, Shaikh A. Multihoming performance benefits: an experimental evaluation of practical enterprise strategies. In *USENIX*, 2004.
8. Goldenberg D, Qiu L, Xie H, Yang YR, Zhang Y. Optimizing cost and performance for multihoming. In *ACM SIGCOMM*, 2004.
9. Guo F, Chen J, Li W, Chiueh T-C. Experiences in building a multihoming load balancing system. In *INFOCOM*, 2004.
10. BT Fon. WiFi for everyone. website: <http://btfon.com>
11. Internap. <http://www.internap.com>
12. Radware. <http://www.radware.com>
13. Anagran. <http://www.anagran.com>
14. Mushroom Networks. <http://www.mushroomnetworks.com>
15. viprinet. <http://www.viprinet.com>
16. Goldenberg DK, Qiu L, Xie H, Yang YR, Zhang Y. Optimizing cost and performance for multihoming. In *ACM SIGCOMM*, 2004.
17. Dai R, Stahl DO, Whinston AB. The economics of smart routing and quality of service. *Group Communications and Charges: Technology and Business Models: 5th COST 264 International Workshop on Networked Group Communications, NGC 2003 and 3rd International Workshop on Internet Charging and QoS Technologies, ICQT 2003, Munich, Germany, 16–19 September 2003: Proceedings*, 2003.
18. Akella A, Maggs B, Seshan S, Shaikh A, Sitaraman RK. A measurement-based analysis of multihoming. In *Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, ACM: New York, 2003; 353–364.
19. Chandra R, Bahl V, Bahl P. Multinet: connecting to multiple IEEE 802.11 networks using a single wireless card. In *INFOCOM*, 2004.
20. Stewart R, Metz C. SCTP: new transport protocol for TCP/IP. *IEEE Internet Computing*, 2001.
21. Stewart R. Stream Control Transmission Protocol. RFC 4960 (Proposed Standard), September 2007.
22. Hsieh H-Y, Sivakumar R. A transport layer approach for achieving aggregate bandwidths on multi-homed mobile hosts. In *ACM MOBICOM*, 2002.
23. Habib A, Christin A, Chuang J. Taking advantage of multihoming with session layer striping. In *IEEE Global Internet*, 2006.
24. Bagnulo M, Nordmark E. Shim6: Level 3 multihoming shim protocol for IPv6. *draft-ietf-shim6-proto-08 (work in progress)*, April, 2007.
25. Tao S, Xu K, Xu Y, et al. Exploring the performance benefits of end-to-end path switching. In *IEEE ICNP*, 2004.
26. Papadopouli M, Schulzrinne H. Connection sharing in an ad hoc wireless network among collaborating hosts. In *NOSSDAV*, 1999.
27. Lad M, Bhatti S, Hailes S, Kirstein P. Coalition-based peering for flexible connectivity. In *PIMRC*, 2006.
28. McKeown N, Anderson T, Balakrishnan H, et al. OpenFlow: enabling innovation in campus networks. *ACM SIGCOMM Computer Communication Review* 2008; **38**(2):69–74.
29. Bychkovsky V, Hull B, Miu AK, Balakrishnan H, Madden S. A measurement study of vehicular internet access using in situ wi-fi Networks. In *MOBICOM*, 2006.
30. Nistnet. <http://www-x.antd.nist.gov/nistnet>
31. Wundsam A. Connection sharing in community networks—how to accommodate peak bandwidth demands. Diplomarbeit, TU Munich, Germany, April 2007.
32. Crovella ME, Bestavros A. Self-similarity in world wide web traffic: evidence and possible causes. *IEEE/ACM Transactions on networking*, 1997.

AUTHORS' BIOGRAPHIES

Vlad Manilici is a member of staff of the German Aerospace Center (DLR). He received his diploma in Computer Science from the Technical University of Cluj, Romania in 1999. He expects to receive his Ph.D. from TU-Berlin (Deutsche Telekom Laboratories) in 2009. His research interests include community networks and flow based routing.

Andreas Wundsam is a researcher and Ph.D. candidate at TU-Berlin/Deutsche Telekom Laboratories. He received his diploma in Computer Science from the TU München, Germany in 2007. His current research interests include network virtualisation, flow based routing and future Internet architectures.

Anja Feldmann is a full professor at Deutsche Telekom Laboratories a unit of Deutsche Telekom and an An-Institut of the Technische Universität Berlin, Germany. From 2000 to 2006 she headed the network architectures group first at Saarland University and then at TU München. Before that (1995–1999) she was a member of the Networking and Distributed Systems Center at AT&T Labs—Research in Florham Park, New Jersey. She has published more than 50 papers and has served on more than 40 programme committees, including as Co-Chair of Sigcomm 2003 and as Co-PC-Chair of Sigcomm 2006 and IMC 2009. She received M.S. degree in Computer Science from the University of Paderborn, Paderborn, Germany, in 1990 and M.S. and Ph.D. degrees in Computer Science from Carnegie Mellon University in Pittsburgh, USA, in 1991 and 1995, respectively. Her current research interests include new network architectures as well as understanding the current Internet and its application for the purpose of performance debugging and intrusion prevention.

Pablo Vidales is a senior research scientist at Deutsche Telekom Laboratories a unit of Deutsche Telekom and an An-Institut of the Technische Universität Berlin, Germany. He joined Deutsche Telekom Laboratories in May 2005 and since then he has been strongly involved in academic and industrial projects aiming to improve wireless communications. He has published over 30 scientific papers and he is co-inventor of more than eight patents in the area of wireless distributed systems, and has served on numerous programme committees. He received M.S degree in Computer Science and M.S in Telecommunications from the Mexico Autonomous Institute of Technology (commonly know as ITAM), and Ph.D. degree in Electrical Engineering from the University of Cambridge, UK, in 2001 and 2005, respectively. His current research interests include delay tolerant networks, QoE and QoS in next generation networks, and future mobile services. The National Researchers Institute in Mexico (SNI) has awarded Dr Vidales the nomination to National Researcher in recognition of his ability to perform scientific research.