

On-the-Fly TCP Acceleration with Miniproxy *

Giuseppe Siracusano
Univ. of Rome Tor Vergata/
NEC Laboratories Europe,
Germany

Roberto Bifulco
NEC Laboratories Europe,
Germany

Simon Kuenzer
NEC Laboratories Europe,
Germany

Stefano Salsano
Univ. of Rome Tor Vergata/
CNIT, Italy

Nicola Blefari Melazzi
Univ. of Rome Tor Vergata/
CNIT, Italy

Felipe Huici
NEC Laboratories Europe,
Germany

ABSTRACT

TCP proxies are basic building blocks for many advanced middleboxes. In this paper we present Miniproxy, a TCP proxy built on top of a specialized minimalistic cloud operating system. Miniproxy's connection handling performance is comparable to that of full-fledged GNU/Linux TCP proxy implementations, but its minimalistic footprint enables new use cases. Specifically, Miniproxy requires as little as 6 MB to run and boots in tens of milliseconds, enabling massive consolidation, on-the-fly instantiation and edge cloud computing scenarios. We demonstrate the benefits of Miniproxy by implementing and evaluating a TCP acceleration use case.

CCS Concepts

•**Networks** → **Middle boxes / network appliances**; *Transport protocols*; *Network experimentation*;

Keywords

Middleboxes, TCP proxy, TCP Acceleration, HTTP proxy, Early SYN Forwarding, Unikernels.

1. INTRODUCTION

Service access times are directly correlated to users' experience and thus to service providers' revenues [18]. Amazon estimates that an increase of delay of 100ms cuts its revenue by 1% [7]. Google measured a 0.74% drop in the number of web searches performed by users when the search service delay was artificially increased by 400ms [27]. Similarly, Bing experienced a reduction of 1.2% in per-user revenue

*This work was partly funded by the EU in the context of the SUPERFLUIDITY project (5G PPP).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HotMiddlebox, August 22–26, 2016, Florianopolis, Brazil

© 2016 ACM. ISBN 978-1-4503-4424-1/16/08...\$15.00

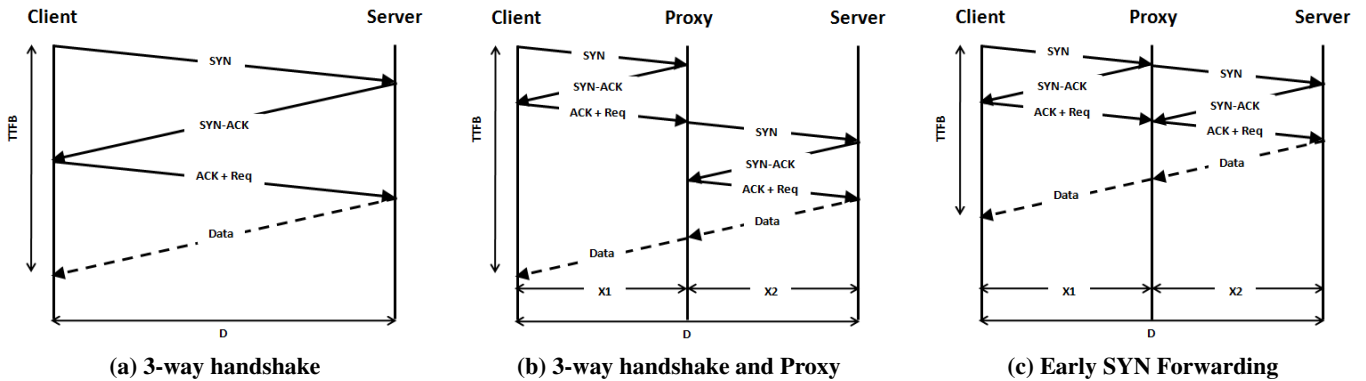
DOI: <http://dx.doi.org/10.1145/2940147.2940149>

when the service delay was increased by 500ms [20]. Given the drastic impact that a few additional milliseconds can have, network performance constitutes a critical element for many Internet services.

In today's networks, latency is dominated by two components: round-trip time (RTT) between the communication's end-points and the number of RTTs required to complete the transfer [7]. The RTT is determined by delay in the physical infrastructure, routing and queuing. In this work, we focus on reducing the number of RTTs required to complete a data transfer. In the case of TCP, optimizations have targeted most of the protocol mechanisms including connection establishment [17, 9], slow start [2] and congestion avoidance [7]. However, since parts of TCP are fundamental to its correct operation, the optimization space is constrained or may require extensive changes to the network infrastructure and protocol stack [3, 27], making deployment harder.

A complementary approach to TCP optimization is the deployment of TCP proxies within the path of an end-to-end connection [11]. A TCP proxy splits a single TCP connection into two connections and, if located in the middle (delay-wise) of the original connection's path, can noticeably speed up end-to-end communication by reducing the feedback-loop time of each TCP connection [9]. However, implementations of this optimization technique have been deployed only in cases in which the requirement for the optimization was known and well established in advance, as is the case for content distribution networks (CDNs) [1, 16]. Dynamic, on-the-fly deployments of TCP proxies for connection acceleration have only been explored, so far, as a service provided by "enhanced" routers [9]. Unfortunately, the deployment of such a solution is difficult, as it requires modification to the routers, and impractical, since it assumes that network flows do not undergo path changes due to routing throughout the lifetime of a flow.

Despite this state of affairs, recent trends are transforming the network into a cloud infrastructure that allows for running a variety of services in a number of different locations [14]; these provide new opportunities for on-the-fly deployment of proxies for TCP acceleration, at the right places in the network and in a timely fashion. To leverage these flexible cloud infrastructures, TCP proxies need to be virtualized,



(a) 3-way handshake (b) 3-way handshake and Proxy (c) Early SYN Forwarding
Figure 1: 3-way handshake, 3-way handshake and Proxy, Early SYN Forwarding

while still offering good performance and scalability.

In this paper we present Miniproxy, a lightweight, virtualized TCP proxy that can support scenarios such as on-the-fly TCP acceleration. Miniproxy is a Xen unikernel [23], is as fast as state-of-the-art GNU/Linux-based proxies, requires only 6 MB of RAM to run and can boot in just 12ms. Using it as a building block, we demonstrate that it is indeed possible to accelerate a TCP connection by deploying one or more proxies on an end-to-end path, even when placement of the proxies requires some deviation from the shortest path. Thanks to its small boot times, Miniproxy allows for the creation of chains of TCP proxies *just-in-time*, by booting instances at locations in the network convenient for TCP acceleration.

2. BACKGROUND & RELATED WORK

This section briefly introduces the technique we leverage to accelerate TCP connections, sometime called *Split TCP*, and presents related work.

2.1 TCP acceleration

In general, a TCP connection comprises three phases: connection establishment (3-way handshake), slow start and congestion avoidance. After connection establishment, throughput increases exponentially during the slow start phase until it reaches a threshold, after which the congestion avoidance phase starts. For short flows, a TCP connection may terminate before reaching such threshold, which is the case for the large majority (about 90%) of flows [2]; as a result, we mainly focus on optimizing the first two phases. In [19] we provide the analytic modeling of the TCP connection establishment time with the detailed analysis of the savings that are obtained with the considered TCP acceleration techniques.

Connection Establishment. For short flows, the 3-way handshake (see figure 1a) may constitute a significant part of the overall TCP flow’s Time To Complete (TTC), i.e., the time it takes to open a connection, transfer all the data and close it. Consequently, the Time To First Byte (TTFB) of a TCP connection, i.e., the time the client waits for receiving the first data packet from the server, requires at least 2 RTTs.

Assume now a scenario with a TCP proxy. The client to server connection establishment requires now two different 3-way handshakes (see figure 1b). Usually, the two 3-way

handshakes are executed in sequence: the first one happens between the client and the proxy, and only after that the second one, between the proxy and the server, takes place.

Adopting the technique proposed in [9], TCP proxies can speedup the connection establishment. In this approach, the client’s TCP SYN packet is forwarded to the next TCP hop as soon as it is received by a proxy (see figure 1c). We refer to this mechanism as *Early SYN Forwarding* (ESF). With ESF, the two 3-way handshakes are partially executed in parallel, reducing the overall connection establishment time and TTFB.

Slow start The duration of the slow start phase is directly proportional to the RTT. Introducing a proxy in the path splits the TCP connection in two. In each of the split connections, the slow start phase starts when the first data segment is received. The data segment is immediately forwarded downstream so that the slow start phases in each part proceed in parallel. Thus, the proxy effectively accelerates the TCP slow start phase. For instance, if the proxy splits the end-to-end delay in even parts, the duration of the slow start phase is halved.

2.2 Related work

Our work follows the trend of using specialized VMs, also known as unikernels [23, 15, 6, 8, 12], for creating virtualized network functions. As opposed to previous work, we focus on transparently accelerating TCP connections and instantiating such servers on-the-fly. More recently, Jitsu [13] and the work in [21] proposed the instantiation of VMs on-demand; we leverage similar mechanisms to show that (virtual) TCP proxies can be instantiated just-in-time in order to improve the performance of TCP flows.

Our work is also related to the vast literature regarding TCP acceleration. Specifically, we leverage the results of [9] and [11] to perform TCP acceleration, though we use a novel TCP proxy implementation, which can be dynamically deployed in virtualized network infrastructures.

The work in [17] and [2] optimizes the TCP connection establishment and slow start phases and so its results are complementary to ours. In particular, one of the additions of TCP Fast Open [17] is the possibility for end-hosts to include some data in a TCP SYN packet. The data can be used to infer the final destination of a proxied connection, when the proxy is explicitly addressed. We developed a new TCP option to

carry the final destination of a proxied connection, in order to allow proxies to use ESF (§4). However, our mechanisms work for both new and re-established connections, while TCP Fast Open can be used only with the latter.

In [16] TCP proxies are deployed in fixed locations to reduce the connection time; instead, we make the case for dynamically-placed proxies.

Another common use of Split TCP is to isolate lossy links [5, 10, 4]. We do not cover these aspects in this paper, however, the observations made in these related works can be leveraged and applied also to our case, if lossy links are comprised in the end-to-end path.

3. SYSTEM OVERVIEW

In this section, we present an overview of a system that leverages Miniproxy to provide on-the-fly TCP acceleration. Although this paper is mainly focused on the implementation of the Miniproxy itself, we believe a description of one of its applications will provide a clearer understanding of the advantages of the proposed technology.

The aim of the envisioned system is to accelerate an end-to-end TCP connection by deploying a chain of TCP proxies. For such an objective, it is critical to properly locate the proxies. In fact, split TCP achieves best performance when minimizing the overall end-to-end RTT (i.e., the sum of the individual connections' RTTs) with an even distribution of the connections' RTTs. However, the best locations are strictly dependent on the actual end-points' relative locations. That is, each end-to-end connection has its own best proxies' locations. Furthermore, for any given couple of end-points, those locations may need to be changed over time because of the changing traffic conditions. In the light of these considerations, it seems evident that a network provider is the best suited actor to provide such a service. However, given that the TCP flows being more suitable for acceleration are those that experience large RTTs, it is neat to assume that the proxies' locations may be distributed among several network providers. Hence it is difficult to deploy this acceleration technique as a network infrastructure service [9], since it would require extended coordination between different network providers.

Luckily, today's Internet offers a set of locations for flexibly deploying TCP proxies: publicly available cloud datacenters. Providers such as Amazon, Microsoft and Google are just the most well-known ones, but a number of national and regional providers can be easily added to the list. Using the cloud to run proxies does not just solve the locations issue, but actually includes a new variable in the picture. Cloud providers offer a flexible utility-based approach to run VMs, with a fine-grained billing model (sometimes taking into account the per-minute activity of each VM's cpu). Clearly, such a cost model strongly calls for a system that is able to run a VM only for the required amount of time, when it is actually required.

Requirements. Achieving such objectives hinges on fulfilling a number of requirements. First, we would like to leverage the availability of cloud deployments and virtualized infrastructure in order to ease deployment and obtain the best acceleration performance out of our proxies; consequently,

the proxies should be virtualized in order to run over such platforms. Second, using cloud deployments requires minimizing the time each proxy (i.e., VM) runs, in order to lower the costs of running such instances when employing a utility cost model. For example, this suggests that proactively deploying the proxies may not be an optimal solution from a cost perspective. Furthermore, for new TCP connections, at least one of the TCP connection's end-points is unknown before the connection is initiated; thus, it is not possible to know, in advance, the best locations for the proxies. Since our gains can only come from the connection establishment and slow start phases, a third crucial requirement is to be able to place the proxies at their optimal locations in very short timescales. Fourth, each pair of end-points may potentially require different proxy locations and those locations can change over time due to varying network conditions; this may require running many proxies at many locations, flexibly scaling their numbers, to optimally accelerate a set of connections. Finally, a big number of proxies means that proxy instances should be as lightweight as possible to reduce the cost of running them and to achieve fine-grained scalability of the system.

Building blocks. The implementation of a system that fulfills the above requirements is technically challenging and requires a number of building blocks. For instance, a building block is a monitoring system that checks the network conditions between the suitable locations; the collected information would be then required by another building block, e.g., an orchestration system, which uses this information to select the best locations to run the proxies. A system that manages and timely creates the proxies is also required, as well as efficient mechanisms to chain together several proxies.

This list could be actually much longer, and it is out of the scope of this work to describe a full-fledged solution. Instead, in the next section we focus on the basic building block for such a system, i.e., a virtualized TCP proxy.

4. MINIPROXY IMPLEMENTATION

Given the above requirements, we would like to leverage the nice properties of VMs such as isolation but without incurring their overheads. To do so, we settle on unikernels [23]: purpose-built, specialized VMs based on minimalistic OSes. Unikernels have a number of advantages including a single address space, so no expensive system calls; low memory footprint (MBs or even KBs) and fast instantiation times (milliseconds compared to seconds for conventional VMs). These properties are important since they let us instantiate our TCP proxies on demand, and potentially a large number of them on the same box.

In particular, we target unikernels on Xen, and use the paravirtualized MiniOS [24] operating system to build on (see figure 2). To implement Miniproxy, we make a number of modifications to `lwip` (a small, open source TCP/IP stack for embedded systems) to handle the actual TCP connections, and develop a MiniOS-based proxy application from scratch that leverages these modifications. We further make use of the optimizations to Xen and MiniOS described in [14] in order to derive even smaller boot times.

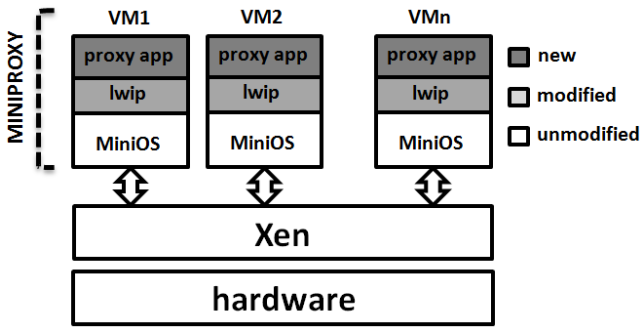


Figure 2: Miniproxy architecture showing modified and new components.

In greater detail, we modify `lwip` to model a TCP proxy as pairs of sockets, one for the connection between the client (or previous hop) and the proxy and one for the connection between the proxy and the server (or next hop). In other words, the first socket is used to receive incoming connections while the second one is used to open a connection towards the server (or the next hop). To implement this in `lwip`, we linked together two PCBs (Protocol Control Block) structures. Our API allows us to instantiate a PCB pair through a new bind function called `tcp_early_syn_bind()`. When a PCB pair is instantiated, each PCB plays a specific role: one listens for incoming connections, while the other one is ready to forward the SYN packet (`outgoing_idle` state). Each of the PCBs performs a separate three-way connection establishment process and invokes an application callback when done. The listening PCB hands over information about the connection with the endpoint that has started the communication; likewise, the `outgoing_idle` PCB includes state to support communication with the target endpoint.

In addition to this change, we added a new callback to `lwip` that is invoked when the SYN is received, allowing applications to instantiate data structures and take decisions about the next hop, if needed. We further increase the TCP IW to 10 segments, the same value used by the majority of the TCP implementations. We set the TCP send buffer size, configurable by the user, to double the TCP receive window, as this resulted in the best performance in our experimental tests.

The other major change to `lwip` has to do with the introduction of a new TCP option that allows us to chain together a set of explicit proxies. Briefly, there are two ways to send traffic through proxies: implicit (or transparent) [9] and explicit [11]. In the implicit case the flow’s destination can be inferred from packet headers, i.e., destination IP address and TCP port, which means that the proxy must be on the path between the client and the packet’s IP destination. The advantage is that all the information to identify the final endpoint of the TCP connection is in the packet’s header. In the explicit case, instead, the IP destination of the TCP connection is the proxy itself, meaning that the proxy has to read the TCP flow’s data to learn the final destination of the connection. For example, when using HTTP proxies, the HTTP header includes the destination server in the URI. Unfortunately, while explicit proxies can be deployed anywhere, they

do not allow for ESF, since they need to complete the first 3-way handshake to read the TCP flow’s data. To overcome this limitation, we introduce a new TCP option that carries information about the source/destination IP addresses and TCP ports inside a SYN packet. The use of this option extends the support for ESF also to explicit proxies.

The TCP option can be introduced directly by the client or by an implicit proxy which is on-path and can redirect the connection to an explicit proxy. Note that at least one implicit proxy on the path is already commonly deployed, e.g., in cellular networks [22, 26]. In this case, the first proxy (likely close to the client), will perform a traditional 3-way handshake, while all the remaining proxies in the chain, if supporting the newly introduced TCP option, can use ESF. Miniproxy supports the new TCP Option, which contains the 4-tuple IP source/destination address and TCP source/destination port.

The implementation of the actual proxy application running on top of MiniOS and our modified `lwip` stack consists of 600 LoC, is event-driven, requires a minimum of 6MB of RAM, and does not need any block devices. Besides performing the TCP acceleration already described, the proxy is also able to parse the new TCP option and apply per-flow policies. Finally, it is worth pointing out that in the current state Miniproxy performs only TCP proxying, but it has been designed for extensibility in order to eventually introduce additional functions, e.g., caching.

5. EVALUATION

In this section we present a preliminary performance evaluation of our Miniproxy implementation. We first perform micro-benchmarks of Miniproxy’s (1) throughput, (2) connection establishment time and (3) boot times, followed by an evaluation of how applying ESF and a chain of proxies affects TCP performance. Each measurement presented in this section is the average across several runs (100). Confidence intervals are not plotted since they are very close to the average.

5.1 Micro-benchmarks

Unless otherwise stated, all tests in this section were run on a server with an Intel Xeon CPU @3.4GHz, 16GB of RAM and a dual port Intel x540 10Gb NIC on Xen 4.4. The server is connected back-to-back to a traffic generator server. Traffic is generated there, forwarded by the box running Miniproxy and sent back to the generator server, which then measures throughput.

Throughput: For the first experiment, we measured the throughput achieved by Miniproxy and compared it to the one achieved by Varnish, a state-of-the-art TCP proxy implementation for GNU/Linux. In this test Miniproxy is configured to run with 8 MB of RAM, while Varnish requires 1 GB. Both proxies run on a VM to which we dedicate a single CPU core, and we instrument the generator to send a single TCP flow. With this setup, Miniproxy consistently outperforms Varnish in terms of throughput by about 5% (1.534 Gb/s versus 1.462 Gb/s for Varnish) while consuming significantly less memory. It is worth noting that the results were without Generic

# Concurrent Conn.	30	70	110	150	190	230
Avg. time (ms)	0,1	1	2	2	3	3

Table 1: Connection establishment time

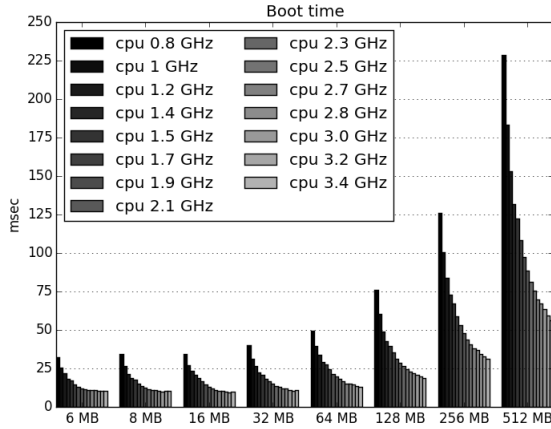


Figure 3: Miniproxy boot times for different CPU frequencies.

Segmentation Offload (GSO [25]): we have already added GSO support to `lwip` and early numbers, without using the proxy code, are encouraging (in the range of 32 Gb/s). We are now in the process of modifying our proxy application to take advantage of GSO.

Connection Establishment: We measure Miniproxy’s connection establishment time when varying the number of simultaneous connections (see table 1). Our test shows that Miniproxy is able to handle an increasing number of simultaneous connections without overly increasing the per-flow connection establishment time (in the worst case we measure a 3ms establishment time when handling 230 simultaneous connection requests).

Instantiation Time: Finally, we measured the boot time of a Miniproxy instance when varying both the amount of RAM allocated to the VM and the CPU frequency. In this test, the RAM size directly impacts the number of simultaneous flows the proxy can handle. The evaluation for different CPU frequencies, instead, is useful in order to show that Miniproxy can run on devices with different capabilities, e.g., Miniproxy could be deployed on home gateways at the edge of the network or on big servers in cloud datacenters. Figure 3 shows the results. Miniproxy requires a minimum of 6 MB to boot, in which case the instance boots in just 12ms for CPU frequencies above 2GHz. For a CPU frequency of 0.8GHz (the minimum our system allows), Miniproxy is still able to boot in just 30ms. For bigger RAM amounts, the boot times go up to 60ms on a 3GHz CPU and 230ms on a 0.8GHz CPU.

5.2 TCP Acceleration with ESF

Having shown that Miniproxy can meet the requirements mentioned in section 3, we now perform a number of tests to verify to which extent we can use it to accelerate TCP connections. In these experiments we assume a RTT of 100ms between client and server [2], a link bandwidth of 100Mb/s and that the proxies apply ESF and evenly split the end-to-end

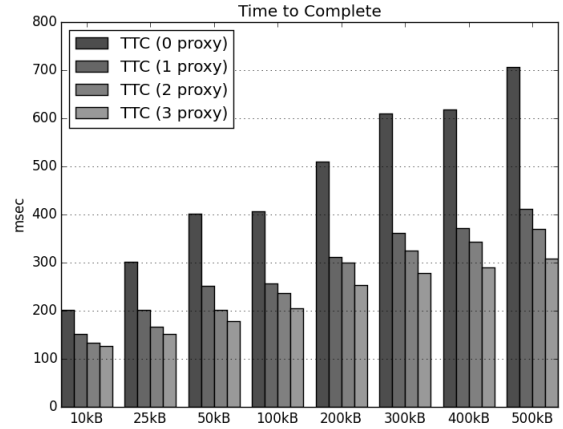


Figure 4: Transfer time when varying the number of proxies for different transfer sizes.

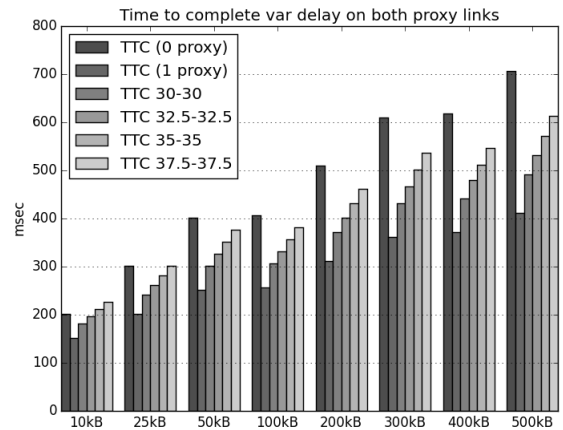


Figure 5: Transfer time with one proxy for different transfer sizes. The proxy introduces additional delay on the client-proxy and proxy-server links.

RTT. The delays are generated synthetically using the Linux Network Emulation (`netem`) tool.

Figure 4 shows the Time To Complete (TTC) for different flow sizes (from 10 KBs to 500 KBs) using a variable number of proxies. For TCP flows of 10 KBs we have $TTC = TTFB$, in fact the initial congestion window of typical today’s TCP is 10 segments (about 15KBs). Hence, for the flows transferring only 10KBs we are only measuring the effect of ESF. The ESF mechanism yields a TTFB reduction of 25%, 33.3% and 37.5% when using 1, 2 and 3 proxies, respectively. For flows that need more than 2 RTTs to complete, we can also see the effects of a faster slow start phase due to the shorter RTT. With respect to the TTC achieved without a proxy, the TTC for 25KB flows is improved by an additional 8%, 10% and 11.5% due to the increased throughput in the slow start phase. In total, we measured a reduction of 33%, 44% and 49% in TTC when using 1, 2 and 3 proxies, respectively. Of course, the longer the duration of the slow start phase, the better the measured acceleration. These results matches the analytical model discussed in [19].

An interesting case is when a proxy is located off-path, which introduces additional delay. To quantify this, we per-

form a test with an RTT of 100 ms (the one way delay is 50ms) and start from the ideal case of an on-path proxy that splits the one way delay in two segments of 25ms each. We then increase the delay on both links from 25 to 37.5ms (in steps of 2.5ms). The introduction of a proxy that splits the connection but increases the end-to-end delay can produce different results. For small increases in the delay, there is still an advantage in the split. However, when the increase is large enough, the TTC may increase to the point of making disadvantageous the use of the proxy. Figure 5 illustrates the effect of the longer delays. It shows that even an increase of 50ms in the overall RTT still produces a reduction of the TTC, when compared to the case with no proxies (for flows that require at least 2 RTTs to complete).

6. CONCLUSION

We presented Miniproxy, a lightweight TCP proxy that can boot in 12ms, while providing high forwarding performance. We showed that Miniproxy can be used to accelerate TCP's connection establishment and slow start phases. Miniproxy's very short boot times enable us to even boot the proxy instances on-the-fly, as the SYN packet of a TCP connection is first received.

Miniproxy is so far a simple proof-of-concept. Going forward, we intend to develop strategies for optimally selecting where to boot Miniproxy instances given a flow's endpoints and a set of available deployment locations. Further, we are also extending our evaluation to include wide area experiments, taking into account different traffic conditions and evaluating the suitability of available cloud datacenters (e.g., Amazon's EC2) as Miniproxy execution environments.

7. REFERENCES

- [1] Akamai. Akamai's state of the internet: Q1 2015 report.
- [2] M. Al-Fares et al. Overclocking the yahoo!: Cdn for faster web page loads. In *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference, IMC '11*, pages 569–584, New York, NY, USA, 2011. ACM.
- [3] M. Alizadeh et al. Data center tcp (dctcp). In *Proceedings of the ACM SIGCOMM 2010 Conference, SIGCOMM '10*, pages 63–74, New York, NY, USA, 2010. ACM.
- [4] A. Bakre et al. I-tcp: Indirect tcp for mobile hosts. In *Proceedings of the 15th International Conference on Distributed Computing Systems, ICDCS '95*, pages 136–, Washington, DC, USA, 1995. IEEE Computer Society.
- [5] H. Balakrishnan et al. A comparison of mechanisms for improving tcp performance over wireless links. *IEEE/ACM Trans. Netw.*, 5(6):756–769, Dec. 1997.
- [6] Erlang on Xen. Erlang on Xen. <http://erlangonxen.org/>, July 2012.
- [7] T. Flach et al. Reducing web latency: The virtue of gentle aggression. In *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM, SIGCOMM '13*, pages 159–170, New York, NY, USA, 2013. ACM.
- [8] A. Kivity et al. Osv—optimizing the operating system for virtual machines. In *2014 USENIX Annual Technical Conference (USENIX ATC 14)*, Philadelphia, PA, June 2014. USENIX Association.
- [9] S. Ladiwala et al. Transparent tcp acceleration. *Comput. Commun.*, 32(4):691–702, Mar. 2009.
- [10] F. Le et al. Experiences deploying a transparent split tcp middlebox and the implications for nfv. In *Proceedings of the 2015 ACM SIGCOMM Workshop on Hot Topics in Middleboxes and Network Function Virtualization, HotMiddlebox '15*, pages 31–36, New York, NY, USA, 2015. ACM.
- [11] Y. Liu, Y. Gu, H. Zhang, W. Gong, and D. Towsley. Application level relay for high-bandwidth data transport. *Proc. of GridNets*, 2004.
- [12] A. Madhavapeddy et al. Unikernels: library operating systems for the cloud. In *Proceedings of the eighteenth international conference on Architectural support for programming languages and operating systems, ASPLOS '13*, pages 461–472, New York, NY, USA, 2013. ACM.
- [13] A. Madhavapeddy et al. Jitsu: Just-In-Time Summoning of Unikernels. In *NSDI*, 2015.
- [14] F. Manco, J. Martins, K. Yasukata, J. Mendes, S. Kuenzer, and F. Huici. The case for the superfluid cloud. In *Proceedings of the 7th USENIX Conference on Hot Topics in Cloud Computing, HotCloud'15*, pages 7–7, Berkeley, CA, USA, 2015. USENIX Association.
- [15] J. Martins, M. Ahmed, C. Raiciu, and F. e. Huici. Enabling fast, dynamic network processing with clickos. In *Proceedings of the second ACM SIGCOMM workshop on Hot topics in software defined networking, HotSDN '13*, pages 67–72, New York, NY, USA, 2013. ACM.
- [16] A. Pathak et al. Measuring and evaluating tcp splitting for cloud services. In *Proceedings of the 11th International Conference on Passive and Active Measurement, PAM'10*, pages 41–50, Berlin, Heidelberg, 2010. Springer-Verlag.
- [17] S. Radhakrishnan et al. Tcp fast open. In *Proceedings of the Seventh Conference on Emerging Networking EXperiments and Technologies, CoNEXT '11*, pages 21:1–21:12, New York, NY, USA, 2011. ACM.
- [18] A. Singla et al. The internet at the speed of light. In *Proceedings of the 13th ACM Workshop on Hot Topics in Networks*, page 1. ACM, 2014.
- [19] G. Siracusano, R. Bifulco, S. Kuenzer, S. Salsano, N. B. Melazzi, and F. Huici. On-the-fly tcp acceleration with miniproxy (extended version). <http://arxiv.org/abs/1605.06285>.
- [20] S. Souders. Velocity and the bottom line. <http://radar.oreilly.com/2009/07/velocity-making-your-site-fast.html>.
- [21] R. Stoescu et al. In-net: In-network processing for the masses. In *Proceedings of the European conference on Computer systems, EuroSys '15*. ACM, 2015.
- [22] Z. Wang et al. An untold story of middleboxes in cellular networks. In *Proceedings of the ACM SIGCOMM 2011 Conference, SIGCOMM '11*, pages 374–385, New York, NY, USA, 2011. ACM.
- [23] Xen Project. The Next Generation Cloud: The Rise of the Unikernel. http://wiki.xenproject.org/mediawiki/images/3/34/XenProject_Unikernel_Whitepaper_2015_FINAL.pdf, April 2015.
- [24] Xen.org. Mini-OS. <http://wiki.xen.org/wiki/Mini-OS>, 2015.
- [25] H. Xu. Gso: Generic segmentation offload. <https://lwn.net/Articles/188489/>, 2006.
- [26] X. Xu et al. Investigating Transparent Web Proxies in Cellular Networks. In *PAM*. Springer, 2015.
- [27] W. Zhou et al. Asap: A low-latency transport layer. In *Proceedings of the Seventh Conference on Emerging Networking EXperiments and Technologies, CoNEXT '11*, pages 20:1–20:12, New York, NY, USA, 2011. ACM.