

How to *truly* improve the Internet’s transport layer

Michael Welzl
University of Oslo, Norway
michawe@ifi.uio.no

I. WHAT’S THE PROBLEM?

The Internet has a long and depressing history of technological failures — developments that were standardized but hardly deployed because the people who should have installed or used them had no good reason to do so. The transport layer, in particular, has hardly changed during the last 20 years. The choice of UDP or TCP does not match the heterogeneity of services that is needed by applications and could be provided by today’s infrastructure, but providing a more flexible choice among transport services is nowadays impeded by the lack of abstraction above the transport layer.

This problem is showcased by the huge number of efforts to adapt TCP to the needs of today’s applications and networks. For example, there are lots of research papers on how to use TCP for multimedia applications (which is already common in practice, e.g. with YouTube), or changing it to become more suitable for them, e.g. by not retransmitting packets. There are even more papers on adapting TCP to environment conditions where it does not work well (e.g. wireless networks). “TCP over everything, everything over TCP” was never the original Internet idea, and it is merely a result of the socket interface that applications have been provided with for decades.

II. SO, CAN THIS BE FIXED?

One obvious solution is to provide a more abstract interface above the transport layer. Underneath that interface, we would need some means to negotiate the right protocol to provide the requested service, and be able to fall back to standard TCP or UDP in case the protocol does not pass through firewalls or is unavailable at the other end, ideally without incurring a lot of overhead. These problems have been discussed before, and solving them seems to be challenging, but feasible.

For such an approach to yield a benefit (which we need for getting it used), some new and better services from other transport protocols must be available, but these protocols will not work across the Internet unless we give the right people the right reasons to enable them.

III. WHAT DO I SUGGEST?

This talk describes the incentive problem (which I call the “transport tussle”); this is about developing new technology in such a way that it is properly aligned with user incentives, and enabling appropriate usage of systems that are nowadays not getting deployed. From this seemingly simple matter of applying common sense to technological developments, major technical challenges arise. It can be deduced that, for eventually changing the transport layer, some intermediate steps must be taken. In particular, we must show that a *transparent* (i.e. unknown to the application) use of new transport protocols yields a significant benefit. We have made a first step in this direction with a simple proof-of-concept implementation of TCP-SCTP mapping using SCTP’s multi-streaming feature, where we can sometimes obtain a massive speedup (in the order of seconds). One major remaining problem is to decide how to use such a system in a way that only yields a benefit and never causes a disadvantage.

As a result of an analysis of:

- the long-term solution (an abstract interface, and what we need in order to realize it)
- the transport tussle
- our TCP-SCTP case study

...we end up with a research agenda, where it is clear which steps need to be taken in order to truly improve the Internet’s transport layer once and for all.