Run, Walk, Crawl: Toward Dynamically Varying Link Capacities

HotNets-XVI Dialogue

Brad Karp and Michael Walfish

BK: I must say that I found the vision of dynamically varying WAN link capacities exciting.

MW: Completely agree. The first thought that I had was why aren't people doing this already? You've worked in wireless networking, where there is bit-rate adaptation. What's the difference between that context and this one?

BK: Wireless links' SNRs vary on short timescales, so the problem is in a sense more pressing in that scenario: if you just pick a modulation that "always works," (that is, in times of poorest SNR) you'll fall far short of the achievable capacity, because the links are better much of the time. On a wired link, SNRs are much more stable over longer time periods—in fact, before I read this paper, I never really thought about SNR variability on wired links, and naively presumed (absent any actual consideration) their SNRs were quasi-perfectly-stable!

MW: And what we learned from this paper is that in fact SNRs on optical links aren't stable, that they change on timescales that—

BK: Wellllll, we know that at *long* timescales there is variability, as the authors' dataset is a time series taken at 15-minute intervals, and they see varying SNRs at that granularity, but what we *don't* yet know, and is an important question to resolve, is what the variability in SNR is at timescales *shorter* than 15 minutes. That is, *if* SNR varies, say, within a 0.5-second interval at times, it may be challenging to adapt modulation on a fiber-optic link that quickly—

MW: Because, as the authors point out, changing modulation takes time, and that time cost has to be paid somewhere?

BK: Exactly. The hardware (laser, transceiver) incurs some delay for reconfiguration, and one presumes that the link is down (or more euphemistically, "idle") during this period. So each change in modulation can potentially offer a performance (higher-bit-rate, high-SNR-requiring modulation) or robustness (lower-bit-rate, low-SNR-tolerant modulation) win, but there's a cost in link idle time for each change. So two big questions remain: (1) What's the "lower bound" (not in a theoretical sense, but in practice) on link reconfiguration time, for example, if we built hardware specifically with the goal of fast reconfiguration in mind? And (2) What's the variability of SNR at fine-grained timescales, and how does it compare with economically achievable hardware for fast link reconfiguration?

MW: Okay, understood. Sounds like (1) and (2) are super-interesting follow-up questions! Teasing out questions of this kind for the community is what HotNets is all about... But for now, let's look into the measurement results some more.

BK: Yeah—thus far we've mostly considered the opportunity to run links "faster," but the authors also claim that there's potential to make links more *robust.* And they present measurements they argue support that claim. Section 2.2 gives statistics on root-cause analysis of link failures. The argument appears to be that 90% of link failures are amenable to *prevention* by dynamic adaptation of link modulation. A tantalizing claim!

MW: Yes, it is tantalizing. But hang on... the 90% is the complement of the obvious outage events: 10%

of events are power issues or fiber cuts, from which there is no recovery. But a very large percentage of the link events that the authors analyze (see Figures 4a,b) have "undocumented causes." It's not at all clear that these unspecified cases in fact correspond to periods in which the link is operational but at degraded SNR; it might be that many of these cases, the link is unusable.

BK: I must agree. There could be an intermittently failing transceiver, for example, or perhaps even a bump of a connector on a switch or ROADM, among who knows how many other non-fiber-cut events that might be consistent with *no* modulation achieving a tolerable bit-error rate.

MW: Wait, but the authors also use their SNR measurement dataset to argue that modulations for lower SNRs can mask what are today link failures, don't they?

BK: They do make this argument, yes. Figure 4c examines the distribution of minimum SNRs correlated with link failures, and the authors note that the points on this distribution admit driving a link at 50 Gbps...which beats a failed link any day. But wait a minute: we know from earlier in the paper that the SNR samples were taken only at a 15-minute granularity, no? So I think this conclusion has the same uncertainty around it as that made earlier about the potential upside from driving links at higher-bit-rate modulations. Here, again, the concern is that on either side of an SNR sample there are quite possibly (and even likely—why would one expect samples on 15-minute boundaries to happen to be the minima per outage?) SNR values lower than the sampled one...in which case 50 Gbps may not have sufficed. Again, finer-granularity measurements in time will let the community figure out whether this claim holds in practice.

MW: We still haven't answered my initial question. Why aren't people doing this already?

BK: While I don't do optical networking for a living, an optical networking colleague tells me that until around 2015, optical transceivers generally supported only a single modulation, then began supporting around 3 modulations whose target SNR regimes were spaced around 3.5 dB apart, and only just recently have transceivers appeared with more modulations targeting more closely spaced SNR regimes. In short, improvements to hardware create the opportunity to exploit what was "excess" headroom on link margins—*if* the adaptation can be done right.

MW: Very interesting! The authors are precisely on top of this trend, then. And their measurements certainly make the case that dynamically varying modulation is worthy of further study and, ideally, reduction to practice.

BK: Yes, I really do hope to see WANs realize material capacity and reliability gains this way.

MW: Me too. Then this paper will have been prophetic at a minimum and possibly a watershed.

BK: Yeah, then we could tell our grandchildren that back when cars had drivers and servers didn't fit in your pocket, WAN optical links used to work at a fixed modulation.