Cross-Layer Simulations of Fast Packet Protection Mechanisms

F. Fidler[(1,2)], P.J. Winzer[(2)], C.P. Lai[(1)], M. K. Thottan[(2)], K. Bergman[(1)]

[(1)] Dept. of Electrical Engineering, Columbia University New York, NY, USA, *ffidler@ee.columbia.edu*
[(2)] Optical Networks Research, Bell Labs, Holmdel, NJ, USA, *peter.winzer@ieee.org*

Abstract *Time-dependent inter-packet and intra-packet BER variations are implemented in the ns-2 network simulator to evaluate fast packet protection based on physical-layer monitoring through pre-FEC BER read-out.*

Introduction

Studying optical networks in a holistic, cross-layer optimized approach and assessing different quality-of-service aware protocols with emerging cross-layer network control schemes[1] requires packet network simulators that incorporate both dynamic physical-layer performance variations as well as data traffic dynamism[2]. Today's discrete-event based packet network simulation tools, such as the widely accepted open-source software[3] ns-2, support many data network architectures and protocols, but generally lack the capability of incorporating realistic physical-layer models for the wide variety of system-specific impairments. Dropped packets are typically simulated using a pre-specified, temporally constant (or slow varying) packet loss ratio, which can be insufficient in the case of fast channel bit error ratio (BER) variations. Such BER variations are generally not considered a problem in today's optical networks due to the large physical-layer margins allocated for ultimate reliability. However, they can be problematic in future networks, which may be engineered with flexibly lower margins and recover from infrequent impairments through dynamic protection mechanisms[1].

In order to interface the behavior of a dynamic physical layer with packet network simulations, we take a parameterized approach that accepts *physical-layer BER variations* as input to a higher-layer packet network simulation. This allows us to incorporate both *simulations* (deterministic or stochastic) and *measurements* of physical-layer performance variations in a unified way, taking into account quasi-static impairments (e.g., loss or chromatic dispersion), moderately fast impairments (e.g., polarization mode dispersion[4] (PMD)), and highly dynamic impairments (e.g., power transients[5], or nonlinear crosstalk between wavelength division multiplexed (WDM) channels[6]).

We incorporate general physical-layer BER variations into the packet network simulator ns-2, enabling cross-layer simulations beyond existing modules specific to the wireless channel[7]; the source code of our software is available on-line[8]. Using our modules, we study packet loss rates using *fast-reroute* and *proactive-protection*[1] cross-layer networking.

BER variations across and within data packets

Figure 1b shows an example of temporal BER variations of a WDM channel, which may be obtained by physical-layer simulations or measurements using a forward error correction (FEC) decoder. Such BER time series form the basis of our packet-level simulations and are assigned to data packets on the corresponding links in the packet simulation. Depending on *(i)* the dynamics of the BER variations, *(ii)* the length of a data packet, and *(iii)* the underlying architecture (packet-switched or circuit-switched core, cf. Fig. 1a), the BER may either be constant over the duration of any single packet or BER variations may result in error bursts *within* a packet.

Packet-by-packet BER variations: If packets are directly transported over the optical infrastructure (IP-over-WDM), most physical-layer impairments will be slow compared to the duration of a packet; e.g., a packet with a maximum transfer unit (MTU) of 1500 bytes on a 10-Gb/s optical link is ~1.2 μs long, shorter than most dynamic optical impairments; it then suffices to assign a single BER value to each packet.

Intra-packet BER variations: Intra-packet BER variations may not only occur for exceedingly fast BER variations but may even be found for relatively slow fluctuations if packets are transported on a time-division multiplexed (TDM) circuit infrastructure (IP-over-SONET/SDH/OTN). Here, the content of a single packet can be spread across many transport frames (cf. Fig. 1b), resulting in a *time-compression effect* that lets relatively slow physical BER variations

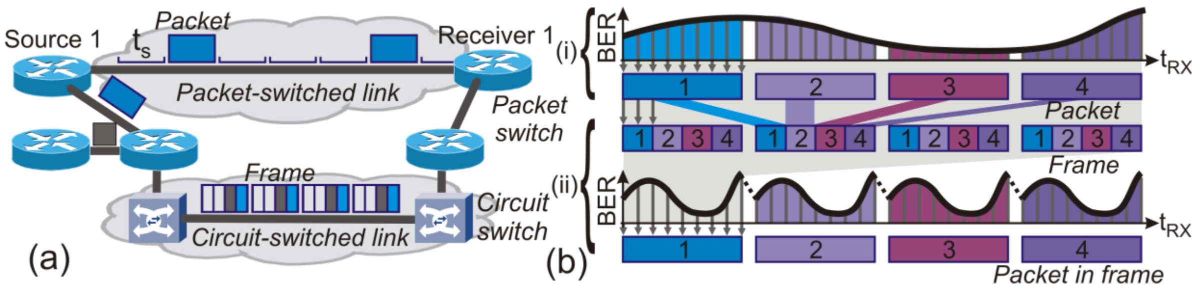

Fig. 1: (a) Packet-switched and circuit-switched networks; (b) mapping of BER variations onto bit errors within a packet (*i*), and time compression of BER variations when splitting a packet across multiple TDM frames (*ii*).

978-3-8007-3173-2 © VDE VERLAG GMBH

appear as fast effective variations over the duration of a packet. For example, an Ethernet jumbo-frame (9000 bytes MTU) from a 1-Gb/s client occupies a time slot t_S = 720 ns on a 100GbE packet interface (IP-over-WDM). Using IP-over-OTN, this packet occupies $8 \cdot 9000/10^9$ = 72 µs, spread (together with other TDM tributaries) over 62 OTU4 frames (at 112 Gb/s), and intra-packet BER variations can, e.g., be caused by optical power transients[5]. The situation is exacerbated for larger packets (e.g., IPv6 jumbograms; MTU < $4 \cdot 10^9$ bytes), for virtual concatenation (VCAT), where a packet can spread over up to 32 ms, or for smaller granularities than OTN (as in SONET/SDH).

Comparison of fast protection mechanisms

Leaving the details of our ns-2 implementation to the documented source-code[8], we now present an application of our simulation package by comparing fast packet protection mechanisms over either a packet-switched or a circuit-switched core, using GbE clients with 1500 bytes MTU and a 100-Gb/s line infrastructure. As soon as the received BER exceeds BER_E, the correction threshold of the underlying FEC ($2 \cdot 10^{-3}$ in our example), *fast-reroute (FRR)* switches the data stream to a protection path. *Proactive protection* (PPT) starts earlier (at a pre-defined threshold $BER_T < BER_E$, here at 10^{-4}), with the goal of near hitless protection for sufficiently slow impairment dynamics[1].

In Fig. 2a, we assume a step-like increase of the BER (inset (*i*)) and simulate the number of lost packets n in the network of Fig. 1 as a function of the slope of this BER step. Proactive protection (green area) offers zero packet loss until the BER step becomes so steep that the time span t_P between BER_E and BER_T is shorter than the round-trip time (RTT), the minimum time required for the protection mechanism to kick in. For fast transitions, the number of lost packets using PPT converges to $n = \text{RTT}/t_S$, as is found for FRR (red area).

In Fig. 2b, we study sinusoidal log-BER variations with period Δt (inset (*ii*)). PPT shows no packet loss for $t_P > \text{RTT}$. Assuming packet durations of $t_S << \text{RTT}$, the two protection mechanisms perform identically for $t_P + t_E = \text{RTT}$, with t_E being the time duration where packet loss occurs (red area in inset (*ii*)). Beyond that point ($t_P + t_E \leq \text{RTT}$), the RTT is too large compared to the impairment dynamics and PPT offers no advantage over FRR. For both schemes, the number of lost packets decreases with increasing BER dynamics, converging to $n = t_E \cdot \text{RTT}/(\Delta t \cdot t_S)$. (The small ripple in the packet loss is due to the interplay of RTT, t_E, and Δt.) The solid blue curve in Fig. 2b pertains to a circuit-switched core. Independent of the protection mechanism, we find the same behaviour as for the packet-switched core up to the point where the effective BER variations seen by the stretched TDM packets become comparable to the stretched packet duration, $8 \cdot 1500/10^9$ = 12 µs in our case. This leads to

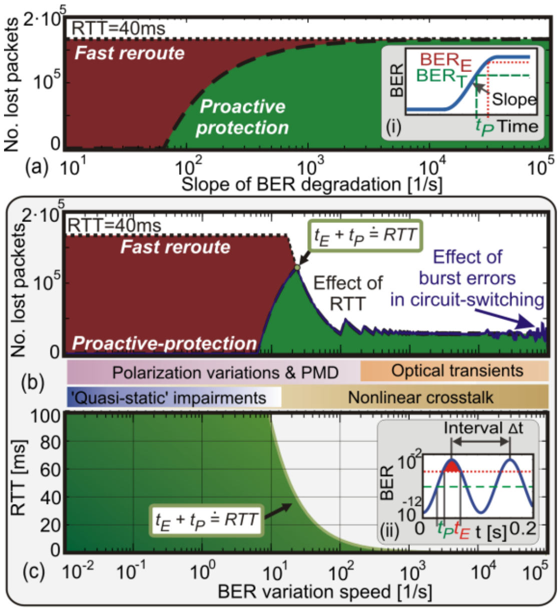

Fig. 2: Number of lost packets (dotted: *FRR*, dashed: *PPT*) vs. (a) slope of BER step (*i*), and (b) BER variation speed $1/\Delta t$ (*ii*); (c) Region where PPT outperforms FRR.

additional variations in packet loss, caused by uncorrectable intra-packet burst errors[9].

Finally, Fig. 2c compares PPT and FRR for various combinations of RTT and BER variation speeds $1/\Delta t$. The green area under the curve $t_P + t_E = \text{RTT}$ denotes the region where PPT outperforms FRR. As expected[1], fast impairment dynamics require short RTTs for PPT to provide an advantage over FRR. In typical optical transport networks (4 ms < RTT < 40 ms), PPT is effective against quasi-static impairments and PMD, but is likely to fail for dynamic impairments such as fast amplifier power transients.

Conclusion

We implemented a simulation tool for the open-source packet network simulator ns-2 that enables cross-layer network simulations based on time-dependent BER series obtained from physical-layer simulations or measurements. Using these modules, we assessed fast protection techniques over circuit and packet transport infrastructures. Comparing proactive-protection to fast-reroute, we studied the impact of network size and physical impairment time scales on the packet loss rate.

We acknowledge valuable discussions with A. Kalmar and S. Trowbridge.

References

1 O. Gerstel et al., Proc. OFC'08, NWD4 (2008).
2 I. Tomkos et al., Proc. ECOC'08, We.3.D.1 (2008).
3 NS-2, [online], (2009), http://isi.edu/nsnam/ns/
4 H. Bulow et al., Proc. OFC'99, 83 (1999).
5 Y. Sun et al., El. Letters **33**, 313 (1997).
6 S. Chandrasekhar et al., PTL **19**, 1801 (2007).
7 N. Baldo et al., Proc. ACM NSTools'07 **321**, (2007).
8 F. Fidler, [*online*], (2009), http://groups.geni.net
9 T. Mizuochi, Proc. OFC'08, OTuE5 (2008).