FAst In-Network GraY Failure Detection for ISPs – *Public Review*

Henning Schulzrinne
Columbia University
New York, NY
hgs@cs.columbia.edu

Internet service providers are used to seeing packet losses and link failures, and packet networks have had mechanisms to deal with congestion-induced packet loss and link failures since their earliest days, with plenty of diagnostic tools to observe and manage them. But network switches and routers can also suffer from more intermittent failures, typically caused by hardware and software bugs, or links with marginal link budgets. Such failures are hard to detect and diagnose since they may only affect a small fraction of the traffic – and the traffic affected may be idiosyncratic, e.g., only dropping packets of a certain length or destination network. Thus, these impairments are often Heisenbugs, problems that shape-shift under observation. The paper calls these packet losses "gray" failures.

This is not just a could-occur-somewhere problem — the authors commendably surveyed 46 ISPs about their experiences with such gray failures and found that this indeed an annoying and time-consuming part of their operational reality. The authors also found numerous examples of vendor bug reports that describe such gray failures. The paper investigates existing tools for detecting and diagnosing packet loss, but the authors find them to be unsuitable for an ISP setting.

The paper provides a first systematic attempt to help ISPs to detect these types of gray failures, using a combination of counters reported by downstream switches that tally up packets for each forwarding table entry or packet header match, and a hash-based tree for aggregate counting.

The paper tests the FANcY algorithms both by ns-3 simulations and in a Tofino switch implementations, using CAIDA traces and random packet loss as metrics. The evaluation finds that gray failures for all but very rare losses on small flows can be detected reliably and without false positives, even as the paper notes the difficulty of representing "modern" ISP traffic.

The reviewers appreciated the practical motivations of the proposed system, consulting with operators, and the careful evaluation by simulation and implementation. Debugging networks is hard, so the reviewers were concerned with the ability of operators to identify flow entries and how well the system can be used to identify the precise cause of the gray failure, given that the subset of the header space covered by an entry is, at least initially, unlikely to correspond to the bug-induced loss pattern. The paper proposes re-routing traffic, but many of the hardware or software bugs are likely to affect all switch ports, making that an unattractive remedy.

As ISP network reliability becomes a prime design objective, the paper offers an additional approach that might motivate others to better understand the less-studied causes of network failures and packet loss, particularly those short of "backhoe fade" and complete switch failures.