Impact of FIB Aggregation on Traffic Offloading

Nadi Sarrar[1*], Stefan Schmid[1], Steve Uhlig[2], Anja Feldmann[1]
{*nadi,stefan,anja*}*@net.t-labs.tu-berlin.de* *steve@eecs.qmul.ac.uk*
[1]*TU Berlin & Telekom Innovation Labs* [2]*Queen Mary University of London* [*]*Student author*

This work seeks to combine two active and so far independent research threads to handle large *Forwarding Information Bases* (FIBs) in Internet routers: *Aggregation* and *traffic offloading*. More specifically, we analyze the characteristics of FIB aggregation and its impact on the data plane in the context of a traffic offloading system which aims at offloading most of the traffic by focussing on the few most heavily used entries in the FIB of an Internet router [2]. We want to find empirical answers, based on ISP as well as IXP traffic traces, to the following four key questions.

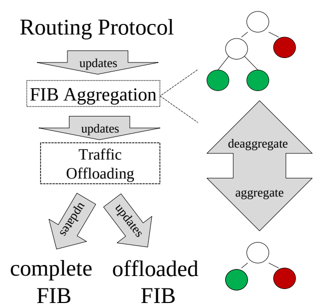

1. Does the number of next-hop routers present in a FIB have an impact on its aggregatability? Two FIB entries – IP prefixes – can be merged without violating routing consistency, if (1) there exists a prefix that exactly covers their combined address space, and (2) the two original FIB entries share the same next-hop. These rules can be relaxed, see [1]. Intuitively, the larger the number of different next-hops present in a FIB, the worse the aggregation ratio. We perform experiments on real routing tables taken from ISP access and backbone routers (tens of next-hops) and routers at an IXP (hundreds of next-hops). Our preliminary results show that, with a provably optimal routing table aggregation algorithm, we can shrink FIBs to at least 60% of their original size, for all three router locations.

2. What is the right trade-off between the compression gain and frequency of changes to the aggregated FIB? In contrast to the offline algorithm used for static compression only, real-world systems such as our traffic offloading system require an *online aggregation algorithm* that also takes into account the number of changes to the aggregated table on routing table updates, as those updates typically impact the data-plane. To achieve this, both existing online aggregation algorithms (SMALTA [3]) as well as our own online algorithm seek to strike the right balance between optimal aggregation and frequency of updates. We will perform simulations with real Internet traffic data to better understand this trade-off and use the insight to identify design criteria for an optimized algorithm.

3. How does the effectiveness of the offloading system change under aggregation? In the traffic offloading scenario, the offloading ratio can potentially be improved by aggregation. Each single offloaded routing table entry will comprise multiples of the original entries and thus have a bigger share in traffic volume. We will perform simulations based on real Internet traffic traces to answer questions like (1) how compressible are the offloaded, top-volume FIB entries, and (2) how much more traffic can be offloaded with the help of prior aggregation? Preliminary results based on IXP data suggest that the fraction of offloaded traffic can be increased by more than 10%, when offloading 2,000 FIB entries[1].

4. How does the number of updates to the offloaded FIB entries change under aggregation? Depending on the specific behavior of any aggregation algorithm, the offloaded FIB entries will likely be aggregated in various different ways. This in turn leads to different numbers of modifications to the offloaded entries over time (*churn*), depending on the algorithm used. We are the first to study the churn of aggregated routing tables specifically for those entries which carry most of the traffic, i.e., the ones that the traffic offloading system uses to offload traffic.

References

[1] DRAVES, R., KING, C., VENKATACHARY, S., AND ZILL, B. Constructing Optimal IP Routing Tables. In *INFOCOM* (1999).

[2] SARRAR, N., UHLIG, S., FELDMANN, A., SHERWOOD, R., AND HUANG, X. Leveraging Zipf's Law for Traffic Offloading. *ACM SIGCOMM CCR* (2012).

[3] UZMI, Z. A., NEBEL, M., TARIQ, A., JAWAD, S., CHEN, R., SHAIKH, A., WANG, J., AND FRANCIS, P. SMALTA: Practical and Near-Optimal FIB Aggregation. In *ACM CoNEXT* (2011).

[1]Based on a snapshot of the IXP's route server's FIB and a 1-day trace of sFlow samples, for bin-optimal and TFO strategies [2].