

CHAPTER 5

Further Development

Throughout this study, we have identified some issues that act as hurdles to an optimal implementation of our TCP Friendly marker. In particular, we notice the fact that the process of token allocation is based on a smoothed estimate of the flow traffic in the past. More specifically, our algorithm relies on the amount of traffic that a flow has introduced to the marker during the immediately previous interval of time between updates (line 12 in Algorithm 1). If such estimate is not accurate for some particular flows, then the token allocation process will suffer from this circumstance. We encountered an optimal interval time³⁰, which we further used in our experimentation. Nonetheless, we cannot provide any guarantees as to whether that interval is also optimal for all edge domain routers in the Internet, and we further suspect that such an interval will depend on the traffic through the marking system on the router. Therefore, we support an evolution to the marker that calculates the amount of error in the estimate (from one interval to the next) and modifies the interval time appropriately in order to minimize such error³¹. We strongly suspect that such adaptive behavior will make the inter-updating time (i.e. interval time) converge rapidly and will guarantee optimal benefits from the usage of the TCP Friendly marker. Nonetheless, we should not discourage its deployment with static interval time (as currently available), since we have seen that the **performance improvements are overly beneficial to be disregarded**.

³⁰ 400ms, for our experiments

³¹ Given some threshold for the error

The original work on the TCP Friendly marker included some simulation analysis regarding the application of the marker on supporting assured services. Their simulations showed that over-provisioning of tokens was in fact effective in offering selected flows the characteristics of assured services. The existence of our implementation could now be used to study the experimental performance of our marker in the arena of assured services, and more specifically could serve to confirm the idea of over-provisioning.

Finally, the structuring of our marking process results into differentiation among only two types of priorities and thereby the use of FRIO. We suggest a further extension to the marker from the sole two classes (IN and OUT), into a larger number of them (and up to 16). We have shown that the identification of key problems of TCP and the targeting of such weaknesses, through prioritizing critical packets, has in fact a great effect on the overall performance of the system. Further segregation of traffic, according to the importance of its arrival for the dynamics of the system, could potentially lead to a much greater improvement. We encourage this particular evolution to our work, especially because the existent implementations of GRED (by Jamal Hadi) and Flow GRED (by us) shorten the road for such potential improvements. Perhaps simulations could confirm our positive views on the further segregation of traffic, and aid in the identification of key issues on which to base the segregation upon.

CHAPTER 6

Conclusions

The experimental results that we have presented throughout this document have proven the validity of previously studied traffic conditioning approaches. We note their feasibility as definite and powerful solutions to the Internet's traffic congestion problem. We have identified the weaknesses of the underlying transport protocol, and presented a running implementation of a system that, using differentiated services concepts, is able to statistically guarantee a significant improvement in performance. The combination of the TCP Friendly marker at the access edges of the network, along with simple packet differentiation in its congested core, is capable of increasing the utilization of the network, reduce the number of timeouts, increase the net throughput as seen by users, and inject more predictability to the servicing of overall traffic.

We have sketched a fundamental architecture for the deployment of our system. The use of the TCP Friendly marker could be used in conjunction with more complex flow classification at the core of the network, for optimal results. Nonetheless, the use of simple RIO will in fact take advantage of all the benefits that are consequence of the use of our marker as edge-domain traffic conditioner tool.

The results of our implementation suggest the following benefits

- **Reductions of two orders of magnitude on the number of timeout instances.**
- **Consistently larger net data throughput for all TCP flows.**

- **Halving of the probability of packet loss across all instances of implementation.**
- **Increments of 200% on the predictability of service by the network.**
- **Improvements immunity to the inappropriate configuration of RED-like parameters.**
- **Improvements independence from version of TCP implementation.**
- **Scalability of the results to any number of TCP flows through a gateway.**
- **Improvements independence from flow classification at the core of the network.**

The actual implementations of the marker and the buffer management schemes have been made available online³². We provide a tool for experimentation within the open source community and enable the immediate deployment of our system, given the portability of the implementation towards proprietary operating systems that support the existent corporate network infrastructure.

³² <http://networks.ecse.rpi.edu/~lomong/tcpfriendly>. Further installation instructions are included in this website