The Incremental Deployability of Core-Stateless Fair Queuing

Yannick Blanpain, Hung-Yun Hsieh, and Raghupathy Sivakumar

School of Electrical and Computer Engineering
Georgia Institute of Technology, Atlanta, GA 30332, USA
gte463w@prism.gatech.edu,{hyhsieh,siva}@ece.gatech.edu
http://www.ece.gatech.edu/research/GNAN

Abstract. In this paper, we study the *incremental deployability* of the Core-Stateless Fair Queuing (CSFQ) approach to provide fair rate allocations in backbone networks. We define incremental deployability as the ability of the approach to gracefully provide increasingly better quality of service with each additional QoS-aware router deployed in the network. We use the *ns2* network simulator for the simulations. We conclude that CSFQ does not exhibit good incremental deployability.

1 Introduction

The growing diversity of Internet applications has motivated the need for supporting service differentiation inside the network. A real time multimedia application streaming live video, and a file transfer application, require very different network-level quality of service. While today's best-effort Internet model will not differentiate between the two applications, the Internet is slowly moving towards a pay-per-use model wherein applications can "buy" the specific network-level service they require. Parameters of such a service can include bandwidth, delay, jitter, loss, etc.

In recent years, the Internet Engineering Task Force (IETF) has developed two different quality of service (QoS) models for the Internet. While the Integrated Services (*intserv*) model provides fine-grained per-flow quality of service, it is not considered as a solution for backbone networks due to its inability to scale to a large number of flows [4]. The Differentiated Services (*diffserv*) model, on the other hand, is scalable at the expense of providing only a coarse-level quality of service that does not support any assurances to individual flows.

Of late, newer QoS models that attempt to bridge the scalability of the diffserv model, and the service richness of the intserv model have been proposed [2–5]. We refer to these new models as *Core-Stateless QoS models* in keeping with the terminology introduced by the authors of [2]. While the core-stateless approaches offer great promise in terms of both scalability and the service models they can support, not much work has been done in terms of evaluating the feasibility of their practical deployment. Specifically, given the enormous size of the present Internet, any solution requiring replacement of routers in the Internet has to be *incrementally deployable*. In the context of quality of service architectures, we define incrementally deployability as *the ability* *to provide increasingly better quality of service with increasing number of QoS-aware routers.*

In this paper, we study the incremental deployability of core stateless fair queuing (CSFQ) [2], a QoS model that attempts to provide rate fairness without maintaining per-flow state at core routers. The contribution of our work is twofold: *(i) We study the performance of CSFQ through simulations, and conclude that it has poor incremental deployability. (ii) Based on the insights gained through our study of CSFQ, we present some "guidelines" for the deployment of an incrementally deployable core-stateless QoS model.*

The rest of the paper is organized as follows: In Section 2, we provide a brief overview of the CSFQ mechanism. In Section 3, we describe the simulation model, and in Section 4 we discuss the simulation results. In Section 5 we conclude the paper.

2 Background

2.1 Core Stateless Fair Queuing

CSFQ attempts to emulate the behavior of fair queuing [1] at core routers without maintaining any per-flow state. Combined with an end-to-end adaptation scheme (like that of TCP), it approximately achieves max-min fairness [4]. We provide a quick overview of the CSFQ mechanism in the rest of the section. CSFQ estimates the *fair share* of each link without maintaining any per-flow state in the core router. The fair share α at a core router represents the share of the output link capacity that is allotted to each flow that traverses the router. In CSFQ, each packet has the rate r - of the flow to which the packet belongs - stamped in its header by the *ingress* edge router. When the packet arrives at a core router, the router drops the packet with a probability of $max\{0, 1 - \alpha/r\}$. If the packet is not dropped, it is accepted for transmission.

If A represents the aggregate arrival rate, F represents the aggregate accepted rate (where the two variables are updated after the arrival of every packet), and C represents the link capacity, the fair share α is updated as follows:

if $(A > C)$ $\quad \alpha_{new} \leftarrow \alpha_{old} * C/F$
else $\quad \alpha_{new} \leftarrow$largest rate of any active flow

The combination of fair share estimation and probabilistic dropping of packets for those flows whose rate exceeds the fair share enables CSFQ to enforce fair sharing of a link without maintaining any per-flow state in the router.

2.2 Incremental Deployability

In Section 4 we evaluate the incremental deployability of the CSFQ mechanism by studying the fairness properties of a network in which only a fraction of the nodes are CSFQ routers. Specifically, we investigate incremental deployability from two perspectives: the core of the network and the edges. In other words, we study both the incremental deployment of CSFQ routers in the core of a backbone network (assuming that all the edge routers are CSFQ-aware) and the incremental deployment of CSFQ-aware routers at the edges of the network. For both cases, we use First In First Out

(FIFO) routers for the non-QoS aware routers. Note that incremental deployment of QoS-aware routers at the edges can also be seen as increasingly more number of flows in the network being QoS-aware flows. In the rest of the paper, we refer to QoS-aware routers as *fair routers* and the non-QoS-aware routers as *legacy routers*. Likewise, we refer to QoS-aware flows as *fair flows*, and the default flows as *legacy flows*.

3 Simulation Model

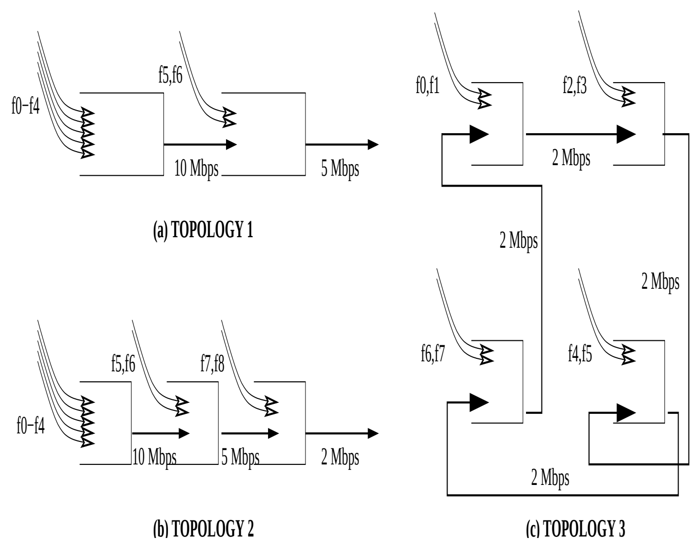

Fig. 1. Small Topologies Used For Simulations

We use the *ns2* network simulator [7] for our simulations. The *ns2* extensions for CSFQ were downloaded from http://www.cs.cmu.edu/˜ hzhang/csfq/. Although several topologies were used for the study, in this paper we illustrate our arguments using results for the 3 simple topologies shown in Figure 1 and the large topology shown in Figure 6. The four topologies are described in the next section. For each of the topologies and for the two scenarios of core and edge deployment, we start from a scenario where all routers use the FIFO scheme. The subsequent scenarios are obtained by incrementally changing one router at a time to use the CSFQ mechanism or be CSFQ-aware, till all the routers in the scenario are changed. For each of the scenarios, we measure the fairness offered by the network as a whole. *We use Jain's fairness-index to demonstrate the fairness achieved among the flows in the network in terms of their end-to-end throughput.* We plot the fairness index against the configuration of the network defined by the number of CSFQ (or CSFQ-aware) routers.

The fairness-index plotted is an average over multiple simulations using the same number of QoS-aware routers, but with different placements for the QoS-aware routers.

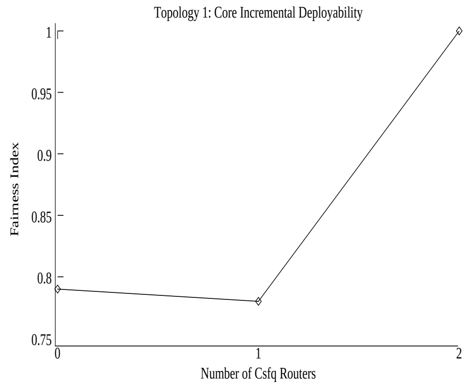

Fig. 2. Topology 1: Core Incremental Deployability

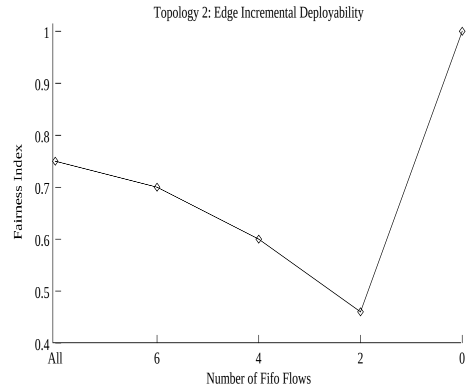

Fig. 3. Topology 1: Edge Incremental Deployability

The flows use UDP and the traffic is generated using Poisson distribution. The labels on the x-axis of the graphs represent the number of QoS-aware routers used. The fairness-index is plotted on the y-axis.

4 Simulation Results

4.1 Simple Topologies

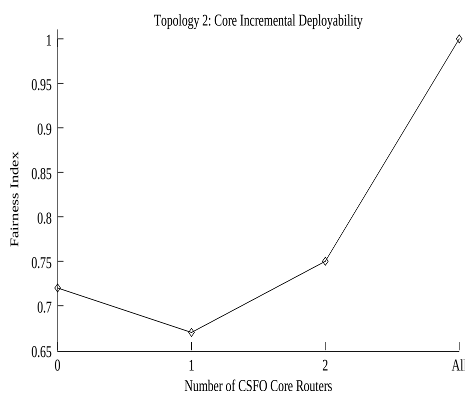

Fig. 4. Topology 2: Core Incremental Deployability

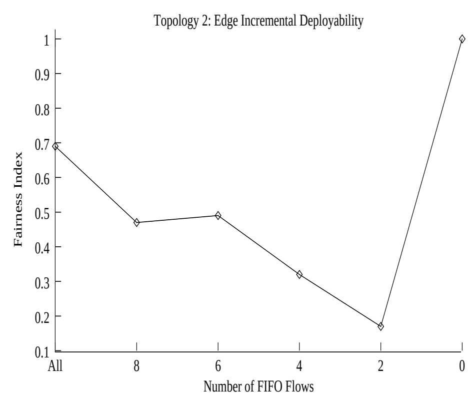

Fig. 5. Topology 2: Edge Incremental Deployability

Figures 2, 4, and 7 show the fairness demonstrated by CSFQ when CSFQ routers are incrementally deployed in the core for Topologies 1, 2, and 3 respectively. Figures 3, 5, and 8 show the fairness when CSFQ-aware routers are incrementally deployed at the edges. As can be seen from Figures 2, 4, and 7, CSFQ does not exhibit good incremental deployment when being deployed in the core of the network. Note that each datapoint in the graphs shown was averaged over multiple simulations for all possible configurations with that many number of CSFQ routers. For example, in Figure 2, the second datapoint represents the average of the fairness indices for the two possible configurations (CSFQ-FIFO and FIFO-CSFQ) of Topology 1 with one CSFQ router. Also, for each possible configuration, simulations were run with varying rates for the flows, and an average taken over all the simulation runs.

Similarly, from Figures 3, 5, and 8, it can be observed that CSFQ does poorly in terms of edge router incremental deployment. As mentioned earlier, with fewer number of CSFQ-aware routers at the edges, the number of CSFQ-aware flows in the network decreases. Hence, the unfairness stems from CSFQ flows sharing links with best-effort flows (legacy flows) in the core of the network. An interesting aspect of the results shown in the figures is the "dip" in the curve for all the scenarios. Specifically, the fairness index goes down with increasing number of fair-routers at the edges and rises only when all the edge routers are fair-routers. This anomaly can be explained as follows: When legacy-flows share a congested CSFQ link with fair-flows, CSFQ will be significantly unfair toward the fair-flows: When the fair share estimate is initially reduced, CSFQ will not observe any reduction either in the arrival rate (because legacy-flows cannot be assumed to adapt to fair share estimates). This will result in a further reduction in the fair share estimate. However, since the rate of the legacy-flows by itself is more than the link capacity, the presence of sustained congestion will ensure that the fair share estimate is cut down to zero, at which stage none of the packets belonging to the fair-flows will be accepted. However, if all the flows on a CSFQ link happen to be legacy-flows, "better" fairness can be expected as CSFQ will not interfere on behalf of any particular flow, leaving it for the indiscriminate (but not unfair) allocation of the drop-tail mechanism to decide the per-flow shares. Hence, it can be observed that the fairness achieved when there are no fair-flows in the network is better than the intermediate cases.

4.2 Large Topology

For the results shown in Figures 9 and 10, the topology shown in Figure 6 was used. All link capacities were set to 10Mbps. Each core-router in the topology was connected to exactly one edge-router. Simulations were performed with 1100 flows. Each edge-router serves as an ingress for 100 flows. The 100 flows at an ingress-router is divided into 10 sets of equal number of flows, each set having a unique egress-router among the other 10 edge routers. The rate of each flow was set to 200 kbps. Although, simulations were performed using other flow rates, we do not present them here for lack of space. The results demonstrate that even for more realistic topologies, CSFQ performs poorly in terms of incremental deployability, both at the core and at the edges.

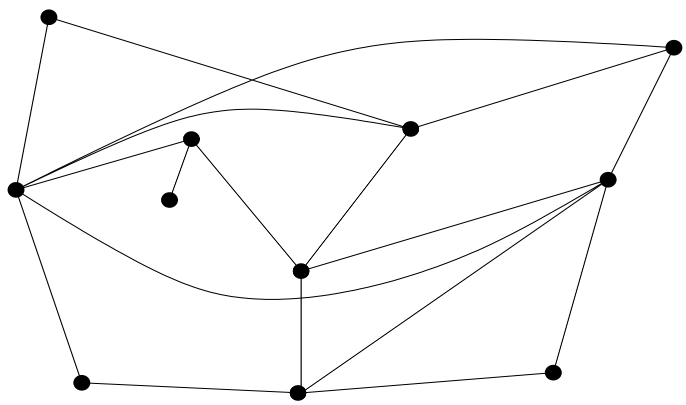

Fig. 6. Topology 4 - Sprint Backbone

4.3 Discussion and Insights

Our simulation results demonstrate that CSFQ is not an incrementally deployable approach to achieve fair rate allocation. However, the following insights can be gathered from conducted study:

Impact of Legacy Core Routers The impact of the presence of legacy-routers on the performance of fair-routers depends upon the specific allocation scheme employed at the fair-routers. In core-stateless approaches, fair allocation of link bandwidth at core routers is achieved by making use of *dynamic state information* about the flows. Edge routers pass the dynamic state to core routers through varying schemes including fields in the packet headers (dynamic packet state) [2], or specialized packets in the flow (dynamic flow state) [4]. The core routers rely solely on the dynamic state carried by the flow, and hence do not perform any per-flow processing. However, in the presence of legacy-routers, such approaches encounter the following problems that adversely affect the fair service they offer:

– Legacy-routers drop packets indiscriminately [1]. Hence, flows traversing such routers will inherently receive unfair service. While this by itself cannot be completely overcome, it leads to unfair service even at fair routers. We elaborate upon this phenomenon next.
– While legacy-routers drop packets indiscriminately, being unaware of the fair allocation scheme, such routers will, in addition, fail to update the dynamic state of flows when dropping packets. Hence, the dynamic state at downstream fair routers can be inconsistent with the actual flow state. When fair routers use such inconsistent dynamic state to perform rate allocation, the allocation will be unfair.

[1] Recall that we have assumed the use of the drop-tail mechanism at legacy-routers.

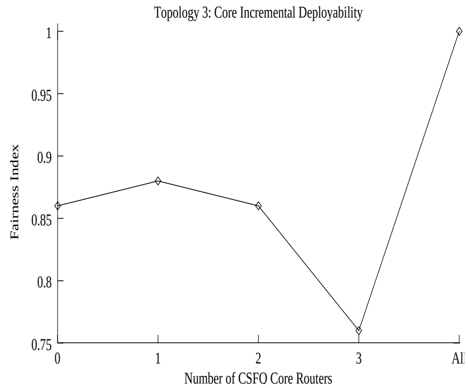

Fig. 7. Topology 3: Core Incremental Deployability

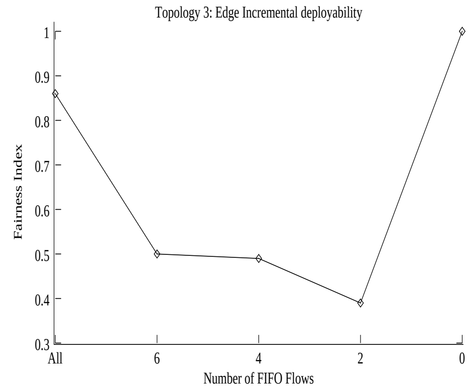

Fig. 8. Topology 3: Edge Incremental Deployability

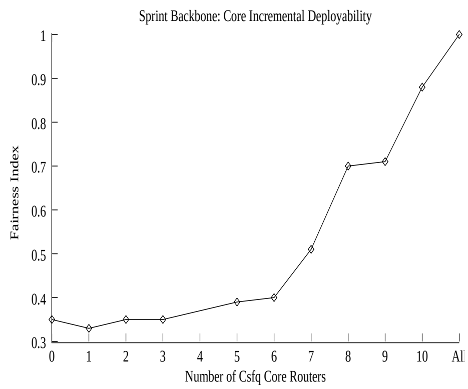

Fig. 9. Topology 4: Core Incremental Deployability

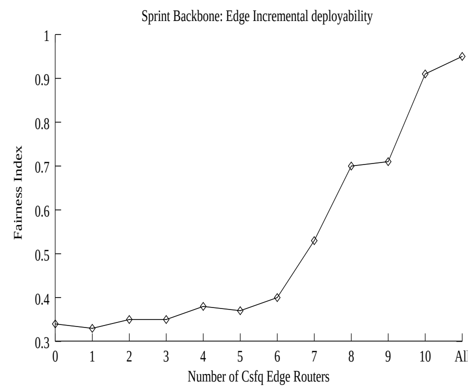

Fig. 10. Topology 4: Edge Incremental Deployability

– While the unfair allocation because of inconsistent state can be plausibly perceived as a transient phenomenon (if edge routers are assumed to adapt to the fair share feedback they receive, it can be shown that flows will eventually receive their fair allocations at the fair routers), this is true only if the unfair allocations at the drop-tail routers remain stable. In other words, given an arbitrarily indiscriminate rate allocation at legacy-routers, that fluctuates with time, the fair share computation at fair routers will fail to converge, causing unfair allocation at fair routers to become a persistent phenomenon.

Impact of Legacy Edge Routers While core routers are responsible for fair allocation, edge routers in core-stateless approaches are responsible for conveying to the core routers, the dynamic state used in their rate allocation schemes. In the event that an edge router is a legacy router, it will fail to convey any such dynamic state information to the core. Hence, the presence of such edge routers will result in *legacy-flows* co-existing with *fair-flows* in the core network. While this does not have any impact on legacy-core-routers (where no fair allocation schemes exist anyway), it obviously has a severe impact on the fairness achieved at fair-routers. Specifically, given that the legacy-flows carry no dynamic state, how should the fair routers treat the legacy-flows? Aggregating all legacy-flows into one logical flow might result in unfair allocations to legacy flows. A traffic engineering solution (wherein, the capacity of the network is partitioned between legacy- and fair-flows) might be possible, but would not be desirable in a pay-per-use service model [8].

The challenge then is to determine dynamically how the capacity at a fair router should be divided between legacy and fair-flows. Once the split is determined, it is sufficient to then provide fairness only among the fair-flows (within their allocation), as the legacy-flows do not expect any fair allocation in the first place. However, it is critical for legacy flows not to be penalized in any way due to the upgrade of a part of the network[2].

5 Summary

We study the incremental deployability of the core-stateless fair queuing (CSFQ) mechanism. Based on our simulations, we conclude that CSFQ is not incrementally deployable. However, to be fair to its authors, CSFQ was not designed to be incrementally deployable [2]. Our motivation for the study was to gain insights that can help in the design of an incrementally deployable core-stateless QoS model. We present some of the insights in Section 4.

6 Acknowledgments

We thank the Yamacraw organization (http://www.yamacraw.org) for their generous support and funding part of this work.

References

1. H. Zhang, "Service Disciplines For Guaranteed Performance Service in Packet-Switching Networks", Proceedings of the IEEE, 83(10), Oct 1995..
2. I. Stoica, S. Shenker, H. Zhang, "Core-Stateless Fair Queuing: A Scalable Architecture to Approximate Fair Bandwidth Allocations in High Speed Networks", Proceedings of ACM SIGCOMM, 1998.
3. I. Stoica, H. Zhang, "Providing Guaranteed Services Without Per Flow Management" Proceedings of ACM SIGCOMM, 1999.
4. R. Sivakumar, T. Kim, N. Venkitaraman, and V. Bharghavan, "Achieving Per-flow Rate Fairness in Core-Stateless Networks", Proceedings of International Conference on Distributed Computing Systems, 2000.

[2] Note that we perceive "unfair allocation" and "indiscriminate allocation" differently, with the former being clearly more undesirable than the latter.

5. T. Kim, R. Sivakumar, K-W. Lee, and V. Bharghavan, "Multicast Service Differentiation in Core-Stateless Networks", Proceedings of International Workshop on Network Group Communication, 1999.
6. Rossi, H. Peter, J. D. Wright, and A. B. Anderson, "Handbook on Survey Research", 1983.
7. K. Fall and K. Vardhan, "*ns* notes and documentation," available from http://www-mash.cs.berkeley.edu/ns/, 1999.
8. Lucent Technologies, "Network Quality of Service: Survey Report," .