



TD(05)013

2nd MCM, Colmar, France, February 3-4, 2005

A new approach to Service Differentiation: Non-Congestive Queueing

L. Mamas and V. Tsaoussidis
Dept. of Electrical & Computer Engineering
Demokritos University of Thrace, Greece
{emamas, vtsaousi}@ee.duth.gr

Abstract

We evaluate a new approach to service differentiation, based on distinguishing congestive and non-congestive data. We initially define as non-congestive data only minimal-size packets, which do not contribute to congestion themselves, but, however, do suffer delays caused by larger packets. Case examples of such data are typical sensor data. Non-Congestive Queueing (NCQ) needs not follow a state-full approach to traffic differentiation. We show that (i) non - congestive data gets a much better service from the network, (ii) congestive data suffers no extra delays; instead, both fairness and efficiency are occasionally improved due to lower contention. Having more flows finishing early may lead to better resource sharing.

Keywords

NCQ, non-congestive, queueing

Working Group 1

A new approach to Service Differentiation: Non-Congestive Queueing

L. Mamas and V. Tsaoussidis
Dept. Of Electrical & Computer Engineering
Demokritos University of Thrace, Greece
{emamas, vtsaousi}@ee.duth.gr

Abstract

We evaluate a new approach to service differentiation, based on distinguishing congestive and non-congestive data. We initially define as non-congestive data only minimal-size packets, which do not contribute to congestion themselves, but, however, do suffer delays caused by larger packets. Case examples of such data are typical sensor data. Non-Congestive Queueing (NCQ) needs not follow a state-full approach to traffic differentiation. We show that (i) non-congestive data gets a much better service from the network, (ii) congestive data suffers no extra delays; instead, both fairness and efficiency are occasionally improved due to lower contention. Having more flows finishing early may lead to better resource sharing.

1 Introduction

Typical scheduling paradigms of packet networks do not match well the requirements of non-congestive applications, which transmit minor data volumes but suffer, however, major queuing delays. Sensor applications are some recent examples that demand a new type of service differentiation. Such applications do not really cause significant delays, raising naturally the issue of whether they deserve a prioritized service or not. For example, a sensor-generated packet may experience zero delay favored by a prioritized scheduling scheme, at a zero cost to other congestive flows. We attempt to exploit this idea in the context of computer networks, and in terms of correctness and practicality. We introduce a new paradigm for service differentiation, based on the principle of packet contribution to queuing delay. That is, we distinguish two traffic classes, congestive and non-congestive; non-congestive traffic gets prioritized service for two reasons: first, a non-congestive flow does not cause significant delay and hence needs not suffer from delays. Second, since it does not cause delays

its service time should eventually have little impact on other flows and therefore, it may well get service first. The whole issue therefore is essentially turned down to exploiting the tradeoff itself. How much is the benefit for the non-congestive applications and how much is the cost for the congestive ones?

Our primary assumption is that non-congestive traffic is generated in the form of small and periodic packet transmission. Certainly, typical sensor applications do fall into that category; however, other applications may fall into this category as well, if we judge solely on the basis of packet length. We discuss this issue later in this paper. However, the concept of service differentiation based on congestive and non-congestive data is powerful indeed: no flow states are required and no packet marking either.

Although the idea sounds straightforward, the system properties and design details reveal interesting dynamics. We first show with simple scenarios that the benefit for non-congestive flows is impressive. In terms of system service, this is translated into increased user satisfaction due to the increased amount of applications finishing much earlier. Notably, system performance is not degraded. Total time of applications system-wide is decreased. Furthermore, the gained performance of non-congestive flows is not always counterbalanced by a corresponding performance loss of congestive flows.

Although the proposed algorithm is still in experimental face we do provide feedback on several issues: What amount of traffic should get prioritized service? Is there a performance loss when no congestive flows exist? Is there a loss when no non-congestive flows exist? However, more details need to be parameterized and evaluated and further analysis is on the works. At this stage of research, our results clearly indicate the applicability of non-congestive queueing (NCQ) for packet networks.

In section 2 we discuss the related work and detail our proposal. In section 3, we provide the pseudo-

code of NCQ and present the basic assumptions, fundamental concepts and projected outcome. In section 4 we discuss the evaluation plan and metrics along with a justification for the plan. Furthermore, we present the results, analysis and justification. We show the impact of the percentage of packets that get prioritized service on system performance. In section 5 we discuss several open issues and address reasonable concerns. In section 6 we summarize our conclusions and future work.

2 Related Work

Service differentiation in computer networks is a topic of research, mainly focused on supporting application requirements for delay and bandwidth. Differentiation is mainly an operation that deals with reallocating bandwidth and controlling delay for the benefit of strict-service-requiring applications and at the cost of more flexible applications. Therefore, service differentiation is based on the principle 'get better service - if you need better service'.

Internet service differentiation has not been designed on the basis of theoretical research; rather, it was driven by the need for supporting real-time multimedia applications over the Internet. Such applications do have strict bandwidth and delay requirements; the flows that generate can describe their specifications in broad or detailed terms, and the network can plan for guaranteed service or otherwise for somewhat better service. Two approaches have gained acceptance recently: According to diffserv [1], the inherent properties of packet-switched Internet are masked with a number of gentle mechanisms, naturally matching Internet's structure, which - one way or another - shape traffic at some functionally-enhanced nodes. Alternatively, with intserv [2] the architecture itself can be redesigned to allow for guarantees through signaling and reservation.

Differentiation in both the aforementioned cases is application-specific and naturally oriented either by some explicit and strict flow characteristics or by some application class. Even in the latter case, associating application types with service classes requires a rather sophisticated implementation, ranging from packet marking, to shaping, scheduling and dropping schemes. Perhaps network engineering would have been different had the pressing demand of application requirements been ignored. For example, a natural principle to lead the design of network services (and consequently the service differentiation policy) could have been the network ability to function, the number of users serviced better without damaging the

rest, or the service offered on the basis of the cost to other applications. It is not unnatural to service first applications that require minimal time for service; in that case the gain for such applications can be significant, while the cost for the other applications may be small.

A similar scheduling discipline has been studied in operating systems where some schedulers select processes based on their completion time, rather than the time they started (shortest job first). Such a service alone may lead to starvation in case the rate of small processes is sufficient to keep the processor busy; processes demanding more time for completion could never get their turn. However, due to the cost of context switch and the limited concurrent presence of processes, this domain had limited scheduling flexibility; our service differentiation scheme should guarantee not only better service for non-congestive data but also a non-degraded service to congestive applications. Thus, only a limited amount of non-congestive data should be able to benefit from the differentiated service; once the cost of better service to some flows impacts the service to other flows, service differentiation should terminate.

A lot has been done in the networking community aiming at controlling traffic based on its characteristics. Controlling is implemented either through scheduling or through dropping policies mainly aiming at penalizing high - bandwidth - demanding flows rather than favoring low - bandwidth - demanding flows. In [6] Floyd and Fall introduced mechanisms based on the identification of high-bandwidth flows from the drop-history of RED. The RED-PD algorithm (RED with Preferential Dropping) [8] uses per-flow preferential dropping mechanisms. Two other approaches that use per-flow preferential dropping with FIFO scheduling are Core-Stateless Fair queuing (CSFQ) [14] and Flow Random Early Detection (FRED) [7]. CSFQ marks packets with an estimate of their current sending rate. The router uses this information in conjunction with the flow's fair share estimation in order to decide whether a packets needs to be dropped. FRED does maintain a state although only for the flows which have packets in the queue. The flows with many buffered packets are having an increased dropping propability.

The CHOKe mechanism [12] matches every incoming packet against a random packet in the queue. If they belong to the same flow, both packets are dropped. Otherwise, the incoming packet is admitted with a certain propability. However, a high-bandwidth flow may have only a few packets in the queue. Authors of [11] continued the CSFQ and CHOKe approaches. Their proposed mechanism keeps a sample

of arriving traffic. A flow with several packets in the sample, has an increased dropping probability. The Stochastic Fair Blue (SFB) [5] uses multiple levels of hashing in order to identify high-bandwidth flows. As the authors state, their mechanism works well only with a few high-bandwidth flows. Anjum and Tassiulas proposed in [3] a mechanism that drops packets based on the buffer occupancy of the flow while ERUF [13] uses source quench to have undeliverable packets dropped at the edge routers. On the other hand, SRED [10] caches the recent flows in order to determine the high-bandwidth flows.

3 Non-Congestive Queueing

NCQ is based on the principle 'Less Impact - Better Service' (LIBS). Although one can generalize the perspective to apply to all incoming packets, we only deal with two classes of packets: very small packets, deterministically set to 100 bytes and long packets that typical Internet applications use for data transfers. A natural question therefore is what if small-packet rate reaches levels, which delay significantly long-packet transmission. We complement the differentiating scheme with a service threshold: Non-congestive traffic cannot exceed a predetermined $ncqthresh$ percentage of prioritized service. We investigate the impact of $ncqthresh$ in the result section. NCQ uses priority queuing to implement priority service. That is, with in the same buffer, each packet is checked for its length, contrasted to a current priority rate and gets priority whenever it satisfies two conditions: (i) length is below 100 bytes and (ii) priority rate is below $ncqthresh$.

The algorithm below shows the pseudo-code for NCQ:

```

For every received packet
Begin
    Count received packets
    if (packetLength<100)
        and
        (non-congestive packets /
         received packets < ncqthresh)
    then begin
        packet gets high priority
    end
    else begin
        Count non-congestive packets
        packet gets normal priority
    end
End

```

Clearly, applications that utilize small packet transfers at high rate may benefit from NCQ. As soon as

the rate of priority reaches $ncqthresh$, no further prioritized service is allowed. We discuss in the open issues section more sophisticated ways to allow better fairness among non-congestive flows, i.e., not to allow a single application to monopolize the priority service.

On the other hand, a typical application could be transformed into a small-packet application intentionally. Since the amount of gain is limited by the $ncqthresh$ and the packet length, the transform should cause that much overhead and extended communication time that naturally the penalty of transformation will be greater than the gain.

4 Evaluation

4.1 Evaluation Methodology

We have implemented our evaluation plan on the ns-2 network simulator. The network topology used as a test-bed is the typical single-bottleneck dumbbell, as shown in Figure. 1. The bw_1, bw_3 links are 10Mbps. The link's capacity (bw_2) is 10Mbps, unless it is explicitly stated otherwise. We used equal number of source and sink nodes. The number of flows ranges from 10 to 100. Due to the deterministic nature of the experiments, statistical validity is not an issue.




Figure 1: Simulation topology

In order to validate our statements, we selected and evaluated the TCP NewReno protocol with both DropTail and Non-Congestive queuing algorithms (with buffer size of 100 packets). We used both FTP non-congestive flows and regular FTP flows¹.

We measured the efficiency of each flow using Goodput: $Goodput = Original_Data / Time$ where $Original_Data$ is the number of bytes delivered to the high-level protocol at the receiver (i.e. excluding retransmitted packets and overhead) and $Time$ is the amount of time required for the data delivery.

We used the system Goodput in order to measure the overall system efficiency. The system Goodput

¹The non-congestive flows have a maximum segment size of 100 bytes and the regular flows a packet size of 1000 bytes

is defined as: $SystemGoodput = \frac{\sum(Goodput_i)}{n}$ where $Goodput_i$ is the Goodput of the i^{th} flow and n the flow number.

Fairness is measured by the Fairness Index, derived from the formula given in [4] and defined as: $Fairness = \frac{\sum(Throughput_i)^2}{n \sum(Throughput_i)}$ where $Throughput_i$ is the Throughput of the i^{th} flow and n the flow number.

We measure application efficiency and user satisfaction based on the worst / average task completion time as well as on the number of completed tasks.

4.2 Evaluation Results

In the following results we used $ncqthresh$ 0.05. The non-congestive flows were the 5% of all flows. As we can see from Figure 2, the system has significant performance gains (upto 16% - in case of 90 flows) in terms of Goodput. This is not unreasonable: the impact of timeouts caused by short packets is more significant for non-congestive flows compared with the impact of long packets; i.e., regardless of the packet length, timeout is the same and extending total time for a small retransmission degrades system throughput. Due to the fact that the non-congestive flows are favored, their average / worst time is decreased (Figures 5, 7). It is very interesting to note that the NCQ has a positive impact also in the system's average / worst time (Figures 4, 6). According to Figure 6, all tasks are completed upto 6.5 seconds sooner (an improvement of about 40% - in case of 100 flows).




Figure 2: System Goodput

As we can see from Figure 8, the Non-Congestive Queueing (NCQ) leads to increased user satisfaction due to the increased number of applications finishing much earlier. Also, since flows are faced earlier with less contention, they seem to be able to share the channel better. Additionally, the system appears fair (Figure 3). As we can also see in Figure 2, less contention leads to increased goodput.

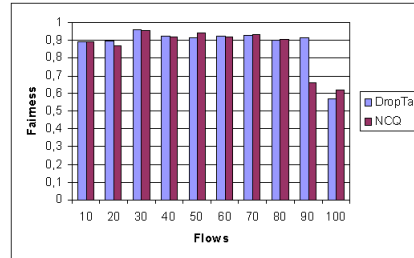


Figure 3: Fairness




Figure 4: Average Time of all Flows

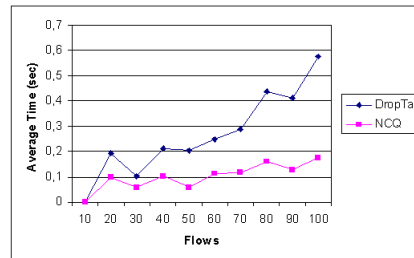


Figure 5: Average Time of Non-Congestive Flows




Figure 6: Worst Time of all Flows

Impact of traffic threshold ($ncqthresh$)

We carried out the same experiment with different values of $ncqthresh$ (0.01, 0.03, 0.05). As we can see from Figure 9, in average, the value of 0.05 is a good choice. Additionally, Figure 9 indicates that increasing further the threshold does no guarantee better service. For example, in Figure 9 flows with $ncqthresh$




Figure 7: Worst Time of Non-Congestive Flows




Figure 8: Tasks Completed

0.03 finish earlier. More detailed evaluation results can be found in [9].




Figure 9: Tasks Completed and Traffic Threshold

5 Open Issues

Although we demonstrated NCQ’s high potential, we do not presently address all concerns necessary to justify its need for deployment. For example, a single application which utilizes small-size packets for data transmission may be congestive indeed. That means, such application will probably monopolize the priority service for themselves. We are working on an extension of the algorithm to assign probabilistically priority service to small packets. The probability per packet decreases as the rate of non-congestive packet exceeds the $ncqthresh$. Probabilistic priority will guarantee fairness among non-congestive flows in a similar fashion to RED’s probabilistic dropping. Alternatively, the $ncqthresh$ is dynamically adjusted,

based on the projected outcome. In another front of research, acks and control packets may benefit from the priority treatment. Control packets prioritization are expected to boost the performance of short-lived flows (mice) increasing fairness compared to long-lived flows (elephants). Acks are expected to increase transmission rate; how far this can happen (considering also the delayed-ack scheme which is widely deployed) and how far it can impact congestion control is an open issue.

6 Conclusions

We have shown that NCQ is a simple but powerful tool for differentiation, particularly beneficial for applications that utilize small rates and short packets, such as typical sensor applications. We discovered that a limited prioritization has a dual impact: it benefits non-congestive flows significantly and reduces contention among the congestive flows, resulting occasionally in better delays and throughput for all applications. NCQ is still in experimental phase. Further analytical and experimental results with more sophisticated mechanisms are underway.

References

- [1] Differentiated services charter. Technical report, Web Page: <http://www.ietf.org/html.charters/diffserv-charter.html>.
- [2] Integrated services charter. Technical report, Web Page: <http://www.ietf.org/html.charters/intserv-charter.html>.
- [3] F. M. Anjum and L. Tassiulas. Fair bandwidth sharing among adaptive and non-adaptive flows in the internet. In *IEEE InfoCom 99*, March 1999.
- [4] D. Chiu and R. Jain. Analysis of the Increase/Decrease Algorithms for Congestion Avoidance in Computer Networks. *Journal of Computer Networks and ISDN*, 17(1):1–14, June 1989.
- [5] W. Feng, D. Kandlur, D. Saha, and K. G. Shin. BLUE: A new class of active queue management algorithms. Technical Report CSE-TR-387-99, April 1999.
- [6] S. Floyd and K. Fall. Promoting the use of end-to-end congestion control in the Internet. *IEEE/ACM Transactions on Networking*, 7(4):458–472, 1999.

- [7] D. Lin and R. Morris. Dynamics of random early detection. In *SIGCOMM '97*, pages 127–137, Cannes, France, september 1997.
- [8] R. Mahajan, S. Floyd, and D. Wetherall. Controlling high-bandwidth flows at the congested router, 2001.
- [9] L. Mamas and V. Tsaoussidis. A new approach to service differentiation: Non-congestive queueing. Technical Report TR-DUTH-EE-2004-14, December 2004.
- [10] T. J. Ott, T. V. Lakshman, and L. H. Wong. SRED: Stabilized RED. In *Proceedings of INFOCOM*, volume 3, pages 1346–1355, 1999.
- [11] R. Pan, L. Breslau, B. Prabhakar, and S. Shenker. Approximate fairness through differential dropping, 2001.
- [12] R. Pan, B. Prabhakar, and K. Psounis. CHOKe - A stateless queue management scheme for approximating fair bandwidth allocation. March 2000.
- [13] A. Rangarajan. Early regulation of unresponsive flows. Technical Report TRCS99-26, July 1999.
- [14] I. Stoica, S. Shenker, and H. Zhang. Core - stateless fair queueing: Achieving approximately fair bandwidth allocations in high speed networks. In *SIGCOMM*, pages 118–130, 1998.