Differentiating Services with Noncongestive Queuing (NCQ)

Lefteris Mamatas, *Member*, *IEEE*, and Vassilis Tsaoussidis, *Senior Member*, *IEEE*

Abstract—We discuss a new packet service paradigm, called "Less Impact Better Service" (LIBS), which is realized through a novel queuing discipline, called "Noncongestive Queuing" (NCQ). NCQ prioritizes small packets when conditions permit and utilizes service thresholds to confine the delay impact of prioritization on congestive applications. We show that LIBS and NCQ satisfy more users with diverse demands on delay and throughput. We obtained both analytical and simulation results, which are very promising. Diversity is simulated by using FTP, sensor, and VoIP traffic.

Index Terms—Service differentiation, packet scheduling, QoS.

1 Introduction

Typical scheduling paradigms of packet networks do not match well the requirements of noncongestive applications, which transmit minor data volumes but suffer, however, major queuing delays. 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 almost-zero delay favored by a prioritized scheduling scheme, at an almost-zero cost to other congestive flows. In the same context, a voice application, although it generates periodic traffic at small rates, experiences delays comparable with other applications that transfer large data volumes.

Typical service paradigms assume resource demand exceeding resource supply, thus focusing on bandwidth sharing among flows. Other service paradigms incorporate a proportional service scheme, focusing on bandwidth allocation in proportion to the demand. However, both perspectives lack a delay-oriented service discipline. Delay-oriented services have been traditionally managed based on delay requirements of some applications, which are eventually reflected in the prioritization during scheduling. Thus, service disciplines are primarily application-oriented and have the inherent property to satisfy some applications more, rather than satisfying more applications.

We exploit here a system-oriented service discipline, which targets in satisfying more users. We depart from two main observations:

1. Noncongestive flows do not cause significant delays and hence should not suffer from delays. We call this service discipline "Less Impact Better Service" (LIBS).
2. As it is with Internet voice where the users tolerate more delays with Internet than with dedicated telephone lines, users of long and congestive applications, such as FTP, tolerate more delays than the users of small-data sensor applications.

In this context, bandwidth alone could not have the central role; instead, efficient distribution of resources needs to be characterized by the delay suffered by each flow in relation to the delay they cause. The latter is occasionally associated with the delay that users tolerate. Our service approach promotes small-sized packets at small rates, which define "noncongestive" traffic. To avoid starvation and also significant delay impact on congestive traffic, noncongestive traffic is confined by corresponding service thresholds. Hence, we analyze the behavior of systems where noncongestive traffic has controlled prioritization without affecting congestive traffic. From a user perspective, applications that utilize small data packets and rates (and are also intolerant to long delays) are satisfied while other applications suffer almost-zero extra delays.

The key idea of Noncongestive Queuing (NCQ) [27], [28] that we discuss here departs from the operational dynamics of gateways: they may service small packets instantly. Noncongestive packets do not cause significant delays and hence should not suffer from delays. Although our approach sounds straightforward, the system properties and design details reveal interesting dynamics. The simplicity of NCQ algorithm reduces implementation and deployment effort. NCQ does not require any modification at the transport protocol or packet marking; a minor modification of the gateway's software is sufficient.

Our primary assumption is that some applications, such as sensors or even VoIP, generate packets in the form of noncongestive traffic. Typically, they transmit periodically small packets. For example, in [12] the authors claim that both the packet size and data rate in many typical sensor networks are very small. In [40], the authors assume a maximum sensor packet size of 44 bytes, while in [37], it is considered that typical packet sizes in a sensor network are 32, 64, 96, and 128 bytes.

- *L. Mamatas is with the Department of Electronic and Electrical Engineering, University College London, Torrington Place, WC1E 7JE London, UK. E-mail: l.mamatas@ee.ucl.ac.uk.*
- *V. Tsaoussidis is with the Department of Electrical and Computer Engineering, Democritus University of Thrace, Vas. Sofias 12, 67100 Xanthi, Greece. E-mail: vtsaousi@ee.duth.gr.*

Manuscript received 1 Nov. 2007; revised 7 July 2008; accepted 25 Sept. 2008; published online 23 Oct. 2008.
Recommended for acceptance by A. Zomaya.
For information on obtaining reprints of this article, please send e-mail to: tc@computer.org, and reference IEEECS Log Number TC-2007-11-0557.
Digital Object Identifier no. 10.1109/TC.2008.197.

0018-9340/09/$25.00 © 2009 IEEE Published by the IEEE Computer Society
Authorized licensed use limited to: University of Thrace (Democritus University of Thrace). Downloaded on December 03,2020 at 09:09:21 UTC from IEEE Xplore. Restrictions apply.

Sensor and voice data have strict requirements in delay and the service received by the network cannot be judged on the basis of achieved throughput. This observation calls for a new metric for application fairness as well, which relies mainly on the delay rather than on throughput. We introduce the *Application Satisfaction Index (ASI)*, which captures the delay fair share per application on the basis of the delay impact of each application on others. Thus, ASI reflects how fairly applications receive service, delay-wise, under diverse delay expectations and impact.

The proposed service paradigm impacts other performance measures as well, such as energy expenditure, which is very significant indeed for energy-limited sensors or VoIP mobile devices. The gains in energy are achieved through reduction of the communication, and hence, application time and their importance varies depending on the device itself, the communication pattern, the network contention, and so forth. We show that NCQ improves energy efficiency of sensor devices and other noncongestive applications, without damaging the dynamics of multiple flow equilibrium and without causing any statistically important *Goodput* losses to the congestive flows.

In Section 2, we discuss the related work. In Sections 3 and 4, we provide the pseudocode of NCQ and present the basic assumptions, fundamental concepts, and projected outcome, based on analytical methods. Furthermore, we explore the impact of different application strategies regarding the size of the transmitting packets, assuming that a LIBS-based approach is deployed. In Section 5, we discuss the experimental evaluation plan and metrics, along with a justification for the plan. Furthermore, we present the results, analysis, and justification and demonstrate the impact of service thresholds on system performance. In Section 6, we discuss several open issues and address reasonable concerns. In Section 7, we summarize our conclusions.

2 Related Work and Discussion

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 basic approaches have gained acceptance: According to diffserv [3], 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 [4], 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 concept 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, the lack of precision in estimating cost-per-process, and the limited concurrent presence of processes, this domain had limited scheduling flexibility; our service differentiation scheme guarantees better service for noncongestive data only as far as the service to congestive applications is not degraded. Thus, only a limited amount of noncongestive data should be able to benefit from our differentiating scheme.

According to [2], the average delay for the system tends to be reduced when customers with short service times are given high priority. The authors give examples from everyday life such as special checkout counters for customers with few items or the waiting lines at copying machines, where people often give priority to others who need to make just a few copies.

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 [14], Floyd and Fall introduced mechanisms based on the identification of high-bandwidth flows from the drop history of RED [15]. The RED with Preferential Dropping (RED-PD) algorithm [26] uses per-flow preferential dropping mechanisms. Two other approaches that use per-flow preferential dropping with FIFO scheduling are Core-Stateless Fair Queuing (CSFQ) [38] and Flow Random Early Detection (FRED) [23]. 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 packet needs to be dropped. FRED does maintain a state although only for the flows that have packets in the queue.

Authorized licensed use limited to: University of Thrace (Democritus University of Thrace). Downloaded on December 03,2020 at 09:09:21 UTC from IEEE Xplore. Restrictions apply.

The flows with many buffered packets are having an increased dropping probability.

The CHOKe mechanism [34] 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 probability. Pan et al. [33] extended the CSFQ and CHOKe approaches. The Stochastic Fair Blue (SFB) [13] uses multiple levels of hashing in order to identify high-bandwidth flows. Anjum and Tassiulas [1] proposed a mechanism that drops packets based on the buffer occupancy of the flow, while ERUF [36] uses source quench to have undeliverable packets dropped at the edge routers. On the other hand, SRED [32] caches the recent flows in order to determine the high-bandwidth flows.

In [18], the authors introduced the Alternative Best Effort (ABE) mechanism. ABE improves performance of delay-sensitive traffic but uses only two possible traffic classifications: delay- and throughput-sensitive. Delay-sensitive applications sacrifice throughput, and vice versa. Traffic Sensitive QoS (TSQ) mechanism [7] allows applications to indicate via marking their preferable delay/throughput sensitivity at packet level. However, their approach does not imply service per packet. That is, although applications may mark all or selected packets only, the basis for marking is not the current network state, which is unknown[1] by the application but rather the specific properties of a packet. The dynamics of such service interactions imply application-level QoS indeed. Application level QoS has therefore several undesirable properties:

- The markings are predetermined and cannot correspond to various possible packet sizes. They inherently produce a "quantization" error.
- Marks correspond to priorities and do not take into account the impact of prioritization. For example, a favorably marked packet will be serviced first even if the following packet will have zero impact to its service.
- Marks are application-level service requests; the system can find an optimal operating point for multiplexed applications only by introducing system-oriented criteria.

We note that packet marking and NCQ are not competitive technologies. The former can complement the latter through a second level of prioritization. For example, different priorities can be assigned via packet marking to the different noncongestive or congestive applications.

3 Our Approach: Noncongestive Queuing

We assume different classes of packets according to their size. NCQ is incorporated into routers to differentiate services according the impact of each traffic class on delay. For example, a class with smaller packets and sending rate receives better service than one with large packets or high sending rate. A natural question therefore is what if small-packet rate reaches levels that delay significantly long-packet transmission. We complement the differentiating scheme with a service threshold: The favored noncongestive traffic cannot exceed a predetermined threshold, called *ncqthresh*, which represents the upper limit of permitted prioritized service. The threshold typically reflects a service percentage for prioritization. However, this percentage corresponds to the number of packets; not the occupied buffer space. Indeed, since service prioritization applies for small packets only, the queue size that corresponds to the prioritized packets, percentage-wise, is much smaller.

On the other hand, a typical application could be intentionally transformed into a small-packet, high-rate application. Since the *ncqthresh* and the packet length confine the amount of gain, the transformation should cause that much overhead and extended communication time that naturally the penalty of transformation will be greater than the gain. In Section 4.2, we calculate numerically the impact of such transformation.

Although the perspective of NCQ is more general, initially, we only deal with two classes of packets: small packets (up to 130 bytes)[2] and long packets that typical Internet applications use for data transfers. The threshold of 130 bytes follows the assumption that typical sensor applications usually range from 32 to 128 bytes [37]. NCQ uses nonpreemptive priority queuing to implement priority service. That is, within the same buffer, each packet is checked for its length, contrasted to the current state of prioritized service rate and gets priority whenever it satisfies two conditions: 1) length is below 130 bytes and 2) prioritized service rate is below *ncqthresh*.

The algorithm below shows the pseudocode for NCQ:

```
for every received packet
begin
   count received packets
    (congestive and noncongestive)
    if (packetLength < 130)
     and
      (favored_packets/
       received_packets < ncqthresh)
    then
     packet receives high priority
     count favored packets
    else
      packet receives normal priority
    end
end
```

In this stage of our work, we do not favor ACKs even though they have a small size. We note that NCQ may occasionally favor a part of a noncongestive data flow. However, a possible slight increase[3] of the reordered packets is counterbalanced by the high number of packet drops that are avoided due to the prioritization.

4 Analysis

4.1 Impact of NCQ

Initially, we attempt to approach numerically the impact of NCQ priority on congestive traffic for any given proportion of traffic classes. We assume two classes of traffic (the noncongestive and congestive) that are formed by a large number of flows. We assume that all packets arriving at the

1. Even if it is measured, the precision and granularity of measurements are dubious.

2. For longer packets, we experimentally determined that the cost of prioritization for the congestive flows starts to be significant.

3. It is not significant in our experiments.

Authorized licensed use limited to: University of Thrace (Democritus University of Thrace). Downloaded on December 03,2020 at 09:09:21 UTC from IEEE Xplore. Restrictions apply.

TABLE 1
Notation Table

Symbol	Description
λ_1	Arrival rate of class 1
λ_2	Arrival rate of class 2
T_{S1}	Average service-time of class 1
T_{S2}	Average service-time of class 2
$\lambda = \lambda_1 + \lambda_2$	Total arrival rate
$u_1 = \lambda_1 T_{S1}$	Utilization of class 1
$u_2 = \lambda_1 T_{S1} + \lambda_2 T_{S2}$	Cumulative utilization
T_{Q1}	Average queuing delay for class 1
T_{Q2}	Average queuing delay for class 2
T_Q	Average queuing delay

bottleneck queue follow a Poisson distribution.[4] Class 1 has priority over class 2. We use a nonpreemptive head-of-line priority system per class. Class 1 has smaller packets (so, average service time too) and lower packet arrival rate ($\lambda_1 < \lambda_2$). We summarize our notation in Table 1.

We define the following:

- **Waiting time.** Waiting time represents the amount of time a packet waits for service in the queue.
- **Service time.** Service time represents the amount of actual service time required by a packet and is proportional to its size.
- **Time-in-system.** Time-in-system is equal to the Waiting Time plus Service Time (in our case, it is the same as Queuing Delay).

The packet-departure rate is equal to the service distribution, because we are using a single server.

In the three different cases of prioritization below, we calculate the average queuing delay for each class and for the system:

1. Class 1 has full priority over class 2.
2. The two classes have the same priority (scheduling without priority).
3. Class 1 has priority over class 2 but only when less than the *ncqthresh* percentage of the total traffic is prioritized.

4.1.1 Case 1: Priority Scheduling

We calculate the average waiting time for each of the two classes as

$$T_{W1} = \frac{\lambda_1 T_{s1}^2 + \lambda_2 T_{s2}^2}{2(1-u_1)}, \tag{1}$$

$$T_{W2} = \frac{\lambda_1 T_{s1}^2 + \lambda_2 T_{s2}^2}{2(1-u_1)(1-u_2)}. \tag{2}$$

Consequently, the total average waiting time is equal to the average of T_{W1}, T_{W2} weighted by the arrival rate for each class

$$T_W = \frac{\lambda_1}{\lambda} T_{W1} + \frac{\lambda_2}{\lambda} T_{W2}. \tag{3}$$

4. It is widely adopted (such as [16], [30], [22], and [10]) that the packet arrival process for highly multiplexed environments tends to a Poisson Distribution.

We calculate the queuing delay for each class and estimate the total average time-in-system as

$$T_{Q1} = T_{W1} + T_{S1}, \tag{4}$$

$$T_{Q2} = T_{W2} + T_{S2}, \tag{5}$$

$$T_Q = \frac{\lambda_1}{\lambda} T_{Q1} + \frac{\lambda_2}{\lambda} T_{Q2}. \tag{6}$$

4.1.2 Case 2: Nonpriority Scheduling

Without a priority queue, the two classes (noncongestive and congestive) have the same average waiting time. In such case, the network utilization of the system is

$$u_{np} = u_1 = u_2 = \lambda_1 T_{S1} + \lambda_2 T_{S2}. \tag{7}$$

The service time is

$$T_{Snp} = \frac{\lambda_1}{\lambda} T_{S1} + \frac{\lambda_2}{\lambda} T_{S2}. \tag{8}$$

The average waiting time is

$$T_{W1np} = T_{W2np} = \frac{u_{np} T_{Snp}}{2(1-u_{np})}. \tag{9}$$

The average time-in-system is

$$T_{Q1np} = T_{W1np} + T_{S1}, \tag{10}$$

$$T_{Q2np} = T_{W2np} + T_{S2}. \tag{11}$$

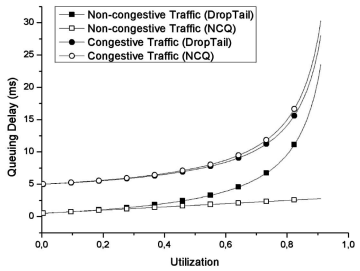
From (10) and (11), we obtain

$$T_{Qnp} = \frac{\lambda_1}{\lambda} T_{Q1np} + \frac{\lambda_2}{\lambda} T_{Q2np}. \tag{12}$$

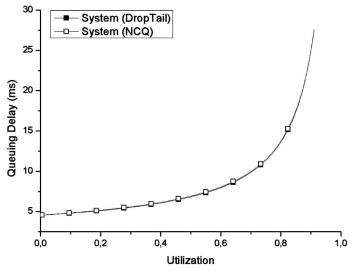
Using (4), (5), (6), (10), (11), and (12), we calculated the average queuing delays for each class as well as for the system, for two different percentages of noncongestive traffic (Figs. 1a, 1b, 1c, and 1d).

In Figs. 1a and 1b, the 10 percent of arriving packets form the noncongestive traffic (class 1) and the 90 percent form the congestive (class 2) traffic. Their service times are 0.5 and 5 ms, respectively. For high utilizations (exceeding 0.6), there is a slight increase in the queuing delay of the congestive traffic for a significant improvement in the average delay of the noncongestive traffic (Fig. 1a). When we increase the rate of the noncongestive packets to 20 percent and for high utilizations (exceeding 0.3), the impact of the prioritization on the congestive traffic appears significant (Fig. 1c). For both percentages of the noncongestive traffic (10 percent and 20 percent), the average queuing delay of the system remains statistically the same (Figs. 1b and 1d).

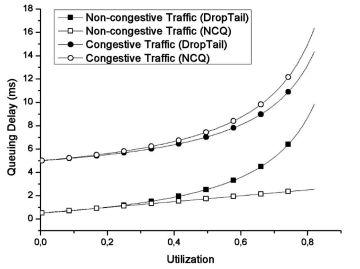
Authorized licensed use limited to: University of Thrace (Democritus University of Thrace). Downloaded on December 03,2020 at 09:09:21 UTC from IEEE Xplore. Restrictions apply.



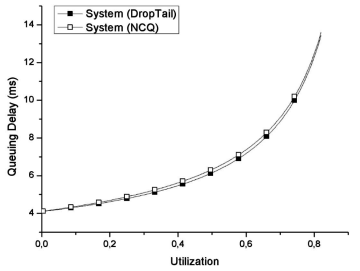
(a)



(b)



(c)



(d)

Fig. 1. Numerical results. (a) Average queuing delay of congestive and noncongestive traffic (10 percent of the packets are noncongestive). (b) System's average queuing delay (10 percent of the packets are noncongestive). (c) Average queuing delay of congestive and noncongestive traffic (20 percent of the packets are noncongestive). (d) Average queuing delay of congestive and noncongestive traffic (20 percent of the packets are noncongestive).

4.1.3 Case 3: Priority Scheduling with Ncqthresh

In the following analysis, we assume that only a portion of the noncongestive traffic is favored. More precisely, *Ncqthresh* represents the percentage of the total traffic that can be favored without any statistically important impact on the congestive traffic and corresponds to the $\frac{\lambda_1}{\lambda} * ncqthresh$ percentage of the noncongestive traffic, which we call *kthresh*. The two priority classes (1 and 2) consist of the *kthresh* percentage of noncongestive traffic and the $(1 - kthresh)$ percentage of the noncongestive traffic plus the congestive traffic, respectively.

We calculate the arrival rates and service times for each class as

$$\begin{aligned}
\lambda_1' &= kthresh * \lambda_1 = \frac{\lambda_1^2}{\lambda} * ncqthresh,\\
\lambda_2' &= (1 - kthresh) * \lambda_1 + \lambda_2\\
&= \left(1 - \frac{\lambda_1}{\lambda} * ncqthresh\right) * \lambda_1 + \lambda_2,\\
\lambda' &= \lambda_1' + \lambda_2',\\
T_{S1}' &= T_{S1},\\
T_{S2}' &= (1 - kthresh) * \frac{\lambda_1}{\lambda_2'} * T_{S1} + \frac{\lambda_2}{\lambda_2'} * T_{S2}\\
&= \left(1 - \frac{\lambda_1}{\lambda} * ncqthresh\right) * \frac{\lambda_1}{\lambda_2'} * T_{S1} + \frac{\lambda_2}{\lambda_2'} * T_{S2},\\
u_1' &= \lambda_1' T_{S1}',\\
u_2' &= \lambda_1' T_{S1}' + \lambda_2' T_{S2}'.
\end{aligned}$$

The average waiting time for each of the two traffic classes becomes

$$T_{W1}' = \frac{\lambda_1' T_{s1}'^2 + \lambda_2' T_{s2}'^2}{2(1 - u_1')}, \tag{13}$$

$$T_{W2}' = \frac{\lambda_1' T_{s1}'^2 + \lambda_2' T_{s2}'^2}{2(1 - u_1')(1 - u_2')}. \tag{14}$$

We calculate the waiting times of each class using the weighted average of the waiting times T_{W1}' and T_{W2}' as

$$\begin{aligned}
T_{W1} &= \frac{kthresh * \lambda_1}{\lambda_1} T_{W1}' + \frac{(1 - kthresh) * \lambda_1}{\lambda_1} T_{W2}'\\
&= kthresh * T_{W1}' + (1 - kthresh) * T_{W2}'\\
&= \frac{\lambda_1}{\lambda} * ncqthresh * T_{W1}'\\
&\quad + \left(1 - \frac{\lambda_1}{\lambda} * ncqthresh\right) * T_{W2}',
\end{aligned} \tag{15}$$

$$T_{W2} = T_{W2}', \tag{16}$$

$$T_{Q1} = T_{W1} + T_{S1}, \tag{17}$$

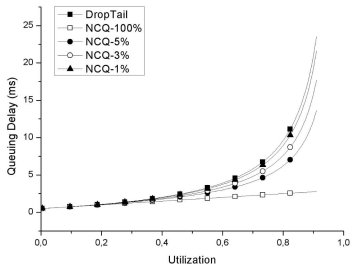
$$T_{Q2} = T_{W2} + T_{S2}. \tag{18}$$

For the system,

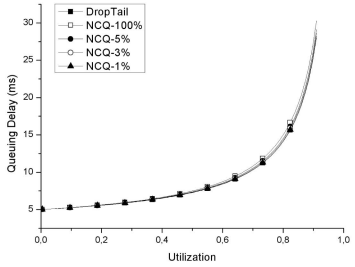
$$T_Q = \frac{\lambda_1}{\lambda} T_{Q1} + \frac{\lambda_2}{\lambda} T_{Q2}. \tag{19}$$

For a low *ncqthresh* value (e.g., 0.01 - 0.05), the impact of the prioritization on the average queuing delay of the congestive traffic is almost zero (Figs. 2b and 2d). Actually, *Ncqthresh* bounds the prioritization of the noncongestive traffic to a limit that is not harmful to the bandwidth exploitation of the congestive applications. In utilizations

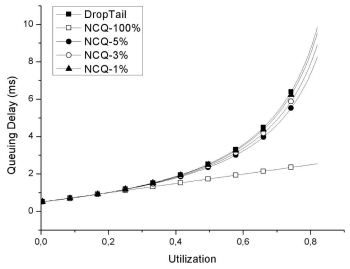
Authorized licensed use limited to: University of Thrace (Democritus University of Thrace). Downloaded on December 03,2020 at 09:09:21 UTC from IEEE Xplore. Restrictions apply.



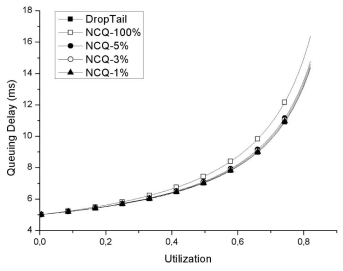
(a)



(b)



(c)



(d)

Fig. 2. Numerical results. (a) Average queuing delay of the noncongestive traffic (10 percent of the packets are noncongestive). (b) Average queuing delay of the congestive traffic (10 percent of the packets are noncongestive). (c) Average queuing delay of the noncongestive traffic (20 percent of the packets are noncongestive). (d) Average queuing delay of the congestive traffic (20 percent of the packets are noncongestive).

below 23 percent, the average queuing delay of the noncongestive traffic remains statistically the same. As the network utilization builds up, there are significant gains for the noncongestive applications in terms of delay and increase more for higher values of *ncqthresh* (Figs. 2a and 2c).

4.2 Design Strategies of Applications

Assuming that a LIBS-based mechanism is deployed, an application may follow alternative strategies in order to have performance gains. Here, we explore analytically such

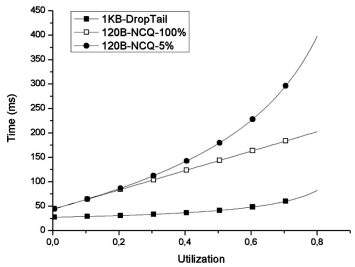


Fig. 3. Numerical results of FTP application.

strategies. We assign a specific task (i.e., to transfer 5 Mbytes) to each of the traffic classes and calculate the service time of each task.

The required number of transmitted packets for a specific class n (NP_n) is

$$NP_n = \frac{Total_Task_Data_n}{Packet_Size_n - Overhead_Per_Packet}. \quad (20)$$

The service time for a specific class n (ST_n) is

$$ST_n = T_{Q1} * NP_n. \quad (21)$$

The following analysis is based on (20) and (21). We assume that each packet has 40 bytes overhead.

4.2.1 FTP Application

We assume two different traffic classes that consist of large packets, in order to approach numerically a similar behavior to bulk-data transfer (FTP). Each class has 1-Kbyte packets. The 10 percent of the arriving packets form the first class and the 90 percent form the second one. An application designer may decide to use smaller packets in order to have performance gains due to the LIBS-based mechanism. Is this effective? We show that such applications should, in general, avoid splitting their packets into smaller ones.

In Fig. 3, we depict the total service time required for the first traffic class to transfer 5 Mbytes, in three distinct cases:

1. The FTP application uses 1-Kbyte packets.
2. The FTP application uses 120-byte packets, the NCQ mechanism has been deployed, but there is no limitation like *ncqthresh*.
3. The FTP application uses small packets (120 bytes) and the NCQ mechanism has been deployed with an *ncqthresh* with a value of 0.05.

According to Fig. 3, there is a significant performance decrease in terms of total service time for the NCQ mechanism without an *ncqthresh* limitation (e.g., 240 percent decrease for a utilization of 0.60). In the case of an *ncqthresh* with the value of 0.05, the performance decrease reaches 371 percent (for 0.60 utilization). Although the small packets are favored from the NCQ mechanism, the increased overhead due to the higher number of required packets overcomes these gains.

4.2.2 Congestive Multimedia Application

In this scenario, we used two traffic classes: 500-byte and 1-Kbyte packets, respectively. We assume that the first class follows the traffic pattern of a Congestive Multimedia Application (uses packets with average size). The 10 percent

Authorized licensed use limited to: University of Thrace (Democritus University of Thrace). Downloaded on December 03,2020 at 09:09:21 UTC from IEEE Xplore. Restrictions apply.

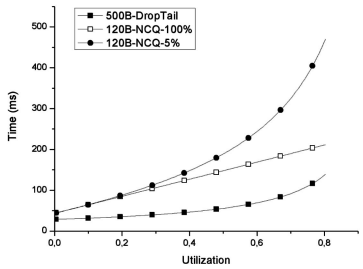


Fig. 4. Numerical results of congestive multimedia application.

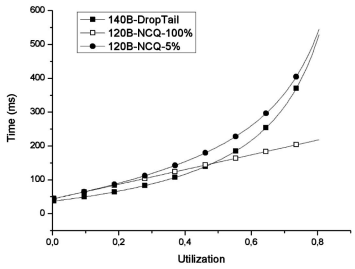


Fig. 6. Task completion time for noncongestive multimedia application.

of the arriving packets form the first class and the 90 percent form the second one. Although in this case we have less extra overhead, the performance gains of the NCQ mechanism are still inadequate to overcome this drawback (Fig. 4). For a utilization of 0.60, the performance decrease in case of NCQ without *ncqthresh* is 143 percent and in case of NCQ with *ncqthresh* of 0.05 is 253 percent.

4.2.3 Noncongestive Multimedia Application

In this scenario, we approach numerically a noncongestive multimedia application. The first class represents the multimedia application and forms the 10 percent of the traffic and the second class represents the FTP application. The multimedia application has a packet size of 140 bytes. In this scenario, we investigate whether it is worth for this application to change its packet size from 140 to 120 bytes.

For both versions of NCQ, we have a significant decrease in the average queuing delay (see Fig. 5) but is inadequate for a significant improvement in terms of task completion time due to the increased overhead (Fig. 6). As we can see from Fig. 6, for the NCQ without *ncqthresh* we have an improvement in terms of total service time for higher utilizations than 0.48 (e.g., 21 percent for a utilization of 0.60). In the case of NCQ with *ncqthresh* of 0.05, we have a performance decrease for all utilizations (e.g., 19 percent for a utilization of 0.60). Consequently, in special cases, such applications may adjust their packets in order to have performance gains.

To summarize, the policy to adjust the size of the transmitted packets to the NCQ threshold is very often inefficient. However, there are cases that an application can be favored from the NCQ mechanism, if it adopts smaller packets. For example, a noncongestive application may slightly reduce its packet size in order to have performance gains from a LIBS-based packet scheduling algorithm.

Furthermore, in case a malicious application tries to monopolize communication by transmitting small packets at high data rates, the extra introduced overhead is not counterbalanced by the gains. A Denial-of-Service (DoS) attack that uses aggressive flows consisting of small packets, in the worst case, would disable the extra prioritization of NCQ. However, if the attack uses large packets, many noncongestive applications would not suffer from the exhaustion of the resources.

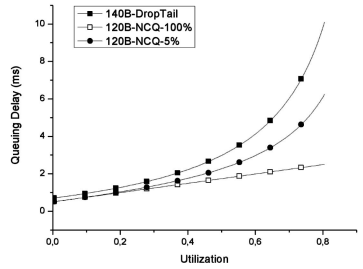


Fig. 5. Queuing delay for noncongestive multimedia application.

5 Evaluation

5.1 Evaluation Methodology

We have implemented our evaluation plan on the ns-2 network simulator [31]. We attempt to address five specific matters:

1. To show by simulation that numerical results do not lack any important parameter (Scenario 1: Impact of NCQ on FTP applications).
2. To show the impact of NCQ on sensor-based applications (Scenario 2: Internetworking with Sensors).
3. To evaluate the applicability of NCQ mechanism for VoIP traffic (Scenario 3: Impact of NCQ on VoIP traffic).
4. To explore whether NCQ can be incrementally deployed (Scenario 4: Incremental deployment of NCQ).

We discuss each scenario along with the corresponding results.

5.2 Evaluation Results

5.2.1 Scenario 1: Impact of NCQ on FTP Applications

In this scenario, we use a simple dumbbell topology, as shown in Fig. 7.

We measure

$$Goodput = \frac{Original_Data}{Time},$$

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 corresponding data delivery. We used

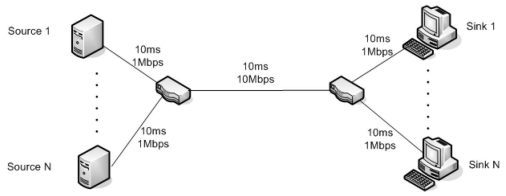


Fig. 7. Simulation topology.

Authorized licensed use limited to: University of Thrace (Democritus University of Thrace). Downloaded on December 03,2020 at 09:09:21 UTC from IEEE Xplore. Restrictions apply.

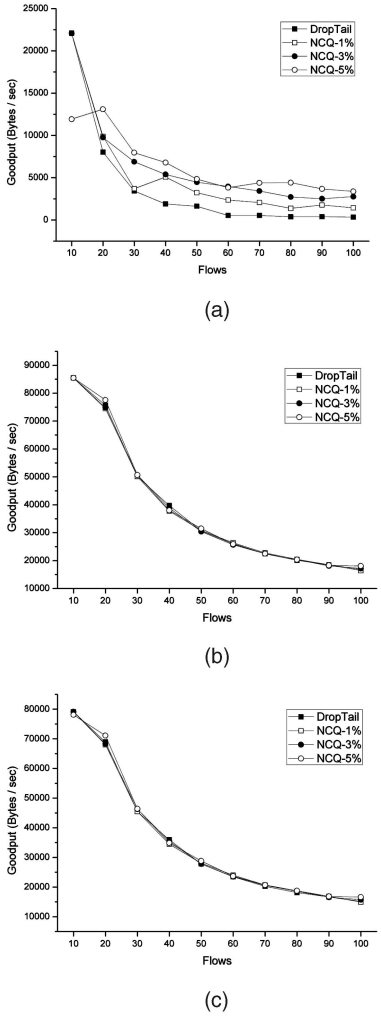


Fig. 8. Impact of NCQ on goodput. (a) Average goodput of noncongestive flows. (b) Average goodput of congestive flows. (c) Average system goodput.

the *Average Goodput* in order to measure the efficiency per flow. The *Average Goodput* for n flows is defined as

$$AverageGoodput = \frac{\sum_{i=1}^{n}(Goodput_i)}{n},$$

where $Goodput_i$ is the *Goodput* of the ith flow, and n is the flow number.

Additionally, we measure application efficiency and user satisfaction based on the *worst* and *average task completion time* as well as on the *number of completed tasks*. Every congestive FTP flow transmits 1-Mbyte data and every noncongestive FTP flow transmits 10-Kbyte data. We assume that a task is completed with the successful transmission of the carried data of the corresponding flow.

Furthermore, we experiment with different traffic thresholds and traffic class proportions to demonstrate the overall system behavior when the NCQ parameters change. We carried out the same experiment with different values of *ncqthresh* (i.e., 0.01, 0.03, and 0.05) in order to evaluate the impact of traffic threshold (*ncqthresh*). As we can see from the following results, in average, the value of 0.05 is a good choice.

At this point, we adjust the number of noncongestive flows to the 10 percent of the total flows. As we can see in

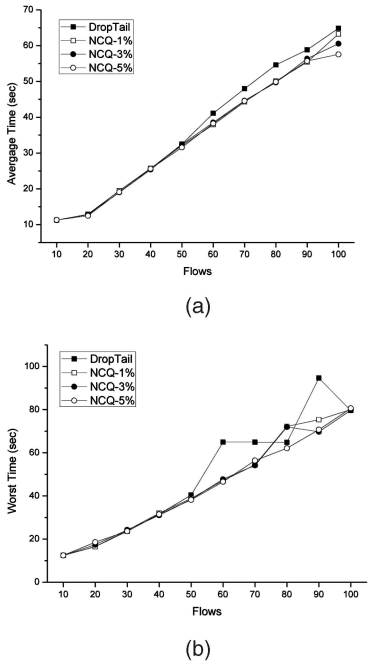


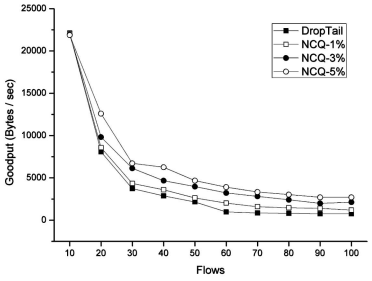
Fig. 9. Impact of NCQ on FTP applications. (a) Average task completion time in the system. (b) Worst task completion time in the system.

Fig. 8a, the noncongestive flows achieve significant performance gains (e.g., 10.43 times—in the case of 80 flows and *ncqthresh* of 5 percent) in terms of *Goodput*. The impact of the prioritization on the congestive flows is not only insignificant but we also notice occasionally a slight performance improvement in terms of *Goodput* for the congestive flows (e.g., 7.6 percent improvement—in the case of 100 flows and *ncqthresh* of 5 percent) and the overall system (Figs. 8b and 8c). This is not unreasonable: the impact of time-outs caused by short packets is more significant for noncongestive flows compared with the impact of long packets; i.e., regardless of the packet length, time-out is the same and extending total time for a small retransmission degrades system throughput. It is very interesting to note that the NCQ has a positive impact also in the system's average/worst task completion time (Figs. 9a and 9b). According to Fig. 9b, all tasks are completed up to 24.8 seconds sooner (i.e., in the case of 90 flows and *ncqthresh* of 3 percent).

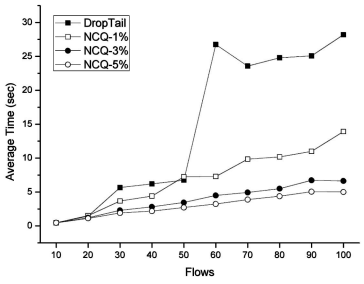
Next, we varied the traffic proportion of congestive and noncongestive flows. As the rate of the noncongestive flows increases (i.e., the number of the noncongestive packets increases), the benefit for noncongestive packets is gradually decreasing. This behavior is not symptomatic: the *ncqthresh* value confines the priority service to guarantee small impact on congestive flows. For example, we can see in Fig. 10a (rate of noncongestive packets 20 percent) that the NCQ algorithm favors now a smaller portion of the noncongestive packets. However, the impact on the average task completion time of the noncongestive flows remains significant (Fig. 10b).

Additionally, in Fig. 11, we depict the number of completed tasks at different time instances for both congestive and noncongestive flows. In this example, we consider the case of 50 flows. We show that, on average, more applications finish earlier. For example, at the

Authorized licensed use limited to: University of Thrace (Democritus University of Thrace). Downloaded on December 03,2020 at 09:09:21 UTC from IEEE Xplore. Restrictions apply.



(a)



(b)

Fig. 10. Impact of NCQ on noncongestive applications. (a) Average goodput of noncongestive flows (20 percent of the packets are noncongestive packets). (b) Average time of noncongestive flows (20 percent of the packets are noncongestive packets).

30.5th second, 7 flows are finished in the case of DropTail and 10 flows in the case of the NCQ-1 percent.

We adjusted the task of the FTP flows to be 100 Kbytes, assuming short-lived flows (e.g., like a typical Web page). As we can see from Figs. 12a and 12b, the noncongestive flows have a significant improvement in terms of goodput without noticeable impact on the goodput of the short-lived FTP flows.

5.2.2 Scenario 2: Internetworking with Sensors

For a more realistic scenario of internetworked sensor applications, we integrated NRL's Sensor Network Extension [11] into ns-2. Our scenario consists of a sensor network (a grid of 25 wireless sensors) and a simple dumbbell wired topology (see Fig. 13). The sensor application periodically transfers data to a wired node (the data collector). The sensor application notifies the data collector about the behavior of a moving phenomenon. The simulated phenomenon has a pulse period of 0.01 seconds. The sensor-generated data and the congestive FTP flows coexist in the same link and cross the same NCQ-enabled

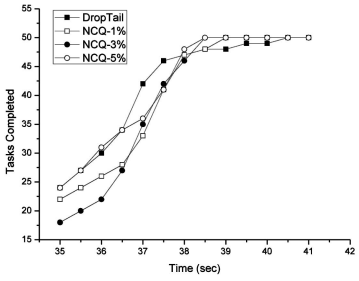
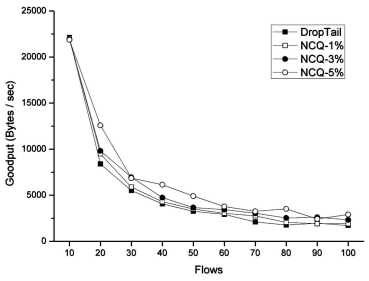
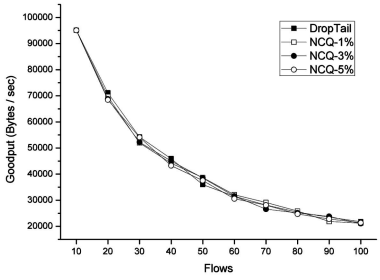


Fig. 11. Tasks completed (50 flows).



(a)



(b)

Fig. 12. Impact of NCQ on noncongestive applications and short FTP flows. (a) Average goodput of noncongestive flows. (b) Average goodput of short FTP flows.

gateway. The number of FTP flows ranges from 10 to 100. We used the TwoRayGround radio-propagation model and the AODV [35] routing protocol. We measured *Goodput* on both congestive (FTP) and sensor-related noncongestive traffic.

We evaluate the energy efficiency of NCQ using the *Energy Potential* (*EP*) [24] index

$$EP = 1 - \left(a \frac{Throughput - Goodput}{Throughput_{max}} + b \frac{Throughput_{max} - Throughput}{Throughput_{max}} \right).$$

The *EP* index takes into account the difference of achieved *Throughput* from *maximum Throughput* ($Throughput_{max}$) for the given channel conditions along with the difference of *Goodput* from *Throughput*, attempting to locate the *Goodput* as a point within a line that starts from 0 and ends at $Throughput_{max}$.

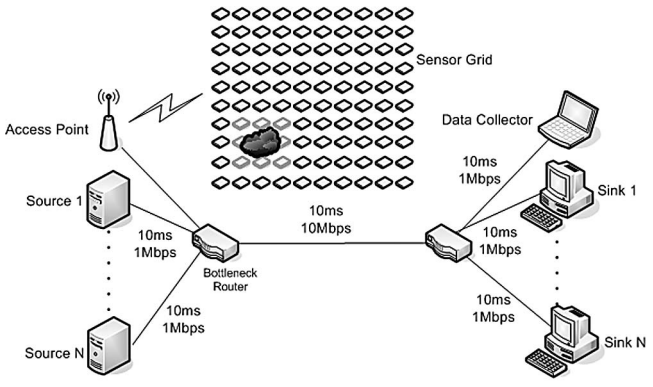


Fig. 13. Simulation topology.

Authorized licensed use limited to: University of Thrace (Democritus University of Thrace). Downloaded on December 03,2020 at 09:09:21 UTC from IEEE Xplore. Restrictions apply.

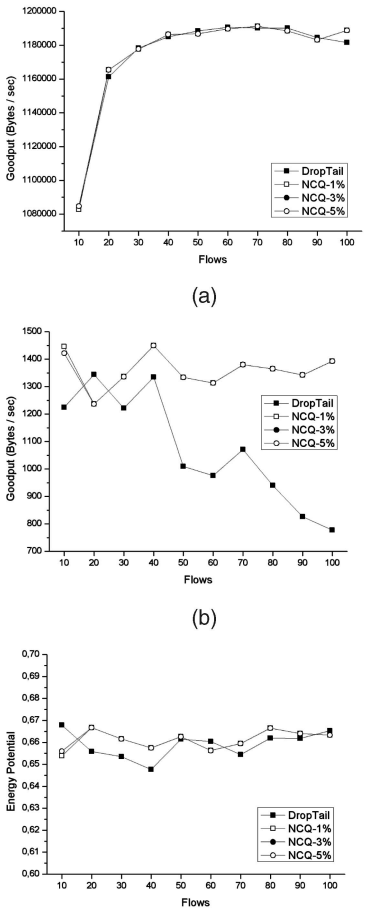


Fig. 14. Impact of NCQ on goodput/energy efficiency for sensor internetworking. (a) Goodput of FTP flows. (b) Goodput of sensor applications. (c) EP (noncongestive flows).

In order to measure fairness in the context of LIBS, we introduce *ASI*. *ASI* is defined as

$$ASI = 1 - \frac{\left| \sum_{i=1}^{n} Delay_i - \frac{Data_i}{TotalData} Delay_{max} \right|}{nDelay_{max}},$$

where n is either the number of active nodes or the number of different traffic classes; $Data_i$ is the total transmitted data of the ith node to the receiver application; *TotalData* is the total transmitted data of all nodes; $Delay_i$ is the average queuing delay of the ith node; and $Delay_{max}$ is the maximum queuing delay of the system. *ASI* ranges from 0 to 1.

Unlike other fairness indices (such as [6], [29], and [39]), *ASI* captures the deviation of the actual delay and the expected delay per flow. Note, however, that expected delay is determined by the factor $\frac{Data_i}{TotalData}$. In this context, *ASI* represents fairness of the *LIBS* architecture, since the expected delay per packet (and in turn, per flow) grows in proportion to the volume of their transmitted packets.

In Fig. 14a, we depict the *Goodput* performance of the congestive flows, and in Figs. 14b, 14c, 15a, and 15b, we demonstrate the *Goodput*, the *EP*, and the *ASI* for both overall system and sensor applications. We argue that the congestive FTP flows are not suffering from any important

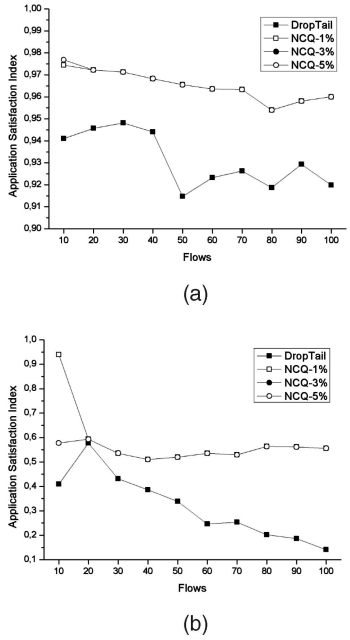


Fig. 15. Impact of NCQ on fairness for sensor internetworking. (a) ASI for the system. (b) ASI of sensor applications.

performance loss (see Fig. 14a), while the performance gains for the noncongestive flows are significant (up to 92 percent) in terms of *Goodput* as well as in terms of energy efficiency (see Fig. 14c). Additionally, the system appears fair according to *ASI* (Figs. 15a and 15b). We note that, due to *Goodput* increase of sensor applications, fairness performance may also be captured by the traditional index of fairness.

As we demonstrate in Figs. 14a, 14b, 14c, 15a, and 15b, the value of *ncqthresh* does not impact the results when the system has more than 20 flows since, from that point onward, the percentage of noncongestive traffic is less than 1 percent. So, every packet generated by the sensor applications is prioritized.

5.2.3 Scenario 3: Impact of NCQ on VoIP Traffic

We simulate VoIP traffic based on the following assumptions: During a conversation, speakers alternate between activity and idle periods. Taking into consideration the ON and OFF periods [5], as well as the heavy-tailed characteristics and self-similarity of VoIP traffic [9], we use the *Pareto* distribution for modeling the call holding times. We configure *Pareto* with a mean rate to correspond to transmission rate of 64 Kbps and the shape parameter is set to 1.5. In accordance with [5], we distribute the ON and OFF periods with means of 1.0 and 1.35 seconds, respectively. We simulate VoIP streams of 64 Kbps (following the widely used ITU-T G.711 [21] coding standard), and we set packet sizes at 120 bytes (i.e., each packet carries 10-ms G.711-encoded speech and a 40-byte packet header).

In the following results, we characterize the quality of voice communication using the R-Factor, which is included in the E-Model [19], [20] (an ITU-proposed analytic model of voice quality). R-Factor captures voice quality and ranges from 100 to 0, representing best and worst qualities, respectively. R-Factor is also associated to the familiar estimated Mean Opinion Score (MOS). MOS is equal to the arithmetic average of opinions, where "excellent" quality is

Authorized licensed use limited to: University of Thrace (Democritus University of Thrace). Downloaded on December 03,2020 at 09:09:21 UTC from IEEE Xplore. Restrictions apply.

TABLE 2
R-Factor, Quality Ratings, and MOS

R-Factor	Quality of voice rating	MOS
$90 < R < 100$	Best	4.34 - 4.5
$80 < R < 90$	High	4.03 - 4.34
$70 < R < 80$	Medium	3.60 - 4.03
$60 < R < 70$	Low	3.10 - 3.60
$50 < R < 60$	Poor	2.58 - 3.10

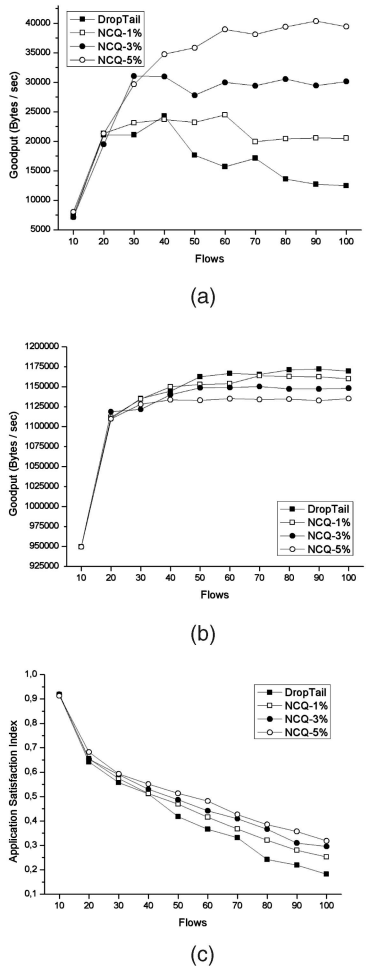


Fig. 16. Impact of NCQ on goodput and fairness. (a) Goodput of noncongestive VoIP application. (b) Goodput of congestive application (FTP). (c) ASI of VoIP application.

given a score of 5, "good" a 4, "fair" a 3, "poor" a 2, and "bad" a 1. R-Factor incorporates several different parameters, such as echo, background noise, signal loss, codec impairments, and others. In [8], the authors simplified E-Model to transport-level measurable quantities and resulted in a more suitable R-Factor formula. Based on the above, we define R-Factor as

$$R = \alpha - \beta_1 d - \beta_2 (d - \beta_3) H(d - \beta_3) - \gamma_1 - \gamma_2 ln(1 + \gamma_3 e), \tag{22}$$

where $\alpha = 94.2$, $\beta_1 = 0.024\ \mathrm{ms}^{-1}$, $\beta_2 = 0.11\ \mathrm{ms}^{-1}$, $\beta_3 = 177.3\ \mathrm{ms}$, d expresses the mouth-to-ear delay, and e is the packet loss rate. For the G.711 codec, $\gamma_1 = 0$, $\gamma_2 = 30$, and $\gamma_3 = 15$.

The R-Factor is related to the MOS through the following set of expressions:

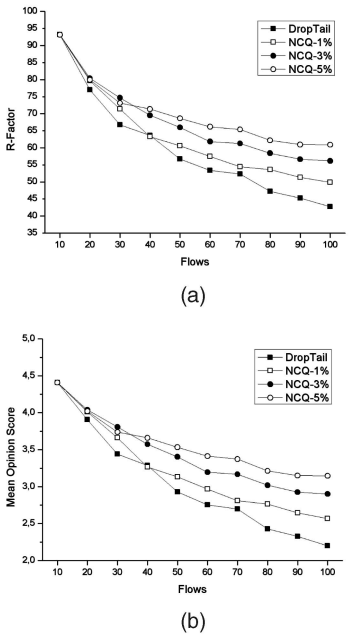


Fig. 17. Impact of NCQ on voice quality. (a) R-factor of VoIP application. (b) MOS of VoIP application.

- **For** $R < 0 : MOS = 1$.
- **For** $R > 100 : MOS = 4.5$.
- **For** $0 < R < 100 : MOS = 1 + 0.035R + 7 \times 10^{-6}(100 - R)$.

For reference, we give the relation of R-Factor to MOS according to Table 2.

In our scenario, we introduce extra delay of 10 ms because of the G.711 voice encoding (according to [8]). We use the simple dumbbell topology in Fig. 7 and adjust the *ncqthresh* to 0.01. We range the number of congestive flows from 10 to 100.

According to Fig. 16a, the VoIP traffic has significant performance gains in terms of *Goodput*. In addition, the NCQ mechanism is not degrading the performance of the congestive FTP flows significantly (Fig. 16b)—especially when *ncqthresh* value is 0.01. Furthermore, fairness among VoIP applications is improved (Fig. 16c).

According to Figs. 17a and 17b, there is a significant improvement in voice quality, while more calls are rated better (Figs. 17a and 17b). For example, in case of 50 calls, all are rated "poor"; NCQ improves their rating to medium (Fig. 17b). Practically, NCQ satisfies more up to 50 users. For more than 80 calls, several calls are rated "poor," exhausting resources unduly.

5.2.4 Scenario 4: Incremental Deployment of NCQ

Here, we evaluate NCQ using a complex topology that incorporates multiple bottlenecks, cross and reverse traffic (Fig. 18). In this scenario, we explore whether NCQ can be incrementally deployed. This analysis is important because the incremental deployment of a new network service or protocol is typically a hard problem, especially when it has to be deployed in the routers [17].

More specifically, we range the number of routers that incorporate the NCQ scheme from 0 to 4, assuming incremental deployment. The *ncqthresh* value is set to 0.03 and we have a random number of applications for both congestive and noncongestive traffic; uniformly random

Authorized licensed use limited to: University of Thrace (Democritus University of Thrace). Downloaded on December 03,2020 at 09:09:21 UTC from IEEE Xplore. Restrictions apply.

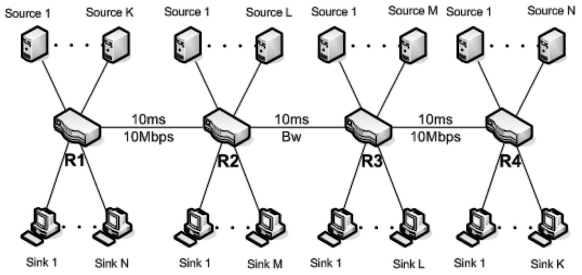
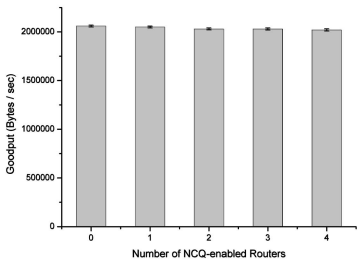
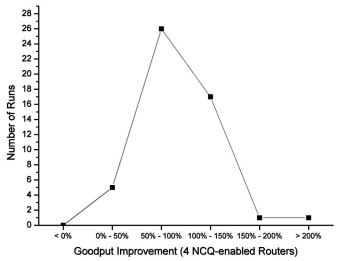


Fig. 18. Complex network topology.



(a)



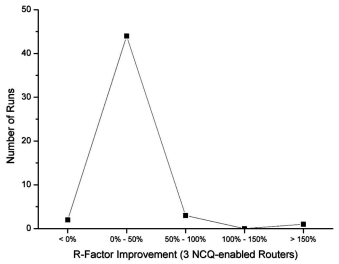
(b)

Fig. 19. Performance results in terms of goodput. (a) Goodput of congestive applications. (b) Goodput improvement of VoIP applications (four NCQ-enabled routers).

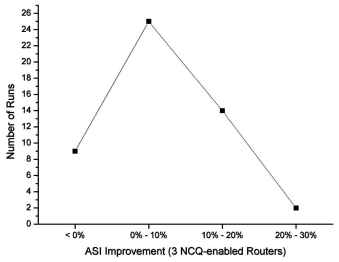
distributed in [1, 50] and [1, 10], respectively. The *bw* value is 10 Mbps. We measure *Goodput*, *ASI*, and *R-Factor*. The queue size is 100 packets and the duration of the experiment is 60 seconds.

We note that this experiment is nondeterministic. Consequently, we performed 50 runs in order to have statistically accurate results. The standard deviation of our measurements is often high due to the nature of our experiment. For example, in the case of randomized contention, one flow receives almost always less throughput than 50 flows. In such results, the standard deviation for the measurements does not decrease by the number of experimental runs. Each time the standard deviation of the differences between the average value and the samples is very small (e.g., 1 to 3 percent), we depict the average values and the confidence intervals (i.e., the above standard deviation). Otherwise, we analyze and interpret our results using frequency tables. More precisely, we group measurements into suitable intervals that allow us to find the most common values. Each figure that uses a frequency table depicts the percentage of performance improvement/decrease of each measure, using as a reference point the experimental results of a scenario using DropTail.

In Fig. 19a, we have only a minor impact on the *Goodput* of the congestive flows for any number of NCQ-enabled routers. However, the *Goodput* of the noncongestive applications is significantly improved (Fig. 19b). When NCQ is fully deployed (four NCQ-enabled routers—Fig. 19b), NCQ achieves more than 100 percent improvement for the 38 percent of the runs.



(a)



(b)

Fig. 20. Voice quality and fairness. (a) R-factor improvement (three NCQ-enabled routers). (b) Improvement in terms of ASI (three NCQ-enabled routers).

The significant improvement in terms of *Goodput* for the VoIP applications is reflected in the voice quality of each call. For example, in the case of 75 percent deployment (three NCQ-enabled routers), in the majority of the runs (i.e., 88 percent), the VoIP flows achieve improvement that reaches 50 percent. For the 4 percent of the runs, the improvement is more than 50 percent. The improvement in terms of *Fairness* is also significant, in the 82 percent of the runs we have an improvement on the *ASI* (Fig. 20b).

More results on the evaluation of NCQ and other similar LIBS-based mechanisms can be found in [25], including complex topologies with different congestive and noncongestive application types.

6 Open Issues

Although we demonstrated NCQ's high potential, we do not presently address all concerns necessary to justify its need for deployment. 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 noncongestive packet exceeds the *ncqthresh*. Probabilistic priority will guarantee fairness among noncongestive flows in a similar fashion to RED's probabilistic dropping [15]. An initial approach is proposed in [25]. Alternatively, the *ncqthresh* may be dynamically adjusted, based on the projected outcome.

In another front of research, ACKs[5] and control packets may benefit from the priority treatment. Control packets prioritization is expected to boost the performance of

5. In this paper, NCQ does not favor ACKs.

Authorized licensed use limited to: University of Thrace (Democritus University of Thrace). Downloaded on December 03,2020 at 09:09:21 UTC from IEEE Xplore. Restrictions apply.

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 under further investigation.

Initially, we assume that noncongestive traffic is formed by small packets at low data rates. Although this approach has significant advantages (requires almost-zero memory state, requires minor implementation effort, and so forth) and is suitable to demonstrate the potential of our approach, it does not cover all the cases of noncongestive traffic (e.g., the noncongestive traffic that is formed by a few long packets). A more sophisticated noncongestive traffic identification algorithm is the subject of a future work.

7 Conclusions

We have shown that NCQ is a simple but powerful tool for service differentiation, particularly beneficial for applications that utilize small rates and short packets, such as typical sensor applications and VoIP. We discovered that a limited prioritization has a dual impact: it benefits noncongestive flows significantly and reduces contention among the congestive flows, resulting in satisfying more users; a leading design issue in today's Internet, which creates hopes for deployment success.

References

[1] F.M. Anjum and L. Tassiulas, "Fair Bandwidth Sharing among Adaptive and Non-Adaptive Flows in the Internet," *Proc. IEEE INFOCOM '99,* Mar. 1999.

[2] D. Bertsekas and R. Gallager, *Data Networks,* second ed. Prentice-Hall, 1991

[3] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, *An Architecture for Differentiated Service,* RFC 2475, 1998.

[4] R. Braden, D. Clark, and S. Shenker, *Integrated Services in the Internet Architecture: An Overview,* RFC 1633, June 1994.

[5] P. Brady, "A Statistical Analysis of On-Off Patterns in 16 Conversations," *The Bell System Technical J.,* vol. 47, pp. 73-91, Jan. 1968.

[6] D. Chiu and R. Jain, "Analysis of the Increase/Decrease Algorithms for Congestion Avoidance in Computer Networks," *J. Computer Networks and ISDN,* vol. 17, no. 1, pp. 1-14, June 1989.

[7] M. Claypool, R. Kinicki, and A. Kumar, "Traffic Sensitive Active Queue Management," *Proc. Eighth IEEE Global Internet Symp.,* Mar. 2005.

[8] R. Cole and J. Rosenbluth, "Voice over IP Performance Monitoring," *ACM SIGCOMM Computer Comm. Rev.,* vol. 31, no. 2, pp. 9-24, Apr. 2001.

[9] T.D. Dang, B. Sonkoly, and S. Molnar, "Fractal Analysis and Modeling of VoIP Traffic," *Proc. 11th Int'l Telecomm. Network Strategy and Planning Symp. (NETWORKS '04),* June 2004.

[10] G. de Veciana, T. Konstantopoulos, and T.-J. Lee, "Stability and Performance Analysis of Networks Supporting Elastic Services," *IEEE/ACM Trans. Networking,* vol. 9, no. 1, pp. 2-14, 2001.

[11] I. Downard, "Simulating Sensor Networks in NS-2," Technical Report NRL/FR/5522-04-10073, Naval Research Laboratory, May 2004.

[12] J. Elson and D. Estrin, "Random, Ephemeral Transaction Identifiers in Dynamic Sensor Networks," *Proc. 21st Int'l Conf. Distributed Computing Systems (ICDCS '01),* pp. 459-468, Apr. 2001.

[13] 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, Univ. of Michigan, Apr. 1999.

[14] S. Floyd and K. Fall, "Promoting the Use of End-to-End Congestion Control in the Internet," *IEEE/ACM Trans. Networking,* vol. 7, no. 4, pp. 458-472, 1999.

[15] S. Floyd and V. Jacobson, "Random Early Detection Gateways for Congestion Avoidance," *IEEE/ACM Trans. Networking,* vol. 1, no. 4, pp. 397-413, Aug. 1993.

[16] S. Fred, T. Bonald, A. Proutiere, G. Régnié, and J. Roberts, "Statistical Bandwidth Sharing: A Study of Congestion at Flow Level," *Proc. ACM SIGCOMM '01,* vol. 31, no. 4, pp. 111-122, 2001.

[17] X. He, C. Papadopoulos, and P. Radoslavov, "A Framework for Incremental Deployment Strategies for Router-Assisted Services," *Proc. IEEE INFOCOM '03,* vol. 2, pp. 1488-1498, 2003.

[18] P. Hurley, J. Boudec, P. Thiran, and M. Kara, "ABE: Providing a Low-Delay Service within Best-Effort," *IEEE Network,* vol. 15, no. 3, pp. 60-69, 2001.

[19] *The E-Model, a Computational Model for Use in Transmission Planning,* ITU-T Recommendation G.107, Dec. 1998.

[20] *General Characteristics of General Telephone Connections and Telephone Circuits—Transmission Impairments,* ITU-T Recommendation G.113, Feb. 1996.

[21] *Pulse Code Modulation (PCM) of Voice Frequencies,* ITU-T Recommendation G.711, Nov. 1988.

[22] A. Kherani and A. Kumar, "Performance Analysis of TCP with Non-Persistent Sessions," *Proc. Workshop Modeling of Flow and Congestion Control,* Sept. 2000.

[23] D. Lin and R. Morris, "Dynamics of Random Early Detection," *Proc. ACM SIGCOMM '97,* pp. 127-137, Sept. 1997.

[24] L. Mamatas and V. Tsaoussidis, "Protocol Behavior and Energy-Saving Potential," *Proc. Sixth Int'l Workshop Wireless Local Networks (WLN '06),* Nov. 2006.

[25] L. Mamatas and V. Tsaoussidis, "Less Impact Better Service (LIBS): A Service Paradigm for Internet Telephony," Technical Report TR-DUTH-EE-2007-16, Democritus Univ. of Thrace, Nov. 2007.

[26] R. Mahajan, S. Floyd, and D. Wetherall, "Controlling High-Bandwidth Flows at the Congested Router," *Proc. Ninth Int'l Conf. Network Protocols (ICNP '01),* Nov. 2001.

[27] L. Mamatas and V. Tsaoussidis, "A New Approach to Service Differentiation: Non-Congestive Queuing," *Proc. ICST First Int'l Workshop Convergence of Heterogeneous Wireless Networks (CONWIN '06),* July 2006.

[28] L. Mamatas and V. Tsaoussidis, "Differentiating Services for Sensor Internetworking," *Proc. IFIP Fifth Ann. Mediterranean Ad Hoc Networking Workshop (Med-Hoc-Net '07),* June 2007.

[29] M. Marsan and M. Gerla, "Fairness in Local Computing Networks," *Proc. IEEE Int'l Conf. Comm. (ICC '02),* June 1982.

[30] L. Massoulié and J. Roberts, "Bandwidth Sharing and Admission Control for Elastic Traffic," *Telecomm. Systems,* vol. 15, no. 1, pp. 185-201, 2000.

[31] S. McCanne and S. Floyd NS-2 Network Simulator, http://www.isi.edu/nsnam/ns/, 1997.

[32] T.J. Ott, T.V. Lakshman, and L.H. Wong, "SRED: Stabilized RED," *Proc. IEEE INFOCOM '99,* vol. 3, pp. 1346-1355, 1999.

[33] R. Pan, L. Breslau, B. Prabhakar, and S. Shenker, "Approximate Fairness through Differential Dropping," *ACM SIGCOMM Computer Comm. Rev.,* vol. 33, no. 2, pp. 23-39, Apr. 2003.

[34] R. Pan, B. Prabhakar, and K. Psounis, "CHOKe—A Stateless Queue Management Scheme for Approximating Fair Bandwidth Allocation," *Proc. IEEE INFOCOM '00,* Mar. 2000.

[35] C. Perkins, E. Belding-Royer, and S. Das, "RFC3561: Ad Hoc On-Demand Distance Vector (AODV) Routing," *Internet RFCs,* 2003.

[36] A. Rangarajan, "Early Regulation of Unresponsive Flows," Technical Report TRCS99-26, Univ. of California at Santa Barbara, July 1999.

[37] H. Sabbineni and K. Chakrabarty, "Location-Aided Flooding: An Energy-Efficient Data Dissemination Protocol for Wireless Sensor Networks," *IEEE Trans. Computers,* vol. 54, no. 1, pp. 36-46, Jan. 2005.

[38] I. Stoica, S. Shenker, and H. Zhang, "Core-Stateless Fair Queueing: Achieving Approximately Fair Bandwidth Allocations in High Speed Networks," *Proc. ACM SIGCOMM '98,* pp. 118-130, 1998.

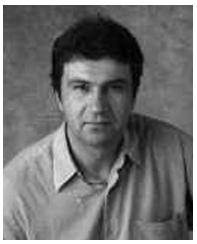
[39] D. Vardalis and V. Tsaoussidis, "On the Efficiency and Fairness of Congestion Control Mechanisms in Wired and Wireless Networks," *The J. Supercomputing,* 2002.

Authorized licensed use limited to: University of Thrace (Democritus University of Thrace). Downloaded on December 03,2020 at 09:09:21 UTC from IEEE Xplore. Restrictions apply.

[40] Y. Zhang, W. Liu, and Y. Fang, "Secure Localization in Wireless Sensor Networks," *Proc. IEEE Military Comm. Conf. (MILCOM '05)*, vol. 5, pp. 3169-3175, Oct. 2005.



Lefteris Mamatas received the PhD degree from Democritus University of Thrace in 2008. He is a research associate in the Department of Electronic and Electrical Engineering, University College London. His research interests are focused on network protocols, service differentiation, and network management. He published 12 conference and seven journal papers. He is a member of the IEEE.



Vassilis Tsaoussidis received the BSc degree in applied mathematics from Aristotle University, Thessaloniki, Greece, the diploma in statistics and computer science from the Hellenic Institute of Statistics, and the PhD degree in computer networks from Humboldt University, Berlin, in 1995. He held faculty positions at Rutgers University, New Brunswick; State University of New York, Stony Brook; and Northeastern University, Boston. In May 2003, he joined the Department of Electrical and Computer Engineering, Democritus University of Thrace, Xanthi, Greece. His research interests lie in the area of transport/network protocols, i.e., their design aspects and performance evaluation. He is an editor in chief for the *Journal of Internet Engineering* and an editor for the *IEEE Transactions in Mobile Computing*, the *Journal of Computer Networks* (Elsevier), *Wireless Communications and Mobile Computing*, and the *Journal of Mobile Multimedia*. He participates in several technical program committees in his area of expertise, such as INFOCOM, International IFIP-TC6 Networking Conference (NETWORKING), IEEE Global Telecommunications Conference (GLOBECOM), IEEE International Conference on Computer Communication and Networks (ICCCN), IEEE Symposium on Computers and Communications (ISCC), IEEE Workshop on Energy-Efficient Wireless Computing and Networking (EWCN), International Workshop on Wireless Local Networks (WLN), and several others. He is a senior member of the IEEE.

▷ **For more information on this or any other computing topic, please visit our Digital Library at** www.computer.org/publications/dlib.

Authorized licensed use limited to: University of Thrace (Democritus University of Thrace). Downloaded on December 03,2020 at 09:09:21 UTC from IEEE Xplore. Restrictions apply.