Admission Control for DiffServ Based Quality of Service in Cut-through Networks

Sven-Arne Reinemo, Frank Olaf Sem-Jacobsen, Tor Skeie, and Olav Lysne

Simula Research Laboratory
P.O.Box 134, N-1325 Lysaker, Norway
svenar@simula.no, frankose@ifi.uio.no, {tskeie, olavly}@simula.no

Abstract. Previous work on Quality of Service in Cut-through networks shows that resource reservation mechanisms are only effective below the saturation point. Admission control in these networks will therefore need to keep network utilization below the saturation point, while still utilising the network resources to the maximum extent possible. In this paper we propose and evaluate three admission control schemes. Two of these use a centralised bandwidth broker, while the third is a distributed measurement based approach. We combine these admission control schemes with a DiffServ based QoS scheme for virtual cut-through networks to achieve class based bandwidth and latency guarantees. Our simulations show that the measurement based approach favoured in the Internet communities performs poorly in cut-trough networks. Furthermore it demonstrates that detailed knowledge on link utilization improves performance significantly.

1 Introduction

Internet has today evolved into a global infrastructure supporting applications such as streaming media, E-commerce, network storage, etc. Each of these applications must handle an ever increasing volume of data demanded as predictable transfers. In this context the provision of Quality of Service (QoS) is becoming an important issue. In order to keep pace with computer evolution and the increased burden imposed on data servers, application processing, etc. created by the popularity of the Internet, we have in recent years seen several new technologies proposed for System and Local Area Networking (SAN/LAN) [3,4,7,14,24]. Common for this body of technologies is that they rely on point-to-point links interconnected by off-the-shelf switches that support some kind of back-pressure mechanism. Besides, most of the referred technologies also adhere to the cut-through or wormhole switching principles - only Gigabit Ethernet is using the store-and-forward technique. For a survey of some relevant networking principles we refer to [6].

IETF has for several years provided the Internet community with QoS concepts and mechanisms. The best known ones are Integrated Services (IntServ) [8], Resource Reservation Protocol (RSVP) [13], and Differentiated Services (DiffServ) [5]. IntServ together with RSVP define a concept based on per flow reservations (signalling) and admission control to be present end-to-end. DiffServ, however, takes another approach assuming no explicit reservation mechanism in the interior network elements. QoS is here realized by giving data packets differentiated treatment relative to the QoS header code information. In the underlying network technologies QoS has to a less extent been emphasised - the key metrics here have mainly been mean throughput and latency. To provide QoS end-to-end, possibly over heterogeneous technologies this means that the lower layers should also have support for predictable transfer including the ability to interoperate with a higher level IETF concept. This issue is being challenged by emerging SAN/LAN standards, such as InfiniBand^{TM} [4] and Gigabit Ethernet [24] providing various QoS mechanisms.

Recently we have also seen several research contributions to this field. Jaspernite et. al. [9,10] and Skeie et. al [15] discuss different aspects of taking control of the latency through switched Ethernet relative to the IEEE 802.1p standard aiming for traffic priorities. Another body of work is tailored to the InfiniBand^{TM} architecture (IBA) [1,2,12]. In [12] Pelissier gives an introduction to the set of QoS mechanisms offered by IBA and the support for DiffServ over IBA. In this approach the presence of admission control is assumed. Alfaro et. al build on this scheme and present a strategy for computing the arbitration tables of IBA networks, moreover a methodology for weighting of virtual layers referring to the dual arbitrator defined by IBA [2]. The concept is evaluated through simulations assuming that only bandwidth sensitive traffic requests QoS. In [1] Alfaro et. al also include time sensitive traffic, besides calculating the worst case latency through various types of switching architectures.

DiffServ is foreseen to be the most prominent concept for providing QoS in the future Internet [11,17]. DiffServ makes a distinction between boundary nodes and core nodes with respect to support of QoS features. Following the DiffServ philosophy no core switch should hold status information about passing-through traffic. Neither should there be any explicit signalling on a per flow basis to these components. This means that within the DiffServ framework any admission control or policing functionality would have to be implemented by boundary nodes or handled by a dedicated bandwidth broker. The core switches are assumed to perform traffic discrimination only based on service class, which is decided by a QoS tag included in the packet header - all packets carrying the same QoS tag will get equal treatment. From that viewpoint DiffServ is apparently a relative service model having difficulties giving absolute guarantees.

None of the previous debated contributions comply with the DiffServ model. In [12] Pelissier, however, discusses interoperation between DiffServ and IBA on a traffic class and service level basis, but refer to RSVP with respect to admission control. The strategy proposed by Alfaro et. al has to recompute the IBA dual arbitrator every time that a new connection is honoured [1,2]. And such a scheme is not associable with DiffServ. Neither is the admission control scheme presented in [29] by Yum et. al, which use hop by hop bandwidth reservations and requires recomputations of the weighted round robin scheduler at every hop towards the destination. In [25] Reinemo et. al. studied the provision of QoS in cut-through networks by adhering to the DiffServ model. The problem was approached without any explicit admission control mechanism, as a pure relative model. Empirically they examined the sensitivity of different QoS properties under various load and traffic mixture conditions, hereunder assessing the effect of back-pressure.

In this paper we endeavour to achieve class based QoS in cut-through networks by use of admission control. More specifically, we extend the concept described in [25] with admission control. However, still in compliance with the Diff-Serv paradigm where service classes, as aggregated flows, are the target for QoS. Three different admission control mechanisms are proposed and carefully evaluated through extensive simulations. Two of the schemes assume pre-knowledge of the network's performance behaviour without admission control, and are furthermore implemented as a centralised bandwidth broker. The third scheme is based on endpoint/egress admission control and relies on measurements to assess the load situation, inspired by Internet QoS provisioning. To the best of the authors' knowledge no detailed admission control methods have been proposed for cut-through networks before.

The rest of this paper is organised as follows. In section 2 we give a description of our QoS architecture and routing algorithm, in section 3 our three admission control mechanisms are described, and in section 4 our simulation scenario is described. In section 5 we discuss our performance results, and finally in section 6 we give some concluding remarks.

2 QoS Architecture

The architecture used in our simulations is inspired by IBA link layer technology [4] and is a flit based virtual cut-through (VCT) switch. The overall design is based on the canonical router architecture described in [6].

In VCT the routing decision is made as soon as the header of the packet is received and if the necessary resources are available the rest of the packet is forwarded directly to the destination link [23]. If the necessary resources are busy the packet is buffered in the switch. In addition we use flow control on all links so all data is organised as flow control digits (flits) at the lowest level.

The switch core consists of a crossbar where each link and VL has dedicated access to the crossbar. Each link supports one or more virtual lanes (VL), where each VL has its own buffer resources which consist of an input buffer large enough to hold a packet and an output buffer large enough to hold two flits to increase performance. Output link arbitration is done in a round robin fashion.

To achieve QoS our switch architecture support QoS mechanisms like the ones found in the IBA architecture. IBA supports three mechanisms for QoS which are mapping of service level (SL) to VL, weighting of VLs and prioritising VLs as either low priority (LP) or high priority (HP). A more detailed description of these QoS aspects can be found in [25].

The routing used is a newly introduced routing algorithm called *Layered shortest path routing* (LASH) [16]. LASH is a minimal deterministic routing algorithm for irregular networks which only relies on the support of virtual layers.

There is no need for any other functionality in the switch, so LASH fits well with our simple approach to QoS. An in-depth description of LASH is found in [16].

3 Admission Control

In this section we propose three different admission control (AC) mechanisms that we carefully evaluate in section 5.

3.1 Calibrated Load Based Admission Control

The *Calibrated Load* (CL) approach is a simple scheme relying on the fact that a BB knows the total rate of traffic entering the network. Our AC parameter is the amount of traffic which can be injected into the network while still keeping the load below saturation. As the rate of traffic entering the network reaches the CL parameter no more traffic will be admitted. In most cases the CL parameter must be decided by measurements on the network in question to find the saturation point - our CL is deduced from measurements performed in [25].

To keep HP and LP traffic separated we use two different CL parameters, one for the total HP traffic and one for the total LP traffic. For HP traffic this can be expressed as follows

$$\sum_{i=0}^{n} L_{HP,i} + P_{HP} < CL_{HP} \ . \qquad (1)$$

Here CL_{HP} is the calibrated load for outgoing HP traffic, $L_{HP,n}$ is the HP load in node n and P_{HP} is the peak rate for the requesting flow. Moreover, the flow is admitted if the total HP load $\sum_{i=0}^{n} L_{HP,i}$ plus the requested increase P_{HP} is below the calibrated load CL_{HP}. LP traffic can be expressed similarly just substituting HP values with LP values. The strength of this scheme is that it is simple, its weakness is that it is inaccurate since it does not take into account the distribution of flows in the network. And from that viewpoint it is less suitable for handling hot spots.

3.2 Link by Link Based Admission Control

Our second scheme is the Link-by-Link (LL) approach. Here the BB knows the load on every link in the network and will consult the availability of bandwidth on every link between source and destination before accepting or rejecting a flow. Compared to the CL approach, this solution assumes both topological and routing information about the network.

For the AC decision we adopt the *simple sum* approach as presented in [28]. This algorithm states that a flow may be admitted if its peak rate p plus the peak rate of the already admitted flows s is less than the link bandwidth bw. Thus the requested flow will be admitted if the following inequality is true [28]

$$p + s < bw \qquad (2)$$

we view p as the increase in peak rate for the flow and s as the sum of the admitted peak rates. As for the *CL* method we deduce the effective bandwidth from the measurements obtained in [25]. Since we are dealing with service levels where each SL have different bandwidth requirements it is natural to introduce some sort of differentiation into equation (2). We achieve this by dividing the link bandwidth into portions relative to the traffic load of the SLs, and include only the bandwidth available to a specific service level bw_{SL} in the equation as follows

$$p + s_{SL} < bw_{SL} \tag{3}$$

where

$$bw_{SL} = bw_{link} * \frac{load_{SL}}{load_{total}} \tag{4}$$

and s_{SL} is the sum of the admitted peak rates for service level SL and bw_{link} is the effective link bandwidth.

3.3 Egress Based Admission Control

Our third scheme is the Egress Based (EB) approach. This is a fully distributed AC scheme where the egress nodes are responsible for conducting the provisions. Basically, we here adopt the Internet AC concept presented by Cetikaya and Knightly in [26]. This method does not assume any pre-knowledge of the network behaviour as was the case with our previous solutions. Also different from the previous approaches is the use of a delay bound as the primary AC parameter. For clarity we give a brief outline of the algorithm below, a more detailed description can be found in [27].

The method is entirely measurement based and relies on that the sending nodes timestamp all packets enabling the egress nodes to calculate two types of measurements. First, by dividing time into timeslots of length τ and counting the number of arriving packets, the egress nodes can deduce the arrival rate of packets in a specific timeslot. By computing the maximum arrival rate for increasingly longer time intervals we get a peak rate arrival envelope $R(t)$, where $t = 0, ..., T$ timeslots, as described in [26]. Second, by comparing the originating timestamp relative to the arrival time, the egress node can calculate the transfer time of a packet. Having this information available the egress node can furthermore derive the time needed by the infrastructure to service k following packets; i.e. a consecutive stream of packets where the next packet in the service class enters the infrastructure before the previous packet has departed the egress node. By doing this for larger and larger k sequences of packets within a measuring interval of length $T\tau$ and subsequently inverting this function we achieve the service envelope $S(t)$, giving the amount of packets processed by the network in a given time interval t. Now repeating this M times, the mean $\overline{R}(t)$ and the variance $\sigma^2(t)$ of $R(t)$, and the mean $\overline{S}(t)$ and variance $\Psi^2(t)$ of $S(t)$ may be calculated. If a flow request has a peak rate P and a delay bound D it may be accepted if the peak rate P plus the measured arrival rate $R(t)$ is less than the service rate allowing for the delay $D, S(t + D)$.

The EB scheme derives its knowledge of the network from measurements of the traffic passing through the egress nodes. It is therefore difficult for the egress nodes to have a complete picture of the load in the network, moreover the packet latency is used to infer the network load utilising the fact that an increased network load will cause increase in latency as well. From that viewpoint it seems difficult to give a service class bandwidth guarantees since it has no concrete knowledge of the network load. The algorithm will admit as much traffic as it can without breaking the delay bound. The key instrument of the scheme is the given delay bound for the different flows, and the efficiency of the algorithm is linked to its ability to limit the service levels to operate within the delay bounds.

3.4 Target for Admission Control

The main findings for the work in [25] are that (i) throughput differentiation can be achieved by weighting of VLs and by classifying the VLs as either low or high priority, (ii) the balance between VL weighting and VL load is not crucial when the network is operating below the saturation level. In general this sets the target for the AC, since as long as we can ensure that the load of the various service classes is below saturation level we can also guarantee that each of these classes get the bandwidth they request. The target for admission control is thus the point where the amount of accepted traffic is starting to become less than traffic offered. The effective bandwidth at this point will be used as a steering vehicle by the CL and LL methods.

Another main finding in [25] is that though the latency characteristics below saturation were fairly good, significant jitter was observed. This problem we challenge by proposing the EB method, where a given delay bound is the requested quality of service. Since this concept is continuously monitoring the end-to-end latency characteristics for all pair of nodes one should possibly expect that delay guarantees could be given.

4 Simulation Scenario

For all simulations we have used a flit level simulator developed in house at Simula Research Laboratory. In the simulation results that follow, all traffic is modelled by a normal approximation of the Poisson distribution. We have performed simulations on a network with 32 switches, where each switch is connected to 5 end nodes and the maximum number of links per switch is 10 in addition to the end nodes. We have randomly generated 16 irregular topologies and we have run measurements on these topologies at increasing load. We use LASH [16] as routing algorithm and random pairs as traffic pattern. In the random pairs scheme each source sends only to one destination and no destination receives from more that one source. The link speed is one flit per cycle, the flit size is one byte and the packet size is 32 bytes for all packets.

The five different end nodes send traffic on one of five different service levels. One service level for each node (Table 1), SL 1 and 2 are considered to be of the expedited forwarding (EF) class in DiffServ terminology. And SL 3 and 4 are considered to be of the assured forwarding (AF) class. SL 5 is considered as best effort (BE) traffic and from that viewpoint is not a subject of AC.

For the CL and LL schemes all simulations where run with an ACT deduced from our measurements in [25]. In the first part of the simulation the send rate is steadily increased by adding more and more flows until admission is denied by the AC scheme. When this happens the current rate is not changed, but the node will continue to try to go beyond the ACT for a fixed number of times before it gives up. For the EB scheme the send rate is increased in the same way, but the AC decision is primarily based on measured latency as described in section 3.3.

Table 1. Services levels

SL	DS[1]	Load %	BW[2]	Pri	SL	DS	Load %	BW	Pri	SL	DS	Load %	BW	Pri
1	EF	10	4	high	3	AF	20	8	low	5	AF	30	1	low
2	EF	15	6	high	4	AF	25	10	low					

5 Performance Results

5.1 Throughput

Recall that the target for the AC is to make sure that the network operates below saturation at all times, since below this point we can guarantee that all SLs get the bandwidth they request. The relative requests for each SL are as shown in Table 1. Figure 1(a) shows what happens in a network without AC when it enters saturation. We are no longer able to give all service classes the bandwidth they request and the HP classes preempts LP bandwidth, i.e. the bandwidth differentiation is no longer according to the percentages in Table 1. In the CL scheme, we see from figure 1(b) that we are successful in keeping the accepted load below the saturation point, even as the offered load goes beyond this point. The bandwidth differentiation does not fail as is the case in figure 1(a), but it suffers slightly as we reaches *high* load. As the load increases the differentiation between SLs in the same class is diminished. Thus, the CL scheme is able to keep the load below saturation. However, it appears that it is too coarse to achieve good bandwidth differentiation between SLs of the same class since it makes its AC decisions based on the total load for a class and not for each SL.

Moving on to the LL scheme (figure 1(c)) we see several improvements compared to the CL scheme. First, we get a sharper bandwitdh cut-off with much

---

[1] The DiffServ equivalent service class.

[2] The maximum number of flits allow to transmit when scheduled.

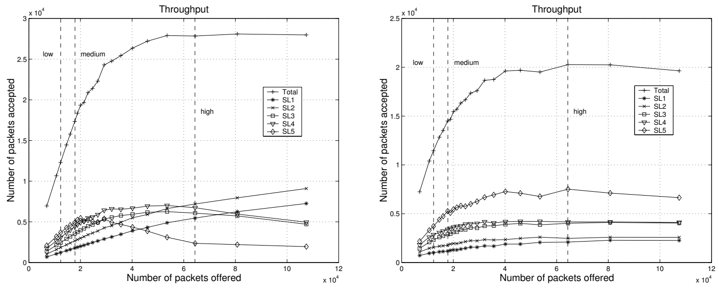

(a) No admission control (b) Calibrated load

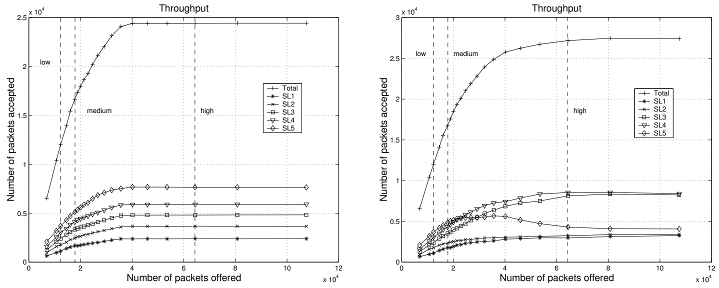

(c) Link by link (d) Egress measurements

Fig. 1. Throughput

less hesitation than for CL. Second, we achieve a differentiation relative to the requests, meaning that we can give bandwidth guarantees. Third, we are able to utilise the network resource better as we get closer to the saturation point. This improvement is probably due to the fact that the LL scheme knows the load of every link in the network and is able to make the AC decision based on the load along the actual source/destination path.

Finally, we have the EB method. It is apparent from figure 1(d) that this method is unable to give bandwidth guarantees, as well as increasing the load beyond the saturation point and admitting too much traffic. Now as the load increases beyond saturation the best effort traffic (SL 5) is reduced as it must make way for traffic on the other SLs. The issue here is that delay is the most significant AC parameter in this scheme and bandwidth requirements have more or less been ignored.

5.2 Latency

Let us now turn our attention to the latency results. Figure 2(a) shows the average latency for increasing load values without admission control. Comparing it with the CL results in figure 2(b) shows that the CL scheme is quite close when we look at the same load values. The average latency for all packets is 436 for CL at the high mark which is a 6% increase compared to the scheme without AC at a corresponding load. A problem with the CL scheme is that the latency

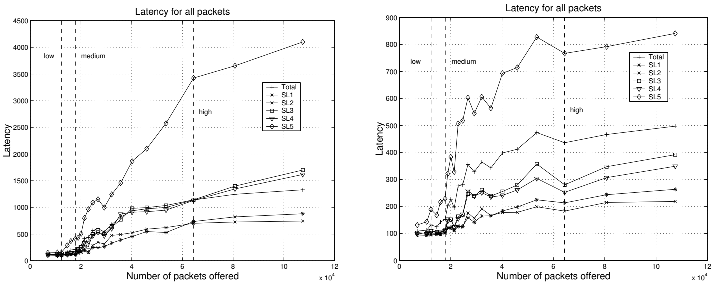

(a) No admission control (b) Calibrated load

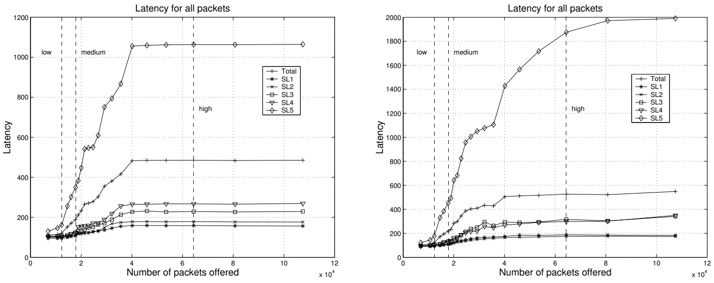

(c) Link by link (d) Egress measurements

Fig. 2. Average latency

values are unstable as the load increases since the estimate of current throughput is to coarse to target the exact rejection point. The LL scheme overcomes this problem as show in figure 2(c). As soon as the LL scheme starts rejecting flows the latency stabilises. The latency values for high load is 485 which is slightly above the CL latency at this particular point. The LL method also gives a more linear increase in latency as we approach the rejection point for new flows. The LL scheme is the better of the two as it gives lower latency to SL 1-4 and higher latency to the best effort traffic in SL 5. Even if it has a higher average latency for all packets (compared to CL) it performs better since the increase in the average is caused by the best effort traffic in SL 5.

The EB scheme uses measured latency as its primary AC parameter. The results are presented in figure 2(d). This scheme produces average results which fall between the CL and LL scheme. Furthermore, it is capable of giving the same latency to SLs fairly independently of weight such as SL 1 and 2, but it is unable to satisfy the delay bound of 100 for SL 1 and 2, and 250 for SL 3 and 4. The achieved latency at the high mark is 187 for SL 1 and 316 SL 3. So even if using a measurement based method we are unable to give hard delay guarantees. It seems very difficult to give delay bounds in combination with good throughput in cut-through networks. To remedy this problem one possible could reduce the throughput by hardening the AC, or we could turn to other means such as modifying the flow control to better handle delay bound traffic.

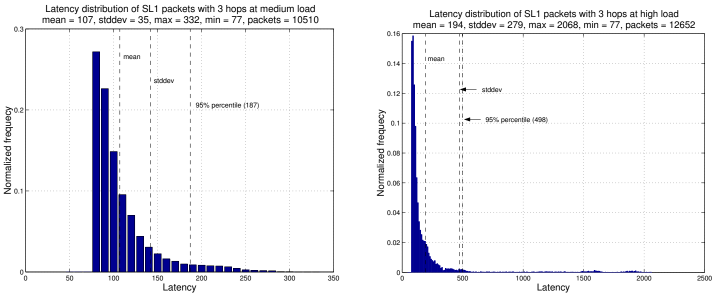

(a) No admission control

(b) Calibrated load

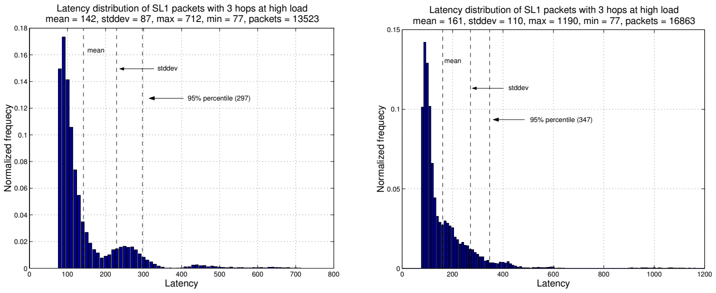

(c) Link by link

(d) Egress measurements

Fig. 3. Latency distribution for packets with 3 hops

5.3 Jitter

Finally lets turn our attention to jitter. Figure 3 shows the latency distributions for packets with a path length of 3 hops. This was the most frequently occurring path length in our simulations. The mean, standard deviation and 95 % percentile are marked with a dashed line in the figures. The distance between the mean mark and the standard deviation mark reflects the standard deviation.

Figure 3(a) shows the latency distribution achieved without AC at the load marked as medium in figure 1(a). Figure 3(b) shows the distribution for the CL scheme at the load marked as high in figure 1(b). Note that the load at this point is about 20% above that without admission control. We see that the CL scheme has quite poor jitter characteristics, which is reasonable if we recall the latency curve we saw in figure 2(b). The mean is 194, the standard deviation 279 and the 95% percentile is 498. Thus 95% of the packets have a latency of 498 or below. Even if the mean value is not too bad, the large standard deviation and the 95% percentile shows that jitter is clearly very high. The LL scheme has better jitter characteristics. From figure 3(c) we see that the histogram has a shorter tail compared to figure 3(b). We have a mean of 142, standard deviation of 87 and a 95% percentile of 297. Which reduces the jitter potential and gives us almost a 40% reduction of packet latency for 95% of the packets. In addition the load for LL at this point is about 7% above the CL load. Thus, better results are achieved at a higher load. For EB scheme in figure 3(d) we see that it is unable to improve on the results from the LL scheme. With a 95% percentile of 357 it has a 30% improvement over the CL. Note that this is achieved at a load 25% above the CL load. Still, even a measurements based delay bound scheme is unable to give good jitter characteristics in cut-through networks.

6 Conclusion

In this paper we propose and evaluate three different admission control schemes for virtual cut-through networks. Each one suitable for use in combination with a DiffServ based QoS scheme to deliver soft real-time guarantees. Two of the schemes assume pre-knowledge of the network's performance behaviour without admission control, and are both implemented with bandwidth broker. The third method is based on endpoint/egress admission control and relies on measurements to assess the load situation.

Our main findings are as follows. First, bandwidth guarantees for aggregated flows are achievable with the use of the Link-by-Link scheme. While the Calibrated Load and Egress Based methods are unable to achieve such good guarantees. Second, latency and jitter properties are hard to achieve regardless of the method used. This is due to the nature of cut-networks and the way flow control affects latency. Strict admission control can be used to improve latency, but at the cost of lower throughput. To achieve a combination of high throughput and low latency modifications to the flow control may be considered.

References

1. F. J. Alfaro, J. L. Sanchez, J. Duato, and C. R. Das. A strategy to compute the InfiniBand arbitration tables. In *Proceedings of International Parallel and Distributed Processing Symposium*, April 2002.
2. F. J. Alfaro, J. L. Sanchez, and J. Duato. A strategy to manage time sensitive traffic in InfiniBand. In *Proceedings of Workshop on Communication Architecture for Clusters (CAC)*, April 2002.
3. N. J. Boden, D. Cohen, R. E. Felderman, A. E. Kulawik, C. L. Seitz, J. N. Seizovic, and W. K. Su. Myrinet – a gigabit-per-second LAN. IEEE MICRO, 1995.
4. InfiniBand Trade Association. Infiniband architecture specification.
5. Differentiated Services. RFC 2475.
6. J. Duato, S. Yalamanchili, and L. Ni, *Interconnection Networks an engineering approach*, IEEE Computer Society, 1997.
7. R. W. Horst. Tnet: A reliable SAN. *IEEE Micro*, 15(1):37–45, 1995.
8. Integrated Services. RFC 1633.
9. J. Jaspernite, and P. Neumann. Switched Ethernet for Factory Communication. In *Proceedings of* 8^{th} *IEEE International Conference on Emerging Technologies and Factory Automation (ETFA'01)*, 205–212, October 2001.
10. J. Jaspernite, P. Neumann, M. Theiss, and K. Watson. Deterministic real-time communication with switched Ethernet. In *Proceedings of* 4^{th} *IEEE International Workshop on Factory Communication Systems (WFCS'02)*, August 2002.
11. K. Kilkki. Differentiated services for the Internet. *Macmillian Tech. Publishing*, 1999.

12. J. Pelissier. Providing quality of service over InfiniBand^{TM} architecture fabrics. In *Proceedings of Hot Interconnects X*, 2000.
13. ReSource ReserVation Protocol. RFC 2205.
14. M. D. Schroder et.al., "Autonet: a high-speed, self-configuring local area network using point-to-point links," SRC Research Report 59, Digital Equipment Corporation, 1990.
15. T. Skeie, J. Johannessen, and Ø. Holmeide. The road to an end-to-end deterministic Ethernet. In *Proceedings of* 4th *IEEE International Workshop on Factory Communication Systems (WFCS'02)*, August 2002.
16. T. Skeie, O. Lysne, and I. Theiss. Layered shortest path (LASH) routing in irregular system area networks. In *Proceedings of Communication Architecture for Clusters*, 2002.
17. X. Xiao and L. M. Ni. Internet QoS: A Big Picture In *IEEE Network Magazine*, 8–19, March/April 1999.
18. J. S. Yang and C. T. King, "Turn-restricted adaptive routing in irregular wormhole-routed networks," in *Proceedings of the 11'th International Symposium on High Performance Computing (HPCS97)*, July 1997.
19. A. A. Chien and J. H. Kim, "Approaches to Quality of Service in High-Performance Networks," in *Lecture Notes in Computer Science*, vol. 1417, 1998.
20. J. Duato and S. Yalamanchili and B. Caminero and D. S. Love and F. J. Quiles, "MMR: A High-Performance Multimedia Router - Architecture and Design Trade-Offs," in *HPCA*, pages 300-309, 1999.
21. B. Caminero, C. Carrion, F. J. Quiles, J. Duato and S. Yalamanchili, " A Solution for Handling Hybrid Traffic in Clustered Environments: The MultiMedia Router MMR," in *Proceedings of IPDPS-03*, April 2003.
22. M. Gerla and B. Kannan and B. Kwan and E. Leonardi and F. Neri and P. Palnati and S. Walton, "Quality of Service Support in High-Speed, Wormhole Routing Networks," in *International Conference on Network Protocols (ICNP'96)*, 1996.
23. P. Kermani and L. Kleinrock, "Virtual Cut-through: A New Computer Communication Switching Technique," in *Computer Networks*, no. 4, vol. 3, 1979.
24. R. Seifert, *Gigabit Ethernet,* Addison Wesley Pub Co., 1998.
25. S. A. Reinemo and T. Skeie and O. Lysne, "Applying the DiffServ Model in Cut-through Networks," in *Proceedings of the 2003 International Conference on Parallel and Distributed Processing Techniques and Applications*, 2003.
26. C. Cetikaya and E. W. Knightly, "Egress admission control," in *INFOCOM (3)*, pages 1471-1480, 2000.
27. J. Schlembach and A. Skoe and P. Yuan and E. Knightly, "Design and Implementation of Scalable Admission Control," in *QoS-IP*, pages 1-15, 2001.
28. S. Jamin and S. J. Shenker and P. B. Danzig, "Comparison of Measurement-Based Admission Control Algorithms for Controlled-Load Service," in *INFOCOM (3)*, pages 973-980, 1997.
29. K. H. Yum, E. J. Kim, C. R. Das, M. Yousif, J. Duato, "Integrated Admission and Congestion Control for QoS Support in Clusters," in *Proceedings of IEEE International Conference on Cluster Computing*, pages 325-332, September 2002.