Chapter 6

Switch Fabric Performance

Three classes of switch fabric have been proposed, each with a range of possible implementation parameters, and these require investigation in order to select a preferred switch design. This switch design will then be investigated in greater detail in the following chapter.

The simplest way to quantify the performance of a particular switch implementation is to measure the normalised average throughput of the switch when saturated with traffic with a uniform random distribution of packet destinations. This is called the throughput at saturation, or sometimes the maximum throughput, of the switch and it gives a useful measure of the capacity of the switch by which different switch designs may be compared. Another useful measure is the mean packet delay, from entry to exit of the switch, for identical traffic sources on each of the switch ports. A simulation model has been developed to compare the throughput at saturation and mean delay performance of the different switch fabrics according to the various implementation parameters summarised in table 6.1.

6.1 Traffic Models

Two traffic models were developed for the comparison of the various switch designs: a saturation model for the evaluation of throughput at saturation and a slotted traffic source for the delay comparison. In the saturation model every input port of the switch is saturated with incoming traffic so that a new packet is always available at every input port on completion of transmission of the packet in service. For switches that do not employ input queue by-pass it was not necessary to model the input queues but merely to supply each input port with a fresh packet whenever it completed a packet transmission. For switches with input queue by-pass, an input queue was modelled on every switch port which was always full. For switches up to size 512×512 the queue was 100 packets long while for larger switches a queue size of 10 packets was considered sufficient. All packets were of the same length and followed a uniform random destination distribution while all output ports acted as a perfect sink.

Parameter	*Range*
Switch Fabric Size	2×2 to 4096×4096
Interconnection Networks	Crossbar Delta Beneš
Degree of Switching Element	2, 4, 8, 16
Multiple Path Algorithms	Searching Flooding Random
Multiple Switch Planes	1 to 4
Port Controllers	Regular Input Queue By-Pass Double Buffered Output De-Luxe

Table 6.1: Switch fabric design parameters.

While the throughput at saturation gives information about the maximum capacity of the switch design, no input buffered switch can effectively be operated at this capacity as the input queues would become permanently full leading to very high delay and high packet loss. To compare the mean delay through different switch structures, traffic sources which produce a load below that of the throughput at saturation must be used on each input port. One such traffic source that has been analysed in the literature is the slotted traffic source, or more correctly referred to as the Bernouilli arrival process [76]. In this model each input port receives a contiguous stream of timeslots, each timeslot being the length of a single packet. Each timeslot may be either empty or filled with a single packet and the load offered by the traffic source is the uniform random probability of any timeslot containing a packet. (A slotted source with a load in excess of the throughput at saturation of the switch fabric becomes equivalent to a saturated source.) All packets are given a random destination distribution and all sources are set to the same value of applied load for each measurement of mean delay.

6.2 The Simulation Model

In order to reduce the amount of computer time required by the simulation model to acceptable proportions one major simplification was made. Each packet set-up attempt was modelled as an instantaneous event. Thus in the model each packet set-up attempt is taken in turn and a complete path traced through the interconnection network. If all nodes in the path are originally free then the set-up attempt is deemed successful and all nodes in the path marked busy for the duration of the transmission of the packet. For flooding algorithms all of the multiple paths to the destination are investigated until a free path is located. If no free path is available the set-up attempt is considered blocked and a new set-up attempt scheduled to occur after the retry delay.

In reality a packet will set up on a stage by stage basis, thus a packet which fails set-up could itself cause blocking during its unsuccessful set-up attempt. The effect of this simplification is to over-estimate the throughput at saturation. A more detailed simulation model which does consider the set-up of packets on a stage by stage basis has also been investigated to evaluate the effect of this simplification. It will be shown that for two-plane delta networks constructed from switching elements of degree 8 or 16 the error introduced by this simplification in the evaluation of the throughput at saturation is in general less than about 2%.

Other simplifications include the modelling of the release of the path on completion of packet transmission as instantaneous. In a real implementation there may be a fixed delay or a stage by stage release mechanism. Also it is assumed in the saturation model that a new packet set-up attempt immediately follows transmission of the previous packet whereas a small recovery delay might actually be required. To assume otherwise, however, would imply modelling the characteristics of one particular implementation which is not the intention of a general comparison of switch fabrics.

The simulation model is a discrete time, event driven simulator. For the measurements of throughput at saturation and delay for slotted traffic a resolution on the time axis of 1/100 of a packet length was found to be sufficient. (Time within the simulation model is generally normalised to the packet length, the packet length being the emission delay of a packet at the speed of the switch fabric.) In effect, no limit was placed upon the number of set-up attempts per packet. The simulation was in general initialised with random time relationships between all packets and run for a time of 200 packet lengths to acquire stability before measurements commenced. The majority of simulations were run for a total of 200,000 packets which, for the measurements of throughput at saturation, yielded a standard deviation of about 0.4% of the mean for the smaller network sizes to about 0.2% for the larger networks (64×64 and greater). The throughput per port at saturation is normalised to represent the proportion of time on average that an output port carries valid traffic. Due to the simplification of the model this figure will include the routing bits of the tag, the label, and any other line code, framing or other overhead bits that a packet must carry. The useful port bandwidth is thus slightly less than that indicated by the measurements of throughput at saturation but the figure proves useful for the purpose of the comparison of switch fabrics.

6.3 The Crossbar Switch Fabric

Throughput at Saturation

The crossbar switch fabric is non-blocking thus it offers the ideal performance against which other interconnection networks may be compared. In the input buffered cross-

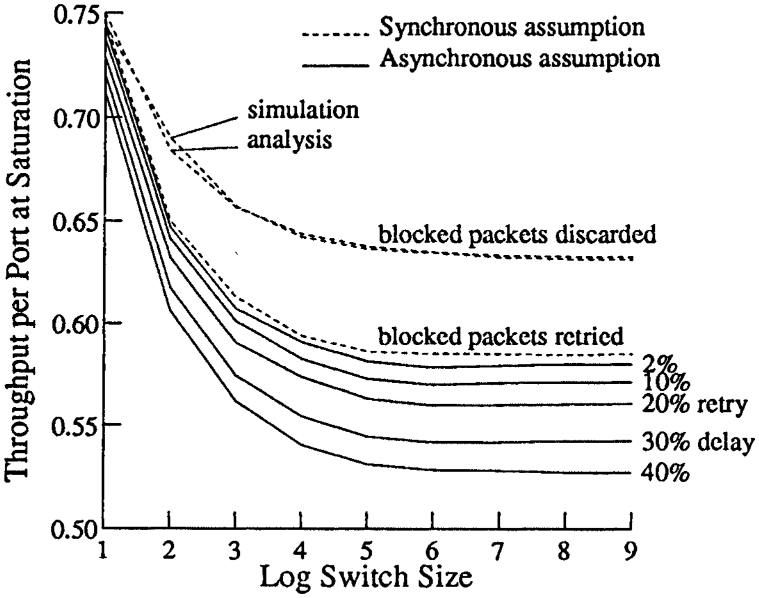

Figure 6.1: Throughput at saturation for the crossbar switch fabric.

bar switch, blocking, (or to be more precise contention,) proceeds solely from the probability of multiple sources attempting to transmit to the same destination at the same time. The throughput at saturation results of crossbar switch fabrics up to size 512×512 are presented in fig. 6.1 under various assumptions. The switch size ($N \times N$) is expressed as $\log_2 N$ and the curves are discrete, points being connected purely for visual convenience. In the synchronous assumption all packets are presented to the switch synchronously and in phase such that contention for all output ports occurs instantaneously and those packets that are successful are then transmitted across the switch fabric. Those packets that are blocked may either be discarded or wait exactly one packet length to be resubmitted in the next timeslot. In the asynchronous assumption there are random time relationships between all packets. If a packet is blocked it waits for the duration of one retry delay and is then resubmitted independently of all other packets. The retry delay is expressed as a percentage of the packet length.

The analysis of [126] gives an expression for the throughput at saturation of a crossbar switch fabric under the assumptions of synchronous packet arrival and blocked packets discarded. The expression is $1-(1-1/N)^N$ and is compared with the simulator output in the upper curves of fig. 6.1. It has an asymptote of $(1-e^{-1}) = 0.632$ for large N and the simulation may be seen to agree very closely with the analytical result. This is hardly surprising as the crossbar switch fabric is a very simple network to simulate. If packets are resubmitted rather than discarded, the synchronous crossbar switch fabric becomes much harder to analyse but both [76] and [71] give the result of $(2-\sqrt{2}) = 0.586$ for the asymptote of the throughput at saturation for large N. Again this agrees closely with the simulation results.

The throughput at saturation for asynchronous operation of a crossbar switch fabric at various values of retry delay is also given in fig. 6.1. Asynchronous operation with a retry delay of zero is equivalent to synchronous operation as the contention for the output ports occurs immediately they become free from the previous packet. In asynchronous operation the throughput at saturation is reduced as the retry delay increases. This is because the greater the retry delay, the greater the probability that an output port spends some idle time after serving one packet before a contending input port resubmits a packet set-up attempt.

For the synchronous crossbar switch fabric with blocked packets retried the throughput at saturation for traffic with a uniform random destination distribution represents the probability (p_a) that any packet will be successful on any set-up attempt. Hence the operation of any input port of the switch may be modelled as a geometric server in which the probability that a packet will require j set-up attempts is $p_a(1-p_a)^{j-1}$. The mean delay across the switch fabric is thus:

$$\sum_{j=1}^{\infty} j p_a (1-p_a)^{j-1} = 1/p_a$$

The delay across the switch fabric was measured with the simulation model and was shown to equal the inverse of the throughput at saturation for both synchronous and asynchronous models.

Mean Delay for Slotted Traffic

Analytical results are also available of the mean delay for slotted traffic for both the input buffered crossbar switch fabric and the output buffered switch. In [76] the mean delay of the input buffered crossbar switch fabric is derived numerically while that of the output buffered switch is given as:

$$\frac{(N-1)}{N} . \frac{p}{2(1-p)} + 1$$

where N is the size of the switch and p is the traffic load (which is the probability that any input timeslot contains a packet). An approximate expression is derived in [71] for the mean delay of the input buffered crossbar switch of large N:

$$\frac{(2-p)(1-p)}{(2-\sqrt{2}-p)(2+\sqrt{2}-p)}$$

These analytical results are plotted in fig. 6.2, for switches of large N, together with the results of the simulation model. The curves from the simulation model and the analysis of [76] are virtually coincident while the analysis of [71] gives an approximation of reasonable accuracy.

The effect of introducing input queue by-pass, and output buffering across a two-plane crossbar switch fabric, is shown in fig. 6.3. (A two-plane design with output buffering is termed 'double buffered' and a double buffered structure with input queue

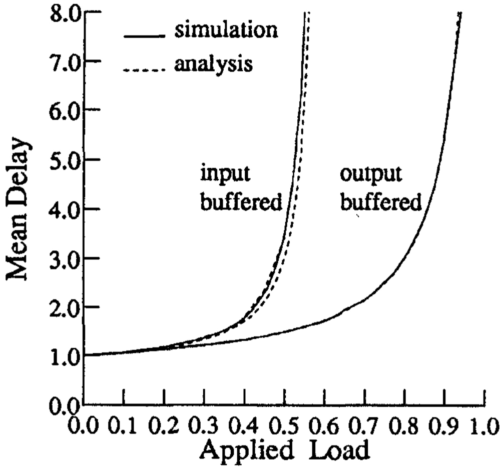

Figure 6.2: Analysis and simulation of mean delay performance for slotted traffic.

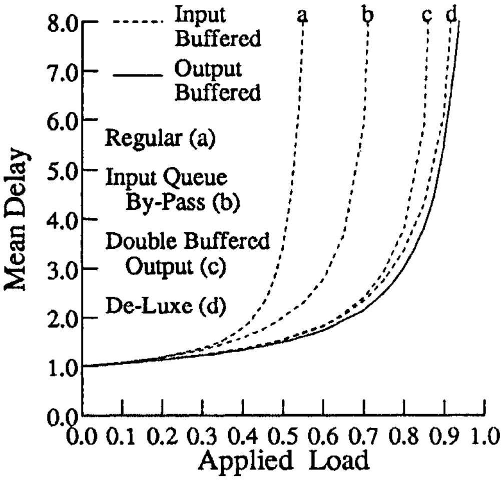

Figure 6.3: Mean delay performance of crossbar switch structures for slotted traffic.

by-pass is referred to as 'the de-luxe model' for convenience. A switch structure with neither queue by-pass nor output buffering is termed 'regular' or occasionally 'pure input buffered'.) These curves apply to the input buffered crossbar switch fabric of large N and are compared to the delay performance of the output buffered switch. For the regular model, each output port was capable of handling only a single packet arriving at any time whereas the double buffered model could handle two. The length of both input and output queues was effectively unlimited. A retry delay of 10% of the packet length between unsuccessful set-up attempts was assumed in the input queue by-pass model. The double buffered curve gives the performance that a two-plane input buffered Batcher-banyan switch fabric might achieve with output buffering across the two planes. It assumes that each input port is only capable of handling at most one packet at a time. (A synchronous Batcher-banyan switch fabric is effectively unable to make use of input queue by-pass.) It is clear that the performance of the de-luxe model approaches very closely that of the output buffered switch but these results assume the use of a non-blocking switch fabric. The performance of the blocking and rearrangeable non-blocking structures will now be presented which are more easily capable of implementation. Finally, a table of the throughput at saturation results of the crossbar switch fabric is given in the appendix for the four classes of switch design with both synchronous and asynchronous operation.

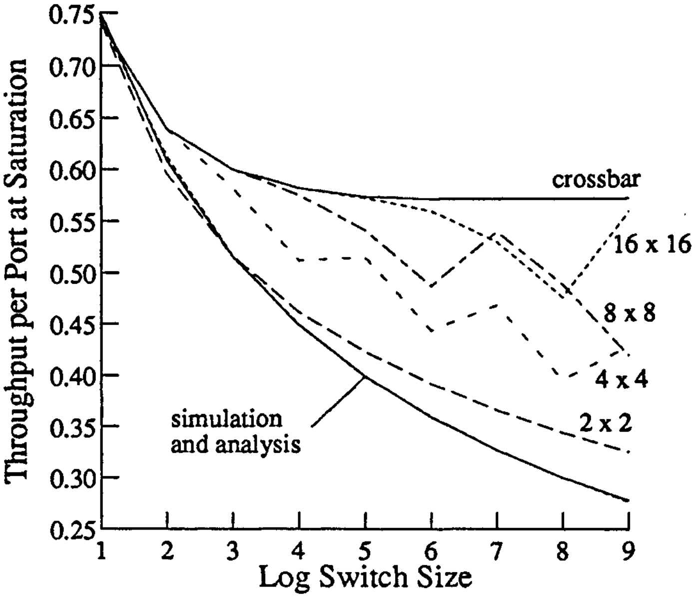

Figure 6.4: Throughput at saturation for single plane input buffered flooding delta networks.

6.4 The Delta Network

Single Plane Delta Networks

The throughput at saturation performance of single plane input buffered delta networks is presented in fig. 6.4 constructed from switching elements of degree 2,4,8 and 16. A flooding algorithm has been used with asynchronous operation, blocked packets resubmitted and a retry delay of 10% of the packet length. The corresponding curve for the equivalent crossbar switch is included for comparison. The perturbations in the curves of degree greater than 2 are due to the presence of multiple paths across the network between the same input and output ports. The minima indicate the pure delta network in which the size of the network is an integer power of the degree of the switching element and only a single unique path connects any input to any output.

For the pure delta network under synchronous operation with blocked packets dropped, analytical results for the throughput at saturation are available from [126] in the form of a recurrence relation:

$$m_i = 1 - \left(1 - \frac{m_{i-1}}{d}\right)^d$$

The throughput at saturation at the output of stage i is given by m_i where $m_0 = 1$ and the delta network is constructed from switching elements of degree d. The analytical results from the above expression agree very closely with the results from the simulation model for synchronous operation with blocked packets dropped, to within 0.4%. This is well within the 95% confidence interval of the simulation results.

The analytical and simulation results for delta networks of degree 2 are compared in fig. 6.4 where it may be seen that the two curves are virtually co-incident.

The throughput at saturation performance of the delta network under synchronous operation with blocked packets retried is approximately 8% lower than with blocked packets dropped. This reduction in performance is of the same magnitude as for the crossbar network, fig. 6.1. For the case of blocked packets retried, the throughput at saturation of the delta network under synchronous operation is less than that obtained under asynchronous operation whereas for the crossbar switch fabric a slightly greater result was obtained, fig. 6.1. This is due to the additional blocking effect of unsuccessful packet set-up attempts which themselves cause blocking within a multi-stage network. This effect is at a maximum in the synchronous model where all set-up attempts occur together. It is approximately proportional to the number of stages of interconnection within the network. For asynchronous operation the magnitude of the effect depends upon the packet length, the retry delay, the size of the switch fabric and the speed with which a packet set-up attempt is detected and removed from the switch fabric. The selection of the retry delay in an asynchronous switch is thus a compromise between the loss of throughput from an increasing retry delay, illustrated in fig. 6.1, and the increase in blocking resulting from an increased number of unsuccessful packet set-up attempts, for decreasing retry delay, in the delta network. A retry delay of 10% of the packet length has been selected as a reasonable compromise for fast packet switching applications. An investigation of the case in which the retry delay was exponentially distributed about a mean value yielded no difference in performance to that of a constant retry delay.

Multi-Plane Delta Networks

The throughput at saturation of a switch fabric constructed from multiple delta networks in parallel is illustrated in fig. 6.5 in which multiple paths are investigated by (a) a searching algorithm and (b) a flooding algorithm. The networks are constructed from switching elements of degree 8 and 1 to 4 parallel switch planes are shown with a retry delay of 10% of the packet length and are compared to the crossbar switch fabric. Both the input and output port controllers are capable of handling only a single packet at a time (i.e. the regular switch structure). It is evident that the use of two switch planes in parallel yields a useful increase in throughput performance beyond that of a single plane but that the use of more than two switch planes offers little incremental improvement in throughput performance.

A maximum of two switch planes in parallel has thus been selected for the regular switch structure. No more than two switch planes have been investigated for the structure in which output buffering is provided across the switch planes as the complexity of such a structure begins to approach that of an output buffered switch. The performance of the switch may be enhanced with much less hardware than a three plane output buffered structure would require. Two possible techniques are the use of input queue by-pass and input port controllers capable of handling two packets simultaneously, one across each switch plane.

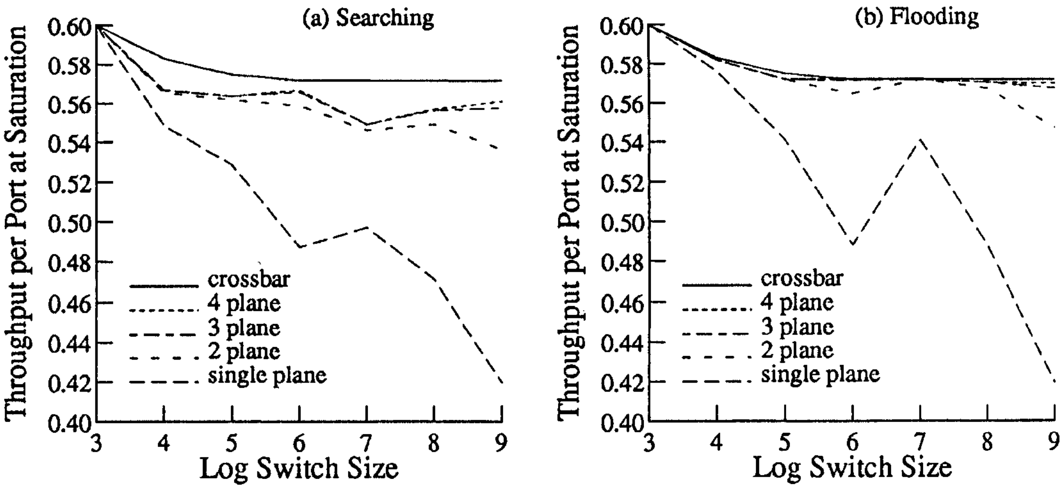

Figure 6.5: Throughput at saturation for multiple delta networks in parallel.

Two basic algorithms exist to select a path through a switch structure that offers multiple paths: flooding and searching. In the flooding algorithm the number of equivalent paths that require investigation changes with the size of the switch fabric thus the switch fabric would require construction with a number of different basic switching elements. A slightly modified flooding algorithm has been investigated that permits the switch fabric to be constructed from identical switching elements. It floods across the parallel planes but searches within each plane and is referred to as the flood-planes algorithm. For both the flood-planes and searching algorithms, improved performance is obtained if a random path is selected for the first set-up attempt, (flood-planes random and search random.) After this, paths are searched in sequence until a free path is located across the switch fabric. The throughput at saturation of the various algorithms is compared in fig. 6.6 for a two-plane regular delta network with switching elements of degree 8 and a retry delay of 10% of the packet length. The flood-planes random algorithm has been selected for further study as it differs from the pure flooding algorithm only marginally and results in a simpler hardware implementation.[1] Also the simulation does not model the interference between failed set-up attempts which is likely to be larger in the pure flooding algorithm than for the flood-planes algorithm.

The throughput at saturation for pure input buffered two-plane delta networks constructed from switching elements of degree 2,4,8 and 16 is shown in fig. 6.7 for a flood-planes algorithm and a retry delay of 10%. Curves for the crossbar switch fabric and the single plane flooding delta network of switching elements of degree 8 are included for comparison.

The effect of input queue by-pass and output buffering on the mean delay performance with slotted traffic is shown in fig. 6.8. It presents the results for a 64×64

---

[1] Further references to the flood-planes and searching algorithms in the context of delta networks imply the flood-planes random and search random variants.

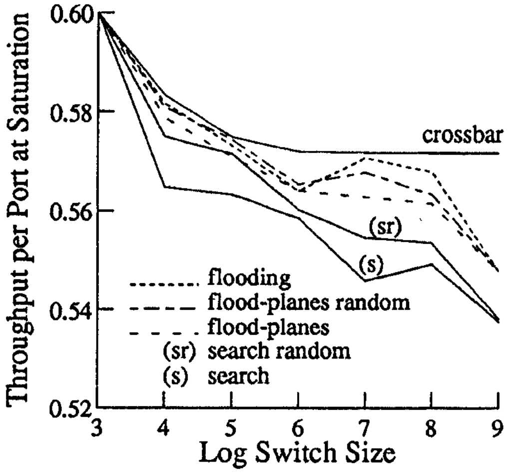

Figure 6.6: Comparison of algorithms to select a free path across the network.

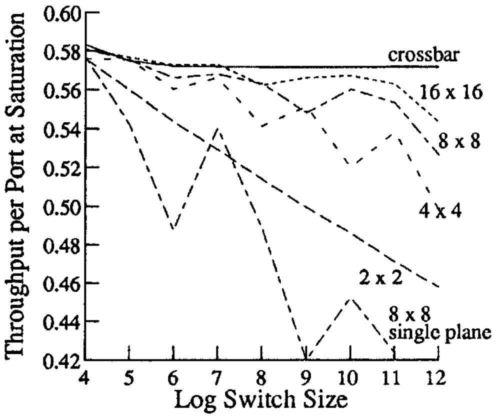

Figure 6.7: Throughput at saturation for two-plane pure input buffered delta networks.

two-plane delta network of switching elements of degree 8 with a flood-planes algorithm. The performance of the single plane crossbar, the two-plane output buffered crossbar, and the output buffered switch are also given for comparison. The curves behave as expected and the de-luxe two-plane delta network, (output buffered with input queue by-pass,) offers a similar performance to the two-plane output buffered crossbar switch fabric. The detailed results of the throughput at saturation performance of both single and two-plane delta networks with input queue by-pass and output buffering are given in the appendix. Delta networks constructed from switching elements of degree 2, 4, 8 and 16 are considered with a flood-planes algorithm and a retry delay of 10% of the packet length.

A More Accurate Model

The simulation results presented so far, for the delta network under asynchronous operation, have ignored the element of blocking caused by unsuccessful packet set-up attempts. A more accurate simulation model was developed to measure the error introduced into the simulation results by this simplification. In the accurate model, packets are set up on a stage by stage basis such that all set-up attempts, partial, successful and unsuccessful, contribute to the blocking within the network. A packet length of 256 bits was selected with an additional routing tag of length $\log_2 N$ bits. It was assumed that the partial path of a blocked set-up attempt would be cleared after a delay of two bit times. The performance of a searching algorithm was modelled as the multiple simultaneous set-up attempts of a flooding algorithm would have consumed an excessive amount of computer time in the simulation.

The percentage error of the simple model expressed with respect to the results of the more accurate model are given in table 6.2 for switching elements of degree 8 and

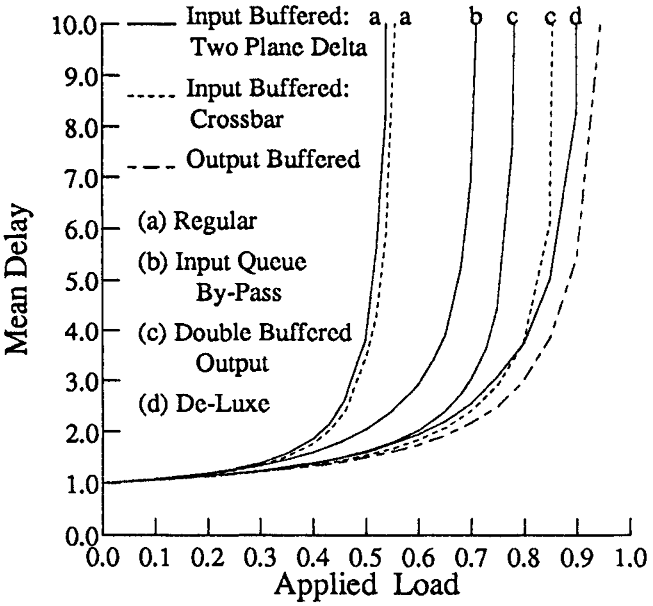

Figure 6.8: Comparison of mean delay performance for slotted traffic of various switch structures of size 64×64.

Size	Single Plane	Two Plane	Double Buffered
8	−0.83	−0.85	−0.34
16	+2.44	+1.38	+0.68
32	+2.32	+0.88	+0.29
64	+2.45	+0.27	−0.20
128	+3.21	+2.61	+0.84
256	+3.21	+1.94	+0.34
512	+4.24	+1.15	−0.01

Table 6.2: Percentage error in throughput at saturation of simple model for delta networks with switching elements of degree 8.

a retry delay of 10% of the packet length. The results for the single plane, two-plane (regular) and two-plane output buffered (double buffered) models are given and the results for the same structures with input queue by-pass are similar. The percentage error in the simple model for single plane delta networks with switching elements of degree 2, 4, 8 and 16 is given in table 6.3.

The simple model slightly underestimates the throughput at saturation of a single stage switch fabric, i.e. a switch fabric consisting of a single crossbar switching element. This is because the simple model assumes that an output port is busy for the entire duration of the packet plus the route bits whereas in the accurate model an output port is not marked busy until the route bits have been consumed. The simple model of the single plane switch is the least accurate as in this structure the density

Size	Degree 2	Degree 4	Degree 8	Degree 16
2	-0.12	–	–	–
4	-0.05	-0.22	–	–
8	$+1.34$	$+1.37$	-0.83	–
16	$+2.79$	$+1.26$	$+2.44$	-0.97
32	$+8.41$	$+3.29$	$+2.32$	$+3.13$
64	$+6.75$	$+4.21$	$+2.45$	$+2.79$
128	$+7.28$	$+4.03$	$+3.21$	$+2.25$
256	$+7.46$	$+5.25$	$+3.21$	$+2.37$
512	$+8.63$	$+4.92$	$+4.24$	$+3.91$

Table 6.3: Percentage error in throughput at saturation of simple model for single plane delta networks.

of unsuccessful packet set-up attempts is greatest. Also the error introduced by the simplified model is greater for delta networks constructed from switching elements of lower degree. In general, for two-plane delta networks constructed from switching elements of degree 8 or more, the error in the simple model for the estimation of throughput at saturation is no greater than 2%.

The comparison between the simple and accurate models has been measured for the searching algorithm. Thus the measurement of the throughput at saturation under a searching algorithm provides an accurate lower bound on the performance that may be expected from delta networks. The measurements taken with a flood-planes algorithm give an upper bound to the expected throughput performance at saturation. The throughput at saturation measurements from both flood-planes and searching algorithms are presented in the appendix for all of the classes of switch structure discussed, with switching elements of degree 2, 4, 8 and 16.

6.5 The Beneš Network

An investigation of the throughput at saturation of the Beneš and the sub-equipped Beneš structures reveals that the searching algorithm yields a much poorer performance than the flooding and random algorithms and is even slightly inferior to the performance of the equivalent single plane regular delta network. This result is perhaps to be expected as the Beneš structures offer many more paths to be searched sequentially than does the delta network. Also the increased number of switch stages increases the possibility of blocking within the switch fabric at high loads. The random algorithm offers a performance which is better than that of the equivalent single plane regular delta network because it always selects a path through the distribution stages of the switch fabric which is known to be free. As might be expected the flooding algorithm offers the best throughput performance which is very close to that of the equivalent crossbar network. The simulation model, however, does not consider the effect of the interference between packet set-up attempts caused by the multiple

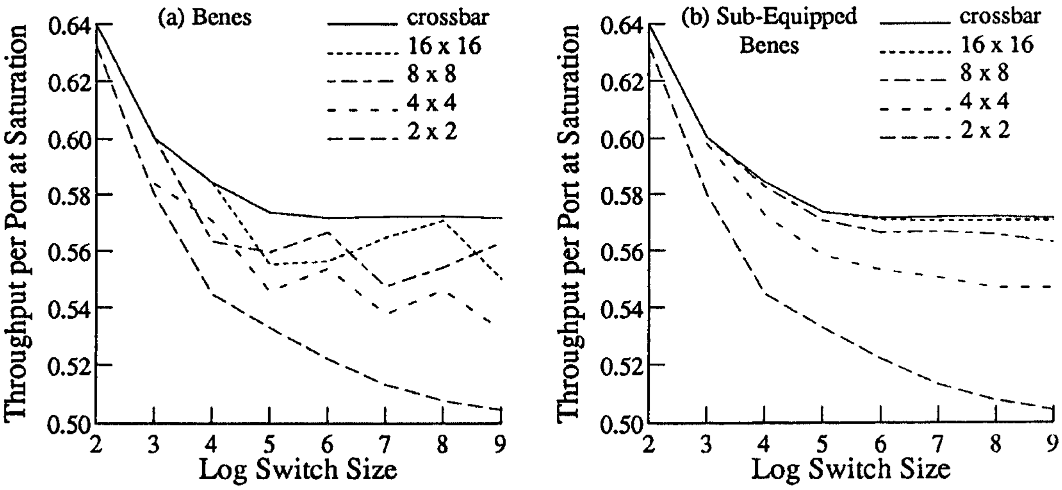

Figure 6.9: Throughput at saturation for flooding Beneš structures.

copies of each packet generated by the flooding algorithm. The results for the flooding algorithm must therefore be considered as an upper bound on the performance of Beneš structures while those of the random algorithm may be considered as a lower bound. The accuracy of the results for the random algorithm is likely to be similar to, if not better than that of the single plane delta network. For applications that require a short packet length the random path selection algorithm is therefore preferred. For longer packet lengths, however, the flooding algorithm may offer improved performance as the relative interference between multiple packet set-up attempts is reduced by the increase in packet length.

Two versions of the flooding algorithm were compared. One version selected at random between all of the free paths to the destination whereas the other was biased to select a path through the uppermost central stages of the switch fabric thus attempting to pack the successful set-up attempts as closely as possible. No difference in throughput performance at saturation was observed between these two variations.

The throughput at saturation for both (a) Beneš and (b) sub-equipped Beneš structures with a flooding algorithm and a retry delay of 10% of the packet length are given in fig. 6.9. The perturbations in the Beneš curve are due to the degree of the switching elements in the first and final stages which may be lower than that of the switching elements in the central stages. The throughput at saturation is a maximum when all switching elements are of the same degree *d* which occurs when the size of the switch fabric is an integer power of *d*.

The detailed results of the throughput at saturation of the sub-equipped Beneš structure for both random and flooding algorithms, with and without input queue by-pass are presented in the appendix. For networks constructed from switching elements of degree 8 and 16 the performance of the flooding sub-equipped Beneš structure is very close to that of the equivalent crossbar network for both regular and input queue by-pass switch designs. Under the random algorithm a Beneš network offers a

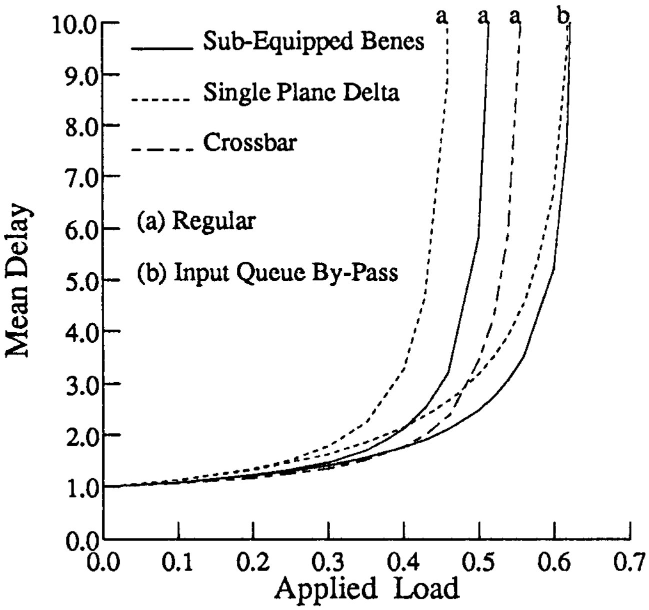

Figure 6.10: Comparison of mean delay performance for slotted traffic of 64×64 sub-equipped Beneš networks against other structures.

performance between that of the single and two-plane regular delta networks whereas with input queue by-pass the Beneš structure offers a performance very similar to that of the single plane delta network with input queue by-pass. For switch structures that feature output buffering across two Beneš switch planes the performance of the crossbar network may be considered as an upper bound and that of the equivalent delta network as a lower bound.

The mean delay performance for slotted traffic of a 64×64 sub-equipped Beneš network of switching elements of degree 8 with a random algorithm is given in fig. 6.10. Both regular and input queue by-pass switch structures are compared to the equivalent single plane delta and crossbar networks. The delay performance of the single plane delta network with input queue by-pass may clearly be seen to approach that of the Beneš structure as the load on the switch increases. This suggests that the improvement in the throughput performance of the Beneš structure which is gained by spatially distributing the incident traffic across the switch fabric is matched by the distribution in the time domain of the input queue by-pass algorithm as the average queue length in the input buffers grows with the load.

Finally, it is unlikely that a Beneš switch fabric will be selected as opposed to a delta network purely on the grounds of its throughput performance at saturation. The performance difference between the two structures is not particularly significant. The Beneš structure is of interest because of its greatly reduced sensitivity to the destination distribution of the incident traffic when compared to the delta network.

6.6 Summary

A simulation model has been developed to compare the throughput at saturation performance of various switch structures according to a number of implementation parameters. Where analytical results have been available, comparison with the results of the simulation model has been shown to yield very close agreement to well within the 95% confidence interval of the simulation results. The performance of a pure input buffered switch with a crossbar switch fabric is shown to be about 58% of that of an output buffered switch. However, the use of input queue by-pass and a two-plane switch fabric with output buffering enhances the performance to become very close to that of the output buffered switch.

For delta networks, the use of a two-plane switch fabric is recommended on the grounds of increased throughput, increased reliability and ease of maintenance. The use of more than two switch planes in parallel is not justified from the point of view of increased performance. The use of switching elements of degree 8 or 16 is preferred against those of degree 2 or 4 because of improved throughput performance and reduced interconnections within the switch fabric. For delta networks the searching algorithm offers a performance only slightly lower than that of a flooding algorithm and a hybrid algorithm which searches within each switch plane but floods across the planes is recommended for its ease of implementation. A two-plane regular delta network offers a throughput performance only slightly inferior to that of a crossbar switch fabric. The introduction of input queue by-pass together with output buffering yields a performance comparable to the two plane crossbar switch fabric with output buffering, and only slightly lower than that of the output buffered switch.

In general, a Beneš switch fabric is unlikely to be favoured above a delta network purely on the grounds of its throughput performance. It is of interest because it reduces the sensitivity of the switch to the destination distribution of the incident traffic. The performance of the Beneš switch fabric lies between that of the equivalent delta and crossbar networks. For applications that require a short packet length the random path selection algorithm is recommended whereas a flooding algorithm may yield improved performance for longer packet lengths.