

BPX 8600 Architecture and Performance

This paper provides guidelines for positioning the Cisco BPX[®] 8600 series wide-area switch when it comes to architectural discussions and analyzing performance figures. It targets any audience that has interest in understanding the BPX architecture. It will cover the following topics:

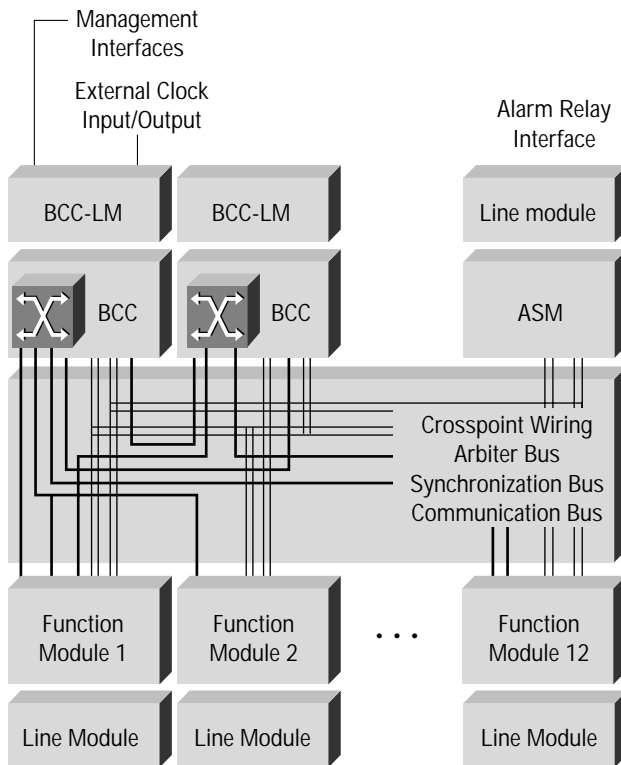
- Introduction
- Crosspoint matrix concept description
- Buffering strategies and blocking performance
- BPX buffering and crosspoint architecture
- The port speed issue: Clos' Rule
- BCC-4
- Asymmetrical crosspoint switching matrix
- Crosspoint arbitration
- Arbitration with BCC-3 and "traditional" function modules
- Arbitration with BCC-4 and new-generation function modules
- BPX switch performance
- Oversubscription
- Multicast
- References

Introduction

Since its introduction in 1993, the BPX 8600 series wide-area switch has been a trend-setting platform in the ATM wide-area switching marketplace, pioneering concepts such as congestion avoidance in ATM networks, service guarantees for multiple classes of service, fair and efficient utilization of network resources, advanced buffer management, fully distributed network topology information, and source-based routing algorithms.

The introduction of the Stratm[™] technology represents, once again, a leap ahead in ATM networking technology. And the flexibility of the BPX architecture allows it to deliver all the benefits of the new Stratm-based function module family (the broadband switch module [BXM] cards) without requiring a dramatic upgrade path. All new function modules are fully compatible and can coexist in any switch with the already-installed modules, thus allowing already-installed BPX nodes to be equipped with Stratm cards. This capability fully exploits the advanced traffic management, large intelligent buffers, and the high port density of the new technology, while providing customers who require additional switch capacity with an easy, nonservice-disruptive growth path.

Figure 1 BPX Broadband Shelf Overview



Available Function and Line Modules

- DS3/E3—45/34 Mbps
 - Three-port trunk card (BNI)
 - Two-port access card (ASI)
- OC 3/STM-1—155 Mbps
 - Two-port trunk card (BNI)
 - Two-port access card (ASI)

Stratm Technology-Based Function Modules

- DS3/E3—45/34 Mbps
 - Twelve- or eight-port trunk or access
- OC 3/STM-1—155 Mbps
 - Eight- or four-port trunk or access
- OC12/STM-4—622 Mbps
 - Two-port trunk or access

The basic architectural characteristics of the BPX broadband shelf are shown in Figure 1—a 15-slot chassis, 12 slots of which are available to support function modules that implement trunk interfaces to other BPX/IGX™/MGX™ or ATM UNI/NNI interfaces. Two slots are reserved for the redundant broadband control cards (BCCs), which combine both the switching fabric and the control subsystem. One slot is used for the alarm status monitor (ASM) card. There is plenty of material describing the hardware architecture of the BPX 8600, but this white paper focuses exclusively on the switching architecture. The BPX switching architecture is based on a crosspoint switch design, which is described in more detail in the following sections.


Crosspoint Matrix Concept Description

The heart of the BPX 8600 is a crosspoint matrix switching fabric, which is basically a space-switching device (a single stage connecting every input with every output). The crosspoint matrix is an independent subsystem that resides on the BCC card. In this section, the “first-generation” BPX crosspoint matrix is used as a basis for discussion.

The primary function of the switch fabric is to pass traffic between interface cards. A crosspoint switching matrix provides users with superior performance compared to that of bus-based products when operating at broadband speeds. The crosspoint switching matrix is a single-element, externally buffered matrix switch. The BCC cards available until BPX Release 8.4.1, such as the BCC-3, are 16 x 16 matrices, where each of the 16 crosspoint matrix ports is a full-duplex-capable, 800-Mbps link. Of the 16 “ports” to the crosspoint matrix, only 14 ports are used: 2 by the redundant BCCs, and the remaining 12 are available for the 12 function modules that can reside on the BPX broadband shelf. Each interface slot in the BPX 8600 is also connected to a redundant switching matrix (which resides on the redundant, hot-standby BCC) with a redundant, full-duplex, 800 Mbps serial interface. Thus, in the event of a control card failure, the redundant card can take over all traffic without cell loss.

A very simplified overview of the operation of the crosspoint matrix is as shown in Figure 2. Every 687.5 ns, the crosspoint matrix arbiter polls the 14 connected cards for the internal destination of the next cell each card wants to transmit. The crosspoint matrix checks all the requests, verifies that there are no conflicts, subsequently configures the crosspoint to serve all requests, and finally grants the cards permission to send that cell to the serial 800 Mbps crosspoint port. The cell is switched to the destination egress card. It is important to note that the function modules also implement onboard arbitration functions, and that the dialog between the function module arbitration logic and the crosspoint arbiter is, in fact, far more elaborate than this short description might infer.

Figure 2 Current BPX 8600 Switch Architecture



Buffering Strategies and Blocking Performance


A question that often arises when the switch architecture is discussed is the “buffering strategy” implemented on the BPX 8600. With a crosspoint architecture, two buffering strategies are possible: input buffering and output buffering. Note that while the ATM literature primarily discusses these two options as being mutually exclusive, practical switch implementations using a crosspoint matrix design always tend to implement a mixture of both, and that rather than

the fundamental buffering architecture, aspects such as arbitration logic and buffer management become crucial to assess the design’s performance capabilities in a real-world environment.

Also, this section only makes the argument on an intuitive basis rather than the theoretical mathematical approach this subject calls for. For interested readers, some pointers to additional information are provided in the reference section at the end of this paper.

Many switch architectures can achieve “nonblocking” behavior; however, it is not within the scope of this white paper to discuss switch architectures in general. The crosspoint architecture (see Figure 3) used in the BPX 8600 is viewed as “nonblocking,” meaning that when (in the example) four cells arrive at the ingress port destined to the different output ports, the crosspoint matrix will always switch them to the destination. This is an intuitive argument (a more formal mathematical analysis comes to the conclusion that “queuing delays for output buffering are largely a function of the arrival process and not affected by blocking in the switch”). The job of a switching element is to make it possible for every input to reach every output when noncontending requests arrive. A crosspoint matrix achieves this task with a very practical architecture, which explains the popularity of this switch design.

Figure 3 Four-Port Crosspoint Architecture



It is very important to understand that the widely used term “nonblocking,” when looking at ATM switch architectures, refers to the treatment of noncorrelated, statistical, “Bernoulli” traffic (a sequence of cells with no relation to each other). This definition is a major limitation of the theoretical approach when investigating blocking behavior, because even introductory ATM literature acknowledges that the assumption that ATM switches deal with uncorrelated “Bernoulli” traffic cell sequences is considered very optimistic (see Reference 1). It is important to understand that the term “nonblocking” has only theoretical relevance for the most part, and that it is far more important to analyze how the switch architecture can handle “real-world” traffic patterns. (Note that we do not make this statement because the BPX switch is not a nonblocking architecture; when dealing with Bernoulli traffic, the BPX 8600 is a nonblocking, 9.6-Gbps architecture).

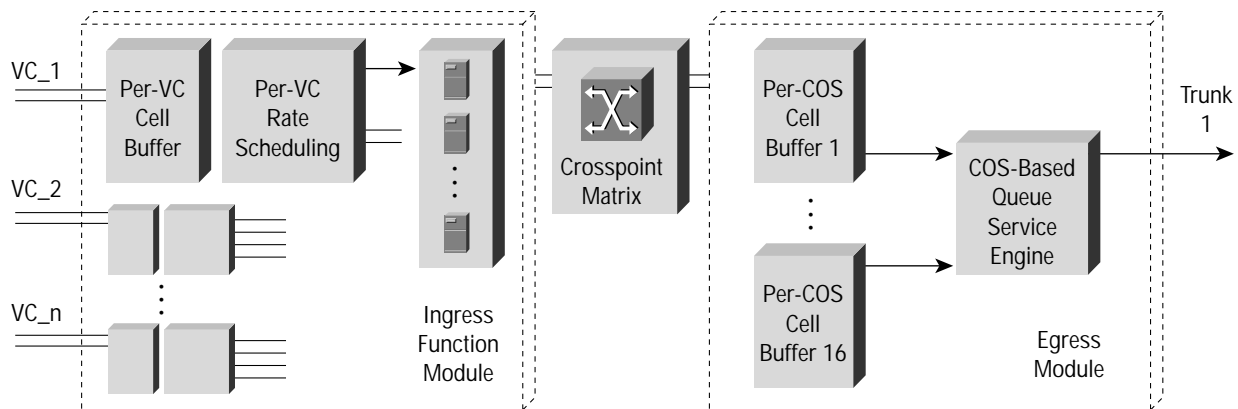
In practical, somewhat simplified terms, the Bernoulli traffic assumption can be used for ports that have thousands of user connections logically multiplexed onto them. Trunks between switches in large networks with many users can be assumed to operate this way. Thus, when it comes to the BPX switch’s traditional trunk card design, the BNI (Broadband Network Interface) card has always relied nearly exclusively on egress buffering (up to 32,000 cells can be buffered for every trunk in the egress direction).

On an ATM User-Network Interface (UNI), though, the assumption that user traffic is going to be uncorrelated Bernoulli traffic becomes a liability for the switch design. The frame-oriented, higher-layer protocols that feed long frames into the convergence, adaptation, and segmentation layers, such as TCP/IP, lead to long bursts of correlated cells that are headed toward the same destination and thus the same output port in the switch fabric. A contention situation now leads to the size of the egress buffer, which now must try to

accommodate these long bursts, and becomes the factor that decides whether an ATM switch architecture is lossy or not and therefore should be regarded as blocking or nonblocking. Consequently, the egress buffer is a critical resource in the switch and in the network, and it becomes critical that intelligent flow control algorithms that rely on feedback messages accurately reflecting the momentary utilization of these resources work on top of these egress buffering architectures to avoid cell loss under high load.

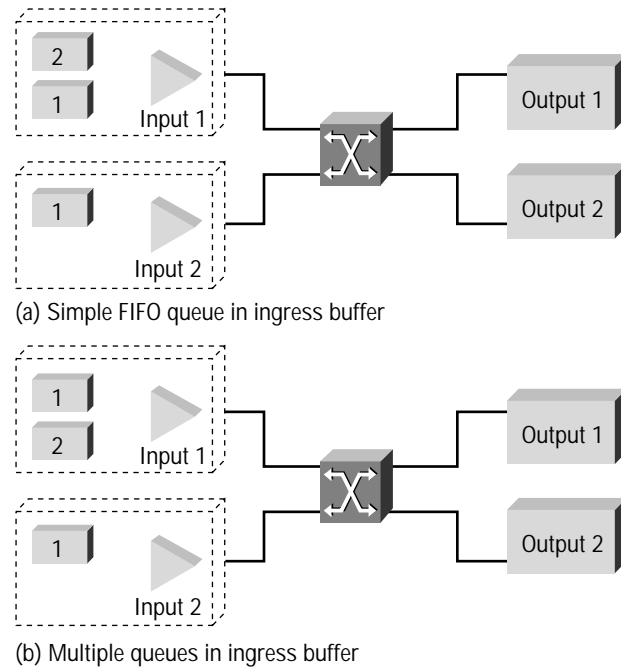
Thus, it is mandatory for an ATM service switch architecture to devise mechanisms to control long, correlated cell bursts on ingress ports without dropping cells other than in the most extreme network overload situations, and without allowing these bursts to flow in an uncontrolled fashion toward the egress buffers. The BPX switch solves this design challenge by implementing large ingress buffering for cells for ports where correlated traffic patterns can be expected. This large ingress buffer space is allocated on a per-connection basis and can support up to 64,000 cells per connection on the first-generation ASI-1 (ATM Service Interface) card for T3/E3 ports, and up to 960,000 cells on Stratm-based function modules. The implementation of per-connection buffers also allows for individual per-connection rate scheduling, using feedback messages from the network that provide accurate information on the network load situation to adjust the rate at which the individual per-connection buffers are emptied. Thus, the ingress buffer addresses two issues: it dramatically enhances the switch’s ability to cope with correlated cell flows, and it provides an optimal architectural framework on which to implement advanced traffic management mechanisms. This architecture (shown in Figure 4) is crucial to provide for robust, frame-over-cell data services (which includes all data services) and is unique to the BPX system architecture.

Figure 4 Simplified “Traditional” BPX Buffering Philosophy Overview



Ingress buffering has a negative reputation for eventually causing head-of-line (HOL) blocking. This is an undeserved reputation, it should be noted, because the association of ingress buffering and HOL blocking can be easily eliminated with a more elaborate ingress buffering strategy. HOL blocking results from cells that come from many connections, and are thus destined to different output ports, being queued in the ingress direction using a simple first-in/first-out (FIFO) algorithm before going to the cell-switching matrix (see Figure 5). As the first example (a) shows, the problem arises when, because of the simple FIFO architecture of the ingress buffer, the two cells destined to Port 1 cause a contention situation, one of the cells has to remain in the buffer. The next cell, destined to Port 2 (which is not experiencing any contention during that cycle), can't be treated because of the strict FIFO principle. The obvious remedy is to implement a series of parallel queues that can be treated concurrently and transmitted to the crosspoint matrix, thus making it possible for the switching matrix to switch a different cell from a different buffer if the destination port of the cell in another queue cannot be treated because of a contention situation. This solution can be seen in Example (b). The arbiter function on ingress Port 1 now can decide whether to treat the cell destined to Port 1 or Port 2 to achieve the best possible utilization of the switching matrix. The BPX 8600 implements this principle wherever the risk of HOL could affect the switch's nonblocking performance. The Stratm-based cards implement 256 individual queues that can transmit their content to other cards in the switch at any time, which together with the distribution of the crosspoint matrix arbitration intelligence, totally eliminates the possibility of HOL blocking in the BPX switch.

Figure 5 HOL-Blocking Elimination with Multiple Ingress Queues and Distributed Arbitration



BPX Buffering and Crosspoint Architecture

All in all, the BPX architecture can best be described as a crosspoint matrix architecture that primarily relies on output buffering on the function modules, but complements it with elaborate ingress buffering mechanisms to be able to cope more efficiently with real-world traffic characteristics. Ingress buffering is required to allow intelligent, distributed arbitration (to allow the arbitration logic to resolve contention situations and utilize the crosspoint matrix as efficiently as possible). As mentioned before, ingress buffering also allows implementation of virtual source/virtual destination (VS/VD) Available Bit Rate (ABR) traffic management within the switch and allows for a safe implementation of oversubscription—a requirement that a cost-efficient ATM service switch architecture has to fulfill in many application environments.

The reason for not providing buffering on the matrix itself is that it does not make a difference where the buffering is provided—it is the buffering architecture and functionality that counts. Implementing the buffering on crosspoint board itself can be a very constricting architectural decision; distribution of the buffers onto the peripheral function modules makes it far easier to implement large buffers.

The Port Speed Issue: Clos' Rule

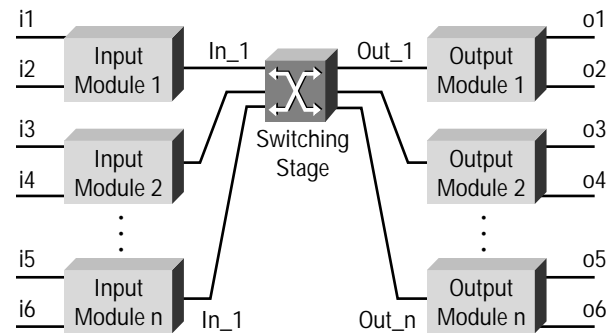
One should also keep in mind that aggregated traffic volume from user ports and the speed of the port in and out of the crosspoint matrix play a crucial role when looking at the blocking behavior of the switch. Most vendors neglect this aspect when presenting and explaining their architectures.

A helpful rule of thumb when discussing this is “Clos' Rule,” which basically converts different switching architectures back into their equivalent space-switching-based architectures and checks whether the rule $k^* = 2k - 1$ is given, which characterizes a nonblocking architecture. As with ATM, converting back to a space-switching architecture is not an easy task; a simple generalization turns Clos' rule to $k^* = 2k$, meaning that if a switch must handle input lines of speed k , the switching stage needs to run at twice that speed to guarantee nonblocking performance. (Note: this is a simplified claim, and with advanced arbitration and ingress queuing, the statement needs to be revised, but it enables an introductory, intuitive discussion.) While most switch architectures do this for T3 speeds, high-density OC-3 cards already suffice to push many an architecture beyond the limits of its “nonblocking” claim, and in fact, OC-12 interfaces turn all existing ATM service switches into blocking architectures. This is not the case for the BPX switch with the next-generation BCC-4, as the 1.6 Gbps allocated in the egress direction, sufficient to fulfill and exceed Clos' rule for BXM cards using only one of the two available OC-12 ports. It is for this reason that for OC-12 trunks, where nonblocking behavior is important, only one OC-12 port is used on the BXM card.

Based on Figure 6, it is possible to make a few intuitive points on why typical switch architectures have major problems delivering low blocking probability with increasing interfaces speeds (and thus increasing traffic load). A typical ATM switch implements an architecture where the In_n port speed is equal to the Out_n port speed; typically around OC-12 speeds, that is somewhat more than 622 Mbps. Now, let us assume that ports $i1$, $i3$ and $o1$ are OC-12 ATM ports running at 622 Mbps. This switch architecture has two major problems:

- If ports $i1$ and $i3$ burst for even a very brief period of time with cells attempting to reach port $o1$, an architecture that relies exclusively on output buffering will drop cells right away, because the Out_1 link running at a lower speed than the aggregated traffic of the two ingress ports cannot accommodate the cells. And because the ingress cards do not have any buffers able to cope with this high-speed burst, cells are dropped right away. So every contention situation for an egress port leads to cell loss, which is a serious architectural problem. All of a sudden the architecture has a desperate need for ingress buffering. But primitive ingress buffering implementations can lead to HOL blocking. The same cell loss effect can take place with high-density cards providing an aggregate throughput that equals or surpasses the OC-12 speed that the Out_n links can accommodate, although with a lower statistical probability.

Figure 6 Switch Architecture and Blocking



- The only OC-12 traffic this architecture can accommodate is simple port-to-port forwarding traffic—from Port $i1$ to Port $o1$, for example. And in this scenario, given the speed of the links involved, the output buffer allocated on the egress card is not used efficiently. All the traffic that the Out_n links forward to the card can be immediately forwarded to the outgoing OC-12 port.

With the enhanced control cards (BCC-4) in Release 9.0, the BPX switch implements a switching architecture with 800 Mbps IN_n links, and 1.6 Gbps (in fact, 2×800 Mbps) for the Out_n links with the new 16×32 crosspoint matrix chip, making the architecture optimally suited to face the challenges of OC-12 traffic switching. The following section focuses on the BCC-4 card and discuss this particular enhancement of the crosspoint matrix in more detail.


BCC-4

The enhanced control card for the BPX switch, or BCC-4, fully unleashes the potential of Stratum technology in nodes equipped with BXM function modules.

This next-generation BCC provides enhanced processing power for general administrative node functions, but its real benefit lies in the fact that it provides the BPX switch with a 16 x 32 switching matrix. Some minor modifications have been made to the arbitration scheme to more efficiently handle multicast traffic.

From an architectural point of view, the BCC-4 card is very similar to the existing BCC-3 control card (see Figure 7). The CPU runs the software subsystem responsible for broadband shelf administration. The utility processor exchanges messages over the Communication Bus with ATM Service Interface (ASI) and Broadband Network Interface (BNI) cards to download software or get statistics data from the cards. An onboard Stratum 3-quality clock can be used for high-quality plesiochronous node operation or distributed as a reference, or the node can use any interface or the redundant BCC clocking port signals as a clocking reference.

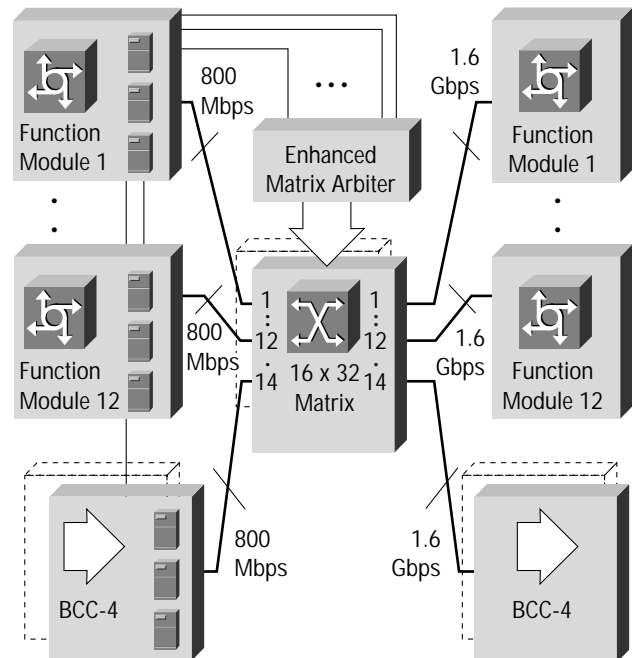
Figure 7 BCC-4 Architectural Overview



Asymmetrical Crosspoint Switching Matrix

The major innovation introduced by the BCC-4 is the asymmetrical crosspoint switching matrix. As Figure 8 illustrates, this only represents a minor modification to the architecture of the BPX switch as it was presented in Figure 2; function modules still transmit their cells to the crosspoint matrix over an 800-Mbps link but in receive direction, a 2 x 800 Mbps (= 1.6 Gbps) link receives cells from the crosspoint matrix. As mentioned earlier when generally discussing switch architectures, this leads to enhanced blocking behavior for high-speed links (OC-12) or high-density cards (such as the eight-port OC-3 BXM card). The switch latency is also improved. The earlier discussion mainly provided an intuitive argument on why this architecture provides better blocking behavior. References [2] and [3] provide an exhaustive analysis of this switch architecture. It is important to state that the combination of advanced arbitration logic on BXM cards and the 16 x 32 crosspoint matrix switch deliver 19.2 Gbps of peak switch throughput.

Figure 8 BPX Switch Architecture With BCC-4



Crosspoint Arbitration

As mentioned in the previous section, the BCC-4 introduces a new, more elaborate arbitration dialog with the function modules, compared with the one used on the BCC-3. The discussion about HOL blocking already stated that the BPX switch has never implemented and will never implement a simple single FIFO queue access scheme to the crosspoint matrix.

Arbitration with BCC-3 and “Traditional” Function Modules

The existing function modules, the ASI and BNI cards, implement three queues (Constant Bit Rate [CBR], Variable Bit Rate [VBR], and multicast) that can send their requests to the arbiter on the existing BCC cards. It is important to note that the arbitration process is distributed; a subset of functions and intelligence resides on the function modules, while overall control functions reside on the crosspoint arbiter of the BCC. The arbiter on the BCC cards is then responsible for ultimately granting crosspoint access as instructed by the arbitration function on the function modules in a way that minimizes contention situations.

Arbitration with BCC-4 and New-Generation Function Modules

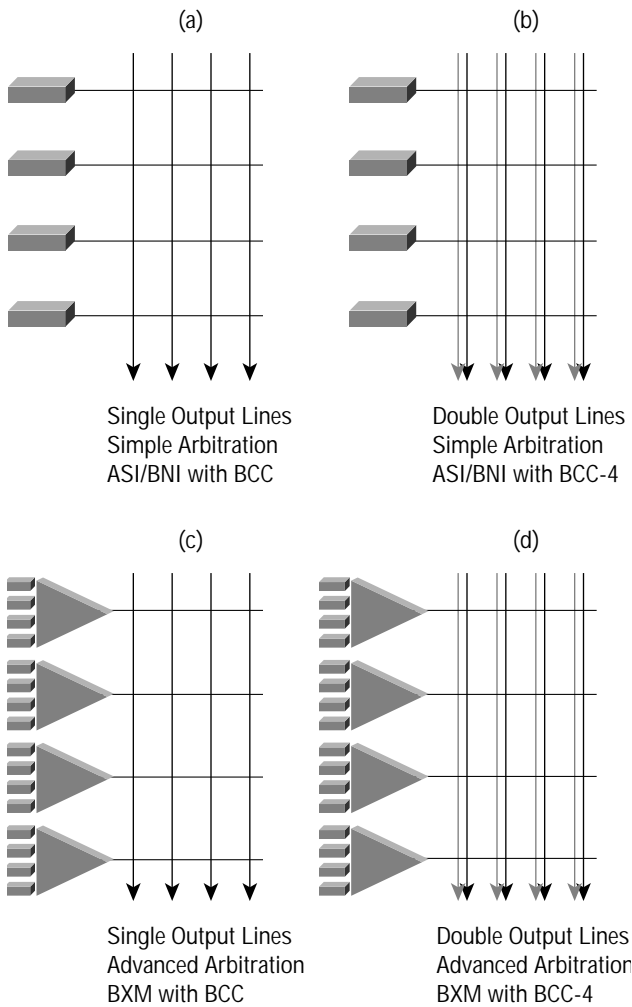
New-generation function modules, the BXM cards based on Stratm technology, implement far more than three queues toward the crosspoint matrix. To guarantee total absence of HOL blocking and minimal intraswitch delay, each class of service has its own queue and each possible intraswitch destination a full set of class-of-service (CoS) queues. Because there are 16 possible destinations from the crosspoint matrix and 16 classes of services supported, the result is 255 individual queues that can issue individual requests to the crosspoint switching matrix arbiter. Effectively, this turns the BPX switch into a 16-plane architecture. The BXM cards implement an advanced onboard arbitration that adapts

to the type of crosspoint matrix and arbiter on the control card; if a BCC-3 with the 16 x 16 architecture is encountered, the BXM arbitration process maximizes the utilization of this switching fabric and eliminates contention situations. If a BCC-4 is encountered, the BXM adapts to maximize the utilization of the 16 x 32 crosspoint switching matrix and to optimally interwork with the advanced arbitration logic on the BCC-4. It is important to emphasize that because BXM cards can be configured to oversubscribe the 800-Mbps link toward the crosspoint switching matrix, advanced arbitration is a key system requirement.

Oversubscription is not to be viewed as a limitation, but rather as an actual strength of the BPX architecture, because it enables cost-efficient service implementation for ATM access. It should also be noted that for backward compatibility reasons with ASI and BNI cards, the BCC-4 also implements full interworking with these function modules, and thus a mixture of all types of function modules and control cards in one switch is fully supported.

Because full interworking between all existing and future cards is ensured in the BPX switch to maximize investment return for customers, four possible combinations of crosspoint type and arbitration are possible. Figure 9 shows these combinations, although for ease of illustration, a 4 x 4-type switch is depicted, and other queuing subsystems in the switch (per-VC, or per-CoS, or per-virtual interface on BXM cards) have been omitted. The “classic” BPX architecture with the BCC-3 and ASI and BNI function modules is seen in Example (a). In Example (b), the support for existing ASI and BNI function modules on a BCC-4 supporting arbiter Type 1 is shown. In Example (c), the initial architecture for the BXM card rollout is shown, where the advanced, 255-queue access algorithm into the 16 x 16 crosspoint switch fabric is represented. Finally, Example (d) shows the architecture that will truly deliver a peak switching capacity of 19.2 Gbps on the BPX switch.

Figure 9 Possible Combinations of Arbitration and Crosspoint Matrix



BPX Switch Performance

From a merely intuitive point of view, the assumption would be that (d) is superior in performance to (c), which in turn is superior to (b), while (b) surpasses (a). A more formal analysis leads to the following results when it comes to nonblocking switch throughput and switch delay behavior:

Crosspoint throughput is not limited to 58.6 percent, as is often claimed. This result applies to the most simple kind of arbitration and to a basic, single-line, symmetric crosspoint fabric. The BPX switch uses advanced arbitration techniques and, with the BCC-4, a double-line, asymmetric crosspoint fabric. A theoretical analysis of crosspoints with various arbitration techniques is given in References [2] and [4], and analysis of asymmetric fabrics can be found in References [3] and [5]. The BPX architecture combines

methods similar to those discussed in the BPX Congestion Avoidance White Paper and this paper. The simulation results presented here complement the theoretical analysis, because they take into account the details of the switch arbitration mechanism and show the distinct performance advantage of using advanced arbitration and double output lines in combination.

What is the nonblocking throughput of the BPX architecture with the different BCC and function module combinations? There are two commonly understood definitions of nonblocking. In academic literature, even the simple crosspoint is classed as nonblocking because of the potential to send cells from all cards simultaneously. However, nonblocking is often used in a more conservative sense to mean the saturation throughput.

Figure 10 Delay Performance





Figure 10 shows the mean delay as a function of the load on the fabric for each of the four fabric and line card type combinations. The fabric delay includes time spent buffered on the input side of the fabric, but does not include delay in buffers on the output side of the fabric or in cell processing pipelines or physical layer framing devices, because these will depend on the type and number of ATM ports considered. The result is a mean across all traffic classes. Real-time traffic classes are given preferential delay treatment, so these numbers are a worst case bound on delay for real-time traffic. If nonreal-time traffic is included in the mix, mean delays for the real-time traffic would be lower than indicated here.

Figure 11 shows the 10^{-6} quantile of cell-delay variation (CDV) as a function of the load on the fabric. Again, it includes the CDV introduced on the input side of the switch because of contention for crosspoint lines. It does not include CDV introduced by port-based queuing on the output side of the fabric. The CDV here is across all service classes. In a scenario with a mix of real-time and nonreal-time services, the CDV for real-time services will exceed that shown here, while the CDV for nonreal-time services may be slightly worse.

These findings might look theoretical, but they have been repeatedly confirmed in lab tests with a HP 1401B High-Power Mainframe ATM tester. At *full* utilization of the architecture, the port-to-port delay across the BPX switch never surpassed 29 ns. The peak-to-peak CDV even under very high load, stayed at less than 2 ns.

Figure 11 Cell Delay Variation Performance



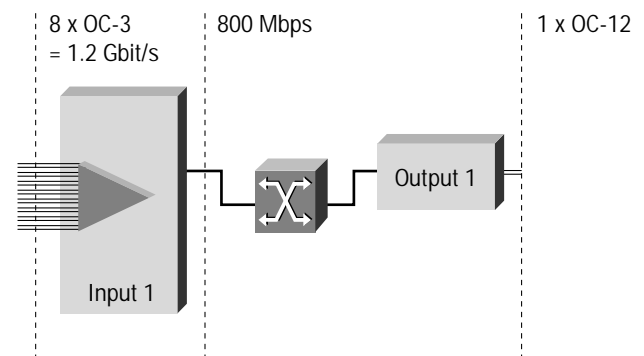
In both cases, the loading model for the switch simulations was Bernoulli traffic with traffic evenly distributed across all input ports, each port having traffic evenly distributed for each destination. The model is commonly applied in the switch performance literature. Of course other loading models may produce slightly different results; however, extensive simulations with a variety of traffic models indicate that the saturation asymptotes shown here are relatively independent of traffic model.

Oversubscription

Although oversubscription has been mentioned in the appropriate section, it is important to stress the fact that the oversubscription achieved with Stratm-based function modules (two-port OC-12 and eight-port OC-3) is not a limitation but a design intent, and is a benefit for service providers who intend to offer cost-effective ATM services (see Figure 12). To support oversubscription without cell loss, large buffers and elaborate arbitration schemes are a must.

Service nodes should behave in a fully nonblocking way for trunks (which is why only one OC-12 trunk or four OC-3 trunks are supported on BXM function modules), but it is mandatory for an access switch to provide aggregation of user traffic toward the switching matrix. A situation where all user ports on one multiport function module show peak activity is a statistical anomaly, and the ingress buffering on the BXM card is able to cope with these extremely rare activity peaks.

Figure 12 Supporting Oversubscription



Multicast

It should also be noted that the enhanced 16×32 crosspoint matrix is optimally suited to accommodate the characteristics of multicast traffic distribution, which is always “output biased” and creates more traffic in the egress direction of the crosspoint matrix. It is, after all, the crosspoint arbiter that replicates the cells from one input port to a number of output ports, and in a second pass, BXM cards can implement logical multicasting to replicate one cell arriving from the crosspoint matrix to different virtual connections (typically different VPs).

References

1. Craig Partridge: *Gigabit Networking*. Addison Wesley, 1993. ISBN 0-201-56333-9.
2. Michael G. Hluchyj, Mark J. Karol: Queuing in high-performance packet switching. *IEEE Journal on Selected Areas in Communications*, Vol. 6, No. 9, December 1988.
3. Soung C. Liew, Kevin W. Lu: Comparison of buffering strategies for asymmetric packet switch modules. *IEEE Journal on Selected Areas in Communications*, Vol. 9, No. 3, April 1991.
4. Rainer Händel, Manfred N. Huber, Stefan Schröder: *ATM Networks*. Addison Wesley, 1995. ISBN 0-201-42274-3.
5. Raif O. Onvural: *Asynchronous Transfer Mode Networks: Performance Issues*. Artech House, 1994. ISBN 0-89006-662-0



Corporate Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters

Cisco Systems Europe s.a.r.l.
Parc Evolic, Batiment L1/L2
16 Avenue du Quebec
Villebon, BP 706
91961 Courtaboeuf Cedex
France
<http://www-europe.cisco.com>
Tel: 33 1 6918 61 00
Fax: 33 1 6928 83 26

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-7660
Fax: 408 527-0883

Asia Headquarters

Nihon Cisco Systems K.K.
Fuji Building, 9th Floor
3-2-3 Marunouchi
Chiyoda-ku, Tokyo 100
Japan
<http://www.cisco.com>
Tel: 81 3 5219 6250
Fax: 81 3 5219 6001

Cisco Systems has more than 200 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the Cisco Connection Online Web site at <http://www.cisco.com>.

Argentina • Australia • Austria • Belgium • Brazil • Canada • Chile • China (PRC) • Colombia • Costa Rica • Czech Republic • Denmark
England • France • Germany • Greece • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia
Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Russia • Saudi Arabia • Scotland • Singapore