

Distributed Switching Architecture Trends and Concepts

E. Lee Sneed
George D. Huensch
John J. Kulzer
Harry Moerkerken

The increasing sophistication of the global telecommunications business is creating new, more demanding requirements for switching systems. To meet these requirements, AT&T must introduce the most up-to-date technologies for advanced capabilities, such as high-bandwidth services, Personal Communication Services (PCS), and increased network intelligence in a timely, cost-effective manner. Switching systems that are adaptable, evolvable, and able to gracefully incorporate these new capabilities and technologies are in the best position to meet market needs. Such switching systems have architectures that are both *evolvable*—able to incorporate new technology and new development processes within their overall structure—and *maintainable*—able to accommodate new features continually without rapid increases in the cost of development or equipment. This paper examines the choices for a switching system architecture by defining the functional components of a switch and examining the choices for implementing them.

Introduction

Today's telecommunications customers have a broad spectrum of needs, ranging from basic access and connectivity to sophisticated value-added services. Building networks that economically respond to this diversity of demands is a key challenge for service providers. The switch represents an enormous capital investment for these service providers, and while new capabilities—media for transport and signaling, network-wide services, etc.—will continue to arise, there is an economic incentive to incorporate them into existing products. Switching systems that are both *evolvable*—able to incorporate new technology and new development processes into their overall structure—and *maintainable*—adaptable to the continuous addition of new functionality without a rapid rise in the development or equipment cost—are in the best position to meet these needs.

In more developed regions, digital switching systems will need to evolve rapidly. For example, in the United States and similarly developed regions, five key initiatives

are rapidly changing the public switched network: the Integrated Services Digital Network (ISDN), intelligent networks, wireless access, fiber-based access and transport, and asynchronous transfer mode (ATM) switching.

ISDN is being deployed as an access platform on which many new value-added voice and data services are being built. It offers effective new message-based capabilities, while reusing the existing access facilities.

Intelligent network service initiatives build on the existing network database architectures used for 800 (Freephone) service, which enables service providers to control switching network elements. These architectures provide for more rapid introduction of network features and the ubiquitous operation of those features across many vendors' switches. Intelligent network service is also emerging as a key building block in the development of Personal Communications Services.¹ In all cases, to support intelligent network needs, switching systems require special control and interface functionality.

Panel 1. Abbreviations, Acronyms, and Terms

APS — Attached Processor System
ATM — asynchronous transfer mode
BSS — Broadband Switching System
CCS — common channel signaling
CFI — control fabric interface
FC — fully centralized
FD — fully decentralized
FS — fully specialized
FU — fully unified
GCNS — gigabit cell network switch
GSM — Groupe Speciale Mobile
ISDN — Integrated Services Digital Network
MC — mostly centralized
MD — mostly decentralized
MS — mostly specialized
MU — mostly unified
PCS — Personal Communications Services
SDH — synchronous digital hierarchy
SONET — synchronous optical network
STM — synchronous transfer mode

Wireless access, widely used to provide “anytime, anywhere communication,” is also being used to offer economical basic service in developing regions. With the introduction of wireless access and tighter coupling with wireless overlay networks, public switched networks will move toward universal personal communications, which provides generalized PCS services across both land-line and wireless access.

Fiber-based access and transport has the potential to increase bandwidth and flexibility to customers, while decreasing cost. Fiber has become the primary inter-switch transport medium in developed regions; new initiatives, such as AT&T Network System’s hybrid fiber-coaxial access architecture, will make fiber-based access an important part of the public switched network. In addition, fiber is a critical element in paving the way for broadband services, which will also depend on new high-speed switching fabrics (e.g., ATM).

ATM technology enables a single, universal type of packet switch fabric to handle voice, data, and video applications simultaneously. It also avoids some problems that synchronous transfer mode (STM) approaches encounter as the number of channels increases.² In developed regions, new information networking needs,

such as increased data networking, high-resolution graphics, visual telecommunications, and multimedia, are promoting the initial deployment of ATM technology in overlay networks.³

Incorporating these key initiatives in central office digital switch designs is a significant challenge, because they affect almost every aspect of switching system architecture. Perhaps the greatest challenge is to introduce change, while preserving the large investment in existing systems. Evolvability, a function of the form of system architecture, is the subject of this paper. Previous papers have described tradeoffs in the distribution of control in digital switching systems,⁴ and in the architecture⁵ of the fabric (switching matrices) and system interfaces (signaling, access, and transport). This paper expands on that work.

Definitions

A telecommunication network can be viewed as a collection of interconnected terminal equipment that provides end users with different forms of telecommunication services, ranging from voice and/or data to image and video. A network comprises one or more *network nodes* connected by some transmission media. Each node, such as a wire center, contains one or more *network elements*, such as switching, operations, or transmission systems.

This paper presents an abstract model of the architecture of a network element, in terms of *control*, *fabric*, and *interface* domains. Each domain is defined by its *function*, not its implementation. A fabric, for instance, performs the function of a fabric, regardless of the components used to implement it.

The *control domain* includes the functions that interpret the switch user’s requirements and issue commands to the fabric and interface domains to carry them out. Control functions include coordination, allocation of resources, sequencing, maintenance of state, fault recovery, status reporting, traffic measurement, billing, and similar administrative activities.

The *fabric domain* includes the functions that provide the communication paths that users request. It also controls *transport*—the encoding, bandwidth, and similar characteristics of the information passing through the switch—and *topology*—the grouping of ports and users into sets that share and exchange information.

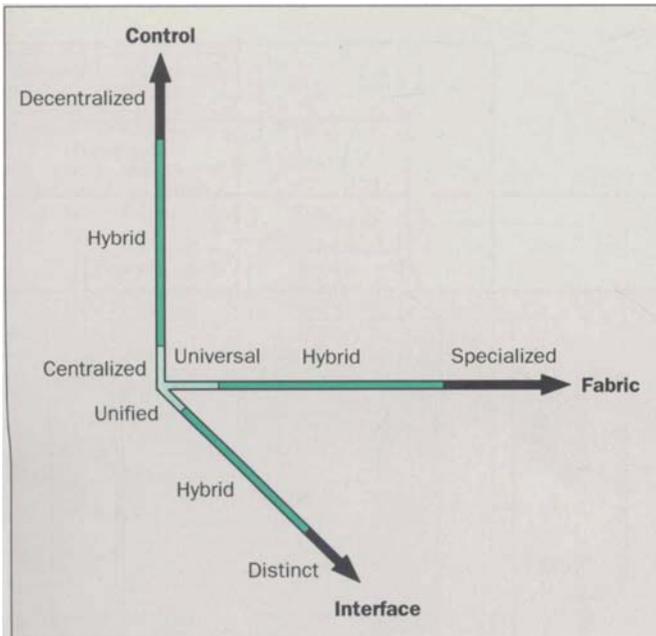


Figure 1. A spatial representation of a CFI continuum. In each domain, diversity increases in proportion to the distance from the origin.

Fabrics may be analog or digital, electronic or optical, narrow- or broadband; topology may be circuit- or packet-switched, synchronous or asynchronous, single- or multistage, blocking or non-blocking, etc.

The *interface domain* is the bridge between the outside of the switch and the inside. Outside the switch, not under its control, are a variety of written and de facto standards; inside is a simpler, generally more proprietary protocol set. The interface domain provides translation between these environments, including voltage and impedance matching, analog-to-digital conversion, line and interoffice signaling, craft and operations system interfaces, and similar functions.

Switching systems, the primary focus of this paper, incorporate all three domains. Intelligent network database systems and operations systems typically contain limited fabrics, or none at all. Transport elements, such as cross-connect systems, often have a control domain with limited functionality.

System Architecture Forms

All switching systems contain control, fabric, and interface components. Various architectural forms exist for implementing each component and for establishing its association with the others. This section introduces an abstract classification scheme that characterizes various architectures.

The CFI Continuum. Reference 4 classifies distribution of control along a continuum of diversity, ranging from fully centralized through hybrid to fully distributed. Reference 5 extends this approach, establishing a sepa-

rate, similar continuum for the fabric and interface domains. In each case, fully centralized implies a single entity that contains all the functionalities required, fully distributed implies a distinct entity for each function, and hybrid implies a combination of these. In the fabric domain, fully centralized equates to a universal fabric, one that can carry all types of traffic; fully distributed represents a specialized fabric for each traffic type. In the interface domain, fully centralized defines a unified interface terminating all types of outside connections, while fully distributed represents a distinct interface for each type.

At this level, the classification scheme can be represented as a three-dimensional space. In this space, each domain (control, fabric, interface) is represented on a separate axis. As Figure 1 shows, diversity within each domain increases with distance from the origin.

The representation of the control fabric interface (CFI) continuum as a three-dimensional space, while intuitive, is limited by its inability to differentiate between physical and functional diversity. A more sophisticated model is necessary.

Physical diversity describes the physical distribution of components. The essential functions of a particular domain may reside in a single physical node, or they may be replicated in many nodes. As the number of nodes and the distance between them increase, the architecture becomes more decentralized. According to this definition, a system can be a geographically distributed network, but can still be centralized, if it needs all (or even most) of its network's nodes to perform its essential functions.

Functional diversity describes the partitioning of a function among system components. In a functionally unified architecture, a single implementation (a single node, or repetitions of a node, if the architecture is physically decentralized) performs all functions in a domain. In a functionally diverse architecture, each function has a separate implementation.

As Figure 2 shows, these two forms of diversity may be used to define a plane, with one axis describing physical diversity and the other functional diversity. The physical diversity axis can be described in terms of two endpoints and two hybrid regions:

- Fully centralized (FC). At this end of the axis, a single physical entity (possibly containing multiple functional components) performs all functions.

- Mostly centralized (MC). In this region, multiple physical entities, located close together, are connected by some mechanism. The degree of coupling between entities determines a particular system's position in this region—a very tightly coupled system (e. g., a shared-memory multiprocessor) is more centralized than a loosely coupled system.
- Mostly decentralized (MD). In this region, the multiple physical entities of the MC region can be separated by significant distances (e. g., miles). Again, a range exists in this division; systems with larger numbers of entities separated by larger distances are considered more decentralized.
- Fully decentralized (FD). At this end of the axis, the entire functional domain is implemented in each network termination or in the terminal equipment.

The functional diversity axis is similar:

- Fully unified (FU). At this end of the axis, a single, general-purpose entity (possibly with multiple, physically distributed copies) performs all of the domain's functions.
- Mostly unified (MU). This region contains several general-purpose entities, each assigned a subset of the domain's function.
- Mostly specialized (MS). In this region, at least some functions are performed by specialized devices, though some may still be performed by general-purpose entities, such as those in the MU region. This is a range—the more functions assigned to specialized devices, the more specialized the system is said to be.
- Fully specialized (FS). At this end of the axis, each of the domain's functions is implemented by a separate, special-purpose entity.

In each domain, the two dimensions of physical and functional diversity produce sixteen distinct forms, ranging from FC/FU, in which a single, centralized unit performs all of a domain's functions, to FD/FS, in which there is a functionally partitioned cluster of specialized units at each termination. The following sections describe how the dimensions of diversity and the domains interact to generate distinct forms.

Control Domain Forms. The control architectures currently in use in switching systems are illustrated in Figure 3a. The central processor architecture uses a single processor to perform all switch control functions, and is designated as FC/FU. Tightly coupled multiprocessors—

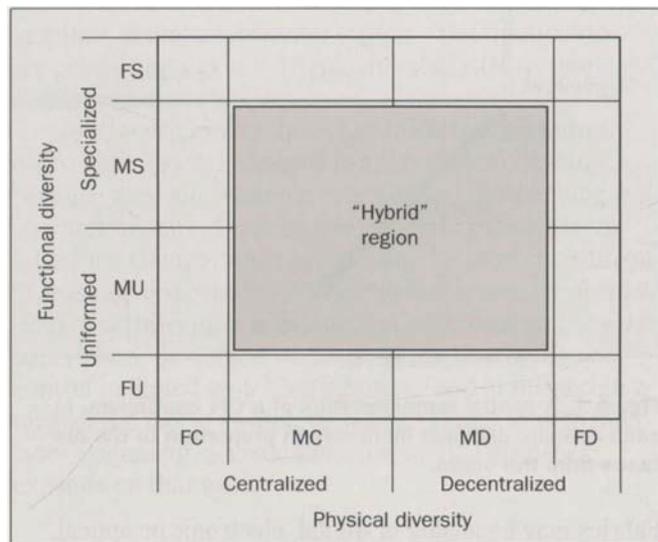


Figure 2. Physical and functional diversity represented as a plane.

those that share memory or incorporate a number of processors on a bus—include symmetric (MC/FU) and functional (MC/MU) forms. Tightly coupled systems—which include special-purpose control elements, such as signal processors—are described as MC/MS. Depending on the degree of functional distribution in the remote nodes, systems containing remote-control elements that can perform a minimal set of functions needed to switch calls are physically MD, typically MD/FU, or MD/MU.

Extremely diverse control configurations are rare in modern switches. The FD form, which has a complete control system at each termination, does not exist in current switching systems, although step-by-step, wireless, and cellular systems contain some aspects of an FD system. The FS form, in which each function has its own control element, is seldom found in electronic switches, but the registers, senders, and markers of crossbar switches exemplify this architecture.

Fabric Architecture Forms. The fabric domain does not currently vary in form as much as the control domain, as shown in Figure 3b. The most common configuration, an expandable fabric that carries only circuit-switched voice, fits in the MC/FU region. In the case of a spatially separate or subtending fabric (i.e., well separated from the host, or central-office-resident main fabric, which we refer to as *remote*), they can be viewed as MD/FU in form.

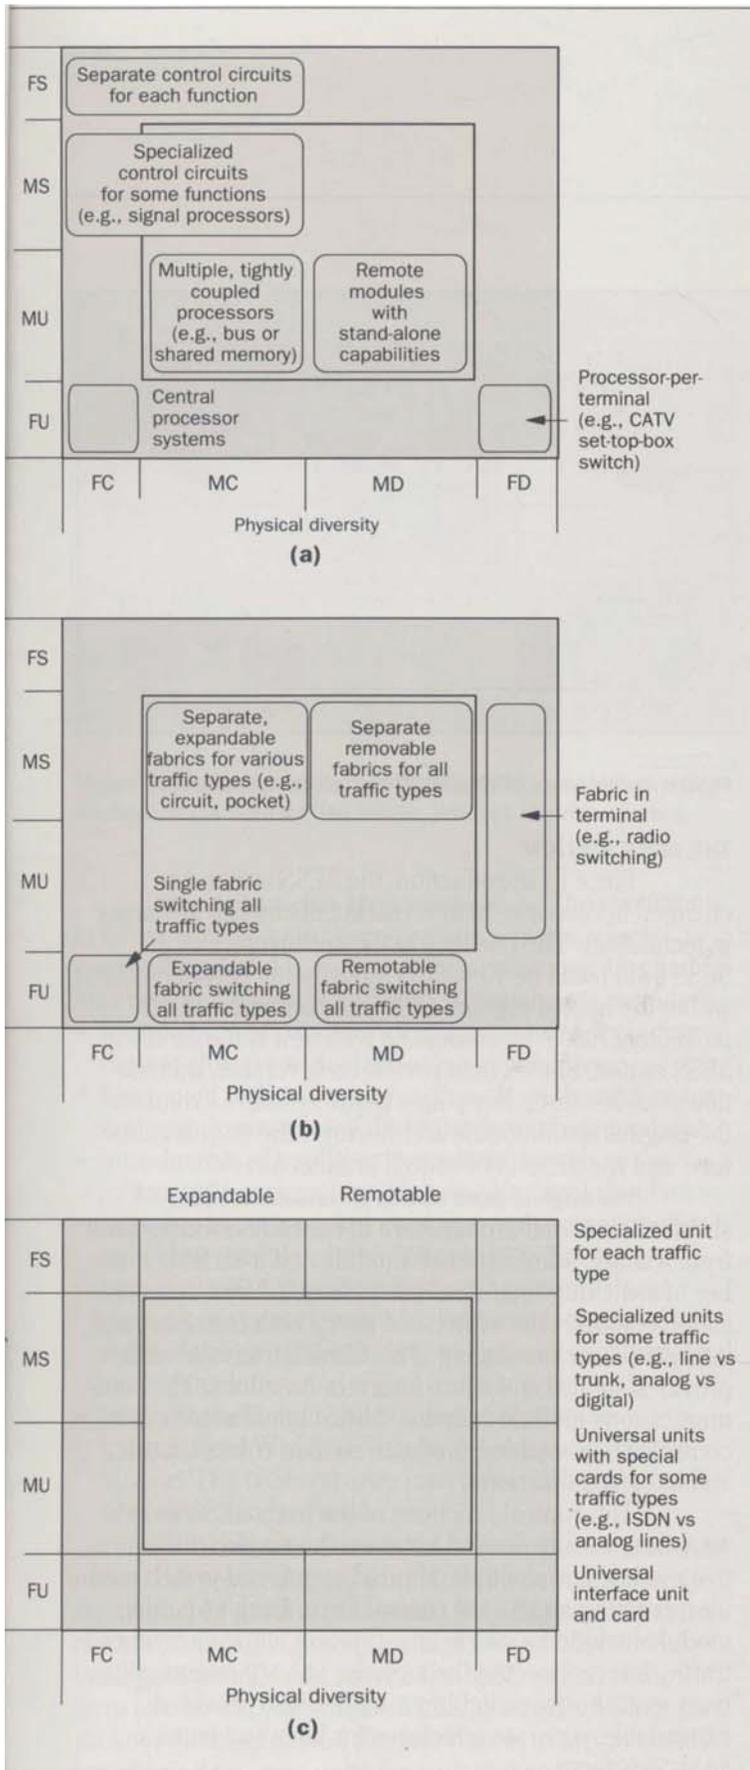


Figure 3. Examples of (a) control, (b) fabric, and (c) interface forms.

New network initiatives—such as packet switching, data, and broadband—have added diversity to fabric architecture. Systems that carry these new kinds of traffic on fabrics separate from narrowband voice may be characterized as MC/MS, or MD/MS, depending on how remote they are from the host, or central-office-resident fabric. The development of universal fabrics based on ATM technology may lead a movement back toward the MU side of the functional diversity spectrum. However, adding new ATM fabrics to existing switching systems to complement their narrowband fabrics will, initially, increase the degree of their fabric specialization. Only in *overlay*, or pure ATM, networks will fabrics become more unified.

Interface Architecture Forms. Figure 3c shows several examples of interface architecture forms. Switching systems interface to such a wide variety of signaling and transmission standards that neither FU nor FS forms are cost-effective. Hybrid, or mixed, forms permit tradeoffs: MU systems may have less overhead in the presence of different types of traffic, while MS systems can be more cost-effective when traffic is more homogeneous.

FC and FD interface configurations also appear unlikely. MC forms are common in more traditional wire-center applications, while MD forms, including Subscriber Loop Carrier systems integrated with the switch, are becoming more common.

For example, a central-office system that has separate, expandable interfaces for trunks, analog lines, and ISDN is classified as MC/MS. Another system, with remotes and the ability to connect both ISDN and analog lines to the same interface unit, might be classified as MD/MU.

CFI Continuum in Switching

The history of switching systems illustrates how the CFI continuum evolves over time. Technology advances expand its size and diversity, allowing new forms and more options in system design, while new user needs are the catalyst for expanding the domains of the continuum.

Historical Perspectives. In the early years of automatic switching, options were limited by the technolo-

gies available, and by the definition of telephony in terms of narrowband voice services. Because only one type of traffic existed, no distinction was made between specialized and universal fabrics or interfaces.

The differentiation between trunks and lines expanded the interface domain, and the development of common-control systems expanded the control domain. The use of general-purpose computers, particularly microprocessors, has further expanded the control domain, and common channel signaling (CCS) and ISDN have expanded the interface domain. The move toward remotes has also expanded the physical dimension of all domains. Most recently, the emergence of new kinds of traffic, such as packet data and broadband, have again expanded interface and fabric domains.

Examples of Switching Architecture Forms. This section describes three systems of historical importance on the CFI continuum: the No. 4A crossbar, the 1/1A ESS™ switch, and the No. 101 ESS switch. Subsequent sections detail the evolution of two additional switches, the 5ESS® switch and the GlobeView™-2000 Broadband Switching System (BSS).

The No. 4A crossbar switch⁶ used a highly refined system of electromechanical common control, a centralized complex that included special units for each switching function (registers, senders, markers, a router, and so on), giving its control domain an FC/FS form. Both the fabric and interfaces were MC/MS, expandable but not remotable, and optimized for trunk-switching applications.

As its control system, the No. 1 ESS, the first commercial stored program control switch,⁷ had a single-processor FC/FU configuration, which evolved to FC/MS by adding a signal processor and Attached Processor System (APS). Its fabric form—MC/MU, and expandable but not remotable—enables it to carry all types of traffic. The system's interfaces are more diverse; its distinct line and trunk units, tone and pulse signaling, and CCS interface through the APS, give it an MC/MS interface form.

The No. 101 ESS (also described in Reference 7) introduced the concept of remoting to electronic switching. This switch used an FC/FU control form similar to the No. 1 ESS, but its remotable modules, including fabric and interfaces, gave both the fabric and interface domains an MD/MU form.

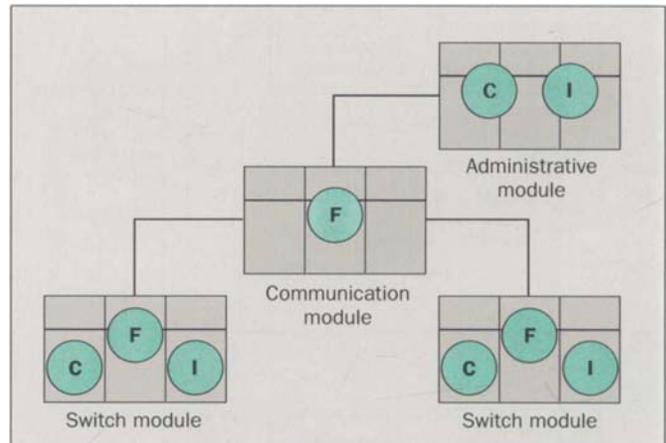


Figure 4. Overview of the 5ESS switch architecture.

THE 5ESS SWITCH

Since its introduction, the 5ESS switch has changed in response to new market needs and advances in technology. The changes in the architecture of the 5ESS switch can be viewed as movements between points within the hybrid regions of the CFI continuum. Each movement has been associated with new releases of the 5ESS switch. Rather than review each version, this section analyzes three key points in the system's evolution: the original multimodule architecture, the ISDN architecture, and the 5ESS-2000 switch architecture.^{8,9}

The Original 5ESS Switch Architecture. Figure 4 shows the general architecture of the 5ESS switch.⁹ Built from a single administrative module and a variable number of switching modules, interconnected by a communications module, the switch can carry both transport and interprocessor messaging. The administrative module provides control and some interface functions. The communications module contains fabric and, in some cases, control. The switching modules contain control, fabric, and interface functions.

The control functions of the early 5ESS switch were functionally divided between the single administrative module and multiple, similarly equipped switch modules, creating an MC/MU control form. Each switching module included a single time-division fabric carrying all traffic. Interconnected through the MC/MU communications module, the switching modules also provided expandable, separate interfaces for lines and trunks, an MC/MS form.

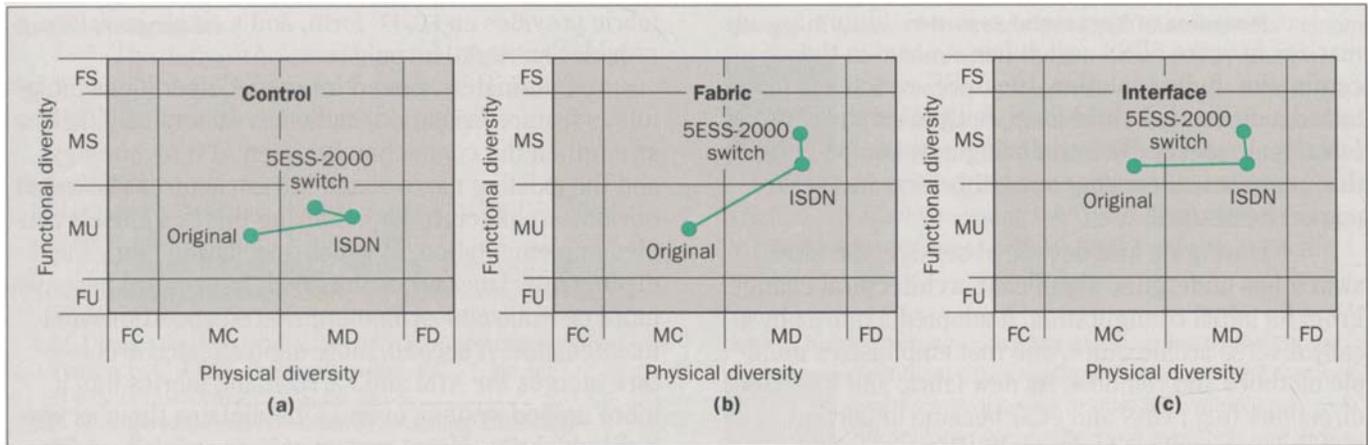


Figure 5. Evolution of the 5ESS switch CFI continuum in the context of (a) control, (b) fabric, and (c) interface forms.

The 5ESS Switch ISDN Architecture. The evolution to ISDN involved a number of enhancements, which gracefully evolved to significantly larger capacities and improved reliability. Significant enhancements included:

- Introduction of remote switching modules, including remote clusters with stand-alone capabilities;
- Movement of line-routing functions from the administrative module into a set of designated switching modules;
- Introduction of intelligent peripherals, such as the Integrated Services Line Unit and Protocol Handler;
- Introduction of a packet switching unit provisioned in switching modules that terminate ISDN lines;
- Introduction of Signaling System 7 (SS7) interfaces through the packet switching unit in international versions of the switch, including an optional capability to support an integrated signal transfer point; and
- Implementation of software that supports the International Wireless Standard, Groupe Speciale Mobile. The GSM software interfaces with remote cell sites over standard links.

The introduction of stand-alone remote modules moved the physical aspects of all three domains into the MD region. Migration of function from the administrative module to switching modules, and the introduction of intelligent peripherals, moved the 5ESS switch control form into the MD/MU region. Adding a distinct fabric for packet switching moved the fabric form to MD/MS, and introducing separate units for ISDN and analog lines

moved the interface form closer to the specialized end of the MD/MS region.

The 5ESS-2000 Switch Architecture. Initially introduced in 1993, the 5ESS-2000 switch architecture couples switching and transmission functions to cost-effectively support synchronous optical network/synchronous digital hierarchy (SONET/SDH) interfaces, reduce the cost of wire center operations, and provide a platform for future technologies and services. Significant evolutionary changes in the 5ESS-2000 switch include:

- A new, expandable, advanced technology switching module, featuring a more powerful processor, expandable fabric and interface units, and SONET/SDH interfaces;
- Development of switch software to facilitate intelligent network control; and
- An expandable, optional digital networking unit, which provides an integrated SONET and, later, SDH interface.

The introduction of a larger switching module reduced the number of modules in a switch of a given size, leading in all domains to a slight movement away from the more decentralized end of the MD region. Otherwise, the control domain remains unaffected at MD/MU. The fabric domain moves farther into the specialized part of the MD/MS region. This is a result of the multiple fabric types being supported; adding new interfaces also moves the interface form closer to the specialized part of the MD/MS region. ATM and SONET/SDH have the potential to be universal fabrics and interfaces. If the market moves to these standards, the 5ESS-2000 fabric and interface subsystems could be configured in the MD/MU region for some applications.

Measures of Successful Evolution. Figure 5 summarizes how the 5ESS switch has evolved in the CFI continuum. In its evolution, the 5ESS switch has maintained a distinctly hybrid form, both in its physical and functional aspects. By introducing new components that complement existing ones, this form has served market needs well.

During its first decade of service, the 5ESS switch has undergone significant architectural change. From its initial configuration, it adopted a more physically diverse architecture, one that emphasizes multiple modules and remotes. As new fabric and interface directions (e.g., ISDN and CCS) became important, the switch became more functionally diverse, providing new functions while protecting the existing base of installed switches. In this process, four of the five key network initiatives described earlier—ISDN, intelligent network, wireless access, and SONET/SDH interfaces—have been incorporated into the 5ESS switch architecture. By the year 2000, 5ESS switch-based ATM switching entities will likely be offering broadband services.

Initially, broadband services will be part of an overlay ATM network. ATM permanent virtual circuits are part of AT&T's GlobeView-2000 GCNS Broadband System. Architecturally, the system contains an ATM fabric and a service management module, which provides operations, administration, maintenance and provisioning. The speed of the ATM fabric can be set at either 2.4 gigabits per second (Gbit/s) or 20 Gbit/s, and has a duplex fault-tolerant architecture. It terminates and supports ATM, frame relay, Switched Megabit Digital Service, and STM clear channel.

The GlobeView-2000 GCNS can also be configured for switched broadband services by adding a switched services module, which provides both switched services and signaling. The switched virtual circuit system inherits all the functionality of the permanent virtual circuit system, and adds switched broadband services. The switched services module supports real-time, on-demand broadband call setup and dynamic allocation of bandwidth. This arrangement is called the GlobeView-2000 BSS. Its control is divided between the switched services module and the service management module, giving an MC/MU form. In the initial application, a single, non-expandable ATM

fabric provides an FC/FU form, and a single interface provides an FC/FU form.

Ultimately, a need for unified operations, ubiquitous feature execution, and other factors will likely strengthen the connection between ATM technology and the existing narrowband infrastructure. The most obvious architecture for achieving this is a loosely coupled implementation, in which specialized fabrics and distinct interfaces are maintained, but control becomes more centralized—promoting shared operations and maintenance. A second, more sophisticated architecture merges the ATM and narrowband fabrics into a more unified system, even as it maintains them as specialized fabrics. Here, control of both specialized fabrics becomes common for at least operations and maintenance functions. In some cases, certain customer feature software may also share common control. In this arrangement, interfaces still remain largely distinct, for both narrowband and broadband capabilities. As the deployment of broadband services increases, so will the desirability of the more sophisticated architectures for integrating ATM with narrowband switching systems. Today's digital switches will thus evolve through a series of architectural changes as the broadband market develops.

In the context of Figure 5, such speculation about the ultimate integration of ATM with the 5ESS switch may drive evolution in a number of possible directions. The precise placement of such points in these domains will be the subject of a future paper.

Summary

This paper presents a model for classifying the architecture of switching systems. In the context of this model, it describes the evolution of some historical switches in general, and of the 5ESS switch in more detail. Today, a variety of architectural forms exist among the many switching systems available in the marketplace. In general, hybrid switching system arrangements in the control, fabric, and interface domains, with balanced levels of functional and physical diversity, seem to be the most practical, flexible, and evolvable forms. Among these forms are many well suited to deal with today's trends toward the use of new technology, implementing increasingly more complex, advanced features.

Acknowledgments

The authors acknowledge the contributions of S-J. Chang, J. I. Cochrane, D. C. Starr, G. H. Zimmerman, and M. J. Zola to the content of this paper.

References

1. "Intelligent Networking: Network Systems," *AT&T Technical Journal*, Vol. 70, No. 3-4, Summer 1991.
2. P. K. Prasanna, R. Levy, and J. Swenson, "Principles and Standards for Broadband ISDN," *AT&T Technical Journal*, Vol. 72, No. 6, November/December 1993, pp. 9-22.
3. M. Demaio and H. J. Kafka, "Broadband Applications and Services Prospectus," *AT&T Technical Journal*, Vol. 72, No. 6, November/December 1993, pp. 23-28.
4. L. G. Anderson, et al., "Distributed Switching Tradeoffs," *Proceedings of the 1987 International Switching Symposium*, paper B1.1, Vol. 1, pp. 26-33.
5. S.-J. Chang, et al., "Evolving Switching Systems into the Information Age," *Proceedings of the 1992 International Switching Symposium*, paper P21, Vol. 1, pp. 263-267.
6. R. Chapuis, *100 Years of Telephone Switching, Part 1: Manual and Electromechanical Switching*, Amsterdam, North-Holland, 1982.
7. R. Chapuis and A. Joel, *100 Years of Telephone Switching, Part 2: Electronics, Computers and Telephone Switching*, Amsterdam, North-Holland, 1990.
8. D. L. Carney, et al., "Architecture Overview: The 5ESS Switching System," *AT&T Technical Journal*, Vol. 64, No. 6, Part 2, August 1985.
9. J. L. Johnson, et al., "AT&T's 5ESS Switch: Early Experience and Details on Planning and Implementing a Large Scale ISDN Corporate R&D Network," *Proceedings of the 1990 International Switching Symposium*, Paper B4.1

(Manuscript approved September 1994)

E. Lee. Sneed is technical manager of the System Architecture Planning Group in AT&T Bell Laboratories in Naperville (Indian Hill), Illinois. His responsibilities include network services architecture, wireless and personal communications services architecture, and U.S. application architecture. He holds a B.S.E.E. from Oklahoma State University, Stillwater, and an M.S.E.E. from the University of Nebraska at Lincoln. Mr. Sneed began his career at AT&T in 1979.



George D. Huensch is director of 5ESS® Switch Global Platform



Architecture and Performance at AT&T Bell Laboratories in Naperville (Indian Hill), Illinois. He is responsible for the system, hardware, and software architecture, as well as the performance and capacity, of the 5ESS switch globally. Mr. Huensch joined AT&T, in 1977, after receiving both his B.S. and M.S. in computer science from Purdue University, West Lafayette, Indiana.

John J. Kulzer, Network Systems architect within the Group



Technical Officer's Organization of AT&T Network Systems, Naperville (Indian Hill), Illinois, is responsible for architectural evolution and cross-product/offer synergy planning within AT&T Network Systems. Mr. Kulzer received B.S., M.S., and Ph.D. in electrical engineering, all from the Illinois Institute of Technology, Chicago. Mr. Kulzer joined AT&T Bell Laboratories in 1969.

Harry Moerkerken is director of Public Switching Development and Information Technology at AT&T-Network Systems International (NSI) in Hilversum, the Netherlands. He joined AT&T in Hilversum in 1983, where he worked in the Development Department of AT&T and Philips Telecommunication Joint Venture. Mr. Moerkerken spent two years at AT&T Bell Laboratories in Naperville, Illinois, developing the 5ESS® switching system. After his return to Hilversum in 1988, he became head of the Public Switching Development Laboratory. In 1991, he assumed the position of director of Product Management Switching in the marketing department of AT&T-NSI. Recently, he became director of Public Switching Development and Information Technology for AT&T-NSI. Mr. Moerkerken received a Ph.D. in molecular physics from the Catholic University, Nijmegen, the Netherlands.