

FASTARSM Platform Gives the Network A Competitive Edge

Chin-Wang Chao
Gene Fuoco
Diane Kropfl

The FASTARSM Operations System Platform is a set of systems that provide AT&T's fast automated restoration capabilities. Since its introduction in March 1992, this platform has provided AT&T with dramatic improvements in network reliability. This paper describes the FASTAR platform architecture and restoration process. It then highlights three key factors in overcoming the many software engineering challenges posed by the development of the FASTAR Operations System Platform: effective software development strategies, effective implementation strategies, and teamwork.

Introduction

AT&T's mission is to be the world's best at bringing people together—giving them easy access to each other and to the information and services they want—any time, any place. While consumers expect the highest standards of quality and reliability in placing their calls, service outages are a fact of life for all vendors, sometimes impacting business customers' operations severely and often frustrating residential customers. AT&T's commitment to network reliability is met through:

- Programs that reduce the frequency of network outages, and
- Capabilities, such as the FASTARSM platform, that automatically detect and correct network problems.

There has been a 10-fold improvement in AT&T's network reliability, with significant reductions in blocked calls and incidences affecting customers, since the introduction of the FASTAR platform three years ago.

With the advent of fiber optic technology, the network capacity has increased. Unfortunately, as a result, so have the consequences of cable cuts. The FASTAR platform improves network reliability and offers a significant strategic advantage for AT&T. Before the introduction of the FASTAR platform, fiber cable cuts were potentially disastrous events.

A fiber cut in the Newark, New Jersey area in November 1988, for example, caused the loss of over 270 DS3 (45 Mbits/s)

paths—or more than 180,000 telephone circuits. Since this was in the pre-FASTAR era, it took more than 15 hours to repair the cable and re-establish service. The FASTAR platform, in contrast, alleviates the impact of these cuts by reducing— from hours to minutes—the time it takes to restore service. During the record flooding that occurred in the summer of 1993, a major backbone cable between Kansas City and St. Louis washed away where it crosses the Missouri River. With the FASTAR technology, it took less than five minutes to restore more than 250 DS3 paths.

Architecture Overview

This section first presents an overview of the FASTAR platform architecture.¹ In the next section, "Restoration Process," we describe the FASTAR platform restoration process. The FASTAR platform architecture consists of the FASTAR Operations System Platform, transmission equipment, restoration test equipment, and a communications network interconnecting the digital facility offices to the central computer sites (see Figure 1).

FASTAR Operations System Platform.

The FASTAR Operations System Platform consists of a set of computer systems that work in concert to manage restoration operations for the AT&T transport network. The primary systems in this platform are:

- The Restoration and Provisioning Integrated Design (RAPID), a central processor that receives and correlates alarm reports, calculates re-routes, controls restoration

Acronyms and Terms Used in This Paper

ASCII—American Standard Code for Information Interexchange
CADS—Central Access and Display System
DACS—Digital Access and Cross-connect System
DEMS—DACS Element Management System
DS3—45 Mbit/s digital signal
DSL—Digital Services Laboratory
DSX—Hardwired, manual digital signal cross-connect frame
EMS—Element Management System
G2 data language—An applications technology that describes complex records as ASCII data
GTP—General telemetry processor
IP—Internet protocol
LTE—Line terminating equipment
TMP—Transport maintenance platform
PINnet—Packet internal network
PMT—Process management team
RAPID—Restoration and Provisioning Integrated Design, the central restoration processor
RTE—Restoration test equipment
RNC—Restoration node controller
RNC-EMS—Restoration Node Controller-Element Management System
SFV—Service failure verification
SRE—Software reliability engineering
SRLG—Shared risk link group
SV—Service verification
T3P, T3R—Protection and restoration channels
TMP—Transport Maintenance Platform
TSID—Test signal identification
TSNOW—Transport Service NOW
X.25—CCITT packet-switching protocol at link and packet layers.

path implementation, directs and coordinates path assurance tests, and supports the roll-back to the original transmission facility after the failure is repaired. During the provisioning process, the Transport Service NOW (TSNOW) platform provisions RAPID with facility layouts, including cross-connect data. During the restoration process, RAPID dynamically creates facility paths to reroute around failed facilities.

- Restoration Node Controllers (RNCs), central office-based intelligent processors for local control of alarm gathering and reporting, and for test and transfer functions. RNCs are deployed, one per node, in all of the primary facility offices in the fiber network. RNCs monitor transmission equipment for alarms and status

changes, filter the data, and report autonomously to RAPID via two X.25 networks, the Packet Internal Network (PINnet) or the Skynet® Clearlink Service network. PINnet—a derivative of Accunet® Packet Service—is a terrestrial network, while Skynet Clearlink Service is satellite based. In addition, RNCs provide local control of continuity tests as needed for path assurance.

- The Transport Maintenance Platform (TMP), among its other operations functions, ensures that traffic is switched to protection channels on the FT Series G fiber transmission facility via commands sent through the general telemetry processor (GTP). A RAPID-to-TMP link provides RAPID with the needed control of this function.
- The Central Access and Display System (CADS), a graphical access system that allows technicians to access RAPID and other systems in the the FASTAR Operations System Platform.

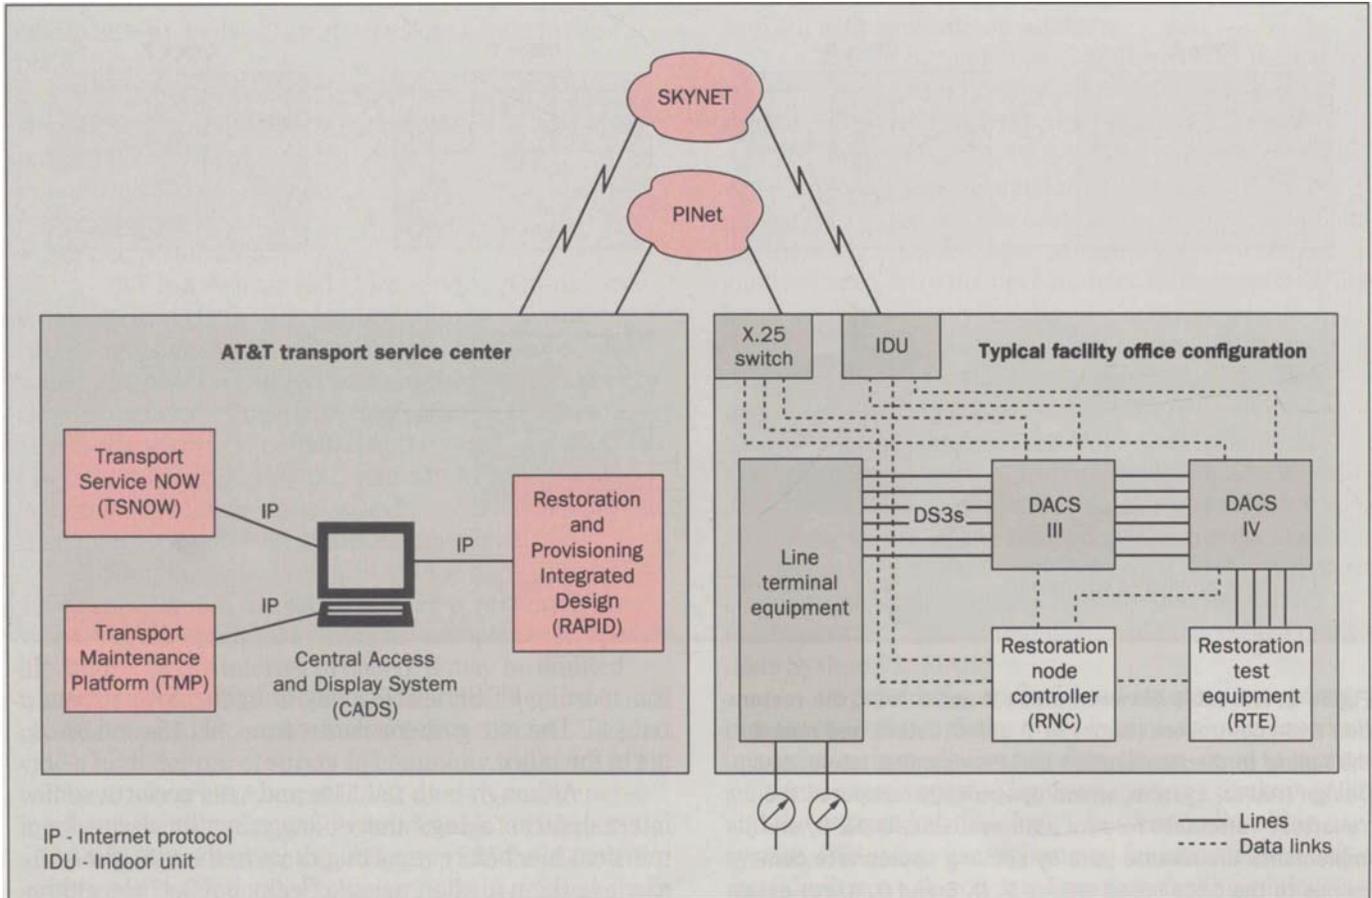
Transmission Equipment. Transmission equipment that comprise the backbone of the AT&T transport network and play a significant role in the FASTAR platform are:

- FT Series G, the predominant fiber transport technology in the AT&T network today. The FT Series G line terminating equipment (LTE) allows immediate access to protection channels for restoration. The LTES also report unprotected transmission failures between LTES, as well as any changes in protection channel status to the RNC.
- Digital Access and Cross-connect System (DACS III), which makes automated restoration possible by providing the remotely controllable, cross-connect switching capability of DS3s, for both path implementation and test access. In addition, the DACS III is responsible for reporting intra-office transmission failures to the RNC.

Test Equipment. The test equipment used for the FASTAR platform restoration is:

- The Restoration Test Equipment (RTE), containing DS3 test signal generators and receivers with a test signal identification (TSID) capability. Its primary function is to provide the means to perform fail-safe continuity tests.

Communications Network. The necessary high reliability and availability of the network between RNCs and DACS IIIs and RAPID is achieved through full redundancy, using two totally diverse, parallel X.25 packet-switched networks, the PINnet and Skynet Clearlink Service networks. Either, in the absence of the other, can support the entire communications burden during any



major restoration event. In addition, the Internet protocol (IP) network provides reliable communications among RAPID and the other systems of the FASTAR Operations System Platform.

Restoration Process

The restoration process consists of three major steps:

1. Gathering and reporting alarms.
2. Implementing restoration paths, including reroute computation.
3. Assuring paths, including service failure verification (SFV) when required, continuity test if needed, and service verification (SV).

Restoration Process Scenario. The FASTAR platform restoration capability is shown in Figure 2. If a cable between office B and C fails, the RNCs at B and C detect and report the failure to RAPID. After RAPID dynamically

Figure 1. The FASTAR Operations System Platform includes the Restoration and Provisioning Integrated Design (RAPID) system, the Transport Maintenance Platform (TMP), the Transport Service NOW (TSNOW) platform, and the Central Access and Display System (CADS); the transmission equipment in the central office, plus the restoration test equipment (RTE) and restoration node controller (RNC); and two communications networks, Skynet® and PINnet.

computes the “shortest” alternate reroute path available, B-D-E-C, RAPID implements the reroute path by sending appropriate commands to DACS IIIs at offices B, D, E and C. Meanwhile, RAPID establishes the backup route via restoration and protection channels on FT Series G between offices B-D, D-E, and E-C. Protection channels are reserved for RAPID by sending commands via the TMP platform (a procedure that is also called protection switch lockout).

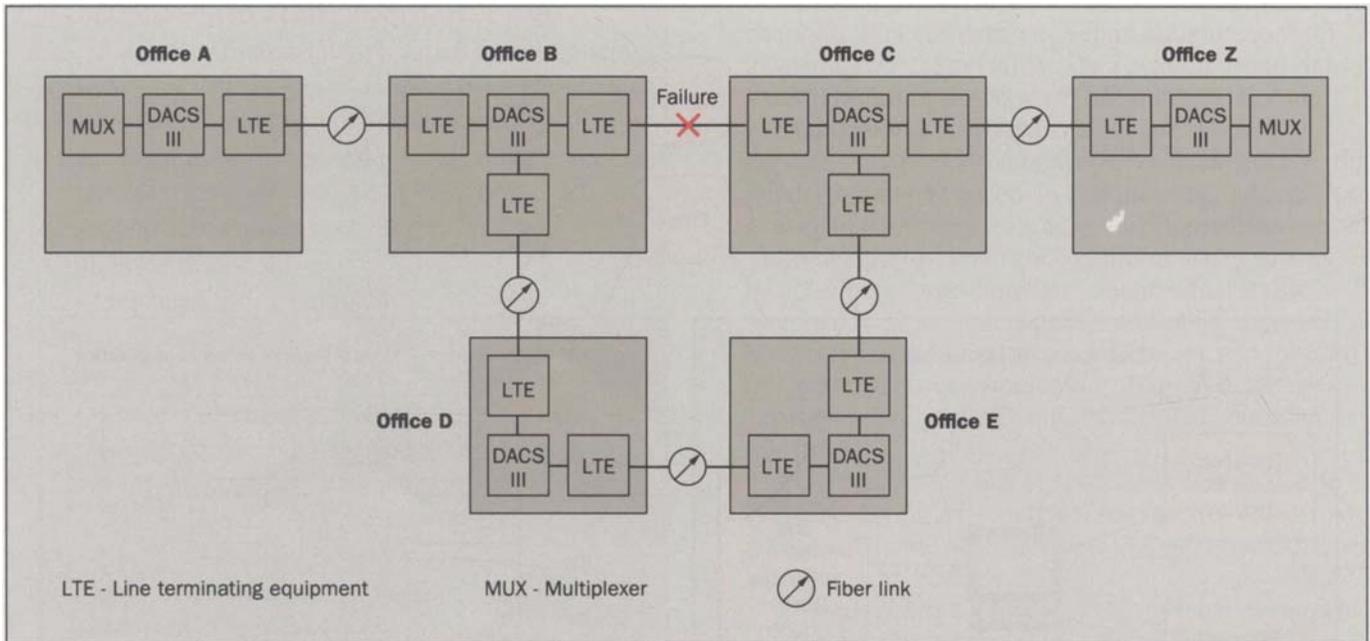


Figure 2. If a cable between office B and C fails, the restoration node controllers (RNCs) at B and C detect and report the failure to the Restoration and Provisioning Integrated Design (RAPID) system, which dynamically computes the “shortest” alternate reroute path available, B-D-E-C, and implements the reroute path by sending appropriate commands to the DACS IIIs at offices B, D, E and C. RAPID establishes the backup route via restoration and protection channels on FT Series G between offices B-D, D-E, and E-C. Protection channels are reserved for RAPID by sending commands via the Transport Maintenance Platform (TMP).

As soon as the reroute path is completed, RAPID transfers service from the original path to the reroute path by sending “roll” service commands to DACS IIIs at offices B and C. These commands, in effect, roll over the service from the original channel to the protection or restoration channel. Upon completing the service transfer, RAPID initiates path assurance by sending instructions to the RNCs at A and Z. The RNCs respond by commanding associated DACS IIIs to bridge the failed DS3 to an RTE test signal receiver. This verifies that the failure has been cleared. The RNCs then report the SV outcome to RAPID.

Alarm Gathering and Reporting. DS3 signal failures between the output of one DACS III and the input of the next downstream DACS III are detected either by the

transporting FT Series G facility or by the downstream DACS III. The RNC gathers alarms from all LTEs and DACS IIIs in the office.

Although both DACS IIIs and LTEs perform some internal alarm “aging,” that is, analyzing the alarms for transient hits before reporting them to the RNC, the RNC may age them further, using a “leaky bucket” algorithm. The RNC then correlates the alarms and reports them to RAPID. A similar process is used to apprise RAPID of any changes in the status of protection channels.

Path Implementation. Because failure reports for different parts of the same failure event may be received from multiple RNCs, RAPID, upon receiving the first report, initiates a sliding “event window” to ensure it has received the complete report before proceeding further. It then prioritizes the list of failed DS3s and, in priority order, calculates reroute paths for them. The real-time router program dynamically determines the alternate routes for each failed DS3, based on the shortest available restoration path and the available restoration capacity in the network at the time of the failure.

RAPID next begins the process of implementing connectivity along the calculated restoration paths. This entails sending cross-connect commands to the DACS IIIs using Skynet and PINnet, and protection switch lockout commands to the TMP platform. Finally, RAPID issues

commands to the DACS IIIs to transfer service to the restoration paths.

Path Assurance. Path assurance is a necessary step in the restoration process. It comprises up to three major steps:

- Service failure verification,
- Continuity test, and
- Service verification.

SFV is a step taken before service is transferred to the alternate path. It is performed to avoid breaking a "good" DS3 signal, even momentarily. In the interest of speed, SFV may be omitted when network database accuracy is regarded as nearly perfect. After a service transfer, SV is necessary to be certain that the reroute succeeded. For either SFV or SV, the RNC commands the DACS III to bridge the incoming signal to an RTE receive test port, and then queries the RTE for status on that signal.

Continuity testing, prior to a service transfer, is an essential part of the process when rerouting live service around a hazard condition, in anticipation of a possible failure. In the interest of speed, it may be omitted when the service has already failed, since the final SV will provide nearly the same assurance after the fact. To provide a high degree of surety in continuity tests, the RTE will have two primary capabilities: a test signal generating feature with the capability to transmit a unique TSID, and a test signal receive capability to read the TSID. RTEs at both ends of the path exchange TSIDs in a "hand-shake" operation, and each reports test status to its local RNC.

FASTAR Platform Software Development Strategies

The FASTAR platform software development strategies include:

- Making effective technology choices,
- Following a standard software development process,
- Using software reliability engineering (SRE) techniques,
- Certifying the platform with real network equipment,
- Providing FASTAR platform simulation and dry-run capabilities, and
- Developing network management capabilities.

FASTAR Platform Technology Choices. The choice of technology for the FASTAR Operations System Platform was influenced by the need to develop a complex, high-reliability, real-time, communications-intensive networking application. Essential components in developing this platform were the hardware and system software,

coupled with applications software.

Vendor partnerships established with the RAPID and RNC development organizations formed a base for successfully developing the platform. Hewlett-Packard was the computer and system software vendor for RAPID, while Motorola was the vendor for the RNC. High-level management support was a key to creating and maintaining these successful vendor partnerships. Our vendors jointly shared, with the development team, responsibility for delivering a highly reliable platform.

Managing the complexity of the system was a critical challenge for the FASTAR applications development team. Both RAPID and the RNC were developed as an application package on top of the UNIX* operating system, using the C and C++ programming languages. Both RAPID and RNC used in-memory databases to quickly and effectively access information about the current state of the network. Both RAPID and RNC relied on the G2 data language, an applications technology that helped the developers manage the complexity of the FASTAR Operations System Platform.²

G2 provided a readable and efficient representation for the FASTAR platform application data by describing complex records as ASCII data. G2 was used to represent data communicated between systems, such as RAPID and the RNC, and also to depict data internally within a system. Client/server applications for event-driven software, built upon the G2 data language, facilitated development and testing.

An example is the reporting of alarms to RAPID. A RAPID client subscribes to alarm reports from RNC servers. Whenever a cable cut occurs, alarm reports are transmitted in G2 format from the RNC to its subscriber clients, in this case RAPID. The alarms in G2 format can be reviewed by users or interpreted by programs to assess system performance during the cable cut.

Another example is the RNC *injection point* capability. An RNC client test program can inject data into an RNC server program at an injection point. The data that, as a result, is outputted from the server program can then be inspected for correctness, either visually or via a computer program. In summary, the G2 language, coupled with client/server event-driven software, provide a powerful foundation for managing the complexity of the FASTAR Operations System Platform.

Standard Software Development Process. The FASTAR platform team responsible for implementing the

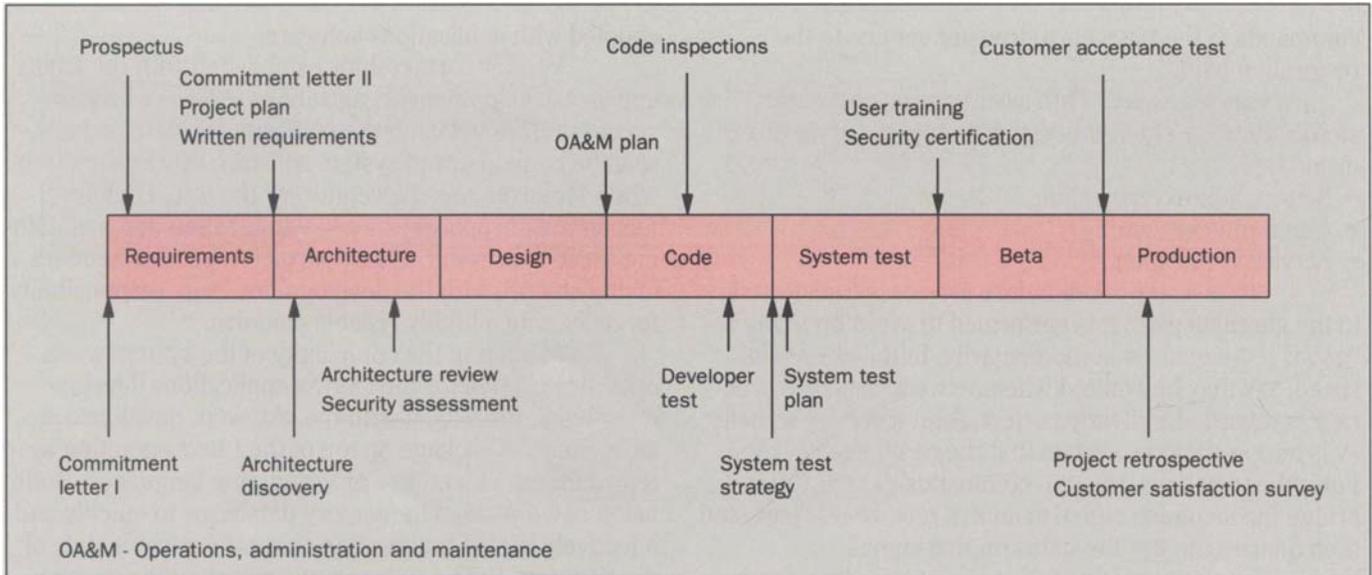


Figure 3. The FASTAR Operations System Platform follows a standard seven-phase software development process. The platform engineers wrote the feature requirements specifications, the developers wrote the architecture documents and the design documents, and the testers validated the entire platform. Beta and user acceptance tests were then performed to certify the system for production.

FASTAR Operations System Platform followed the standard software development process used throughout the Operations Technology Center in the Network Services Division. This is a formal, structured process for software development. Documented practices and procedures for the center's software projects define the steps, from requirements through production, for developing operations systems. Process management teams (PMTs) composed of technical staff and managers define, champion, and improve each of these steps in the process.

The standard process consists of the seven phases shown in Figure 3. FASTAR *systems engineers* wrote feature requirements specifications to translate into detailed requirements the FASTAR objective to restore 100 DS3s in less than five minutes. FASTAR *developers* wrote architecture documents to describe the framework for building each of the systems within the context of the requirements. They also wrote the design documents that describe each of the system's capabilities, and detail how the requirements will be implemented within the

framework of the system architectures. FASTAR *testers* validated the entire FASTAR platform, including all hardware and software components, during system test. At the completion of the rigorous system test phase, beta and user acceptance tests were performed to certify the system for production.

FASTAR Platform Software Reliability Engineering.

We applied software reliability engineering to RAPID and the RNC to ensure that these systems met stringent reliability and robustness requirements.³ SRE is a technology and practice that sets a quantitative, reliability objective for a software system, and then designs and tests to meet that objective.⁴ We found SRE instrumental to the success of the FASTAR platform by ensuring that reliability objectives for RAPID and the RNC had been met prior to introducing these systems into the AT&T network. We modeled the FASTAR SRE initiative after a successful program pioneered by the system testers for the TMP platform.^{5,6} Key steps in applying SRE were to derive an operational profile, develop a simulator test environment, and execute a load/stress/stability test program.⁷

Derive Operational Profile. An operational profile characterizes the operating conditions for the system—that is, how the software will be used in the field.⁸

We derived the initial version of the FASTAR operational profile early in the requirements phase, using a computer simulation model of the FASTAR platform. This first profile focused solely on cable cut scenarios. The

FASTAR simulator could model any potential fiber cut in the AT&T network, and was an invaluable tool for engineering and analyzing system performance. We defined the typical and worst-case cable-cut incidents from simulation results. A typical cut affects 150 DS3s, and most cuts affect less than 300 DS3s. We derived the frequency of various cable-cut scenarios from historical data, and estimated that there are, on average, about 10 major cable cuts in the AT&T fiber network a year.

As development proceeded, we recognized that in order to adequately test the system, we needed to define a richer operational profile that described all the system inputs and outputs. Background events were added to the profile and defined as all non-cable cut activities that could occur on the system. These non-cable cuts included such activities as database provisioning, maintenance operations, network management operations, and single DS3 path failures caused by component equipment failures. While, as noted, major cable cuts occur approximately 10 times during a year, background events occur continuously. We viewed the operating scope of RAPID as the entire AT&T network, and the operating scope of the RNC as one digital facility office. In order to take these important differences into account in our test program, distinct operational profiles were defined for each of the two systems.

In the operational profile for RAPID, we were concerned with the overall FASTAR platform profile, and the impact of receiving data from over 250 RNCs, maintaining communications to over 250 DACS IIIs, and processing commands from 50 simultaneous users. Special attention was paid to collecting data that described how users actually interacted with the system. For the RNC, one operational profile was developed that represented all the offices and focused on the worst case load in the largest AT&T digital facility offices.

An operational profile was derived for each of the systems, based on discussions with customers and log data collected from beta test versions of the system. Differences in the two operational profiles were captured by analyzing event frequencies. The frequency-of-event occurrences in RAPID dwarfed that of the RNC. For RAPID, on average, 10 major cable cut events are expected in a year. For the RNC, any single system can expect to experience less than one cable cut event per year. The difference in background events is even more dramatic. About 40,000 background events are executed on RAPID every

day. This is due to the many users who interact with the system for training, and who use the simulation and dry-run capabilities described in a later subsection, "Simulation and Dry-Run Capabilities." For the RNC, 500 background events are performed each day.

Develop Test Environment. Since it was not economically viable to test with the RAPID and RNC operational profile using real network facilities in a lab environment, we implemented a test environment that gave us the capability to simulate the AT&T transport network's restoration capabilities. This simulated test environment handled all the inputs and outputs identified in the operational profile, and gave us the ability to test the impact of worst-case cable-cut scenarios on RAPID and the RNC in the safety of our lab.

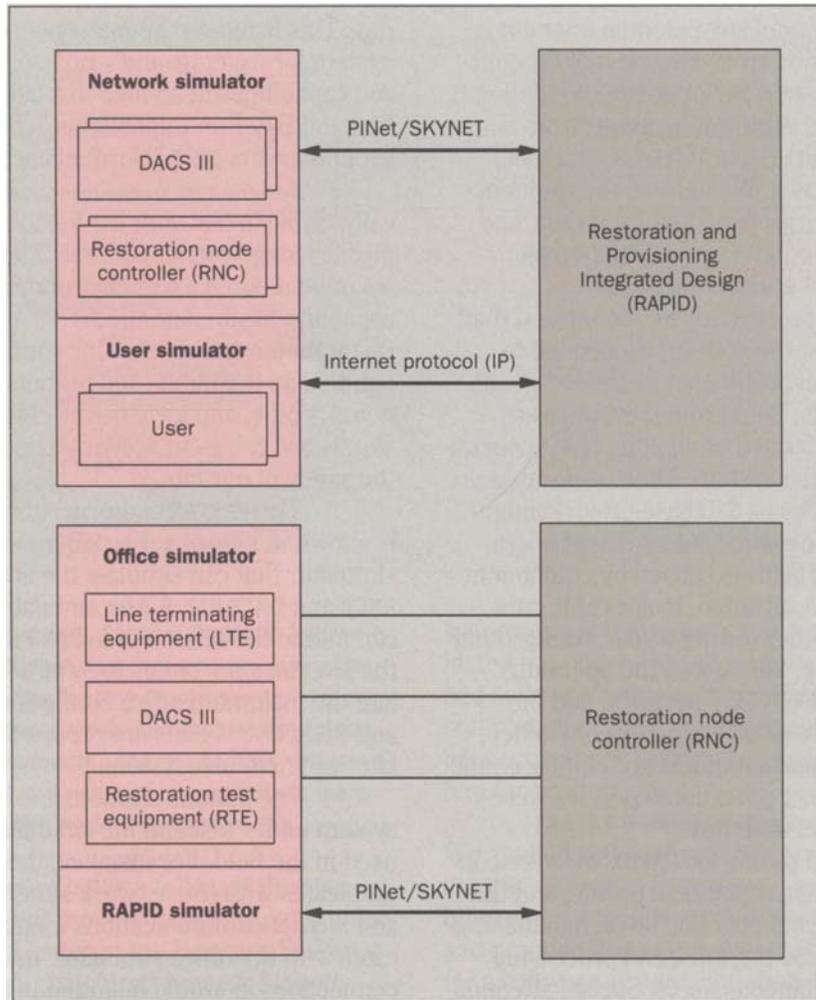
The FASTAR platform simulator test environment is shown in Figure 4. For RAPID, we developed a *network* simulator that can simulate the actions of more than 250 RNCs and DACS IIIs. A *user* simulator was developed that can mimic the actions of over 50 simultaneous users. For the RNC, an *office* simulator was developed that can simulate the maximum office configuration of LTES, DACS IIIs, and RTEs. RNC results are reported to either a RAPID simulator or RAPID test system.

The communications media used between the system under test and the simulator is the same as that used in the field. For instance, the RAPID under test communicates with the network simulator over the PINnet and Skynet communications media. The RNC communicates with the office simulator, using the same direct connections as would be found in an actual digital facility office. Both the RAPID and RNC simulation environments were designed with a programmable interface. This allowed the development of automated test scripts to drive the test runs.

For the RNC, we developed automated test scripts in the office simulator environment to simulate the RNC's operational profile. Scripts simulating cable cut alarms from LTES and DACS IIIs, for example, provide the load to test the RNC. The RNC reports alarms to a simulated RAPID, and performs service verification requests by interacting with simulated RTEs.

For RAPID, we implemented automated test scripts in both the network and user environments to simulate RAPID's operational profile. For instance, user scripts mirror the operations of the various RAPID users, and alarm scripts simulate cable cut alarms from RNCs.

Figure 4. The FASTAR platform simulator test environment includes a network simulator that can simulate the actions of more than 250 RNCs and DACS IIIs, a user simulator that can mimic the actions of over 50 simultaneous users, and an office simulator that can simulate the maximum office configuration of LTEs, DACS IIIs, and RTEs. RNC results are reported to either a RAPID simulator or RAPID test system.



RAPID interacts with the simulated DACS III network to implement cross-connects and transfer service, and also interacts with simulated RNCs to perform service verifications.

Execute Load/Stress/Stability Test Program. During our load/stress/stability test program, we tested RAPID and the RNC under the same conditions that would be encountered in the field. *Load tests* mirrored the operational profile of the two systems. *Stress tests* executed worst-case cable-cut scenarios on each of the systems. *Stability tests* verified each system, with its operational profile, for a minimum of 100 hours of continuous operations. Often, stability tests ran for more than one month.

Since it was not feasible to execute the test program for a year to simulate the expected 10 cable cuts,

we designed accelerated test runs. For the RNC, the load was accelerated by a factor of 24 during stability tests. This allowed us to execute an entire day's activities in one hour. Hence, an entire year's operational profile could be executed in the test environment in a matter of weeks. For RAPID, the load was accelerated by a factor of 1.25 during stability tests and by a factor of 50 during stress tests.

DS3 paths correctly reported by the RNC, and DS3 paths correctly rerouted by RAPID, were the metrics used for tracking results. When we started obtaining load/stress/stability test results, 90% of the DS3 paths were correctly reported and rerouted. We continued to improve the system, run tests, measure results, and fix defects until our reliability results approached 100%. At

this point, the SRE results helped guide us in making the critical decision on when to release the FASTAR platform into the AT&T network.⁹

In the second major release of the RNC, we applied experimental design techniques for measuring software reliability of systems with highly critical but infrequent events.¹⁰ The steps followed were:

- Specify distinct modes of system operation: background load only mode, alarm reporting mode, service verification mode, combined alarm reporting and service verification mode, and single DS3 failure mode.
- Apply experimental design techniques to ensure a balanced mode-execution sequence.
- Execute the software under accelerated load conditions at acceleration settings of 10, 100, and 200.
- Collect software reliability performance data, failure data, and central processing unit usage, within each mode of operation at each acceleration load setting.
- Statistically extrapolate reliability to unaccelerated load.
- Derive operational reliability by weighting the extrapolated mode reliabilities by the probability of the mode occurring during field operations.

The experimental design approach to software reliability testing proved to be a valuable technique to help carry out testing within time and cost constraints.

Digital Services Laboratory Certification. Digital Services Laboratory (DSL) certification ensured that the FASTAR Operations System Platform would work with real network equipment. The purpose of this testing was to verify the overall operation of the FASTAR platform architecture. The DSL group performed this testing in its lab and during beta test in the field.

The laboratory test bed included real network equipment, such as LTEs and DACS IIIs, and was used to perform both integration and system tests. Using real equipment in a lab environment ensured that as much testing as possible could be performed prior to the field tests.

Live field tests were conducted by the DSL group using a test bed of seven offices with reserved facilities in Florida: actual DS3 facilities that did not carry live traffic. Since the Florida test bed was designed so that it could be operated and managed remotely, the DSL group was actually able to trigger the FASTAR platform restorations on the reserved facilities from a lab in Holmdel, New Jersey. The group did this by "breaking" the fiber with a device that remotely cuts off the optical transmission in a fiber. Problems that could only be identified in an actual

office environment were found and fixed during this stage of testing.

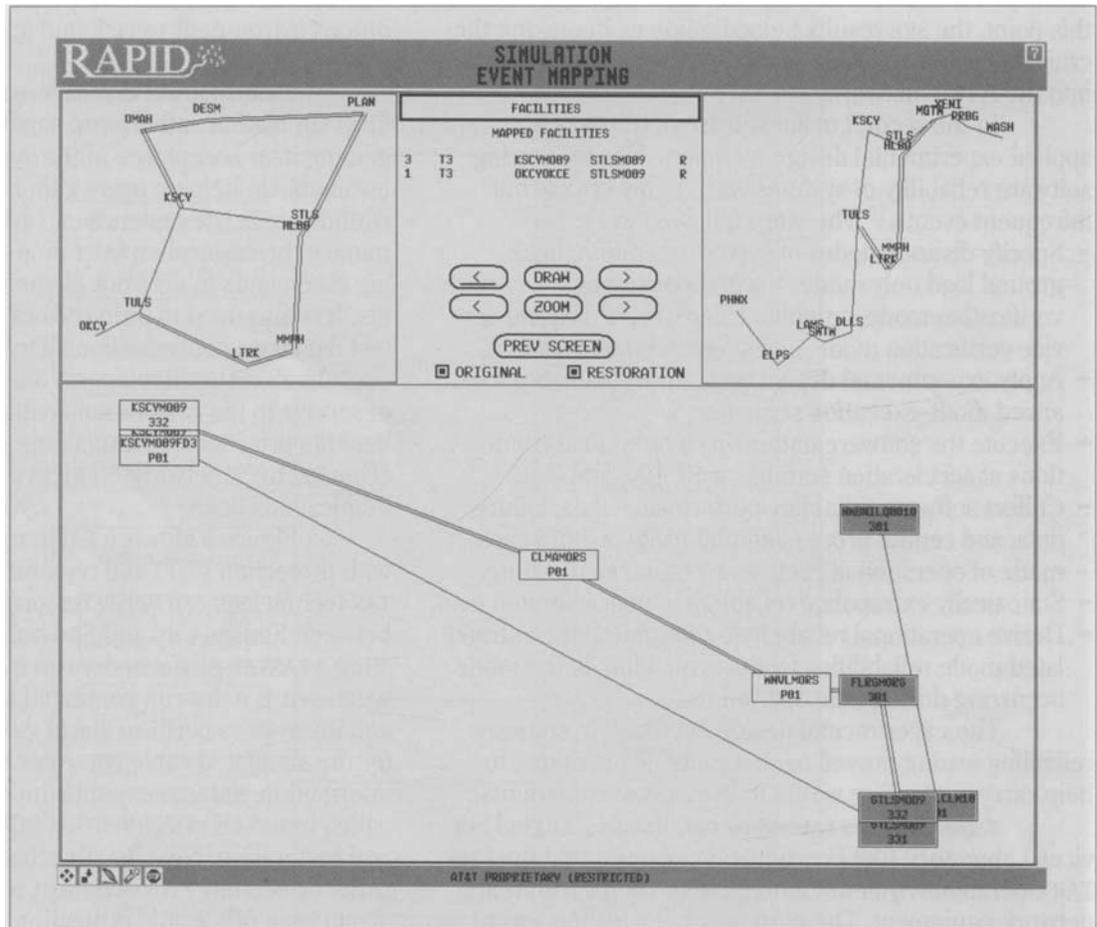
Simulation and Dry-Run Capabilities. FASTAR platform simulation and dry-run capabilities were crucial to gaining user acceptance of the system. These tools were invaluable in helping users gain hands-on experience and confidence in the system's capabilities. A simulation mimics the restoration operations without actually issuing commands to network elements, such as the DACS IIIs. It is also used to train restoration technicians, and to test data base configurations. Dry-runs can be used to execute all restoration operations, except for the transfer of service to the restoration facility. In addition to the benefits provided by simulations, dry-runs allow technicians to check network element configurations and communications links.

Figure 5 shows a RAPID restoration network map with protection (T3P) and restoration (T3R) channels. FASTAR technicians can verify restoration capabilities between Kansas City and St. Louis, for instance, by initiating a FASTAR platform dry-run from a RAPID screen. RAPID sends a dry-run command to the appropriate RNCs, and these RNCs perform alarm gathering and reporting for the simulated cable cut. After receiving the failure information, RAPID computes the optimal DS3 restoration paths, based on priority order, just as it would during a real restoration. Next, RAPID actually implements these paths by sending cross-connect commands to the DACS IIIs, in each office, and protection switch lockout commands to the TMP platform. However, no actual service transfer is performed to reroute traffic. The facility management group reviews the results of the restoration and takes action on any unexpected events. This provides a proactive mechanism for monitoring the FASTAR platform restoration capabilities.

FASTAR Platform Network Management. An effective network management platform for maintaining the many FASTAR components is essential to the ongoing success of the FASTAR platform. The basic network management design strategy is to consolidate and correlate the status and performance data supplied by various system elements for centralized display, and to isolate problems to root cause through fault-management analysis. In addition, the restoration process team uses the performance data to fine tune the FASTAR platform performance.

The FASTAR network management architecture consists of a set of element management systems

Figure 5. This illustration shows the simulation of a cable failure between Kansas City and St. Louis on a terminal screen of the Restoration and Provisioning Integrated Design (RAPID) system. Single lines depict service DS3s and double lines show the rerouted traffic on restoration and protection channels. The large lower map shows the reroute of Kansas City-St. Louis traffic through three digital central offices, CLMAMORS, WNVLMOORS, and FLRSMOORS. The two upper maps show the reroute paths of DS3 traffic originating from points beyond Kansas City and St. Louis.



(EMSS), which collectively monitor and manage the network elements and their connecting links. Two of these are the DACS Element Management System (DEMS) and the RNC Element Management System (RNC-EMS). The use of DEMS and RNC-EMS to centrally monitor the DACS IIIs and RNCs was critical to establishing user confidence prior to the cutover of the FASTAR platform into the AT&T network. Also, in addition to their primary roles in the FASTAR platform, RAPID and TMP each provide element manager functions. RAPID periodically polls each DACS III to check the health of PINnet and Skynet links, and to verify its ability to log in to the DACS III. TMP gathers transmission equipment alarms.

DEMS is an EMS designed to assist the Transport Service Center with the management of DACS IIIs for the FASTAR platform. RAPID forwards any pertinent status change information to DEMS, which triggers DEMS to

isolate the problem by collecting and heuristically correlating alarm data from several sources. The fault management feature in DEMS is an expert system that correlates alarms, telemetry, and diagnostic queries. Its function is to present to the user the root cause of a failure and a prescription for its resolution.

RNC-EMS is an EMS designed to centrally monitor the health of the RNCs and their connections to office equipment. The RNC application monitors its communication with local equipment, such as DACS IIIs, LTES, and RTEs. The RNC-EMS monitors the RNCs directly, over PINnet and Skynet, and indirectly monitors the LTE, DACS III, RTE, and their links to the RNC via components residing in the RNC. Among the functions that RNC-EMS support are data collection, alarm collection and thresholding, performance metrics, remote software download, and administration.

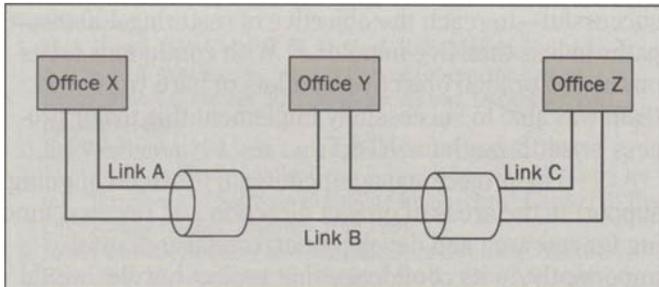


Figure 6. A cable cut between two offices will be reported if the FASTAR platform alarm reporting covers at least one end of every system riding the span. A cut on Link A between offices X-Y would be covered if office X was maintained by the FASTAR platform, even if offices Y and Z were not. If a cut occurred between offices Y-Z, the DS3s on link B would be reported by the FASTAR platform, while those on link C would not.

FASTAR Platform Implementation Strategies

Many hurdles had to be overcome to build this complex platform. To minimize the risks and take advantage of the opportunities for continuous learning and improvement, the FASTAR platform was developed and deployed in phases.

Phased Development. FASTAR platform development occurred in two phases, proceeding from the *mechanized* restoration phase to the *fully automated* restoration phase. Prior to the FASTAR platform, technicians performed restorations manually at hardwired cross-connect systems known as DSXs.

In the first phase, DACS IIIs were deployed throughout the network to replace most DSXs. Technicians began using the manually triggered restoration capabilities of the RAPID system by initiating restorations for DS3 failures where DACS III systems were deployed. As major restorations were performed successfully, the project members gained confidence in the system. Improvements to the system were identified and made before proceeding to the next phase.

In the second phase, the RNCs were deployed to perform the alarm gathering and reporting. With automated detection and reporting of DS3 failures, the restoration time objective plummeted to less than five seconds to restore 100 DS3s. Even with fully automated restorations, the system was designed so that technicians could intervene if problems occurred during the restoration process.

DACS III Cutover. Transformation of the network to DACS IIIs was fundamental to realizing the FASTAR Phase 1 restoration capabilities. As DACS IIIs were deployed and DS3 service was cut over to the DACS IIIs, the ability to perform manually triggered restorations with RAPID improved over time.

Deployment and cutover of DACS IIIs in 250 nodes in the network spanned almost 12 months. Careful planning and flawless execution was required to ensure that the transition to DACS IIIs did not cause any service interruptions. To minimize risk, cutovers were performed during off-peak hours and special procedures were implemented. One method of performing the cutover to the DACS III entailed using a “dangler” method, in which DS3s were accessible from either the DSX or the DACS III. This process provided a fallback strategy if problems occurred during the cutover.

Database Integrity. One significant system improvement that was made during Phase 1 was enhancing RAPID database integrity to ensure that the RAPID facility database mirrored the actual network facility database. In order to ensure the necessary high level of accuracy for the RAPID facility database, the TSNOW provisioning team worked with the RAPID team to establish metrics for database quality, and to implement process changes and software modifications to achieve the stated metrics. The goal for RAPID database accuracy was set at 99.95%, the level of database quality required for transitioning to the Phase 2 automated-restoration process. In fact, the team achieved a 99.97% database accuracy level.

FASTAR Platform Deployment Planning. FASTAR deployment planning covered criteria for bringing offices online for the FASTAR platform. The first year goal was to certify and turn up 210 offices in a carefully phased, office-by-office schedule that would provide maximum incremental restoration benefits.

To serve that goal, our initial plan was to prioritize all the potential cuts of cable spans in the network, according to a strategic vulnerability model. End offices of the systems affected by the highest priority spans were to be targeted for first deployment. We quickly recognized, however, that gradual deployment based solely on that strategy would create temporary exposure to a subtle set of problems, in which lower-priority DS3s might be restored to the detriment of higher-priority DS3s. To handle this, the strategy for deployment was modified to ensure that each group of links with shared

risk during a cable cut—referred to as a shared risk link group (SRLG)—would receive at least minimal coverage as soon as any part of it became covered.¹¹

An example SRLG is shown in Figure 6. The economics of long-haul fiber systems dictate that rights-of-way be shared wherever possible among multiple links. The links sharing the span do not, in general, all share the same pair of terminal offices, but all share the risk of simultaneous failure when the span is cut. To deal with the problem in an orderly way, we divided the set of links into disjointed subsets of SRLGs.

A cable cut is covered for the FASTAR platform alarm reporting if at least one end of every system riding the span is covered. In Figure 6, a cut of span X-Y would be covered if office X was monitored by the FASTAR platform, but offices Y and Z as yet were not. But in that case, given a cut of span Y-Z, the DS3s on link B would be reported by the FASTAR platform, while those on link C would not.

FASTAR Platform Cutover. With the successful completion of beta and user acceptance testing, the FASTAR platform was ready for deployment. A network validation plan was implemented to methodically validate each office prior to its turn-up as an office covered by the FASTAR platform. As RNCs were deployed, validated, and turned-up, FASTAR platform capabilities were brought on line, and the ability to perform fully automated FASTAR platform restorations improved with the cutover of each RNC.

During the initial cutover of the Florida beta site to full-production, FASTAR platform capabilities were planned so that we could fall back to Phase 1 manual restoration capabilities if an unforeseen problem occurred. The fallback strategy was never needed. In March 1992, the first network event occurred in Florida, after the FASTAR platform cut over. A cable was dug up with a backhoe by a contractor. The first 100 DS3s were restored in eight minutes—just three minutes over our objective—with fewer than 7000 blocked calls. With the confidence gained from this successful restoration, the team quickly moved forward in implementing the FASTAR platform at more than 250 nodes in the AT&T network.

FASTAR Platform Teamwork

The FASTAR platform team was composed of a diverse group of talented individuals, representing many disciplines from different parts of the business. The team's common goal was to make the FASTAR platform

successful—to reach the objective of restoring 100 DS3 paths in less than five minutes.¹² With continuous focus on the restoration objective, and lots of hard work, the team was able to successfully implement this major process breakthrough for AT&T.

The project management team provided ongoing support in the areas of project direction and reviews, funding for research and development, capital and, most importantly, in its confidence that project hurdles would be overcome and the FASTAR platform would be a success. The FASTAR system engineers and developers worked hand-in-hand with customers to implement customer requirements, and solve development and deployment issues. The DSL team supported the project by testing, and providing test environments, with actual network equipment. The restoration process management team defined and implemented metrics, such as restoration cycle time and the number of DS3s restored. They understood the need for the FASTAR platform and its impact on AT&T's success in the marketplace. Finally, the users of the FASTAR platform in the AT&T Transport Service Centers were invaluable, and became its biggest champions.

Striving for an environment with teamwork required constant communication among team members. Of course, as with all complex projects, many conflicts arose when dealing with tough issues. These conflicts were resolved through a well-defined and documented conflict-resolution process implemented by project management. Always, the team focused on the question—"What do we have to do to make the FASTAR project successful?"

The FASTAR team was recognized with the "1993 Bell Laboratories President's Quality Award" for its achievements in developing the FASTAR Operations System platform.

Acknowledgements

While there are too many people to name here who were involved in the success of the overall FASTAR Operations System Platform, some of the key teams were led by the following people: Harry Burns, Hossein Eslambolchi, Jane Fitzgerald, Joe Gottfrid, Bob Guenther, Bill Lentz, Bob Mallon, John McCanuel, Jerry Presley, Tom Ryan, Jerry Shrimpton, Jim Weythman, Paul Wolfmeyer, and Simon Zelingher.

* UNIX is a registered trademark of UNIX Systems Laboratory

References

1. H. S. Burns, C-W Chao, P. M. Dollard, R. E. Mallon, H. Eslambolchi, P. A. Wolfmeyer, "FASTAR Robust System for Fast DS3 Restoration," *Globecom '91 Conference Record*, December 1991, pp. 1396-1400.
2. Jim Weythman, "G2 Data Language Manual - Issue 2," 1994.
3. John D. Musa, Bruce Juhlin, Gene Fuoco, Diane Kropfl, Nancy Irving, "Handbook of Software Reliability Engineering, Chapter 3: The Operational Profile," McGraw Hill, 1994.
4. John D. Musa, Anthony Iannino, Kazuhira Okumoto, "Software Reliability: Measurement, Prediction, Application," McGraw Hill, 1987.
5. Willa K. Ehrlich, John P. Stampfel, Jar R. Wu, "Application of Software Reliability Modeling to Product Quality and Test Process," *Proceedings of the 12th International Conference on Software Engineering*, Nice, France, March 1990, pp. 108-116.
6. W. K. Ehrlich, A. Iannino, B. S. Prasanna, J. P. Stampfel, J. R. Wu, "How Faults Cause Software Failures: Implications For Software Reliability Engineering," *Proceedings of the 1991 International Symposium on Software Reliability Engineering*, Austin, Texas, May 1991, pp. 223-241.
7. Harry S. Burns, Eugene Fuoco, Diane A. Kropfl, "FASTAR and the Engineering Behind the Magic," AT&T Reliability Information Forum, Holmdel, New Jersey, October 1992.
8. John D. Musa, "Operational Profiles in Software-Reliability Engineering," *IEEE Software*, March 1993, pp. 14-32.
9. Willa Ehrlich, Bala Prasanna, John Stampfel, Jar Wu, "Determining the Cost of a Stop-Test Decision," *IEEE Software*, March 1993, pp. 33-42.
10. W. Chen, W. K. Ehrlich, M. E. Engel, D. A. Kropfl, P. Verma, "Testing Restoration Node Controller Performance Under Highly Critical But Infrequent Event Occurrences," AT&T Networks and Systems Performance Symposium, September 1993.
11. H. S. Burns, C-W Chao, P. M. Dollard, R. E. Mallon, H. Eslambolchi, P. A. Wolfmeyer, "FASTAR Operations in the Real AT&T Transport Network," *Globecom '93 Conference*, December 1993.
12. "FASTAR Finish," *AT&T Focus Magazine*, October 1992, p. 24.

(Manuscript approved March 1994)

Chin-Wang Chao is the technical manager of Transport Operations Planning in the Network Operations Systems Planning Division of AT&T Bell Laboratories, in AT&T's Network Services Division (NSD) in Middletown, New Jersey. His group is responsible for restoration strategy and maintenance operations of operations systems planning, and requirement specifications to support the AT&T digital network and emerging technologies. He joined the company in 1986. Mr. Chao has a B.S.E.E. degree from Tatung Institute of Technology in Taiwan, an M.S. degree in computer science from Monmouth College in West Long Branch, New Jersey, and an M.S. degree and Ph.D., both in electrical engineering, from Purdue University in West Lafayette, Indiana.



Gene Fuoco is a district manager of the Transport Network Operations Systems Development Division of AT&T Bell Laboratories, in AT&T's Network Services Division (NSD) in Alpharetta, Georgia. His organization is responsible for developing the RAPID central controller, the PPS plant protection system, and other transport-related applications. He joined the company in 1969. Mr. Fuoco has a B.S. degree in mathematics from St. John's University in Philadelphia, Pennsylvania.



Diane Kropfl is technical manager of the Transport Network Operations Systems Development Division in AT&T Bell Laboratories, in AT&T's Network Services Division (NSD) in Middletown, New Jersey. She is responsible for developing the restoration node controller (RNC) for the FASTAR platform, and the Integrated Network Management System (INMS) for joint Pacific undersea cable operations with KDD from Japan. She joined the company in 1983. She has a B.S. degree in business administration, and an M.S. degree in computer science, both from the University of Delaware in Newark.

