DISASTER RECOVERY APPLICATIONS FOR SATELLITE COMMUNICATIONS SYSTEMS[†]

Hans Kruse, Tony Mele, Sara Young, Chris Washburn
J. Warren McClure School of Communication Systems Management
Don Flournoy
School of Telecommunications
Ohio University
Athens, Ohio 45701

ABSTRACT

We report on a series of experiments conducted jointly by Ohio University and The Huntington National Bank. Our aim was to determine if ACTS-based data communications services can serve as redundant facilities in an otherwise terrestrial network. We examined both technical compatibility and network management issues. In addition, Ohio University researchers conducted interviews with business resumption planners at several large corporations. As a result of these activities, we have concluded that the ACTS system is fully compatible with the terrestrial network. We expected that the transmission delay inherent in geostationary systems would have a negative impact on the response time provided to the end user. We confirmed that this is the case, and give quantitative results for a number of circuits in our test network. We can confirm full compatibility between the ACTS system and terrestrial network management systems. We suggest opportunities for closer ties between the network management system and the satellite service. With these improvements, users could take advantage of the fast circuit setup times in the ACTS system. Finally, we project possible applications for satellite based data communication services, and suggest appropriate service parameters for such services.

BACKGROUND

Figure 1 shows a portion of the network backbone for The Huntington National Bank. All lines shown are T1 (1.544Mbits/sec) lines. Lower speed tail circuits (typically 9.6kbps and 56kbps) connect banking offices to the nearest backbone location. This network shows several features which are typical for large financial institutions.

The network is designed primarily to provide access to a single data center. Increasingly the backbone also carries LAN to LAN traffic. Currently, much of the LAN traffic is also destined for the data center (either for mainframe access , or for servers co-located with the mainframe). The majority of the traffic on the network consists of IBM proprietary protocols, mostly SNA (System Network Architecture) over SDLC (Synchronous Data Link Control). The older Bisync protocol is being used for some applications, and an increasing amount of TCP/IP traffic is being handled.

Since so many financial networks are centered around a critical data center, business resumption plans and network protection mechanisms tend to focus on protecting access to this critical site. We briefly discuss data center protection below. Network protection is discussed in the next section on Applications. For a data center, several layers of protection are used:

[†] Portions of this paper have been drawn from the Final Report submitted to the ACTS Experiments Office [1] .

Physical Access Protection

Central data processing sites usually employ tight physical security. Guards, security cameras, and locks keyed to magnetic-strip ID cards are common. More elaborate methods such as retina scanners may be used if that level of protection is deemed to be warranted by the potential loss of business caused by a data center disruption.

Data centers are usually housed in unmarked buildings, and in some cases they are located in remote areas to permit the establishment of a physical security perimeter.

Data Center Redundancy and Backup

Critical components in the data center are made redundant where possible. Service contracts with equipment providers usually specify response times and require that the service provider have spare equipment available for immediate deployment.

Utilities (mostly electric power and possibly cooling systems) are backed up with in-house standby systems. Critical data bases are backed up on removable media, and stored off-site. These backups may take place weekly, daily, or even hourly, depending on the volatility of the data.

Access control and redundancy cannot, however, protect against the partial or total loss of the data center resulting from natural or man-made catastrophic events such as floods, hurricanes, or terrorist activities. As a final layer of protection, the data center operator establishes an alternate data center site. Most financial institutions do this by contracting with a service provider who specializes in the operation of standby data centers (“hot sites”). In the case of Huntington Bank, this provider is Sungard Recovery Systems in Philadelphia. Figure 2 shows the typical setup for such an arrangement. The operational network in figure 2a has one data center at the top of the (hypothetical) network diagram. If the data center becomes disabled as shown in figure 2b, the data center operator “declares” an emergency and activates the business resumption plan. Typically, backups of data bases and programs are shipped to the hot site, and company personnel is dispatched to activate critical applications on the hot site computers, typically within a day. On-demand data circuits (such as AT&T’s ACCUNET Reserve™ T1 service or equivalent offerings from other carriers) are activated to connect the company network to the hot site. Depending on the extent of damage at the original data center site, additional applications may be activated at the hot site as time progresses.

APPLICATIONS

The focus of our project is to improve network availability in order to lessen the probability of a company having to activate a hot site as a result of catastrophic network failure. We will briefly describe typical network redundancy designs, and then discuss two groups of applications addressed by the tests conducted in this experiment.

Figure 3 shows a typical network redundancy design. The operational network backbone in figure 3a provides at least two network connections at each site. Every effort is made to make these two connections truly independent. If possible, two different local access vendors are used. Physical cable routes are examined, and a minimum physical separation of routes is enforced. Multiple long-distance carriers, or diverse routes from one carrier, are used. In theory, a failure event will effect only one of the routes such as shown in figure 3b. If all bandwidth utilization are kept below 50%, traffic from the failed route can be redirected to the remaining one. This redirection can be accomplished automatically using T1 multiplexers or network routers.

In light of this redundancy design, we examine two possible satellite applications.

Redundant Circuits

Depending on the actual location of the data center, achieving true diversity between two terrestrial routes can be very difficult and expensive. Diverse routing is fairly easy in urban areas with high concentration of business communications users. Rural and remote areas pose the most difficult problem. In many cases, special construction projects are needed to implement redundant routes.

Figure 4a shows a satellite circuit as the redundant route. During normal network operations, the satellite circuit could carry a portion of the data traffic. Based on the results of our tests, it is more likely that the satellite circuit will be idle (i.e. not use bandwidth in the satellite network), since protocol performance and response times are degraded for interactive applications (see also [2]). During a terrestrial route failure, the satellite circuit must activate quickly enough to prevent the loss of user sessions. In the case of SNA and TCP/IP, the permissible startup delay is around 10 sec.

As more critical applications are migrated to LAN-based server systems, more than one site in the network will become “critical”. Earth stations at each critical site, and at least two other sites in the network would allow the establishment of redundant paths if any one of the critical sites lost its network connection. These satellite circuits could also be used to provide additional network capacity if one network site has to assume processing responsibility for a failed site (i.e. a LAN file server).

Due to protocol performance issues which will be discussed later in this paper, the satellite circuits may not be suitable to carry interactive SNA or TCP/IP traffic during normal network operations. Satellite circuit usage will have an extremely low duty factor if activation occurs only during network outages. Additional circuit activations could be used to carry out network-based file backup operations, for example from a primary to a redundant data base server.

Network Backup

The previous section was focused on redundant circuits with fast activation times. Figure 4b shows a configuration which resembles the one used for data center backup. We refer to this as “network backup”.

In this application, earth stations are maintained by a service provider at a central facility, and at critical customer sites in “stand-by” mode. During a network failure, the service provider activates the earth stations and a satellite link to the critical site that has lost network access. Terrestrial on-demand circuits are activated from the service provider to network access sites on the customer network, thereby making the critical site available to the network users. Activation times will probably be somewhat slower than in the redundant case. User session losses will probably occur.

As in the redundant case, the satellite links may also be used for non-interactive traffic to improve satellite system utilization, provide additional network backup services, and reduce the effective cost of the stand-by equipment.

The network backup approach permits a “rapid deploy” variant, in which earth stations are dispatched after a failure occurs. Applications of this approach exist in emergency communications following severe weather or earthquake damage, and in networks which involve temporary sites.

EXPERIMENT DESIGN

The experiment conducted by Ohio University and The Huntington National Bank was designed to test the feasibility of the applications described above. We wanted to determine ACTS circuit compatibility with terrestrial equipment, system reliability, and protocol performance.

Figure 1 indicates the sites where ACTS earth stations were located, namely the data center in Columbus, Ohio, a backbone network node in Parma, Ohio, and the campus of Ohio University in Athens, Ohio. The Ohio University and the data center sites were used for initial tests “outside” the actual production network, while production network tests were conducted between Columbus and Parma.

Figure 5 shows an equipment block diagram for each site. User equipment attached to the multiplexer included IBM front-end processors (at the data center), cluster controllers (in Athens and Parma), 9.6kbps and 56kbps tail circuits (in Parma), and video conferencing equipment (in Columbus and Parma). A low-speed (9.6kbps or 56kbps) channel between the multiplexers was used to monitor bit error rate performance during all tests.

TECHNICAL COMPATIBILITY

Circuit Level

The ACTS supplied T1 circuit was found to be fully compatible with the terrestrial network equipment. This result is of course not surprising since the earth station interface is an industry standard switching system (the Redcom MSP switch). It is significant that the earth stations in the ACTS system use a digital interface to the satellite. Traditional analog satellite links introduce “jitter” into the T1 link which must be corrected using buffers in the user equipment. These buffers are usually not configured, and often not even provisioned in terrestrial networks. Since the ACTS system does not require such buffers, it can provide fully compatible redundant circuits; the user equipment does **not** need to be configured for satellite use.

From the perspective of a network engineer, the ACTS earth stations function as central offices. All framing is terminated at the earth station. This includes the ESF (Extended Super Frame) facilities data link, and the ESF performance registers. Remote CSUs must therefore be controlled using in-band loop and reset codes.

Network Management

In our experiment, the ACTS T1 line was entered into the network management database used to manage the multiplexer network. No fundamental problems were encountered. We did notice that the management system we were using was not well suited for on-demand T1 circuits (both satellite and terrestrial). Special procedures had to be developed to prevent numerous critical alarm indications when an on-demand circuit is taken out of service.

We used two sets of procedures for this purpose. The first procedure relies on setting maintenance loop-backs on the CSUs; this satisfies the network management system, but can negatively impact the operation of the earth station in stand-by mode (see “Clocking”, below).

The second procedure involved the explicit removal of all circuits that were using the T1, either before shut-down, or right after the deactivation of the circuit. This appears to be the preferred method. It would be the most natural configuration if the ACTS circuit is used as a redundant circuit. We note in passing that our efforts to automate this procedure were hampered by shortcomings in the scripting language of the network management system.

We also note that the ACTS earth stations do not provide circuit statistics in the ESF performance registers. While we did not have an opportunity to test this, it would appear that a network management system using ESF performance reports would not reflect the true satellite circuit quality. In the case of the ascom Timeplex T1 multiplexers, all circuit quality data is derived from an in-band test channel in the proprietary frame structure used by the multiplexers. In this case, all reports correctly reflect the observed end-to-end circuit quality.

Clocking

Please refer to Figure 5 for the following discussion. In the ACTS system, the ACTS spacecraft always represents a clock boundary, since the earth stations and the spacecraft use independent time sources. The user has the option of either clocking the attached user equipment using the earth station’s internal clock, or to time the terrestrial interface equipment in the earth station using the attached user network. This latter configuration will be most common in a mixed satellite/terrestrial site. Earth station based clocking is used primarily at isolated user sites.

We verified the correct operation of both clocking modes. No significant issues arose in the configuration which uses the earth station as the master clock.

When the earth station interface is timed from the user network, a new clock boundary appears inside the earth station. This boundary exists between the interface switch (i.e., the Redcom MSP) and the digital signal processor which interfaces to the satellite modem. The ACTS earth stations contain a buffer at this boundary to protect against timing differences between the earth station and the user network. If a timing discrepancy exists over an extended period of time, buffer capacity is exhausted, and the buffer is reset; this is often referred to as a buffer slip. Table I shows the time between buffer slips for a range of clock differences.

We were able to verify the operation of these buffers since the terrestrial network drifted substantially from the nominal T1 rate during our tests. While the buffers provide protection for the times indicated in Table I, we did have difficulty recovering the circuit after a buffer slip.

Our problems were attributable in part to a software problem in the earth station. However, even when the buffer recovery operation was carried out correctly, we noticed problems in the network. The current buffer recovery method involves discarding all buffer content, and re-initializing the buffers. This generates a T1 error event of sufficient duration to cause the T1 multiplexers to take the T1 circuit out of service. Recovery of the circuit then requires significant time since the multiplexer has to verify acceptable bit error rate performance before permitting data flow to resume. Since we could not modify the buffer recovery mechanism, we cannot offer data on alternative recovery schemes.

A different problem appeared when we were using a CSU loop-back to silence alarms in the multiplexers. When a CSU loop is in place towards the terrestrial network, the earth station is effectively isolated form its clock source. If the redundant circuit is the only application for the earth station, this is not a problem. However, care must be taken to provide an alternative clock source if the earth station is to provide other circuits while the terrestrial connection is looped. The current interface equipment in the earth station does not provide this option.

CIRCUIT QUALITY

Bit Error Rate and Signal Strength Results

During an extended operations period from 2/7/94 to 2/11/94, we focused our data collection on Bit Error Rates and Signal Strength trends. Weather information from the nearby Parkersburg airport was collected during the first, second, and last day of operation. Figure 6 shows the barometric pressure, and the temperature and dew points on these days. The distance between the temperature and the dew point lines is an indication of relative humidity (the lines will coincide at 100% humidity). The transit of a major storm front can be clearly seen from this data.

In figure 7, we show the receive signal strength at the Athens earth station. The effect of the storm on the signal strength is clearly visible. A peak in signal strength seems to coincide with the driest period recorded. The signal strength graph also shows short-term fluctuations of about 3db in signal strength. These do not seem to correlate with any weather observations we were able to collect.

Figure 8 shows – for the same time period – measured bit error rates. This graph is a compilation of many types of measurements. All data were recorded in remote loop-back mode. Some points are based on circuits which traversed the satellite more than once. Bit error rates for those cases have been normalized to a single-hop equivalent.

Points in the data series labeled “Normal” were taken while the earth stations were operating with the rain fade selection mechanism enabled. Some points are from coded operations as determined by the rain fade parameters. All of these data points indicate end user error rates during production-mode operations.

The data points marked “Uncoded” were taken with The Athens earth station forced to operate in the uncoded mode, regardless of signal strength. The data at the end of the time period shown clearly show the BER differences between coded and uncoded operation. Under normal conditions, the end-user will not experience error rates above the threshold set for the onset of error correction.

The vertical dashed line in figure 8 indicates the time of a partial equipment (High Power Frequency Doubler) failure. There is no effect of the component on the receive signal strength, and a change in coding parameters insured that there was minimal effect on normal operations. Uncoded operation did show the effect of the weakened transmit path.

Overall, the week of operation analyzed here represented a period of fairly severe weather. Over the remainder of the experiment, bit error rates remained well below the ones shown here.

SYSTEM RELIABILITY

Analysis of Circuit Outages During Phase I Extended Operations

The analysis in table II applies to the first extended (7 day) operations period. It illustrates some of the issues that had to be resolved. It must be stressed, however, that an availability computation for an experimental system like ACTS cannot predict availability of an equivalent commercial system.

Outage Analysis and Determination of Availability:– The numbers in table II translate into 14.7 hours of downtime during 119 hours of operation, or about 12% downtime.

The outages are distributed over four failure modes as follows:

4% are attributed to the spacecraft related TEW failure mode,
2% were weather related (the earth station dropped out of the network due to excessive rain fade; since these failures software changes at the NCS have improved the speed at which circuits can be restored),
2% were planned downtime for software and hardware maintenance, and
4% were outages without known causes (3% were in the single failure on 2/9; circuit restoration has been automated since then so that long downtime after network restoral should not reoccur).

Using standard measures for network reliability we report a MTBF (Mean Time Between Failures) of 468 minutes, with a MTTR (Mean Time To Repair) of 62.9 minutes (this translates into an availability figure of 87%).

Two actions were taken during the week to improve the system availability. A software change enabled us to establish T1 (24 channel) circuits while the earth stations are in coded mode. This change allowed us to implement an automated (i.e., unattended) restart procedure at the earth stations which sets up the T1 circuit as soon as the earth stations have been acquired after an outage.

A very rough estimate would suggests that we can reduce the observed outages by about a factor of 2 with these changes (including the assumption that planned outages can be avoided during a production run).

Analysis of System Outages During Phase III

Operations during phase III suffered from a persistent T1 clock mismatch problem. For that reason, on-site monitoring equipment registered frequent circuit outages (about one per 3 hours), with lengthy restoration sequences. Worse yet, since the T1 VSAT buffers did not properly reset after about 10 buffer slips, some restoration periods are excessively long since manual intervention was required in that case to restore the circuit.

We are therefore not able to report availability figures based on the on-site equipment. A review of the ACTS system logs permits a very rough estimate. It must be noted that the logs do not always appear complete, and often restoration periods must be estimated. Nevertheless, we believe that the analysis shown in table III represents a reasonable upper bound on system availability during the time period of May 16 to May 20, 1994.

System availability computed on this basis was 93%. This agrees very well with our earlier assertion that outages could be reduced to 1/2 of what was observed in February. Again, it must be stressed that this is an upper bound estimate. Overall circuit availability was probably somewhat lower even if the T1 clock slip problem had been eliminated, because bit error rates occasionally did disrupt the circuit even though the earth station remained synchronized. These types of outages are very difficult to isolate in the on-site logs, and the clock slip problem affects the circuit restoration time (i.e., the TTR (Time To Repair)).

PROTOCOL PERFORMANCE

Response Time Results

Large amounts of time in this experiment were devoted to the measurement of response times in a standard IBM SNA/SDLC environment. We used a consistent response time script for all measurements except for the production circuit tests in phase III. In the script, the experimenter requested the display of a large text file at the test terminal. The file was displayed one screen at a time; requests for the display of the next screen were made as soon as the previous screen output was complete. At the end of the file, the process was repeated. For each scenario, the script was run for a minimum of one hour. Local test equipment at the terminal site was used to measure the response time. Response time estimates were also obtained from the host-resident Netspy program. Netspy consistently reported higher response times at the 56kbps speed than at 9.6kbps. We cannot explain this behavior, although we suspect that it is caused by a faulty estimation algorithm. Since Netspy is located at the host it must estimate the transit delay of frames. Unless otherwise noted, we report the data collected at the terminal site.

Table IV shows our tests for both 9.6kbps and 56kbps lines. In most cases the line in question had only the test controller with one attached terminal on it; the lower half of the table reports data for two production circuits with multiple controllers and a large number of concurrent users. We set up normal “single-hop” circuits, as well as 3-hop circuits which simulate multiple satellite links in a network. In this paper we show results from the single hop cases only.

Unless otherwise noted, results are shown for cases when the multiplexers report the receive bit error rate at both sites to be “Very Good”, which translates to a bit error rate around 10^{-6}.

We draw several conclusions from the data in table IV The rain fade compensation appears to have no effect on response time. Degradation in the bit error rate causes a significant increase in response time. Circuits with heavy device loads (the production circuits) experience a much more severe response time degradation compared to the lightly loaded test circuit. Worst case response times, which we attribute to transmission errors or traffic “bursts”, degrade very heavily.

Overall, we show a minimum increase in the response time of 1.1 sec when the circuit is moved to ACTS during periods of good BER performance, with larger increases expected on heavily loaded circuits.

These response time increases are tolerable in a disaster recovery context, but may cause problems in normal production operations.

CONCLUSIONS

Domestic Applications

Backup Systems: – In the US, users correctly perceive the terrestrial network as highly redundant[3] . Backup systems for networks are unattractive since they are slow to react, and since redundant system with fast response time can be obtained economically. Domestic use of satellite circuits for network protection is therefore expected to take the form of redundant circuits, which are discussed in the next section.

Satellite backup systems will be attractive for temporary locations and emergency communications. Financial services are increasingly demanded for large public events in remote locations, such as outdoor concerts, fairs, or “Eco-Tourism” events. A rapid-deploy satellite system is well suited for this application.

The past few years have shown that the terrestrial network is temporarily vulnerable to natural disasters such as hurricanes or earthquakes[4-9] . While restoration typically requires only a few days after a disaster, financial, banking, and insurance systems are needed quickly to aid in the recovery. A satellite trunking system with local wireless (e.g. cellular) communications can serve this need[10] . The ACTS system design is fully capable of supporting this type of application.

Redundant Systems: – We have outlined earlier in this paper that most terrestrial data networks rely on redundant designs to protect against single circuit failure. The portion of the network which is the hardest to protect is the local loop, i.e. the connection between the user site and the first telephone company switching office in the circuit path. In urban areas it is usually neither difficult nor expensive to obtain multiple local loops which are physically separated so that a simultaneous failure of both circuits is highly improbable.

In rural areas redundant local loops can be obtained only at considerable cost. In addition, redundant local loops protect only against random failures, not against natural disasters or deliberate disruptions of the network connection (e.g., acts of terrorism or industrial sabotage). Protection against loss of network connectivity in these cases is more easily and economically obtainable using satellite links.

It is therefore our expectation that a domestic market for satellite data communications capacity will exist for users in rural or remote areas, as well as for users who have a need to protect their installations from deliberate attacks on the terrestrial network. Due to the requirements placed on redundant circuits, the earth stations will be permanently installed. Acquisition and installation costs of earth stations will not necessarily drive purchase decisions, but they must be competitive with startup costs for other redundancy options.

International

Compared to the domestic network, circuits to, and among, non-US sites are much less reliable. In addition, transport and interface standards are not always compatible with US equipment. We therefore expect some market potential for permanent production installations of satellite circuits into international sites without adequate terrestrial capacity.

More importantly, we expect a significant market for redundant facilities using satellites. The mean time to repair for circuit outages in non-US networks is often much higher than those in the US. This is not driven by technical issue, but rather by work rules (e.g., no 24-hour coverage of network control centers). In certain situations, a market may exist for backup facilities, using rapid-deploy earth stations.

It is interesting to note that the current network reliability requirements are driven by the need of international sites to access central data base resources as US data centers. With increasing deployment of truly distributed data bases, the need for access by US user to international sites will become more important and drive up the demand for network reliability.

Requirements

Bandwidth: – We expect a minimum requirement of one or more DS1s (1.544Mbps) for a site that needs redundant circuits. This requirement will increase as the networks in question carry an increasing proportion of image and video traffic.

Duty Factor: – The duty factor for a redundant circuit will be very low. Domestic redundant facilities are often not used except during scheduled tests, for the entire lifetime of an installation. Duty factors for international sites may be somewhat higher.

In the case of terrestrial redundant circuits, users either accept the low duty factor, or add low-priority traffic onto the redundant circuit. Often load-balancing is used since terrestrial redundant circuits are not always on-demand facilities.

In the satellite case, these approaches may not work. The satellite circuit is most likely an on-demand circuit for which load-balancing is not desirable due to the cost structure. Terrestrial protocols for interactive applications are not well suited to the satellite environment as we have shown here. The satellite circuit will most likely not be seen as a viable production facility for this type of application. Satellite service providers will need to offer incentives for non-interactive use of the capacity to improve duty factors.

An interesting opportunity surfaced during our discussion with business resumption planners. With the increasing importance of LANs (which we argued earlier enhances the need for redundant links to remote sites) the need for reliable backup of LAN data bases increases[11] . Users may be receptive to the use of on-demand circuits to retrieve data bases from remote sites directly to data vault facilities. Other applications to be looked at include overnight data base replication to remote sites, distribution of multi-media training material, and video conferencing.

Price: – Based on our interviews with major users, justification of redundant and backup facilities is mainly based on the overall expected business impact of an outage. Successful disaster recovery service providers point out vulnerabilities, and help assess the potential losses. The effect of a major systems outage in terms of actual losses, loss of future business opportunities, and potential legal liability is so large that protection will be put in place. Users will look for the least expensive option which provides adequate protection.

Potential satellite services will therefore be evaluated by the user based on price only if an equivalent terrestrial solution exists. Otherwise the satellite service will be evaluated based on its capability to provide protection. In the latter case, reliability becomes the main issue. In almost all cases the user has a “last-resort” option of constructing user-owned facilities (e.g. microwave). These options provide a “benchmark” for — or a limit on — satellite service pricing and reliability requirements.

In many user organizations, the data processing and networking groups will present available options to the end-user department. The final decision is made there. The cost of implementing a redundant facility must in some sense be “reasonable” compared to the perceived value of the application to the end-user. In our observations we found little use of formal cost/benefit analysis.

Users will be most sensitive to initial startup costs (e.g., equipment purchases). Monthly fixed fees are more easily justified if some of the additional applications discussed earlier can be used to spread the cost. For recovery applications usage fees are almost immaterial due to the expected low duty factor.

REFERENCES

[1] Kruse, H. Disaster Recovery Via ACTS: Final Report to the NASA ACTS Experiments Office. Ohio University 1994; Technical Report OU-CSM-HK94-TR002.

[2] Kruse, H. Performance of Common Data Communications Protocols Over Long Delay Links: An Experimental Examination. in 3rd International Conference on Telecommunication Systems Modeling and Design. 1995. Nashville, TN.

[3] Fagerstrom, R. and J. Healy, The Reliability of LEC Telephone Networks, IEEE Communications Magazine 31 (1993) 44-48.

[4] Cocheo, S., Don’t take phone service for granted (banks’ disaster recovery planning), ABA Banking Journal (1990) 76-77.

[5] Johnson, S., AT&T Network Failure Raises Questions: A Review of the News Events, in Disaster Recovery Journal. 1991, p. 40.

[6] Johnson, S., The Impact of Recent Disasters, in Disaster Recovery Journal. 1992, p. 8-9.

[7] Olson, R.A., Earthquake Protection for Data Systems, in Disaster Recovery Journal. 1992, p. 35-42.

[8] Piersall, J.R., Small Business -- Big Risk, in Disaster Recovery Journal. 1992, p. 34-35.

[9] Zorpette, G., Keeping the phone lines open, IEEE Spectrum 26 (1989) 32-36.

[10] Kruse, H., et al. Comments on the Use of ACTS Technology For On-Demand Communications Using Rapid Deploy Earth Stations. in 15th International Communication Satellite Systems Conference and Exhibit. 1994. San Diego.

[11] Nasr, F.J., Enhance Your Mainframe D/R Plan To Cover Micro Computers, in Disaster Recovery Journal. 1992, p. 50-54.

TABLES

Table I: Predicted intervals between buffer reset events for different T1 rate mismatches
(T1 VSAT MSP timing derived from the terrestrial circuit)

T1 Rate Mismatch (Bits per sec)	Average hours between buffer resets
1	6.83
2	3.41
3	2.28
4	1.71
5	1.37
6	1.14
7	0.98
8	0.85
9	0.76
10	0.68
11	0.62
12	0.57
13	0.53
14	0.49
15	0.46

Table II: Analysis of system availability for 2/7/94 to 2/11/94

Date/Time	Duration	Code[1]	Notes
2/7 0900			Start of run
2/8 0358	30 min	TEW	SF sync lost; abort 4005 followed by 6001 and 7004
2/8 1100	30 min	Fade	T1 sync lost between multiplexers, satellite circuit remained up.
2/8 1313	27 min	Fade	Network connection lost (unknown abort code); Eb/No is poor (13.7).
2/8 1615	9 min	Other	Network connection lost (unknown abort code), Eb/No is 18.2, circuit setup very slow.
2/8 1946	104 min	Fade	Network connection lost, abort code 4003. Extreme rain fade (Eb/No 12-14 variable, worse for short periods of time). Unable to re-acquire. Circuit setup requires uncoded operation. After acquisition succeeds, circuit setup is very slow (3 min delay).
2/9 0429	211 min[2]	Other	Network connection lost, abort code 7001. Monitoring procedure at Huntington fails, circuit outage is not reported. Circuit remains down until 0800. Circuit setup attempt fails (>5 min delay).
2/9 0800	55 min	Planned	System maintenance shutdown. Circuit setup delay is 1 min.
2/9 1346	60 min	TEW	Network connection lost, abort code 4005. No delay in resetting the circuit.
2/9 2229	60 min	Other	Network connection lost, abort code 7001. No delay in resetting circuits.
2/10 0552	120 min	TEW	Network connection lost, abort codes 4005, 6001, 7004. Extreme (5 min) delay in resetting circuit.
2/10 2027	28 min	Planned	System maintenance shutdown. No delay in resetting circuit.
2/10 2119	19 min	TEW	SF sync lost, abort codes 4005, 6001, 7004. 2.5 min delay in resetting circuit.
2/11 1215	90 min	TEW, Planned	Network connection lost, SF sync lost, abort code 4005. Software reload at the LeRC NCS performed during the outage.
2/11 2223	37 min	TEW	Network connection lost; SF sync lost.

---

[1] TEW: An outage caused by an interference problem on the ACTS spacecraft;
Fade: Weather related outage due to a loss of the transmitted signal in extreme rain conditions;
Planned: Scheduled outage due to the need for software or hardware maintenance at the LeRC MCS;
Other: Cause of the outage is unknown.

[2] Network outage was 15 minutes. Outage detection and recovery were delayed due to a procedural breakdown.

Table III: Observed system outages based on the ACTS error log for May 16, 1994 to May 20, 1994. Month, Day, and Time define an observed event.
TTR is Time to Repair (estimated from restoration of Sync and/or the circuit).
TBF is defined as the Time Between Failures.

Month	Day	Time GMT	Hour	Min	TTR	TBF	Notes
5	16	0000	0	0			Start of Observation
5	16	1415	14	15	10	855	No Record of restart time, TTR estimated
5	17	0737	7	37	47	1042	Sync Loss
5	19	1625	16	25	116	3408	Sync Loss
5	19	1835	18	35	3	130	Sync Loss
5	20	1416	14	16	309	1181	Sync Loss and Circuit Setup Rejection
			Average:		97	1323.2	
						93%	Availability

Table IV: Measured response times for 9.6kbps and 56kbps circuits. Measurements are reported in chronological order. Tests after 4/22/94 were made on production circuit with actual banking traffic; all tests prior to that time were conducted on single-controller test circuits with scripted test traffic.

	Response Time (sec)
9.6kbps circuit, 1/10/94 run.	1.50
56kbps circuit, 1/19/94 run.	1.55
9.6kbps circuit, 2/10/94 run. Slightly higher BER and coded operation	1.63
56kbps circuit, 2/10/94 run. Slightly higher BER, uncoded operation.	1.51
56kbps circuit, 2/10/94 run immediately before the result shown above, coded operation.	1.49
56kbps circuit, 2/10/94 (morning). This was a period where the multiplexers reported “Fair” receive BER at ES #10 (around 10^{-4})	1.68
9.6kbps circuit re-test on 4/22/94	1.63
9.6kbps terrestrial circuit baseline at the same time as the measurement above	0.48
56kbps circuit re-test on 4/22/94	1.50
56kbps terrestrial circuit baseline at the same time as the measurement above	0.37
Production circuit #9, 9.6kbps, during terrestrial operation on 5/16/94 (Netspy report)	1.64
Worst case response time for the case above	3.50
Production circuit #9, ACTS test, 5/17/94	4.94
Worst case response time for the case above	11.60
Production circuit #80, 56kbps, during terrestrial operation on 5/17/94 (Netspy report)	0.55
Worst case response time for the case above	2.40
Production circuit #80, ACTS test, 5/18/94	1.94
Worst case response time for the case above	13.30

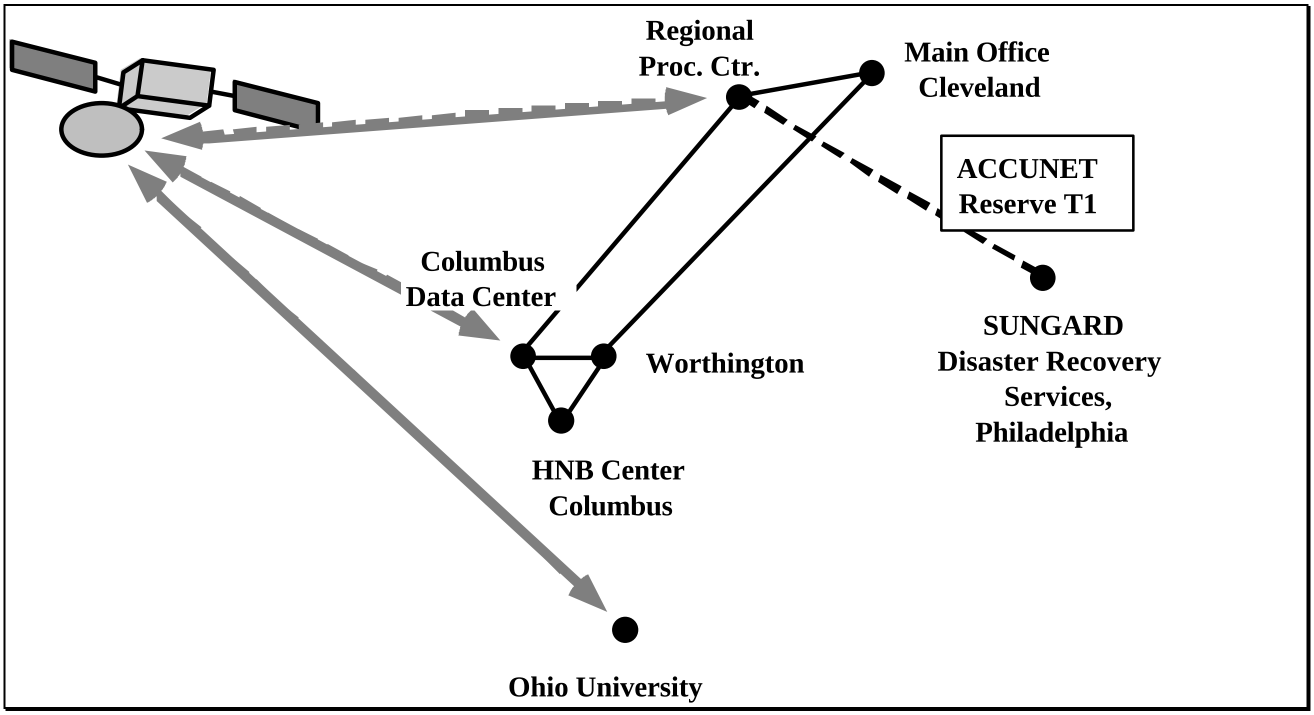

Figure 1. — Partial network diagram for The Huntington National Bank. Only backbone T1 links are shown. Backbone links that were not relevant to the experiment have been omitted. The connection to SUNGARD is a terrestrial on-demand facility.

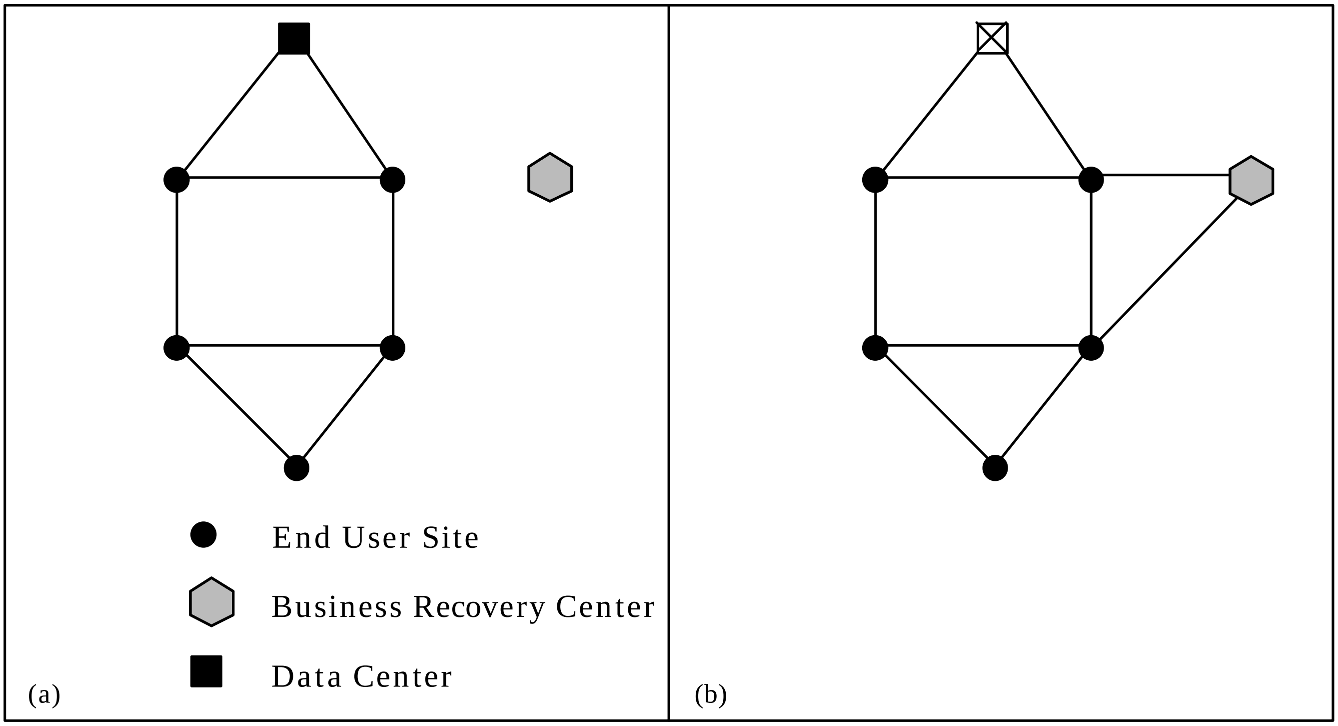

Figure 2. — Disaster recovery using a hot site. Normal operation is shown in (a). Operations during periods of data center failure are depicted in (b).

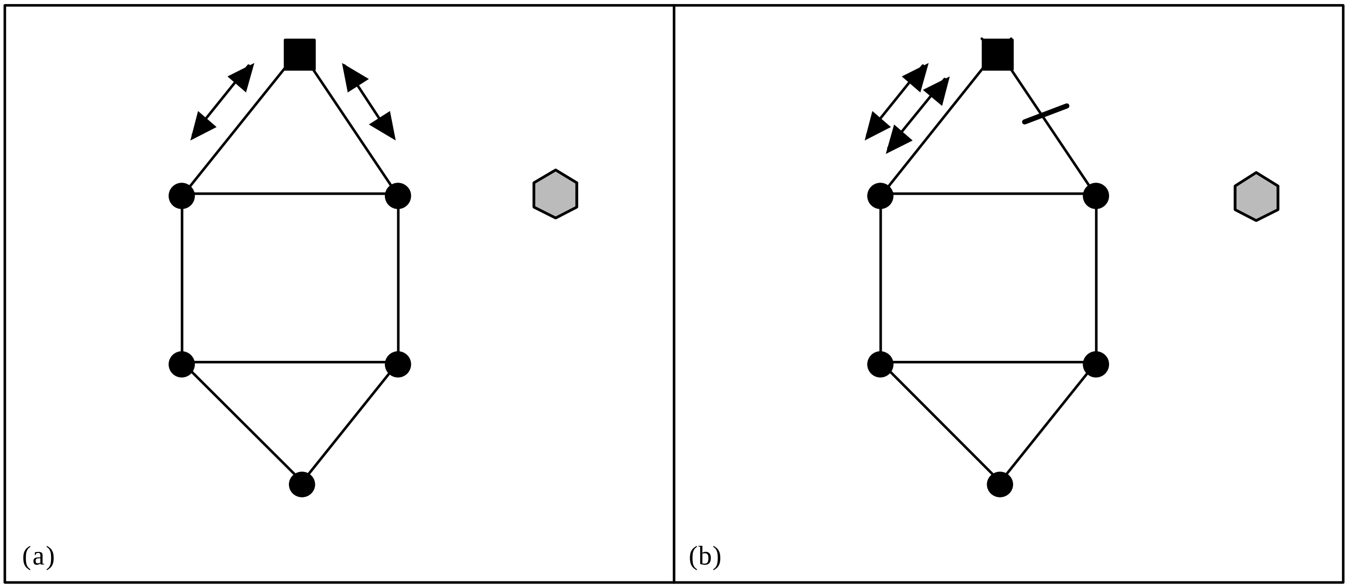

Figure 3. — Use of redundant circuits in a mesh data network. Part (b) of the figure shows the redirection of normal traffic flow during a single-link failure.

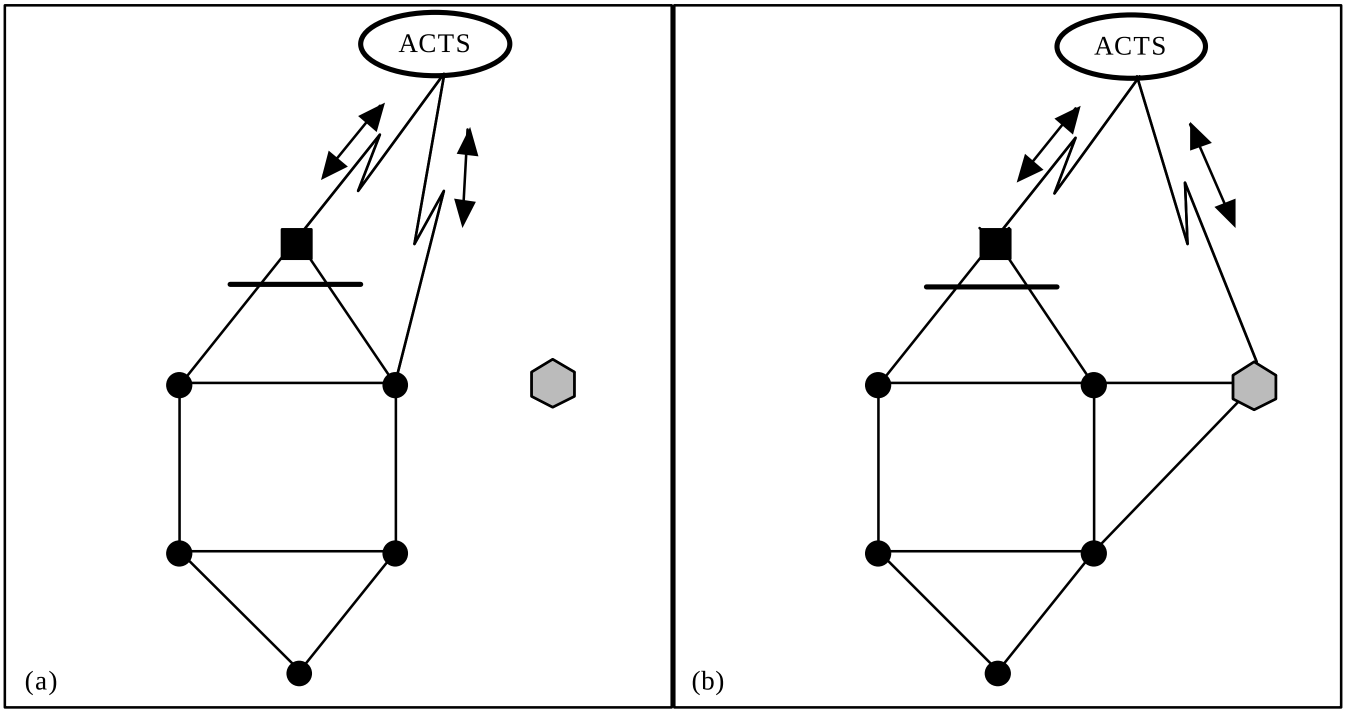

Figure 4. — Use of satellite circuits for redundant (a) or backup (b) operations.

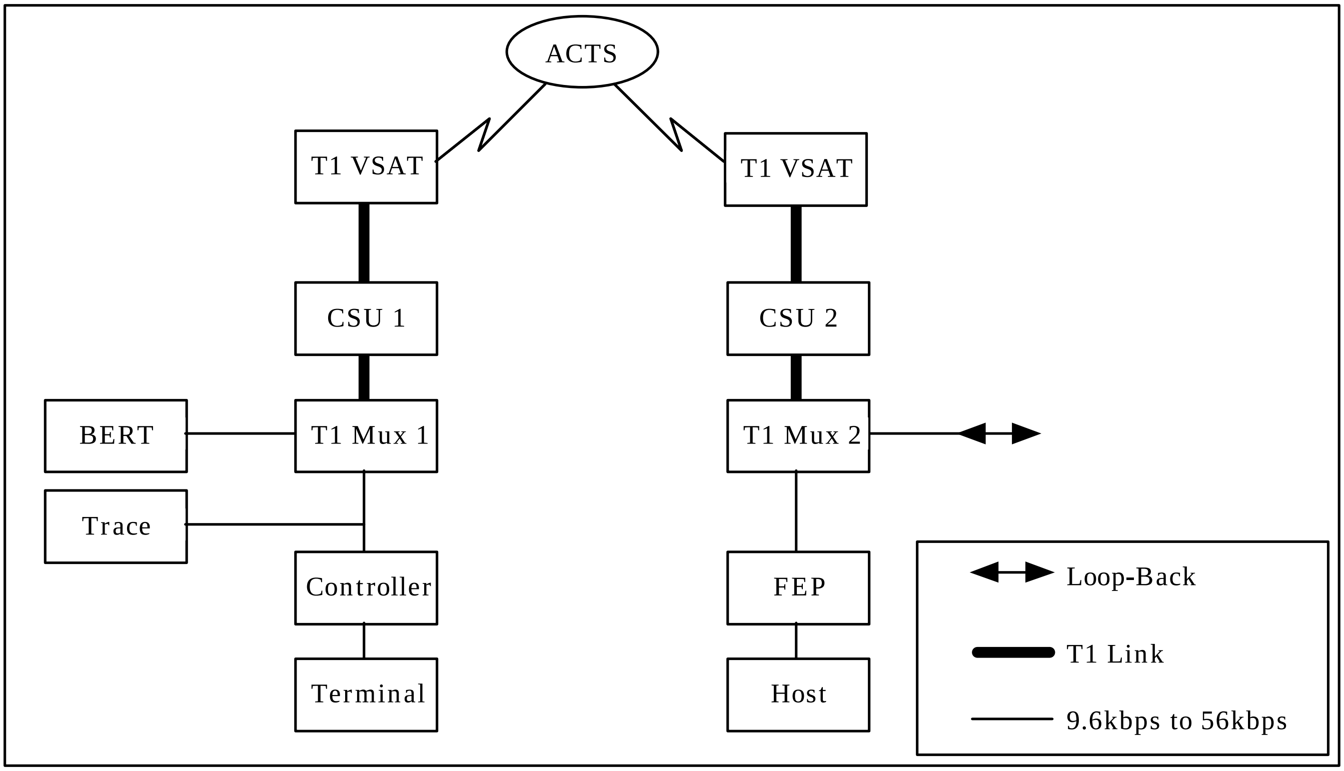

Figure 5. — Equipment block diagram for the error rate and response time testing. (CSU–Channel Service Unit; FEP–Front End Processor; BERT–Bit Error Rate Tester).

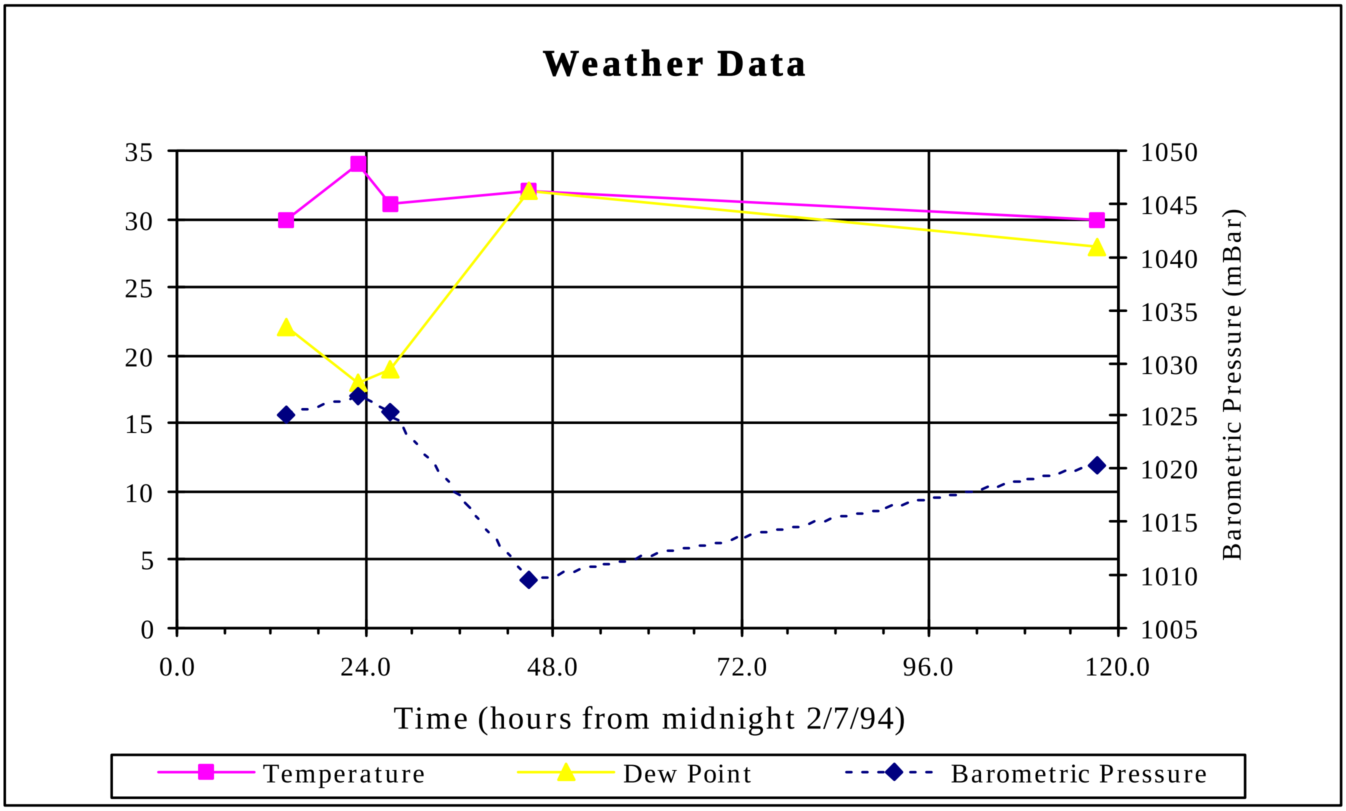

Figure 6. — Weather observation data (Parkersburg, WV airport) for February, 1994.

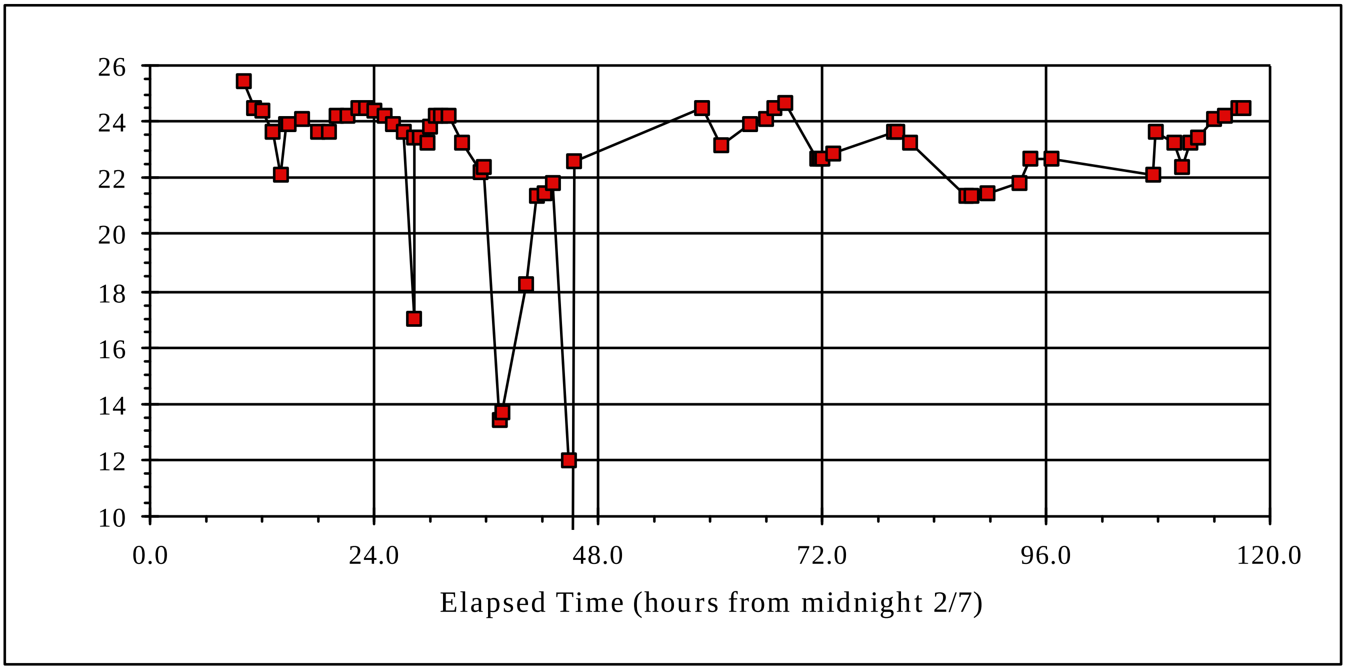

Figure 7. — Receive signal strength (Eb/No) at the Athens Earth station; February 1994.

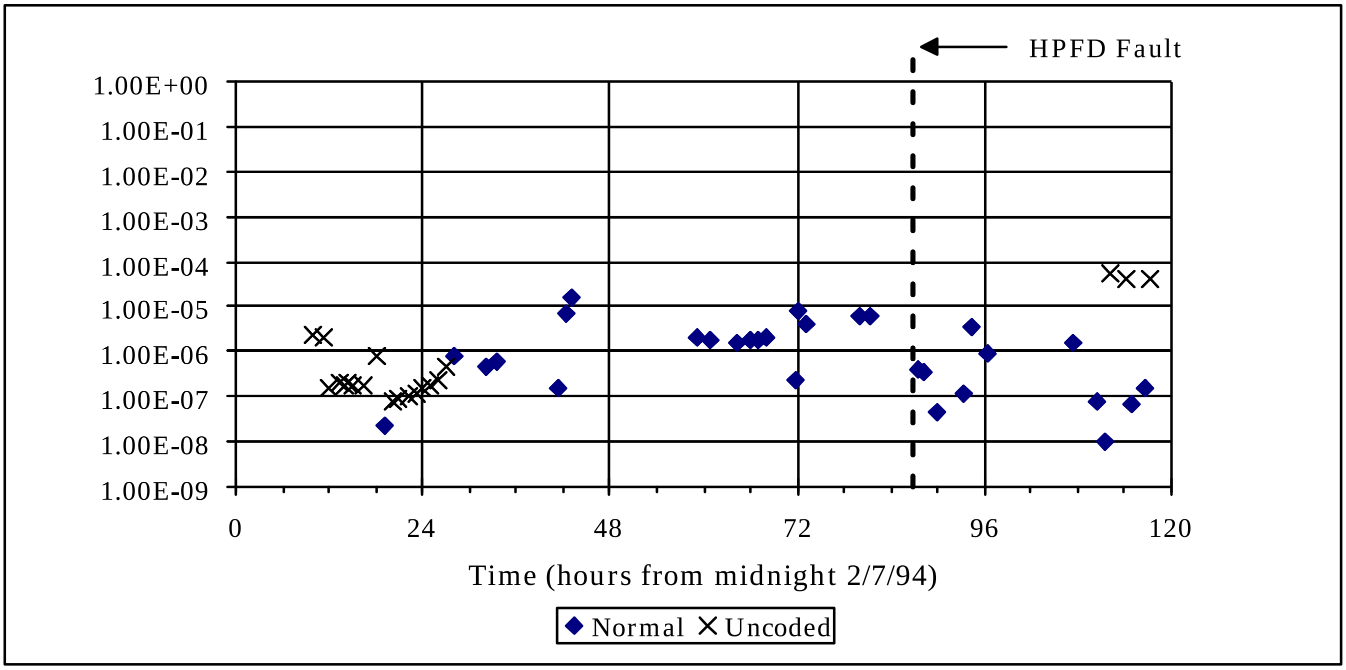

Figure 8. — Bit Error Rate performance measured at the Athens earth station. Measurements are made in loop-back mode and include the full round-trip path.