# DEPLOYING INTERNET SERVICES OVER A DIRECT BROADCAST SATELLITE NETWORK: CHALLENGES AND OPPORTUNITIES IN THE GLOBAL BROADCAST SERVICE

Robert Lindell and Joseph Bannister
USC Information Sciences Institute, Marina del Rey, California, USA

Cheryl DeMatteis, Michael O'Brien, and James Stepanek, and Michael Campbell
The Aerospace Corporation, El Segundo, California, USA

Fred Bauer
SRI International, Menlo Park, California, USA

## Abstract

*The use of direct-broadcast communication satellites and low-profile, portable receive antennas is an effective means to provide communications to military forces in the field, enabling access to a wide range of information sources. Supported by the Defense Advanced Research Projects Agency (DARPA) under the Battlefield Awareness and Data Dissemination (BADD) program, a team was assembled to develop and demonstrate network protocols for the Global Broadcast Service (GBS), a system currently being under development and acquisition by the Air Force Space and Missile Systems Center (SMC). GBS will provide broadcast video, voice, and data from stateside and other information sources to users in a theater of operations. This paper presents the evolving architecture of the GBS Internet and our team's experience with implementing the access, routing, transport, and middleware protocols used in the GBS Internet.*

## Introduction

Two powerful currents in communications and networking have recently mingled together. The integration of direct-broadcast hardware (ground stations, satellites, and "set-top boxes") and Internet software (routing protocols, end-to-end transport, and a large installed base of inexpensive network-centric applications) makes a great deal of good sense for field commanders and troops, who stand to have powerful applications supported by the delivery of information from remote sources. The integration of direct-broadcast and Internet technologies, however, requires the careful solution of technical problems. This paper discusses these problems and reports on our efforts to solve them.

We review early influences on the GBS network architecture and protocols, including work done for a Joint Warrior Interoperability Demonstration (JWID95) and commercial direct-broadcast communication systems. We then discuss the design of the BADD GBS network architecture, which is intended to use open Internet technology that permits interoperation with existing and future systems. At the same time, the GBS Internet architecture exploits special capabilities of GBS, such as native multicast. We discuss also the technical challenges overcome, such as how to utilize effectively links that are unidirectional, how to implement connectionless Internet services over connection-oriented Asynchronous Transfer Mode (ATM) links, how to achieve high throughput with "long fat pipe" (high-latency, high-speed) links, and how to implement desirable services [such as World Wide Web (WWW) browsing and periodically "pushed" data distribution] efficiently.

## Background

Motivated by the success of commercial direct-

*This work was supported by the Defense Advanced Research Projects Agency and the Aerospace Sponsored Research program.

broadcast satellite systems, such as DirecTV and PrimeStar, the Air Force Space and Missile Systems Center (SMC) is engaged in the acquisition of a GBS system, which will provide in-theater communications over satellite. The GBS satellite is expected to be a geosynchronously orbiting communications platform that receives an uplink stream, reamplifies it, and transmits the stream over an earth-coverage or spot-beam downlink antenna. Numerous receiving units are able to pick up the satellite's transmissions by pointing at the satellite a small dish antenna of about 18 inches' diameter and selecting the proper frequency.

The data rate of the GBS is nominally 20 megasymbols per second using Quadrature Phase Shift Keying (QPSK) yielding 2 bits per symbol or 40 megabits per second (Mb/s). After the overhead of Reed-Solomon and convolutional/Viterbi coding the effective data rate seen by the user is about 20 to 30 Mb/s. The receiver has only a receive capability and must therefore use an alternate channel for sending traffic (e.g. a terrestrial link). Indeed, this unidirectional characteristic of the GBS presents a challenge to system designers and implementors, to whom a bidirectional channel would permit a more-familiar milieu for deploying Internet protocols.

## JWID

In support of JWID95, we developed and fielded a transport layer protocol for unidirectional file transfer over a GBS system. Data files — generally large imagery files — were broadcast to receivers using a best-effort approach. To enhance reliability, data could be replicated during the transmission of a file to increase the probability of successful reception. This can be viewed as a simplistic form of forward error correction. Partially received files would be cached and repaired if a subsequent rebroadcast of the file was performed. Given the low bit-error rate (BER) of the system, coupled with these two forms of error correction, correct file reception was achievable with a high rate of success.

If the transport of files were handled in a serialized fashion, large imagery files could delay the transmission of time-critical information. Our transport protocol provided multiple file transfers to occur simultaneously and provided a mechanism to allocate bandwidth across each session. This was accomplished by including a packet scheduler on the sending end to manage the GBS link. Since there was no feedback, on account of the lack of a communications backchannel, the packet scheduler was also used to provide sender rate throttling to prevent overruns at the receiving end.

The transport protocol was implemented using UDP/IP and ATM Adaptation Layer Class 5 (AAL5) services. Both were tested using the GBS system and were shown to provide reasonable utilization of the link bandwidth for file transfers.

Given the success of this transport protocol during JWID95 and the subsequent deployment to Bosnia, one may conclude that this was a good solution for GBS. On the contrary, we would argue that this solution, while adequate for a time-critical demonstration, lacks many characteristics that are important for the final system. First, our solution was proprietary and offered no interoperability with other systems. Second, the protocol was strictly unidirectional, which had many limitations. If the protocol were bidirectional, it would have been possible for end users to request file transfers directly or request retransmission of lost information. Other services, inherently more interactive, such as WWW service or remote terminal access, could not be implemented. Finally, the unidirectional file transfer service we implemented lacked the flexibility that one would expect from such a service. This service required attended operation controlled by a user with an out-of-band channel (often a telephone). Such an approach would require many well-known applications that employ the standard file transfer (FTP) or hypertext transfer (HTTP) protocols to be rewritten. Given the sheer quantity of existing applications, this did not seem to us to be a very attractive option. For this reason we decided that an alternate approach was needed — one based on the use of existing Internet services and the large suite of applications that use them.

## BADD

The BADD program is demonstrating a testbed where advanced battlefield awareness applications can be forward deployed in the field using wireless high-speed GBS data links to access, cache, and process important data sources which are geographically distant from the warfighter. The military has an extensive communications infrastructure which is used in part to develop a large world-wide military Internet. The communications capability of GBS, coupled with the network technology developed during the BADD project, is demonstrating how GBS satellite communications can augment the existing military Internet by providing low-cost and high-bandwidth wireless network links which can be used effectively for certain transport services. In this paper, we will focus on the data dissemination part of the BADD project and, more specifically, how GBS technology can be used as a network component which is a piece of the larger military Internet.

The testbed is comprised of a centralized satellite uplink facility in the Washington, D.C., area which is connected to other important data sources throughout the country using a high speed networks such as the Defense Information Service Network's Leading Edge Services (DISN-LES). Fixed and mobile reception units, known as WarFighter Associates (WFA), receive, cache, and process the data in the field. A WFA consists of a small parabolic dish antenna, commercial set-top box receiver, KG decryption unit, ATM switch, and a UNIX workstation. This testbed configuration is shown in Figure 1. The current BADD infrastructure uses encrypted ATM as the link layer for the Internet Protocol. The usage of ATM was in part motivated by the quality-of-service support offered by ATM. It is envisioned that the GBS data links would be able to offer differential services to different classes of users and applications. ATM was also chosen for the underlying link because the receiving switch could perform on-the-fly filtering of cells based on a particular virtual-connection identifier.

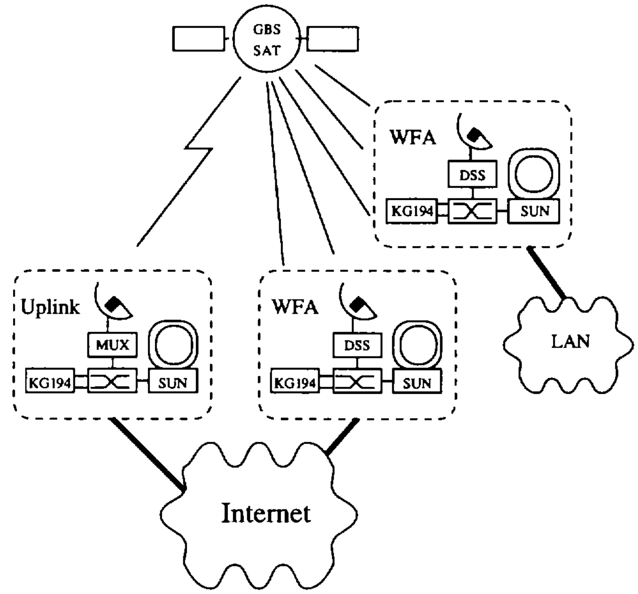


Figure 1: BADD GBS Network Testbed

## The GBS Internet

To develop an IP network over the BADD GBS system, The Aerospace Corporation and SRI International teamed together to develop a working prototype. Aerospace focused on the issues of IP routing, while SRI adapted and tuned a transport protocol for reliable multicast (TRM) for the BADD system [Sab96]. In addition, SRI demonstrated the utility of this multicast transport protocol by modifying an existing WWW proxy application to use TRM [Bau97, Lev97].

The Internet community has recently seen an increased interest in the usage of duplex links with asymmetric bandwidths. This has resulted in part from the emergence of technologies such as cable modems and asymmetric digital subscriber links (ADSL). The GBS system provides an additional twist. Since GBS is in fact a unidirectional broadcast link, an entirely different communications infrastructure must be coupled with GBS to provide full-duplex service. This generally leads to asymmetric bandwidth and, more importantly, asymmetric routing. Current IP routing protocols such as the routing information protocol (RIP) and the open shortest path first (OSPF) protocol do not support these types of links. For example, OSPF initiates its routing algorithm by sending "hello" packets to its neighbors and listening for responses on the reverse link. Recent interest by other research groups in this area has led to formation of an Internet Engineering Task Force (IETF) working group on UniDirectional Link Routing (UDLR).

In the short term, two different solutions for using RIP over the BADD network have been tested. The first solution was developed by researchers at INRIA who proposed simple changes to RIP, OSPF, and the distance-vector multicast routing protocol (DVMRP) [Dur96]. Their solution marked the router advertisement requests that were sent over unidirectional links so that neighboring routers would know to encapsulate the response over the secondary reverse link. The other solution was developed at Aerospace and is known as the Virtual IP Relay (VIPRE). The VIPRE solution would capture outgoing traffic on a unidirectional incoming interface at the link layer. This IP traffic would be encapsulated and relayed to the other host by means of a tunnel through the reverse path. This technique would make unidirectional interfaces appear to be bidirectional at the network layer and permit the usage of routing protocols such as RIP to run unmodified. This hiding of the topology at the network layer can have undesired consequences. Primarily, although IP routing protocols such as RIP run unmodified, they also make the false conclusion that the reception site is a neighbor of the uplink site (removed by one IP hop) even though the actual path of the tunnel may be a complex route through a large IP "cloud". Only the forward link from the uplink to the reception site is really one hop, although both ends look similar at the network layer using VIPRE.

Network links that have both high bandwidth and high delay, so called Long Fat Networks (LFNs), are known to lower the effective throughput which can be achieved using TCP. This is because in LFNs many bits of information can be in flight on the link at a given instant in time. Since TCP employs a flow-control scheme in which the receiver explicitly asks the sender for the next sequence of bytes, one runs into trouble when the "bit length" of the link exceeds the amount of data that TCP can request at any time. Effectively, the LFN's pipe runs dry before the next sequence of bytes can be requested. The GBS system, with its 20 Mb/s rates and 400 millisecond delay, falls into the category of a LFN. Internet standards for TCP-LFN or TCP with large windows have been developed as a solution to this problem [Jac92], yet few implementations exist. An unsupported implementation of TCP-LFN for SUN Solaris was used successfully to demonstrate effective utilization of the link. Standard TCP, with 64KB windows, was able to transmit only 1.3 Mb/s over a GBS link. In comparison, a TCP-LFN implementation, with larger windows, was able to achieve much higher transfer rates as shown in Figure 2. While TCP-LFN implementations were able to achieve high transfer rates in the absence of link errors, other solutions will be necessary to maintain adequate TCP performance in the presence of errors. Standard TCP/IP, which has support for Internet congestion control, will behave poorly if a single packet is lost as part of transmitting the large windows necessary to effectively utilize a GBS link. Loss of a segment will cause TCP to assume congestion has occurred, which triggers a mechanism which will wait for the link to drain and then "go back $n$" segments and begin retransmission of a potentially large part of the window. Figure 2 shows that a Bit Error Rate (BER) of $10^{-6}$ can devastate the throughput of TCP-LFN. This behavior of TCP could be improved with the implementation of selective acknowledgements (SACK) [Mat96]. With selective acknowledgements, only those lost packets would be retransmitted, although more information needs to be transmitted on the acknowledgement communications link to signal this information. This may be an undesirable tradeoff if the acknowledge link bandwidth is quite low, which might be the case if a handheld military radio modem were used. We plan to test the effectiveness of SACK for the BADD GBS system and then determine what supporting acknowledge link rates would be feasible when using SACK.

| Round Trip Latency (ms) | TCP Window Size (KB) | Max TCP BW (Mb/s) | Observed BW With Given Bit Error Rate (Mb/s) | | |
|---|---|---|---|---|---|
| | | | 0 | $10^{-8}$ | $10^{-6}$ |
| 0 | 64 | $\infty$ | 65.7 | 65.6 | 51.4 |
| 50 | 64 | 10.5 | 9.9 | 9.6 | 3.8 |
| 150 | 64 | 3.5 | 3.5 | 3.5 | 1.3 |
| 300 | 64 | 1.7 | 1.8 | 1.7 | 0.7 |
| 0 | 1024 | $\infty$ | 64.9 | 63.1 | 51.3 |
| 50 | 1024 | 167.8 | 64.1 | 63.0 | 3.6 |
| 150 | 1024 | 55.9 | 52.0 | 42.8 | 1.3 |
| 300 | 1024 | 28.0 | 25.3 | 21.0 | 0.7 |
| 600 | 1024 | 14.0 | 11.6 | 9.8 | 0.3 |
| 0 | 2048 | $\infty$ | 65.4 | 64.2 | 53.8 |
| 50 | 2048 | 335.5 | 64.4 | 63.4 | 3.8 |
| 150 | 2048 | 111.8 | 62.3 | 56.2 | 1.3 |
| 300 | 2048 | 55.9 | 49.0 | 34.0 | 0.7 |
| 600 | 2048 | 28.0 | 21.8 | 15.1 | 0.4 |

Figure 2: TCP LFN Performance Results

Using TCP/IP, standard WWW services can be provided over this GBS Internet. Information repositories can be created to allow the warfighter to access weather data, maps, and other useful information. Information needed by one warfighter, such as weather imagery, is likely to be of interest to other warfighters in the area. Thus, information retrieved over the WWW by the warfighter would probably benefit from the inherent multicast capability of GBS to transmit information to more than one reception unit simultaneously and efficiently. Using the caching features already available in WWW proxy applications, this multicast information could be retained in a local cache until the information is actually requested by a user on a reception unit which is distinct from the unit which made the initial request for the information over the GBS WWW system. To prototype this capability, we need to substitute a reliable multicast transport protocol for TCP, as the underlying transport protocol used by HTTP in the GBS Internet. TRM was chosen for this purpose and integrated with an existing WWW proxy application to create a multicast-capable proxy. To facilitate the usage of off-the-shelf WWW servers and browsers, these TRM-based proxies were inserted at each end of the GBS link. The proxy at the uplink would talk to other WWW servers using standard HTTP over TCP except over the GBS link, where it was capable of talking both standard HTTP and HTTP over TRM. Likewise the receiving-end proxy would talk standard HTTP to a browser, but would forward requests using HTTP/TRM over the GBS link. This configuration hides the use of multicast and TRM from both server and browser applications. In addition, it is possible for the proxies to provide caching of information both in a central fashion at the uplink and locally at each reception unit. If multicast data in a local cache was never accessed, it could be removed when necessary. A prototype of this WWW system using HTTP over TRM was demonstrated on the BADD testbed.

## Future Work

Additional work needs to be performed to expand the multicast-based application suite available over the GBS Internet. This would include multicast-based FTP services as well as MBone-type teleconferencing applications. Multicast routing protocols such as DVMRP need to be tested using VIPRE or modified similar to that proposed by INRIA.

Resource reservation and quality-of-service support need to be developed for the GBS links. Our team plans to investigate the usage of the ReSource ReserVation Protocol (RSVP) on the BADD testbed. In addition, quality-of-service capabilities could be mapped to the underlying ATM link layer to provide efficient enforcement of resource-allocation limits. Similarly, the application of RSVP with other packet-scheduling algorithms could be tested using link layers such as Ethernet.

## Conclusion

The core system has been implemented and includes standard Internet protocols adapted for unidirectional routing by means of new routing protocols and a scheme for using reachback links called Virtual Internet Protocol with Reachback (VIPRE). The so-called Long Fat Network Transmission Control Protocol (LFN TCP) and the Transport Protocol for Reliable Multicast (TRM) are used to support reliable, flow-controlled data exchange, the latter allowing for data to be delivered simultaneously to a number of receivers. Other services, such as Web access, are supported by caching schemes that place frequently requested information at sites in the theater of operations to promote low response times and reduce congestion on the satellite links. The system is operational and in use by military commands. Extensions to the core system, including improved network management, enhanced security, and support for quality of service, are currently being designed, implemented, and evaluated.

## References


[Bau97] Bauer, F., Brown M.J., A Reliable Multicast Transport Protocol for a Global Broadcast Service-based Network, *Proceedings of Milcom 97*, Monterey, California, Nov. 1997.

[Dur96] Duros, E., Huitema, C., Handling of unidirectional links with RIP, *Work in Progress*, INRIA, Sophia-Antipolis, Mar. 1996.

[Jac92] Jacobson, V., Braden, R., Borman, D., TCP Extensions for High Performance, RFC 1323, LBL, USC/ISI, and Cray Research, May 1992.

[Lev97] Levin, L. J., Multicast Aware Proxy Servers, *Proceedings of Milcom 97*, Monterey, California, Nov. 1997.

[Mat96] Mathis, M., Mahdavi, J., Floyd, S., Romanow, A., TCP Selective Acknowledgment Options, RFC 2018, PSC, LBNL, and Sun Microsystems, Oct. 1996.

[Sab96] Sabata, B., Brown M.J., Denny, B.A., Heo, C., Transport Protocol for Reliable Multicast: TRM, *Proceedings of the IASTED International Conference on Networks*, Orlando, Florida, pp. 143–145, Jan. 1996.