Title

Dynamic Access for Satellite-Based Networks

Mike Rupar, Naval Research Laboratory Code 5554

Abstract

Dynamic Access can be defined as a process of using application-layer and instantaneous network traffic characteristics to dictate the allocation of physical layer (the satellite bandwidth) assets to a particular node in the satellite network. It consists of control techniques, processes, and procedures effectively plan, monitor, control, and configure/reconfigure those assets in instantaneous dynamic response to user priority and utilization while optimizing the overall throughput of the system. The methodology described herein is a non-proprietary system developed at the Naval Research Laboratory (NRL) to support more effective use of satellite assets in the US military.

It is predicted that the application of Dynamic Access to satellite network management would result in an overall increase in throughput as great as 20%. An application of Dynamic Access will allow for real-time traffic assignments adjusting to both time-critical and/or priority needs as required, while allowing lower-tier traffic to utilize idle bandwidth when available.

The benefits of a Dynamic Access control system on a satellite-based network are: increased throughput to users based on better utilization of unused network and satellite channel capacity while increasing the cost efficiency for satellite assets both owned and leased (e.g., DSCS, Intelsat) and maintaining critical prioritization and QoS for users. These benefits add up to provide more bits per dollar.

Introduction

The past few years have seen significant progress in networked shipboard communications with the augmentation of commercial satellite connectivity from Inmarsat-B (between 32 and 64 kbps per ship) to Challenge Athena (up to 1.544 Mbps per ship). The insertion of advanced local- and wide-area networks and the Advanced Data Network System (ADNS) have helped to dismantle the legacy “stovepipe” system infrastructure that has dominated shipboard communications since the advent of radio. The US Navy currently makes use of more than thirty commercial satellites and ten government/DoD satellites to satisfy its mission requirements throughout the world. The majority of these assets are “bent-pipe” transponders, which serve as a radio-frequency relay in the sky and perform no traffic management. Larger deck combatants (such as carriers and flagships) are connected as never before, allowing modern applications to better support the fleet’s warfighting and management capabilities.

Small-deck combatants, aircraft, and forward-deployed forces are currently communications-limited compared to the larger combatants. Even where it exists today, HDR network connectivity for this class of communications-disadvantaged user is very expensive due to the static nature of which its use is provisioned. Forward-deployed forces with a HDR requirement currently transport as many as 2 trailers and require up to an hour to setup a T1-class

service. Once the circuit is set up, its static configuration requires the dedicated use of satellite resources regardless of the utilization of the link. Thus, the expense of the current system is realized not only by the cost of the equipment but also by the lost opportunity for other nodes to communicate using that unused bandwidth.

One difficulty when operating with the communications-disadvantaged users is that in order to support their smaller aperture terminals, the effective throughput of the transponder must be limited. This is because much of the satellite resource (both power and bandwidth) is allocated to support the smaller aperture. This limits the resource sharing capability with the other nodes that wish to communicate. Another common situation is that satellite resources are often statically assigned to the larger platforms using a "platform prioritization" scheme that is not rapidly re-configurable. This leaves little flexibility in the satellite resource for reallocation to support the disadvantaged user.

The management of satellite assets for the Department of Defense is conducted through the Defense Information Systems Agency (DISA). DISA allocates bandwidth to specific users at sea and on land based on mission requirements for that user and other users in theatre, the satellite terminal capabilities of the user, as well as available satellite resources for the operational area. Such allocations are assigned after a formal and lengthy process (on the order of weeks) that does not lend itself well to the dynamic nature of a battlegroup or of a crisis situation. These bandwidth assignments are static, and often do not meet the peak requirements of the platform in question. There also may be significant periods where minimal and/or low-priority traffic is carried over the link.

The Transformational Communications Study (TCS) currently on-going in the DoD calls for "packet-switched" traffic between assets, and would be best served by a operational structure supporting dynamic access to the satellite bandwidth. Most current and evolving network applications (intelligent sensors, combat systems, coordination and management, and distance learning) often do not require a constant bit-rate service, and often transmit data in a burst-like fashion. Statically assigned channels that support this type of traffic are often under-utilized, up to 60%. However with the current scheme of "platform prioritization" and using systems that are not rapidly re-configurable, few commanders would be willing to give up their allocation for other node's use.

It is critical that as satellite networks evolve, they be provided with a means for allowing the management of network traffic more precisely and on an interval that more closely matches the instantaneous priority and utilization of the network. ONR guidance states that "Toward this end, these IP technologies need to… address issues related to multiplexing mission-critical military multimedia traffic and for managing the total RF bandwidth available on naval platforms and within a theatre in order to enhance communications efficiency and affordability as envisioned by JMCOMS/ADNS."[1]

Technical Background

NRL has been supporting fleet applications and demonstrations with satellite and wireless connectivity since 1995 [2]. NRL has long been an advocate of the interrelationship between the layers within the OSI model when architecting satellite networks. The Information Technology Division at NRL has a history of developing multi-layer solutions for satellite and wireless connectivity. NRL has continued to investigate and research evolving methodologies for improving network performance and examining and testing algorithms to improve TCP and data link performance over satellite channels [3] through participation in various commercial (the IETF) and governmental partnerships (the project-118x and project-154 team ventures with NASA Glenn Research Center).

There are different methods used to assign satellite resources to particular users. The two most common methods employed in MILSATCOM are Demand Assigned Multiple Access (DAMA) and Fixed Access. In fixed access a particular "slot" (either frequency when using frequency-division multiple-access or time when using time-division multiple-access) is manually and statically configured for a particular user in the network. Once configured, that satellite resource is no longer available for use by others whether or not there is information to send. In DAMA, "slots" are allocated according to a request received via an order wire system. Users make requests for resources and a number of "slots". If available, the system assigns "slots" that are made available to that user for the duration of that circuit. In these instances there is significant wasted bandwidth within the system that could be dynamically shared with other users.

Within the Internet Engineering Task Force (IETF), the research emphasis is predominantly based on static, fixed-bandwidth networks [4]. These areas of research address situations in which the offered load to a network exceeds the network's capacity. In these situations, the network layer is controlled based on offered load and the bandwidth is assumed fixed.

The proposed methodology to maximize channel throughput is to automatically directing the use of physical layer assets (channel capacity) to a satellite network based on application parameters, channel loading and real-time network requirements to maximize channel throughput. The fleet currently assigns traffic in a static, fixed bandwidth manner, and even within the IP framework established fixed bandwidth (circuit emulation) to serve particular priority applications, such as VTC.

A program of record exists that can be used to improve the utilization of the satellite channel within a particular node at sea. The Advanced Digital Networking System (ADNS) has been implemented on naval platforms to allow for a greater utilization of all available bandwidth allocated to the platform. It is router-based and handles most (but not all) of the satellite-based traffic from a Naval platform. However, significant amounts of bandwidth on SHF and Challenge Athena channels are still parsed out in fixed amounts, and both these fixed allocations and the intelligent router-based ADNS traffic are passed through time-based multiplexers which further partitions the bandwidth in a non-dynamic fashion.

Packets at each layer in the protocol stack, particularly the transport, network and link layers, carry information that describes the traffic being carried, and/or could be used to store

such information. A PC-based routine that has access to packets at one or more of these layers could be employed to access the priority and QoS requirements of the information being carried by that packet, and could also be used to access particular application requirements for bandwidth in the near future. Locating such a routine at a hub or teleport and monitoring traffic going to and coming from the users in the fleet would enable network-wide decisions to be made in real time, resulting in significant improvements in bandwidth utilization for all.

Realization of such a capability is complicated by the nature of applications which, while requiring a particular level of bandwidth to pass peak traffic, often operate at much lower throughput. This applies to commercial networks as well- off-the-shelf applications typically run far below peak traffic levels for a majority of their use. The peak levels for these application programs are often on the order of hundreds of kbps. To effectively re-distribute unused bandwidth, one would need to recognize the ebb and flow of traffic for particular applications, responding in a rapid fashion when the program required the full allocation for peak bandwidth.

Such a routine would also need to be cognizant of the capacity of the various transponders (not necessarily on the same satellite) being utilized. The configuration of the satellite transponder often varies from specification, due to additional output power backoff added by the satellite vendor (for commercial services), atmospheric conditions, and degradation and of the satellite transponder. As a transponder is loaded with more and more users, its output power is distributed among those users, and not necessarily in an even fashion: optimum network performance would dictate power control mechanisms are enlisted on all network users. Such controls are already the norm for TDMA and CDMA networks.

Technical Approach

A methodology for Dynamic Access as it applies to a satellite-based hub-and-spoke network configuration has been developed at NRL. Dynamic Access can be defined as a process of using application and instantaneous network traffic characteristics to dictate the application of physical layer (satellite bandwidth) assets to a particular node in the network.

The project will be to develop a method for determining the bandwidth applied to a particular application at a node based on the following conditions:

- Available bandwidth/radiated power at the transponder
- Available network capacity/radiated power at the satellite terminal
- Priority of the user/application
- Quality of Service (QoS) of the application
- Condition of the overall network as determined at the hub

There are two methodologies to be developed here: first, to allow the hub to sense both networked throughput and “application qualifiers”, (i.e., QoS and prioritization of the application being run) from each spoke (to/from each user terminal). Secondly, to develop the capability for the hub to reconfigure the network based on channel utilization, addition/subtraction of new applications/users, and demands for prioritized access. The first capability will reside

predominantly in the network layer and higher while the second capability will take place predominantly in the physical layer. The dynamic access methodology as a totality will be a seamless integration of these two areas of focus.

Transponder Management

A critical requirement for Dynamic Access is a thorough and continuous monitoring of the status of the transponder assets supporting the satellite network. The condition of the transponder system on the satellite is of primary concern. If it is a "bent-pipe" satellite, acting as a straight or hard-limited amplifier, then issues such as phase and amplitude linearity are a concern. The transponder must be characterized for gain and linearity performance during normal operation and when compression and saturation are reached. For a commercial satellite vendor who is interested in supporting other users on opposite polarizations, the total gain of the transponder assembly is often throttled back to ensure that the system is not overdriven. This "output backoff", or OBO, is a degradation of the overall performance of the network, and will result in less throughput through the satellite.

Management of the bent-pipe transponder requires a continuous monitoring of the overall power carried, the integrity (e.g., BER) of the carriers on the transponder and the location of the various users within the satellite footprint, which affects their performance directly, and therefore the overall capability of the network. This is most effectively conducted at the hub, but will require inputs from the various user terminals ("spokes") in the network. Real-time decisions on the total throughput possible through the satellite will dictate the ceiling for resources to be allocated to the individual users/traffic.

Transponder management is complicated by weather events, particular rain cells, which can severely affect the link characteristics from all or perhaps just one of the users in the network. (A rain fade at the hub would affect all users regardless of their location.) Weather events are the most dynamic variable in assessing the capacity and integrity of the transponder, and the only one that cannot be controlled.

Other variables that can affect the performance of the transponder are the power and polarization of the individual carriers sharing the bandwidth. Individual signals that run too hot can upset the effective power distribution on the transponder, and carriers running too cold will break particular threads in the network. Therefore, some sort of power control must also be established, preferably at the hub, to ensure the peak performance of the network.

If the satellite supporting the network is a processing satellite, there are some additional issues for any dynamic access. A processing satellite will have packet-level control and monitoring capabilities, and would be an excellent platform to implement a Dynamic Access capability. This is particularly true for networks of the future, which will be mesh configurations as well as hub-and-spoke. Nevertheless, some sort of control from the command center (whether a hub or not) would be required, particularly if the flexibility of the processing transponder is limited either in dynamic range (for overcoming weather events) or in it's ability to command nodes in the network to reconfigure.

The Network Executive

The focal development for NRL's Dynamic Access is a "Network Executive", a routine that operates at the hub of a satellite-based network, and directs the management of the individual links to the user terminals. It is the Network Executive that will make decisions regarding the allocation of bandwidth to individual links, and then implement those decisions. A graphic of this is shown in Figure 1, below. The Executive will reside on a PC at the hub, with connections to all elements of the communications stack.

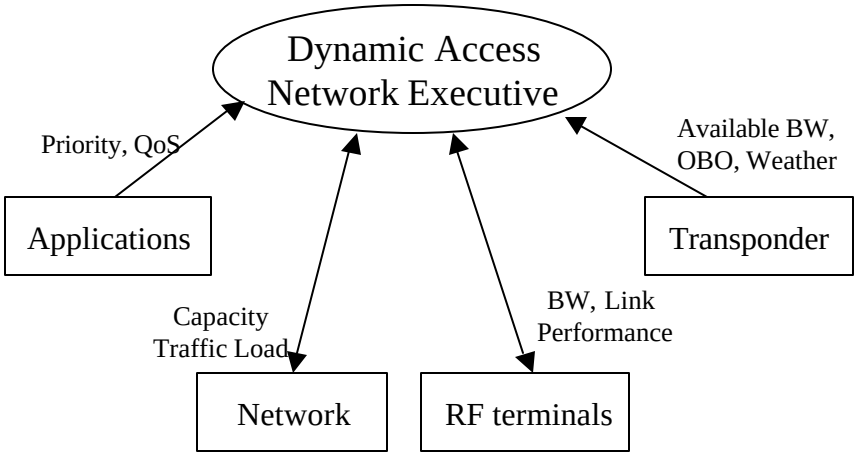

Figure 1: Representation of the Network Executive Routine and it's integration into a satellite-based network. The RF terminals include the hub and the user terminals (or spokes) in the network.

The Network Executive (or just "Executive") receives a number of inputs based the condition of the traffic requirements of the network (discussed below), and on the capacity of the transponder. The executive will act to reconfigure the network based on two different classes of stimuli: the condition of the physical network of the links through the transponder, and on network/traffic requirements.

Physical link events, such as weather and link performance for individual user terminals, need to be accounted for automatically by the executive, and act on a rather static rule set. Priorities and trade-offs are set for maintaining effective connectivity to individual terminals, the bandwidth to each terminal, and the capacity load on the transponder itself. If a particular node is affected by a weather event, the Executive needs to sense the change and act accordingly to increase power, change the bandwidth or change the coding for that carrier. If a power control algorithm outside of the network is implemented (e.g., a modem-based power control that maintains a particular Eb/No with set limits), that too must be monitored and tailored to meet the needs of the overall network.

Network/traffic events are conditions where a sudden surge or drop in bandwidth to support a particular application (or entire node) occurs. The challenge is to respond quickly

enough so that the priority users (or QoS requirements) are not compromised. The Network Executive will utilize existing diagnostics within the network to assess such dynamics in traffic load, and based on a rule set, reconfigure the assets to accommodate the surge (or drop) in bandwidth necessary to support the application(s). One such event could be a video-teleconference: allocating bandwidth to support this application continuously would result in wasted capacity when the VTC is not being conducted. However, when the VTC is initiated, the bandwidth must be established to support that as soon as possible (based on it's priority relative to other traffic throughout the network). Dynamic Access can act to accommodate the surge in traffic.

Implementation

The following is an outline for implementing a dynamic access capability on a hub-and-spoke network.

Transponder management

A routine to assign physical carriers to a transponder requires the following:

1. Apriori knowledge of the power capacity of the transponder
2. Monitor and control of all modems and output power for each user terminal and the hub participating in the network
3. Performance feedback from the user terminals (assuming a hub-based routine).
4. A framework for carrier assignment

NRL has developed a routine that will implement these components. This routine allows a controller to reconfigure carriers for an entire network from the hub. The hub routine will be able to monitor the link performance at both ends of each spoke, and compare that to link budget calculations to identify potential discrepancies. In addition, a continuous assessment of the capacity of the transponder is conducted to monitor performance. This will involve the cooperation of the satellite controller (whether DoD or commercial) to determine the peak power capacity of the transponder, as well as any potential non-linearities, particularly at the edges of the transponder. The aggregate carrier power across the entire transponder will be continuously measured and compared to peak. In addition, beacons from the satellite must be monitored to insure that weather events at the hub do not affect the accuracy of the transponder power values.

Network/Traffic Management

A routine to manage the requirements to support an application or network with sufficient bandwidth requires the following:

1. The ability to receive and respond to requests for additional capacity from the LAN on ship.
2. The ability to monitor the type of network traffic being generated, Quality of Service (QoS) contracts (CBR, reliability, etc.), the priority and volume of

offered network load, and the availability of alternate channels (e.g., UHF) to accommodate an off-loading of low-priority traffic.

3. The ability to assign an ultimate priority to applications and manage the network in a manual fashion if the need arises
4. An emphasis on the use of standards-based protocols where they exist, such as the use of existing QoS protocols for identification of priority.

NRL primarily focused on standard protocols for TCP, IP version 4, and ATM to identify methodologies for identifying, monitoring and directing application traffic, although there has been some consideration for IPv6. One necessity is to keep the latency at a minimum – the goal is not to add appreciable delay to the routing and processing of traffic. Methods for the forwarding of bandwidth dynamics to the network executive as rapidly as possible have been identified.

Testing and Evaluation

The testing of particular components of this methodology has been limited to controlled experimentation at NRL in Washington DC. This methodology is currently under review by SPAWAR PMW-176, the Joint Systems Integration Command (JSIC), and the commercial providers Loral Skynet and Panamsat. The results are as follows: (to be finished by Nov.)

References

1. Communications Technology FY00 Tentative Guidance, from the Science and Technology Requirements Guidance (STRG) N091, August 1998.

2. Michael A. Rupar, Adrian S. Eley, Timothy L. Krout, "Demonstration of High Data Rate and Medium Data Rate VSAT Communications using the Global Broadcast System (GBS) Transponder," NRL Memo report NRL/MR/5550-97-7921, March 29, 1997.

3. T. Krout, M. Solsman, "The Effects of Asymmetric Satellite Networks on Protocols", MILCOM 98 Proceedings, October 18-21, 1998.

4. Jamalipour, Abbas, "Broad-Band Satellite Networks – The Global IT Bridge," IEEE Proceedings, Vol. 89, No. 1,pp 88-104, January 2001