PROVIDING UNSIGNALED CRITICAL SERVICES FOR A FUTURE SATELLITE PACKET NETWORK

Kathleen Nichols
Pollere LLC
Menlo Park, CA 94025

Max Williams
ITT Corporation
Herndon, VA

ABSTRACT

A Future Satellite-Based Packet Network (FSPN) includes highly mobile users operating in demanding environments. Packet-switched traffic encompasses applications that are robust in the face of these challenges but also includes applications whose operational utility requires a relatively low delay, jitter, and loss performance. On-path, in-band signaling protocols are inappropriate for many terminals. For these, minimum delay, low loss transfer characteristics are provided by Unsignaled Critical Services (UCS) for which aggregate usage levels are communicated out-of-band. UCS need not be strictly static; time-of-day constraints are available, allocations can be invoked through network management interfaces, or short-term network managment interaction can change allocations.

UCS are implemented through the use of IP Differentiated Services, specifically the Virtual Wire Per-Domain Behavior, but includes resource allocation interactions between the IP network layer and the Per-hop Behavior provided by the link layer in a manner unique to satellite systems. The network is further constrained by the use of encyrption which limits packet inspection techniques to the in-the-clear packet header fields only. The architecture of UCS is described, issues for users and providers of UCS are explored, and issues in the network-link layer interactions are examined. Current work on UCS is focused on these issues.

BACKGROUND

Packet networks make it possible to assign resources flexibly, responding the varied needs of military missions. In circuit networks resources are assigned in a three stage process of dimensioning, provisioning, and signaling, where the former two are long-term procedures, on the order of months to years, and the latter is short-term and has specific and limited capabilities. There is a large gap between mission provisioning, at the scale of months, and signaling individual calls in real-time. Many of today's military operations fall into that gap. Packet networks open up a wider range of allocation approaches that can be matched to specific military use patterns and can function under demanding conditions where signaling in infeasible. Packet networks have a capacity planning step, similar to circuit system dimensioning, where a particular configuration of network devices is planned that results in bounds on resources available. Unlike a circuit system, assignment of those resources to particular missions and terminals can take place along a range of time scales, from capacity planning time to formal requests in advance of use to an off-path messaging of FSPN NM entities immediately prior to use. Further, unlike circuit signaling, the requester need not be located at any of the affected terminals.

A future Satellite-Based Packet Network (FSPN) must deliver services to terminals that span a range from highly mobile and cost-constrained to fixed-location, high complexity hardware. These terminals are owned by various organizations but contain both network edge elements for FSPN and for the attached networks, thus are jointly managed. For highly mobile terminals in particular connectivity may be intermittent and location may change throughout a period of service. Despite these difficulties, mobile terminals have applications which require relatively (with respect to satellite network constraints) low-latency, bounded-delay communications services. In addition, applications may be multicast and, for security purposes, minimal information about the importance level of the traffic can be exposed. Use of in-band signaling protocols (e.g., RSVP-over-DiffServ, ARSVP) for such a service is inappropriate due to their poor match to most of the intended uses of FSPN (which include UAVs and mobile units maintaining radio silence) as well as security, cost, and network considerations of maintaining a signaling stream in the presence of link intermittency, terminal mobility, changing membership of the multicast groups, and fine-grained tracking of resource changes. Instead a low-latency, bounded-delay, low loss (no congestive loss) service of a peak bandwidth for specific terminal pairs, lists of terminals, or unspecified terminals can be negotiated in Service Requests and specified in a Service Level Agreement (SLA) with specific Service Level Specifications (SLSs). SLAs for these services are established “off-line,” i.e., outside of run-time, along with limits on their use. A sophisticated deployment can mix quasi-static service levels (changing only on time-of-day) with out-of-band notification methods such as secure web interface, SIP to network management entities, and other nework management interactions. These Service Classes are called Unsignaled Critical Services (UCS) and their construction is described in this paper.[1]Though allocation may be relatively static between FSPN and an attached network and only at the granularity of the unencrypted packet data available to FSPN, attached networks may dynamically control the allocation by restricting its use in a number of ways, including per-flow or per-session signaling that is strictly local to the attached network and makes use of packet data that is opaque to FSPN.

UCS is based on a decade of commercial practice for VoIP and video conferencing traffic. A low-latency service that is a forerunner of VW [RFC2638] was implemented in demonstration testbeds at SuperComputing ’97 and as part of the DOE 2000 effort. Deployments of VoIP in enterprise networks [CSCO] and ISP circuit replacement services [XIAO, TLKMP, IPQ] all follow essentially this methodology. While the mechanics of delivering low-latency packet services are well-established for terrestrial networks, the challenge for UCS is to adapt this body of work to the FSPN environment.

ADAPTING UCS TO THE FPSN

FSPN uses a Differentiated Services architecture [RFC2475] with Per-Domain Behaviors [RFC3086] to implement the edge-to-edge treatments that appear as distinct service classes in the SLAs. A Per-Domain Behavior is the expected treatment that an identifiable group of packets will receive from "edge-to-edge.” A PDB is used to provide each externally differentiable edge-to-edge treatment that is realizable and required. Differentiated Services use in-the-clear packet header fields as well as ingress terminal information. SLAs are maintained by monitoring and policing these identified packets to hold them to particular rate and temporal characteristics per DSCP (DiffServ Codepoint). Boundary nodes use DiffServ edge functions to control the initial entry of DSCP marked packets to the domain. Utilizing these mechanisms UCS are admission controlled to ensure there are no congestive packet drops.

[1]“Unsignaled” refers to no use of on-path signaling and the fact that off-path signaling is optional.

Figure 1 shows the FSPN domain boundaries and subdomains. The FSPN’s DiffServ edge is in the terminal; this is the locus of enforcement on ingress traffic covered by SLSs. Since terminals may be compromised, ther is an additional enforcement boundary at the satellite edge of the space domain. PDB use is tracked by a Bandwidth Broker (BB) [RFC2638], such as MultiService Forum’s Bandwidth Manager specification [MSF], Juniper’s SDX product, or Operax’s bandwidth manager product. FSPN use of a distributed BB, one hop from each satellite, permits localized, robust operations. Terminals receive resource allocations from the BB and their attached networks use the terminal to perform edge functions or have their own Border Routers (BR) attached to the terminal. When a terminal attaches to FSPN, the BB sets the terminal’s Traffic Conditioners (TCs) to achieve the appropriate edge conditioning. The BB also sets TCs at the interior Space Domain boundary to admit only conforming aggregates. Statistics are taken at the terminal and reported to the BB and used to monitor FSPN’s performance against the SLAs.

Figure 1. DiffServ Boundaries and Sub-domains

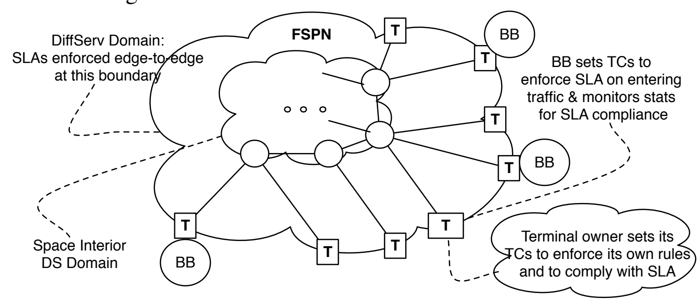

UCS is an aggregate service with packets identified by encrypted source id, particular ingress, encrypted destination id, marked with DSCP as per SLA and sent shaped and rate-limited (which may be performed by the sender or at the terminal). The service can appear as circuit emulation, specifying a particular rate to be delivered with bounded delay subject to incoming packet streams obeying peak rate and packet size limits. The bounds provided by the SLAs make it possible to size playout buffers to reproduce a constant rate stream. Application uses for such a service can encompass voice, including push-to-talk, sensor-to-shooter applications or control of UAV systems.

FSPN is provisioned for particular amounts of each PDB across the space backbone and amounts are further limited on access links. Attached networks submit Service Requests (SRs) with specific parameters; FSPN checks authorizations, identifies the PDB that meets the SR, and returns an SLS indicating the DSCP to be used, rate and burst limits at entry, and any other relevant parameters (e.g. packet size limits, source/destination limits, time interval). To configure the service FSPN sets its edge devices to police to the SLS. FSPN dictates a rate limit and other traffic characteristics; attached networks comply or packets are dropped. Particular destination ids may be designated or may be wildcarded. Terminals may provide shaping to rate. Packet size may be policed and special agreements may cover packet remarking. Run-time events may cause deallocation of configured services which can be communicated through near real-time Network Management Compliance and Status (NMCS) reports. This approach has advantages over on-path signaling, including the ability to decouple the service requester from the terminals receiving service (particularly useful for intermittent connections), local response, and the ability to provide a service that is not path-specific and path-fragile, allowing IP routing to robustly provide the best route.

The Virtual Wire (VW) PDB [VWPDB] provides mathematical and network background for constructing the PDB needed to provide Unsignaled Critical Services. The basic concept is to ensure that the traffic aggregate runs in exact flow balance; if a new packet is clocked into the domain every T_{vw} time units the network guarantees that at least one packet departs during the same interval. Flow balance guarantees that no backlog of packets can build up in the domain thus the jitter seen by any packet is strictly bounded to T_{vw}(figure 2). Particular challenges of adapting VW's PHB requirements to FPSN are discussed in later sections.

Figure 2. VW PDB across a Network Domain

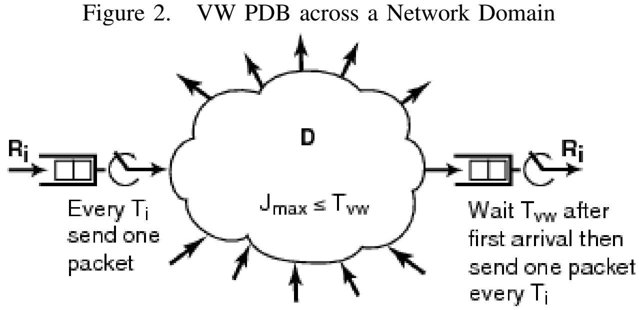

T_{vw} is roughly analogous to a circuit frame. As figure 3 shows, it is only necessary to keep each packet within its "frame" as it transits the network domain; the interpacket spacing may vary.

The amount of VW bandwidth that can be offered while

Figure 3. Packet timing using VW PDB

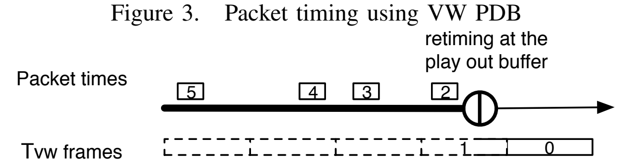

meeting the desired delay bound is computed during capacity planning. Provisioning results in particular settings of packet schedulers in routers and hard limits on the amount of VW that can be parceled out in response to SRs. Requests on behalf of a terminal or group of terminals must fall within the SLA. The BB will further check current conditions and relative priorities before making an allocation. Allocations are recorded in access profiles that are held in the BB and configured into the traffic conditioners of routers when terminals and their service profiles are active.

Provisioning, Allocating, and Using UCS

A VW aggregate on the order of 500 Mbps is expected on the space domain internal links, with allocations made in multiples of some Basic Allocatable Rate (e.g., 32 kbps). Allocations from the VW PDB must fit within the space domain's capacity but must also be supportable on the space-ground "tail" links they traverse, a more stringent limit which will be partially discussed in the next section. FSPN's space domain is topologically quite simple, so VW allocation in the space domain could reduce to allocation against a 500 Mbps total. Though possible to bookkeep full paths between individual terminals, the efficacy of such an approach is doubtful as the space path is simple and a large amount of usage is expected to be multicast. The initial approach is simple allocation with future study of a more complex allocation strategy. (Allocation strategies can be tuned after deployment.)

SRs are used to allocate VW during capacity planning. The SR will also contain information that guides planning in how to allocate the VW PDB, specifically number of terminals involved and whether the usage will be point-to-point, have multiple specific endpoints, or be completely general. In Planning, there is extensive information from the SR about priority and other usage plans. Not all allocations will be committed as the associated terminals may not be attached immediately. Some of these allocations may be committed (i.e., configured onto network elements) in the course of normal network management updates. Further, some capacity may remain to be allocated and committed in response to short-term events. Capacity planning sets limits on Traffic Aggregates while allocation designates a portion of a Traffic Aggregate for an identifiable (and permitted) substream for a specific duration (and may include overallocation).

Traffic conditioners (TCs) at the edge are an integral part of Network Management (NM). Configurations may be changed at run-time in response to security threats, environmental conditions, or operational priority. Short time scale fluctuations in total rate available within a beam are handled via Layer 2 mechanisms. Shortfalls that persist are reported through NM which may trigger reconfiguration.

The BB keeps a record of allocations and commitments for each PDB for each source and is kept updated on the current network map and state. Records are updated at time or event intervals, pushing out configurations of new allocations, de-allocating expired allocations, producing necessary messages and configurations for allocations changed due to network events or pre-emptions. Allocation records are to be multiply indexed (e.g., by source destination, link) and indicate commitment. These stateful records are used to set classifiers and TCs at the ingress and these may be refreshed periodically. At time intervals or on events configurations are pushed out, expired allocations are deallocated and any necessary NMCS is generated. Changes in network state may cause reallocation.

Requests for allocation can be evaluated at planning time or at other times so long as NM has determined they are within an existing SLA. Requests are evaluated against the ingress link, the space domain, and the egress link, if link specific. Short-term requests may be made via any method that NM finds useful, including secure web interface, SIP messages, or any secure network-to-network interaction. FSPN NM then calls upon the BB to make an allocation and instantiate the commitment.

When allocations are more general, an overallocation factor is used. A conservative model ensures that the sum of all allocated VW Basic Rates is less than the total VW rate of a space domain link. This ensures that terminals can move and multicast can be employed without readjustment of rates. This rule can be relaxed (an Erlang model may be developed) with measurement and experience. Allocations on uplinks and downlinks use a “hose” model [HOSE]. The parameters used will evolve with measurement and experience but the allocation model is being developed to take these factors into account.

Multiple Networks and Partitioning

Encryption of packets means FSPN has no way of matching packets to specific user-side flows and cannot identify flows beyond the aggregate information of the in-the-clear packet header information and packet ingress. (Nor should it, to meet security requirements.) Service networks always have more specific information about individual packets and flows that is needed to enforce fine granularity decisions on packet traffic. The FSPN controls access to its resources at an aggregate level; the entering traffic must be consistent with the SLA in order to ensure the integrity of all SLAs provisioned on FSPN. The service network ensures that its traffic conforms to the SLA by discarding, delaying, or remarking packets as needed. The attached network uses its knowledge both of its own current network objectives and of the nature of each flow to enable this process. An attached network may use signaling, e.g. RSVP, to control access to its UCS virtual network, but the signaling is interior to its network (though signaling messages maybe be sent through the satellite network in an encrypted packet, opaque to FSPN) see figure 4.

Figure 4. User Network using UCS

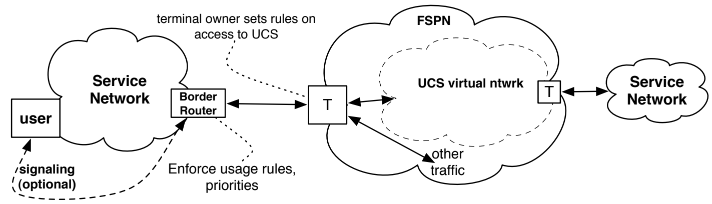

When multiple networks share a terminal many configurations are possible but all are based on the premise that FSPN has an enforceable, monitorable SLA with the terminal owner; the terminal may have enforceable, monitorable SLAs with the attached networks. Figure 5 shows the case where three networks are attached through a BR pass through an encrypter, and then a terminal. The attached networks, the BR, and the terminal may all have the same administrative owner (like a large enterprise model) or some of the networks may have a different owner and have an SLA with the administrative owner of the Border Router (like a Tier 2 provider). The BR’s traffic conditioning is set to properly discriminate between the attached networks according to the agreement with them and its packet schedulers are set to conform to the terminal SLA with the FSPN. For UCS, the BR will ensure that the packets identified for UCS (by DSCP or including more extensive packet header or port information) are conforming to the individual agreement the terminal owner has with each attached network. Once these packet flows are queued for output in the BR, the packet scheduler for the queue used by the PHB ensures that they are sent in compliance with the terminal-FSPN SLA (the combined rates should be less than or equal than the terminal-FSPN SLA rate but it may be necessary to perform some shaping of the output stream).

Figure 5. Multiple Networks

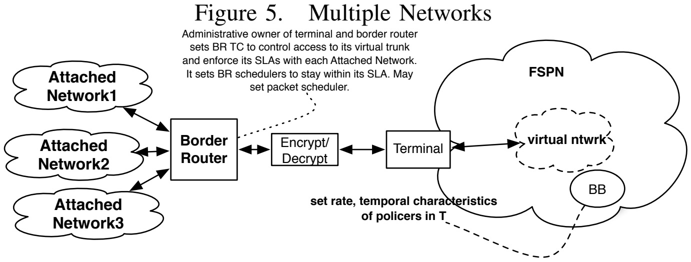

As an alternative, multiple networks may be attached directly to a terminal. In this case, the terminal owner can set the TCs for each ingress port (and, may set packet schedulers for the egress back toward the attached networks) to satisfy the agreements or SLAs the terminal has with each attached network. The sum of these agreements must be accomodated within the terminal to FSPN SLA.

Details of Bounding FSPN PHB Delay

In terrestrial networks, link layer characteristics are generally negligible compared to the Layer 3 effects; the Layer 3 queue scheduling dominates and gets the bulk of focus. But FSPN turns that upside down, introducing delay factors that dominate the characteristics of a hop: long propagation delays, processing delays for interleaving/encoding and variable delays due to the TDMA framing structure. Per Hop Behaviors (PHBs) over packet-based SATCOM links must include link characteristics in their model, figure 6. The Delay Bound (DB) PHB [RFC3248] provides support for VW by providing a specifiable bound on delay variation for DB packets arriving at a node at a rate of no more than R. This section derives MAC layer requirements that bound the delay and delay variation over these SATCOM links order to support the DB PHB.

The delay due to TDMA access delay (d_{mac}) is potentially the largest delay contribution from the lower layers and also a source of delay variation. d_{mac} is the variable time packets must wait upon arrival at the terminal uplink interface before being processed and transmitted in the next frame with a timeslot assignment. Multiple timeslots may be assigned to a terminal in a frame. To provide link security, timeslots are scrambled within the frame so complete transmission of all timeslots assigned to a particular terminal is not assured until the end of the frame. If all packets arriving since the last frame have a timeslot assignment in the next frame, the upper bound on d_{mac} is the time between the start of consecutive frames.

Figure 6. FSPN PHB Model

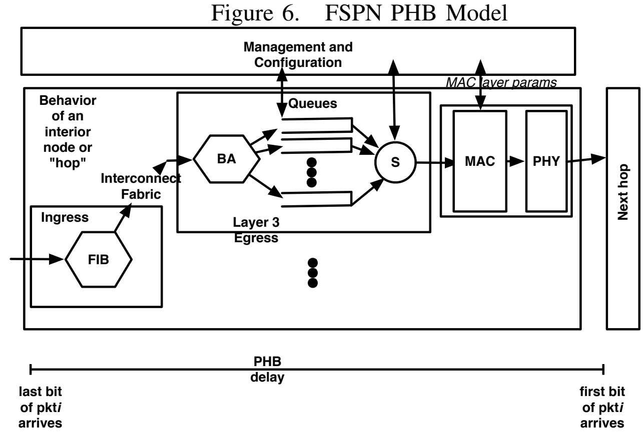

The value of d_{mac} for a packet depends on its arrival time at the terminal interface with respect to the next uplink frame with an assignment for the terminal, leading to delay variation in the traffic. The maximum delay variation is defined as the difference in wait times between the packet with the longest wait and the packet with the shortest wait and is upper bounded for a terminal by the maximum distance between consecutive frames with timeslot assignments in an epoch (J_{mac}).

For example, if a terminal is assigned to transmit in frame f_n and the next frame with a timeslot assignment is frame f_n+4, the distance between them is 4 frames. This results in a variable delay on packets from 0 to 4 frame times depending on when they arrive at the terminal uplink interface. As shown in the Figure 7, the set of packets that will be transmitted up to the payload during frame f_n+4 arrive at the terminal uplink interface during frames f_n-2 and f_n+1. The first packet to arrive at the beginning of frame f_n-2 has to wait up to the full 4 frame times before being processed for transmission in frame f_n+4 while the last packet to arrive in this set arriving at the end of frame f_n+1, has very little wait time (~0) before processing. The range of wait times results in a delay variation of 4 frame times across this link from the lower layers.

Figure 7. TDMA Access Delay Effects

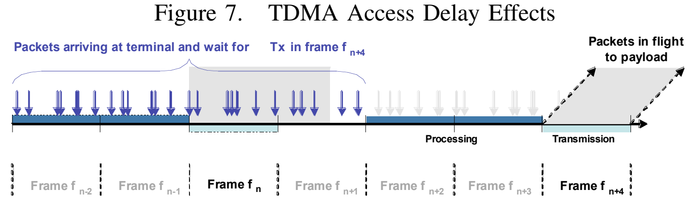

To ensure DB behavior across these SATCOM links in support of the VW PDB, namely low bounded delay (d_{mac}) and delay variation (J_{mac}) as defined above, certain requirements must be met as to the rate assigned to the terminal as well as the distribution of frame assignments within the superframe. The next section derives those MAC layer parameters for the DB PHB. Requirements apply to both the uplink and downlink which have similar lower layer delay contributions but with different magnitudes reflecting the point-to-point nature of the uplink versus the multiplexed, broadcast downlink channel.

The DB PHB must be configured at each hop to provide an available rate sufficient for the maximum amount of VW traffic aggregate that may transit it. To ensure this rate on Bandwidth-on-Demand (BoD) SATCOM links, a constant rate allocation (CRA) should be provided to the link in support of this PHB. Because of the large BoD response times for GEO SATCOM links, the request/grant mechanisms typical of terrestrial BoD networks such as those using 802.16 and DOCSIS which allocate rates dynamically based on actual demand are inappropriate for supporting the DB PHB in this case. Once a terminal is logged on and authorized to receive the VW service specified in its SLA, a MAC_{DB} CRA should be granted for the uplink and downlink based on the total configured VW allocation for the terminal ingress and egress respectively.

Rate allocations to terminals are based on a superframe time period (about a half second) and are defined as the total number of bytes granted divided by the superframe period. Simply allocating a rate equal to a terminal's configured VW rate without regard to where assigned timeslots are within the epoch is not sufficient to ensure the strict DB PHB bounds on delay and jitter. The number of timeslots required and their distribution within the superframe will affect the delay and delay variation experienced across the link. To support DB PHB requirements, the configured rate must be enforced over a timescale equivalent to the desired bound on delay variation. Therefore, sufficient timeslots must be assigned so as to meet the configured rate over timescales considerably smaller than a superframe - on the order of 10s of ms. This means that both the number of timeslots and their distribution within the superframe are important in support of the DB PHB at the MAC layer.

[VWPDB] defines the configured VW rate (R_j) for a VW flow as S_j bytes entering the domain over an interval of T_{vw} (which is constrained by the TSAT jitter). Due to the SATCOM link's TDMA framing, the lower bound on T_{vw} (in which a transfer rate can be ensured during a particular superframe) is equivalent to the maximum distance between consecutive frames with timeslots assigned to a terminal (T_{mac}). T_{vw} is set to its lower bound T_{mac}, to minimize jitter bound. To enforce a bound on jitter (J_{mac}) at the lower layers for the DB PHB that is equal to T_{mac}, the capability must exist to transmit at least S_{mac} bytes every T_{mac} time period such that $S_{mac} = T_{mac} \cdot R_{VW}$ where R_{VW} is the configured rate for ingress VW for a terminal. With frame time T_f, the previous equation requires a minimum of $T_{mac} \cdot R_{VW}$ bytes be allocated to a terminal every $T_{mac} \div T_f$ frames.

Uplink allocations are made in quanta of timeslots that consist of a specific number of user bytes depending on the mode assigned (burst rate, code rate and modulation order). The number of user bytes per timeslot for a particular mode i is $TSsize_{modei}$. To ensure at least S_{mac} bytes may be transmitted by a particular terminal over any T_{mac} time period in an superframe, the number of timeslots (TSs) that must be assigned over any $T_{mac} \div T_f$ frame period is:

$$NumTS_{modei} \geq \lceil S_{mac}/TSize_{modei} \rceil$$

In terms of the configured virtual wire rate for the terminal, applying equation (1) to the previous equation results in:

$$NumTS_{modei} \geq \lceil T_{mac} \cdot R_{vw}/TSize_{modei} \rceil$$

assuming no non-VW packets are allowed to use the link. When other types of traffic share the link, configuration must ensure that not only S_{mac} VW user data bytes can be transmitted every T_{mac}, but also an additional MTU data bytes, where MTU is the maximum allowable packet size for non-VW traffic sharing the link. This accounts for the possibility of an MTU-sized non-VW packet in the transmission stream blocking arriving VW packets from being transmitted in the next available transmission opportunity. Then, to maintain J_{mac}, equation 2 is modified to yield the following:

$$NumTS_{modei} \geq \lceil (T_{mac} \times R_{VW} + MTU)/TSize_{modei} \rceil$$

The following is an example uplink timeslot allocation for a terminal requesting 1 Mbps of virtual wire PDB. The target bound on jitter (J_{mac}) is 20ms due to the lower layers. T_f is 10ms so we set the constraint on TS spacing (T_{mac}) to 2 frames. The mode selected has a transfer rate of 1000 bytes per timeslot. The link may contain packets from other TAs with a maximum packet size of 1500 bytes. Thus, $NumTS_{modei} \geq \lceil (0.02 \cdot 10^6 + 1500 \cdot 8)/(1000 \cdot 8) \rceil =$ 4. There are several ways in which this could be achieved including 2 timeslots every frame or 4 timeslots every other frame. The effective rate allocated to the link is 1.6 Mbps.

At the receiving end of the SATCOM link, the packets are decoded in groups of packets and forwarded to the router closely spaced in time. For the uplink, the variable wait time in the terminal essentially bunches up the packets that arrive during the TDMA wait time into a large burst. The burstiness of a TA is an important characteristic in the ability of a PHB to maintain QoS performance objectives. For the VW PDB, bursts are not permitted on ingress to the FSPN domain boundary in order to maintain the strict QoS bounds. The TDMA framing process that is a part of the DB PHB in FSPN will allow burst conditions to develop in the domain interior on the VW TA. The size of the bursts created in a frame is a function of the traffic arrival rate, burst size of the mode assigned and the number and distribution of timeslot assignments. The burst size created by this PHB is upper-bounded by the max distance between consecutive timeslots times the configured DB rate across the hop. Factors that mitigate the effects of this PHB burstiness on the overall provisioning of the VW PDB are the relative size of the burstiness to the bandwidth of the next hop as well as queue scheduling mechanisms that can smooth out the burstiness (e.g. rate-based scheduling as opposed to priority scheduling).

SUMMARY

UCS employs IP Differentiated Services to create a relatively low delay for FSPN but the specific challenges of the network and its use must be considered. These include mobility of the network's terminals, variability of the links to those terminals, terminals that must have radio silence periods, and terminals that are logged off of the network for periods of months, reappearing at new locations. On-path signaling introduces cost and overhead, requires explicit participation of endpoints, and can be at odds with the IP routing. UCS is ideal for FSPN as resources can be allocated and committed across a continuum from static with planning cycles to immediate responsiveness to network management interactions. Resources can be requested by any authorized agent and resource allocations can respond rapidly to to high-level changes in overall needs and policies. Two major considerations for UCS in FSPN are the need to capture lower layer effects in the characteristics of the Per-hop Behaviors and the need to define appropriate, tailorable, and evolvable methods of allocating UCS that are responsive across a range of time scales.

REFERENCES

[RFC2475] "An Architecture for Differentiated Services", S. Blake, D. Black, M.Carlson, E.Davies, Z.Wang, W.Weiss, www.ietf.org/rfc/rfc2475.txt, December, 1998.

[RFC2638] "A Two-bit Differentiated Services Architecture for the Internet", K. Nichols, V. Jacobson, and L. Zhang, www.ietf.org/rfc/rfc2638,ps, July, 1999.

[RFC3086] "Definition of Differentiated Services Per-domain Behaviors and Rules for their Specification", K.Nichols and B.Carpenter, RFC 3086, www.ietf.org/rfc/rfc3086.txt, April, 2001.

[RFC3248] "A Delay Bound alternative revision of RFC 2598", G. Armitage, B. Carpenter, A. Casati, J. Crowcroft, J. Halpern, B. Kumar, J. Schnizlein, www.ietf.org/rfc/rfc3248.txt, March, 2002.

[VWPDB] K. Nichols, V. Jacobson, K. Poduri, "A Per-Domain Behavior for Circuit Emulation in IP Networks," ACM Computer Communications Review, Vol 34 No. 2, April 2004, pp 71-84.

[CSCO] Enterprise QoS Solutioin Reference Network Design Guide, Version 3.3, Cisco Systems, November 2005.

[MSF] "Bandwidth Management in Next Generation Networks," http://www.msforum.org/techinfo/reports/MSF-TR-ARCH-005-FINAL.pdf

[IPQ] "Inter-provider Quality of Service White Paper," Draft 1.1, Quality of Service Working Group, MIT Communications Futures Program, November 17, 2006, www.cfp.mit.edu

[XIAO] A Practical Approach for Providing QoS in the Internet Backbone, Aug. 2002 XiPeng Xiao, Thomas Telkamp, Victoria Fineberg, Cheng Chen, Lionel M. Ni Published in IEEE Communications Magazine, Dec. 2002

[TLKMP] T. Telkamp, "Traffic Characteristics and Network Planning," NANOG-26, Eugene, OR, October 27-29, 2002, http://nanog.org/mtg-0210/telkamp.html

[HOSE] N.G. Duffield et al, "A Flexible Model for Resource Management in Virtual Private Networks," Sigcomm, 1999.