INTERNATIONAL JOURNAL OF SATELLITE COMMUNICATIONS AND NETWORKING
Int. J. Satell. Commun. Network. 2004; **22**:1–17 (DOI: 10.1002/sat.785)

Transport protocols in multicast via satellite

Gun Akkor‡, Michael Hadjitheodosiou*,† and John S. Baras§

Center for Satellite and Hybrid Communication Networks, University of Maryland, College Park, MD 20742, U.S.A.

SUMMARY

In a wide variety of broadband applications, there is a need to distribute information to a potentially large number of receiver sites that are widely dispersed from each other. Communication satellites are a natural technology option and are extremely well suited for carrying such services because of the inherent broadcast capability of the satellite channel. Despite the potential of satellite multicast, there exists little support for multicast services over satellite networks. Although several multicast protocols have been proposed for use over the Internet, they are not optimized for satellite networks. One of the key multicast components that is affected when satellite networks are involved in the communication is the transport layer. In this paper, we attempt to provide an overview of the design space and the ways in which the network deployment and application requirements affect the solution space for transport layer schemes in a satellite environment. We also highlight some of the issues that are critical in the development of next generation satellite multicast services. Copyright © 2004 John Wiley & Sons, Ltd.

KEY WORDS: multicast transport protocols; satellite networks; taxonomy

1. INTRODUCTION

In a wide variety of broadband applications, such as software updates, distributed computing, and multimedia content distribution, there is a need to distribute information to a potentially large number of receiver sites that are widely dispersed from each other. Communication satellites are a natural option for this because of the inherent broadcast capability of the satellite channel. Also, a satellite-based infrastructure can, in many cases, be established to offer widespread service provision with greater ease and simplicity than an infrastructure based on terrestrial broadband links, when the latter is not available. Hence, while much of the broadband communication today is carried via terrestrial links, satellites will come to play a greater and more important role, especially for point-to-multipoint services.

*Correspondence to: Michael Hadjitheodosiou, Center for Satellite and Hybrid Communication Networks, A W Williams Building, University of Maryland, College Park, MD 20742, U.S.A.
† E-mail: michalis@isr.umd.edu
‡ E-mail: akkor@isr.umd.edu
§ E-mail: baras@isr.umd.edu

Next-generation satellite communication systems utilizing higher frequency bands such as the Ka-band, spot-beam technology and on-board processing are currently under development (e.g. Hughes SPACEWAY [1] by Hughes Network Systems). Ka-band is very desirable for satellite communication systems, because it offers abundant bandwidth. The use of spot-beam and on-board processing technologies enable the use of small, low-power, low-cost user terminals that offer two-way direct communication. These systems are likely to play an important role in the global communication infrastructure [2].

Despite the potential of satellite multicast, there exists little support for multicast services over satellite networks. Although several multicast protocols have been proposed for use over the Internet, they are not optimized for satellite networks. Therefore, efficient integration of next-generation satellite systems and multicast services requires the study and adaptation of these protocols. One of the key multicast components that is affected when satellite networks are involved in the communication is the transport layer. Both Internet Engineering and Internet Research Task Forces (IETF and IRTF) have been involved in a research effort to identify the design space for a general purpose multicast transport protocol and standardize certain protocol components as *building blocks* [3, 4]. Their efforts show, however, that it is arguably impossible to achieve a *one-size-fits-all* design recipe, due to the diverse range of multicast applications and the variety of multicast network topologies.

In this paper, we focus on transport protocol issues, and we attempt to provide an overview of the design space and the ways in which the network deployment and application requirements affect the solution space for transport layer schemes in a satellite environment. An earlier survey of protocols, functions and mechanisms for multi-point communication can be found in Reference [5], while Reference [6] gives a taxonomy of multicast transport protocols and classifies them according to important common features. This work has been taken forward in the context of satellite networks in Reference [7].

We maintain a similar taxonomy to that established in References [6, 7] because it is also consistent with the IETF standardization efforts [3, 4, 8, 9]. However, we first discuss the characteristics of satellite networks and identify which key blocks of a general multicast transport protocol are affected by certain network deployment scenarios. We also pay special attention to issues related to reliable dissemination of data since it imposes several additional requirements on the design of protocols. In Section 3, we present a survey of current alternatives on the design of multicast transport protocols and discuss their applicability to our network scenarios. In Section 4, we give recommendations on the desirable futures of protocols for satellite multicast. Section 5 concludes the paper.

2. CHARACTERISTICS OF SATELLITE NETWORKS

Satellites systems are becoming an integral component of broadband networks. They have several characteristics that are particularly attractive in this case, such as breadth of broadcast 'reach', ubiquitous access, low-cost global coverage, large and most importantly flexible capacity. At the same time, characteristics such as the relatively long propagation delay and the complications of a wireless channel, constitute serious shortcomings. Finally, a few other characteristics such as on-board switching, regeneration and spot-beam technology offer opportunities and represent challenges that can transform some of the shortcomings into strengths.

In this section, we provide a description of the technology challenges that must be overcame to permit the successful design of transport layer schemes for satellite networks.

2.1. Constraints imposed by the channel

Spectrum congestion of frequency bands (L, S, C, or Ku bands) that had been conventionally used for satellite services is leading to the use of higher frequency bands such as the Ka-band (20–30 GHz). Ka-band is very desirable for multimedia communications because it offers wider bandwidth. Such large bandwidth segments are unavailable at lower frequencies, such as Ku-band (12–18 GHz) and C-band (4–6 GHz), which were until recently the bands used for fixed satellite service (FSS) communications. Most very small aperture terminal (VSAT) and direct broadcast satellite television (DBS TV) systems in operation today use portions of the Ku-band, while most Ka-band systems are scheduled to start offering customer service in the near future [2, 10].

However, Ka-band has one major disadvantage. Rain and atmospheric attenuation present a significant challenge to transmission of signals at Ka-band frequencies [10, 11], as the molecular water vapour absorption resonance frequency is located at the center of the band, at 22.3 GHz. The resulting signal fading causes not only just random bit errors but also longer bursty errors that can cause the transmission operation to be completely 'off' for short periods of time, even though the channel would be of extremely high quality for most of the operation. Moreover, such extreme attenuation usually occurs in cells of small dimension (2–8 km in diameter) [12] compared to the footprint of a narrow spot-beam (650 km in diameter at a geosynchronous orbit). Therefore, at any time instant, user stations experience very different channel conditions, leading to a *heterogeneity* in the reception quality. Continuing demand for additional bandwidth has forced commercial satellite system designers to consider even higher frequency bands, namely the V-band (40–75 GHz). Some military satellite systems already operate in this frequency range. These higher frequencies offer additional challenges to the designer such as more severe multi-path fading and scattering of transmitted signals. Clearly, satellite channels suffer from higher error rates and bursty error patterns than terrestrial wireline networks, and transport protocols have to have adequate mechanisms to maintain fiber-like quality-of-service (QoS).

Another setback for satellite networks is the long propagation delays associated with satellite channels. The physical distance of a communications satellite from the source and destination imposes a significant propagation delay on every transmission. Besides causing problems for real-time delay-sensitive applications, this delay can adversely affect the performance of certain transport protocols, such as transmission control protocol (TCP). In geostationary earth orbit (GEO) and medium earth orbit (MEO) systems, the propagation delay (260 and 100 ms, respectively) is much higher than in low earth orbit (LEO) systems (10 ms), but in LEO constellations the need to route a signal through multiple satellites imposes delay, too, and might also increase the variance of the delay.

Unreliability of the satellite channel, *long propagation delays* and the *heterogeneity* in the channel quality impose a significant challenge in the design of several reliable multicast transport protocol blocks. These challenges are discussed in detail in Section 3.

2.2. Constraints imposed by the network topology

We consider two of the most common topologies for multicast service support involving broadband satellites:

(i) Satellite networks can be deployed as a *backbone* for interconnection of geographically distributed high-speed local area networks (LAN). In this scenario, LAN are connected to the satellite backbone through one or more gateway nodes that have satellite uplink capability (Figure 1(a)). This network topology gives rise to a hierarchical structure: satellite and gateway nodes acting as an overlay network for the LAN(s). In general, LAN(s) may also have access to a terrestrial core network. This network architecture provides unique opportunities for multicast distribution of data. In this scenario, the forward (downlink) feed may be used to efficiently distribute content to many sites, while the terrestrial links are used for transmission of out-of-band control information, user feedback, and data retransmissions.

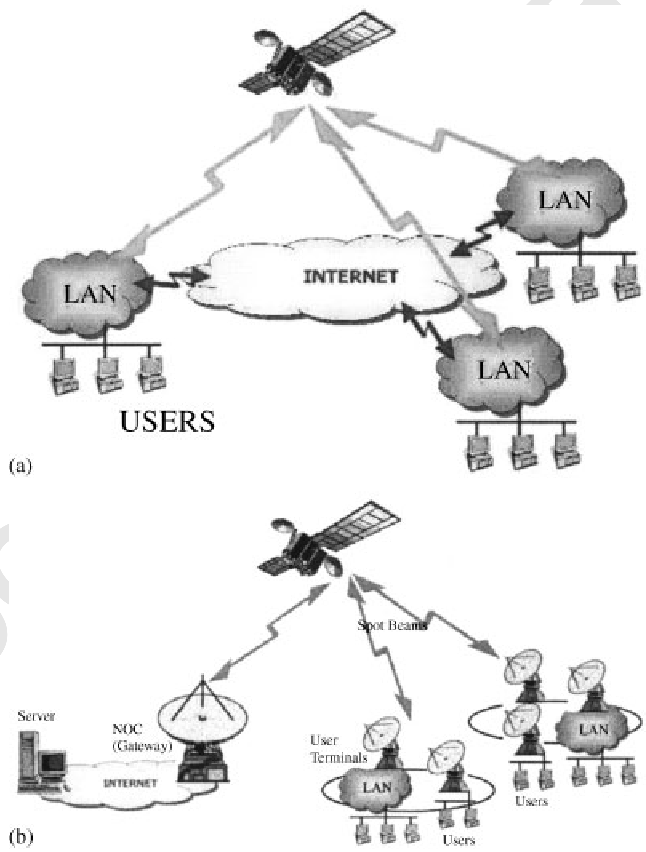

Figure 1. Satellite network topologies. (a) Case I: backbone deployment. (b) Case II: direct-to-home deployment.

(ii) A *direct-to-home* (or *direct-to-business*) deployment, where the network consists of independent ground terminals with direct uni- or bi-directional connection to the satellite. In this scenario, network has a star topology and user terminals have no direct access to other networks. Ground terminals access the terrestrial core network through a gateway node located at the so-called *network operations center* (NOC) (Figure 1(b)). This architecture imposes additional challenges on the transport protocol, especially if reliable delivery of the data is required, because a satellite return channel is required and the transmission of user feedback, out-of-band control information, and data retransmission have to go through the satellite channel.

Network topology is one of the most important constraints, because it strictly limits the mechanisms that can be used for reliability, flow and congestion control. We discuss the limitations and strengths of current design alternatives for a given network topology in the next section.

3. TAXONOMY OF MULTICAST TRANSPORT PROTOCOL BUILDING BLOCKS

A multicast transport protocol operates through inter-working of several different components. The topology of the deployed network and the requirements of the target application constrains the design space of these components. In this section, we provide a taxonomy of the design features that are desirable for a multicast transport protocol. Some of these features are required for providing reliability.

3.1. Loss detection and feedback

To provide reliability, a protocol needs to identify the packets that failed to reach a given destination. This is achieved through loss-notification (feedback) packets returned to the designated source(s)[¶] (e.g. sender) by the receivers. Traditionally, this feedback has been in one of the following forms:

- *Positive acknowledgments* (ACK): Receivers return ACK packets to the designated source, indicating which packets have been received.
- *Negative acknowledgments* (NACK): Receivers return NACK packets to the designated source, indicating only packets that are missing by a receiver.
- A hybrid approach (i.e. both ACK(s) and NACK(s)).

For multicast transport protocols, ACK-based loss-notification has been shown to lead to the *ACK implosion* problem. The problem arises when a large number of multicast receivers return an ACK packet for every packet they receive correctly, causing a serious network congestion around the links of the source. Another potential problem is that the source is required to keep track of the size, and the current state of the reporting receiver-set, in order to identify the last data packet correctly received by all receivers. In a large scale multicast application, the memory and processing load becomes prohibitive.

[¶]We adopt the term *designated source* in place of the sender because, unlike the unicast communication, some multicast protocols use intermediate nodes to collect and aggregate feedback information.

NACK-based feedback alleviates some of the problems. The performance comparison study presented in Reference [13] confirms that NACK-based multicast transport protocols deliver better performance than their ACK-based counterparts in wireline terrestrial networks. In NACK-based feedback, receivers detect missing packets by checking for gaps in the packet sequence numbers and report to the designated source. The source neither needs to know the size of the receiver-set at any point in time, nor keeps track of the current state of every receiver in its group. Also, the number of NACK packets is expected to be less than that of ACK packets at low error rates. A disadvantage is that it is more difficult for the source to know when it can free the transmission buffers, since NACK packets may be lost in transmission.

For a satellite multicast application, the potential problems related to loss detection and user feedback are more pronounced. Due to high error rates over the satellite links and the potentially large number of receivers, even NACK-based feedback may lead to an implosion problem. Moreover, long propagation delays make retransmission of data in response to user feedback inefficient, if not impossible. In a *backbone* deployment, existence of terrestrial links allows the use of supporting mechanisms such as feedback aggregation [14–18] and feedback suppression [17–20], which have been primarily developed with wireline terrestrial networks in mind. In a *direct-to-home* deployment, however, the applicability of these algorithms are limited as we will further elaborate in Section 3.3.

3.2. Loss reduction

A packet recovery mechanism is an essential component of a reliable transport protocol. Automatic repeat request (ARQ) is a well-known technique for this purpose. In ARQ, sources respond to loss notification reports by retransmitting the missing packets. In a satellite multicast application, pure-ARQ for packet recovery turns out to be inefficient for several reasons. Long propagation delays associated with satellite links make the delay incurred during this repeat-request process unacceptable for many delay sensitive applications (e.g. video streaming). The frequent and deep fades observed in satellite links result in high error rates and bursty error patterns. Therefore, even if the feedback mechanism is efficient in returning the loss information back to the sources, in a typical wide area satellite deployment, considerable network bandwidth and processing time would be spent by retransmitting the different packets lost at different receivers. Several early papers exist on use of ARQ in satellite multicast applications [21–24].

Forward error correction (FEC) coding is a well-known technique for protecting data against corruption [25–27]. FEC coding involves addition of redundant (parity) data to the original stream either at the physical bit level or at the packet level. In some FEC schemes, redundant data are added to the physical bit stream. This may improve the satellite link's effective bit error rate (BER), reducing the number of errors without an increase in the satellite transmitted power (EIRP). However, in the case of IP multicast protocols, the network layers will detect corrupted packets and discard them or the transport layers can use packet authentication to discard corrupted packets. Therefore, an alternative application of FEC codes to multicast protocols is as an erasure code [8, 9, 27,28] at the packet level. Erasure codes allow generation of n encoding packets from k original data packets. In such cases, the original k data packets can be reconstructed as long as k out of the transmitted n encoding packets are received correctly. Use of FEC erasure coding offers some solution to a number of problems.

One immediate benefit of packet level FEC coding is the reduction in the number of lost packets. This minimizes the number of feedback reports as well as the need for retransmissions, and improves scalability. FEC coding has an additional benefit for satellite multicast services. Unlike the unicast communication, where the retransmission exclusively benefits the receiver that made the request, in multicast, the protocol has the option to either unicast or multicast the retransmitted packet. Multicast of retransmitted packets would benefit multiple receivers in case of correlated loss. However, it may as well degrade channel utilization in case of uncorrelated loss. Because of the broadcast nature of the satellite channel, all packets are received by all members of the receiver-set. Transmitting parity packets in place of corrupted data packets significantly improves the channel utilization, since a single parity packet can repair loss of different packets at different receivers of the same session. Also, the protocol becomes more scalable in the number of receivers. The protocol does not need to know which packets have been lost by the receivers of the session but only the maximum number of packets lost by any receiver of the session. Therefore, the feedback from a single receiver has been reduced [28–31].

Packet level FEC coding is a valuable component in the context of reliable multicast delivery, because FEC encoding symbols can be useful to all receivers in the recovery process even when different encoding symbols are received at different receivers. Therefore, there is an agreement among the research community on the use of packet level FEC coding, especially over error-prone and high latency channels, where retransmission of packets in response to individual receiver requests is inefficient. Many of the recent reliable multicast transport protocols integrate some form of packet level FEC coding [32–34] in their design. We will elaborate more on the integration of packet level FEC coding in Section 4.3.

3.3. Feedback suppression and aggregation

Several existing multicast protocols adopt feedback suppression and feedback aggregation schemes to control the number and flow of feedback packets to avoid feedback implosion problem. Applicability of these methods in a satellite multicast application depends heavily on the architecture of the deployed network.

Feedback aggregation is applicable in networks where a logical or physical hierarchy is possible. Feedback aggregation relies on intermediate entities to collect, filter-out and combine feedback information towards the source. In a *backbone* scenario, it is possible for intermediate LAN routers to aggregate feedback packets from receiver nodes and to report only the aggregated feedback information to the satellite gateways. Gateway nodes can further aggregate reports from several routers to pass only a single feedback information back to the source. This aggregation towards the root of the hierarchy cuts back on the number of feedback reports and prevents receivers from contacting the source directly, enabling the protocol to scale over large set of receivers. In this scenario, it also possible to use the terrestrial network for transmission and aggregation of feedback information [14–18]. In a *direct-to-home* deployment, feedback aggregation would not be possible because the star topology does not lend itself into a hierarchy. Even if the number of interested users inside the local domain is small or feedback aggregation techniques are used in the local domain, feedback implosion remains to be a challenge due to large number of terminals.

Feedback aggregation schemes require additional functionality at the intermediate routers of the network [35]. This raises questions about whether the router extensions can be standardized and automatically configured. Instead of controlling the volume of feedback information by aggregation, feedback suppression algorithms control the number of feedback packets forwarded by every user. In terrestrial networks, if feedback reports were multicast to the network, then receivers could listen to the feedback reports of others to suppress their own feedback reports if the received report is for the same missing packet. Coupled with random timers and statistical back-off, this type of feedback suppression would, in an ideal setting, let only one feedback report to be returned to the sender for every missing packet [17–20]. However, multicast forwarding of feedback reports must be scoped by use of *time-to-live* (TTL) fields to minimize the flooding of the network by feedback packets. When terrestrial links are used for transmission of feedback information, this form of feedback suppression is applicable in satellite networks. It is not, however, applicable for *direct-to-home* deployment since we assume there is no direct connection among receivers. Even though the satellite can relay the feedback packet by broadcasting it to all the receivers, the delay incurred during the co-ordination may be unacceptable in many cases. However, other possible implementations of feedback suppression is possible if, for example, a representative (or a set of representatives) among the users could be selected to transmit its feedback for the group. These possibilities are further discussed in Section 4.1.

3.4. Packet recovery

In unicast communication, feedback reports are returned to sender and retransmissions are initiated by the sender. In order to avoid feedback implosion and traffic concentration at the sender and to reduce the packet recovery latency, several multicast protocols adopt *local recovery* methods [14–16, 18, 20]. Local recovery allows designated nodes to buffer data packets, and initiate retransmissions on behalf of the sender. Therefore, receivers first try to recover missing packets through these nodes. If the missing packet can not be recovered through the use of designated nodes, then the retransmission request is forwarded to the sender. In networks where a logical or physical hierarchy is possible, intermediate routers can buffer the packets forwarded through them and initiate retransmission up on receiving a request. Coupled with feedback aggregation (Section 3.3) *router-assisted* local recovery is shown to improve the scalability and the performance of reliable multicast protocols [31, 36–39].

Local recovery of packets are particularly important in satellite networks because of the long propagation delays associated with satellite links. In a *backbone* deployment, satellite gateways are in a very good position to assist in the local recovery together with other intermediate routers. However, as it is the case with feedback aggregation, *router-assisted* local recovery requires support from intermediate network elements and its availability depends on existence of router extensions [35, 40]. A generalization of *router-assisted* local recovery is possible, if receiver nodes are also allowed to respond to retransmission requests. In this case, feedback packets need to be multicast to the (sub)set of receivers. However, efficient scoping of feedback and repair packets must be implemented to avoid flooding of network. In a *direct-to-home* deployment, all receivers receive packets directly from the satellite and there are no intermediate routers between satellite and the receivers. Therefore, it is difficult to implement local recovery for this type of networks.

3.5. Congestion and flow control

The Internet congestion control has been achieved by the widespread use of TCP protocol. Multicast protocols need to have suitable techniques to avoid congestion and to determine a 'fair' share of the available resources with respect to TCP network traffic. Therefore, most of the recent studies and proposals for congestion control of multicast traffic in the Internet try to be

TCP-friendly by adopting the behaviour of TCP algorithm as the standard of fairness. In the context of satellite multicasting, there are two problems: (i) adaptation of TCP behaviour to the multicast communication, and (ii) adaptation of TCP behaviour to satellite networks. We will first discuss the issues related to having a TCP-like behaviour in the context of multicast communication, and then elaborate on problems related to satellite context.

TCP protocol uses a window-based congestion control algorithm, where the size of the congestion window determines the number of outstanding (i.e. not yet acknowledged) packets in the network. The window size is regulated in response to ACK packets returned to the sender. Lack of an ACK packet acts as a loss (and/or delay) indication causing the window size to be decreased, otherwise it is increased. However, it is difficult to extend window-based regulation to multicast communication. Applying a single window size to the whole session has been shown to restrict the throughput more severely than required by the bottleneck path [41] because, the throughput is dictated by the receiver with the longest round-trip delay (RTT) rather than the bottleneck receiver (i.e. the receiver with the smallest throughput). Use of per packet ACK(s) for regulation of the session rate is another problem for multicast applications, because it causes, as in the case of receiver feedback for reliability, feedback implosion problem at the sender. Another fundamental problem in multicast congestion control is the *loss path multiplicity problem* [42]. Unlike unicast congestion control, where the sender regulates its transmission rate based on the loss indications from a single receiver, a multicast source receives loss indications from multiple receivers, reflecting diverse conditions in various parts of the network. For example, if a source receives multiple loss indications from a set of receivers that are behind the same bottleneck link, and reduces its rate in response to each such loss indication, then it would be over-compensating for a single loss. Therefore, loss indications have to be appropriately combined when making a single rate control decision [43, 44]. The multicast source must regulate its rate according to the most congested path, or equivalently, according to the lossiest receiver in the multicast group.

Some of the previous proposals on the issues of error recovery and feedback implosion (Sections 3.1–3.3) can be incorporated into congestion control schemes. Multicast transmission control protocol (MTCP) [45] is based on a multi-level logical tree, which is used to aggregate ACK-based feedback packets, where the root is the sender and other nodes in the tree are receivers. Internal nodes, referred to as *sender's agents* (SA) monitor the congestion level of their children and compute the minimum bandwidth available from the sender to their children by maintaining a TCP-like congestion window. When sending an ACK to its parent, a SA includes in the ACK packet, a summary of the congestion level of its children. The sender regulates its rate based on its own summary. The pragmatic general multicast congestion control (PGMCC) [46, 47] uses a window-based control that mimics the TCP behaviour, but runs it between the sender and the bottleneck receiver avoiding some of the problems addressed in Reference [41]. The bottleneck receiver, referred to as the *acker*, is chosen dynamically based on receivers' loss rate reports embedded into NACK-based feedback reports. ACK-based feedback is only used between the *acker* and the sender for congestion control.

Difficulties in implementing window-based congestion control have lead to rate- and equation-based control protocols [48–50]. These protocols calculate the rate at the sender based on the long-term throughput equation for TCP protocol. The equation gives the expected rate of a TCP flow as a function of the steady-state loss event rate, round-trip time, and the packet size. Receivers measure their loss rate and include this information in the feedback packets returned to the sender. The sender uses the feedback messages to measure and estimate the round-trip time to the receivers and uses the rate equation to derive the acceptable transmission rate. However, estimation of round-trip time without causing excessive traffic at the sender presents a formidable task, and is the key component of equation-based control protocols.

Use of single rate congestion control algorithms for multicast services necessarily forces the sender to adjust its rate according to the bottleneck receiver. Heterogeneity of the receiver-set penalizes the receivers with higher throughputs. To avoid this problem, several authors have suggested layered organization of the transmitted data [51, 52]. In layered multicast transmission, sender distributes the data using layers with increasing bandwidth. Receivers add or drop layers based on the perceived bandwidth of the path to the sender. This allows each receiver to match its desired rate. In order to use layered transmission as a congestion control algorithm, however, receivers behind the same router have to act in a co-ordinated manner, because the action of a receiver dropping a layer (due to perceived congestion) would not be effective unless all the receivers sharing the bottleneck drop the same layer. Moreover, if a receiver causes congestion on the bottleneck link by adding a new layer, another might interpret the resulting losses as a consequence of its current level and drop a layer. Therefore, fairness of the protocol requires that all of them have a similar behaviour [52]. As another point, a layered organization is only possible if the data to be transferred supports it. Layered organization is more effective for video and audio-conferencing applications, where data can be organized into layers of increasing quality [53]. However, for bulk data transfer, this is not possible. By using a proper arrangement, a similar approach could be used [54–56], if the data can be processed beforehand. For data generated on the fly, it becomes more complex to find a suitable data organization without wasting network capacity. Layered transmission is less suitable for satellite multicast, since all receivers see the same effective link rate.

All current proposals try to achieve TCP compatibility because, TCP protocol is the *de facto* standard for end-to-end congestion control in the Internet. However, deployment of satellite networks, either as an access network to Internet or a core network for global services, poses new questions on the performance of TCP-like protocols over hybrid or pure satellite networks. TCP protocol, initially designed for terrestrial networks with low link error rates, assumes all loss indications are due congestion in the network (i.e. overflow of router buffers). Therefore, sender decreases its rate each time a packet loss is detected. For satellite networks, this causes unnecessary performance degradation when the losses are due to link errors. To overcome this problem, several solutions have been proposed in the literature. One solution is to differentiate between losses due to congestion and losses due to link errors. In *link corruption notification* algorithm [57], which is a part of space communication protocol standard transport protocol (SCPS-TP) favoured by NASA and U.S. DoD for space-to-ground downloads, receivers continuously monitor their channel state and enter the link-corrupted state if the channel condition is bad. The sender behaviour depends on the channel condition:

- If the channel condition is good, the cause of packet losses is assumed to be due to network congestion as in traditional TCP.
- If the channel condition is bad, the cause of packet losses is assumed to be due to link errors, and the rate is controlled using a open-loop token bucket.

In the proposed *rapid recovery* [58] algorithm, at first, all packet losses are assumed to be due to network congestion as it is the case in traditional TCP. However, in order to probe the availability of network resources, the source transmits certain number of dummy packets. If the packet loss was due to link errors, the dummy packets would reach the receiver and be acknowledged, increasing the transmission window rapidly. However, these ideas are not directly applicable to multicast communication, since receivers at different locations would experience different channel and network conditions and therefore, in the decision making process, the sender is forced to combine multiple indications in making a single rate control decision. More details on these and similar approaches could be found in Reference [59].

The rather long propagation delays usually associated with satellite links could impose performance constraints for TCP-like protocols, since the congestion control window is regulated per round-trip time. Especially, TCP's slow start mechanism has been shown to cause performance problems in satellite networks, and several solutions have been proposed primarily for unicast communication. *TCP-connection-splitting* protocols split a TCP connection between the sender and the receiver into three separate connections—one between the sender and the satellite gateway, the other between the satellite gateway and the user satellite terminal, and a third one between the user satellite terminal and the user machine. In this case, a specialized protocol tuned to the satellite environment may be used over the satellite-only hop, while traditional TCP runs on other connections. In this approach, over the satellite-only hop, packet losses are attributed to channel conditions, and there exists a bandwidth allocation problem rather than a congestion problem, since the satellite uplink/downlink bandwidth is fixed and has to be shared amount all active connections. Over the TCP connections, on the other hand, packet losses are due to congestion only, and are dealt with accordingly. These and similar issues are addressed by the research community in several occasions and more details can be found in References [58–61]. We discuss some of the specialized protocols for the satellite environment in Section 4.

4. CHALLENGES AND RELATED RESEARCH IN SATELLITE MULTICAST TRANSPORT PROTOCOLS

In this section, we outline some of the current directions in the development of next generation satellite multicast transport protocols. Some of the issues outlined in this section have already been identified as the key problems for multicast services by the research community. However, in the case of satellite multicast services, these problems have either different roots or different solution spaces. Therefore, in this section, we revisit these questions and try to provide design guidelines.

4.1. Avoiding feedback implosion

Current design alternatives for feedback implosion avoidance primarily depend on whether an alternative terrestrial path is available for transmission of user feedback. When terrestrial links are available, user feedback could be sent to the source over the terrestrial links. Feedback implosion could be avoided on these paths by using one of the existing feedback scoping alternatives discussed in the previous sections. However, in areas where no such path is available, it is not possible to send user feedback over the terrestrial links, and for reliable communication, we have to consider a satellite system with bi-directional links. In this case, uplink channel resources have to accessed and shared by all receivers. In Reference [62], the proposed algorithm uses a round-robin time division multiple access (TDMA) scheme on the uplink channel. The disadvantage is that, it does not scale well with increasing number of receivers, since more uplink bandwidth has to be allocated for feedback transmission. In Reference [63], the authors group receivers into clusters, where only the cluster heads transmit feedback over the satellite link. Although this scoping reduces the number of users that have to access the uplink channel, the approach assumes that receivers inside a cluster are able to communicate each other through secondary means.

We are concerned that both the requirement for terrestrial links, as well as the requirement for secondary communication between the users are restrictive and in general contradicts the reasons for deployment of satellite networks. The feedback implosion and the uplink bandwidth sharing problem should be studied without these assumptions. Some of the more recent work on this topic try to approach the problem in this direction. To our knowledge, the approach described in Reference [64] is one of the first proposals to solve the feedback implosion problem over satellite networks without any secondary communication between the users. The disadvantage of this approach is that the feedback resolution requires multiple round trip times over the high latency satellite channel. In Reference [65], the authors propose forming receiver clusters based on their spatial correlations, e.g. rain fall precipitation data, which is a prime factor that degrades service quality. Receivers in the same cluster use different return channels, while receivers in different clusters may use the same channel since there is a low probability of receivers located away from each other sending feedback simultaneously. In Reference [66], a proposed feedback implosion suppression algorithm allocates receivers into return channels based on a knapsack algorithm that tries to minimize the number of total return channels allocated for feedback transmission.

We believe that feedback implosion avoidance is one of the key problems that has to be addressed in the context of satellite networks, even though several proposals exist for use over terrestrial wireline networks.

4.2. Packet recovery

Packet recovery is another problem for which there are several available solutions in the context of terrestrial wireline networks, but has to be revisited in the context of satellite networks. The design alternatives for packet recovery also depend heavily on the existence of a terrestrial connection among the receivers. In Reference [63], authors use local recovery of missing packets among the members of a receiver cluster. Only when the packet can not be recovered within a cluster, its is retransmitted over the satellite. This approach reduces the recovery latency compared to schemes which have to go through a repeat-request process over the high latency satellite channel.

Missing packets could be retransmitted over the satellite channel, if the receivers are experiencing correlated losses. This approach would improve the transmission efficiency because of the broadcast nature of the satellite channel. Retransmissions could also be unicast to individual receivers which are requesting the packets over a terrestrial link, if losses are uncorrelated. The satellite reliable distribution protocol (SRDP) [34], combines both approaches and uses three separate phases to complete the transmission of the packets. During the first phase, packets are transmitted over the satellite channel. After collecting feedback from the receivers, in the second phase, the packets needed by a large number of receivers are retransmitted over the satellite channel. Only after the second phase, receivers are allowed to individually receive and request packets from the source.

We believe that the retransmissions over a terrestrial or a secondary network are restrictive and that all transmissions could be handled over the satellite channel. However, retransmission of individual packets in response to individual requests over the satellite channel clearly does not make good use of the broadcast channel and causes high recovery latency. Therefore, we think that, integration of packet-level FEC schemes for preventing losses as well as for assisting in the recovery process, is the key for successful multicast protocols. We discuss this issue in the following section.

4.3. Integrated FEC

Packet level FEC is very effective at reducing the repair traffic per packet loss. However, use of traditional block FEC codes such as Reed-Solomon codes that use the algebraic properties of finite fields, have limitations in terms of the computationally feasible block and symbol sizes. A small block size means that the sender may run out of unique encoding packets before all receivers could accumulate enough number of packets for successful decoding. Such receivers would have to request original packets individually from the sender, adding to recovery latency.

These limitations could be avoided by codes capable of working over large block sizes, such as Tornado codes [67]. Tornado codes require slightly more than k out of the n encoding packets to recover the k source packets, i.e. there is a small reception overhead. Another type of FEC codes, which we refer to as expandable FEC codes, can generate as few or as many unique encoding packets as required on demand for the same group of k source packets. Luby Transform (LT) codes [68, 69] are an example of large expandable FEC codes with a small reception overhead. We believe that such codes are key to the design of successful multicast protocols. They make very good use of the broadcast nature of satellite channel and remove the restriction of a fixed block length, since additional encoding packets may be transmitted as long as there are receivers that have not yet accumulated enough encoding packets to recover the k source packets.

It is important to note that FEC techniques trade bandwidth efficiency to improve the probability of successful transmissions. Satellite channels, as many other wireless communication systems, are time-variant causing the transmission operation to be completely 'off' for short periods of time, even though the channel would be of extremely high quality for the rest of the time. Hence, fixed protection against the worst case scenario is likely to be inefficient. Therefore, there has been a lot of effort to come up with adaptive schemes that dynamically adjust the amount of parity generated [70–74]. We think that designs that take into account channel estimations and receiver feedback in determining the number of encoding packets will be the next research direction. Future systems could couple these schemes with on-board buffering and processing to improve the efficiency of the protocols. These possibilities are discussed in the next section.

4.4. On-board buffering and processing

Implementing on-board processing and buffering has several merits that would improve the performance of satellite services, including but not limited to multicast. For satellites operating around 8 GHz and above, on-board buffering of messages and automatic repeat request (ARQ) can be used to minimize the effects of rain outages. This is particularly advantageous at higher frequencies and at lower elevation angles, both of which significantly increase the excess path loss due to rain. The benefits are more pronounced in case of multicast services. By performing

ARQ on-board the satellite, it is possible to minimize the number of feedback packets that needs to be propagated back to the source. This would improve both the bandwidth efficiency of the multicast service and the packet recovery latency.

On-board processing allows regeneration, duplication and encoding of packets. Implementing on-board processing would benefit multicast services by allowing adaptive FEC coding and parity (re)transmission. It also allows isolating the uplink and downlink which would be beneficial in a multicast scenario since, unicast uplink communication between the source and satellite is likely to have different parameters than the satellite downlink. For example, at the downlink, more redundancy may be needed to accommodate the requirements of the receiver with the worst channel quality than the uplink channel. Isolation of the uplink and downlink, in this case, would allow a higher code rate at the uplink, and a lower rate code at the downlink. Coupled with on-board buffering, it makes it possible to dynamically change the transmission parameters of the downlink as the multicast membership changes, while keeping the uplink communication the same.

4.5. Flow control and congestion

There is an effort in the research community to provide a ubiquitous end-to-end congestion control algorithm for hybrid satellite-terrestrial networks [58, 75]. This requires careful modification of the TCP protocol to match the characteristics of the satellite channels. Another approach is to split a connection at the terrestrial-satellite network interface, and to run TCP between the end nodes and the satellite network gateways, while running a customized algorithm in the satellite-only core network. The latter approach has several advantages in practice. It allows satellite provider to optimize the traffic flow inside the core network to achieve high utilization. Also, it makes it possible to push the congestion to the edges of the satellite network and let the TCP protocol's congestion control mechanism to take care of the congestion in the terrestrial portion of the network. Inside the core network, the problem is reduced to a flow control problem rather than a congestion control problem since the satellite bandwidth is fixed and is to be shared among all the active connections. Hence, the satellite-only tier is isolated from the rest of the Internet, and flow control can be implemented by a simple window control mechanism. There is also no need to differentiate between losses due to congestion and losses due to link losses.

Isolating the satellite network from the rest of the Internet would require additional functionality at the gateway nodes [59]. However, we believe that satellite providers would be willing to go with this design, since the implementation will allow them to have full control of the traffic flow in their networks. Adding to this the challenges of having a TCP-like behaviour for multicast services, we believe that, having a customized algorithm for the satellite network has considerable advantages.

5. CONCLUSION

In this paper, we have presented a taxonomy and survey of various design alternatives for supporting multicast services, and discussed how the design space of various multicast transport protocol components are constrained in the context of satellite networks. Our classification is based on the IETF *building blocks* for multicast transport protocols, but highlights which key components of a general transport protocol are affected by the two most common satellite network deployment scenarios.

We also outlined some of the issues that are critical in the development of next generation satellite multicast services. Some of these problems, such as feedback implosion avoidance and packet level FEC coding at the transport layer, have counterparts in terrestrial networks, but they have to be addressed separately while taking into consideration the unique characteristics of satellite networks. We believe that efficient solutions to these problems and development of new technologies, such as on-board processing and buffering, would demonstrate the true value of satellite networks in the global communication infrastructure.

REFERENCES

1. Fitzpatrick EJ. SPACEWAY system summary. *Space Communications* 1995; **13**:7–23.
2. Farserotu J, Prasad R. A survey of future broadband multimedia satellite systems, issues and trends. *IEEE Communications Magazine* 2000; **38**(6):128–133.
3. Handley M, Floyd S, Whetten B, Kermode R, Vicisano L, Luby M. The reliable multicast design space for bulk data transfer. Internet Society Request for Comments, August 2000, RFC2887.
4. Whetten B, Vicisano L, Kermode R, Handley M, Floyd S, Luby M. Reliable multicast transport building blocks for one-to-many bulk-data transfer. Internet Society Request for Comments, January 2001, RFC3048.
5. Diot C, Dabbous W, Crowcroft J. Multipoint communication: a survey of protocols, functions and mechanisms. *IEEE Journal on Selected Areas in Communications* 1997; **15**(3):277–290.
6. Obraczka K. Multicast transport protocols: a survey and taxonomy. *IEEE Communications Magazine* 1998; 94–102.
7. Koyabe MW, Fairhurst G. Reliable multicast via satellite: a comparison survey and taxonomy. *International Journal of Satellite Communications* 2001; **19**(1):3–23.
8. Luby M, Vicisano L, Gemmell J, Rizzo L, Handley M, Crowcroft J. Forward error correction building blocks. Internet Society Request for Comments, December 2002, RFC3452.
9. Luby M, Vicisano L, Gemmell J, Rizzo L, Handley M, Crowcroft J. The use of forward error correction in reliable multicast. Internet Society Request for Comments, December 2002, RFC3453.
10. Gargione F, Iida T, Valdoni F, Vatalaro F. Services, technologies, and systems at Ka band and beyond: a survey. *IEEE Journal on Selected Areas in Communications* 1999; **17**(2):133–144.
11. Davarian F. Ka-band propagation research using ACTS. *International Journal of Satellite Communications* 1996; **14**(3):267–282.
12. Enjamio C, Vilar E, Gremont B, Fontan FP. Considerations on short distance diversity gain and dynamic rain fade. In *Proceedings of AIAA ICSSC*, May 2002, Online Technical Information.
13. Towsley D, Kurose J, Pingali S. A comparison of sender-initiated and receiver-initiated reliable multicast protocols. *IEEE Journal on Selected Areas in Communications* 1997; **15**(3):398–406.
14. Yavatkar R, Griffoen J, Sudan M. A reliable dissemination protocol for interactive collaborative applications. In *Proceedings of ACM Multimedia*, 1995; 333–344.
15. Paul S, Sabnani KK, Kristol DM. Reliable multicast transport protocol (RMTP). *IEEE Journal on Selected Areas in Communications* 1997; **15**(3):407–421.
16. Lehman LH, Garland SJ, Tennenhouse DL. Active reliable multicast. In *Proceedings of INFOCOM*, vol. 2, April 1998; 581–589.
17. Speakman T *et al*. Pragmatic general multicast (PGM) reliable transport protocol specification. Internet Society Request for Comments, December 2001, RFC3208.
18. Kasera SK, Bhattacharyya S, Keaton M, Zabele S, Kiwior D, Kurose J, Towsley D. Scalable fair multicast using active services. In *Proceedings of DANCE*, May 2002; 333–343.
19. Macker JP, Adamson RB. The multicast dissemination protocol toolkit. In *Proceedings of IEEE MILCOM* 1999; vol. 1, November 1999, 626–630.
20. Floyd S, Jacobson V, Liu C, McCanne S, Zhang L. A reliable multicast framework for light-weight sessions and application level framing. *IEEE Journal on Selected Areas in Communications* 1997; **15**:407–421.
21. Calo SB, Easton MC. A broadcast protocol for file transfer to multiple sites. *IEEE Transactions on Communications* 1981; **29**:1701–1707.
22. Mase K, Takenaka T, Yamamato H, Shinohara M. Go-back-N ARQ schemes for point-to-multipoint communications. *IEEE Transactions on Communications* 1983; **31**:583–589.
23. Gopal IS, Jaffe JM. Point-to-multipoint communication over broadcast links. *IEEE Transactions on Communications* 1984; **32**:1034–1044.

24. Sabnani K, Schwartz M. Multidestination protocols for satellite broadcast channels. *IEEE Transactions on Communications* 1985; **33**:232–240.
25. Blahut RE. *Theory and Practice of Error Control Codes*. Addison-Wesley: MA, 1984.
26. MacAuley AJ. Reliable broadband communication using a burst erasure correcting code. In *Proceedings of ACM SIGCOMM*, Philadelphia, PA, September 1990.
27. Rizzo L. Effective erasure codes for reliable computer communication protocols. *ACM Computer Communication Review* 1997; **27**(2):24–36.
28. Nonnenmacher J, Biersack EW, Towsley D. Parity-based loss recovery for reliable multicast transmission. *IEEE/ACM Transactions on Networking* 1998; **6**(4):349–361.
29. Rizzo L, Vicisano L. A reliable multicast data distribution protocol based on software FEC techniques. In *Proceedings of HPCS*, June 1997; 116–125.
30. Rubenstein D, Kasera S, Towsley D, Kurose J. Improving reliable multicast using active parity encoding services. In *Proceedings of IEEE INFOCOM*, March 1999; 1248–1255.
31. Rhee I, Joshi SR, Lee M, Muthukrishnan S, Ozdemir V. Layered multicast recovery. In *Proceedings of IEEE INFOCOM*, vol. 2, 2000; 805–813.
32. Donner A, Bovelli S, Shabdanov S. Reliable multicast based on DVB-RCS. In *Proceedings of AIAA ICSSC*, May 2002, Online Technical Information.
33. Ernst H, Shabdanov S, Donner A, Scalise S. Reliable multicast for land mobile satellite channels. In *Proceedings of AIAA ICSSC*, April 2003, Online Technical Information.
34. Tommasi F, Molendini S, Tricco A. Design of the satellite reliable distribution protocol (SRDP). In *Proceedings of IEEE GLOBECOM*, December 2003, Online Technical Information.
35. Radoslavov P, Papadopoulos C, Govindan R, Estrin D. A comparison of application-level and router-assisted hierarchical schemes for reliable multicast. In *Proceedings of IEEE INFOCOM*, 2001; 229–238.
36. Kasera SK, Kurose J, Towsley D. A comparison of server-based and receiver-based local recovery approaches for reliable multicast. In *Proceedings of IEEE INFOCOM*, March 1998; 988–995.
37. Nonnenmacher J, Lacher M, Jung M, Biersack EW, Carle G. How bad is reliable multicast without local recovery? In *Proceedings of IEEE INFOCOM*, vol. 3, 1998; 972–979.
38. Jianhua H. Performance evaluation on the satellite based reliable multicast with and without local recovery. In *Proceedings of ECUMN*, 2000; 311–318.
39. Kumar SK, Bhattacharyya S, Keaton M, Kiwior D, Kurose J, Towsley D, Zabele S. Scalable fair reliable multicast using active services. *IEEE Network* 2000; **14**(1):48–57.
40. Rubenstein D, Maxemchuk NF, Shur D. A centralized, tree-based approach to network repair service for multicast streaming media. In *Proceedings of NOSSDAV*, June 2000, Online Proceedings, www.nossdav.org/2000/
41. Golestani JS, Sabnani KK. Fundamental observations on multicast congestion control in the Internet. In *Proceedings of IEEE INFOCOM*, vol. 2, March 1999; 990–1000.
42. Bhattacharyya S, Towsley D, Kurose J. The loss multiplicity problem in multicast congestion control. In *Proceedings of IEEE INFOCOM*, vol. 2, March 1999; 856–863.
43. Bhattacharyya S, Towsley D, Kurose J. A novel loss indication filtering approach for multicast congestion control. *J. of Comp. Commun.* 2001; **24**(5–6):512–514.
44. Thapliyal P, Li J, Kalyanaraman S. LE-SBCC: loss-event oriented source-based multicast congestion control. *Multimedia Tools and Applications Journal* 2002; **17**(2–3):257–294.
45. Rhee I, Balaguru N, Rouskas GN. MTCP: scalable TCP-like congestion control for reliable multicast. In *Proceedings of IEEE INFOCOM*, vol. 3, 1999; 1265–1273.
46. Rizzo L. PGMCC: a TCP-friendly single-rate multicast congestion control scheme. In *Proceedings of ACM SIGCOMM*, August 2000; 17–28.
47. Rizzo L, Iannaccone G, Vicisano L, Handley M. PGMCC single rate multicast congestion control: protocol specification. Internet Draft, June 2003, work in progress, draft-ietf-rmt-bb-pgmcc-02.txt.
48. Widmer J, Handley M. Extending equation-based congestion control to multicast applications. In *Proceedings in ACM SIGCOMM*, August 2001; 275–286.
49. Widmer J, Handley M. TCP-friendly multicast congestion control: protocol specification. Internet Draft, July 2003, work in progress, draft-ietf-rmt-bb-tfmcc-02.txt.
50. Macker JP, Adamson RB. A TCP-friendly, rate-based mechanism for NACK-oriented reliable multicast congestion control. In *Proceedings of IEEE GLOBECOM*, vol. 3, 2001; 1620–1625.
51. McCanne S, Jacobson V, Vetterli M. Receiver-driven layered multicast. In *Proceedings of ACM SIGCOMM*, August 1996; 117–130.
52. Vicisano L, Crowcroft J, Rizzo L. TCP-like congestion control for layered multicast data transfer. In *Proceedings of IEEE INFOCOM*, vol. 3, March 1998; 996–1003.
53. Li X, Paul S, Ammar M. Layered video multicast with retransmissions: evaluation of hierarchical rate control. In *Proceedings of IEEE INFOCOM*, vol. 3, March 1998; 1062–1072.
54. Vicisano L. Notes on a cumulative layered organization of data packets across multiple streams with different rates. UCL Computer Science Research Note, May 1998, RN/98/25.

SAT : 785

55. Roca V. Packet scheduling for heterogeneous multicast transmissions. In *Proceedings of PfHSN*, August 1999; 101–116.
56. Donahoo MJ, Ammar MH, Zegura EW. Multiple channel multicast scheduling for scalable bulk-data transport. In *Proceedings of IEEE INFOCOM*, vol. 2, March 1999; 847–855.
57. Durst RC, Miller GJ, Travis EJ. TCP extensions for space communication. In *Proceedings of ACM Mobicom*, November 1996; 15.
58. Akyildiz IF, Morabito G, Palazzo S. TCP-Peach: a new congestion control scheme for satellite IP networks. *IEEE/ACM Transactions on Networking* 2001; **9**(3):307–321.
59. Akyildiz IF, Morabito G, Palazzo S. Research issues for transport protocols in satellite IP networks. *IEEE Personal Communications* 2001; **8**(3):44–48.
60. Bharadwaj VG, Baras JS, Butts NP. Internet service via broadband satellite networks. In *Proceedings of SPIE*, February 1999; 169–180.
61. Zhou X, Liu X, Baras JS. TCP over GEO satellite hybrid networks. In *Proceedings of IEEE MILCOM*, vol. 1, October 2002; 29–34.
62. Wang C, Leung VCM. Performance evaluations of SRMTP for reliable multicasting over satellite networks. In *Proceedings in WCNC*, vol. 3, March 2003; 1813–1818.
63. Cao G, Wu Y. Reliable multicast via satellite. In *Proceedings of IIT*: *Coding and Computing*, April 2001; 183–188.
64. Cho S, Akyildiz IF. A poker-game-based feedback suppression algorithm for satellite reliable multicast. In *Proceedings of IEEE GLOBECOM*, vol. 3, November 2002; 2930–2934.
65. Ichihara E, Kikuchi K, Tsuchida T, Kawazoe K, Kazama H. Reliable IP-multicast protocol for bi-directional satellite communication systems. In *Proceedings of AIAA ICSSC*, April 2003, Online Technical Information.
66. Akkor G, Baras JS, Hadjithedosiou M. A feedback implosion suppression algorithm for satellite reliable multicast. In *Proceedings of IEEE GLOBECOM*, December 2003, Online Technical Information.
67. Luby M, Mitzenmacher M, Shokrollahi A, Spielman D. Effcient erasure correcting codes. *IEEE Transactions on Information Theory* 2001; **47**(2):569–584.
68. Luby M. Information additive code generator and decoder for communication systems. U. S. Patent No. 6,307,487, October 2001.
69. Luby M. LT codes. In *Proceedings of IEEE Symposium on Foundations of Computer Science*, November 2002; 271–280.
70. Linder H, Miloucheva I, Clausen HD. A FEC based multicast transport protocol for multimedia applications in satellite environments. In *Proceedings of IPCCC*, 1997; 419–425.
71. Gossink DE, Macker JP. Reliable multicast and integrated parity retransmission with channel estimation considerations. In *Proceedings of IEEE GLOBECOM*, vol. 6, 1998; 3608–3613.
72. Yoon J, Bestavroz A, Matta I. Adaptive reliable multicast. In *Proceedings of IEEE ICC*, September 1999; 1542–1546.
73. Cho S, Goulart A, Akyildiz IF, Jayant N. An adaptive FEC with QoS provisioning for real-time traffic in LEO satellite networks. In *Proceedings of IEEE ICC*, June 2001; 2938–2942.
74. Manousakis K, Baras JS. Reliable multicasting for flat hierarchy networks based on adaptive air caching. In *Proceedings of MILCOM*, October 2002; 1295–1299.
75. Akyildiz IF, Fang J. TCP-Peachtree: a multicast transport protocol for satellite IP networks. *IEEE Journal on Selected Areas in Communications* 2004; to appear.

