Traffic Behavior of Scalable Multicast: Self-similarity and Protocol Dependence

Öznur Özkasap[a] and Mine Çağlar[b]

[a]Department of Computer Engineering, Koç University, 34450 Sarıyer, Istanbul, Turkey

[b]Department of Mathematics, Koç University, 34450 Sarıyer, Istanbul, Turkey

The development of high-speed networks and the expansion of the Internet have increased both geographical extent and participant population of applications such as videoconferencing, multimedia dissemination, electronic stock exchange, and distributed cooperative work. The key property of this type of applications is the need to distribute data among multiple participants together with application specific quality of service needs which fact makes multicast protocols an essential underlying communication structure. In this paper, we analyze traffic characteristics of two scalable multicast protocols, namely Bimodal Multicast (Pbcast) and Scalable Reliable Multicast (SRM), each having different approaches for loss recovery and providing reliability. Particularly, our simulation studies demonstrate that epidemic approach of Bimodal Multicast generates a more desirable traffic than SRM with lower overhead traffic and transport delays. SRM delays show long-range dependence and self-similarity whereas Bimodal Multicast delays are short-range dependent. Self-similarity and long-range dependence are ubiquitous in wide area networks, which lead to adverse consequences in network performance. We elaborate on the protocol mechanisms as the underlying factor in our empirical results. The intrinsic relation of these mechanisms to traffic characteristics is explored.

1. INTRODUCTION

The development of high-speed networks and the expansion of the Internet have increased both geographical extent and participant population of applications such as videoconferencing, multimedia dissemination, electronic stock exchange, and distributed cooperative work. The key property of this type of applications is the need to distribute data among multiple participants together with application specific quality of service needs, which fact makes multicast protocols an essential underlying communication structure. For instance, *scalability* of multicast communication as the participant size as well as the network size increases is a significant issue. In the case of an increase in participant and network size, the intention is to have overhead traffic of a multicast protocol remain almost constant or grow very slowly.

In this paper, we study a recent scalable multicast protocol, namely Bimodal Multicast (Pbcast), in comparison to Scalable Reliable Multicast (SRM) protocol [1]. Bimodal Multicast emerges as both reliable and scalable, and provides remarkably stable delivery output [2]. SRM is also scalable, but having best-effort reliability can be problematic in the presence of low levels of system wide noise or by transient elevated rates of message loss [3,4]. On the other hand, multicast transport protocols offering strong reliability guarantee such as atomicity, virtual synchrony, delivery ordering, and network-partitioning support have limitations in terms of scalability and throughput stability. Bimodal Multicast's reliability guarantees are midway between the very strong virtual synchrony and the much weaker best-effort guarantees.

The *loss recovery* mechanism of most reliable multicast transport protocols are either pure receiver-initiated or a hybrid of sender and receiver-initiated approaches. In the case of large-scale multicast applications, pure sender-initiated approach is impractical as it may cause ACK implosion. In the receiver-initiated approach, upon detecting message losses, receivers request their retransmission by generating negative acknowledgements (NACKs). On the other hand, Bimodal Multicast provides an epidemic loss recovery mechanism as a novel approach, which has promising outcomes in terms of robustness and overhead [2]. In particular, SRM serves as a reference protocol here with its receiver-initiated nonhierarchical feedback control mechanism for loss recovery that is based on the principles of TCP, the unicast transport prevalent in the present networks.

In this paper, we focus on the traffic that scalable multicast protocols generate. In particular, our simulation studies demonstrate that epidemic approach of Bimodal Multicast generates a more desirable traffic than SRM with lower overhead traffic and transport delays. SRM delays show long-range dependence and self-similarity whereas Bimodal Multicast delays are short-range dependent in all cases we study. We elaborate on the protocol mechanisms as the underlying factor in our empirical results. The intrinsic relation of these mechanisms to traffic characteristics is explored. Our results can be considered toward the general problem of integration of multicast communication to the Internet. The ultimate aim is to discover and develop multicast protocols that not only feed well-behaved traffic discretely into the existing networks, but also can cope with the existing self-similar traffic and its adverse consequences. It has already been shown that Bimodal Multicast imposes constant loads on links and routers if configured correctly [5]. In Section 2, we review the related work on multicast traffic research, and state the aspects that motivate this study. In Section 3, Bimodal Multicast and SRM are reviewed. We present our simulation results in Section 4, discuss them in Section 5 along with the performance implications. Finally, the conclusions and future work appear in Section 6.

2. RELATED WORK AND MOTIVATION

Analyses of fine-grained measurements over the last decade have shown that network traffic is often bursty on a wide range of time scales with significant correlation across arbitrarily large time lags. These characteristics, called self-similarity (SS) and long-range dependence (LRD) respectively, imply significant queuing delays and degraded network performance. Analysis reveals that long-range dependence exists in the network traffic as well as self-similarity [6]. It is well known that Hurst parameter denoted by H is a measure of persistence of correlations in traffic when it takes values in (0.5, 1) and indicates long-range dependence. Self-similarity in the presence of long-range dependence has adverse consequences on network performance.

Most efforts in multicast communication have focused on developing new protocols and applications. These protocols have been compared with respect to several performance measures such as scalability, reliability and congestion control. However, the nature of the traffic stream generated by each type of protocol particularly with respect to self-similarity has not been studied extensively. Reliable traffic analysis is crucial for different performance evaluation tasks such as dimensioning of buffers, bandwidth allocation decisions, and network planning and design [7,8].

We are actually motivated by the fact that both the network and transport level multicast protocols are still evolving as the deployment of multicast on large-scale has been slower than expected. Although SS property of TCP traffic has been analyzed for unicast communication, multicast traffic has not been incorporated from this perspective. Our preliminary study [9] relies on the comparison of traffic generated by two multicast transport protocols. Through the message delay analysis, we empirically see that traffic belonging to SRM is LRD whereas that of Bimodal Multicast is short range dependent under identical settings. This is the basis of our comparative study in the present paper for several parameters of the multicast network. The questions we ultimately aim to answer are as follows:

- What is the best protocol that discretely feeds well-behaved traffic into the existing networks?
- Which multicast protocol(s) can best cope with the existing SS traffic and its adverse performance consequences in WAN/Internet?
- What are the principles and mechanisms behind these protocols that qualify them to be superior/better?
- Can these principles be also used for the modification and/or design of existing protocols for unicast communication as well as wireless data communication?

In support of the above, it has been shown that the transport layer mechanisms are important components in translating heavy-tailed file size distributions at the application layer into link traffic self-similarity. Larger time scales like minutes and hours are affected by application and human causes, whereas TCP is capable of shaping the traffic at the time scales of few milliseconds to tens of seconds [10,11]. The smaller time scales are also significant as they could be relevant for traffic engineering purpose of real systems with finite buffer sizes. There exist recent efforts to incorporate the transport layer to mathematical models of traffic [11,12]. In this paper, we provide empirical evidence of our claim that self-similarity is protocol dependent. The multicast protocols under investigation are reviewed next.

3. REVIEW OF MULTICAST PROTOCOLS

Bimodal Multicast [2] is a novel option in the spectrum of multicast protocols that is inspired by prior work on epidemic protocols [13], Muse protocol for network news distribution [14], and the lazy transactional replication method of [15]. Bimodal Multicast is based on an epidemic loss recovery mechanism. It has been shown to exhibit stable throughput under failure scenarios that are common on real large-scale networks [2]. In contrast, this kind of behavior can cause other reliable multicast protocols to exhibit unstable throughput.

Bimodal Multicast consists of two sub-protocols, namely an optimistic dissemination protocol and a two-phase anti-entropy protocol. The former is a best-effort, hierarchical multicast used to efficiently deliver a multicast message to its destinations. This phase is unreliable and does not attempt to recover a possible message loss. If IP multicast is available in the underlying system, it can be used for this purpose. Otherwise, a randomized dissemination protocol can play this role. The second stage of Bimodal Multicast is responsible for message loss recovery. It is based on an anti-entropy protocol that detects and corrects inconsistencies in a system by continuous gossiping. The two-phase anti-entropy protocol progresses through unsynchronized rounds. In each round:

- Every group member randomly selects another group member and sends a digest of its message history. This is called a 'gossip message'.
- The receiving group member compares the digest with its own message history. Then, if it is lacking a message, it requests the message from the gossiping process. This message is called 'solicitation', or retransmission request.
- Upon receiving the solicitation, the gossiping process retransmits the requested message to the process sending this request.

SRM [1] is a reliable multicast protocol which is inspired by the principles of IP multicasting, application level framing (ALF), and the TCP/IP architecture design. The protocol necessitates the basic IP delivery model and forms reliability on an end-to-end basis. Similar to TCP that adaptively sets timers or congestion control windows, SRM algorithms dynamically regulate their control parameters such as request and repair timers, based on the observed performance within a session. The protocol aims to scale well both to large networks and sessions, and exploits a receiver-based reliability mechanism.

In SRM, each group member multicasts low-rate, periodic session messages that report the sequence number state for active sources, or the highest sequence number received from every member. As well as the reception state, the session messages also contain timestamps that are used to estimate the distance from each member to every other. Members utilize session messages to determine the current participants of the session. In addition to state exchange, receivers use the session messages to estimate the one-way distance between nodes. The session packet timestamps are used to estimate the host-to-host distances needed by loss recovery mechanisms. The random delay before sending a request or repair packet is a function of that member's distance in seconds from the node that triggered the request or repair. Repair requests and retransmissions are multicast to the whole group. A lost packet ideally triggers only a single request from a host just downstream of the point of failure.

As a comparison, in Bimodal Multicast, if a process detects a message loss, it requires a unicast request and repair message to recover the loss. In the case when one or both of these control messages get lost on a noisy link, additional control messages are required. For SRM protocol, on the other hand, in order to guarantee reliable delivery, a process multicasts request message to the whole group when it detects a message loss. Request and repair timers are exploited to suppress duplicate requests and repairs for the same message loss. A corresponding repair message in response to a request is similarly in the form of multicast to the whole group. This feature of SRM's loss recovery mechanism makes its background overhead and bandwidth requirements to increase as a function of group size, whereas Bimodal's background overhead is scalable and does not increase with the group size. Recent studies [2,3,4] have shown that, for the SRM protocol, random packet loss can trigger high rates of overhead messages. In addition, this overhead grows with the size of the system whereas it remains almost constant for Bimodal Multicast [16].

4. SIMULATION RESULTS

4.1. Simulation Settings and Method of Analysis

We use our Bimodal Multicast protocol model [16] and existing module for SRM both implemented on ns-2 network simulator [17]. The simulation scenarios consist of transit-stub topologies with node number N ranging from 20 to 120. The sender is located on a central node and receivers are located at all other nodes on the network. Transit-stub topologies approximate the structure of the Internet that can be viewed as a collection of interconnected routing domains where each domain can be classified as either a stub or a transit domain. Stub domains correspond to interconnected LANs and the transit domains model WAN or MANs [18]. We used gt-itm topology generator for producing transit-stub topologies [19]. One of the sample transit-stub topologies consisting of total 40 nodes is shown in Fig.1. A certain link message drop probability is set on every link that forms a randomized system-wide noise. Each link has a bandwidth of 1.5MB. The operating parameters, namely the group/network size and the system-wide message drop rate are varied. This scenario primarily focuses on the impact of randomized message loss over traffic generated by the protocol. We obtain our results from several runs of simulations, each run consisting of a sequence of 35000 multicast data messages transmitted by the sender with the rate 50 messages (each with size 210 bytes) per second. This can be considered as a large file being multicast at constant bit rate to all receivers.

Our approach is to consider the delay of packets over the network. Recently, packet delay measurements over the Internet are used to trace the conditions of the network between an origin and destination pair [20,21,22]. In our simulations, such measurements represent traffic at the transport level. A typical receiver with a fixed distance (3 hops) and the farthest receiver from the sender are fixed for analysis. The delay of a message is calculated as the difference of the deployment time at the sender from the receive time at a receiver. In most cases the delay measurements form a stationary sequence; the exceptions are explained below. We estimate the Hurst parameter from the stationary part of the delay sequence. We use the wavelet estimation method and nonstationarity analysis tool of Veitch and Abry [23], through the authors' Matlab implementation with Daubechies wavelets. The nonstationarity tool splits data into equally spaced blocks and compares the mean and Hurst parameters in these blocks in order to detect nonstationarity. The interarrival distribution of the data messages and the throughput at a typical receiver are also analyzed as other performance measures.

4.2. Results

We compare SRM and Bimodal Multicast in the same simulation settings with the same sequence of random numbers. Initially, several independent runs of ns were obtained to observe the effect of randomness in our results. For the statistical precision of our results for long-range dependence, each run lasts for $35000 \geq 2^{15}$ messages. With such a long sequence, independent runs with different seeds showed almost no random variation in the estimated Hurst parameters and other statistics of performance. That is why, we have chosen to generate 3 independent random topologies for each group size and report their average statistics here. This has introduced randomness in our experiments while reinforcing our results as all topologies for a fixed group size showed similar performance characteristics within some variation.

The most striking result of this paper is displayed in Fig. 2. We have started with small groups of size 20 and went up to 120 all on a transit-stub topology. The Hurst parameter estimates are given in Fig. 2 for Bimodal Multicast and SRM at two levels of system-wide drop rate, namely 1% and 2%. Both protocols behave similarly up to group size 80; they generate short-range dependent traffic with H values around 0.5. When the group size increases to 100 or more, SRM delays show long-range dependence with H values statistically significantly greater than 0.7. Bimodal Multicast continues to produce short-range dependent traffic for groups of size 100 and 120. In fact, as indicated in [9], on a tree topology of much larger network, namely 500-nodes where 300 are group members, Bimodal still produces short-range dependent traffic with H value 0.54 whereas SRM traffic has H value 0.65. Our current results hence detect the threshold group size for which SRM's scalability ceases under system wide noise.

The load that group size asserts on the network in the SRM case is apparent in Fig. 3. In group sizes smaller than 100 and for all cases of Bimodal Multicast, the delay sequence is stationary starting from the very first message. However, for SRM in left plot of Fig. 3 where N=100, there is a transient period where the delays are huge. The protocol then decreases the delay and stabilizes the throughput (not shown here). Our Hurst parameter estimate corresponds to the latter stationary segment for which the wavelet-scaling diagram is also provided in Fig. 3. For the cases where H is around 0.5, this plot would have a slope around 0. Here, we observe two levels of pattern in the correlation structure. For small octaves, there is a nonlinear scaling until octave 5 and then there is a linear scaling after octave 7, which is relevant for long-range dependence. Equivalently, the pattern changes at a time scale corresponding to about $2^6 = 64$ messages, that is about the order of 1 second in view of the message rate 50/sec. This is in accordance with the results of TCP traffic examined at the transport level in [10]. The smaller time scales might represent the effect of SRM's control actions at the granularity of time-to-live, TTL, and request and repair timers (in analogy to RTT and retransmission timer in the case of TCP), spanning time scales from 0.04 seconds to 1 second. In contrast, the larger scales show the overall effect on the network due to congestion and they span time scales from 1 second to an hour. Although the delay decreases somewhat after the transient period in Fig. 3, it does not do so for N=120 yielding a stationary sequence from the start. In addition to the Hurst parameter, the mean of delay is also strikingly different for the two protocols.

The mean delay for Bimodal multicast is in the order of 0.03 second, whereas it is in the order of 0.70 second for SRM.

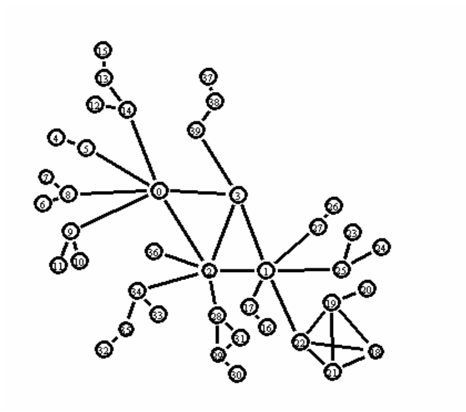

Figure 1. Sample 40-node transit-stub topology.

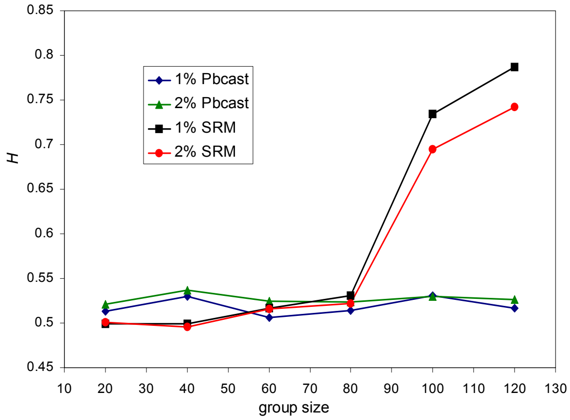

Figure 2. Hurst parameter as a function of group size for Bimodal Multicast and SRM at two levels of system-wide drop rate.

The interarrival distributions reflect the performance implications of the traffic patterns. In Fig. 4, the interarrival distribution of the long-range dependent sequence for SRM is found to be Exponentially distributed, whereas it is found to be Normally distributed for the corresponding Bimodal Multicast traffic, which is short-range dependent. The difference in the standard deviation of the two distributions is remarkable although they have the same mean. Overall, we see that the interarrival distribution is approximately Normally distributed when the network is not pressured. This is true for Bimodal Multicast in all cases, and in smaller size groups for SRM. In the SRM case, as the group size increases, the distribution becomes right skewed (long right tail) and an Exponential distribution fits well. There is a similar difference for their throughputs, which is not documented here. Although the mean throughput is the same, the variance is significantly smaller for Bimodal Multicast.

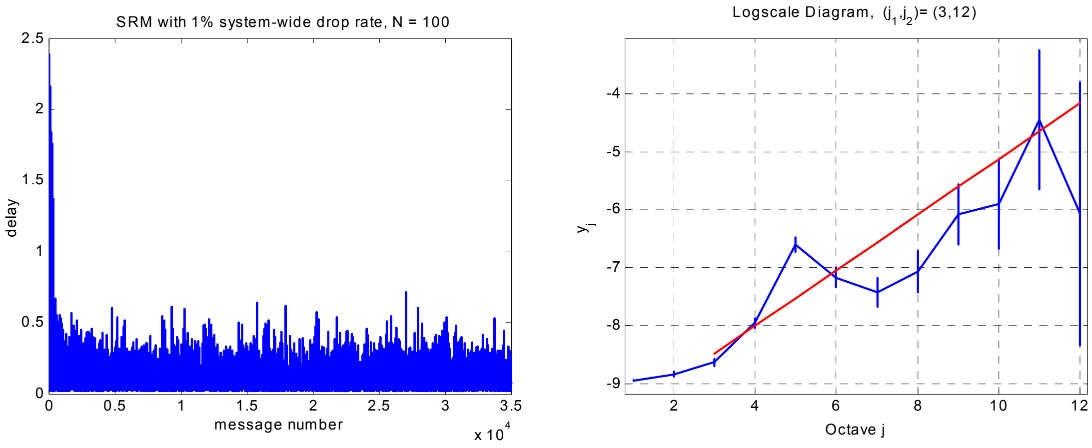

Figure 3. A sample delay sequence for the farthest receiver in a group of size 100 with 1% system-wide drop rate in the case of SRM (left); and the corresponding scaling diagram (right) used for the estimation of H, which is obtained as a result of the wavelet transform of the stationary part of the delay sequence.

The scalability of Bimodal Multicast is remarkable, only at a negligible cost in reliability. At most 1 or 2 losses are encountered in 35000 messages for all simulations. The throughput decreases and the variance of the interarrival increases only slightly as the group size and/or the drop rate increases. The Hurst parameter on the other hand is very stable in response to doubling of drop rate or increase of the group size. Bimodal Multicast provides stable throughput in the sense that it has smaller variance than SRM. On the other hand, SRM makes utmost effort for reliability as no losses have been encountered in our simulations. This comes at a cost of longer delays, slightly lower throughput than Bimodal Multicast (significantly lower for N=120), more variable interarrivals, and most importantly self-similar traffic patterns. To demonstrate this pattern with respect to bigger group sizes, the scaling diagram is given finally in Fig. 5 for SRM (left), and for comparison for Bimodal Multicast (right). The trivial scaling for Bimodal Multicast is apparent. In comparison to Fig. 3 (right), as the group size increases, the small time scaling becomes linear rather than nonlinear. The difference between the scaling graphs for SRM in Fig. 3 and 5 is very similar to the difference in the graphs of only TELNET and FTP traffic of 1990 Bellcore traces and 1994 Bellcore traces, respectively. In [10], this difference has been attributed to the increasing WWW traffic from 0% to 10% in 4 years, that is, to the application types. In our case, it is due to increasing group size and directly affected by the protocol.

5. PERFORMANCE IMPLICATIONS AND DISCUSSION

The offered load to the network is the same, but the resulting traffic patterns are different with the two reliable protocols. Bimodal Multicast continues to produce short-range dependent delays whereas SRM delays become long-range dependent with $H>0.7$ when the group size is over a threshold value. Such a threshold is expected to depend on the parameters of the network, such as noise, link capacity and buffer size; yet occurs at a moderate group size, $N = 100$, in our settings, which is common to many current distributed applications. The mean message delay is found to be much larger for SRM, in the order of 50 times of that for Bimodal multicast in presence of LRD. It is known that under long-range dependence, the queue lengths decay more slowly in contrast to short-range dependent traffic. This has been shown analytically for various self-similar models in addition to simulation studies [24].

Another striking result verified with both approaches is that the utilization factor cannot be practically improved by enlarging buffers [25,26]. Instead, increasing link bandwidth has the effect of decreasing queuing delay more drastically under self-similar and long-range dependent traffic conditions. On the other hand, when H is not very large ($H<0.7$), or short range dependent exists strongly along with LRD, self-similarity is less significant for performance such as buffer occupancy [27]. Some recent studies focus on the queuing behavior of hybrid traffic [see e.g. 28].

In SRM, a process multicasts a request message to the whole group when it detects a message loss in order to guarantee reliable delivery. Request and repair timers are exploited to suppress duplicate requests and repairs for the same message loss. A corresponding repair message in response to a request is similarly in the form of multicast to the whole group. This feature of SRM's loss recovery mechanism makes its background overhead and bandwidth requirements to increase as a function of group size. As a result, there is zero loss in our simulations but at the expense of self-similar traffic with long-range dependence. The delays are high and loss recovery mechanism works, but it imposes such a pattern that implies low performance as a result of $H > 0.5$. This can be explained by the congestion of the network due to increased overhead, which is triggered by the loss recovery mechanism of SRM. In addition to multicasting of request and repair messages, which brings extra load, the timers which try to balance this fact may not be adequately set in the case of congestion. The estimates that these timers are based on have a small probability of being incorrect, and in the case of SRM as the size of the system increases, the absolute likelihood of mistakes rises causing the background overhead rise.

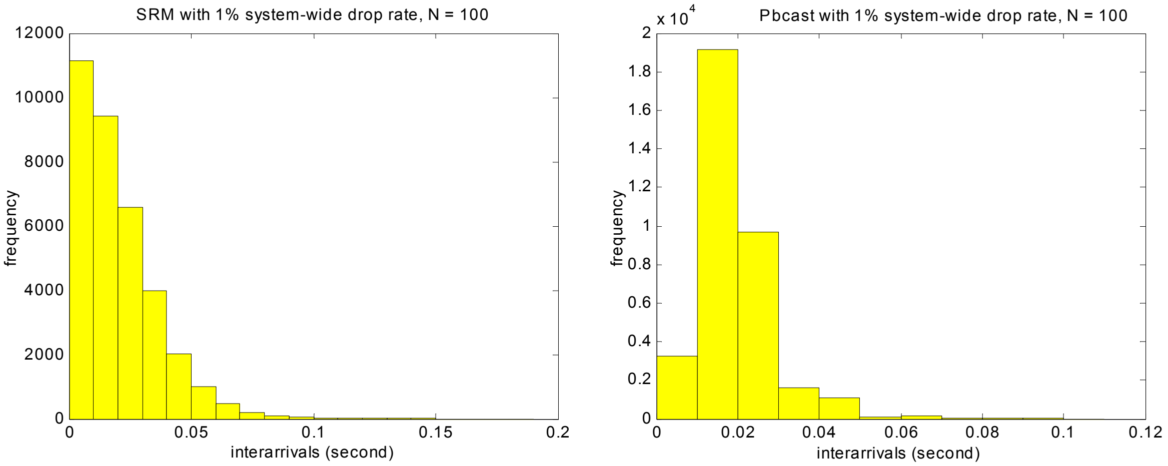

Figure 4. The message interarrival distribution for the receiver in Figure 2 with SRM (left) and the same distribution when Bimodal Multicast is used (right). Exponential distribution fits well in the case of SRM whereas the interarrival spacings of multicasts are roughly Normal in Bimodal multicast with a much tighter dispersion.

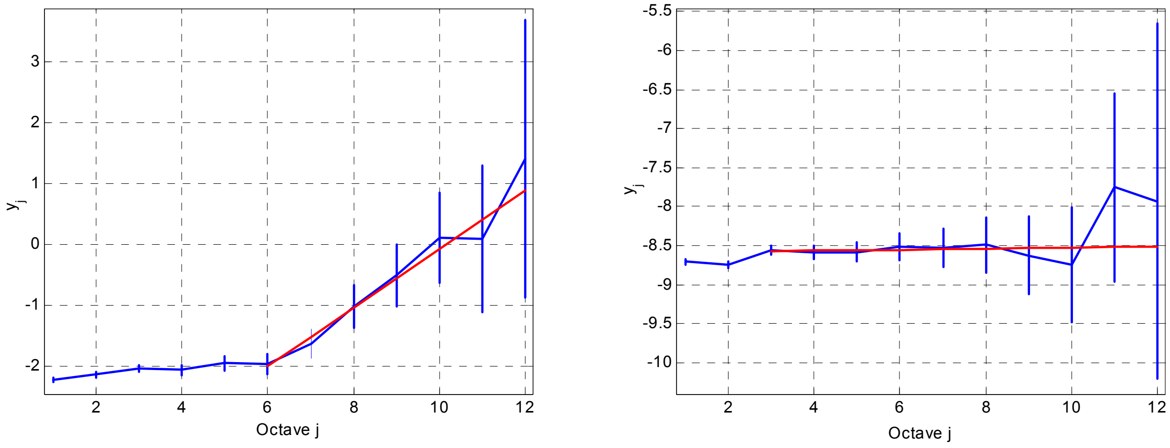

Figure 5. Scaling diagram used for the estimation of H, which is obtained as a result of the wavelet transform of the delay sequence for the farthest receiver in a group of size 120 with 1% system-wide drop rate in the case of SRM (left), and Bimodal Multicast (right). Since the large time linear scaling, which is relevant for long-range dependence becomes prominent in SRM case, we use the larger scales for the estimation of H.

We have captured the effect of SRM not only at small time scales but also showed that a self-similar pattern is induced on the network delays at large time scales as well. Structural models of data traffic account for large time scaling through the effect of the application/user layer, which has also been verified by numerous empirical and simulation work. Recently, the effect of TCP's retransmission mechanisms, in addition to its congestion control, on self-similarity over small to medium time scales has been analyzed in [11]. Our present work demonstrates that SRM having a similar loss recovery mechanism to TCP can induce self-similarity also at those time scales. Ubiquitous presence of self-similarity in data networks might be due to wide spread availability of TCP.

Bimodal Multicast is based on epidemic paradigm for loss recovery. If a process detects a message loss in the end of a gossip round, it requires a unicast request and repair message to recover the loss. Bimodal's background overhead is scalable and does not increase with the group size. Hence, its traffic does not manifest self-similarity with long-range dependence. Although our simulations mimic an infinite size file being transferred with constant bit rate, a heavy load at the application layer, the traffic of Bimodal Multicast is short-range dependent in all group sizes. The anti-entropy protocol uses gossiping, which is a spatial mechanism in contrast to temporal timer mechanism of SRM. This helps explaining the difference in the traffic patterns. Temporal behavior such as heavy-tailed session durations (as a result of user behavior, like OFF times, or heavy-tailed file size distributions) directly translates to long-range dependence. In the case of Bimodal Multicast, the burden is distributed spatially by the gossiping mechanism. See [29] for a similar space versus time stretching analogy established for UDP and TCP, respectively. UDP shows less self-similar characteristic due to lack of reliability compared to TCP. Bimodal multicast provides high level of probabilistic reliability guarantees as well as producing short-range dependent delays and highly scalable stable throughput.

6. CONCLUSIONS AND FUTURE WORK

This paper contributes to identifying better protocols for the design of future's multicast communication. We compare Bimodal Multicast to SRM because the latter is similar to TCP, which is prevalent in the Internet and is known to show more SS than its non-flow controlled counterpart, UDP. Our main conclusions are 1) Bimodal Multicast generates short-range dependent delays which are scalable in number of users like its other superior performance properties scalability and reliability, 2) we confirm that a timer-based loss recovery mechanism generates self-similarity with long-range dependence and hence adverse performance consequences when used, also for multicast communication.

We have identified Bimodal Multicast as the better protocol that discretely feeds well-behaved traffic into the existing networks. We have argued that its novel epidemic loss recovery approach facilitates this outcome. We aim to demonstrate this result by a mathematical proof and also confirm for link level traffic as future work. What happens when Bimodal Multicast mixes with self-similar traffic of WAN/Internet is an open question. We expect that performance results for the multiplexing of short-range and long-range dependent traffic will be verified. We intend to run simulations in the presence of self-similar background traffic. Finally, the principles of epidemic communication can be investigated for the modification and/or design of existing protocols for unicast as well as wireless data communication.

REFERENCES

1. Floyd, S., Jacobson, V., Liu, C., McCanne, S. and Zhang, L., 1997, A Reliable Multicast Framework for Light-weight Sessions and Application Level Framing, IEEE/ACM Trans. Networking, 5(6), 784-803p.
2. Birman, K.P., Hayden, M., Ozkasap, O., Xiao, Z., Budiu, M. and Minsky, Y., "Bimodal Multicast", ACM Transactions on Computer Systems, 17: (2), 41-88, 1999.
3. Liu, C., Error Recovery in Scalable Reliable Multicast, Ph.D. dissertation, University of Southern California, 1997.
4. Lucas, M., Efficient Data Distribution in Large-Scale Multicast Networks, Ph.D. dissertation, Dept. of Computer Science, University of Virginia, 1998.
5. Vogels, W., van Renesse, R., Birman, K., "Using Epidemic Techniques for Building Ultra-Scalable

Reliable Communications Systems", Workshop on New visions for Large-Scale Networks: Research and Applications, Vienna, VA, March 2001.

6. Leland, W. E. , Taqqu, M. S., Willinger, W. and Wilson, D. V. On the Self-Similar Nature of Ethernet Traffic (Extended Version), IEEE/ACM Trans. On Networking, 2, pp. 1-15, 1994.
7. Caglar, M., Krishnan, K.R., Saniee, I., Estimation of Traffic Parameters in High-Speed Data Networks. Proceedings, Sixteenth International Teletraffic Congress Edinburgh, UK, 7-11 June, 1999.
8. Caglar, M., Krishnan, K.R. Parametric Resampling Analysis of Traffic Measurements for Capacity Management. Proceedings, ITC Specialist Seminar on IP Traffic Measurement, Modeling and Management, Monterey, CA (USA) 18-20 September, 2000.
9. Ozkasap, O., Caglar, M., Multicast Network Traffic and Long-Range Dependence. Proceedings, IASTED, International Conference on Advances in Communications, July, 2001.
10. Feldmann, A., Gilbert, C., Willinger, W., and Kurtz, T.G., "The Changing Nature of Network Traffic: Scaling Phenomena". Computer Communication Review, 28: 5-29, 1998.
11. Sikdar, B., Vastola, K., "On the Contribution of TCP to the Self-Similarity of Network Traffic", Proceedings of the 2001 Tyrrhenian International Workshop on Digital Communications: Evolutionary Trends of the Internet, Taormina, Italy, September 17-20, 2001.
12. van Foreest, M., Mandjes, M., Scheinhardt, W., "Modeling TCP using feedback fluid queues", unpublished, 2001.
13. Demers, A., Greene, D., Hauser, C., Irish, W., Larson, J., Shenker, S., Sturgis, H., Swinehart, D. and Terry, D., 1987, Epidemic Algorithms for Replicated Database Maintenance, Proceedings of the Sixth ACM Symposium on Principles of Distributed Computing, Vancouver, British Columbia, 1-12p.
14. Lidl, K., Osborne, J. and Malcome, J., 1994, Drinking from the Firehose: Multicast USENET News, USENIX Winter 1994, 33-45p.
15. Ladin, R., Lishov, B., Shrira, L. and Ghemawat, S., 1992, Providing Availability using Lazy Replication, ACM Transactions on Computer Systems, 10(4), 360-391p.
16. Ozkasap, O., Scalability, Throughput Stability and Efficient Buffering in Reliable Multicast Protocols, Technical Report, TR2000-1827, Department of Computer Science, Cornell University.
17. Bajaj, S., Breslau, L., Estrin, D., et al., Improving Simulation for Network Research, 1999, USC Computer Science Dept. Technical Report 99-702.
18. Ken Calvert, Matt Doar and Ellen W. Zegura. Modeling Internet Topology. IEEE Communications Magazine, June 1997.
19. GT-ITM topology generator. http://www.isi.edu/nsnam/ns/ns-topogen.html
20. Borella, M.S., Brewster G.B., Measurement and Analysis of Long-Range Dependent Behavior of Internet Packet Delay, Proceedings of IEEE INFOCOM 1998, Volume: 2, pp. 497-504.
21. Li, Q., Mills, D.L., On the long-range dependence of packet round-trip delays in Internet, ICC 98. Volume: 2 , 1998, pp: 1185 –1191.
22. Andren, J.; Hilding, M.; Veitch, D., Understanding end-to-end Internet traffic dynamics, GLOBECOM 1998. The Bridge to Global Integration, Volume: 2 , 1998, pp: 1118 -1122
23. Veitch, D. and Abry, P., A Wavelet Based Joint Estimator of the Parameters of Long-Range Dependence, IEEE Trans. on Information Theory, 45(3), 1999, 878-897.
24. Willinger, W., Paxon, V., Reidi, R., Taqqu M., Long-Range Dependence and Data Network Traffic. Long-range Dependence: Theory and Applications, P. Doukhan, G. Oppenheim and M. S. Taqqu, eds., Birkhauser, 2001.
25. Norros, I., On the Use of Fractional Brownian Motion in the Theory of Connectionless Networks. IEEE Journal on Selected Areas in Communications, 13: 953-62, 1995.
26. Park, K., Kim, G.T., Crovella, M.E., On the Effect of Traffic Self-similarity on Network Performance, Proceedings of the SPIE International Conference on Performance and Control of Network Systems, November, 1997.
27. Heyman, D.P., Lakshman, T.V., "What are the Implications of Long-Range Dependence for VBR-Video Traffic Engineering?", IEEE Trans. Networking, 4: (3), 301-317, 1996.
28. Borst, S., Mandjes, M., van Uitert, M., "GPS Queues with Heterogeneous Traffic Classes", INFOCOM 2002, June 23-27, New York, USA, 2002.
29. Park, K., Kim, G.T., Crovella, M.E., On the relationship between file sizes, transport protocols, and self-similar network traffic. Proc. Fourth International Conference on Network Protocols, October, 1996.