Live Network Streaming with Utilities and Cost

Ymir Vigfusson[§], Annie Liu[†], Winnie Liu[†], Yuanchao Zhu[†], Gunnar S. Sigurbjornsson[§], Kristjan V. Jonsson[§], Qi Huang[‡], Ken Birman[‡], Daniel A. Freedman[‡†]

[§]School of Computer Science, Reykjavik University, Reykjavik, Iceland
[‡]Department of Computer Science, Cornell University, Ithaca, New York, USA
[†]School of Electronics Engineering and Computer Science, Peking University, Beijing, China
{ymir,kristjanvj,gunnars10}@ru.is, {ken,qhuang,dfreedman@cs.cornell.edu}

ABSTRACT

The growth in Internet traffic associated with video streaming and sharing of videos is so rapid that it may soon dwarf all other forms of Internet content. One reason for this is that only some forms of content can be cached (such as movies). Data generated in real time such as by live video broadcasts (e.g. sports games, IPTV or new episodes of popular TV shows), chat systems, immersive virtual reality applications and games typically can't be cached at all, and in today's systems, each client may pull such information on its own point-to-point (TCP) stream directly from the data center, even if large numbers of clients share interest in at least some aspects of the data. We propose a new system called GRADIENT *aimed at reducing the load on providers of such and enabling scalable, bandwidth-sensitive streaming service for heterogeneous consumers. The core of the system is an overlay networking architecture intended to run directly on a content hosting platform, and which optimizes aggregate bandwidth use by transforming in-flight data to match the ideal stream quality – expressed as an economic utility – of the consuming client.*

1. INTRODUCTION

Recent years have seen skyrocketing demand for Internet bandwidth, increasingly dominated by real-time streaming of short-lived content; video is common, but in many forms: movies, Internet radio, podcasts and others. If trends continue then Internet video alone will generate almost 10 exabytes per month by the end of 2012, accounting for nearly 50 percent of all Internet traffic [2]. Faced with a competitive landscape, ISPs and content providers are exploring technologies to help satisfy the growing demand alongside the purchase of expensive infrastructure.

Reducing the bandwidth consumption of simultaneous replicated content is a challenge which usually leverages two main tools: caching of content and multicasting. Some forms of video content, such as downloads of unencrypted movies or films where many users will share the same encryption key, can be cached. For these cases, a wide variety of caching options exist [5, 7], of which is the Akamai content distribution network (CDN) is arguably the most famous [9]. Likewise, multicast techniques can reduce the overall network traffic by taking advantage of the packet replication and forwarding within the network infrastructure. However, the deployment of the efficient network-level multicast (IP multicast [4]) over the wide-area Internet has failed [11] and so more expensive application-level overlays are generally used [8].

The devices used by content subscribers have become increasingly heterogeneous – mobile devices, for instance, are projected to consume over 7.6 exabytes of video per month in 2016 [3], implying that a range of subscription rates and policies must be applied over the user base [6]. Even if multiple users are streaming the same event, such as watching the opening ceremony of the Olympics, a smartphone user will need a differently transcoded version than the people watching via Internet television (IPTV) or on their laptops. As another example, different consumer groups may desire different local ads or sub-titles to be embedded into their video streams. Even further, avatars in a virtual world can be viewed as subscribers to updates about objects in their vicinity, and may want more detailed updates for objects that are closer to them in this world. While this growing heterogeneity of device types, research on CDNs has generally assumed a homogeneous population of end-users [1, 13, 10]. Thus most current systems assume multiple video streams to be sent from the source at different resolutions or that a single high-quality stream is transcoded by the receiver who then incurs cost for last-mile traffic owing to unnecessarily detailed video.

We pose the following question: *How can we deliver live dynamic content (such as video broadcasts, or financial stock data) over the Internet to large number of heterogeneous users simultaneously while balancing bandwidth costs, traffic rates and end-user utility?* Here, live content refers to content streams that must be transmitted to multiple receivers simultaneously, such as a live broadcast, an IPTV channel, ticker updates for financial stocks or object updates in a virtual world. In particular, we are not focused on streams with a pause or rewind functions or the video-on-demand problem.

2. THE GRADIENT CDN

To make progress towards the research question, we propose a novel networked content delivery system called GRADIENT to address the complex caching and multicasting issues associated with live streaming of dynamic content to a heterogeneous user population. The systems architecture consists of one or more content providers which together form a cooperative network of GRADIENT CDN nodes. The CDN nodes form a dynamic overlay over which the content is delivered, and for our initial prototypes we will look at spanning trees. The concept of CDN nodes is general. An architecture in which CDN servers are hosted by ISPs to reduce redundant incoming bandwidth is a logical scenario, and another example is that GRADIENT nodes may as well be integrated into set-top boxes [8]. Our approach to the problem resembles content-hosting, and in fact the expected deployment model would employ a geographically distributed set of ISPs or small data centers of the kind operated by today's CDN providers. However, whereas today's content-hosting sites cache objects (files, web pages) [9], our focus is on content that cannot be usefully cached. The GRADIENT project aims to exploit and develop two techniques that improve on existing CDNs: in-network data transformation, and algorithms to balance bandwidth costs with end-user utilities. Our design is focused on modularity and incremental deployment.

In-network data transformation. Dynamic content has substantial levels of redundancy, even when user interests are relatively heterogeneous. Widespread use of streaming video occurs when Internet users watch major events online (such as SuperBowl or the Olympics), and like television users, such clients have little tolerance for lagged data. Here, large numbers of users have essentially the same needs, but since they may access the streams from a variety of devices (with different screen sizes and resolutions, or different connectivity properties). The current solution is to provide each user with a direct connection to a content-server due to the lack of robust multicast technologies. Similar issues arise for newscasts of fast-breaking events, transmission of financial data and virtual on-line worlds.

Our insight is that a data-rich channel can be transformed on-the-fly in-network to create the dynamic content for end-users. For instance, we could add personalized advertisements, subtitles or encryption keys to IPTV broadcasts; filters or aggregates to financial stock updates, or reduce the update rate for distant objects in the virtual world to which the user has subscribed. The same mechanism will also allow the system to tailor to heterogeneous devices, e.g. transcoding a high definition broadcast to adapt its resolution to serve a population of heterogeneous devices – from cell phones to tablets to IPTV – lowering overall bandwidth costs without affecting viewing experience.

Specifically, the in-network transformations will be applied with pluggable “serverlets” designed to execute within the CDN nodes of GRADIENT. The serverlets encapsulate application-speciïňĄc details such as the stream data format, client attributes, and the ways to transform a the data-rich objects into more specialized ones. Open issues include understanding what kinds of content may be subject to such transformation and which dynamic content is not, to assess the effect of the transformation on quality and traffic rates, how transformation should be meaningfully expressed and used by content providers, and to learn how computationally intensive such transformation methods can be without overloading the nodes.

Balancing bandwidth costs with end-user utility. The GRADIENT content delivery system is currently designed to use a spanning-tree overlay, similar to most multicast network architectures, with virtual links connecting GRADIENT CDN nodes. The question is to determine what nodes the in-network processing and connecting to the diverse end-users should be done. Specifically, we need to optimize the overlay to deliver the exact stream quality demanded by users while minimizing bandwidth costs. To do so, we propose to apply an economics framework, considering two primary inputs in determining the optimal network overlay. On the one hand, we consider the cost for network edges to carry traffic, similar to standard bandwidth pricing. On the other, we leverage the perceived utility by end-users for receiving the stream at a given quality.

As posed, the exact solution for this optimization problem is intractable – it is NP-complete. Nevertheless, we have developed algorithms that give an approximate optimal solution. For instance, in the case of video streams whose quality and traffic rates can be downgraded by GRADIENT CDN nodes, we have derived a primal-dual approximation algorithm which produces a solution whose total cost (the difference between total network traffic costs and aggregate end-user utilities) is within a factor of 5.986 of the optimal in the worst case [12]. In Figure 1 we see that the algorithm has lower total cost compared to a single stream source and a minimum spanning tree streaming protocol in a simulation based on a collection of AS

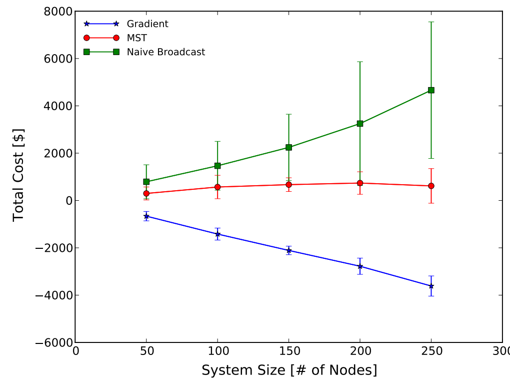

Figure 1: ***Simulation****. The* GRADIENT *optimization is effective compared to a centralized source and a minimum spanning tree (MST) protocol even as system sizes scale up. Error bars represent one standard deviation over 25 trials.*

sub-graphs. The details are deferred to a full report on GRADIENT.

3. CONCLUSION

A number of interesting open questions remain the focus of our continued investigation: How diverse are the classes of content that are amenable to our in-network transformations? How do we best assess the effect of such transformations on stream quality? How should these transformations be expressed, designed, and utilized by the originating content providers to best balance content-domain specificity with ease of development? How can our overlay respond to churn among GRADIENT nodes — realistically low in many common cases such as video streaming, but higher in alternative deployment scenarios? And, finally, how do we ensure that the computational intensity of our transformations do not place too much load on our GRADIENT overlay nodes? In summary, GRADIENT contributes a novel platform for continued study and progress to ever more effective delivery mechanisms.

References

[1] M. Castro, P. Druschel, A. M. Kermarrec, A. Nandi, A. Rowstron, and A. Singh. SplitStream: high-bandwidth multicast in cooperative environments. In *SOSP*, page 313. ACM, 2003.

[2] Cisco. Approaching the Zettabyte Era. *Cisco Visual Networking Index*, page 23, 2008.

[3] Cisco. Global Mobile Data Traffic Forecast Update, 2009-2014. *White Paper, CISCO Systems Inc*, 9, 2010.

[4] S. Deering and D. Cheriton. Multicast routing in datagram internetworks and extended LANs. *ACM Transactions on Computer Systems (TOCS)*, 8(2):85–110, 1990.

[5] K. P. Gummadi, R. J. Dunn, S. Saroiu, S. D. Gribble, H. M. Levy, and J. Zahorjan. Measurement, modeling, and analysis of a peer-to-peer file-sharing workload. In *SOSP*, pages 314–329. ACM, 2003.

[6] D. Horn, E. Cheslack-Postava, and T. Azim. Scaling virtual worlds with a physical metaphor. *IEEE Perv. Comp.*, 8(3):50–54, 2009.

[7] C. Huang, J. Li, and K. W. Ross. Can internet video-on-demand be profitable? In *SIGCOMM*, pages 133–144. ACM, 2007.

[8] M. May, C. Diot, P. Le Guyadec, F. Picconi, J. Roussel, and A. Soule. Service Hosting Gateways-a Platform for Distributed Service Deployment in End User Homes. In *SIGCOMM*, pages 476–477. ACM, 2011.

[9] E. Nygren, R. K. Sitaraman, and J. Sun. The Akamai network. *ACM SIGOPS Operat. Sys. Rev.*, 44(3):2, Aug. 2010.

[10] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker. A scalable content-addressable network. *ACM SIGCOMM Comput. Comm. Rev.*, 31(4):161–172, 2001.

[11] Y. Vigfusson, H. Abu-Libdeh, M. Balakrishnan, K. Birman, R. Burgess, G. Chockler, H. Li, and Y. Tock. Dr. Multicast: Rx for data center communication scalability. In *EuroSys*, pages 349–362, New York, NY, USA, 2010. ACM.

[12] Y. Vigfusson, K. Birman, D. Freedman, Q. Huang, K. V. Jonsson, and G. Sigurbjornsson. Brief announcement: Live streaming with utilities, quality and cost. In *ACM PODC*, 2012.

[13] S. Zhuang, B. Zhao, A. Joseph, R. Katz, and J. Kubiatowicz. Bayeux: An architecture for scalable and fault-tolerant wide-area data dissemination. In *NOSSDAV*, pages 11–20. ACM, 2001.