

# Packet Loss Characteristics of IPTV-like Traffic on Residential Links

Martin Ellis

Department of Computing Science  
University of Glasgow  
Email: ellis@dcs.gla.ac.uk

Colin Perkins

Department of Computing Science  
University of Glasgow  
Email: csp@cspkins.org

**Abstract**—Packet loss is one of the principal threats to quality of experience for IPTV systems. However, the packet loss characteristics of the residential access networks which carry IPTV are not widely understood. We present packet level measurements of streaming IPTV-like traffic over four residential access links, and describe the extent and nature of packet loss we encountered. We discuss the likely impact of these losses for IPTV traffic, and outline steps which can ameliorate this.

## I. INTRODUCTION

As IPTV deployment worldwide gathers momentum and numerous broadband ISPs consider providing IPTV as part of a “triple-play” package, it is becoming important to understand the behaviour of the underlying networks, and how it affects IPTV performance. It is widely appreciated that the users’ perceived Quality of Experience (QoE) is likely to be a key factor in choosing an IPTV service. As such, it is essential to understand the threats to acceptable QoE, especially those effected by the transition from traditional broadcast networks (e.g. satellite, DVB) to converged IP networks.

One such threat is network congestion, which manifests itself as queueing delays and packet loss, causing media playout disruption at the receiver. Line noise on the access link can also become visible as packet loss at the IP layer. No matter what the cause, it is vital to maintain acceptably low rates of packet loss to avoid impairments to the media quality. Moreover, it is not just the overall loss rates, but also the characteristics of the losses that are important, with different loss patterns having different implications for error control and recovery techniques. Despite the importance of understanding the effects of packet loss and jitter on IPTV, the characteristics of the residential networks which will carry it are not well known, with few measurement studies having been published.

In this paper, we present a measurement study of the packet loss characteristics of IPTV-like traffic on residential links. To the best of our knowledge, these are the first packet level measurements of streaming video-like traffic over residential networks. Our primary contribution is an analysis of the differences between ADSL and cable links, showing how receivers can compensate for their different loss characteristics to carry IPTV services with minimal packet loss disruption.

The remainder of this paper is structured as follows. Section II outlines the benefits of measuring residential networks, followed in Section III by a description of our measurement

methodology. Section IV presents the measurement results, and discusses their implications. Section V outlines related work. We conclude in Section VI.

## II. THE NEED FOR MEASUREMENTS

The design of any networked application can greatly benefit from insight into the performance of the network which carries it, and IPTV is no exception. Measurements of IPTV systems taken under varying conditions can be used to study the typical values of various performance and QoE metrics. An understanding of the expected behaviour can provide the ability to recognise deviations from the normal, laying the foundations for automatic fault diagnosis techniques; detecting infrastructure failures, above-normal bandwidth usage, or possible attacks on the network.

In the long-term, performance measurements are also useful for the planning of new networks and service deployments (e.g. capacity planning and provisioning). Moreover, the use of measurement traces to drive realistic network simulation provides an immensely valuable tool. Using this approach, network operators can investigate architectural and policy alternatives, while protocol and application designers can assess the impact of their designs on network behaviour.

In particular, anyone involved in the development and evaluation of error-repair techniques will greatly benefit from understanding the loss characteristics of the target network (and more specifically, the characteristics of the loss bursts). For example, does packet loss occur independently as single events, or are they grouped together into longer “runs” of losses? The latter often indicates that packet losses are being caused by transient congestion, while the former may indicate physical layer problems (e.g. high noise levels on the link).

## III. MEASUREMENT METHODOLOGY

Two options are available to measure characteristics of IPTV traffic on residential links. Firstly, it is possible to directly monitor an IPTV service using a passive network monitor located in a subscriber’s home. This gives accurate measurements for that particular service, but is not useful for evaluating the ability of the subscriber’s link to support other services, or other transmission rates. Alternatively, one can transmit active probe traffic across the residential link (mimicking the behaviour of IPTV flows) and monitor the reception. This

Link	Rate	Time				Duration
<i>adsl1</i> 27/06-18/07	1Mbps	Hourly at :50				1 min
	2Mbps	03:15	10:15	15:15	20:15	10 mins
	4Mbps	05:15	12:15	17:15	22:15	10 mins
	6Mbps	05:35	12:35	17:35	22:35	10 mins
<i>adsl2</i> 07/07-13/07	1Mbps	Hourly at :30				1 min
	2Mbps	04:15	11:15	16:15	21:15	10 mins
<i>cable1</i> 27/06-04/07	1Mbps	Hourly at :30				1 min
	2Mbps	04:15	11:15	16:15	21:15	10 mins
<i>cable2</i> 16/07-22/07	1Mbps	Hourly at :05				1 min
	2Mbps	04:10	11:10	16:10	21:10	5 mins
	4Mbps	04:20	11:20	16:20	21:20	5 mins
	6Mbps	04:30	11:30	16:30	21:30	5 mins
	8.5Mbps	04:40	11:40	16:40	21:40	5 mins

TABLE I: Measurement Schedule. Links *adsl2* and *cable1* have 2Mbps capacity; *adsl1* is 8Mbps; *cable2* is 10Mbps.

gives precise control over the traffic characteristics, allowing a wider range of scenarios to be explored. We adopt this latter approach, since IPTV services are rapidly evolving, to allow us to understand the range of possible behaviour.

We sent active probe traffic to four residential hosts, connected via different service providers and access link types, during June and July 2009. Links *cable1*, *cable2*, and *adsl2* were measured for one week each, while three weeks were spent measuring link *adsl1*. The schedule for the measurements is shown in Table I; rates were chosen to cover the full spectrum of IPTV traffic, from lower-rate standard definition flows (1–2Mbps) through to high-definition content at 4Mbps and above (differing access link capacities mean that not all rates were achievable on every link)<sup>1</sup>. The sender was a well-connected machine located on our university campus, generating constant bit-rate RTP [1] streams over UDP/IP. Receivers were Soekris net5501 embedded computers running FreeBSD 7.2, hosted in the homes (i.e. at the end of a residential link) of four volunteers, and running a custom monitoring application that logs headers of incoming packets for offline analysis. In total, approximately 76 million packets were logged.

End-to-end measurements were taken since this allows a range of links to be probed without involving the residential network operator. This models over-the-top services, such as the BBC’s iPlayer, but is potentially pessimistic for managed intra-domain IPTV services. However, we note that many backbone networks are over-provisioned [2], and previous studies of IPTV flows across a residential ISP’s distribution network [3] have reported excellent performance up to the last-mile, implying that the residential link is the bottleneck that determines end-to-end performance, making these results of interest to managed IPTV service operators.

#### IV. RESULTS

The key factors that affect QoE for residential IPTV flows are the bandwidth of the access links, which directly influences

<sup>1</sup>A problem with the measurement tools caused traces scheduled between 1430 on 10th July and 2230 on 11th July to fail. Additionally, an ISP-imposed rate cap of 2.5Mbps was in effect on link *cable2* on the evening of the 19th July; these traces have been discarded to avoid skewing the results.

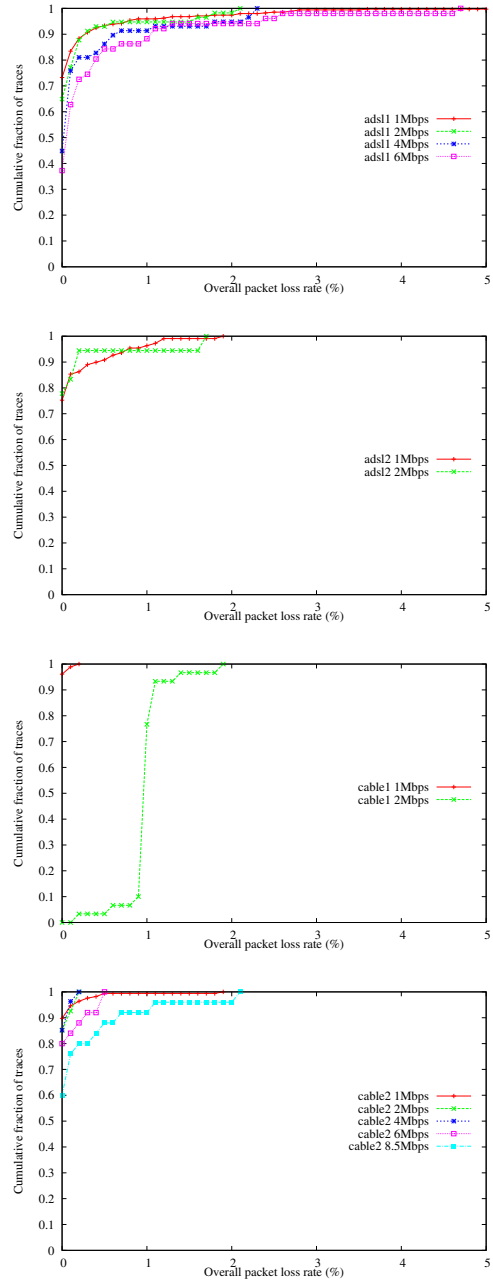


Fig. 1: Cumulative Distribution of Packet Loss Rates

the media encoding, and the packet loss and variation (jitter) in end-to-end delay. In the following, we consider packet loss and timing behaviour for different links and transmission rates.

The cumulative distribution of packet loss rates for each link at each supported rate is shown in Fig. 1. The observed loss rates are extremely low, despite running over unmanaged best effort packet networks: over 85% of all traces experienced less than 1% packet loss (except for the 2Mbps traces on link *cable1*, which we will discuss later). We note that the packet loss rate is directly correlated with the transmission rate, with higher rate flows seeing greater packet loss rates on all links (this is discussed further below).

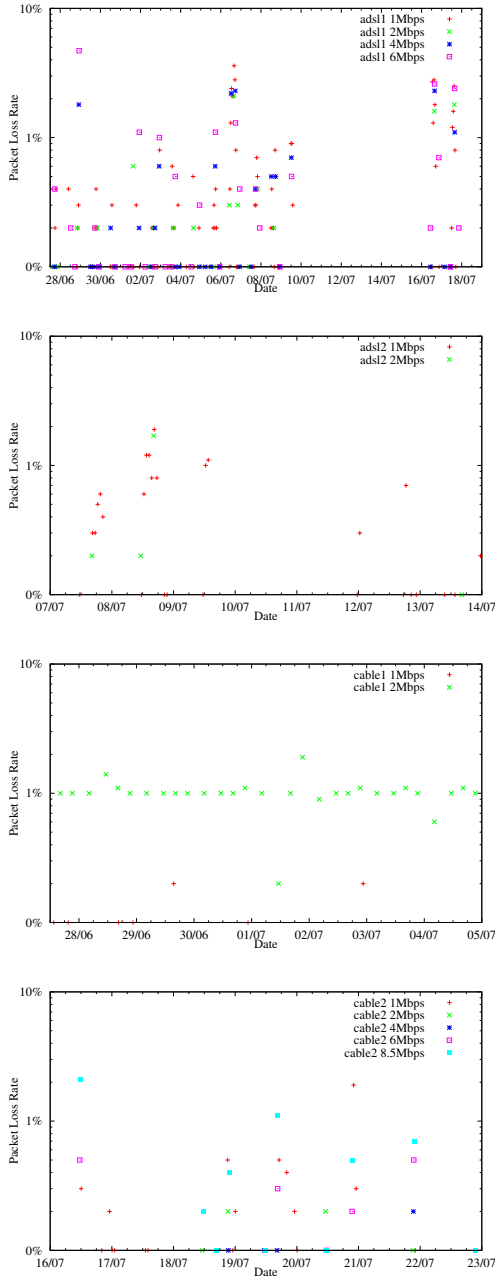


Fig. 2: Packet Loss Rates Over Time

The variation in average packet loss rate over time is shown in Fig. 2. We observe no clear time-of-day or day-of-week patterns of variation in the average loss rate (again, excepting the 2Mbps traces on link *cable1*). We would expect a network operating near capacity to show regular patterns of increased loss during peak hours, so this suggests that the networks studied are not overloaded.<sup>2</sup>

The patterns of loss visible in the 2Mbps *cable1* traces are quite different to the other traces (e.g. the consistent 1%

<sup>2</sup>Lack of space stops us from giving details, but we see temporal variation in jitter as expected. The networks are moderately busy, but not overloaded.

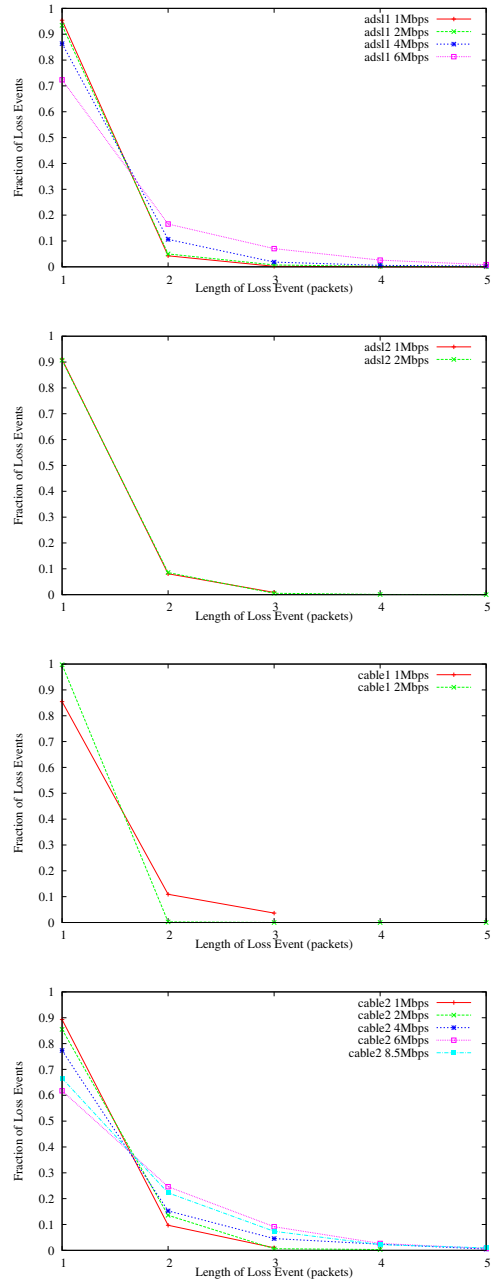
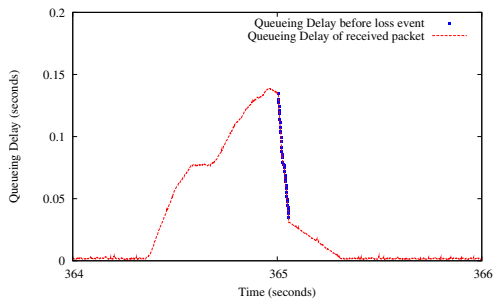


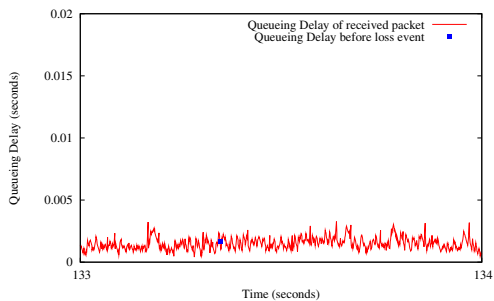
Fig. 3: Distribution of Loss Burst Lengths (packets)

loss rate seen in Fig. 2). Our hypothesis is that these losses are due to traffic shaping being employed by the ISP to enforce the customer's 2Mbps bandwidth allocation, which our measurement stream is *slightly* exceeding. The impact on the performance of our traces is obvious, and provides a strong argument for implementers to make their over-the-top IPTV services sufficiently responsive to adapt to such long-term limitations in network performance.

The probabilities of the packet loss burst durations are shown in Fig. 3. We observe that the overwhelming majority of loss events comprise a single packet loss, and that the likelihood of longer bursts drops rapidly with burst length.



(a) Congestive Loss (with drop-tail queue)



(b) Non-Congestive Loss

Fig. 4: Congestive & Non-Congestive Loss

The distribution of loss event durations is dependent on the transmission rate, with higher rate flows showing an increased number of longer bursts of loss than lower rate flows. The *adsl1* trace shows this most clearly, although the effect is visible in all the traces.

It should be noted that for different rates (and therefore different spacing between packets, for consistently sized packets), a burst of  $n$  packets represents a different length of time. Therefore, although higher rates typically suffer from longer loss bursts, this does not necessarily translate into longer periods of loss at the receiver (i.e. it affects the amount of loss to be repaired, but not necessarily the time available to perform that repair before the media playout deadline).

Detailed analysis of the traces shows that three distinct types of loss event occur. The first, shown in Fig. 4a, is characterised by a steady increase in queuing delay, followed by a rapid decrease back to the baseline delay, with numerous packet loss events occurring on the downward slope. This behaviour is entirely consistent with congestive loss due to excessive data rates in a drop-tail queue: the bottleneck router queue fills, overflows, then gradually empties as packets are discarded. Fig. 4b shows another category, where loss events appear to occur unpredictably and independently of the queuing delay, indicative of non-congestive loss caused by noise on the access link. A third category, which tends to appear in the cable traces, is less easy to explain; we hypothesise that this is due to active queue management on the cable links.

Considering the queuing delay of the packet immediately before the loss event, we observe that congestive losses are characterised by a higher delay than the average; the delay for non-congestive losses is close to the average. This allows us

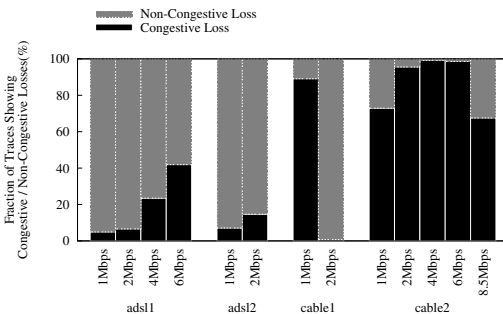


Fig. 5: Congestive vs. Non-Congestive Losses

to characterise the loss events for each link and transmission rate, as shown in Fig. 5. A clear distinction is visible between the ADSL and cable links: the ADSL links have a significantly higher fraction of non-congestive loss. Since *adsl1* and *adsl2* were users in different parts of the city, connected via different ISPs, this suggests that the difference observed between ADSL and cable is not an artefact of a single network or DSLAM, but rather a difference in the access link technology (perhaps older, poor quality, phone lines, compared with more modern, purpose built, cable TV networks). Within the ADSL traces, congestive loss occurs more frequently at higher rates, as expected. The cable links generally show much lower non-congestive loss (except *cable1* 2Mbps<sup>3</sup>), although the higher rate flows unexpectedly experience more *non*-congestive loss. We suggest that this is due to shaping using active queue management on the link to discard excessive traffic without increased queuing delay; further analysis is needed.<sup>4</sup> The implications of these results are discussed in Section VI.

Besides packet loss, the other limiting factors for IPTV flows are timing jitter and reordering. Over the course of the experiments, we observed that only 173 packets were reordered (out of over 76 million packets sent). In fact, only twelve traces showed reordering (of the 1036 logged). Since it occurs infrequently, and on a small scale, we conclude that for this type of traffic, packet timing variation causing reordering is not a significant issue.

## V. RELATED WORK

Previous IPTV measurement studies have focused on issues such as channel changing dynamics [4], system performance diagnosis and trouble-shooting [5], and video-on-demand systems [6]. To date, we believe ours is the only study to have considered packet level performance of IPTV flows across the end-to-end network path, including the residential access links. Previous work on backbone network performance with IPTV traffic [7], [3], [2] provides a useful comparison with the full end-to-end performance captured in our results, and supports our assertion that the edge network is generally the

<sup>3</sup>As earlier, this is likely due to traffic shaping to enforce the rate limit.

<sup>4</sup>The *cable2* flows at 1Mbps also show high rates of non-congestive loss. Review shows that one trace experienced extremely high rates of non-congestive loss, perhaps due to policing, while the other 165 traces had low ( $\leq 1\%$ ) rates of non-congestive loss.



bottleneck for IPTV flows. This allows us to consider the principal component of our results as edge-network effects.

Dischinger *et al.* [8] conducted a large-scale study of residential access link performance, measuring link bandwidth, latencies, and loss rates using ICMP/TCP-RST packet trains to large number of unaware hosts. To the extent that the data is comparable, due to differing methodologies, their packet loss results are generally compatible with ours, although their data shows a clearer time-of-day variation in packet loss rates (possibly reflecting differences in load between providers).

Our study uses similar techniques to conduct the measurements as earlier work on packet loss. We use a similar probing approach as Moon *et al.* [9], while our analysis of loss bursts is similar to that of Yajnik *et al.* [10].

The effect of DSL packet loss on IPTV traffic has been discussed in some detail [11], [12]; these papers also discuss error recovery strategies using FEC and retransmission. It is noted that FEC techniques (e.g. the 1-D/2-D parity FEC scheme currently under development in the IETF [13]) are most effective correcting solitary loss events (e.g. those caused by random noise), but that the presence of longer bursts of loss, or too many single packet loss events in a short period of time will prevent FEC-based recovery. In this case retransmission-based approaches may be more effective for error-recovery, allowing a “block” of lost or damaged packets to be resent from a retransmission server. Such servers will likely already be present in a managed IPTV network, to provide rapid channel change functionality.

## VI. DISCUSSION & CONCLUSIONS

We presented initial measurements of IPTV-like traffic sent to users over residential access links. With the caveat that our sample size is small, and our results therefore not necessarily representative across network operators, we are encouraged to see that the paths measured are generally well behaved, with low levels of packet loss and essentially no reordering.

The relatively low levels of packet loss observed suggest that delivering over-the-top IPTV services to residential users is both feasible, and will give acceptable media quality. Our measurements were conducted over the best-effort Internet services of residential users, with no support from their ISP, and no service differentiation measures in place. This is in contrast to an ISP-managed IPTV service, where the RTP media flows are likely to receive preferential treatment over other best-effort traffic in the network. We therefore consider that our measurements are likely underestimates of IPTV performance in a managed network, and they give us reason to believe that an appropriately engineered network can indeed provide the level of service required for IPTV.

The loss patterns we observed in these experiments imply that a single, FEC-based, approach to error recovery for IPTV traffic will likely not be sufficient to repair all errors. Non-congestive loss events, which comprised around half of all loss events on the ADSL links we measured, but only a small fraction of the loss events on the cable links, were primarily isolated single packet losses, which are readily recoverable

using parity FEC [13]. However, we observed that congestive losses generally occur in bursts after a queue overflow, with sufficient frequency to overwhelm FEC-based protection. Our results suggest that systems which detect the increase in delay due to queue build-up and reduce the bottleneck traffic rate (e.g. receiver driven layered media coding) are desirable to avoid queue overflows. If this is not possible, the combination of retransmission for congestive loss events, with FEC for non-congestive loss, seems compelling.

Our results demonstrate that an appropriately provisioned residential IP network can provide acceptable quality for IPTV without additional quality of service (QoS), provided: 1) an appropriate degree of low-level FEC is used to correct non-congestive losses in the access links; and 2) the system is responsive to, and repairs the effects of, transient congestion (note that the retransmission mechanisms used to repair congestive loss can also be used to support rapid channel change [14]). A managed, QoS-enabled, service might be able to avoid the need for congestion response, but will still be subject to non-congestive losses requiring FEC on the IPTV flows.

Ongoing work will study further links, and consider online algorithms to distinguish congestive and non-congestive loss.

## VII. ACKNOWLEDGEMENTS

Thanks to Jörg Ott, Ali Begen and Dimitrios Pezaros for their comments on early drafts of this paper, and Ross McIlroy, Stephen Strowes, and Morag Stark for hosting the measurement boxes. This work was supported by the UK EPSRC and Cisco Research.

## REFERENCES

- [1] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, “RTP: A Transport Protocol for Real-Time App.” IETF, July 2003, RFC 3550.
- [2] C. Fraleigh, S. Moon, B. Lyles, C. Cotton, M. Khan, D. Moll, R. Rockell, T. Seely, and C. Diot, “Packet-Level Traffic Measurements from the Sprint IP Backbone,” *IEEE Network*, November/December 2003.
- [3] K. Imran, M. Mellia, and M. Meo, “Measurements of Multicast Television over IP,” in *IEEE LANMAN*, 2007.
- [4] M. Cha, P. Rodriguez, J. Crowcroft, S. Moon, and X. Amatriain, “Watching Television Over an IP Network,” in *ACM IMC*, 2008.
- [5] A. Mahimkar, Z. Ge, A. Shaikh, J. Wang, J. Yates, Y. Zhang, and Q. Zhao, “Towards Automated Performance Diagnosis in a Large IPTV Network,” in *ACM SIGCOMM*, 2009.
- [6] Y. J. Won, J. W.-K. Hong, M.-J. Choi, C.-K. Hwang, and J.-H. Yoo, “Measurement of Download and Play and Streaming IPTV Traffic,” *IEEE Commun. Mag.*, 2008.
- [7] M. Cha, G. Choudhury, J. Yates, A. Shaikh, and S. Moon, “Case Study: Resilient Backbone Design for IPTV Services,” in *Proceedings of the IPTV Workshop, International World Wide Web Conference*, 2006.
- [8] M. Dischinger, K. P. Gummadi, A. Haeberlen, and S. Saroiu, “Characterizing Residential Broadband Networks,” in *ACM IMC*, 2007.
- [9] S. B. Moon, J. Kurose, P. Skelly, and D. Towsley, “Correlation of Packet Delay and Loss in the Internet,” Department of Computer Science, University of Massachusetts, Tech. Rep. 98-11, 1998.
- [10] M. Yajnik, J. Kurose, and D. Towsley, “Packet Loss Correlation in the MBone Multicast Network,” in *IEEE GLOBECOM*, 1996.
- [11] A. C. Begen, “Error Control for IPTV over xDSL Networks,” in *IEEE CCNC*, 2008.
- [12] N. Degrande, D. D. Vleeschauwer, and K. Laevens, “Protecting IPTV Against Packet Loss: Techniques & Trade-offs,” *Bell Labs Tech. J.*, 2008.
- [13] A. Begen, “RTP Payload Format for Non-Interleaved and Interleaved Parity FEC,” IETF, May 2009, Work in Progress.
- [14] B. VerSteeg, A. Begen, T. VanCaenegem, and Z. Vax, “Unicast-Based Rapid Acquisition of Multicast RTP Sessions,” IETF, June 2009, Work in Progress.