
On the Use of RTP for Monitoring and Fault Isolation in IPTV

Ali C. Begen, Cisco
Colin Perkins, University of Glasgow
Jörg Ott, Aalto University

Abstract

Since the first operational IPTV networks have been deployed, service providers and operators have struggled to make their subscribers happy and satisfied with their services. To keep them as customers in the long term, they have been looking for ways to identify impairments to the perceived quality of experience. It is now well understood that this can only be achieved if the service providers have virtual eyes throughout their networks. In this article we provide an overview of the Real-Time Transport Protocol and its application to IPTV. We describe the monitoring and reporting features offered by RTP, and emphasize how they can be used to enhance subscriber QoE.

Packet-based multimedia communication has come a long way since its early exploration in the 1970s (for audio over Ethernet) and the 1980s (for audiovisual conferencing over IP), and is nowadays a commodity for millions of users. The most prominent example for the success of IP-based multimedia is IP telephony, where commercial development started in the 1990s, and which has led to the conversion — although operators usually prefer to talk of convergence — of many time-division multiplexing (TDM)-based telecommunication networks into IP networks, from the core to the enterprise and increasingly to the home. With increasing availability of broadband access, the entertainment sector has also started to embrace the idea of running everything over IP, leading to the notion of *triple play* (Internet access, telephony, and IP television [IPTV]) offered in an integrated fashion by Internet service providers (ISPs). Non-ISP-based offerings have followed the trend with a multitude of web-based multimedia streaming offerings including Internet radio stations, TV broadcasters delivering selected programming, and user-generated content platforms such as YouTube.

While the types of services, and the control protocols used to access and interact with them, may differ, the common theme across all services is the delivery of real-time (multi)media contents over a packet-switched network. The Real-Time Transport Protocol (RTP), Internet Engineering Task Force (IETF) Request for Comments (RFC) 3550,¹ was developed to serve this purpose, drawing on the experience of the early experimentation. While RTP's original target was multiparty multimedia conferencing in multicast networks, its design was carefully crafted for much broader applicability — RTP accommodates point-to-point IP telephony on one end of the scale, and broadcasting to millions of users on the

other. The broad acceptance and deployment of IP telephony solutions using RTP prove the former. In this article we make the case for the latter.

After an overview of the IPTV landscape, we review the basics of RTP, introduce the technical concepts of RTP specifically applied to IPTV, and present a new architecture for robust IPTV distribution. We then apply the reporting mechanisms offered by RTP to realize monitoring and fault tolerance, and subsequently show how service providers can use these mechanisms for troubleshooting and increasing customer satisfaction. We conclude with a brief summary and point to future developments. Our contribution is in outlining a multicast RTP-based IPTV distribution architecture that integrates *scalable* RTP Control Protocol (RTCP) feedback, feedback aggregation and fault isolation, rapid acquisition, and media repair to enhance subscriber quality of experience (QoE).

IPTV in the Real World

IPTV is the delivery of entertainment-quality video over an IP network. However, it is not just about how video content is carried in the network and delivered to consumers. The vision behind IPTV encompasses more than this: it extends the reachability of content to any IP-connected device and makes it available any time to consumers, something they cannot get from traditional off-the-air, satellite, or cable services.

Digital video services have historically been classified into two categories, broadcast and on demand. In broadcast services the model is one-to-many, and a single copy of a program is sent to all subscribers regardless of whether they are watching or not. This is called digital broadcast. In contrast, if a single copy of a program is sent to only the subset of the subscribers interested in it, this is called switched digital video (SDV). SDV services require two-way communication and multicast support in the distribution and access networks. In

¹ All RFC documents are freely accessible from <http://www.ietf.org/rfc.html>

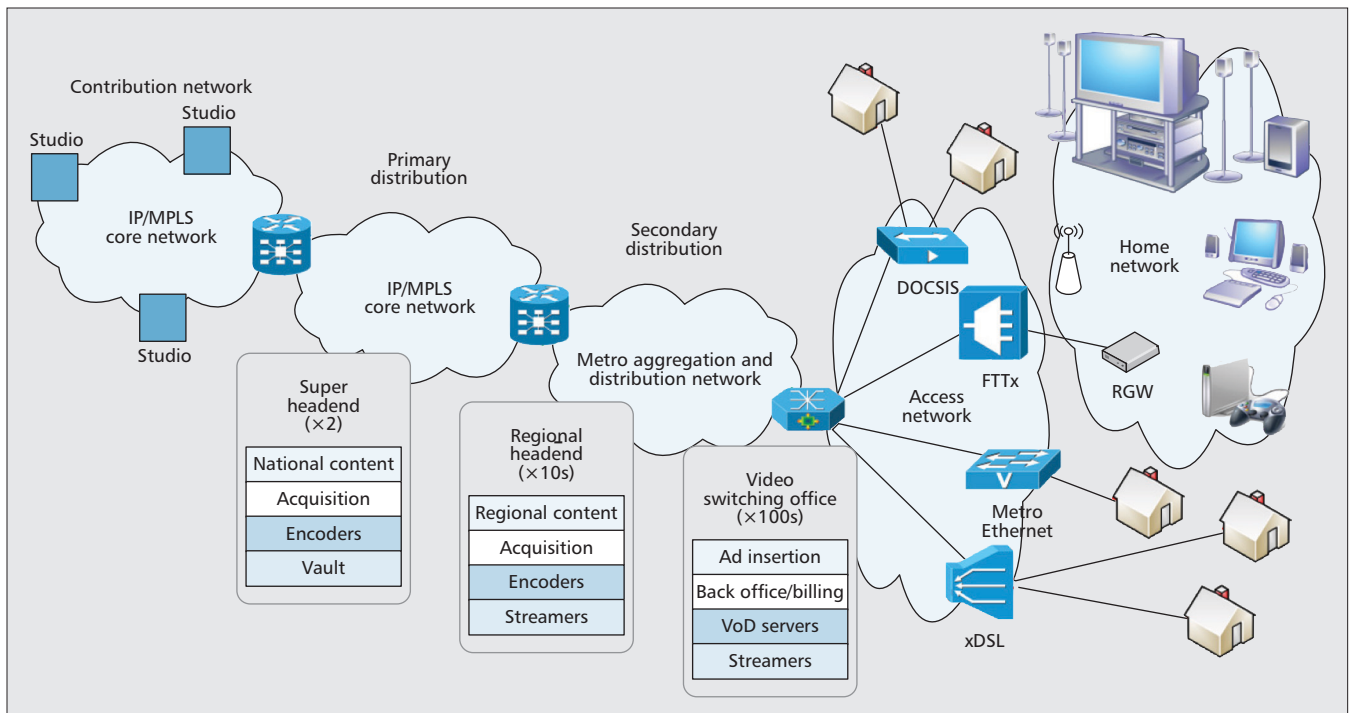


Figure 1. IPTV distribution.

on-demand services, where the model is one to one, a single copy of a program is sent to a single subscriber who controls the streaming interactively.

Historically, off-the-air and satellite providers have been able to provide only digital broadcast services, although more recently satellite providers are mimicking on-demand services using residential broadband connections. On the other hand, cable providers have been a big promoter for on-demand services and are starting to support SDV at larger scales to satisfy the increasing bandwidth demand for on-demand services.

In addition to these managed digital broadcast and on-demand services, a new service, unmanaged over-the-top (OTT) video, is emerging. In this service content owners, resellers, and distributors use consumers' broadband connections to deliver content on demand. The content may be sourced from one or more servers, as part of a content delivery network (CDN), or from other consumers via peer-to-peer connections. However, any abnormalities in the consumer's connection reveal themselves in this service as it runs on top of the best effort Internet without any special treatment from ISPs. Thus, packet loss, delay, and jitter must be dealt with to keep the QoE at a satisfactory level.

Unlike OTT services, IPTV is a managed service where the provider guarantees to meet certain QoE constraints. From a distribution viewpoint, it is similar to an SDV service, and uses IP multicast for efficient distribution. The subscriber's IP set-top box (STB) joins the respective multicast session(s) to receive the stream(s) only for the channel in which the subscriber is interested. For on-demand services, the IP STB establishes a unicast session with a content server inside the provider's network.

Content Acquisition and Distribution

Before content can be distributed, it must be captured and acquired. Production studios are often responsible for this service and carry the contribution video in uncompressed form (or using lossless compression), requiring very high-speed links. Following contribution, content is delivered to service providers, typically using satellites, fiber, or off the air after lightweight compression. This is called primary distribution. In

the service provider network, super headends (SHEs) ingest the national-level content and deliver it to video hub offices (VHOs) (regional headends) over a core IP/multiprotocol label switching (MPLS) network. VHOs aggregate national, regional, and local content with on-demand services, and are connected to video switching offices (VSOs) over metro aggregation networks. VSOs distribute content to the customer premises over access networks — cable (e.g., hybrid fiber coax [HFC]), metro Ethernet, fiber, and digital subscriber line (DSL) technologies. This is called secondary distribution. Since access networks have lower link capacity, further compression is applied, resulting in bit rates of 2–4 Mb/s for standard-definition and 6–20 Mb/s for high-definition content. Figure 1 shows a typical distribution network. While IPTV is currently being rolled out predominantly by wireline providers, it is naturally agnostic to the layers below IP. We expect to see IPTV deployments from fiber and cable providers in the near future.

Due to the nature of the IPTV streams in core and aggregation networks, the on-time delivery performance for IPTV packets in these networks must be almost perfect. A number of technologies are available to service providers to engineer networks that meet the required service-level agreements, including differentiated services (RFC 2475) and integrated services (RFC 1633). However, packet loss can still occur due to network reconvergence events (link or node failures or recovery) or lower-layer errors. Service providers, depending on the nature of their networks, may address these issues through various technologies such as fast IP convergence, MPLS traffic engineering fast reroute, and multicast only fast reroute, or loss recovery approaches such as spatial or temporal redundancy, forward error correction (FEC), and retransmission [1, 2].

IPTV service providers strive to deliver every packet intact and on time to subscribers; however, they often have almost no control over the networks their subscribers run in their homes. Some might be running wired, wireless, or both, while some others might use power-line connections. There are several complications that may occur due to poor wiring, interference from external noise sources or other nearby transceivers,

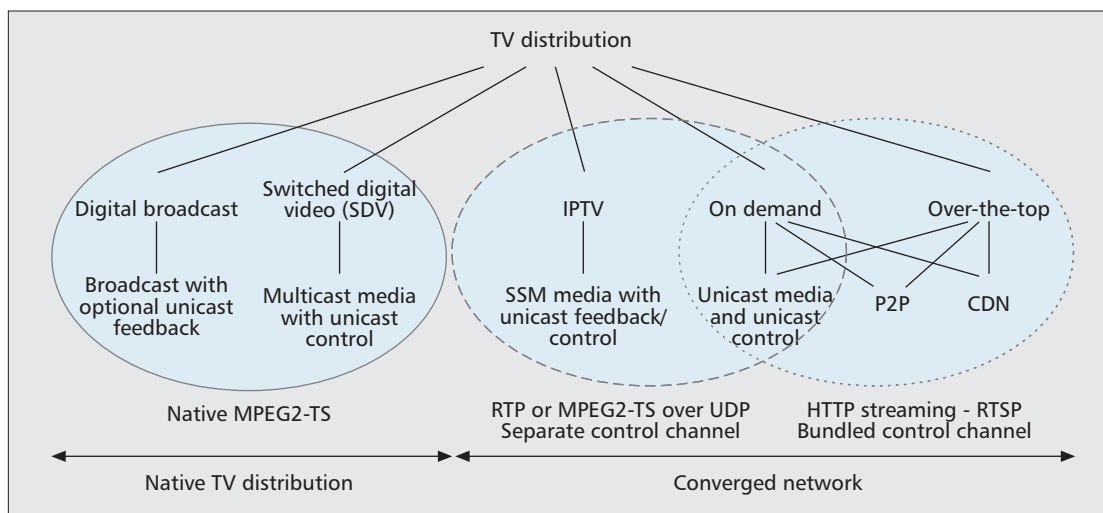


Figure 2. Taxonomy of TV distribution architectures and protocols.

and devices themselves. Often, such complications will hamper the video transmission, significantly degrading QoE. In this relatively unmanaged environment it is challenging for the service provider to identify and pinpoint problems. Yet, as we discuss later, with well engineered tools, the service provider may gain visibility into subscriber premises to diagnose such problems.

IPTV Delivery Architecture and RTP

The distribution models described previously can be categorized as shown in Fig. 2. Considering converged networks, using IP as a substrate, OTT and on-demand services have tended to adopt an architecture derived from Web protocols, streaming either directly over HTTP or via the Real-Time Streaming Protocol (RFC 2326), running media and control over a TCP connection to ease Network Address Translation (NAT) traversal. In contrast, IPTV systems — which are our focus — have tended to adopt multicast UDP-based media transport, with a separate unicast control channel, for scalability, with direct involvement from ISPs in service management.

The high-level architecture of an IPTV system is designed to support large numbers of receivers. It comprises a source-specific multicast (SSM) session for each channel, with receivers joining multicast sessions for the channels in which they are interested. A separate unicast signaling protocol — carefully designed for scalability — conveys reception quality reports and provides an electronic program guide, interactive services, and other feedback. The use of IP multicast rather than an application-layer overlay (whether peer-to-peer or a content distribution network) is desirable to ensure low-latency routing that can be integrated with network quality of service (QoS) mechanisms and managed by the network operator, and to provide resilience to receiver churn.²

The naive choice for media transport in such a network is to map the MPEG2 Transport Stream (MPEG2-TS) containing the content directly onto a sequence of UDP datagrams for delivery via the multicast session. This is a relatively small conceptual change from traditional distribution models, which reduces the barrier to entry and has proven attractive to many operators. Especially, traditional video dissemination channels and well-provisioned QoS-controlled IP networks appear

quite similar from a distance. However, we argue that MPEG2-TS over UDP provides extremely suboptimal performance, since MPEG2-TS was not designed or optimized for IP networks. Significant benefits can be achieved by adopting a transport, such as RTP, which was designed for use in an IP environment.

The RTP Framework

The RTP framework is an Internet standard published by the IETF in July 2003. It has seen global deployment in the voice-over-IP (VoIP) and telephony markets as an essential component of the Third Generation Partnership Project IP multimedia subsystem (3GPP IMS) and many enterprise VoIP systems, but there is nothing in the protocol that is VoIP-specific. Indeed, RTP evolved from multicast conferencing experiments in the early 1990s [3] and was explicitly designed for multimedia applications where scalable *group communication* was a requirement. As a result, it is well suited to IPTV.

RTP comprises a data transfer protocol and an associated RTP Control Protocol (RTCP), augmented by various extensions, profiles, and media payload formats. The data transfer protocol identifies the media synchronization source (SSRC) and any contributing sources, and the type of payload data, and provides sequence numbers and timestamps to allow receivers to detect packet loss and recover media timing. RTCP conveys source description data, reception quality reports, codec control back-channel messages, and timing information to provide interstream synchronization.

Media data are encapsulated in RTP using *payload formats*. These specify how a (continuous) media stream is broken down into discrete units and how these are spread across packets. The packetization should ensure that, since IP packets may be lost, the generated packets are independently decodable; that is, packet N does not depend on the reception of packet $N - 1$, or that any such dependencies are identified and minimized through appropriate framing. Depending on the codec used, an additional payload-specific header may be included in RTP packets for this purpose, for example, to repeat essential metadata in every packet for robustness (several standard payload formats include FEC within the media stream; and the IETF FECFRAME Working Group is currently developing additional FEC schemes). Furthermore, the packetization shall avoid IP-layer fragmentation to minimize the risk and impact of loss. RTP derives much performance from this intelligent use of application-level framing.

There are mappings to run RTP over UDP and TCP. For

² Peer-to-peer (P2P) IPTV distribution is a rapidly evolving area with few standards, and fewer well-known best practices. A comprehensive review of this area is outside the scope of this article.

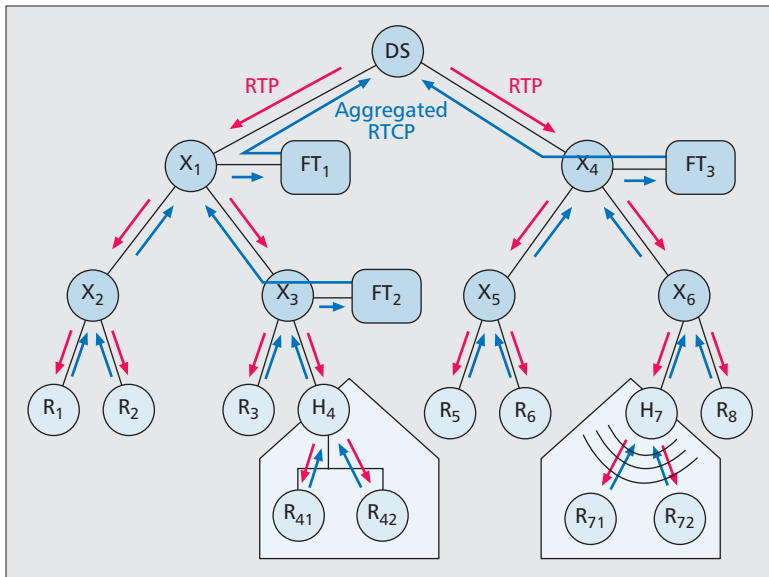


Figure 3. Network topology view of the RTP-based IPTV distribution architecture.

IPTV deployments, RTP over UDP is appropriate, since it allows the use of IP multicast. An RTP session is usually created for each type of media, and each session runs on its own UDP port (i.e., audio and video are sent as separate sessions, on different ports, with individually tailored FEC, synchronized using RTCP).³ RTCP traffic generally runs on a separate port from the media, although a recent extension (RFC 5761) allows them to be multiplexed on the same port.

By default, RTP assumes a back channel that distributes RTCP packets to all session participants. An extension to support SSM sessions was recently published (RFC 5760), allowing RTCP reports to be unicast to a feedback target, which takes responsibility for redistributing them to the other receivers (either in full or using a summary report). The feedback target is typically collocated with the distribution source, and thus can send to the SSM group used to distribute the media data; for large-scale systems such as IPTV, hierarchies of feedback targets can be deployed. Strong authentication and signaling extensions are used to identify the feedback target and avoid packet-bombing attacks.

RTCP Reception Quality Feedback

The basic RTCP Reception Report (RR) conveys transport-level statistics, the fraction of packets lost since the last report, the cumulative number of packets lost, the highest sequence number received, an estimate of the interarrival jitter, and information to allow senders to calculate the network round-trip time. The eXtended Report (XR) (RFC 3611) provides more detailed feedback in several categories:

- **Loss and duplicate RLE reports** use run length encoding to allow receivers to report on loss or duplication of individual data packets, providing the sender with a complete view of which packets arrived.
- **Packet-receipt times reports** allow the receiver to report reception times of individual packets, giving a more detailed view of packet timing than standard jitter reports.

³ A halfway house between native MPEG2-TS and native RTP delivery is to run a complete MPEG2-TS over RTP (RFC 2250). This gets some of the benefits of RTP (e.g., RTCP monitoring), but is poorly matched to the overall RTP framework since it bundles audio, video, and control into a single flow, and cannot adapt the framing to match the network characteristics.

- **Round-trip time reports** contain information to allow receivers to calculate the round-trip time to a sender.
- **Statistics summary reports** allow receivers to summarize the number of packets lost or duplicated, and the minimum, maximum, mean, and standard deviation of the jitter and IP time-to-live (TTL) of received packets.

The detailed loss, duplicate, and receipt times reports were developed to allow network monitoring and fault isolation using *network tomography* based on reception quality reports for multicast RTP streams [4, 5]. Further extensions are under development to convey application-level metrics.

The time between RTCP packets — the reporting interval — is varied to limit control traffic to a small fraction (configurable, typically 5 percent) of the media bandwidth, and randomized ± 50 percent to avoid report synchronization. Average reporting intervals are on the order of several seconds for small groups, scaling to several minutes as group size increases. The reporting interval adaptation algorithms have been shown to scale to

sessions with tens of thousands of participants, and to cope well with flash crowds [6].

There are, however, various scenarios in which it is desirable to send RTCP reports more frequently. The Audio/Video Profile with Feedback (RTP/AVPF) (RFC 4585) extends the standard A/V Profile (RFC 3551) to allow RTCP reports to be sent early provided the overall RTCP bandwidth allocation is respected. There are three uses for this:

- **Retransmission** of lost RTP data packets (RFC 4588), driven by negative acknowledgments sent in early RTCP reports
- **Rapid acquisition of multicast sessions** [7] using RTCP to control rapid unicast delivery of reference information needed to synchronize with an ongoing multicast RTP session
- **Codec control** messages (RFC 5104) to convey requests for full intra-coded frames or select the reference picture, and signal changes in the desired temporal/spatial trade-off and maximum media bit rate

We expect IPTV systems to make heavy use of the rapid acquisition and retransmission extensions to support rapid channel change [8].

Use of RTP for IPTV

The IPTV system architecture we propose, which is under consideration in the Internet standards community and is beginning to see some commercial deployment, combines the aforementioned RTP mechanisms to create a (usually hierarchically structured) dissemination and reporting system as depicted in Fig. 3. The distribution source (DS) originates the IPTV media stream that is distributed via SSM. Receivers (R) obtain multicast channel information (multicast address, source address, port number, local or global feedback target [FT]) via out-of-band means (e.g., an electronic program guide using the Session Description Protocol, RFC 4566) and join the channel. They receive RTP and RTCP packets from the DS and send RTCP feedback to their (local) FT. Feedback targets aggregate the received information and forward summaries upstream; summary information is sent to receivers by either the DS or their local FTs. Standard IP routers (X) deliver SSM traffic downstream and unicast feedback upstream; the home gateway (H) is a standard ADSL/cable router, augmented with SSM support and monitoring functions described below.

Receivers use regular RRs and selected XRs to provide

detailed statistics. Feedback targets may also join the multicast session, receive the media stream, and provide feedback on their own reception statistics, providing measurement probes inside the infrastructure. The RTP session uses the Extended RTP profile for RTCP-based feedback (RFC 4585), allowing receivers to provide instant feedback in case of missing packets due to packet losses or channel-change operations. In either case, the FT system buffers RTP media packets that are later used to repair missing packets. In case of a channel change, not just a single packet is sent, but an RTP burst is generated and pushed to the receiver faster than real time to prime the receiver and allow it to begin playback virtually immediately [7, 8].

The combination of RTP, its existing standard extensions, and the extensions we have proposed for SSM with unicast feedback, rapid acquisition of multicast sessions, and multiplexed data and control provides a powerful and scalable architecture for robust IPTV delivery in ISP environments [4]. In the following we describe how it can be leveraged to enhance QoE and reduce subscriber churn.

Reducing Subscriber Churn

In today's competitive world, consumers are often offered similar entertainment packages from different service (satellite, cable, or IPTV) providers at similar pricing levels. However, consumers are not concerned about the specifics of *how* TV programming is delivered, but rather about the service quality, or more specifically, the QoE, which is by far the primary criterion in their provider selection.

The subscribers' QoE depends on various factors. The video encoding quality, which is largely a function of the codec and bit rate, is clearly essential; as is the network QoS, which is heavily dependent on error-free and on-time delivery of the packets. Various application and user-level metrics, both objective and subjective, must also be considered. For example, channel changing in digital broadcast or IPTV is slower than in analog broadcast and can have a large impact on subscriber satisfaction; similarly, depending on whether the trick modes in a VoD session are responsive enough, subscribers' perception of the service quality may vary. Measuring and quantifying subscribers' QoE is more difficult than measuring network or media QoS, since it relates to human perception, but is *critical* for IPTV service providers who are constantly working to reduce subscriber churn and gain customers. In an environment with fierce competition, the ultimate ammunition for a service provider is to offer a better QoE than its rivals. To do this, the provider has two basic requirements: figuring out what and where the problem is, and reporting the impact of the problem and the effectiveness of its solution. In other words, the service provider must have eyes throughout its network, from source to displays, to derive a proxy measure for subscriber QoE from concrete performance metrics.

Of course, stationing eyes in various locations, monitoring high-bit-rate video streams, and processing large amounts of traffic do not come for free. Especially if specialized equipment, probes, and analyzers are desired, the initial cost of deployment may increase. However, where RTP is already supported in the network — because of its improved performance over plain UDP transport — we argue that the majority of the deployment cost can be avoided since the provider will not need most, or perhaps any, of the specialized measurement devices that would have been needed to monitor UDP-based streams. Thus, with RTP and RTCP support in the network, capital expenditures will be reduced, while manageability will improve since RTCP provides metrics that directly relate to media flows.

Enhancing QoE Using RTCP Feedback

In traditional quality monitoring checkpoints, humans inspect each TV channel to make sure the service is operating normally in every headend location. When a problem occurs, the origin of the problem can usually be isolated and identified after a series of manual checks. During this time, service disruption may occur for subscribers. This will result in an increased number of calls to technical support and increased cost to the service provider until the problem is fixed. Automation in traffic monitoring and data processing is strongly desirable for IPTV, and by leveraging RTCP, we assert that vendors can offer automation tools to handle the monitoring and processing of large amounts of traffic in a scalable manner, and gather data from endpoints as well as network elements.

RTCP offers periodic feedback for reception statistics, yielding coarse loss and delay/jitter statistics from plain receiver reports and selected details using extended reports. Immediate feedback containing loss reports allows for (local) repair to remedy occasional packet losses and improve user-experienced quality. However, this will not help against longer-term problems (e.g., due to component failures or overall changes in the load situation inside a network region leading to insufficient provisioning). This is where permanent real-time monitoring of RTCP can help service providers locate and (manually) fix their infrastructure or routing.

An entity — an IPTV sender or operator's network element — monitoring RTCP reports will gain immediate insights into the characteristics of individual communication paths. Since many receivers are usually connected to a single access concentrator (e.g., a DSL access multiplexer [DSLAM]), statistics about the path to this DSLAM come from multiple sources, allowing improper reports (e.g., due to malfunctioning software) to be weeded out. Reports from multiple paths can be combined to obtain insights into the health of a certain region of the network. With reference to Fig. 3, for example, trends in packet loss reports point to certain network elements, paths, or regions (imminent to) experiencing congestion: reported higher loss rates from R_1 and R_2 indicate problems somewhere below X_1 toward X_2 (or at X_2) if reports from R_3 and H_4 are fine. Stream monitoring in a home gateway allows differentiating between problems inside the provider network and home installations, especially when users watch IPTV over WLANs or use improper networking equipment: if extensive jitter and losses are reported by R_{71} or R_{72} , but not by H_7 , the operator knows the problem is at the user site, and can react accordingly.

Furthermore, individual path segments may be monitored by comparing reports by entities at different *distances* along the paths; obviously, the topological resolution improves as more monitoring probes get added. Referring again to Fig. 3, reports from FT_1 and FT_2 in addition to the receivers enable pinpointing the origin of loss or jitter to an individual segment.

Sophisticated analysis is possible using network tomography to compare the loss/delay patterns seen across different endpoints to isolate fault(s). For example, one can infer, with high probability, internal loss rates [9] and sources of delay variation [10] in a multicast group given the knowledge of the network topology using the detailed reception quality feedback provided by RTCP XR. This is often sufficient to identify the particular network element or link at fault with high confidence. Additionally, it appears that it may be possible to use these same tomography techniques to infer the location of faults in layer-2 networks; for example, to locate faults in ADSL distribution networks that are invisible to the IP layer

by correlating performance across multiple subscribers connected to the same DSLAM.

Rapid fault diagnosis using network tomography requires the network operator to share details of their network topology with the IPTV service provider. In environments where this is not acceptable, in-network RTCP monitoring enhances the ability of a network operator to pinpoint fault locations. For example, certain network elements can join the multicast session as receivers and process the stream as a normal RTP receiver. In some scenarios these elements can benefit from specialized analysis hardware and software for detailed application-specific quality assessment. This approach can provide great visibility in the network; however, it is usually very costly to deploy, and care should be taken to avoid alterations in the multicast distribution trees. In other scenarios the monitoring elements can capture the ongoing traffic and send it to remote locations that centrally perform sophisticated analysis [11]. This helps reduce deployment costs at the expense of increased network traffic. Basic statistical measurements are also possible in network elements that support flow monitoring, as discussed in RFC 4445.

In addition to sending RTCP feedback, RTP-capable network elements at critical locations can be temporarily or permanently configured to trap and clone packets of a particular RTP stream, encapsulating them in a new RTP stream with sequence numbers and timestamps generated according to which packets are received when at this element, and send the new RTP stream to a central facility where further analysis can be carried out. If only network-level measurements are desired, the RTP payloads can be discarded in the new RTP stream while sequence numbering and timing information is retained. Compared to the RTCP reports collected from multiple receivers along the media distribution path, this method allows higher-resolution timing analysis and application-level monitoring.

Service providers can choose a combination of the above mechanisms that best matches their operational demands; they can incrementally deploy different mechanisms, such as starting with end-to-end monitoring and adding network-based probes or more sophisticated report processing as they see fit, to arrive at a solution that precisely matches their needs. Such a solution can evolve over time as new RTP-based services are deployed on the converged networks of the future, forming an important technical part of a holistic solution for managing subscriber QoE.

Concluding Remarks

Appropriate use of RTP with RTCP feedback as outlined in this article allows service providers to rapidly identify and isolate problems. Intelligent monitoring agents can even identify potential problems in a proactive manner and take the necessary steps to avoid them in the first place. Compared to

human inspections, these tools greatly reduce the operational expenditures of service providers. With proper embedding in the operational and business environment, we expect that, over time, providers will see improved customer QoE and reduced subscriber churn, and subsequently a high return on their investment.

References

- [1] C. Filsfil and J. Evans, "Deploying DiffServ in Backbone Networks for Tight SLA Control," *IEEE Internet Comp.*, vol. 9, Jan. 2005, pp. 66–74.
- [2] R. Doverspike *et al.*, "Designing a Reliable IPTV Network," *IEEE Internet Comp.*, vol. 13, May 2009, pp. 15–22.
- [3] S. Casner and S. Deering, "First IETF Internet Audiocast," *ACM SIGCOMM Comp. Commun. Rev.*, vol. 22, no. 3, July 1992, pp. 92–97.
- [4] A. C. Begen, C. Perkins, and J. Ott, "On the Scalability of RTCP-Based Network Tomography for IPTV Services," *IEEE CCNC*, Las Vegas, NV, 2010.
- [5] R. Cáceres, N. Duffield, and T. Friedman, "Impromptu Measurement Infrastructures Using RTP," *IEEE INFOCOM*, New York, NY, 2002.
- [6] H. Schulzrinne and J. Rosenberg, "Timer Reconsideration for Enhanced RTP Scalability," *IEEE INFOCOM*, San Francisco, CA, 1998.
- [7] IETF Internet draft, "Unicast-Based Rapid Acquisition of Multicast RTP Sessions," Feb. 2010; <http://tools.ietf.org/html/draft-ietf-avt-rapid-acquisition-for-rtp>.
- [8] A. C. Begen, N. Glazebrook, and W. Ver Steeg, "Reducing Channel-Change Times in IPTV with Real-Time Transport Protocol," *IEEE Internet Comp.*, vol. 13, May 2009, pp. 40–47.
- [9] R. Cáceres *et al.*, "Multicast-Based Inference of Network-Internal Loss Characteristics," *IEEE Trans. Info. Theory*, vol. 47, Nov. 1999, pp. 2462–80.
- [10] N. G. Duffield and F. Lo Presti, "Multicast Inference of Packet Delay Variance at Interior Network Links," *IEEE INFOCOM*, Tel Aviv, Israel, 2000.
- [11] A. Mahimkar *et al.*, "Towards Automated Performance Diagnosis in a Large IPTV Network," *ACM SIGCOMM*, Barcelona, Spain, 2009.

Biographies

ALI C. BEGEN [M] (abegen@cisco.com) is with the Video and Content Platforms Research and Advanced Development Group at Cisco. His interests include networked entertainment, Internet multimedia, transport protocols, and content distribution. He is currently working on architectures for next-generation video transport and distribution over IP networks, and is an active contributor in the IETF in these areas. He has a Ph.D. in electrical and computer engineering from the Georgia Institute of Technology. He is a member of the ACM.

COLIN PERKINS [SM] (csp@cspcrkins.org) is a lecturer in computing science at the University of Glasgow. His research focuses on the design of robust, scalable, adaptive distributed multimedia systems for use on IP networks, including congestion control, signaling protocols, and the use of high-level languages in systems design and implementation. He is a member of the IETF Transport Area Directorate, and has chaired the IETF Audio/Video Transport and the Multiparty Multimedia Session Control Working Groups, and contributed to MPEG. He has a D.Phil. in electronics from the University of York, and is a member of the ACM and IET.

JÖRG OTT [M] (jo@netlab.tkk.fi) is a professor in the School of Science and Technology at Aalto University. His research interests are in Internet technologies, protocol design, and protocol and system architectures for present and future networks. His research focuses on robust and adaptive interactive (multimedia) communications and content distribution, as well as on communications in (challenged) wireless networks, particularly delay-tolerant networking. In the IETF he has been co-chair of the Multiparty Multimedia Session Control Working Group since 1997 and co-chaired the SIP Working Group from its foundation in 1999 to October 2002. He has a Dr.-Ing. in computer science from TU Berlin and is a member of the ACM.