Quantifying Video-QoE Degradations of Internet Links

Mukundan Venkataraman, *Member, IEEE*, and Mainak Chatterjee

***Abstract*—With the proliferation of multimedia content on the Internet, there is an increasing demand for video streams with high *perceptual* quality. The capability of present-day Internet links in delivering high-perceptual-quality streaming services, however, is not completely understood. Link-level degradations caused by intradomain routing policies and inter-ISP peering policies are hard to obtain, as Internet service providers often consider such information proprietary. Understanding link-level degradations will enable us in designing future protocols, policies, and architectures to meet the rising multimedia demands. This paper presents a trace-driven study to understand quality-of-experience (QoE) capabilities of present-day Internet links using 51 diverse ISPs with a major presence in the US, Europe, and Asia–Pacific. We study their links from 38 vantage points in the Internet using both passive tracing and active probing for six days. We provide the first measurements of link-level degradations and case studies of intra-ISP and inter-ISP peering links from a multimedia standpoint. Our study offers surprising insights into intradomain traffic engineering, peering link loading, BGP, and the inefficiencies of using autonomous system (AS)-path lengths as a routing metric. Though our results indicate that Internet routing policies are not optimized for delivering high-perceptual-quality streaming services, we argue that alternative strategies such as overlay networks can help meet QoE demands over the Internet. Streaming services apart, our Internet measurement results can be used as an input to a variety of research problems.**

***Index Terms*—Digital video broadcasting, Internet, quality of experience (QoE).**

I. Introduction

An ESTIMATED $4.2 billion from revenue generation is predicted for Internet video by 2011, with an annual growth rate of 36% [6], [10]. About 30% of the Internet traffic already carries multimedia content. As customers spend more and more time watching videos online, they are increasingly becoming unsatisfied by low-bit-rate videos available on YouTube and are embracing high-definition (HD) and switched-digital (SD) streaming services. To attract customers and prevent subscriber churn, network service providers need to deliver streaming content with high *perceptual* quality, or quality of experience (QoE). QoE differs from traditionally used objective quality assessment models like quality of service (QoS) in trying to assess quality from a user's point of view. Research has consistently shown that QoS models do not correlate very well with subjective perception [14]. Providing streaming services with high QoE over a best effort and shared infrastructure such as the Internet, however, is nontrivial.

The Internet is organized as an interconnection of independent autonomous systems (ASs). Each AS is under the purview of an Internet service provider (ISP), and ASs peer with each other to cooperatively forward packets. Routing in the Internet is a process of finding a series of paths traversing one or multiple ASs to reach a destination. Intradomain routing policies, ISP-peering policies, as well as delay and jitter distributions of Internet links are hard to obtain because ISPs often consider such information proprietary. As a result, the quality of links both inside an AS and peering links used to exchange traffic between ASs are largely unknown. Since the quality of a video stream is as good as the quality offered by the worst link along its path, understanding link-level degradations will help characterize the extent to which various factors in the Internet affect video-QoE. This will help us gain an insight into designing future protocols, policies, and supporting architectures to meet the rising multimedia demands. Video-QoE is known to be affected by three key network events: loss, delay, and jitter. While path-inflation and loss characteristics have been studied in the past, no prior work has investigated the jitter levels of the Internet or the combined effect of various factors on perceptual video quality at the link level.

In this paper, we systematically study both *intra-* and *inter*-ISP links that collaborate to perform present-day Internet routing, their respective video-QoE capabilities, and ways to improve them. Our study involves 51 ISPs with a dominant presence in either the US, Europe, or Asia–Pacific. We start by tracing all globally prefixed IP addresses for six days from 38 PlanetLab [34] vantage points in the Internet to extract ISP topologies. IP level paths obtained from the trace are converted to AS-level paths, which reveal a rich collection on intra-ISP and inter-ISP peering links. We actively probe these links from vantage points close to these links to measure their response times and relative loading. We present 24-h case studies of both an intra-ISP link run by Level3 and a peering link between Sprint and Qwest that are representative of a large fraction of the discovered links. Raw network statistics are mapped to QoE capabilities of these Internet links using an objective function. Finally, we study the combined interaction of traversing multiple links by analyzing an unoptimized playout buffer of an end-to-end transmission of 150 000 packets between the University of California, Los Angeles (UCLA) and Carnegie

Manuscript received November 10, 2010; revised May 12, 2011; accepted June 17, 2011; approved by IEEE/ACM Transactions on Networking Editor C.-N. Chuah. Date of publication October 21, 2011; date of current version April 12, 2012. A preliminary version of this paper appeared at the IEEE Global Communications Conference (GLOBECOM), Miami, FL, December 6–10, 2010.

The authors are with the School of Electrical Engineering and Computer Science, University of Central Florida, Orlando, FL 32816 USA (e-mail: mukundan@eecs.ucf.edu; mainak@eecs.ucf.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TNET.2011.2167684

Mellon University (CMU), Pittsburgh, PA, separated geographically by 2400 mi.

Our major finding in this paper can be summarized as follows.

1) Internet routing policies are not well suited for streaming services, and much of the pathology has to do with using AS path length as a routing metric and BGP.
2) While ISPs are internally well-connected, intradomain paths show greater delay variations in short timescales than inter-ISP peering links. This suggests that ISPs either ignore load balancing or employ strategies that create transient load oscillations. These oscillations are enough to bring down video-QoE.
3) Though inter-ISP peering links are inflated in terms of delay, we observed lesser load fluctuations in these links promising higher QoE. This implies significant cooperation among ISPs in peer link selection.
4) The combined effect of various policies maps current Internet QoE to be just about "acceptable."
5) While the Internet may not be completely multimedia ready, alternative approaches such as overlay networks are highly conducive to streaming services.

The rest of this paper is organized as follows. Section II describes our data collection methodology in detail. Section III discusses our approach to infer video-QoE from raw statistics. An overview of Internet streaming requirements as well as characteristics of discovered links are presented in Section IV. Section V studies intra-ISP links with a representative case study, while Section VI deals with inter-ISP peering links with a respective case study. Section VII presents results from our end-to-end packet transmission experiment. Discussions are presented in Section VIII, where we look at potential fixes. A case for overlay networks is presented in Section IX, and related work in Section X. We conclude the paper in Section XI.

II. Data Collection Methodology

This section describes our data collection and correlation methodology in detail. Our data collection is broadly divided in two phases: 1) Phase-1: extract topology information of the ISPs we choose to study; and 2) Phase-2: perform active tracing to estimate loss, delay, and jitter inflation. In summary, phase-1 produced a data set of more than 20 million traces from 38 PlanetLab [34] vantage points over six days to discover close to 50 293 router IP addresses that belong to the 51 ISPs we chose to study. We inferred an ISP's topology by classifying a router as either a point of presence (POP) or a backbone router. In the second phase, active probing was used to measure network event inflation between endpoints that could either be in the same ISP or use multiple ISPs to reach one another. Based on endpoint proximity in terms of ISPs traveled, we isolate the impact of network events for both intradomain and interdomain routing.

A. Choosing ISPs

We chose ISPs with great care to make our results as representative of Internet behavior as possible. We chose ISPs who are dominant players in their respective zone, which both ensures that a chosen ISP carries a great portion of that zones traffic while ensuring geographical diversity. We look for the following criteria in choosing ISP: diversity, degree (peering links), and size. We also ensured an interesting mix of *tiers*, where tier-1 represents ISPs closest to the Internet core, and tier-3 represents ISPs farthest from it. We chose 19 of the 22 tier-1 ISPs, while keeping a diverse mixture of tier-2 and tier-3 ISPs. Table I shows the list of 51 ISPs used in this study sorted by their tier.

TABLE I
LIST OF ISPS USED IN OUR STUDY GROUPED AS TIER-1 (TOP), TIER-2 (MIDDLE) AND TIER-3 (BOTTOM)

ISP Name	AS number	Primary Zone	Degree
ATT	7018	US	1490
Verizon	701	US	2569
Qwest	209	US	887
Level3	3356	US	539
Savvis	3561	US	806
Global Crossing	3549	US	585
Genuity	7405	US	622
Globix	4513	US	455
Sprint	1239	US	1735
Verio	2914	US	538
Williams Comm	7911	US	234
XO	2828	US	184
Colt	8220	Europe	161
DTAG	3320	Europe	111
Eqip	3300	Europe	67
Genuity-europe	7176	Europe	90
Open Transit	5511	Europe	172
Telia	1299	Europe	256
Hong Kong Telecom	4637	Asia-Pacific	201
AOL	1668	US	156
ATT-Disc	7170	US	60
Accretive	11608	US	124
Allegiance	2548	US	216
Broadwing	6395	US	231
Cogent	16631	US	187
Epoch	4565	US	84
Gnaps	1784	US	63
Internap	10910	US	113
PNW-GPOP	101	US	28
SW-Bell	7132	US	112
TWTelecom	4323	US	277
ATT-EMEA	2686	Europe	62
Concert	5400	Europe	117
Easynet	4589	Europe	86
GATel	13129	Europe	53
ITDNet	9070	Europe	20
PSI	174	Europe	46
Tiscali	3257	Europe	326
UUNet-europe	702	Europe	587
Vianw	5669	Europe	36
IIJ	2497	Asia-Pacific	165
One Data Network	4725	Asia-Pacific	63
VSNL	4755	Asia-Pacific	49
Soul	9942	Australia	114
Espire	6467	US	30
Exodus	3967	US	43
MFN	6461	US	498
Oregon-GPOP	4600	US	4
Univ. of Oregon	3582	US	5
ATT-AP	2687	Asia-Pacific	24
Singapore Telecom	7543	Asia-Pacific	8
Telstra	1221	Australia	66

B. Phase-1: Extracting ISP Topologies

An ISP typically consists of a backbone network and various POPs. The POPs peer with gateway backbone routers to connect traffic to and from spoke cities, and the backbone routers typically route data between cities. We discover an ISP's topology by identifying these POPs and backbone routers. As an input to our analysis, we use used traceroute data collected from 38 vantage points. The vantage points were spread across the globe, especially in the three dominant zones (US, Europe, and Asia–Pacific). From these vantage points, we traced to all of the 127 000 globally prefixed IP addresses extracted from the BGP tables of RouteViews [36], which peers with 60 large ISPs. Tracing typically took six consecutive days on any vantage point.

Traceroute data was processed to produce hop count, DNS names of routers (when available), and router IP addresses. The IP addresses found were then matched with the prefix advertised by our list of ISPs in Routeviews to find routers that belong to one of them, producing 50 293 unique matches. Of these routers, we discovered 19 832 POPs.

We inferred the AS numbers for every IP address based on prefix advertised by ISPs in RouteView tables. This converted IP-level traceroute data to AS-level paths to reveal intra-ISP and inter-ISP peering links. A given link is intra-ISP if the source–destination pair belongs to the same ISP. Likewise, when a source–destination pair belongs to different ISPs, the link denotes an inter-ISP peering link.

We go one step further in processing these AS-level paths to derive *city*-level paths. We use the undns [35] tool to assign DNS names to cities. POPs were assigned to their respective cities by matching the DNS names with the ISPs naming convention. DNS names usually have an airport code, city, and/or state abbreviation to denote their location. For example, the DNS name `cr1.st6wa.ip.att.net` indicates a router in Seattle, WA (`st6wa`) run by ATT, while `te3-1.ccr01.lax01.atlas.cogentco.com` denotes a router in Los Angeles (`lax`) run by Cogent. City-level paths give us a benchmark (speed-of-light direct fiber) of the expected network events on an idealized link connecting hosts to quantify degradation [26].

C. Phase-2: Studying Network Links

Once we discover all pairs of intra- and inter-ISP links, we perform active probing from various vantage points to measure response times and loading level of the link at various times of the day. We analyze the effects of traversing multiple links using the probe train experiment.

1) Active Probing: Since it is practically impossible to run custom programs on arbitrary routers and hops all over the Internet to perform measurements, we use an alternative way of measuring these links from vantage points. A vantage point that discovers a link of the form $A \rightarrow B$ from trace collection (usually less than 3 hops away to preserve accuracy [2]) sends out back-to-back TTL-limited probes, three at a time every 50 ms, toward the destination that produced $A \rightarrow B$. It sends out probes such that the first set of probes expire as soon as they arrive at A (TTL), and the second set of probes expire when they arrive at B. Both routers A and B send out an ICMP TTL expired message to the vantage point. The difference in RTTs provides a reasonable estimate to the level of loading experienced between A and B [3], [26]. We ensured that the replies came from the intended routers A and B to guard against routing changes. Also, the probes were sent in groups of three to break any synchronization with a routers maintenance tasks (usually 60 s [22], [26]).

In addition to the above probes, we also use King [33] to estimate latencies between routers in the Internet. Since King works by recursive DNS queries, it does not always estimate the latency between every router pair. Hence, we include King estimates only when available. In Sections V and VI, we present delay and jitter distributions taken from these measurements. We also convert these raw statistics to an anticipated video-QoE from taking that link. To analyze the time-of-day effects of traversing such links, we perform two representative case studies of an intra-ISP between Tampa, FL, and Houston, TX, run by Level3 measured from a vantage point in Orlando, FL, and a peering link between Qwest and Sprint measured from a vantage point in Berkeley, CA.

2) Probe Train: In Section VII, we conduct a probe train experiment, where a source in UCLA sent out a train of 512-B-sized high-frequency UDP packets to a destination in CMU, mimicking a high-quality streaming application operating at 30 frames per second or higher. For more than 150 000 packets sent, we measure the amount of "intact" information in an unoptimized playout buffer at every time slot, as we analyze the effects of traversing multiple ISPs to reach a destination. We also extrapolate results based on subjective perceptions of low- and high-motion MPEG-2 clips.

D. Measurable Parameters

The three most important measurable network-level parameters are delay, jitter, and loss. We provide a formal definition of all these parameters and outline our methodology for recording them.

1) Delay: We rely on two primary sources for delay estimation: active probing and the King [33] tool. King directly reports latency between arbitrary nodes, and when available, we include them. We use active probing (described above) to estimate latency of links we measure using PlanetLab vantage points close to these links.

2) Jitter: We report jitter as the variance of delay (sometimes also referred to as delay-jitter). More formally, we express jitter (J) from N delay samples (D_i for $1 \leq i \leq N$) as $J = \sqrt{\sum_{i=1}^{N} \frac{(D_i - \mu)^2}{N}}$, where $\mu = \frac{1}{N} \sum_{i=1}^{N} D_i$.

3) Loss: Missing packets directly contribute to loss. We measure and express loss rate as a percentage, which denotes the number of packets lost for every 100 packets samples. For example, a loss rate of 1% denotes one packet lost for every 100 packets sampled.

E. Caveats and Data Completeness

A trace-driven study such as this is vulnerable to errors, failures, and information misinterpretation. We look at possible caveats in our data collection methodology and discuss steps taken to address them.

1) Data Set Consistency: The data set should be correct, albeit redundant or in excess. Though we use a small subset of ISPs, we ensured that these ISPs are major traffic carriers in their zone. In fact, 87% of our traceroutes traversed at least one of these ISPs. During phase-1 of our data collection, three of the vantage points failed during the six-day event, and a power outage at our host node (web-server) led to an incomplete overnight's worth of data collection, making us start afresh. However, we found that using multiple vantage points to collect data sufficiently thwarted these failures. To measure the completeness of our data set, we measured the number of

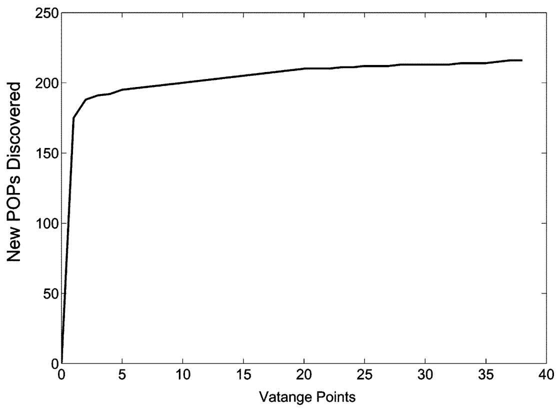

Fig. 1. New information added by additional vantage points to trace data.

new POPs discovered by the addition of each vantage point to infer ISP topologies [26]. Fig. 1 shows that a single node alone accounts for a majority of POPs learned, with additional vantage points adding marginally more information. With our usage of 38 vantage points, we are well into the knee of the curve. We also ensure that the routers we associate using prefix matching indeed belong to the ISPs that we infer them to be by comparing AS numbers obtained from `aslookup` [32]. We take a conservative approach by using the common subset of ISP-router matches from these two methods (determined to be about 83%).

2) Transient Routing Changes: Transient instabilities that happen on any one given day in the Internet do not necessarily paint a correct picture of the state of affairs. To account for transient routing changes or node failures, both phase-1 and phase-2 of data collection were spread across multiple days. It has been shown that such transient routing changes do not last more than a day [19], [20], [22], [26].

3) Spurious Links: Certain false links can appear due to transient routing changes and TTL-based traceroute path discovery. We identify spurious links by applying the speed-of-light criterion: Links that promise lower delay than possible are spurious.

III. Estimating Video-QoE of Links

A variety of investigations have been performed to infer and predict the QoE of a video stream. Unlike the R-Score in the VoIP industry, there is no universal consensus on a video-QoE estimation model. However, researchers have consistency claimed two things: 1) QoE is a much stronger indication of human perception over QoS models; and 2) QoE is affected by loss, delay, and jitter as far as networking is concerned. To make our video-QoE inference as general as possible, we present both uninterpreted statistics of loss, delay, and jitter as well as provide a mapping of QoE. We present three approaches to map trace data to QoE: one using an objective function to return a numerical score (Sections IV and V), inferring subjective perception of clips encoded using MPEG-2 (Section VII), and by estimating the amount of "intact" information at a receiver playback buffer (Section VII).

1) Objective Function Model: Objective mapping functions often return the inferred QoE as a number, or a mean opinion score (MOS). MOS ratings loosely classify a video stream as good, acceptable, poor, or anything in between. We use one such model recommended by the ITU-T [13]. This model takes into account the network delay of audio and video packets, the audio–video sync during playout, the encoding bit rate and frame rate, as well as the type of receiver to predict video quality (on a scale of 1–5). Raw network measurements are used as an input to the model, which returns the anticipated quality degradation by using that link. The defaults assumed in estimating QoE using this model are shown in Table II. In addition to the above, we also assume audio to be of the highest quality (not impaired).

TABLE II
Default Parameter Values Used in the Objective QoE Function From ITU-T [13]

Video Metric	Default Value
Codec Type	MPEG-2
Video Format	VGA
Key Frame Interval	1 sec
Video Display size	9.2 inches

2) Playout Buffer Analysis: We also look at the contents of a receivers "playback" buffer to estimate video-QoE. Ideally, if every host had a direct fiber optic connection to every other host in the network, streaming packets would arrive at a destination in the exact order and time interspacing that they were sent at. In such a case, the destination would consistently have a buffer's worth of data to ensure a smooth playout. Hence, the amount of information *absent* at a receiver buffer is a strong indication of what the network did to the video stream. Jitter causes packets to arrive badly out of sequence long enough to be considered missing, while loss directly contributes to missing information. We measure the amount of missing information in 100 ms of buffering at the receiver. We infer perceptual quality based on information available for playout.

IV. Internet Video Streaming

We begin with a brief review of Internet routing requirements for streaming services and present a high-level overview of the Internet paths studied in this paper. We base this in the context of multimedia streaming over the Internet and briefly discuss video buffering and its effects on preserving interactivity.

A. Typical Internet Route

Internet routing is a process of finding a series of ASs to traverse from a source node to a destination node. The ASs themselves are organized in a hierarchy of tier-1, tier-2, and tier-3 ISPs and are typically characterized as carrier ISPs (tier-1) or stub ISPs (tier-2 and tier-3). Carrier ISPs peer with stub ISPs to route traffic between them, and this relationship is dictated by *provider–customer* contractual agreements between the carriers and stubs. Most source-to-destination pairs in the Internet traverse a combination of an ascending tier path (lower to higher tiers) and a corresponding descending tier path, in that stub ISPs are typically reached via an intermediate tier-1 ISP. A few exceptions to this rule include the presence of exchange points and specific peering relationships between stub ISPs.

We begin with a high-level characterization of an Internet path in traversing from source to destination in terms of number of hops and the distribution of intra- and inter-ISP links we observed from our traceroute (Table III). The table shows that a

TABLE III
CHARACTERIZATION OF INTERNET ROUTES STUDIED

Path Metric	Measured Value
Mean hop count	7.3
Intra-ISP links	72%
Inter-ISP (peering) links	28%

majority of links used to traverse the Internet are *intra*-ISP links, which account for 72% of all the links we studied. This implies that degradations caused within an ISP can have a more profound impact on video quality than those due to peering links. An average Internet route today involves traversing about three ASs and consists of a majority of intra-ISP paths and a few peering links, one each when an AS boundary is crossed.

B. Video Buffering Versus Interactivity

Streaming services over the Internet present the opportunity of significant user interaction, perhaps more than cable-based services. For example, users could go beyond channel changes to participate in opinion polls or providing feedback. Network anomalies that degrade video quality are loss and jitter. Loss directly contributes to missing information. Jitter can cause packets to arrive out of order, sometimes enough to render a received packet useless. Almost all receivers implement a playout buffer to counter network jitter, stalling playback until a buffer worth of information is received. The size of the playout buffer can have a significant impact on user interactivity: Each time a user flips a channel or requests new material (forwarding, rewinding, etc.), the buffer contents are flushed, and information from the new stream is rebuffered before playout. Channel changes apart, real-time streaming (e.g., live broadcasts) necessarily requires smaller buffer sizes to preserve a near real-time viewing experience.

For services such as IPTV or VoD, interactivity is mostly presented as end-user "zapping," which includes channel changes, forwarding, rewinding, etc. Zapping delays of more than 2 s are perceived as poor by end-users. Network service providers typically target round-trip delays of less than 500 ms [8], [9]. Zapping behavior is closely tied with user attention span and browsing habits. A recent investigation into six months of end-user browsing habits reveals that user attention span can be quite low: Over 60% of channel changes happen within 10 s, and a user's favorite channels are nonsequential, but browsing habits are predominantly sequential [5]. This means that the amount of interactivity presented by the average user can be significant, and a consistent loss in interactivity can lead to the user perceiving a service to be poor. On the same lines, users who perceive zapping times to be too high show an unwillingness to switch channels, which further degrades their perception of quality. Hence, an ability to provide consistent zapping experience is crucial to prevent subscriber churn.

V. INTRA-ISP ROUTING POLICIES

We begin by analyzing intra-ISP links and their relative impact on video-QoE. Intradomain links interconnect various routers that belong to an ISP, which are typically POPs in the same or different city.

We perform active probing on a set of 800 intra-ISP links from our 38 vantage points. Probing was continuously performed for an hour on each session, with at least 12 sessions per day for six days. Probing at various times of the day smoothens rate fluctuations that happen throughout the day, and probing for multiple days ensures we smoothen transient routing changes.

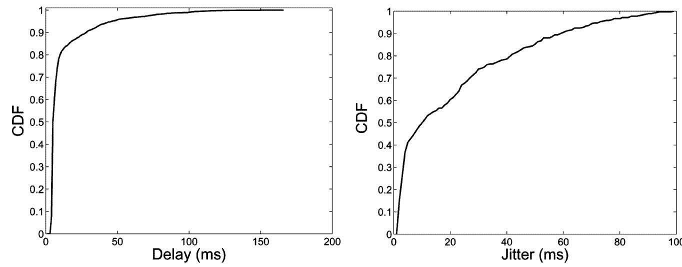

(a) (b)

Fig. 2. Characterizing *intra*-ISP links. (a) Delay distributions and (b) jitter distributions for 800 intra-ISP links from six days of active probing.

Our results offer surprising insights into intradomain routing policies and their impact on video-QoE. While we find that ISPs are internally well connected, there seems to be significant load *imbalance* in network links over short timescales. This suggests that ISPs either do not perform extensive load balancing among their internal links, or the applied strategies are topology-insensitive, resulting in load oscillations. Also, it is well known that RIP and OSFP, the most popular intra-ISP routing protocols, are load-insensitive. Load imbalance creates excessive jitter, which is known to have the greatest impact on video-QoE.

A. Delay and Jitter Distributions

Active tracing reveals delay and jitter distributions for the set of intra-ISP links we study. Delay distributions are shown in Fig. 2(a), which shows that a majority of path lengths are fairly good (< 50 ms), with a great percentage (70% of all links) less than 25 ms. This means that ISP are internally well connected. Though a majority of paths connect city pairs in close proximity, we found certain intra-ISP links connecting city pairs across countries and even continents. This was particularly true for tier-1 ISPs. We did not see any relative difference in the mean delay values based on geographic dominance of an ISP, which seems to indicate that ISP interconnection topologies are consistent (and possibly replicated) across continents.

The jitter distributions [Fig. 2(b)], however, show a surprisingly high value for some points. Jitter levels of more than 20 ms are known to have a very adverse effect on perceptual quality [4]. The plot shows that approximately half the links we study have jitter levels more than 20 ms. Despite the fact that intradomain routing is under the complete purview of an ISP, high jitter indicates that ISPs either do not perform any kind of load balancing, or perform load balancing that results in oscillations. In fact, the jitter distribution of intra-ISP links was observed to be higher than inter-ISP peering links. We also observed higher jitter values for many tier-1 ISPs links, which could be due to the fact that tier-1 ISPs have a large number of nodes making traffic engineering rather complex.

B. Case Study: Level3's Link Between Tampa and Houston

We wondered why intradomain routing would produce such high levels of jitter and sought out to investigate the true causes of it. We chose a representative link to perform sustained active probing for 24 h to determine if there was a change in the loading level of that link. A representative link would have mean delay and jitter levels below the knee of the curve of Fig. 2(a) and (b), which maps about 85% of all links discovered. We chose an intra-ISP link run by Level3 between Tampa (FL) and Houston (TX) from a vantage point in Orlando (FL). We probed this link continuously for 24 h on July 23 and 24, 2009, and took measurements in bins of 2 min each. Each bin represents the result of 240 probes, with a 24-h study accounting for 172 800 probes.

Delay variations in the 24-h period are shown in Fig. 3(a). The plot shows that delay values vary depending upon the time of the day, where it peaks between 3:00 and 9:00 PM for that day. We also discovered two significant outages, both before 9:00 AM on July 24, where a good fraction of the probes were lost, whiles others made it in more than 100 ms. Delay values show fluctuations indicating jitter.

Jitter deviations are shown in Fig. 3(b). Jitter patterns correlate with the delay patterns, in that they peak between 3:00 and 9:00 PM and are reasonably high during the outage times. Jitter values at 95% confidence interval tail the mean values closely. The plot shows consistent jitter at around the 10-ms range, which is enough to degrade video streaming.

To indeed verify that jitter at this link is because of a fluctuation in traffic directed to this link, we measured the median loading of the router at Tampa that connects to Houston. To measure loading, we recorded the router timestamp on the ICMP packets. Variations in the time difference at the destination router with no variations at the source are indications of loading at the source link [26]. Fig. 3(c) shows the median loading at the router in 2-min bins. Indeed, there are significant fluctuations in the link load in very short timescales.[1] Also, loading continues to peak at around 6:00 PM and was considerably high during the outage times.

MOS variations for this study are shown in Fig. 3(d), which are obtained from the ITU-T model [13]. The results suggest that MOS projections are overall in the "good" range, with jitter spikes ensuing in consistent MOS ratings of less than "acceptable." MOS dropped to "poor" for an extended period of time at around 6:00 PM. Consistent jitter spikes, such as in this case, have subtle long-term effects on QoE. Studies have also shown that subjects have a "forgiveness" effect, where users rate a video clip relatively high if the playout is smooth after a brief initial loss in quality. However, with a regular (almost periodic) degradation of quality, users would generally rate the video sequence lower on a longer timescale than the one we chose. The average MOS was around 2.46, which suggests quality just above "acceptable."

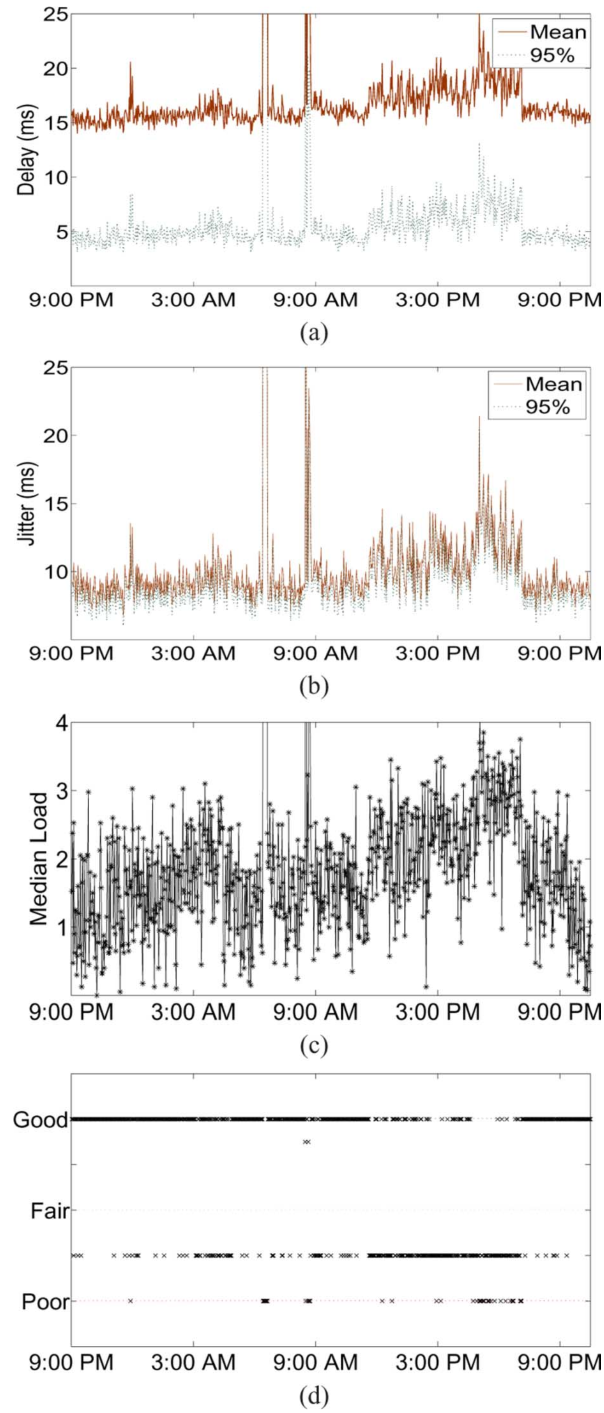

Fig. 3. 24-h observation of an intra-ISP link by Level3 connecting Tampa and Houston. (a) Delay distribution (mean and 95% confidence interval). (b) Jitter distributions (mean and at 95% confidence interval). (c) Median loading of the peering link. (d) Estimated QoE at various times of the day.

VI. Inter-ISP Routing Policies

We next study the peering links used by ISPs to exchange traffic between one another. ISPs often use multiple peering links between one another, and the policies used by ISPs to choose peering links to exchange traffic is largely proprietary and driven by a multitude of factors. We measured 1100 such links that connect routers of different ISPs using active probing to measure the relative loading of each such link.

Our findings indicate that there is significant cooperation among ISPs in distributing load across links, and that peering links are well balanced and even. This was particularly the case for ISPs that used a large number of peering links between one another. Peering links were overall good for tier-1 ISPs compared to cases when a tier-3 ISP was involved. Peering links did not show fluctuations based on geographic presence.

[1]With an OC3 link, median queuing delay of 1 ms corresponds to 40 packets in the queue [22].

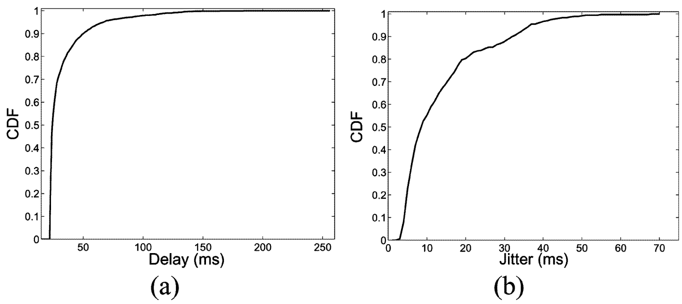

Fig. 4. Characterizing *inter*-ISP peering links. (a) Delay distributions and (b) jitter distributions for 1100 inter-ISP links from six days of active probing.

However, peering links connecting multiple continents were relatively suboptimal.

A. Delay and Jitter Distributions

Delay distributions for peering links are shown in Fig. 4(a). The plot shows that at least half the links studied have delay distributions less than 25 ms, and about 30% are in between 25 and 75 ms. Delay distributions are typically higher than intra-ISP links, suggesting that peering links are relatively less than intra-ISP links and that peering links often connect distant routers. Certain peering links connect distant cities and even continents, which also accounts for higher delay distributions.

Jitter variations are shown in Fig. 4(b), which shows that more than roughly half of the links studied are in the "good" range (<10 ms), and a greater majority of the remaining in the "acceptable" range. This means that delay variations of inter-ISP links are less than delay variations of intra-ISP links, suggesting that there is significant load balancing in peering link selection. We found jitter values to be less for peering links of tier-1 ISPs than between links of tier-3 ISPs.

B. Case Study: Peering Link Between Sprint and Qwest

To provide a basis to compare a typical intra-ISP link to a peering link, we chose a peering link that is representative of the delay and jitter distributions of Fig. 4(a) and (b). We chose a peering link between Sprint and Qwest in California, measured from a vantage point in Berkeley. We actively probed this link for 24 h on July 23 and 24, 2009, for a total of 172 800 probes.

Delay distributions over various times of the day are shown in Fig. 5(a). Delay averaged around 15 ms for most of the study and peaked around late afternoon to early evening. However, the relative difference in delay values at short timescales is not too high. Jitter distributions for 24 h is shown in Fig. 5(b). Jitter fluctuations are low at night and most times of the day, reaching a peak at around 6:00 PM.

We also measured the load on the links origin router at Sprint [Fig. 5(d)]. We compared ICMP router timestamps to look for fluctuations at the Qwest router with little fluctuations at the Sprint router. However, we found that the link was evenly loaded for most of the day, with some fluctuations at around 6:00 PM. An even loading such as in this case strongly indicates that traffic is evenly distributed across peering links used between Sprint and Qwest, a trend that we predominantly saw for most of the peering links we studied.

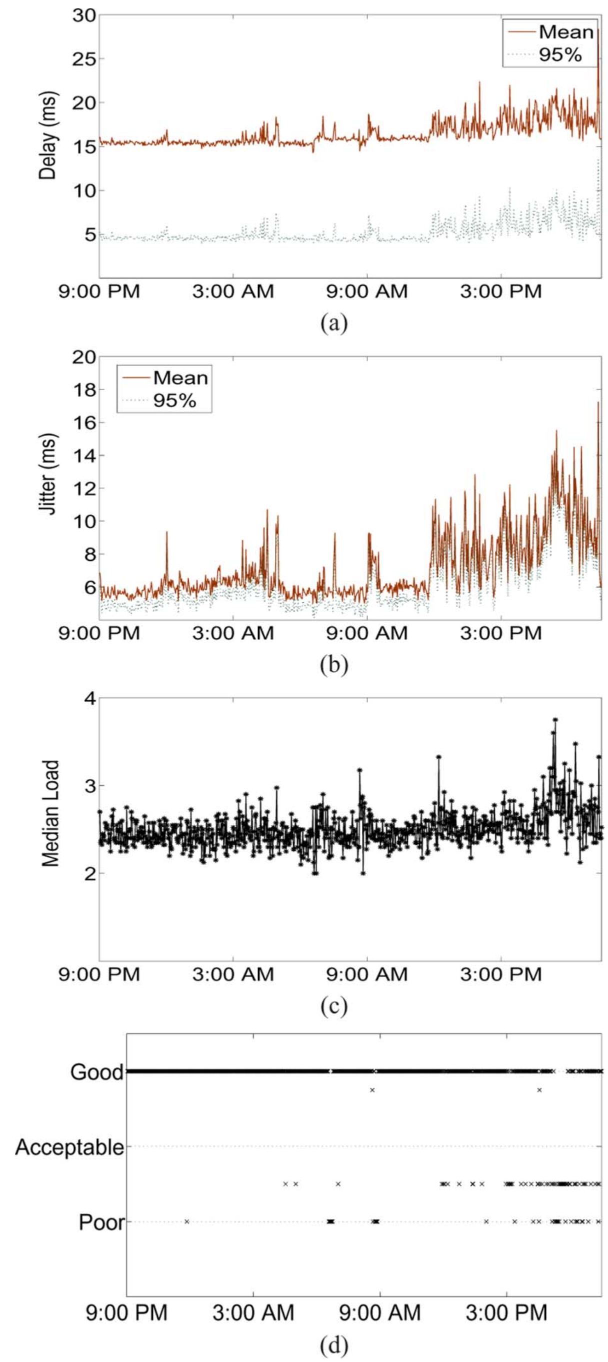

Fig. 5. 24-h observation of a peering link between Sprint and Qwest in California. (a) Delay distribution (mean and at 95% confidence interval). (b) Jitter distributions (mean and at 95% confidence interval). (c) Median loading of the peering link. (d) Estimated QoE at various times of the day.

The combined effect of these on video-QoE is shown in Fig. 5(d). A vast majority of the projections are in the "good" range. Of the points not in the good range, almost all of these are close to "acceptable" with no incidents of "poor" quality. The overall MOS for this study was 2.8, much closer to "good" than the 2.46 of the Level3 link.

VII. Playout Buffer Analysis

We next piece together the combined end-to-end effect of choosing a combination of links to traverse the Internet. We

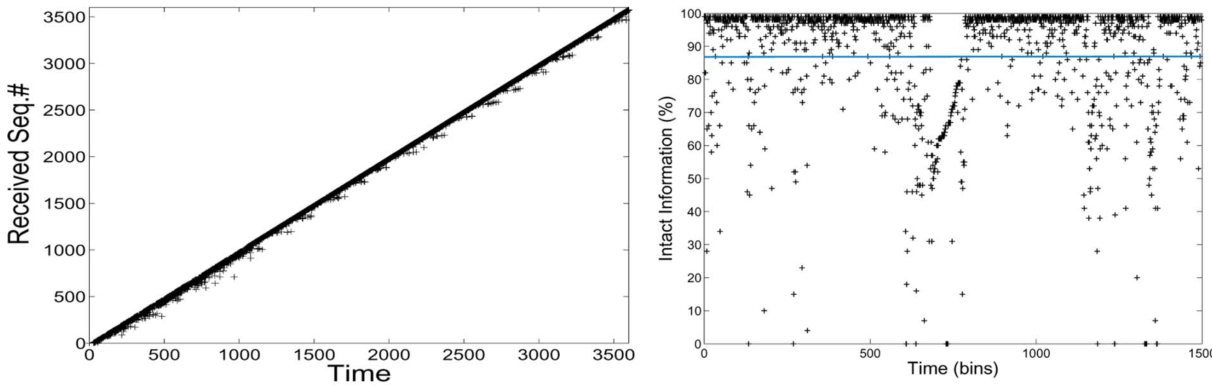

(a) (b)

Fig. 6. (a) Received sequence number at destination for the first 3600 packets. The x-axis denotes the time since the first arrival of a packet, while the y-axis indicates the anticipated sequence number. (b) Contents of the unoptimized playout buffer in various time slots. Horizontal bar represents the mean.

diverge from both active probing and an objective QoE model by next looking at the receiver playout buffer of an end-to-end transmission that mimics a real-time, high-quality multimedia stream. We round off this study by placing the results in the context of an MPEG-2 stream and discuss the impact of degradation for low- and high-motion clips.

We sent 150 000 packets from a source in California (UCLA) to a destination 2400 mi northeast in Pennsylvania (CMU). The packets were 1024 B each and were sent out at the rate of 1 packet every 10 ms (or 100 packets per second). High-data-rate applications typically transmit more than 30 frames per second, which amounts to 120 packets per second assuming each frame takes 4 packets to encode.

The destination receives packets and places them in a "playout" buffer, which stalls for a certain time waiting for a batch of packets to arrive. A playout buffer of 100 ms is recommended for our sending data rate [4]. The buffer is unoptimized and nonadaptive and does not vary buffer limit based on the current reception count. Packets that do not arrive in its expected time slot are *discarded*, but packets that arrive out of order in a given time slot are received correctly. Certain enhancements to buffering strategies often adaptively increase the buffering limit in the face of little or no packet reception, stalling playback at such times [7]. However, since our goal is to understand the raw capabilities of the Internet without optimizations, we underestimate buffering by choosing this playout buffer.

A. Packet Reception

We plot the reception time for each packet for the first 3600 packets received at the destination. Fig. 6(a) shows the time that each packet was received, with time starting from the first packet received at destination. If the source–destination pair had a hypothetical direct optical fiber link between them, the data points would completely overlap the $y = x$ line. Though we observed little loss ($< 0.1\%$), we see that packets frequently arrive beyond their expected time.

B. Playout Buffer Contents

We now look at the contents of the playout buffer at various time slots. We measure the amount of "intact" information available in the playout buffer. We measure the intactness by calculating the percentage of correctly received sequence numbers. The playout buffer infers the expected set of sequence numbers based on the slot number since the time of the first packets reception.

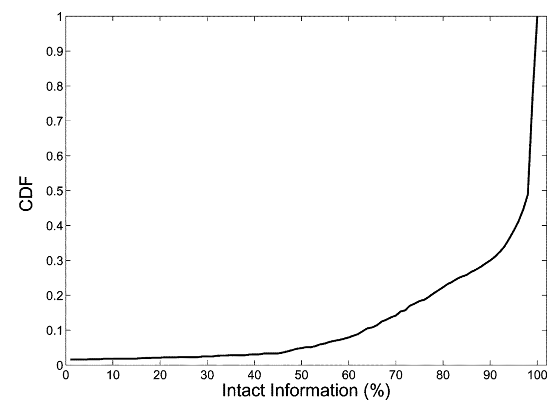

Fig. 7. CDF of the amount of "intact" information in the buffer.

Fig. 6(b) shows the intact information percentage at every time slot. Each time slot is of 1-s duration, which amounts to 1500 time slots for all of our packets. The amount of content in the buffer shows great variation in time. Research has shown that even a marginal loss in information manifests as user dissatisfaction, and more than 20% loss certainly degrades video [16]. The plot also shows that the buffer contains less than 40% of required information at regular intervals of time, even dipping to less than 10% every 4 min in the playout. This means that video quality will frequently dip to low ratings almost periodically. Consistent loss in video quality manifests as a strong dissatisfaction. The mean amount of information in the buffer was calculated to be 86.1%.

Fig. 7 shows the cumulative distribution function (CDF) of the information available in the playout buffer. The plot shows that the buffer has less than 90% of intact information more than 30% of the time, and it has 100% intact information less than 20% of the time. The combined effects of multiple links that induce jitter are hence not merely additive.

C. MPEG-2 Playout Perspective

Streaming content in IP networks is commonly transported as a data stream encoded using the MPEG standard and

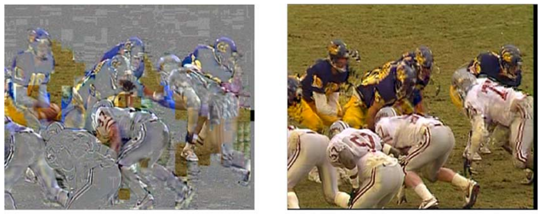

(a) (b)

Fig. 8. QoE versus QoS. Both clips experience the same loss rate (QoS), but the QoE is very different: (a) key frame corrupt (low QoE); (b) non-key frame corrupt (higher QoE).

transported via the real-time protocol (RTP) over a UDP/IP stack. MPEG encodes video streams as a series of Intra (I), Predictive (P), and Bidirectional (B) frames. I-frames carry a complete video picture and, as such, provide reference to the following B- and P-frames for decoding an MPEG stream. P-frames predict the frames to be coded using a preceding I- or P-frame. Lastly, B-frames use the previous or next I-frame for motion compensation.

A typical video dissemination model consists of a head-end that either acquires content from a third party or transmits stored content. The video content is normally transmitted as individual frames. Each frame is typically fragmented into multiple IP packets for transport over the Internet. The packets at the receiver are likewise reassembled from an ingress transmission queue to recreate these frames at the destination. Note that some techniques like error correction, error concealment, and prefetching frames are also directed toward maintaining superior perceptual quality.

The frames are packed into a group of pictures (GOP), where each GOP consists of an I-frame at the start and a series of B- and P-frames that use it as a reference. Depending upon the motion complexity inherent in a clip, the structure of a GOP can be very different: Low-motion clips (like a news program) have larger I-frames and a handful of P- and B-frames to complete a GOP, while high-motion (sports clip) clips have smaller I-frames and relatively larger P-frames for motion compensation.

1) High-Motion Clips: Quality degradation due to an I-frame loss is dominant [e.g., though screenshots in Fig. 8(a) and (b) experience the same loss rate, a corrupt I-frame renders Fig. 8(a) barely perceptible]. However, because of the inherent dynamism, P-frames tend to be larger and more informative. Hence, loss of a P-frame draws a slightly more adverse reaction from the subjects than with low-motion clips. Increased loss of P-frames continues to degrade quality. Because scenes change frequently, I-frames tend to be shorter and less probable candidates for loss during outages. Increased loss rates within an I-frame does not seem to degrade quality any further [8], [9].

2) Low-Motion Clips: Loss of P- or B-frames seems to make little difference in perceptual quality. The difference in perceptual quality for the loss of an I-frame over a P-frame, however, is drastic. The ensuing loss in perceptual quality is similar to Fig. 8(a) for a corrupt I-frame. However, low-motion clips have longer GOP structures with more P-frames, which enables more compression. Because of their long GOP structures, the *duration* of on-screen degradation tends to be longer than for high-motion clips. Low-motion clips also have larger I-frames, which increases their odds of getting impacted during an outage.

VIII. Discussions

We piece together the above studies to infer insights about the way the Internet operates and suggest several ways to overcome these shortcomings. We look at a variety of causes that both contribute to these findings and are affected by them. We take a closer look at intra- and interdomain policies of BGP, which governs route selection, and suggest workarounds. We analyze other solutions, such as increasing capacity, and multihoming as alternatives to providing better streaming quality. Lastly, we take a close look at the prospect of using overlay networks for deploying multimedia networks with one example deployment.

A. BGP and Path Selection

BGP is the *de facto* routing protocol for inter-AS and intra-AS routing. It manages reachability information shared between two ASs and allows for diverse networks to interconnect and become the Internet. BGP summarizes, and often hides, internal topological details about an AS to prevent routing oscillations in a process called "route dampening."

1) Intra-domain (iBGP): Intradomain routing is largely governed by iBGP sessions, which are an overlay of nodes over the OSPF substrate. While ISPs are known to perform certain optimizations intended toward load distribution, routing instabilities within an AS often lead to short-term anomalies. Short-term traffic fluctuations are typically not captured or accounted for and are clearly not a metric for intradomain routing. BGP convergence times are of the order of minutes, and it has been shown that topology sensitive load balancing is hard with BGP. Prevalence of mechanisms such as early or late exit and latency-based route selection within an AS largely contribute to load imbalance. Our results call for even traffic distribution within an AS, which could help improve streaming quality. ISPs are known to be internally well connected, and discovering and utilizing redundant paths to distribute load could very well improve the quality of intra-ISP links [28].

2) "AS-Path Length" Based Routing: BGP determines the choice of ASs used to traverse from source to destination. Route selection is typically performed by comparing AS-path lengths advertised by address prefixes. Research has shown that about 11% of all new paths learned by BGP are unreachable [19], and paths are often inflated in terms of distance traveled [26]. Fig. 9 captures the effect of using AS-path lengths as a routing metric for our probe train experiment between UCLA and CMU. Internet route selection took the packets from California (far west) to Louisiana (east), back to Houston (west), then Atlanta, GA (northeast), and finally Washington, DC (north) before making it to CMU. Clearly, the choice of ISPs traversed badly mismatches a direct series of links between these two geographic locations. Since our findings indicate that the quality or choice of peering links is not a factor for ISPs, we argue that inefficiencies in route selection arise from using AS-path length as a routing metric.

Given that a great majority of Internet traffic will carry multimedia content, we argue for BGP advertising "geographical" information at BGP speaking routers. That way, even an inexact reachability information at a router can be compensated

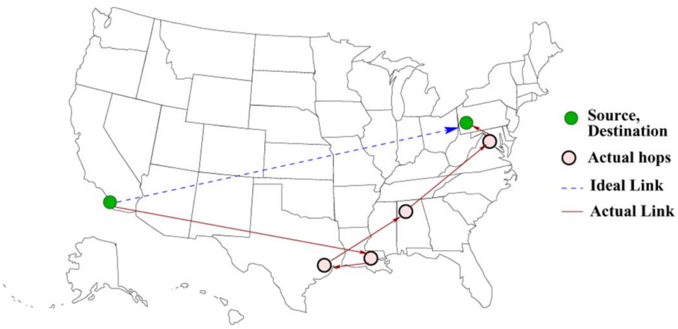

Fig. 9. Ideal path versus actual Internet path for the probe train experiment.

by allowing video streams to "flow" toward the destinations geographical zone. Though far fetched into the future, we believe geographic coordinates embedded in BGP route advertisements can certainly enhance streaming applications.

B. Increasing Capacity and Expected QoE as a Metric

1) Increasing Capacity: A popular suggestion to improve performance is to increase capacity. However, increasing capacity with high jitter levels in the network only *exacerbates* the problem. This is because an increase in capacity provides higher peak levels for the traffic to reach. When links are unbalanced, jitter typically increases with network capacity. In fact, load balancing is known to have often increased the capacity of existing infrastructure [15].

2) Expected QoE as a Routing Metric: Intuitively, if inter-AS path selection was based on *expected QoE* rather than AS-path length alone, streaming content would deliver higher perceptual quality. However, this is not entirely practical for the following reasons: 1) it is not possible for every router to probe each of its links continuously to create a picture of expected QoE; 2) fluctuations in expected QoE may cause route flapping; and 3) changing the behavior of the entire Internet is known to be extremely slow. For example, QoS mechanisms such as IntServ and DiffServ were envisioned to be implemented at all network elements from source to destination. The actual adoption, however, was sporadic and uncoordinated.

IX. Case for Overlay Networks

While change in the Internet may be slow, we look at alternative approaches that can provide some imminent solutions. We look at the case of network overlays tuned for multimedia. Information "detouring" [25] and overlay networks [1], [27] have been shown to be effective workarounds around Internet failures. Overlay networks are a collection of nodes in different autonomous systems that maintain a *virtual* link to one another. A virtual link is the normal IP-level path between these two nodes, which the nodes probe continuously to monitor signs of degradation. Overlay nodes exchange this reachability information with each other periodically to expose other *redundant* paths in the Internet across different ASs that BGP cannot advertise. If the normal path between these two nodes experiences outages, nodes can switch to alternate paths that reach the destination. It has been shown that at most one such "reroute" is enough to route around outages [1]. We bring out benefits of using an overlay-based approach to route around IP failures using measurements from a measurement overlay we deployed.

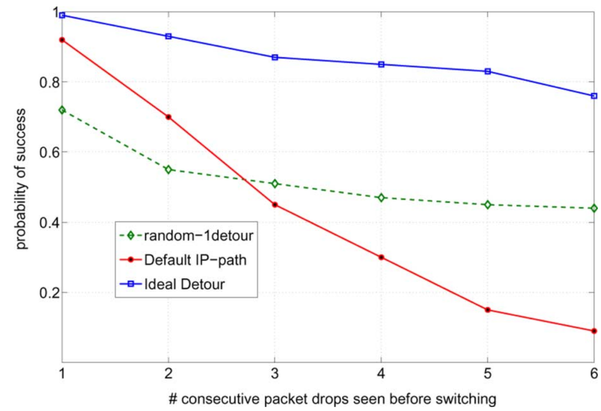

Fig. 10. Recovering from perceptual degradations: a case for overlays that avoid faulty ASs.

A. Methodology

We analyze weeklong measurements of a large number of Internet paths all over the world to understand the benefits of quick redirections to preserve video-QoE and avoid transient fluctuations within an ISP. We created an overlay of 32 nodes deployed in geographically diverse locations, including the US, Europe, and Asia–Pacific.

1) Experimental Setup and Video Clips: Between January 22–29, 2010, every node streamed 1024-B UDP packets every 5 min to a randomly selected destination. The stream mimics the IP-packet trace of a randomly selected high- or low-motion clip from a set of five clips used in the previous round of study. We passed the name of the clip and the type of frame the packet carries in the packet payload, creating an IP-trace of the clip at the receiver. For any of our overlay with N nodes, the source indirectly probed the destination via the $N-2$ other intermediaries while streaming packets to the destination. This probing is performed only when transmitting key frames within the clips. For low-motion clips, we mark the I-frame alone as the key frame. For high-motion clips, we mark both the I- and P-frames as key frames. We record the receiver trace at the destination and the probe responses from the intermediaries at the source to analyze offline the suitability of alternate paths during outages.

B. Benefits of an Overlay Network

A potential workaround to transient degradations that occurs within an ISP is to switch to a route that *avoids* the ISP in question, usually involving a one-hop detour through another node that avoids this ISP [25]. To determine the benefits of switching, and in particular switching early, we measure the probability of restoring key frames following a degradation in Fig. 10. This plot shows the probability of the next key frame being received successfully *after* observing a certain number of consecutive drops in a key frame. We plot this probability for up to six consecutive packet losses observed for the default IP-path, a random-1 path, and the "ideal" detour path. The random-1 path is derived by attempting to detour with *any random* intermediary from our set of 32 nodes. The "ideal" detour path is the most optimum path to take following a degradation, usually obtained by extensive background monitoring of alternative paths.

In effect, we base the default-IP path against a measurement free (random-1) and measurement-based informed detours.

After only two successive drops, the probability of the default IP-path restoring the next key frame seems to diminish to around 0.42. The "ideal" detour maintains a higher recovery probability of more than 80% for up to five consecutive drops. The plot also shows that even random-1 is able to provide higher returns than the default IP-path when three or more packets are lost in succession. This leads us to believe that an early detection and subsequent detour can help recover from degradations.

In general, overlay networks can solve the following problems in terms of raising perceptual quality: 1) while Internet routers cannot individually monitor expected-QoE along links, a small set of overlay nodes can easily achieve this with low overhead probing, and reroute packets based on expected QoE; 2) discover and expose redundant paths that cannot be exposed by BGP, providing greater routing options especially around ISPs that experience transient degradations; 3) implement policy routing and allow customized media dissemination; and 4) converge faster than BGP and recover from outages quickly. In other words, we believe overlay networks can provide a flexible way of implementing policy-based networking changes that cannot be immediately brought about in the Internet.

C. Internet and Overlay Networks

We round off this discussion with some potential drawbacks of overlay networks. Specifically, we discuss the impact of overlay networks on traffic engineering and violations of ISP-peering relationships due to overlay networks.

Overlay networks as a concept has been well investigated in the past decade. Overlay networks, much like peer-to-peer networks, can disrupt or introduce completely new traffic patterns that ISPs do not currently account for. Since overlay networks take control of route selection at the application level, this often disrupts the assumed control at the IP-layer. Overlay networks actively divert traffic on paths the application would never have taken if it assumed the default-IP routing.

An ISP manages traffic by employing a variety of traffic engineering techniques such as load balancing, differential link weights, and proprietary routing policies—all of which are violated by overlay networks. A study of the interaction dynamics between the two layers of control—ISP and overlay routing—was performed in [17]. The authors conclude that current ISP policies are in fact inadequate to deal with emerging overlay services. Another work by Wang *et al.* [31] studied the economic implications of overlay networks on ISP-peering relationships. They conclude that ISPs can indeed take steps to better manage traffic and prevent freeloaders. However, such steps could also reduce the ISPs' collective bargaining power toward subscribers. To conclude, while the discussions of the impact of ubiquitous overlays on the Internet is an open debate, our results do indicate that overlay networks can provide a workaround to raising Internet video-QoE.

X. Related Work

Internet pathologies have been well investigated in the past [18], [23]. The first such large-scale study revealed that Internet path selection was often suboptimal with persistent routing loops, failures, and network unreachability [23]. Even after a decade since that paper, researchers have consistently found that Internet outages are both unpredictable and, worse, can go undetected for a while [1], [19], [20], [25]–[27]. Recent work continues to highlight path inflation in the Internet [26], and likewise, BGP anomalies and misconfigurations [19]. As far as streaming services are concerned, VoIP performance study on the Internet backbone links reveal that a significant number of paths yield poor quality [21].

Internet QoS was aimed at enabling streaming services. The IETF standardized IntServ and DiffServ router mechanisms to improve quality of streaming content, which required changes to every router in the Internet. Given the scale of the Internet, as well as the diversity of various ASs it comprises, these changes could not be completely coordinated. Even after years of slow adoption, there has been no significant performance enhancements in terms of video quality, as the Internet continues to operate on a "best-effort" delivery model. OverQoS [27] investigated the prospect of enhancing QoS using overlay networks. QoS mechanisms operate with a notion of providing service guarantees to enhance application performance. However, service guarantees alone are not sufficient to raise *perceptual* quality. Perceptual quality is best characterized by QoE, which attempts to infer quality from a user's perspective. QoS-based quality assessments have often found to be grossly inaccurate at predicting user experience and, as such, are not applicable in evaluating video quality [12], [14], [24], [29].

There has been a recent surge of interest in subjectively assessing quality of a video stream. Even at a time when there is little or no consensus on the exact nature of such a model, the basic idea has been to provide a subjective interpretation of various events on a video stream, such as loss, delay, jitter, and so on. Performance bounds for various factors to ensure high perceptual quality was investigated in [4]. However, much of the prior work has investigated Internet anomalies and video-QoE independent of one another.

XI. Conclusion

Network service providers need to raise video-QoE to attract customers and prevent churn. QoE has been found to be the most significant measure of human satisfaction and has long been used to characterize customer experience with vendors in a wide variety of domains. Understanding and improving QoE has had a profound impact on the long-term success of many organizations.

This paper presented the first empirical observation of Internet links from a QoE perspective. We studied intra- and inter-ISP links of 51 major ISPs spread across the US, Europe, and Asia–Pacific from 38 vantage points in the Internet. We closely studied these links for six days to infer their suitability for streaming services.

Our findings offer surprising insights at link-level fluctuations in the Internet.

1) We find that intradomain links are poorly engineered, showing significant loading fluctuations that can degrade video quality. This could largely be attributed to BGP, which makes topology sensitive load balancing hard.
2) Contrary to popular belief, ISP peering links are well engineered and quiet suitable for carrying streaming content. This implies that ISPs are not limited by their choice or

quality of peering links. Rather, AS-path lengths are insufficient as a routing metric.
3) The overall effect of Internet route selection maps the current QoE to just about "acceptable."
4) Overlay networks could overcome a majority of Internet's shortcomings in delivering multimedia—both because they can be deployed with relative ease compared to reengineering the Internet, and provide alternative routes that avoid ASs that experience transient fluctuations in delivering packets.

REFERENCES

[1] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris, "Resilient overlay networks," in *Proc. 18th ACM SOSP*, Banff, AB, Canada, Oct. 2001, pp. 131–145.
[2] A. Akella, S. Seshan, and A. Shaikh, "An empirical evaluation of wide-area Internet bottlenecks," in *Proc. ACM SIGMETRICS*, San Diego, CA, Jun. 2003, pp. 316–317.
[3] K. Anagnostakis, M. Greenwald, and R. S. Ryger, "cing: Measuring network internal delays using only existing infrastructure," in *Proc. IEEE INFOCOM*, Mar. 2003, vol. 3, pp. 2112–2121.
[4] P. Calyam, M. Sridharan, W. Mandrawa, and P. Schopis, "Performance measurement and analysis of H.323 traffic," in *Proc. PAM*, Antibes Juan-les-Pins, France, Apr. 2004, pp. 137–146.
[5] M. Cha, P. Rodriguez, J. Crowcroft, S. Moon, and X. Amatriain, "Watching television over an IP network," in *Proc. ACM IMC*, Vouliagmeni, Greece, Oct. 2008, pp. 71–84.
[6] "Cisco visual networking index: Forecast and methodology, 2008–2013 Cisco Inc., San Jose, CA, 2009 [Online]. Available: http://www.cisco.com
[7] K. Fujimoto, S. Ata, and M. Murata, "Adaptive playout buffer algorithm for enhancing perceived quality of streaming applications," in *Proc. IEEE GLOBECOM*, Taipei, Taiwan, Nov. 2002, vol. 3, pp. 2451–2457.
[8] J. Greengrass, J. Evans, and A. C. Begen, "Not all packets are equal, part I: Streaming video coding and SLA requirements," *IEEE Internet Comput.*, vol. 13, no. 1, pp. 70–75, Mar. 2009.
[9] J. Greengrass, J. Evans, and A. C. Begen, "Not all packets are equal, part II: The impact of network packet loss on video quality," *IEEE Internet Comput.*, vol. 13, no. 2, pp. 74–82, Mar. 2009.
[10] M. Goodman, "Internet video forecast: Broadband emerges as an alternative channel for video distribution," Yankee Group, Boston, MA, 2006.
[11] R. Govindan and V. Paxson, "Estimating router ICMP generation delays," in *Proc. PAM*, Mar. 2002, pp. 1–8.
[12] "Subjective video quality assessment methods for multimedia applications International Telecommunication Union, Rec. ITU-T P.910, 1999.
[13] "Opinion model for video-telephony applications," International Telecommunication Union, Rec. ITU-T G. 1070, 2009.
[14] R. Jain, "Quality of experience," *IEEE Multimedia*, vol. 11, no. 1, pp. 95–96, Mar. 2004.
[15] S. Kandula, D. Katabi, B. Davie, and A. Charny, "Walking the tightrope: Responsive yet stable traffic engineering," in *Proc. ACM SIGCOMM*, Philadelphia, PA, Aug. 2005, pp. 253–264.
[16] S. Kanumuri, P. C. Cosman, A. R. Reibman, and V. A. Vaishampayan, "Modeling packet-loss visibility in MPEG-2 video," *IEEE Trans. Multimedia*, vol. 8, no. 2, pp. 341–355, Apr. 2006.
[17] R. Keralapura, N. Taft, C. Chuah, and G. Iannaccone, "Can ISPs take the heat from overlay networks?," in *Proc. ACM HotNets III*, San Diego, CA, Nov. 2004, pp. 1–7.
[18] C. Labovitz, R. Malan, and F. Jahanian, "Internet routing instability," *IEEE/ACM Trans. Netw.*, vol. 6, no. 5, pp. 515–528, Oct. 1998.
[19] R. Mahajan, N. Spring, D. Wetherall, and T. Anderson, "Inferring link weights using end-to-end measurements," in *Proc. ACM SIGCOMM Internet Meas. Workshop*, Marseille, France, Nov. 2002, pp. 231–236.
[20] R. Mahajan, N. Spring, and T. Anderson, "Understanding BGP misconfiguration," in *Proc. ACM SIGCOMM*, Pittsburgh, PA, Aug. 2002, pp. 3–16.
[21] A. Markopoulou, F. Tobagi, and M. Karam, "Assessment of VoIP quality over Internet backbones," in *Proc. IEEE INFOCOM*, New York, NY, Jun. 2002, vol. 1, pp. 150–159.
[22] K. Papagiannaki, S. Moon, C. Fraleigh, P. Thiran, F. Tobagi, and C. Diot, "Analysis of measured single hop delay from an operational backbone network," in *Proc. IEEE INFOCOM*, New York, NY, Jun. 2002, vol. 2, pp. 535–544.
[23] V. Paxson, "End-to-end routing behavior in the Internet," *IEEE/ACM Trans. Netw.*, vol. 5, no. 5, pp. 601–615, Oct. 1997.
[24] M. H. Pinson and S. Wolf, "A new standardized method for objectively measuring video quality," *IEEE Trans. Broadcast.*, vol. 50, no. 3, pp. 312–322, Sep. 2003.
[25] S. Savage, T. Anderson, A. Agrawal, D. Becker, N. Cardwell, A. Collins, E. Hoffman, J. Snell, A. Vahdat, G. Voelker, and J. Zahorjan, "Detour: A case for informed Internet routing and transport," *IEEE Micro*, vol. 19, no. 1, pp. 50–59, Jan. 1999.
[26] N. Spring, R. Mahajan, and T. Anderson, "Quantifying the causes of path inflation," in *Proc. ACM SIGCOMM*, Pisa, Italy, Sep. 2006, pp. 113–124.
[27] L. Subramanian, I. Stoica, H. Balakrishnan, and R. Katz, "OverQos: An overlay based architecture for enhancing Internet QoS," in *Proc. USENIX NSDI*, San Francisco, CA, Mar. 2004, p. 6.
[28] R. Teixeira, K. Marzullo, S. Savage, and G. Voelker, "In search of path diversity in ISP networks," in *Proc. ACM IMC*, Miami, FL, 2003, pp. 313–318.
[29] M. Venkataraman, M. Chatterjee, and S. Chattopadhyay, "Evaluating quality of experience for streaming video in real time," in *Proc. IEEE GLOBECOM*, Honululu, HI, Dec. 2009, pp. 1, 6.
[30] M. Venkataraman and M. Chatterjee, "Case study of Internet links: What degrades video QoE?," in *Proc. IEEE GLOBECOM*, Miami, FL, Dec. 2010, pp. 1–6.
[31] J. Wang, D. Chiu, and J. Lui, "A game–theoretic analysis of the implications of overlay network traffic on ISP peering," *Comput. Netw.*, vol. 52, pp. 2961–2974, Jul. 2008.
[32] "AS number lookup utility tool," 2009 [Online]. Available: http://www.enc.com.au/itools/aut-num.php
[33] "King: A tool to estimate latency between any two Internet hosts, from any other Internet host," 2002 [Online]. Available: http://www.mpi-sws.org/~gummadi/king/
[34] "PlanetLab," [Online]. Available: http://www.planet-lab.org/
[35] "Undns tool, router hostname to location decoder," 2010 [Online]. Available: http://www.scriptroute.org/source/
[36] "The 'RouteViews' project at University of Oregon 2009 [Online]. Available: http://www.routeviews.org/

Mukundan Venkataraman (M'10) received the B.E. degree in computer science and engineering from Visveswaraiah Technological University, Karnataka, India, in 2003, and is currently pursuing the Ph.D. degree in electrical engineering and computer science at the University of Central Florida, Orlando.

He has been a Research Staff Member with the Pervasive Computing Team at SET-Labs, Infosys Technologies, Ltd., Banglore, India. His research interests include networked multimedia systems and social networks.

Mr. Venkataraman is a recipient of the Presidential Doctoral Fellowship at the University of Central Florida. He won the Best Student Paper Award in the IEEE Asia–Pacific (Region 10) at Penang, Malaysia, in 2003, and the Best Student Paper awards by IEEE Bangalore (B. R. V. Varadan Paper Award) consecutively in 2002 and 2003.

Mainak Chatterjee received B.Sc. degree (Hons.) in physics from the University of Calcutta, Kolkata, India, in 1994, the M.E. degree in electrical communication engineering from the Indian Institute of Science, Bangalore, India, in 1998, and the Ph.D. degree in computer science and engineering from the University of Texas at Arlington in 2002.

He is currently an Associate Professor with the School of Electrical Engineering and Computer Science, University of Central Florida, Orlando. His research interests include economic issues in wireless networks, applied game theory, resource management and quality-of-service provisioning, ad hoc and sensor networks, CDMA data networking, and link-layer protocols.

Dr. Chatterjee serves on the executive and technical program committees of several international conferences.