Control-Plane Observability Through Data-Plane Performance

Gomathi Ramachandran
Amazon Web Services
Networking Science
Seattle, USA
ramgomat@amazon.com

***Abstract*—Traditionally data-plane measurements have been used to understand application performance and to detect specific impairments with high confidence. Control plane effects on data-plane performance were often incidental findings, especially for operational measurements in traditional IP networks where highly multiplexed streams were serviced by higher speed, highly protected, optical circuits. As we move to larger and more dynamic loads and to SDN networks where the routing logic is moved further away from the data-plane, it becomes more important to understand and use the connection between control-plane effects on data-plane performance in addition to forwarding plane impairments to manage network performance. We illustrate this connection with examples of how some common control-plane changes affect data-plane metrics. These observations motivate improvement in control-plane protocols.**

***Index Terms*—Wide area networks, Traffic Engineering, Data-plane, Control-plane, Active Measurements, Latency**

I. INTRODUCTION

Traditional Internet Service Provider (ISP) networks have highly multiplexed loads due to the aggregation of smaller loads from customer premises into larger core links. Fixed customer connections to the network also made these loads highly static in terms of geography. As we moved to support cloud loads, we see that cloud loads can be more dynamic, and less multiplexed. The larger loads are dynamic in time as single instances can produce highly bursty demands with little ingress smoothing. They are also more dynamic in a spatial sense, as the cloud enables applications to be rapidly spun up in different regional data centers, shifting cross-network traffic to different geographical areas. This changes the assumptions under which Traffic Engineering (TE) was originally designed – to move multiplexed traffic under failures or traffic events to uncongested paths, to make the network routing congestion-aware. As the assumptions have changed, so has the frequency and duration of TE use on the network. Since cloud burstiness is ubiquitous, demand changes are common causing frequent path optimizations to adapt to these changes. The result is that while in the past TE worked well for the over-provisioned, multiplexed ISP network, it now behaves less optimally with the new loads and this can change the customer experience.

Loads that contain repeated short bursts compared to path reoptimization times for congestion avoidance can cause repeated path changes, or path churn. If the bursts are in a tunnel with steady-state traffic, the path movement does not benefit the bursty traffic and will disrupt the smoother traffic if the possible paths are of sufficiently different latency. Separately we find that under constrained capacity conditions, path reoptimizations can take multiple minutes during which congestion can form. In a cloud-served WAN the path optimizations are frequent due to the changing loads, and could lead to a higher probability of poor tunnel packing under constrained conditions.

Our paper will discuss the changes in underlying assumptions for the use of traffic engineering in the Wide Area Network (WAN) as networks have moved from traditional ISP loads to cloud loads. We discuss the way congestion-aware networks use overlay routing mechanisms to steer traffic away from congested paths to optimize performance in some dimension. We use a representative Multi-Protocol Label Switching (MPLS) with Traffic Engineering (TE) network with Resource Reservation Protocol (RSVP) to explain the processes involved control-plane decisions. We then show examples of this behavior from a cannonical network and connect these to cases that might be seen in a real network. This set of protocols optimizes for low-loss at the expense of longer end-to-end latency. We therefore posit that using data-plane measurements of latency can provide visibility into these control-plane decisions.

We suggest that data-plane information, properly understood, can be used to assess control-plane stability and provide input directly to operators and overlay SDNs to improve traffic placement outcomes. As network controllers move further away from the forwarding nodes, this feedback from the data-plane can be a critical part of maintaining stability in a controller-routed network. We would also like to motivate vendors to revise existing TE protocols to account for this dynamic traffic behavior.

II. TRADITIONAL AND CONGESTION-AWARE IP NETWORKS

Traditional IP networks are built on connectionless Internal Gateway Protocols (IGP) like OSPF [1] or IS-IS [2]. These older protocols are congestion unaware – all traffic is sent down the shortest cost path without regard for the amount of traffic the path can transmit. These networks handled failures using a combination of (1) maintaining high reliability for

Amazon Web Services funded this work.

the links and spans of the network via optical protection or by using high reliability layer 1 networks, (2) extensively diversifying the physical link paths so as not to lose large amounts of capacity under any failure and (3) over-engineering capacity on the next shortest path to carry traffic under a rare failure. This forces a high cost-burden on the network – highly-reliable layer 1 facilities or protected facilities can be close to double the cost of non-protected facilities. Similarly diverse physical links are complicated to design (to avoid all common points of failure), can lead to high-latency paths and can be expensive to procure as the diverse path is often routed around geographic features. Engineering the network so that the second best path can take all of the traffic under failure leads to large over-provisioning of expensive circuits that are only used under failure, leading to inefficient use of capacity. Ultimately, these methods lead to high cost networks that raise the cost to customers.

New congestion-aware protocols were developed to improve network cost-efficiency without reducing customer performance. Modern congestion-aware IP networks use TE to manage sudden bursts of demand or to deal with failure-induced capacity constraint. An example of a network using this formalism is a Multi-protocol Label Switching (MPLS) [3] network. This type of network often uses Resource Reservation Protocol (RSVP) [4] and Traffic Engineering (TE) to redirect traffic to available non-congested paths with minimal loss under changes in demand. TE reoptimization [5] is used to groom the tunnel path when the tunnel bandwidth changes or a better path is possible. In addition, for link failures, Fast Reroute (FRR) [6] was designed in conjunction with MPLS-TE to protect flows from loss during convergence by using pre-computed bypass paths.

These new protocols caused a change in the way IP networks were managed, instead of planning for carrying all of the traffic under failure on the next-best path one could use longer paths that had available capacity. This reduced the cost of the network but also resulted in dropping less traffic as compared to waiting for the shortest path to converge or finding capacity to carry all of the demand. Conversely, however, paths were less deterministic when TE was in use, as paths other than the next-best path are used, and therefore latency was also less deterministic.

III. CHANGING ASSUMPTIONS: INTERNET SERVICE PROVIDER VS. MODERN WAN NETWORK

Traditional IP networks were designed for enough capacity on the shortest paths to meet demand. Internet Service Provider (ISP) backbones had a high level of multiplexing of demands to the core network – traffic enters on lower bandwidth links, is aggregated, then sent to the much larger bandwidth core links. In the early 2000s customers bought capacity in T1 (1.54 Mbps), T3 (45 Mbps), OC3 (155 Mbps) or even OC-48 (2488 Mbps) increments while the backbone had links of 40-100 Gbps. This aggregation multiplexed the incoming traffic. Forward-looking capacity planning took these multiplexed demands and scaled them upwards for forecasts. As the time to acquire or build fiber, light the fiber and provision router ports could take six to eight months, planning horizons were long. In addition, customer applications were unlikely to sustain large flows as computing bandwidth was still small and the link speed to the end customer was smaller than the core link bandwidths, hence customer demands were smoothed. It was during this time that congestion-aware protocols came into existence. The underlying assumption for these protocols was that they would be used under failures to move traffic to uncongested paths for short periods of time.

However, as the cloud has become a dominant player for enterprise and business applications cloud architecture has changed some of the underlying assumptions for WAN networking. The first is that data center computational improvements allow customers to sustain larger loads. The new Amazon Elastic Compute Cloud (Amazon EC2) C7g instances, powered by the latest generation AWS Graviton3 processors offer up to 200 Gbps of network bandwidth of which 100 Gbps can be used for cross-region transport [7]. These new instances are ideal for running the most demanding network-intensive workloads, such as network virtual appliances, data analytics, and CPU-based artificial intelligence and machine learning (AI/ML) inference. The cloud also allows customers to switch their applications from one geographic location to another in minutes, far more dynamically than moving a physical connection. The use of non-blocking fabrics has allowed these dynamic flows to directly enter the network with little ingress smoothing. The type of events that the cloud serves are also larger and more dynamic. As an example, for the Warner Bros. Discovery Channel presentation of the Olympic games (2020 and 2022), AWS streamed this data to 175 million users over 50 countries [8]. Scalability was essential for the Olympic Games coverage because Warner Bros. Discovery needed a solution that could stream dozens of simultaneous video feeds. Additionally, because live sports can have weather delays or run long, Warner Bros. Discovery needed the flexibility to scale based on immediate needs rather than a schedule.

Thus, cloud loads are far more dynamic, making traffic placement a more dynamic problem. As fiber acquisition and deployment still take four to six months to deployment, the timescale is far longer than changes in traffic demands and it is not reasonable to scale up physically for these demands. Instead, TE is now used to place dynamic demands regularly, instead of being used for failures and rare events. This has changed the underlying assumption for TE use.

A. *Effect of Changing Assumptions*

TE expected that the state of the network was traffic-smoothed with abrupt changes due to failures that occur rarely. TE was added to improve the timescale of mitigation from the time taken to bring up the failed circuit (days) or change the IGP weights manually (hours) to automatically finding the shortest uncongested path through the network in minutes. In the cloud-served WAN network, while failures do occur, the majority of Label Switched Path (LSP) changes are due to re-optimization of the tunnels due to tunnel bandwidth changes. The decision to re-groom the tunnel is based on the traffic in the tunnel (estimated by the average over a historical look-back period, often sixty seconds), the bandwidth adjustment threshold (a minimum difference that would trigger re-computation of path) and the auto-bandwdth timer (usually set to five minutes). The timescales of TE are not conducive to elastic loads where burst duration could be much less than five minutes. These timescale differences lead to unexpected path behavior in the modern network. While early observations of MPLS on traditional networks showed sub-optimal behavior of tunnel placement with respect to latency [9] and suggested improvements by optimizing TE parameters, with modern loads, this effect is intensified and a static approach is unlikely to improve behavior.

The TE protocol is also stateless, hence decisions made in placing the tunnel can repeat under bursty traffic common to the new traffic loads. This can happen at a higher frequency than expected – causing fluctuations in latency potentially affecting latency-sensitive applications.

TE has difficulty placing the LSP tunnels (bin-packing) with these changing loads under constrained conditions. Even if there is enough bandwidth to place all tunnels, under constrained conditions it may take more than one round of re-optimization of the tunnels leading to congestion-induced loss and latency. The fact that we need to do this re-optimization more often leads to an increased likelihood of observing poor behavior.

Finally the extended use of TE affects network planning and monitoring. One, this masks capacity demand increases that were previously easily trended by utilization; two, since most congestion-aware protocols prioritize loss, latency is the variable that shows the effects of these protocols on network behavior. Traditional methods for network design use traffic carrying as the behavior to be optimized, the continual use of TE means that latency optimization under failure or load becomes a more important planning goal.

IV. PRIOR WORK

Operators have used data-plane observations to manage the network and to drive improvements often incidentally seeing control-plane effects within the data. Scientists have used data-plane observations to examine control-plane effects on applications and the data-plane. Paxson used traceroutes to show routing loops in traditional IP networks [10]. Early work by Ramachandran *et al* used the data-plane measurements to compare flow-healing under OSPF as compared to MPLS-TE with FRR [11], Rada *et al* used latency to show path change in the AWS network [12], as did Pucha *et al* across networks (BGP and IGP) [13]. The last two studies correlated the change in latency to path changes in the network. Rada *et al* showed that these path changes were associated with TCP packet reordering [12] and other application impairments. While there is precedence for using data-plane observations to infer control-plane changes this is not the primary reason that measurements are designed for operational networks, they are designed for detection of data-plane impairments. We intend to show that there are specific control-plane phenomenon that cause additional data-plane impairments in congestion-aware networks. This understanding can be used to help remediate the response to these effects either by operators or by customers using an overlay SDN controller. This also motivates vendors to improve their protocols to deal with traffic dynamism more gracefully, either internally or through standardization in the networking community (RFCs).

V. MPLS-TE WITH RSVP AND FRR PRIMER

The network routing system in most operational congestion-aware networks is based on two distinct layers of routing. The underlay routing is based on the Dijkstra shortest path algorithm to build a consensus for the shortest paths in the network. These paths define the best next-hop for a given packet. Protocols such as Open Shortest Path First (OSPF) or Internal System to Internal System (IS-IS) use messages from one router to another that advertise their directly connected links (link adjacencies) and their link costs C_{jk} and withdraw these advertisements when the links are lost. The costs can be based on distance, or other metric that determines the potential of using the link – where a lower cost is preferred. All routers listen to these updates and compute a topological understanding of what paths exist and are of lower cost to reach destinations (based on IP prefixes) in the network. This then is programmed into the Routing Information Base (RIB) and Forwarding information Base (FIB) of the router. In Fig 1 if the costs are lower from A to Z via B then A will send packets that need to reach Z to B as the next hop.

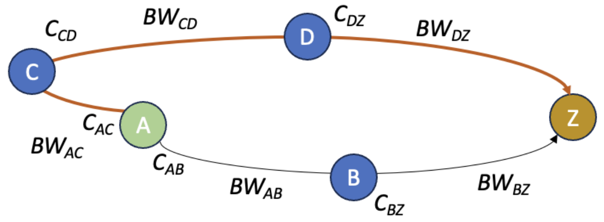

Fig. 1. A representative network with costs and link bandwidth.

Consensus building takes time and depends on the number of links and routers, the size of the routing table (number of prefixes to route). In a traditional ISP network the tables should converge in seconds, IP address aggregation schemes have been developed to bring convergence to sub-second times [14]. However, as the network scales and the number of IP addresses increases this slows down the speed of convergence to multiple seconds with a long tail to convergence for all paths (often ten to fifteen seconds).

While academic proposals have been made to make IGP congestion-aware using dynamically allocated weights (see [15] or [16] for example), this has not been used in operational networks due to the difficulty of path interpretation. Hence MPLS-TE with RSVP was used to add another layer of routing to the network. In MPLS-TE with RSVP each packet is placed into an end-to-end tunnel encapsulating the original packet with an MPLS-TE header based on the exit from the network. The router that sets up the tunnel, called the head-end router is the first MPLS aware router that the packet encounters. The head-end sets up the end-to-end path by looking at the RIB table for the shortest path for the tunnel, then requests the measured tunnel bandwidth $R_{AZ}(t)$ for each link traversed from each router in the path. If all requests are accepted until the end, the tunnel is set up and the packets are sent through the tunnel. If a request fails, the next best path is now queried in a similar fashion. In Fig 1 if the tunnel traffic is too large for the reservable bandwidth on B-Z, with B sending a 'bandwidth unavailable' message back to A then A will try the path A-C-D-Z. Regrooming of the path without failure happens on a regular basis – when the auto-bw timer τ expires (usually set to five minutes) the head-end re-signals the path with a new measured tunnel bandwidth $R_{AZ}(t+\tau)$ if the bandwidth difference is greater than the bandwidth gap set, then the process of searching for a path with enough bandwidth repeats. This timer is a critical determinant for path changes in traffic under sunny-day (no failure) conditions – the addition of demand to the path changes the *routing* at predictable intervals. See Table I for the systems parameters that govern this system.

TABLE I
SYSTEM PARAMETERS

C_{jk} = Unidirectional cost of the link
BW_{jk} = Unidirectional reservable bandwidth of the link
τ = Auto-bandwidth timer
$R_{AZ}(t)$ = Head-end measured bandwidth of tunnel A-Z at time t

For failures new mechanisms were needed to heal traffic quickly – rather than waiting for the entire network to come to a new set of best paths. At the IGP layer the router with the affected link will withdraw the link adjacency for this link which causes all routers to recompute the shortest paths and the best next hop – this time to recompute is called IGP convergence. During IGP convergence not all routers have the same idea of network topology, in IGP networks this causes routing loops that will drop packets as the Time To Live (TTL) of the packets expire from repeated looping (see Casner's early description of 'blender events' [17]). Fast Re-Route (FRR) in MPLS was developed to protect against link failures using a precomputed path that avoids any common physical path for each link to avoid any fiber cuts that would take other potential circuits down. When the link goes down, the router with the failed circuit (point of local repair) shunts traffic to this bypass path (see Fig. 2) in a congestion unaware manner without waiting for a new underlay routing table. This eliminates TTL discards [11]. The router then signals to the head-ends for all tunnels that the circuit is on bypass. The head-ends will try to re-signal a new path and move tunnels if successful. If another path with sufficient capacity is not available it will stay on the bypass until it finds a new path. If bandwidth continues to be unavailable the tunnel that cannot find reserveable bandwidth will stay on the bypass. The process of getting all of the paths off the bypass to the best paths after the failure is called RSVP-TE convergence. The head-end will continue to re-signal paths until it can place the tunnel, this can lead to race conditions during failures as multiple tunnels signal for available capacity.

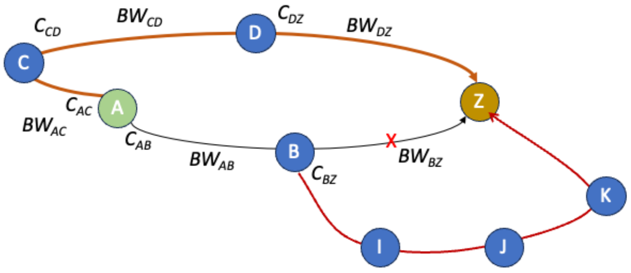

Fig. 2. A Fast Re-Route bypass for link from B to Z avoids all physical paths in common with the link.

VI. EXPERIMENTAL SYSTEM

For this paper we use a representative network that uses IS-IS in the IGP plane with MPLS-TE with RSVP for the overlay routing. The auto-bandwidth timer τ is set to five minutes and reservable capacity for each link is set to 75% of the link bandwidth. Fig. 1 shows how the costs and bandwidth nomenclature. In this network though costs are uni-directional we use the same values for the two directions. There are 13 nodes in the network A, B, C, D, E, F, G, H, I J, K and Z. We study path behavior from A to Z where BW_{AB}, BW_{BZ} are much less than all of the other link bandwidths. In Fig. 3 we show all the nodes and connectivity of the system. There are 4 potential paths between A to Z – (1) A-B-C (lowest cost), (2) A-C-D-Z the next cost, (3) A-E-F-G-H-Z (highest cost) and (4) A-C-D-H-Z (third best choice). In some cases we reduce the available bandwidth of the link by drastically increasing traffic on that link. We take multiple latency measurements from A-Z, these are aggregated to 30 second values. Our bypass path for a B-Z failure is set to path B-I-J-K.

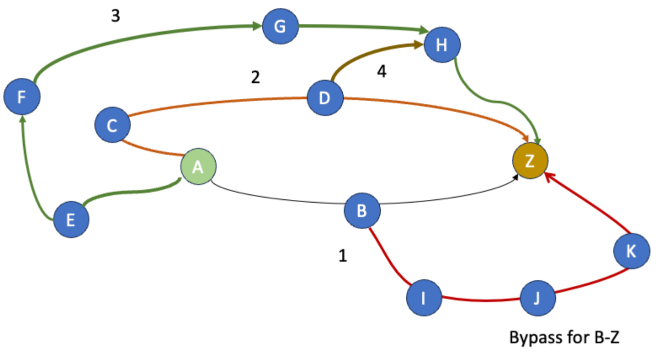

Fig. 3. Toy network nodes and connectivity.

VII. DATA-PLANE VISIBILITY OF CONTROL-PLANE EVENTS

In this section we describe some data-plane impacts that can happen when the control-plane responds to conditions in the network. We then show what this response looks like from data-plane measurements on a representative network. We describe path routing instability, delayed RSVP-TE convergence and traditional congestion. The data-plane response is often clearly distinguishable. If further, one has certain control-plane output one can unambiguously identify cause.

A. Path Churn

Figure 4 shows the mechanism for path changes due to short-lived demand increases. This causes paths to shift from the shortest possible path to a longer latency path to accommodate the increased demand. Every time the auto-bw timer expires the head-end of the tunnel will signal the path again with the new demand – searching for the best path for the larger amount of traffic in the tunnel. If the current path does not have enough bandwidth, the tunnel will take a longer path that has the bandwidth. If the demand increase is short-lived, the tunnel reverts to the shorter path as soon as the demand reduces, which primes the tunnel to move to the longer path when demand increases. This causes repeated changes in latency as the path moves away then back to the shortest path as the demand grows and shrinks. When TE was originally implemented this was recognized as a risk, but it was considered to be a rare one as TE was used to mitigate failures in ISP networks, where demand was highly multiplexed. While tunnels could have been expected to move as demand rose during the peak, this was expected to be stable over hours. The new paradigm of large, highly dynamic flows has been a case where cloud network behavior has changed the original conditions for TE.

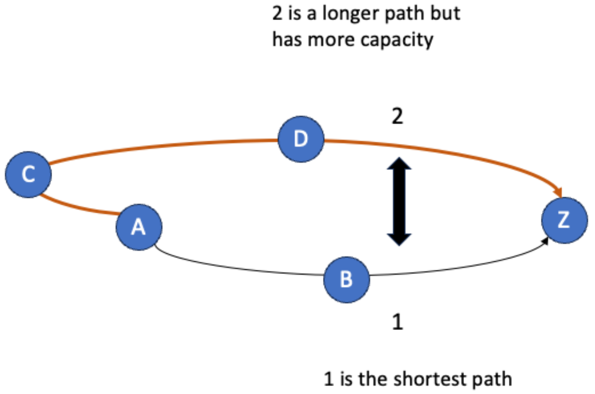

Fig. 4. The path A-B-Z is the lowest cost path, but has less capacity than path A-C-D-Z. An increase in demand can force tunnels to move to the longer path that has more capacity

We observed this in a toy network with two paths between A-Z – the lowest cost path was 14.5 ms round trip (RT) while the other was at 16.5 ms RT. We added bursty traffic in different blocks in one direction of the round-trip (RT) path and measured the round trip latency with test probes. In a later step we added steady-state traffic in both directions to force the path to go to the longer path both ways. When demand increases the traffic goes through the longer path in the direction that there is more traffic but returns via the shortest path (less traffic) initially (see Fig. 5). Since we added bursty traffic the tunnel was able to move back to the original path as demand decreased. Note the switching between one path latency and another. Then we added enough traffic in both directions to cause TE to move to the longer path in both directions. The data is highly aggregated (30s aggregation of latency probe packets) hence there are intermediate values seen as the traffic switches between paths, with a small amount of queuing as the traffic shifts. A close up in Fig. 6 shows that the time for the shorter transitions is about seven minutes – approximately the auto-bandwidth timer, τ, of 5 minutes plus the statistical bandwidth estimation period of a minute at each end, showing the cause is indeed TE moving the path in response to a higher demand on the tunnel. In this case it briefly moves to the shorter path again before moving back to A-C-D-Z, then finally returns to A-B-Z. The ideal behavior would have been for the protocol to recognize that repeated changes happened in the recent past and to have held the tunnel on the higher latency path until it had stabilized. This phenomenon can be observed in real networks when very large but relatively short flows enter the network. As the bandwidth of the tunnel is estimated based on an average of one minute prior, when this large flow arrives, even if it is no longer present the bandwidth estimate for the tunnel is adjusted upwards and moved, meantime the flow has ended and now the bandwidth is reduced. Even a single large flow can do this – for example when there is a short cross-region large log retrieval from an instance that has large capacity. Another example are DDOS flows, this is a new way of disrupting traffic for many flows without targeting interfaces.

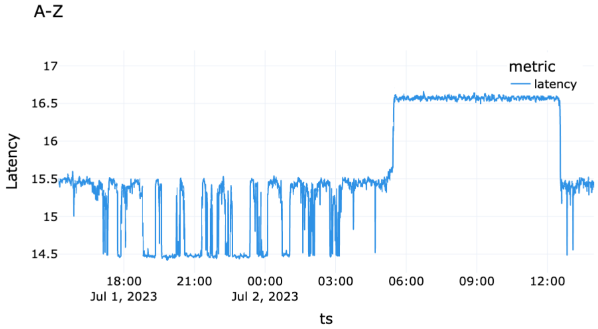

Fig. 5. An example of path churn in a representative network.The shorter path A-B-Z is 14.5 ms, the longer path through A-C-D-Z is 16.5 ms

B. Delayed RSVP-TE Convergence

Another control-plane effect seen in the data-plane happens when convergence is delayed due to path-hunting in RSVP-TE (Fig. 7). This occurs when the tunnels are being re-optimized

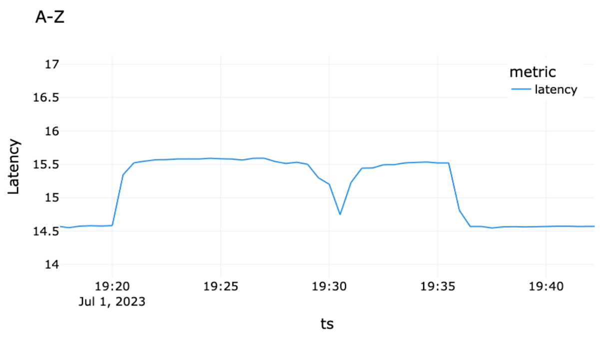

Fig. 6. Close up of one of the increases, shows a flat-topped behavior with a top of about seven minutes which briefly moves to the shorter path then back to the longer path for about six minutes.

(either post-failure from a bypass tunnel or due to increased demand) and cannot find the bandwidth to place the tunnel. Even if there is ultimately enough capacity to place all tunnels, if a large number of tunnels compete for constrained capacity, the head-end tries to place these tunnels on the shortest path and will receive bandwidth unavailable messages from the constrained link on the path based on the packing algorithm for tunnel placement which has a dependence on tunnel size. This results in prolonged path hunting, iterating through a number of paths that lack the required capacity. A symptom of this is the presence of bandwidth unavailable messages for the tunnel from the router on the with insufficient bandwidth on the interface. This is a race condition for the limited amount of head-room needed as tunnels hold on to their current allocation as they seek to place a better tunnel, a 'make-before-break' process. This can also happen for conditons when there are no failures, when loads are moved within the data center, leading to the flows being hashed to different tunnels between the same cross-regional pair. It does not change the amount of traffic offered overall, but each tunnel now has a different bandwidth.

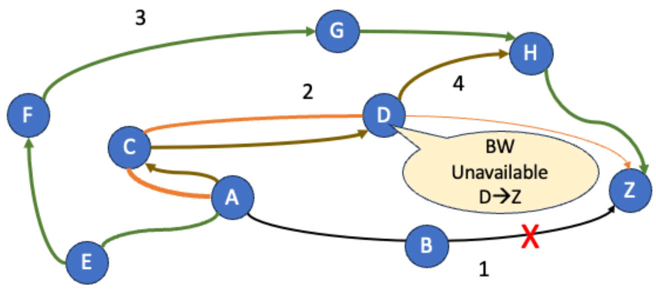

Fig. 7. Delayed RSVP TE convergence is caused by a failure of auto-bw re-optimization resulting in multiple tunnels re-signaling for the same constrained capacity on the next shortest path. In this case A-C-Z is the shortest path but it cannot take the capacity requested, the tunnels try A-C-D-Z but D sends a 'BW Unavailable' message on D-Z. This results in additional time to place all of the re-signalled tunnels.

In the toy network we recreated these conditions by failing a link and then changing the bandwidth of individual tunnels such so that all links were close to 70% utilized (near the reservable bandwidth) but just capable of carrying the traffic when the link B-Z was failed. We observed the presence of bandwidth unavailable messages and saw about 15 ms of added latency in the transition due to congestion before the tunnel moved to the next-best path (15.5 ms). Fig. 8 shows the latency for the A-Z tunnel that is originally on the A-B-Z path but moves to the next best path A-C-D-Z after approximately seven minutes. While not shown, when demand was not high the transition was a smooth step function from the lower latency path to the higher as would be expected from MPLS-TE with RSVP and FRR.

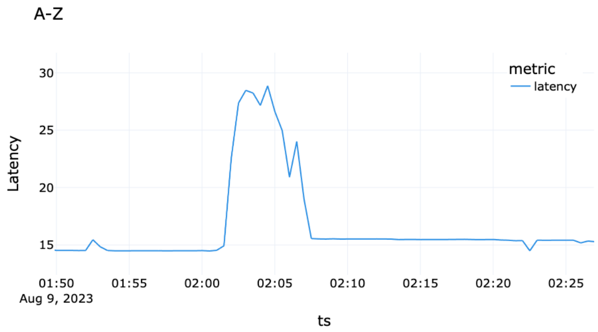

Fig. 8. An example of delayed convergence in a representative network.

C. Congestion

Congestion induced queuing, seen in Fig. 9 is the third example and the most traditional cause of data-plane performance degradation. This causes traditional delay variation (due to fluctuations in queue-depths) as well as longer timescale delay variation – if router buffers are large. It happens when there is more demand coming in than capacity can serve (regardless of the path chosen) and there is a buffer in the router to hold the queued packets until they can be transmitted. This change is usually far larger than a path change as the buffers can be hundreds of milliseconds. Sometimes loss is an accompanying signal of interface congestion, though if there are large buffers, the larger impact to customer applications can be the higher latency. Typically this happens for stub sites with a fixed number of paths, or under catastrophic failures.

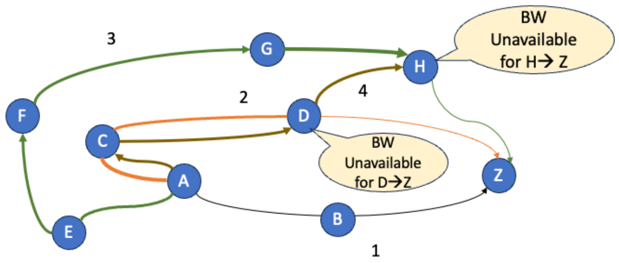

Fig. 9. Congestion-induced latency happens when there is no uncongested path to move to and the traffic stays on the original path despite the lack of capacity. In this case no path to Z can take the demand.

In Fig. 10 we show the results from the toy network where we added traffic to tunnels between D-Z and H-Z so there is no reservable bandwidth to place a tunnel, then added sustained flows with some bursty traffic to the tunnel from A to Z. The tunnel seeks to move to another path but cannot as no path has enough bandwidth. In the end it stays on A-B-Z and congests. In this case we had large buffers available which gives rise to both the large additions and the jaggedness of the latency. Note that hardware can affect this, if we used very small buffers latency would reach the buffer extent and plateau, in the case of small buffers loss becomes the dominant performance metric for congestion again.

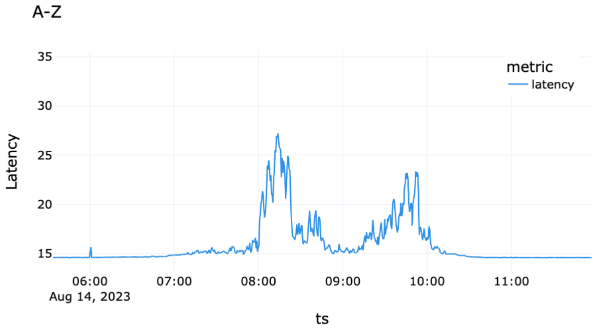

Fig. 10. Congestion-induced latency in a representative network with large buffers, note that the latency is highly variable.

VIII. UNDERSTANDING THE DATA-PLANE RESPONSE TO CONTROL-PLANE ACTIONS

While we have shown some examples of data-plane visibility of control-plane events for a representative MPLS-TE network, this approach is extensible to any congestion-aware network. As SDN controllers become more prevalent the global effects of these controllers is difficult to interpret while simultaneously they could have a larger impact on global traffic [18]. At a higher layer, SDN overlays have also sought to use data-plane performance input to improve overlay decisions [19], but control-plane impairments on the infrastructure are largely opaque to them. Data-plane visibility from an end-to-end view, on the other hand, is available to both network operators and customers.

For operators, the key is to find control-plane signals such as bandwidth unavailable messages or link down events to correlate with data-plane synthetic measurements that are done as representative customer measurements. The data-plane measurements provide representative customer impact of the control-plane. Neither one on its own represents the whole customer experience – synthetic measurements may not cover specific paths or prefixes and will show the most general impacts, control-plane convergence data (e.g., time of last prefix convergence, or the Record Route Object that shares the computed path) only provides the point in time view of the router, and often gives the higher bound (last prefix convergence) or the path without performance attributes. Both are incomplete in themselves, together they provide a powerful tool for tracking performance and frequency of underlying cause. Operationally, one could use these signatures to identify the cause of the issue and to determine the next actions to remediate the situation. For example highly bursty traffic that causes path churn could indicate DDOS attacks that need mitigation. However, some types of true customer bursty load could be due to poor placement of services, for example a large log retrieval for processing across a long cross-region distance. Ideally, the loads should be processed *in situ* with summaries being sent across the WAN. Delayed RSVP-TE convergence on the other hand could indicate insufficient headroom to place traffic and while long term remediation will usually require protocol improvements or placement of additional capacity, in the short term operators can redirect traffic, investigate increases in demand and suggest Internal Gateway Protocol (IGP) weight changes to change the shortest path.

For customers, especially those who use an overlay SDN, understanding the causes of data-plane latency events can improve the overlay response. A simple example would be not to use short-term latency decreases to inform traffic placement if there is evidence of path churn, and instead prefer a higher latency path for stability reasons, especially for non-interactive application workloads. A more complex example would be to change Border Gateway Protocol (BGP) exit routing to another Autonomous System (ASN) if extensive impairments are seen.

For vendors, this should motivate new ways of dealing with significant traffic dynamism in their congestion-aware protocols. We can think of a few major areas of improvement. The first is adding statefulness to the routing protocol, TE needs a dampening behavior to prevent repeated traffic swings, this requires the protocol to remember states for the tunnels. It could do this by maintaining a frequency domain analysis for each tunnel to inform the decision to move the tunnel. TE could also benefit by adding network information to its decisions – for example, taking in the stability of a link being used, instead of just the cost. The third and most difficult area are better bin-packing methods that take into account highly constrained conditions in the network. One idea is to increase the co-ordination of tunnels with the same regional source and destination to prevent race conditions for reserving the same link bandwidth. The difficulty is the enumeration of all possible ways that the network conditions could trigger sub-optimal bin-packing to develop an approximate bin-packing algorithm [20]. It is possible that an algorithm that takes data-plane information as an input might be a better way to improve bin-packing for RSVP-TE convergence.

IX. CONCLUSION

We have talked about how cloud networks have changed the assumptions for the behavior of congestion-aware protocols such as TE in wide area networks. The traffic demands are far more dynamic in cloud networks than was expected when these protocols were designed. At the same time the way congestion-aware protocols respond to failures and traffic demands by minimizing packet loss has made latency a new window for observing the behavior of the control-plane. We have shown in an MPLS-TE RSVP toy network, that these changes can be observed in latency measurements and that these can be shown as a response to control-plane actions to place traffic in a congestion-aware manner. We postulate that operators and customers with overlay SDNs could use active latency measurements to probe control-plane behavior, and by understanding the ways that congestion-aware networks behave, could mitigate this behavior more effectively. We hope vendors are motivated to reduce flow instability by moderating their traffic placement algorithms.

ACKNOWLEDGMENT

The author thanks Dr. Marina Thottan, Jorge Vasquez, Dr. Amogh Dhamdhere and Prof. Mythili Vutukuru for their comments.

REFERENCES

[1] J. Moy, ”OSPF Version 2”, STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, https://www.rfc-editor.org/info/rfc2328.

[2] ISO/IEC 10589:2002 ”Telecommunications and information exchange between systems – Intermediate System to Intermediate System intra-domain routeing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)”. November 2002.

[3] A. Viswanathan, E. C. Rosen and R. Callon, ”Multiprotocol Label Switching Architecture”, IETF Request for Comments 3031, https://www.rfc-editor.org/info/rfc3031, 61 pages, January 2001.

[4] D. O. Awduche, Lou Berger, D. Gan, T. Li, V. Srinivasan, and G. Swallow. ”RSVP-TE: Extensions to RSVP for LSP Tunnels.” RFC 3209, December 2001.565

[5] JP Vasseur, Raymond Zhang, and Yuichi Ikejiri. Reoptimization of Multi-protocol Label Switching (MPLS) Traffic Engineering (TE) Loosely Routed Label Switched Path (LSP). RFC 4736, November 2006.570

[6] A. Atlas, G. Swallow, and P. Pan. Fast Reroute Extensions to RSVP-TE for LSP Tunnels. RFC 4090, May 2005.567

[7] http://aws.amazon.com/ec2/instance-types/c7g/ , http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-network-bandwidth.html

[8] Warner Bros. Discovery Presents the Olympic Games to Millions of Viewers Supported by Innovations Built Using AWS, http://aws.amazon.com/solutions/case-studies/warner-bros-discovery-case-study/

[9] A. Pathak, M. Zhang, Y. Charlie Hu, R. Mahajan, and D. Maltz. 2011. ”Latency inflation with MPLS-based traffic engineering”, In Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference (IMC ’11). Association for Computing Machinery, New York, NY, USA, 463–472. https://doi.org/10.1145/2068816.2068859

[10] V. Paxson. End-to-end routing behavior in the internet. IEEE/ACM Transactions on Networking, 5(5):601–615, 1997.586

[11] G. Ramachandran, L. Ciavattone and A. Morton, ”Restoration measurements on an IP/MPLS backbone: The effect of Fast Reroute on link failure,” 2011 IEEE Nineteenth IEEE International Workshop on Quality of Service, San Jose, CA, USA, 2011, pp. 1-6, doi: 10.1109/IWQOS.2011.5931348.

[12] W. Reda, K. Bogdanov, A. Milolidakis, H. Ghasemirahni, M. Chiesa, G. Q. Maguire, and D. Kostic.” Path persistence in the cloud: A study of the effects of inter-region traffic engineering in a large cloud provider’s network. SIGCOMM Comput. Commun. Rev.,590 50(2):11–23, may 2020.591

[13] H. Pucha, Y. Zhang, Z. Morley Mao, and Y. Charlie Hu. Understanding network delay changes caused by routing events. In Proceedings of the 2007 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems, SIGMETRICS ’07, page 73–84, New York, NY, USA, 2007. Association for Computing Machinery.

[14] P. Francois, C. Filsfils, J. Evans, and O. Bonaventure. ”Achieving sub-second igp convergence in large ip networks. SIGCOMM Comput. Commun. Rev., 35(3):35–44, jul 2005.

[15] H. Tahilramani Kaur, T. Ye, and S. Kalyanaraman,”Minimizing Packet Loss by Optimizing OSPF Weights using On-line Simulation,” Proceedings of 11th IEEE/ACM International Symposium on Modeling, Analysis and Simulation of Computer Telecommunications Systems (MASCOTS 2003), Pages 79-86, Orlando, Florida, October 2003.

[16] A. Basu, A. Lin, and S. Ramanathan. 2003. Routing using potentials: a dynamic traffic-aware routing algorithm. In Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications (SIGCOMM ’03). Association for Computing Machinery, New York, NY, USA, 37–48. https://doi.org/10.1145/863955.863962

[17] S. L. Casner, C. Alaettinoglu, and C. Kuan. ”A fine-grained view of high-performance networking.” NANOG-22, 05 2001.

[18] Z. Xu, F. Y. Yan, R. Singh, J. T. Chiu, A. M. Rush, and M. Yu. 2023. Teal: Learning-Accelerated Optimization of WAN Traffic Engineering. In Proceedings of the ACM SIGCOMM 2023 Conference (ACM SIGCOMM ’23). Association for Computing Machinery, New York, NY, USA, 378–393. https://doi.org/10.1145/3603269.3604857

[19] P. T. Anh Quang, S. Martin, J. Leguay, X. Gong and X. Huiying, ”Intent-Based Routing Policy Optimization in SD-WAN,” ICC 2022 - IEEE International Conference on Communications, Seoul, Korea, Republic of, 2022, pp. 4914-4919, doi: 10.1109/ICC45855.2022.9838293.

[20] C. Barnhart, H. Krishnan, P.H. Vance, (2001). Multicommodity Flow Problems . In: C.A. Floudas, P.M. Pardalos (eds) Encyclopedia of Optimization. Springer, Boston, MA. https://doi.org/10.1007/0-306-48332-7-316