

Control Exchange Points: Providing QoS-enabled End-to-End Services via SDN-based Inter-domain Routing Orchestration

Vasileios Kotronis¹, Xenofontas Dimitropoulos^{1,2}, **Rowan Klöti**¹, Bernhard Ager¹
Panagiotis Georgopoulos¹, Stefan Schmid³



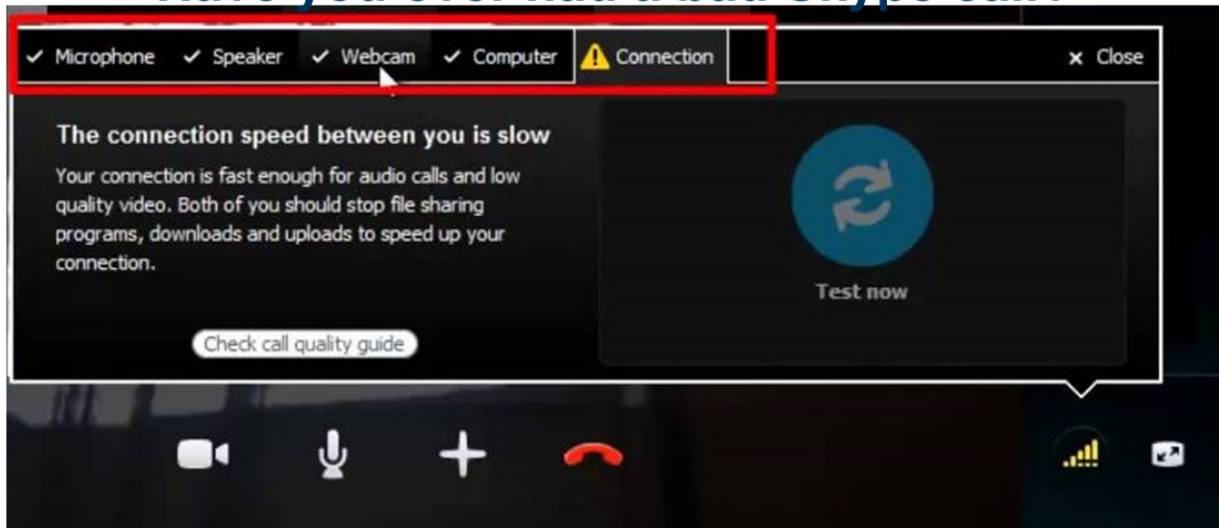
Friday, 21 February 2014

ONS 2014 Research Track

1

- This paper was a joint work with Vasileios, Xenofontas, Bernhard, Panagiotis and Stefan

Have you ever had a bad Skype call?



- Have you ever had a Skype call drop?
- Have you ever had poor quality video and lag?

Would you perform a duet over the Internet?



http://www.internetsociety.org/sites/default/files/pdf/accepted/32_LOLA.pdf

- What do you think about playing a piano duet with someone – over the Internet?
- Careful coordination is required, so latency must be very low.
- This has been implemented – the implementation requires that uncompressed video be sent!

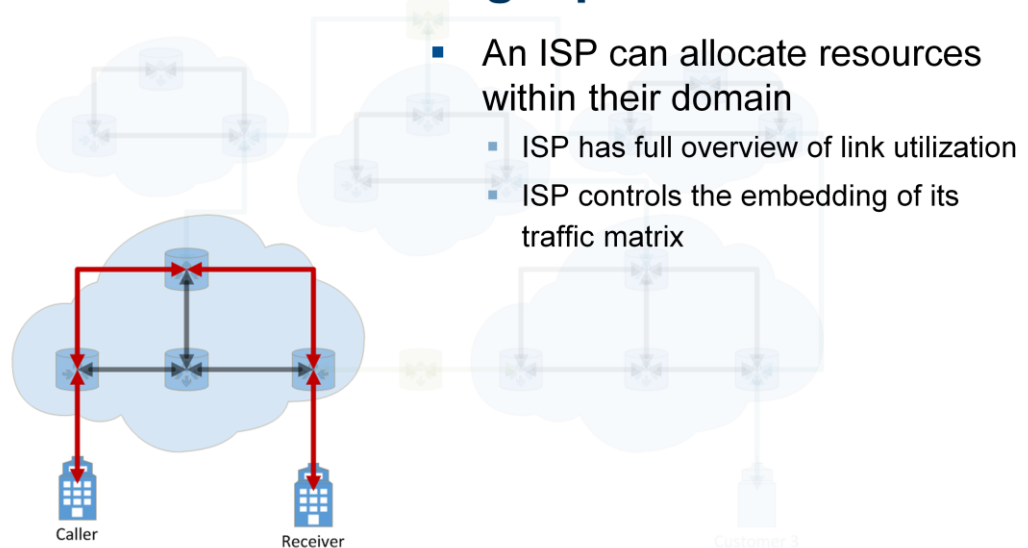
What about having telesurgery?



<http://news.bbc.co.uk/2/hi/science/nature/1552211.stm>


- What about having an operation performed from another continent?
- Not only do we require uncompressed video, but also high resolution.
- Not to mention high availability!

Possible for a single provider



- QoS guarantees are possible if a caller and a receiver have the same provider.
- The provider has the overview of its link utilisations and solutions exist to allow it to embed its traffic matrix to meet guarantees.

What if the endpoints have different providers?



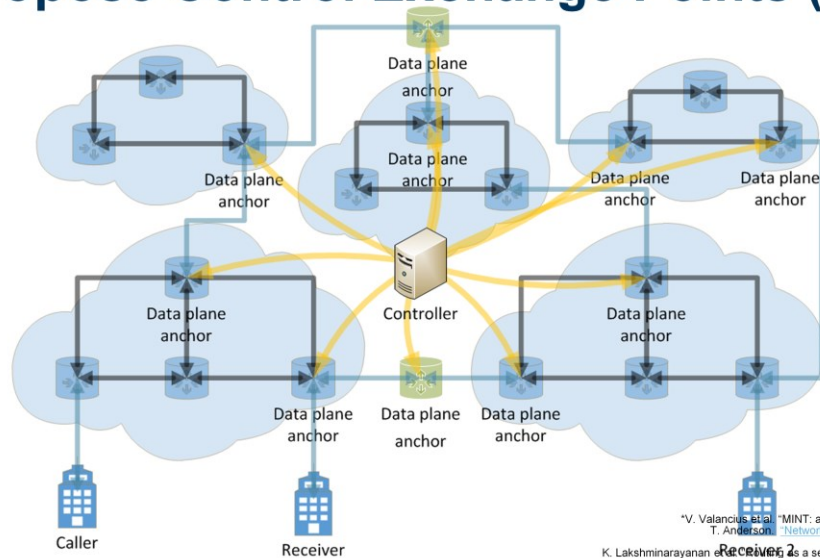
- If a caller and receiver have different providers, then QoS is very difficult with the current Internet.

Inter-domain routing limits us

- No end-to-end guarantees for:
 - Availability
 - Latency
 - Bandwidth
 - ...
- Current inter-domain routing (BGP) does not allow this
 - BGP focuses on *reachability*, not *Quality of Service*
 - We can't replace BGP (easily)

- Between domains, there are no guarantees of bandwidth or latency, or even end-to-end connectivity.
- The current inter-domain routing protocol BGP focuses on reachability rather than QoS.
- Replacing BGP is far from straightforward.

We propose Control Exchange Points (CXPs)




Friday, 21 February 2014

ONS 2014 Research Track

8

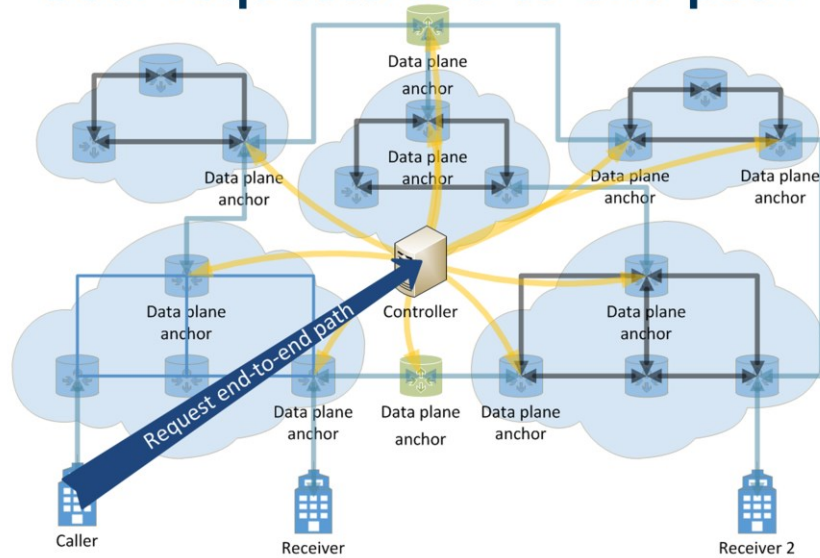
- We propose Control Exchange Points
- Control Exchange Points provide application-specific SDN-based peering.
- Each Control Exchange Point consists one logically centralised controller and many data plane anchors which are operated by the controller
- Data plane anchors are situated on the edge of ISP networks.
- The solution allows for multiple competing CXPs as well as the coexistence with existing inter-domain routing.
- Some of the concepts of this paper have been discussed in prior work.

ISPs announce pathlets




- ISP participating in the scheme announce pathlets to the controller.
- The pathlets connect between data plane anchors and are annotated with performance guarantees.
- The ISPs are paid for the use of the pathlets.

User requests end-to-end path



- A user who wants an end-to-end path with QoS guarantees sends a request with the endpoints and the required guarantees to the controller.

Controller stitches pathlets together




Friday, 21 February 2014

ONS 2014 Research Track

11

- The controller computes several valid paths using the offered pathlets and selects one.
- The remaining paths may be kept as a backup.

Controller monitors guarantees




Friday, 21 February 2014

ONS 2014 Research Track

12

- After establishing the path, the controller continues to monitor the performance of the path via instrumentation at the data plane anchor and possibly at the endpoints.

Controller detects guarantee violations




Friday, 21 February 2014

ONS 2014 Research Track

13

- The controller can detect violations of QoS guarantees, or even a completely failed link.

Controller chooses alternative route



Friday, 21 February 2014

ONS 2014 Research Track

14

- When QoS guarantees are violated, the controller performs a dynamic rerouting using an alternative path.
- The caller and receiver need not notice the change of route.

Best location for CXP data plane anchors?

- Good path diversity
- Maximal coverage of potential users
- Well-connected deployments
- High bandwidth and availability
- Provider neutrality

- The question remains: where to deploy these data plane anchors?
- We need a central location with good path diversity and high bandwidth and availability, but without depending on a single provider.

IXPs have the desired properties!


- Internet Exchange Points are public peering points
- They can have hundreds of providers participating
- They exchange up to Tbps of traffic
- They are independent of individual members

- But what about path diversity?
 - Coverage?

- We propose to use Internet Exchange Points for data plane anchors.
- They have hundreds of participants (DE-CIX has over 600).
- They exchange terabits of traffic every second but yet are not controlled by individual ISPs.

- But what about path diversity and coverage?

Path diversity and coverage



	LINX	DE-CIX Frankfurt	Terremark	AMS-IX
LINX	512	209	51	241
DE-CIX Frankfurt	-	516	47	275
Terremark	-	-	151	48
AMS-IX	-	-	-	629

Euro-IX: <https://www.euro-ix.net/>
CAIDA: <http://www.caida.org/data/overview/>

Friday, 21 February 2014

ONS 2014 Research Track

17

- We investigated path diversity and coverage.
- The plot on the left shows how many IP addresses we can service from a given number of IXPs.
- The green curve is for providers attached directly to IXPs.
- The blue curve is for providers attached to IXPs and their customers over one hop.
- The red line is total number of announced IPv4 addresses.
- We can service a large portion of the Internet address space from just a modest number of IXPs and, if we allow for a single customer-provider hop, virtually the entire Internet.
- On the right you can see a table showing just how many members the largest IXPs have in common.
- With dozens and even hundreds of members connecting between the IXP pairs, the path diversity available is large.

Current and future work

- Simulation
 - Refinement of path embedding algorithm
 - Sensitivity analysis
- Investigate CXP deployment strategy
- Emulation of system
 - SDN northbound APIs between ISPs-IXPs and CXPs
 - SDX* concept suitable for deploying CXP anchors

*N. Feamster et al. "SDX: A Software Defined Internet Exchange", in ONS, (2013).

- We are currently working on the simulation of path embedding, in order to refine the algorithm and perform sensitivity analysis.
- We are investigating deployment scenarios.
- We will be developing an emulation of the working system, in order to test APIs, as well as considering the Software Defined Exchange concept introduced by Nick Feamster et al.

Conclusions

- Interdomain Quality of Service
- Control Exchange Points
 - Data plane anchors deployed in IXPs or other peering points
 - ISPs offer pathlets between data plane anchors
 - Controller stitches these to obtain end-to-end tunnels
 - Controller monitors and reroutes
- Questions?



- In this talk we have proposed a solution for interdomain QoS.
- We presented the concept of Control Exchange Points, which employ SDN to offer dynamically routed end-to-end paths with performance and connectivity guarantees.