


MODΣ

MODE SD-CORE

PURE MATH BACKBONE HITS THE
THEORETICAL LIMIT OF ROUTING
PERFORMANCE

OVERVIEW

Mode offers the world's first wholly autonomous global virtual network, Mode SD-CORE, which provides SD-WAN solutions with an affordable QoS connectivity option. Mode SD-CORE empowers enterprises to instantly create, elastically use, and transparently manage an unlimited number of their own micro-segmented, SLA-backed QoS private networks – at a business internet price point. This white paper examines the inner workings of Mode SD-CORE, and presents the results of extensive tests of the core Mode HALO algorithms, network architecture, and virtual routers.



BACKGROUND

Closed-loop control is a central component in the delivery of most modern dynamical systems, represented by aircraft, automobile engines, spacecraft, and large distributed power grids. Engineers begin the design of these complex systems by modeling the governing differential equations, referred to as the “characteristic equations.” They then apply control systems theory to attain optimal performance using real-time measurements. The application of a feedback controller to a dynamical system is called “closing the loop.” In a block diagram of such a system, the measured outputs (e.g. velocity, altitude, stress, strain) are fed into the controller, which subsequently adjusts a set of control inputs to achieve desired system performance.

Packet-switched routing networks are dynamical systems that to date have not benefited from closed-loop control. Unfortunately, for many decades the derivation of the underlying closed-form solution for the characteristic equations remained elusive. In the absence of preferred closed-loop control, very basic, heuristic protocols like BGP and OSPF were developed to control these networks.

A breakthrough by researchers at Cornell University ended this decades-long quest, revealing the characteristic equations that define any packet-switched network. These researchers then implemented their breakthrough as HALO, the world's first distributed real-time control system for packet switched networks (https://people.ece.cornell.edu/atang/pub/15/HALO_ToN.pdf). Mode was founded by these same researchers, who designed a commercial version of HALO (Mode HALO), and an autonomous global virtual network implementation (Mode SD-CORE).

It is easiest to gain an understanding of the Mode HALO algorithm by examining the results of a series of experiments.

THE FIRST TEST

The National Science Foundation (NSF) supported the original research, and ultimately facilitated an evaluation of the initial implementation of HALO on the NSF network testbed, GENI (Figure 1). In this experiment, the researchers deployed a network across the United States, using the shared compute resources at GENI Points of Presence (PoPs) to create virtual routers (connected as shown in Figure 2). They then used a randomly generated, high-demand traffic matrix to set the communication rates among PoPs, while the routers at each PoP ran HALO. The resulting network path diversity required traffic at each PoP to be apportioned in some non-trivial manner, for optimal bandwidth use and minimized overall network delay in the face of dynamic traffic spikes.

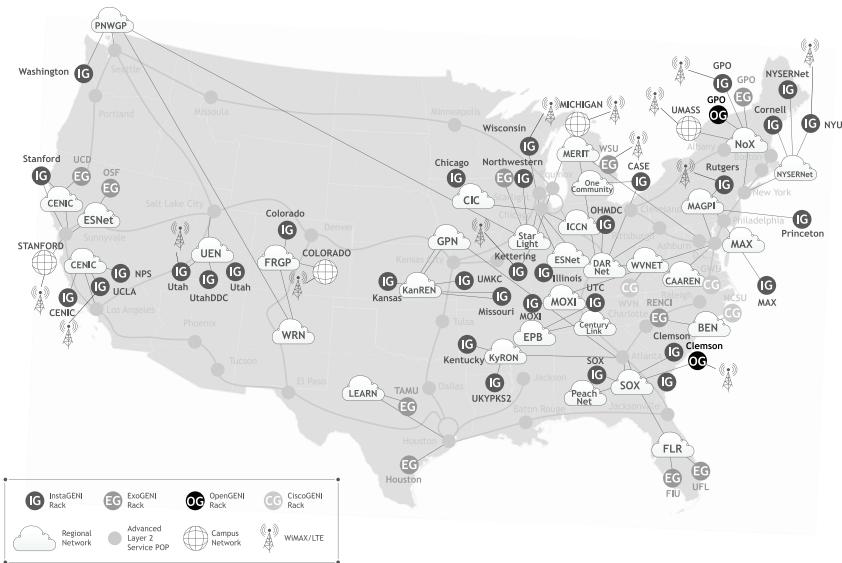


FIGURE 1 – NSF GENI UNDERLAY

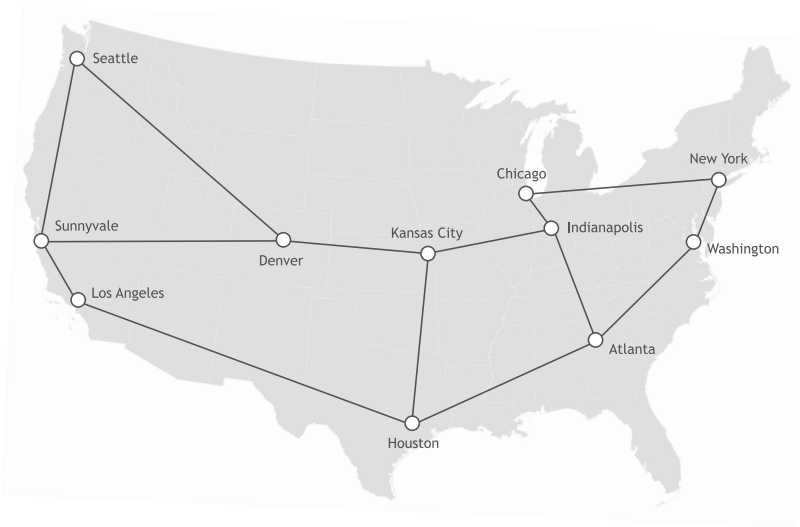


FIGURE 2 – MODE HALO VIRTUAL ROUTER OVERLAY

As shown in Figure 3, HALO supported ~300% the throughput at the lowest possible delay between hosts in New York and Sunnyvale, when compared with the prior state-of-the-art used by network operators to handle traffic spikes. Unlike heuristic protocols, HALO was able to quickly adapt to dynamic traffic changes (Figure 4) without prior knowledge. This inherent dynamism is the key to the performance, flexibility, and reliability of Mode HALO.




FIGURE 3 – MODE HALO SUPPORTS HIGHER THROUGHPUT AND PRESERVES OPTIMAL LATENCIES




FIGURE 4 – MODE HALO RAPIDLY ADAPTS TO DYNAMIC TRAFFIC CHANGES, DRIVING TOWARDS OPTIMAL LATENCY

AT&T SDN NETWORK DESIGN CHALLENGE

A subsequent victory at the AT&T SDN Network Design Challenge further validated the operational advantage revealed by the NSF experiments. The AT&T challenge was to provide an optimal solution on a prototypical carrier network (Figure 5) in the face of rapidly rising dynamic traffic demand, with an emphasis on efficiency and cost-effectiveness. The Mode team leveraged Mode HALO to deliver a near-optimal solution in approximately thirty seconds. The next-closest competitor required nearly a day. Mode HALO had proven unique in its ability to handle the traffic demands and scale of one of the world's largest networks in an efficient and sustainable manner.

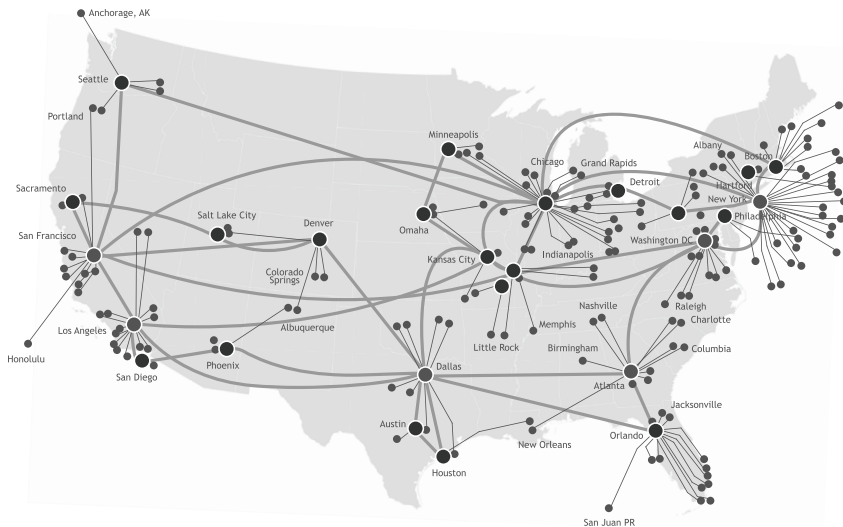


FIGURE 5 – AT&T US NETWORK MAP

DEPLOYING AT SCALE

The next test was designed to confirm the ability of Mode HALO to scale to hundreds of locations, while handling Tbps of traffic. The test setup used an actual customer network topology, with over 1,000 routing nodes, and an asymmetric link capacity of approximately 1 Gbps. The first set of tests introduced traffic changes designed to overload individual link capacity. Even at this scale, Mode HALO maintained its ability to converge and adapt rapidly (Figure 6 shows link utilization vs. time, revealing the response of the test setup to multiple, impactful traffic changes).




FIGURE 6 – MODE HALO QUICKLY ADAPTS AT SCALE BY EFFICIENTLY USING >90% OF NETWORK CAPACITY, MINIMIZING LINK SATURATION

Test results from all cases confirmed that even in large-scale networks with widely varying (and unplanned) traffic changes, Mode HALO was able to adapt and respond in real time, optimizing link utilization and system throughput.

Another set of tests was used to reveal differences between Mode HALO and best-practices Shortest Path Routing among a varying set of PoPs with a theoretical maximum throughput of 40 Gbps. A uniformly random traffic pattern was generated and input to both the Mode HALO and Shortest Path Routing solutions.




FIGURE 7 – SHORTEST PATH LINK UTILIZATION AT VARIABLE LOADS

In the case of the Shortest Path test, the network was able to achieve 12-13 Gbps (33%) of the total network theoretical capacity of 40 Gbps. Link utilization for this configuration is shown in Figure 7. This figure highlights the problem areas encountered by the Shortest Path algorithm, including many unused links, and visible choke points which throttled traffic (please note that in Figure 7 and in the following Figure 8, for ease of exposition, only a subset of routing locations from the core of the network are shown).

Figure 8 shows link utilization for the same test, this time using Mode HALO. Mode HALO delivered almost 36 Gbps of traffic – 90% of theoretical capacity, and more than 2.7x the throughput of the Shortest Path implementation. The numbers and colors in the network graph show the percentage splits of traffic flows between each router node. In the Mode HALO case, only 3 of the network links remain unused. In a large-scale network, Mode HALO provides a consistent performance improvement over the prior state-of-the-art due to its ability to use network capacity optimally.




FIGURE 8 – MODE HALO LINK UTILIZATION AT VARIABLE LOADS

BUILDING A GLOBAL NETWORK: MODE SD-CORE

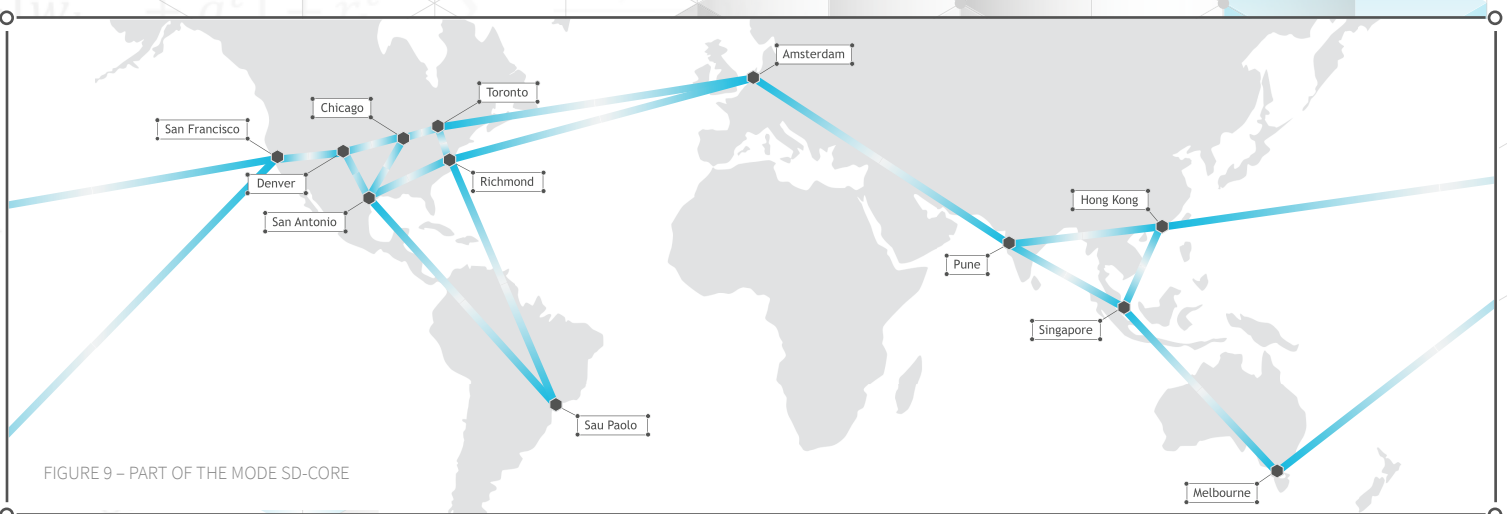
The logical next step was to build a global network controlled by Mode HALO, leveraging its inherent efficiencies to offer SLA-backed reliability at actual business-internet pricing. Unfortunately, legacy routers are not designed for dynamic control algorithms. In the face of this constraint, the Mode team accelerated the development of virtual routers used in the aforementioned experiments. The result of this design effort was a virtual, carrier-grade router that could be dynamically provisioned, modified, and controlled. Next, the Mode team leveraged cloud resources around the world to build the routers, connecting them together to realize the world's first, completely autonomous, global, virtual network (Figure 9): Mode SD-CORE.

Each virtual router in the Mode SD-CORE network is carved out of the resources available in a standard blade server. Each router consumes two CPU cores – one for the containerized control plane, and one for the Open Virtual Switch (OVS)-based data plane.

Currently, a single virtual router using a small fraction of the available server CPU capacity delivers over 6 Gbps of throughput. Extrapolating to the typical multi-core server used in deployments, Mode achieves line rate throughput from the NICs (typically 40 Gbps).

Every virtual router in the Mode SD-CORE is able to support approximately 100K flows/sec, while storing 100K flow rules in the flow table – sufficient for most enterprise private networks. For more demanding enterprise users, Mode has leveraged architectural innovations that guarantee high availability and multi-tenancy with Mode Router, and developed techniques to combine these individual virtual routers into larger routers which can scale to match any demand.

Another Mode SD-CORE differentiator is an orchestration system which enables the spin up of a secure, global network within sixty seconds. The combination of Mode HALO and this orchestration system allows customers to change their network capacity within seconds, or spin up entirely new networks in the same time – all with zero perturbation of existing networks. The Mode SD-CORE portal supports new network creation and deployment, holistic network management, and real-time performance analysis.




SUMMARY

Prior to the Mode HALO breakthrough, network state of the art meant you could have reliability, or affordability – but never both – in a single network. The multiple technological breakthroughs outlined in this white paper represent a massive paradigm shift in network management. Their underlying scientific achievement gives Mode the unique ability to deliver an elastic, secure, SLA-backed, cloud-friendly QoS network – Mode SD-CORE – that offers greater than 99.99% availability, at the price of a regular business internet connection.

ABOUT MODE

Mode operates the world's highest-performing software-defined core network (SD-CORE), built around a real-time network control breakthrough. Mode was co-founded by two Cornell computer scientists widely recognized for their discovery of the characteristic equations that define modern packet-switched networks, and their subsequent implementation of a mathematically optimal routing system, Mode HALO. Mode SD-CORE powers SD-WAN, Cloud Access, Unified Communications (UC), and Ultra Low Latency applications. Mode SD-CORE enhances any enterprise WAN by providing reliable, QoS connectivity in combination with cloud elasticity and business-internet pricing. Mode is based in San Francisco, and backed by GV, NEA and the NSF.



APPENDIX

Mode SD-CORE Data Plane Throughput Scaling

Mode individual server/PoP tests, as well as large-scale throughput tests, support the following data plane scale factors:

- 100s of Tbps of network capacity
- 1000+ router nodes managed
- Individual server scale of 40 Gbps

Mode SD-CORE System Response

Using the aforementioned tests, Mode has demonstrated the following system responsiveness:

- <150 ms network response to traffic changes resulting in real-time ability to reroute flows in less than TCP session timeout intervals
- <1000 ms overall network convergence to routing optimality after significant traffic changes or catastrophic breakage

Mode Router

Mode tests demonstrate the ability to instantiate and tear-down global networks in fewer than 60 seconds.

Mode SD-CORE Network Utilization

Mode achieves near-optimal link utilization. In the test cases above (for both national and global WANs) Mode Router can deliver ~3x throughput with minimal latency by using all available links at high capacity. Mode SD-CORE has consistently demonstrated the ability to run at steady state link utilization 200%-300% greater than that of the the closest competitor.