

Why choose when you can have both: Programmable data planes meet programmable optics

Chris Misa¹, Matthew Nance-Hall, Reza Rejaie,
Walter Willinger, Ramakrishnan Durairajan



<https://onrg.gitlab.io>

1. cmisa@cs.uoregon.edu

Motivation

Networking for AI compute does not end at the data center.



Cloud data centers

“Daily API calls [to OpenAI] surpassed 2.2 billion in 2025, **up from 1.3 billion** last year.”

- <https://sqmagazine.co.uk/openai-statistics/>


Envisioned usecases:

- Hyper-personalized content;
- Ask LLMs questions based on smartphone camera;
- XR-base realtime interaction...

- <https://www.ericsson.com/en/reports-and-papers/mobility-report/articles/genai-impact-on-mobile-network-traffic>


Motivation

Networking for AI compute does not end at the data center.




Motivation

Networking for AI compute does not end at the data center.




Motivation

Networking for AI compute does not end at the data center.





The Challenge

CSP WANs carry AI compute + wide range of other applications.




The Challenge

CSP WANs carry AI compute + wide range of other applications.




The Challenge

CSP WANs carry AI compute + wide range of other applications.




PDP-Only?

Select and drop low-prio traffic with prog. data plane (PDP)?




PDP-Only?

Select and drop low-prio traffic with prog. data plane (PDP)?




OTP-Only?

Use optical topology programming (OTP) tune link capacity?




OTP-Only?

Use optical topology programming (OTP) tune link capacity?




PDP + OTP?

Carefully blend runtime capabilities of PDP and OTP.



PDP + OTP: ShapeShifter


Carefully blend runtime capabilities of PDP and OTP.



PDP + OTP: ShapeShifter

Key challenge:
Unify PDP and OTP
programming
models.


PDP + OTP: ShapeShifter



Initial Evaluation

Dataset: DARPA SEARCHLIGHT

Topology: Heavy dumbbell




Task: MCF-based TE

Simulate linear scaling
demand-factor for all traffic
classes.

Metric: High-prio throughput

Initial Evaluation


Dataset: DARPA SEARCHLIGHT
Topology: Heavy dumbbell



Task: MCF-based TE


Simulate linear scaling
demand-factor for all traffic
classes.

Metric: High-prio throughput



Initial Evaluation

Dataset: DARPA SEARCHLIGHT
Topology: Heavy dumbbell




Task: MCF-based TE

Simulate linear scaling
demand-factor for all traffic
classes.


Metric: High-prio throughput

Baseline MCF and PDP don't
scale beyond ~2X demand



Initial Evaluation

Dataset: DARPA SEARCHLIGHT
Topology: Heavy dumbbell




Task: MCF-based TE

Simulate linear scaling
demand-factor for all traffic
classes.


Metric: High-prio throughput

OTP-only approach is limited by
inability to separate traffic classes.



Initial Evaluation

Dataset: DARPA SEARCHLIGHT
Topology: Heavy dumbbell




Task: MCF-based TE

Simulate linear scaling
demand-factor for all traffic
classes.


Metric: High-prio throughput

Combined PDP+OTP scales up to
7.4X with no throughput degradation.



Takeaway and Next Steps

- Delivering high-quality AI to end-users requires innovation along the entire network path.
- We need:
 - Data / optical plane coordination;
 - Control plane abstractions unifying forwarding and topo.;
 - Large-scale testbeds.



Thanks! Questions?

Chris Misa

cmisa@cs.uoregon.edu

chrismisa.com



<https://onrg.gitlab.io>



Internet Society
Foundation