

One Primitive for all, all for one: Enabling Dynamic Datacenter Load Balancing

Antonio Marsico¹, Gianni Antichi², Theophilus Benson³, Marco Savi¹,
Andrew W. Moore²

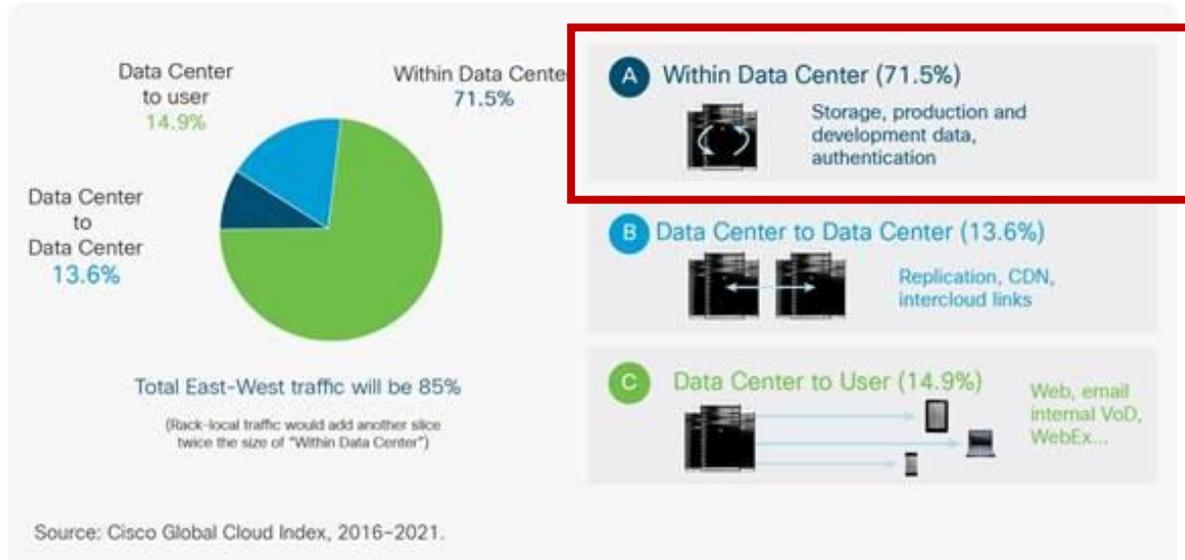
¹FBK CREATE-NET, ²University of Cambridge, ³Brown University

Coseners 2018

July 6th 2018

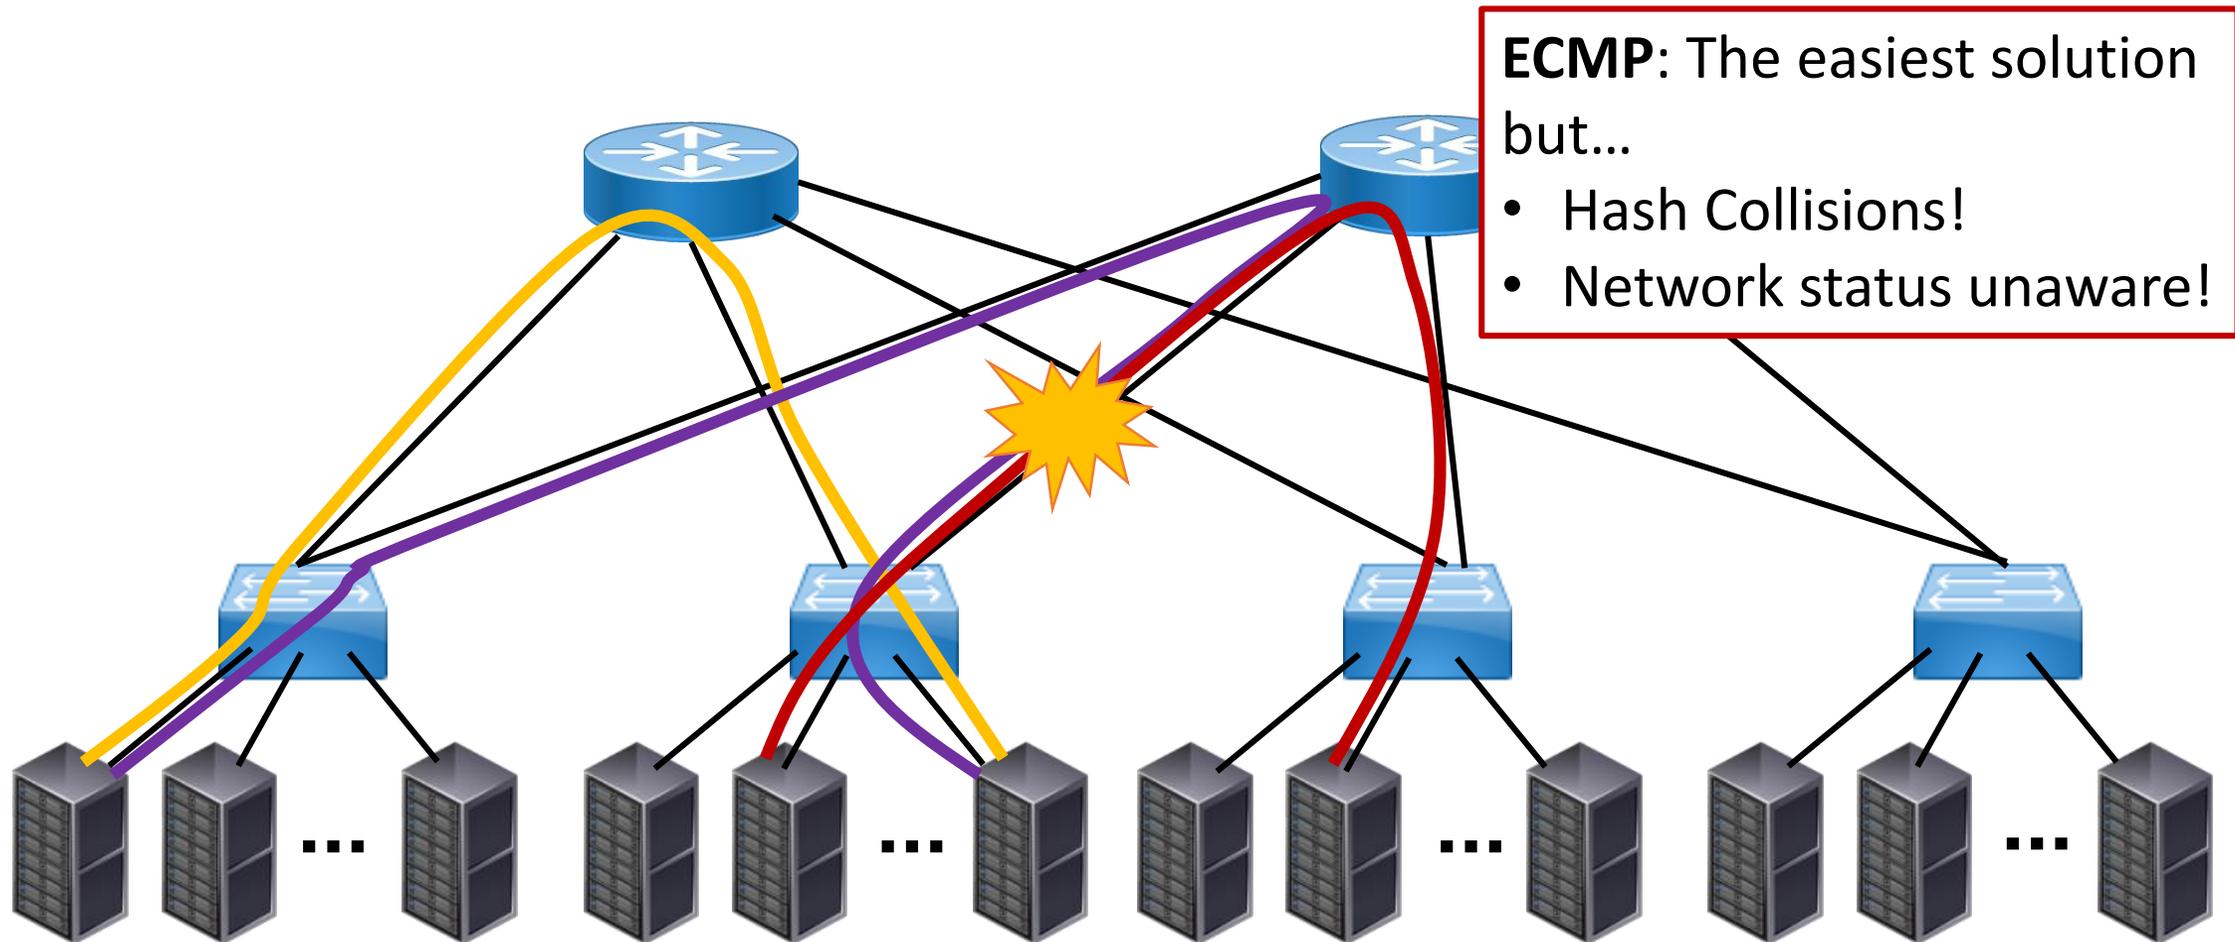
Introduction: Today's Data Center traffic

- A **growing** number of **applications** generating a **huge amount of traffic**



Intra-DC Load Balancing

Objective: *Increasing the utilization* of all the paths with equal costs between a SRC and a DST



Overcoming ECMP drawbacks

Load Balancing Solution	Traffic Granularity	Measurement Primitives	Designed for
Let it Flow	Flowlet	None	Asymmetric topologies
Hedera/MicroTE	Flow	Heavy hitters	Optimizing elephant flows
DRILL	Packet	Queue occupancy	High network load (> 80%)
Local Flow	Flow (selective splitting)	Link congestion	Symmetric topologies
Conga	Flowlet	Link congestion	Sym/Asym topologies
Hula	Flowlet	Link congestion	Sym/Asym topologies

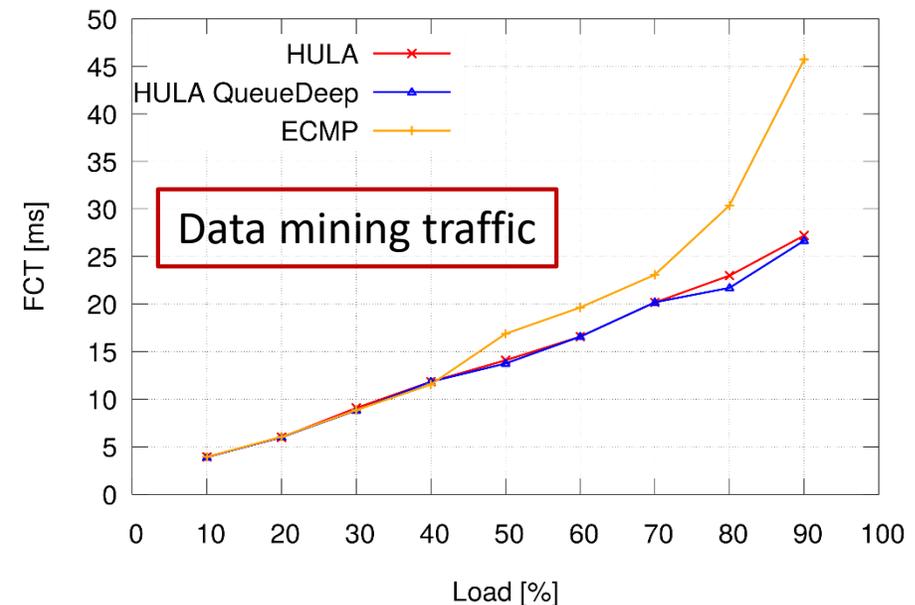
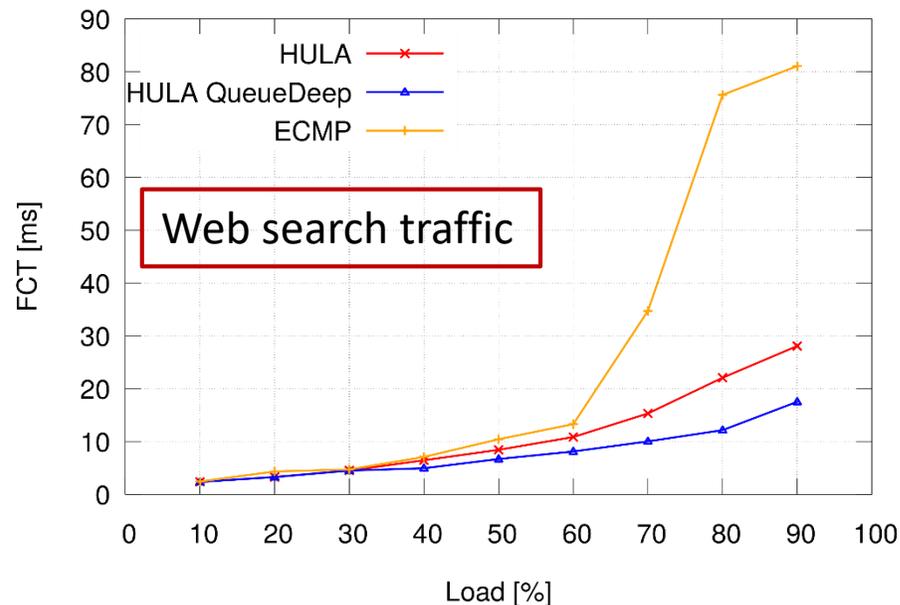
Overcoming ECMP drawbacks

Load Balancing Solution	Traffic Granularity	Measurement Primitives	Designed for
Let it Flow	Flowlet	None	Asymmetric topologies
He		Heavy hitters	Optimizing elephant flows
DR		Queue occupancy	High network load (> 80%)
Lo	(ng)	Link congestion	Symmetric topologies
Conga	Flowlet	Link congestion	Sym/Asym topologies
Hula	Flowlet	Link congestion	Sym/Asym topologies

Takeaway #1:
There *is no one measurement* primitive that is *superior* to the others

Example: HULA¹

- By design, **HULA** exploits *link congestion* information to make routing decisions
- What happens when we modify the measurement primitive?



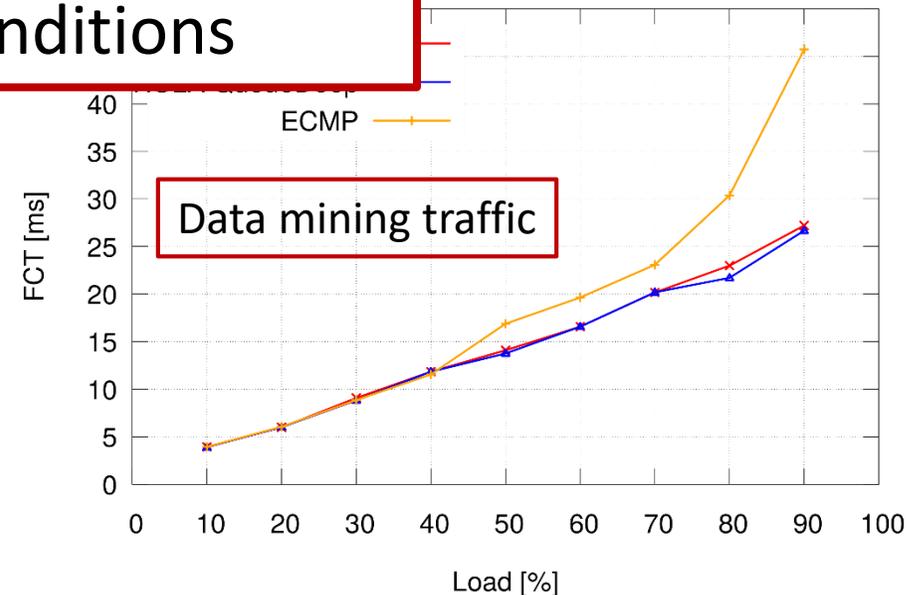
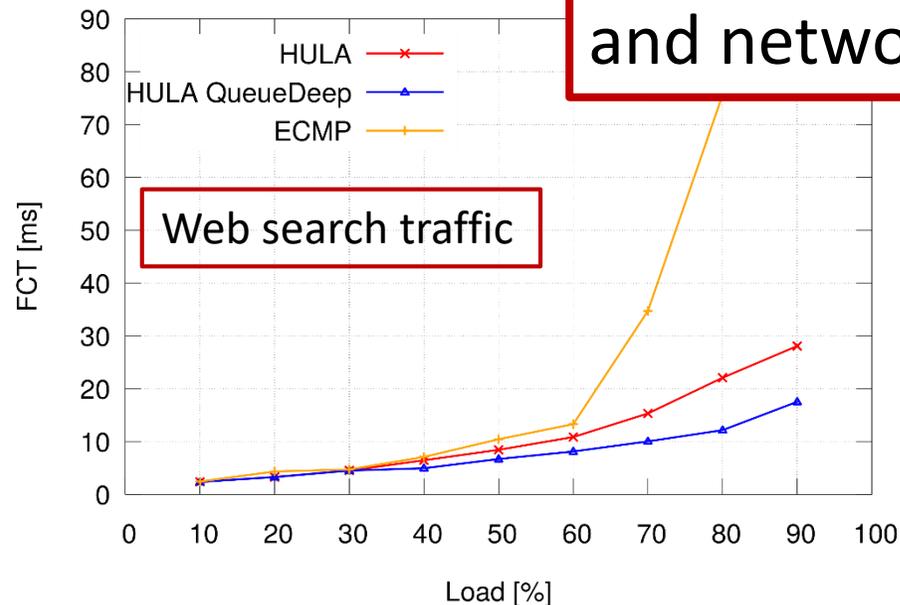
[1] N. Katta et al., "HULA: Scalable Load Balancing Using Programmable Data Planes", SORS 2016

Example: HULA¹

- By design, **HULA** enables **flexible** measurement decisions
- What happens when the application, topology, and network conditions change?

Takeaway #2:

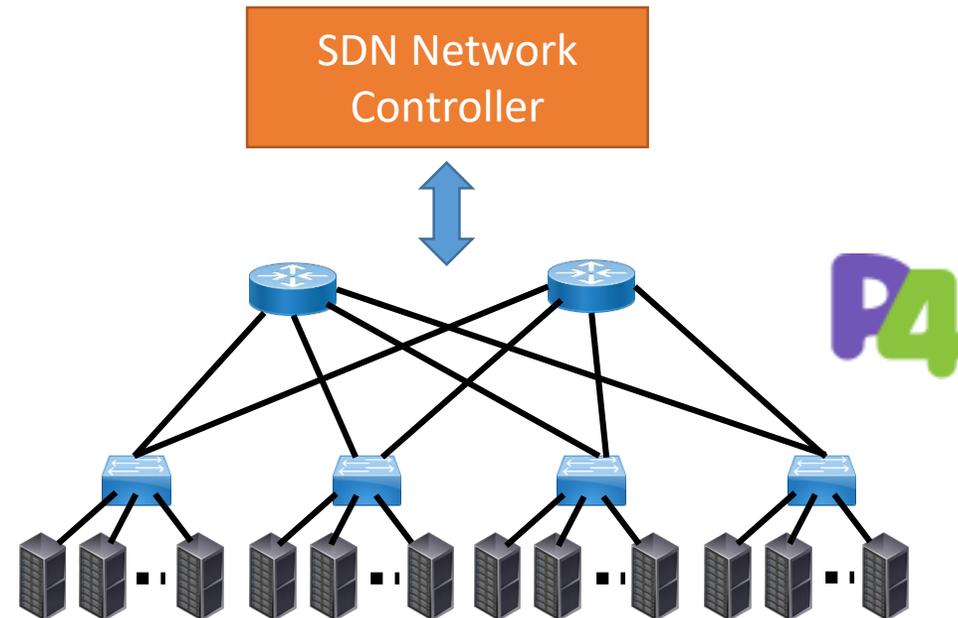
The *choice* of measurement primitive *is predicated on* the application, topology, and network conditions



[1] N. Katta et al., "HULA: Scalable Load Balancing Using Programmable Data Planes", SORS 2016

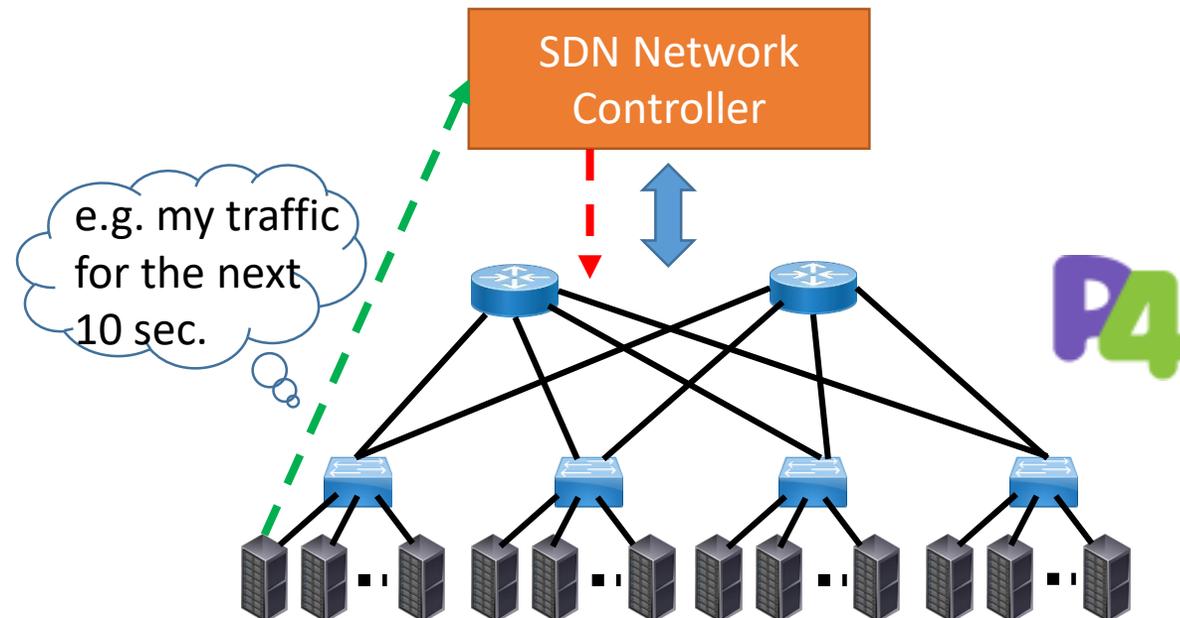
One measurement does not fits all scenarios

- **Data planes** of data center devices must provide a **flexible** and **extensible** substrate to **support** various type of measurements



A possible solution – Interaction between Applications and Networks

- **Applications/Hosts** and **Networks** could **cooperate together** to **improve** the network performance
- Exploiting the best measurement primitive based on the application traffic



Food for thoughts

- **Challenges** for a **real implementation** (some examples):
 - P4 itself cannot generate probes to sense the network (e.g. HULA)
 - Communication between applications/hosts and controller
 - Modifying the data plane configuration
 - Handling a partial deployment scenario: not all the data center devices can be substitute at the same time!

Thank you for your kind
attention!

amarsico@fbk.eu

@toto_120