DANIELE DE SENSI*, TIZIANO DE MATTEIS, KONSTANTIN TARANOV, SALVATORE DI GIROLAMO, TOBIAS RAHN, TORSTEN HOEFLER

Noise in the Clouds: Influence of Network Performance Variability on Application Scalability

* NOW AT SAPIENZA UNIVERSITY OF ROME, ITALY

HPC in the cloud

	Provider	CPU	Net. Bw.	Network	Transp. Prot.
Cloud	AWS	2x18C Intel Xeon Platinum @ 3GHz	100 Gb/s	Fat Tree	SRD
	Azure	2x22C Intel Xeon Platinum @ 2.7GHz	100-200 Gb/s	Non-blocking Fat Tree	InfiniBand
	GCP	2x15C Intel Cascade Lake @ 3.1GHz	100 Gb/s	3:1 blocking Fat Tree	TCP/IP + Intel QuickData
	Oracle	2x18C Intel Xeon Gold @ 3GHz	100 Gb/s	Non-blocking Fat Tree	RoCEv2
On-Premise	Daint	2x18C Intel Xeon E5-2695 v4 @2.1GHz	82 Gb/s	Cray Aries (Dragonfly)	FMA
	Alps	2x64C AMD EPYC 7742 @ 2.25GHz	100 Gb/s	HPE Cray Slingshot (Dragonfly)	RoCEv2
	DEEP-EST	2x12C Intel Xeon Gold @3.2GHz	100 Gb/s	Mellanox Infiniband EDR (Fat Tree)	InfiniBand



Target

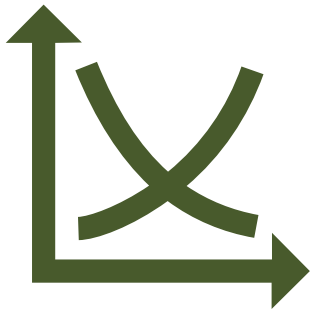

Compare **network performance**
of cloud HPC vs on-premise HPC

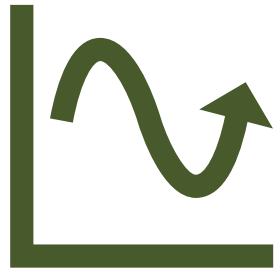

Analyze **network noise**
of cloud HPC vs on-premise HPC
and its impact **at scale**

Network Performance

Bandwidth and latency

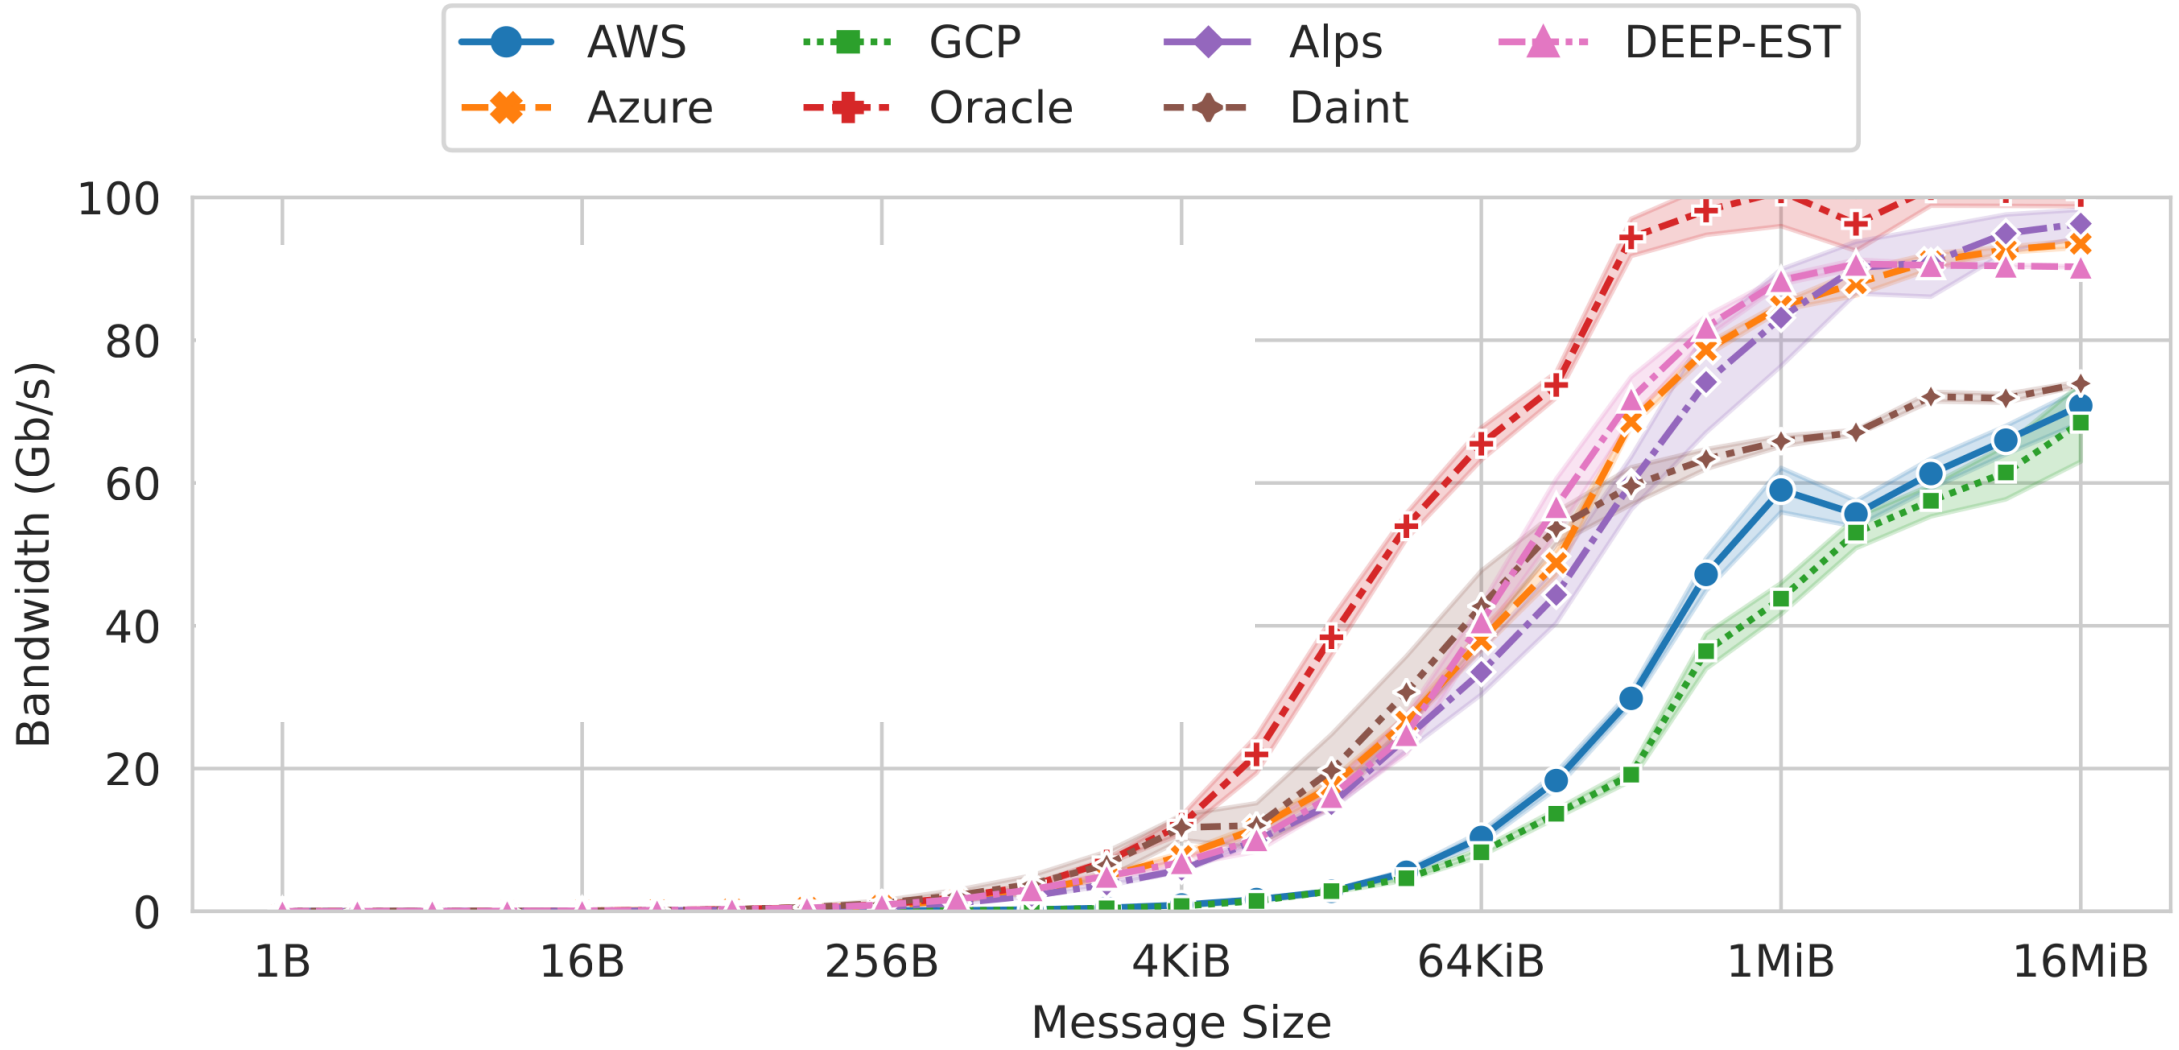

HPC vs. normal instances

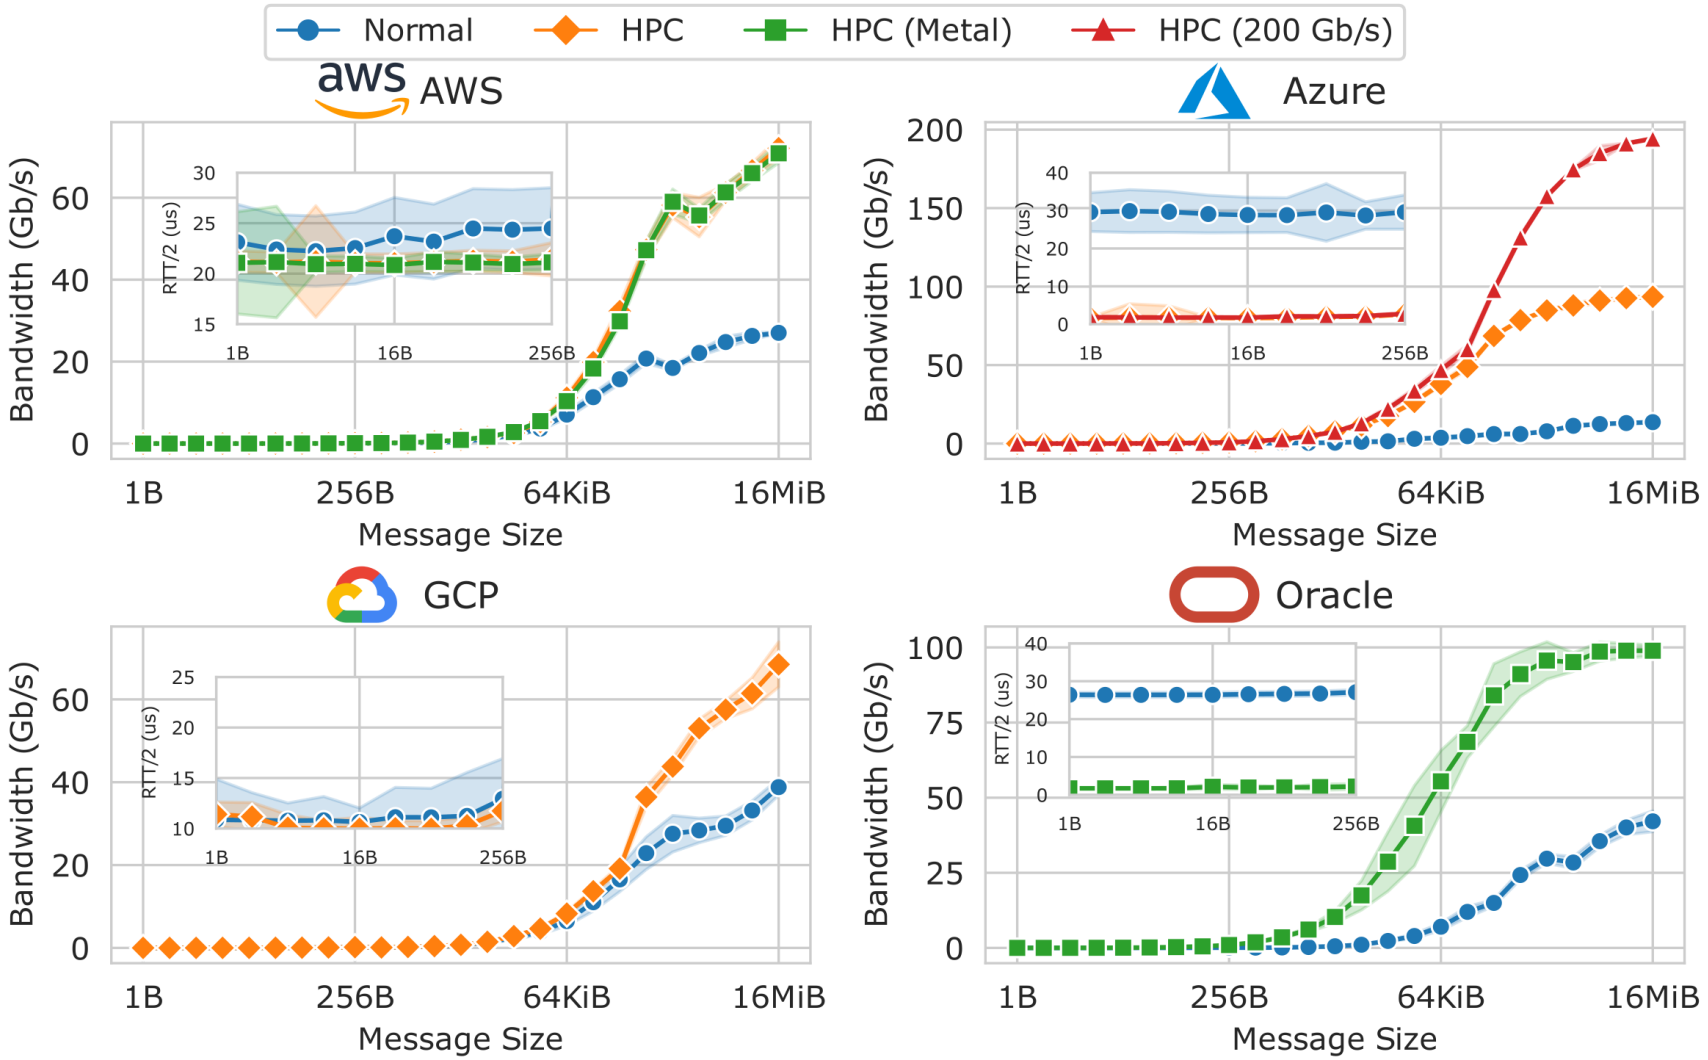

Network Noise

Bandwidth noise

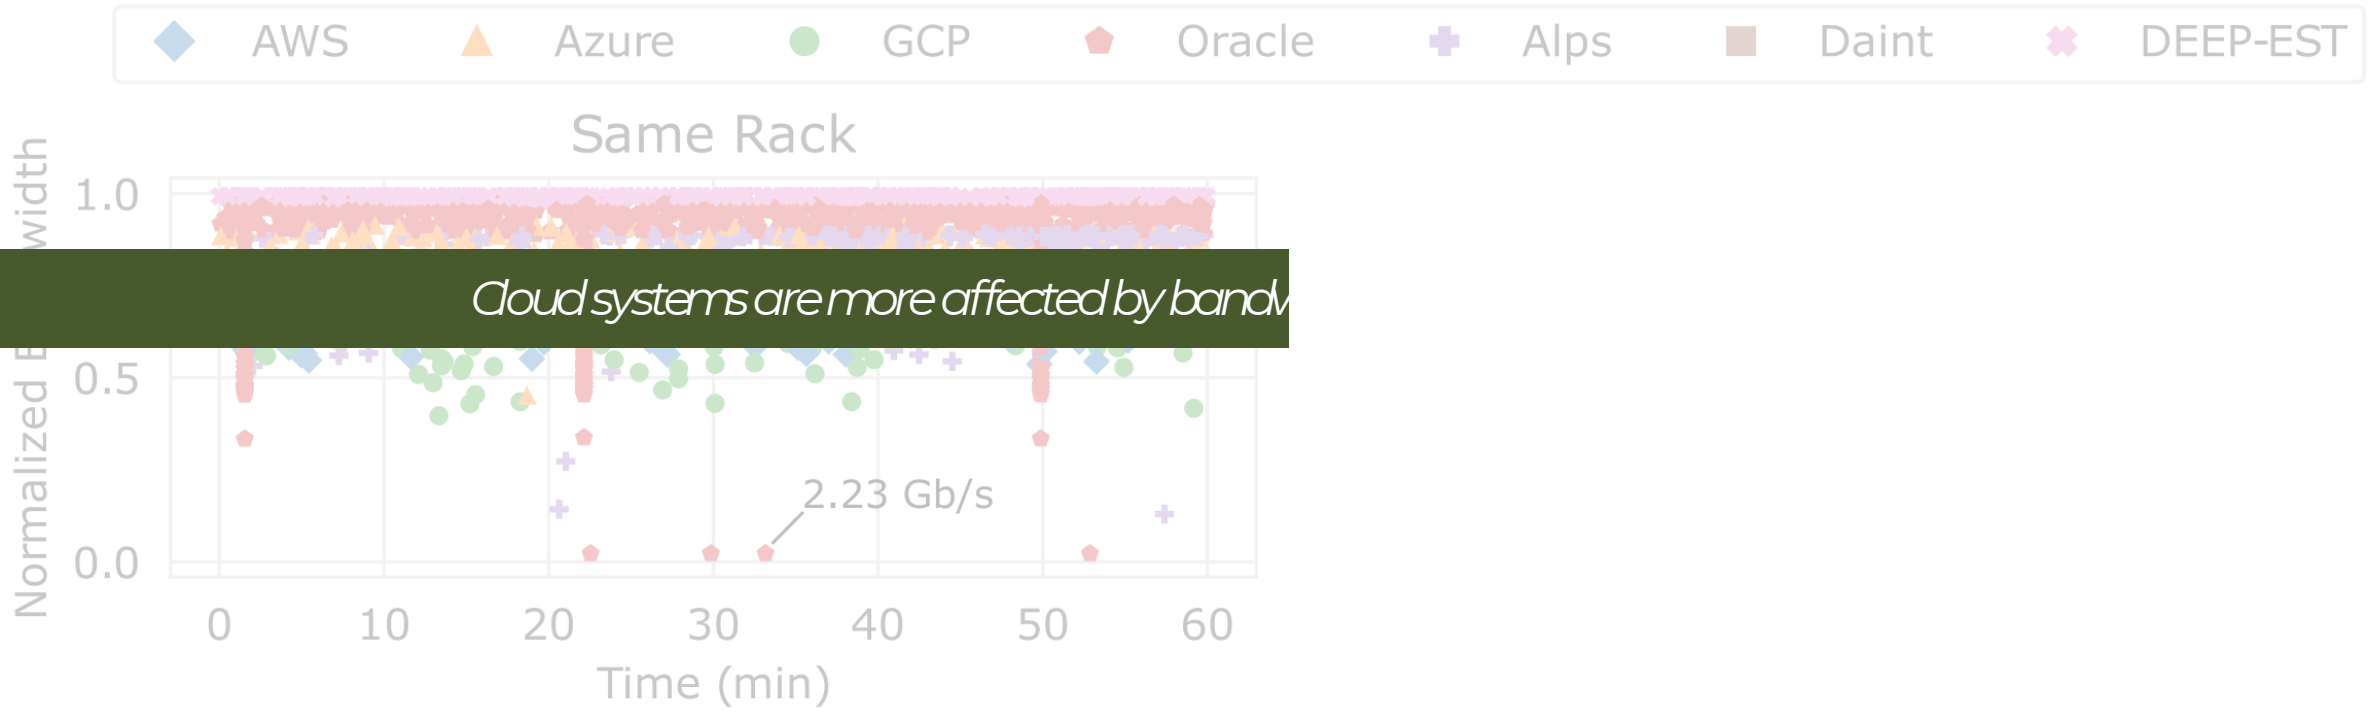

Cloud systems are more affected by bandw

Noise impact at scale

Noise

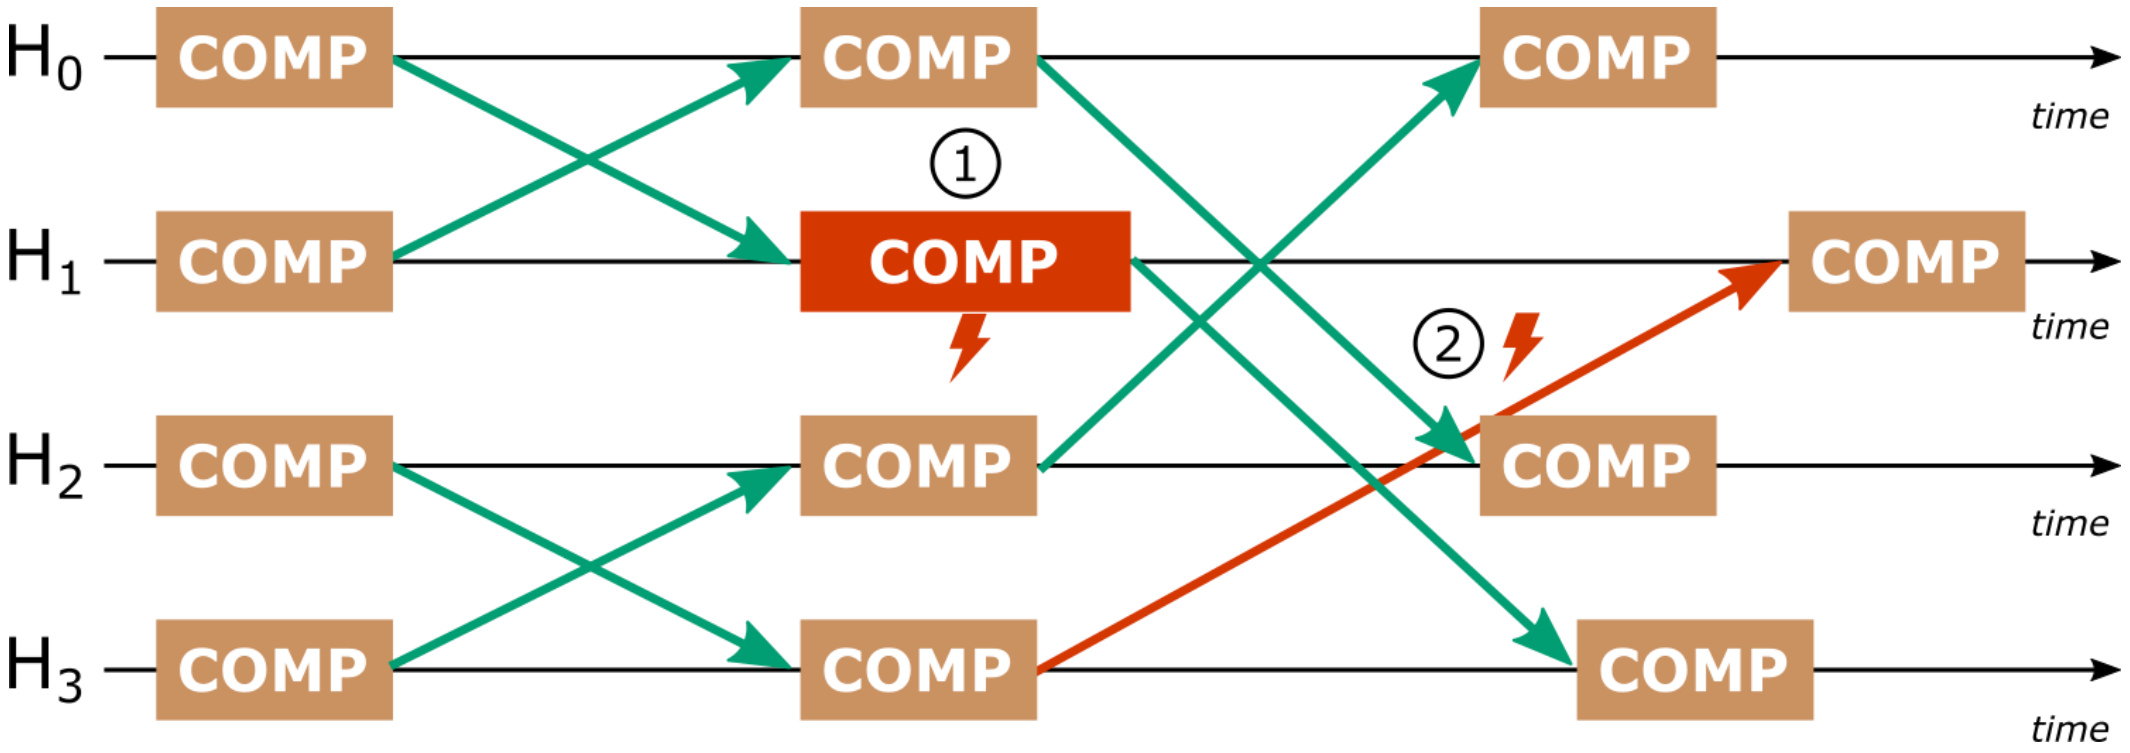
H0
H1
H2
H3
COMP
COMP
COMP
COMP
COMP
COMP
COMP
COMP
COMP
COMP
①
②
time
time
time
time

Methodology

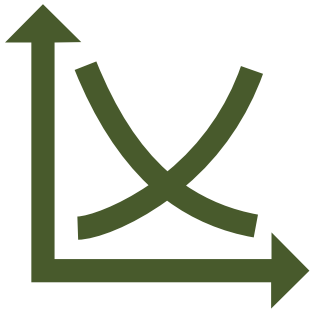

Get network performance measurements...

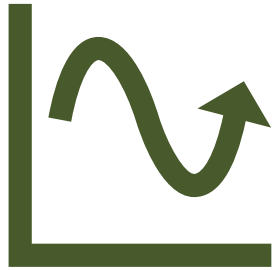

...and OS and network noise measurements...

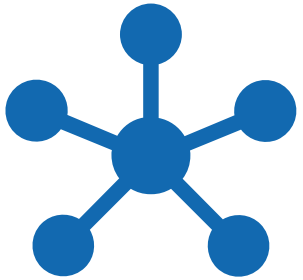

...and simulate performance at scale

Simulation validation (16 nodes)

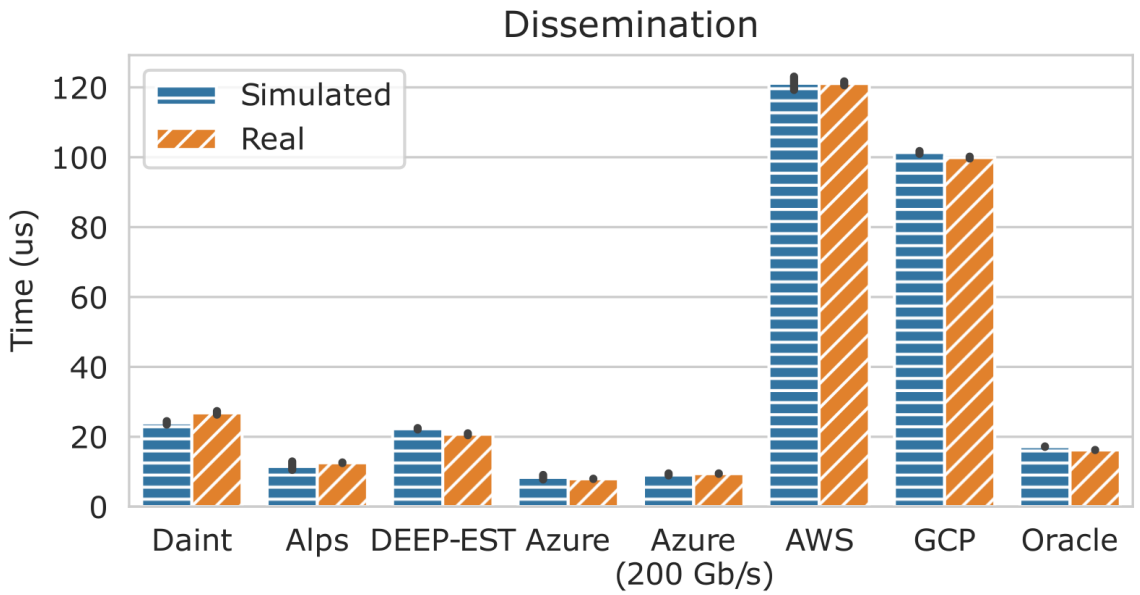

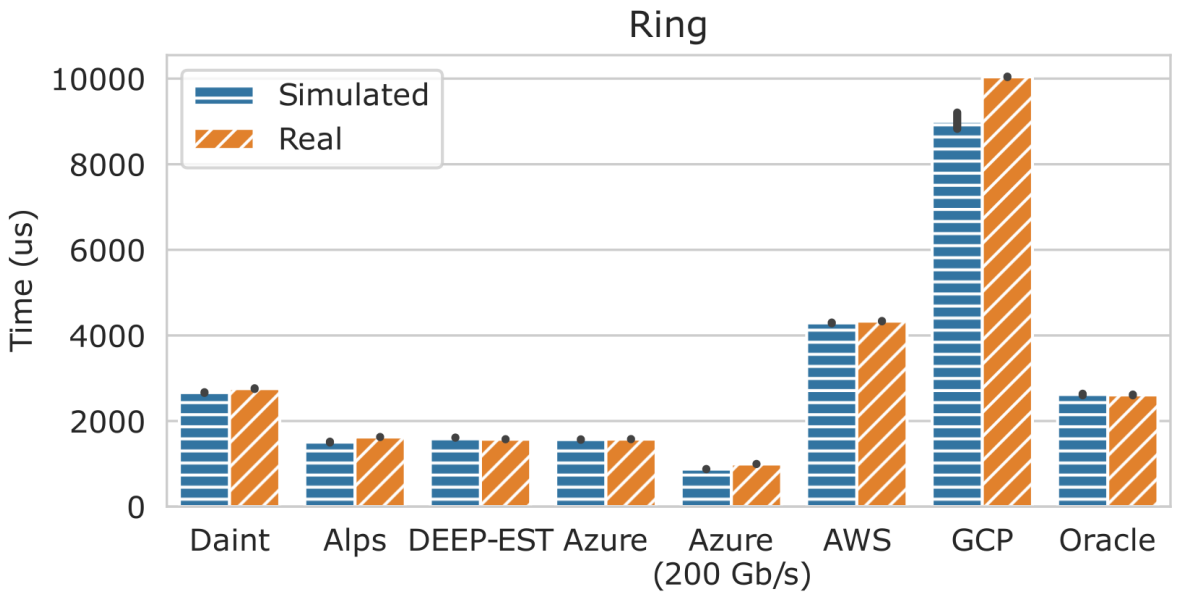

512 MiB Ring collective simulation

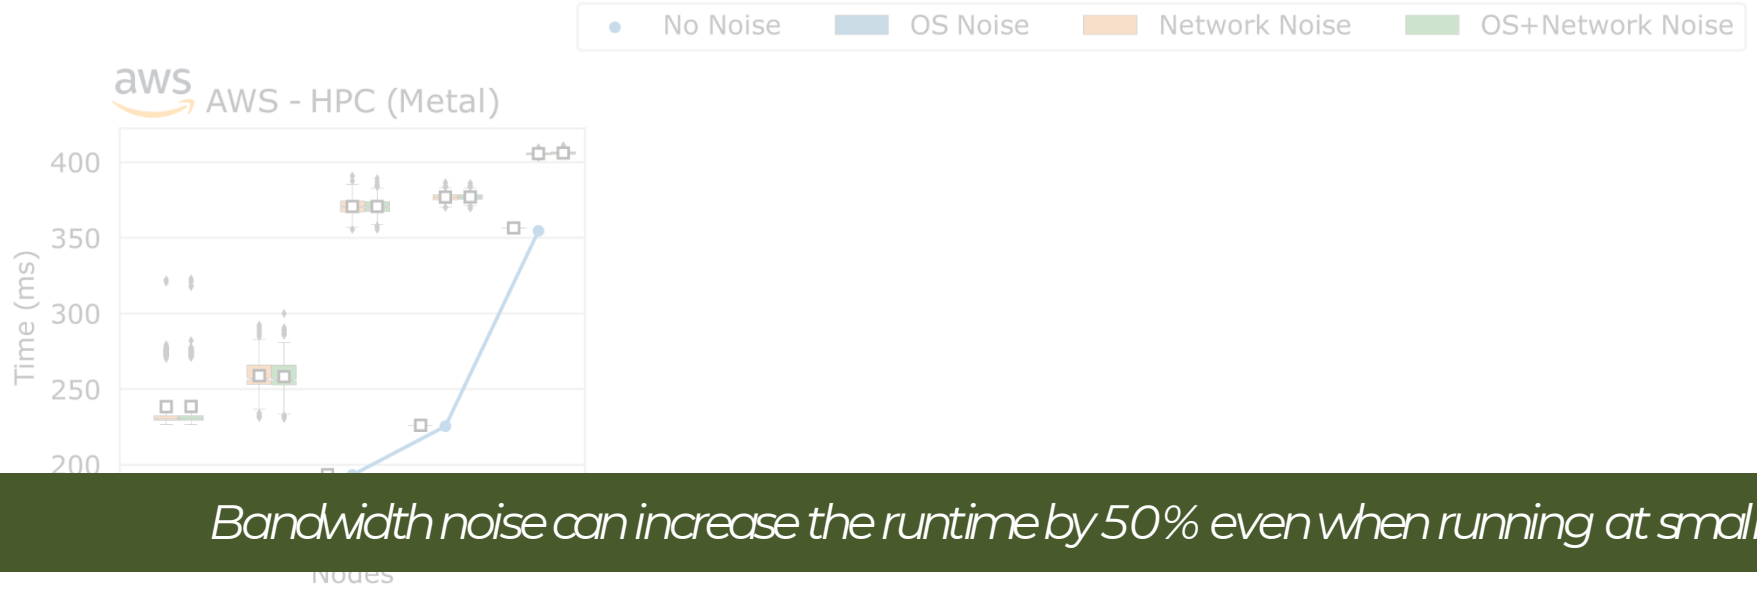

Bandwidth noise can increase the runtime by 50% even when running at small scale (4 nodes)

Impact of bandwidth noise on monetary cost

Dummy application: 8192x8192 matrix multiplication followed by 512MiB allreduce (**20% of time** spent in **communication**)

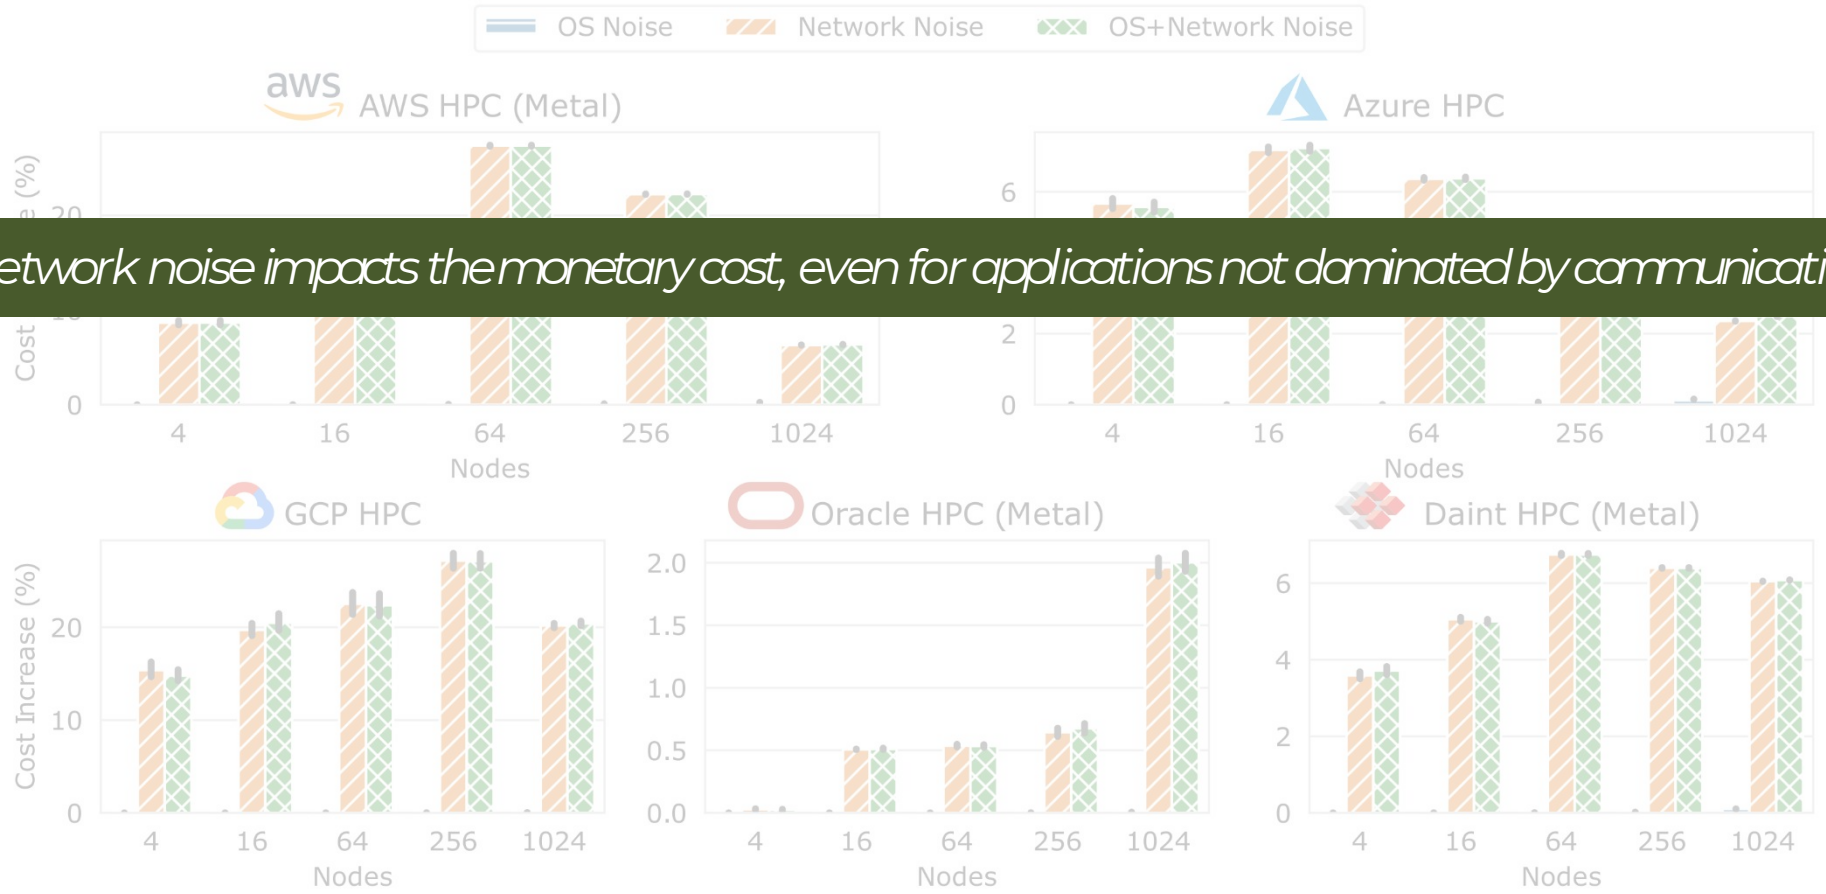

Network noise impacts the monetary cost, even for applications not dominated by communication

Conclusions

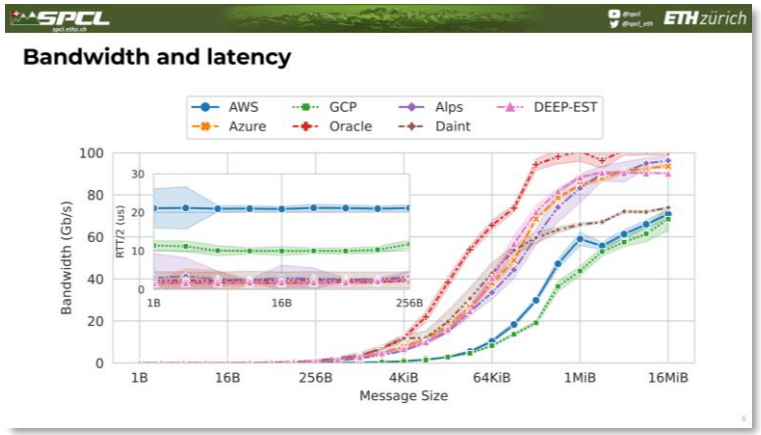

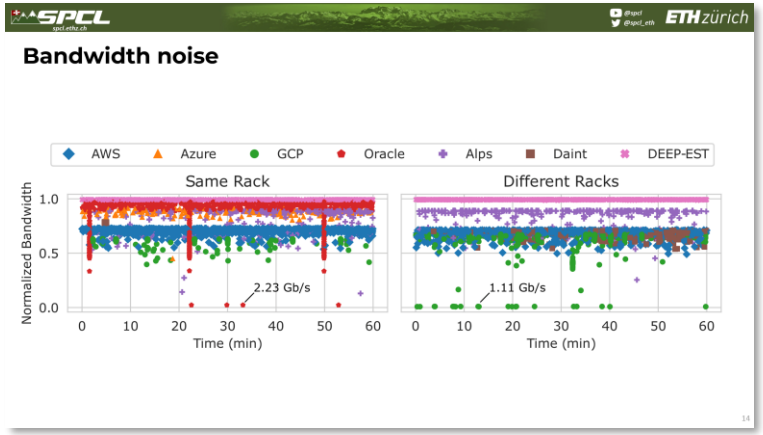

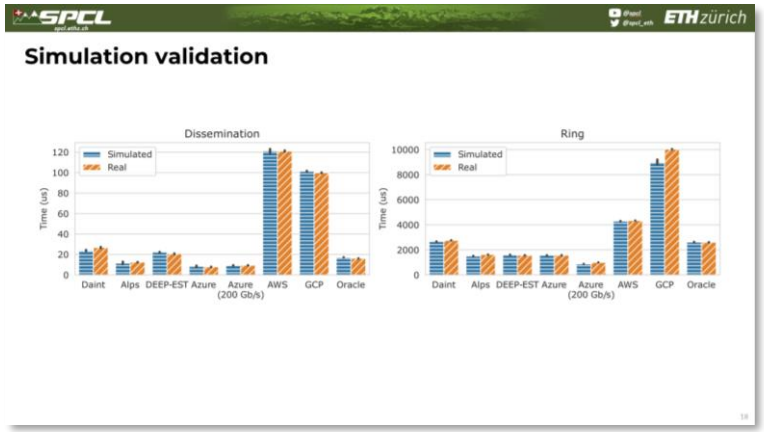

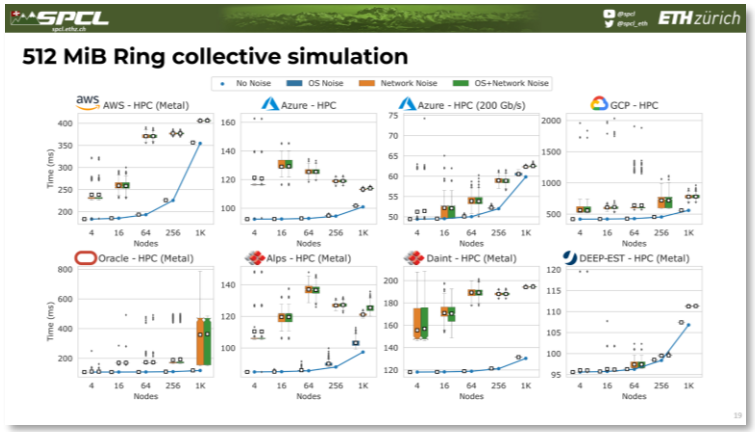

More of SPCL's research:

youtube.com/@spcl — 150+ Talks

twitter.com/spcl_eth — 1.2K+ Followers

github.com/spcl — 2K+ Stars

… or spcl.ethz.ch

https://github.com/DanieleDeSensi/cloud_noise
https://github.com/DanieleDeSensi/cloud_noise_data