

Exploring GPU-to-GPU Communication: Insights into Supercomputer Interconnects


Daniele De Sensi, Lorenzo Pichetti, Flavio Vella, Tiziano De Matteis,
Zebin Ren, Luigi Fusco, Matteo Turisini, Daniele Cesarini, Kurt Lust,
Animesh Trivedi, Duncan Roweth, Filippo Spiga, Salvatore Di Girolamo, Torsten Hoefler




Motivation



**Systems are getting larger
(up to 100,000 GPUs)**




**Nodes are getting denser
(e.g., 8 GPUs on the same node)**



**Maturity of the software stack
(MPI, NCCL, RCCL, etc...)**

Systems analyzed




Alps (#6 in Top500)

- NVIDIA H100 GPUs
- HPE Cray Slingshot
- Dragonfly

Early-access

Systems analyzed

200 Gb/s per direction




Alps (#6 in Top500)

- NVIDIA H100 GPUs
- HPE Cray Slingshot
- Dragonfly

Early-access

100 Gb/s per direction




Leonardo (#7 in Top500)

- NVIDIA A100 GPUs
- NVIDIA InfiniBand
- Dragonfly+

Systems analyzed

200 Gb/s per direction




Alps (#6 in Top500)

- NVIDIA H100 GPUs
- HPE Cray Slingshot
- Dragonfly

Early-access


100 Gb/s per direction



Leonardo (#7 in Top500)

- NVIDIA A100 GPUs
- NVIDIA InfiniBand
- Dragonfly+

200 Gb/s per direction



LUMI (#5 in Top500)

- AMD MI250X GPUs
- HPE Cray Slingshot
- Dragonfly

How to Transfer Data Between GPUs

- **Trivial staging** (GPU->Host->Host->GPU)
- **Device-device copy** (peer access)
- ***CCL** (i.e., NCCL/RCCL)
- **GPU-Aware MPI**



We used the software versions suggested by each supercomputing center

Experiments Executed


	Point-to-point	Collective Operations (allreduce and alltoall)
Intra-Node		
Inter-Node		

Experiments Executed


	Point-to-point	Collective Operations (allreduce and alltoall)
Intra-Node		
Inter-Node		




Intra-Node Point-to-Point



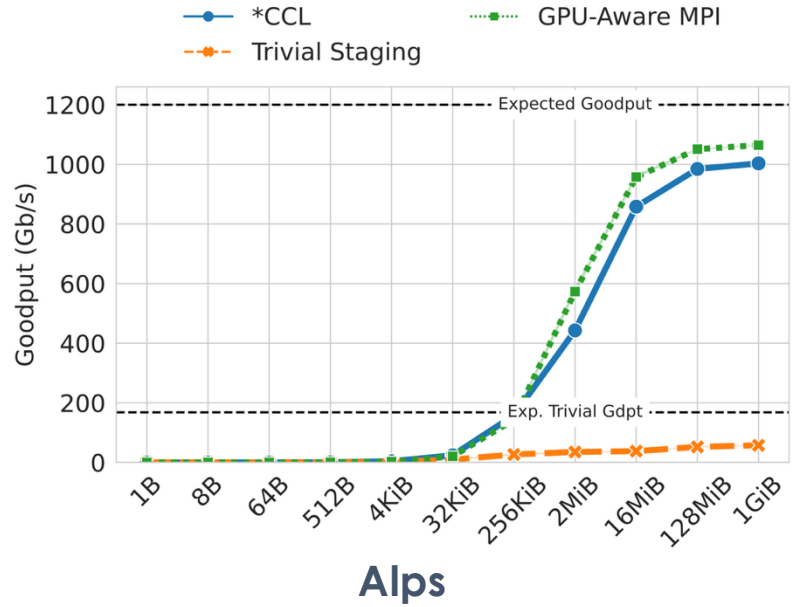
Intra-Node Point-to-Point



Intra-Node Point-to-Point




Intra-Node Point-to-Point





Intra-Node Point-to-Point

Device-device copy
not enabled at the time of writing

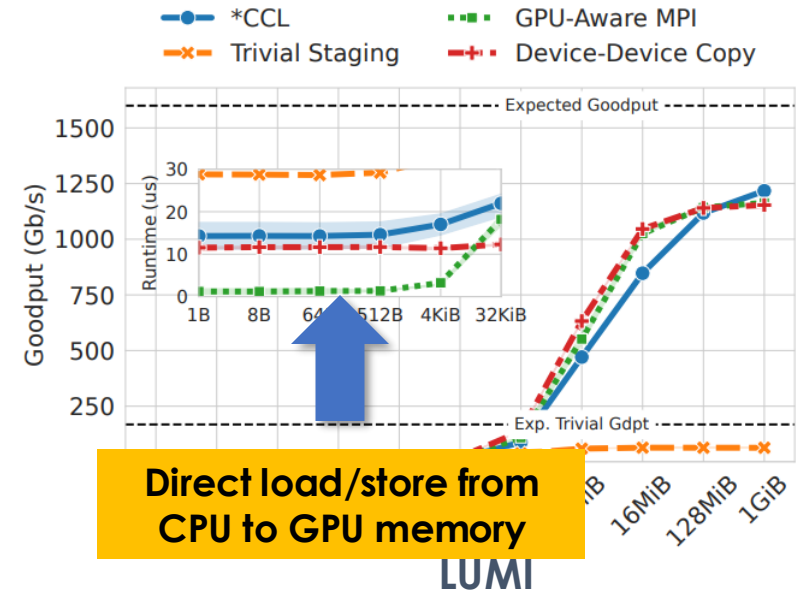




Intra-Node Point-to-Point



Beware: Different y-axis ranges across the plots


Intra-Node Point-to-Point



Beware: Different y-axis ranges across the plots

Intra-Node Point-to-Point

MPI and *CCL out-of-the-box performance is **sub-optimal** (i.e., up to **6x performance improvements** after tuning)





For small point-to-point intra-node transfers,
MPI can be 10x faster than *CCL

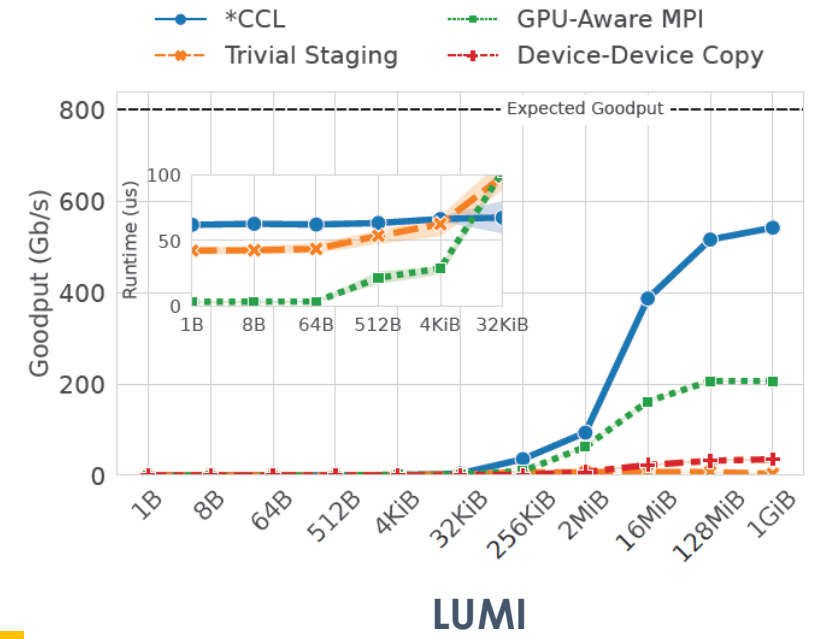


Beware: Different y-axis ranges across the plots

Experiments Executed

	Point-to-point	Collective Operations (allreduce and alltoall)
Intra-Node		
Inter-Node		



Intra-Node Allreduce



Intra-Node Allreduce



***CCL outperforms MPI up to 2x on collective operations**



Alps

Leonardo

LUMI

Experiments Executed

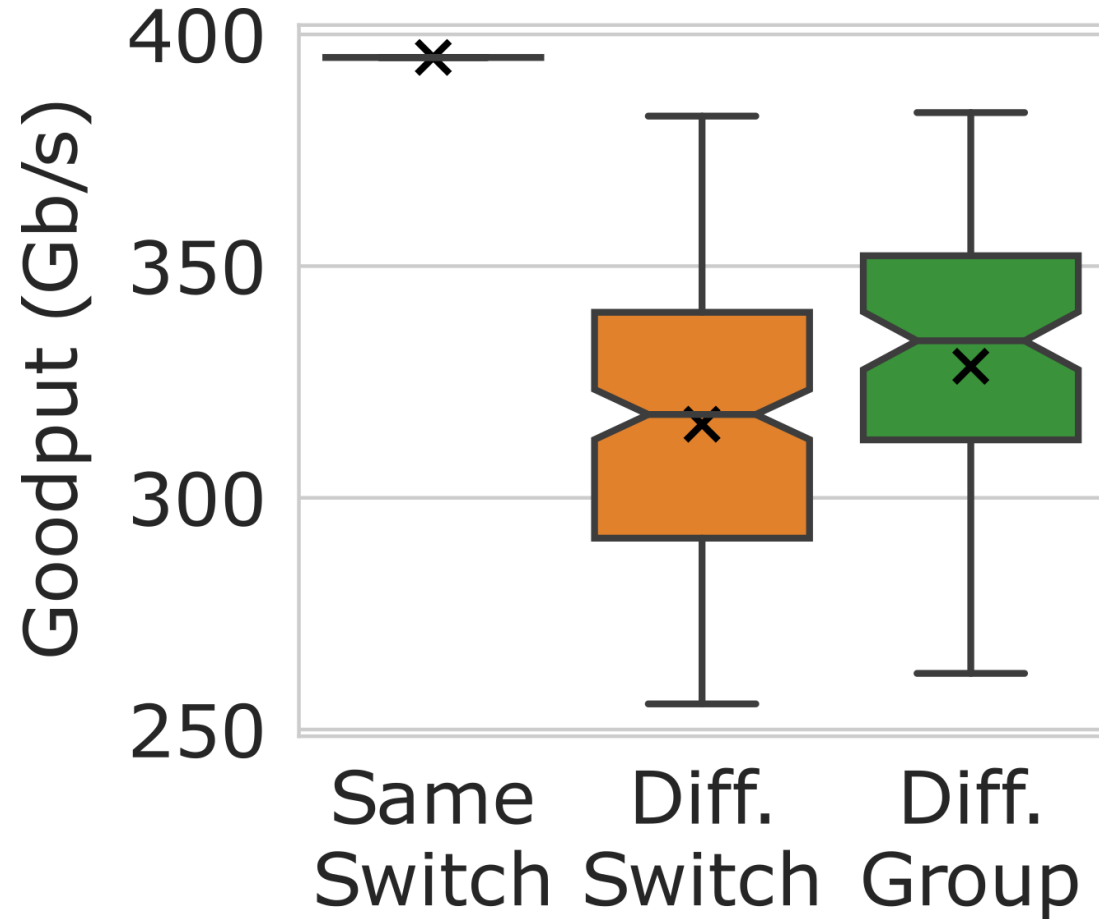
	Point-to-point	Collective Operations (allreduce and alltoall)
Intra-Node		
Inter-Node		



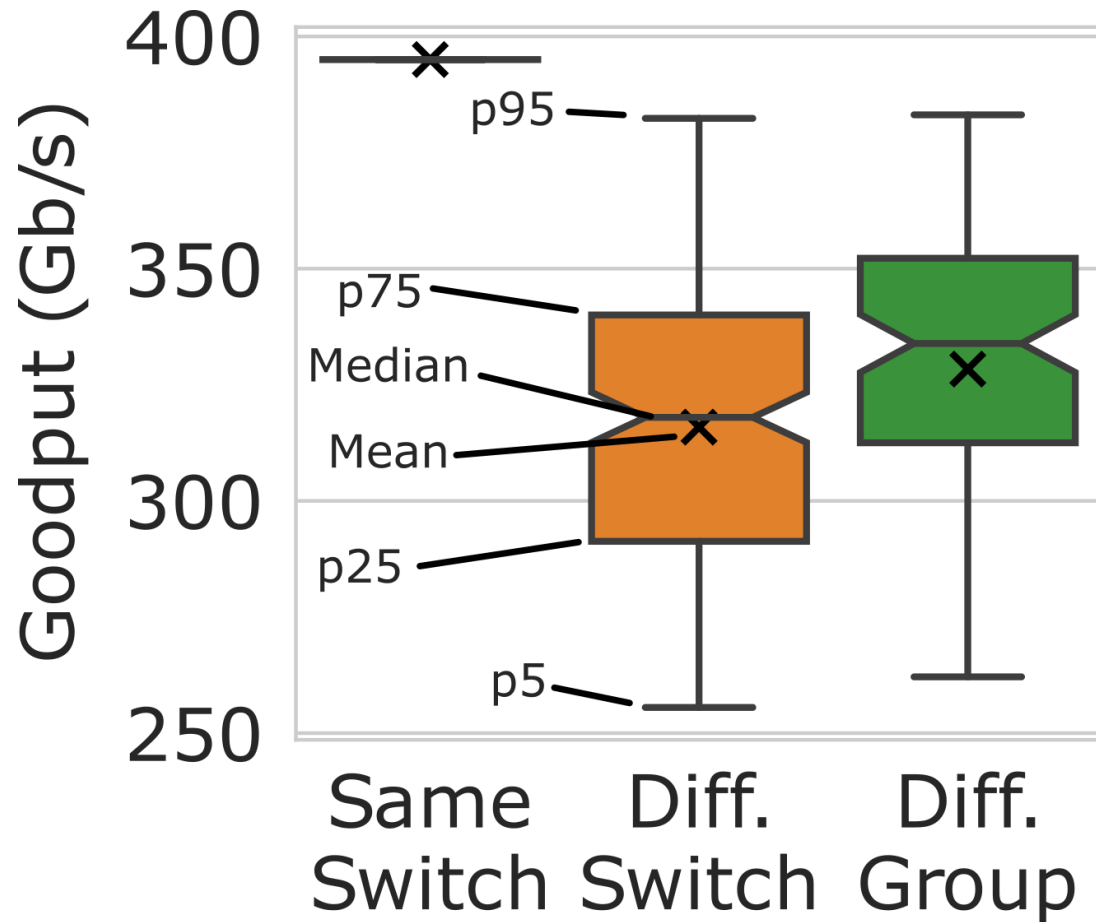
Inter-Node Point-to-Point

MPI outperforms *CCL up to 10x on small transfers,
and **up to 3x** on larger transfers

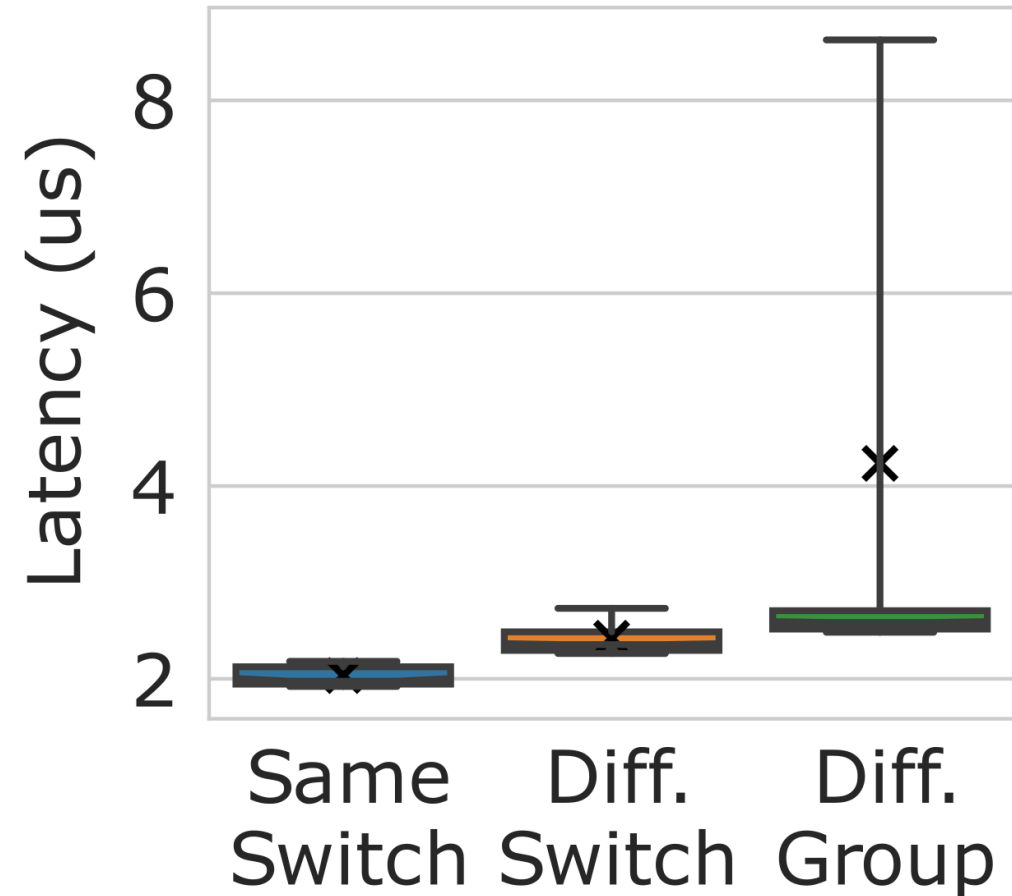
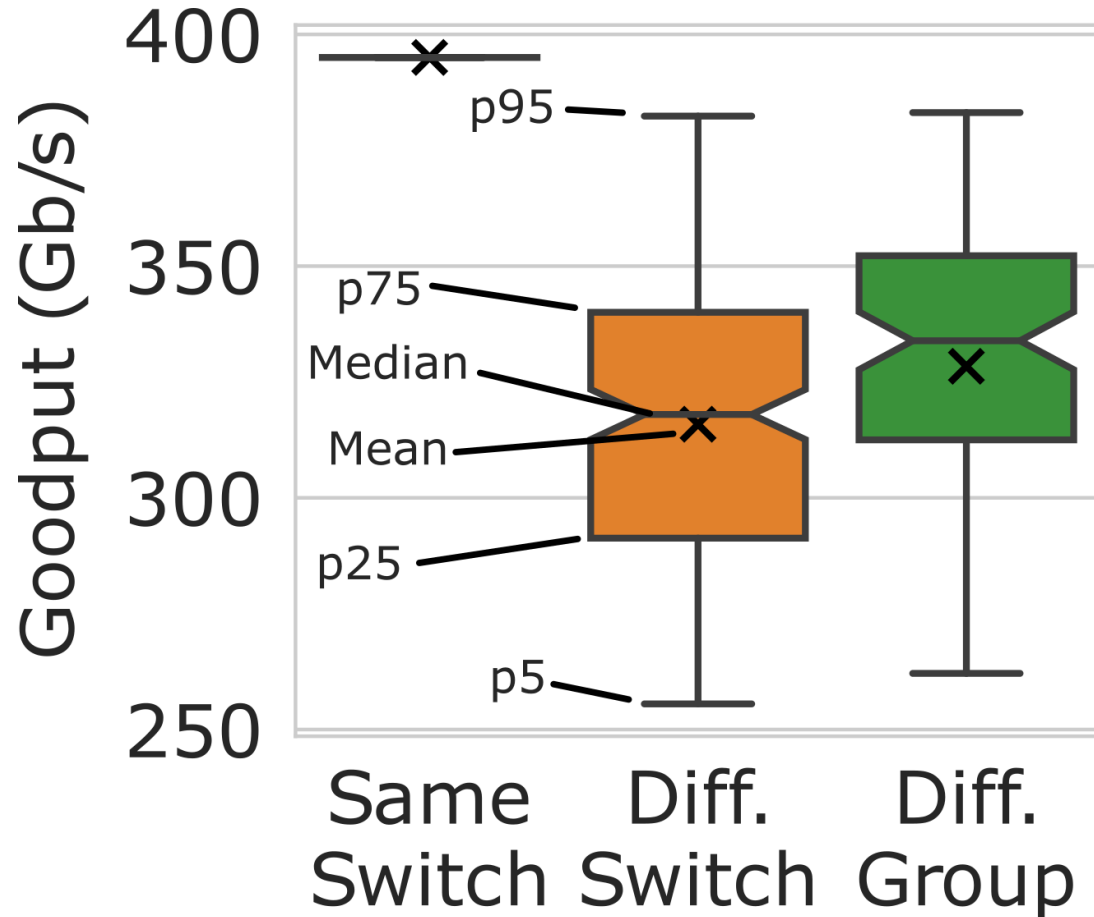
Leonardo Inter-Node Point-to-Point (for Different Distances)



Leonardo Inter-Node Point-to-Point (for Different Distances)



Leonardo Inter-Node Point-to-Point (for Different Distances)



Leonardo Inter-Node Point-to-Point (for Different Distances)



Leonardo inter-node **average latency increases up to 2x** due to network noise. Performance drops up to **50%** for collectives at scale

On Alps and LUMI, **GPU's network location** has a marginal impact

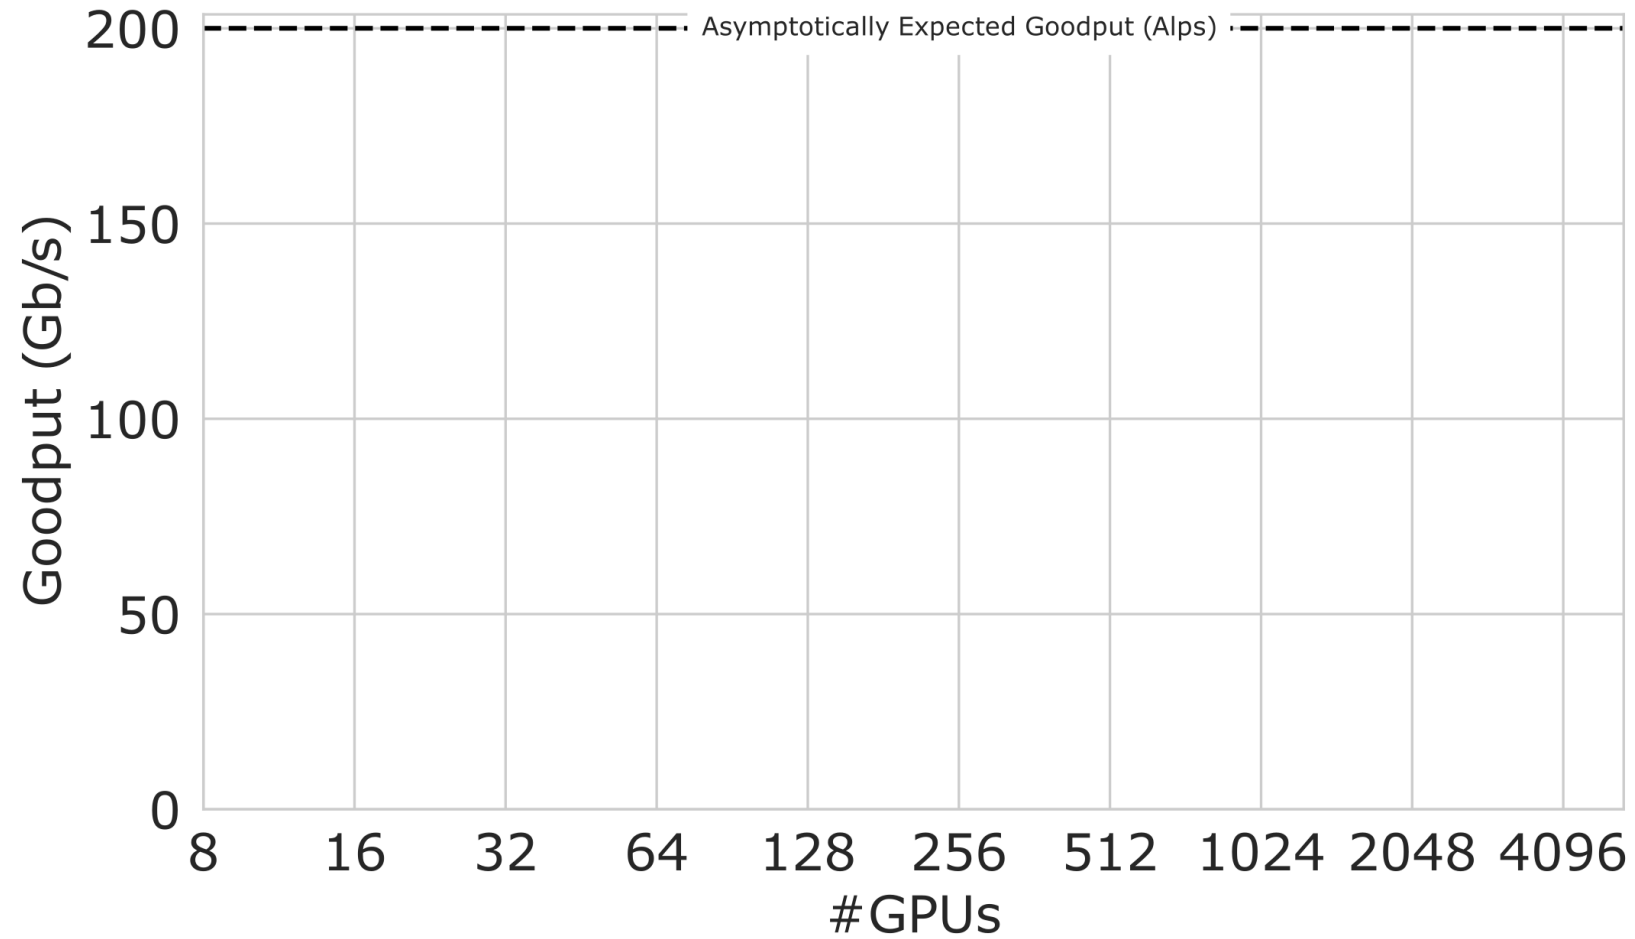
More **details** in the paper

Experiments Executed

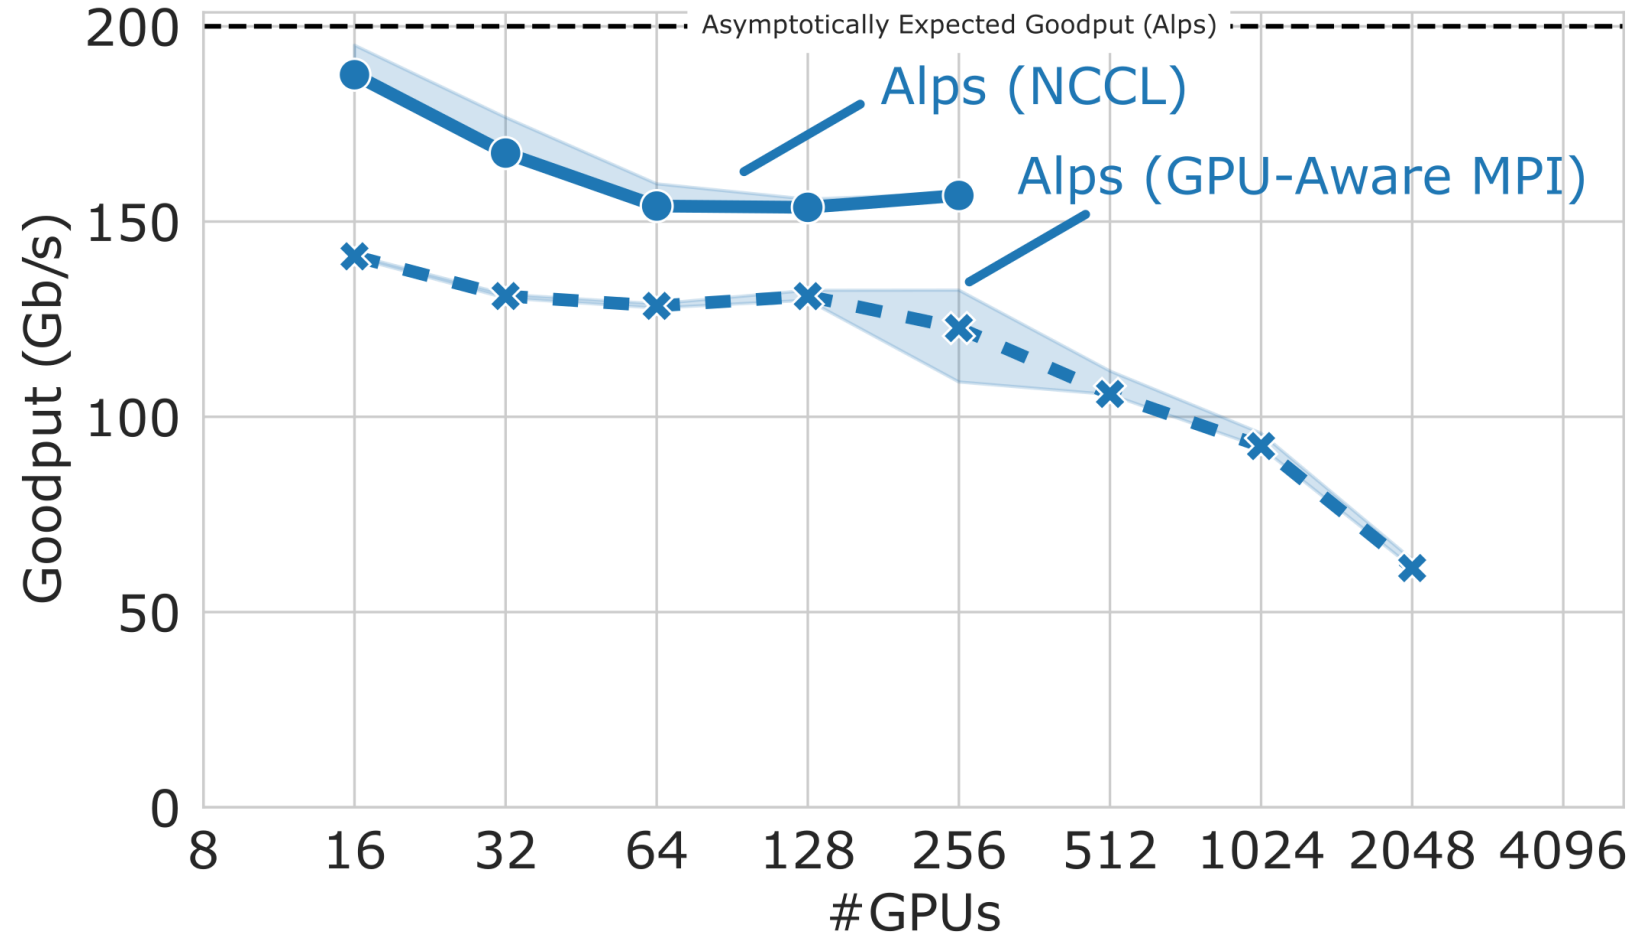
	Point-to-point	Collective Operations (allreduce and alltoall)
Intra-Node		
Inter-Node		




Inter-Node Alltoall




Inter-Node Alltoall




Inter-Node Alltoall




Inter-Node Alltoall



Inter-Node Alltoall

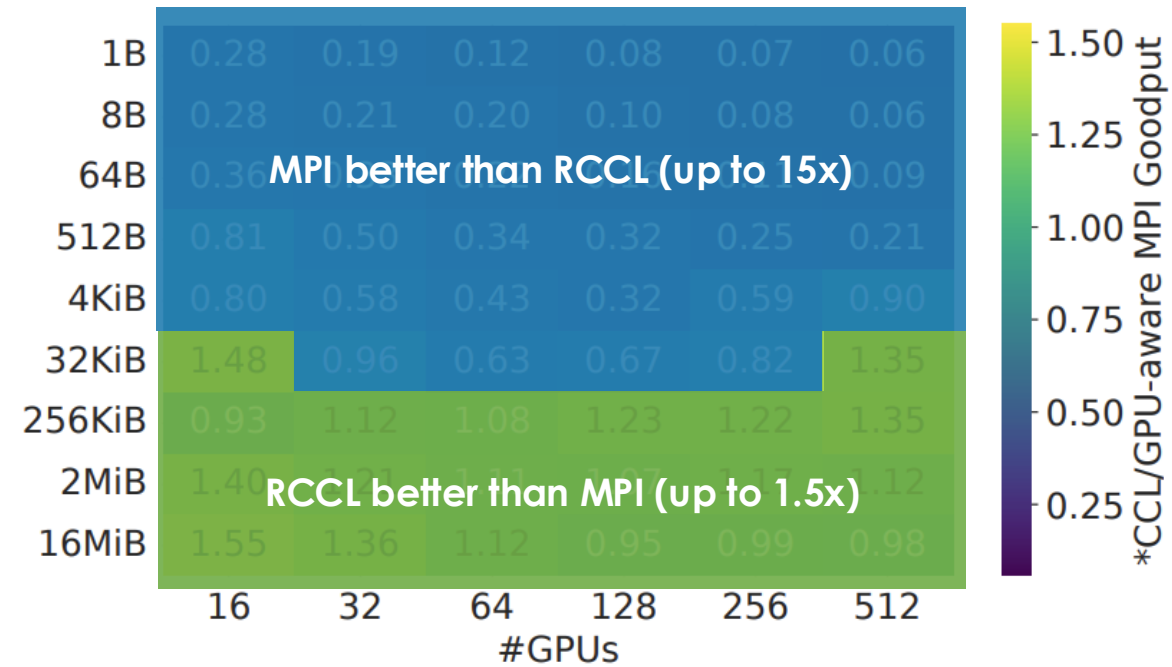


Inter-Node Alltoall




On large-scale collectives, *CCL outperforms MPI
However, we observed instabilities for the alltoall

MPI vs. *CCL



Alltoall



Allreduce

Conclusions & Main Take-Home Messages



To sysadmins and users:

Non-negligible **tuning** effort required to achieve good performance



To software designers:

MPI or *CCL? It depends (collectives vs. point-to-point, etc...)



To sysadmins and system architects:

Adaptive routing (at least on **DF+**) might require further **optimization**



desensi@di.uniroma1.it

<https://danieledesensi.github.io>

