Sandia National Laboratories

Exceptional service in the national interest

TEMPI: An Interposed MPI Library with a Canonical Representation of MPI Datatypes

Carl Pearson[1], Kun Wu[2], I-Hsin Chung[3], Jinjun Xiong[3], Wen-Mei Hwu[4]

[1]Sandia National Labs / [2]University of Illinois Electrical and Computer Engineering / [3]IBM T. J. Watson Research / [4]Nvidia Research

HPDC - Jun 24, 2021

Work completed at University of Illinois prior to joining Sandia National Labs

SAND2021-6846 C

Outline

- Distributed GPU stencil, non-contiguous data
- Equivalence of strided datatypes and minimal representation
- GPU communication methods
- Deploying on managed systems
- Large messages and MPI datatypes
- Translation and canonicalization
- Automatic model-driven transfer method selection
- Interposed library implementation

OLCF Summit Node

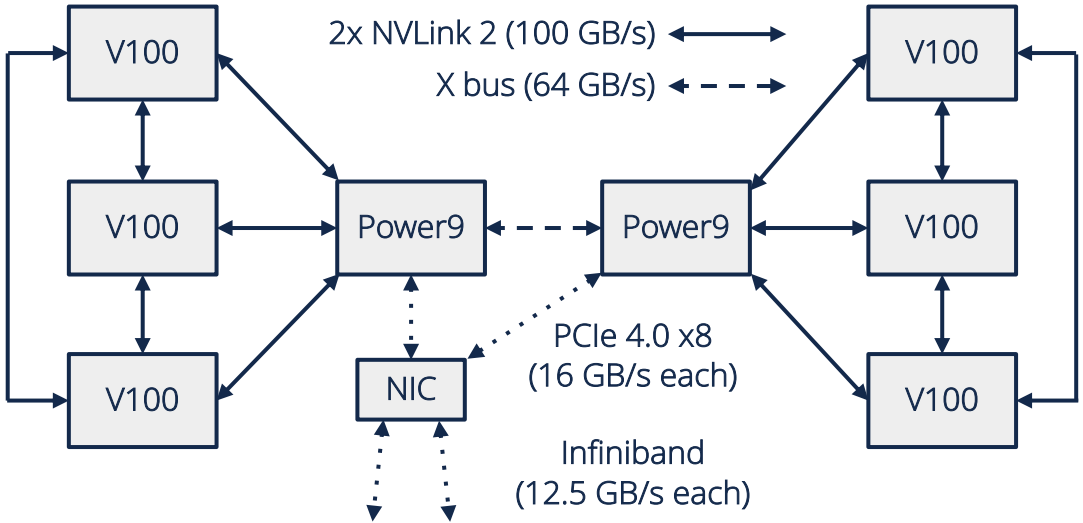

Summit Node
(bidirectional bandwidth)

Stencil Communication and non-contiguous Data

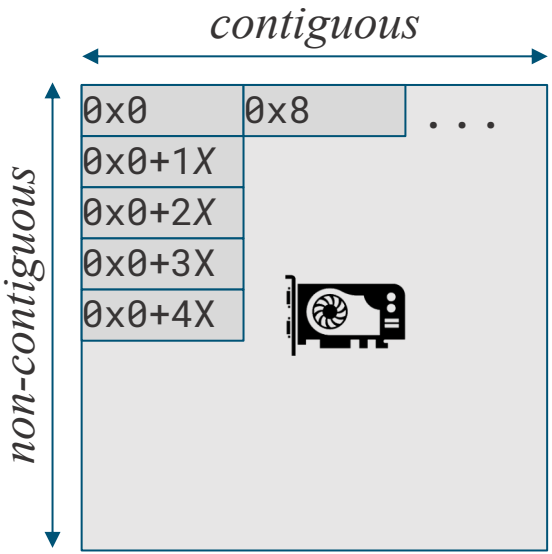

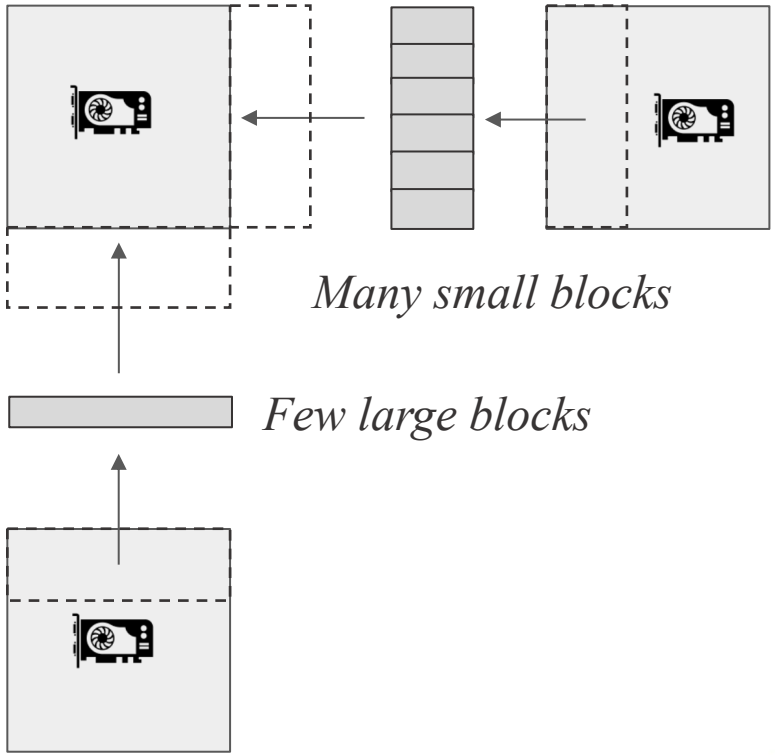

OLCF Summit Latency vs Transfer Size

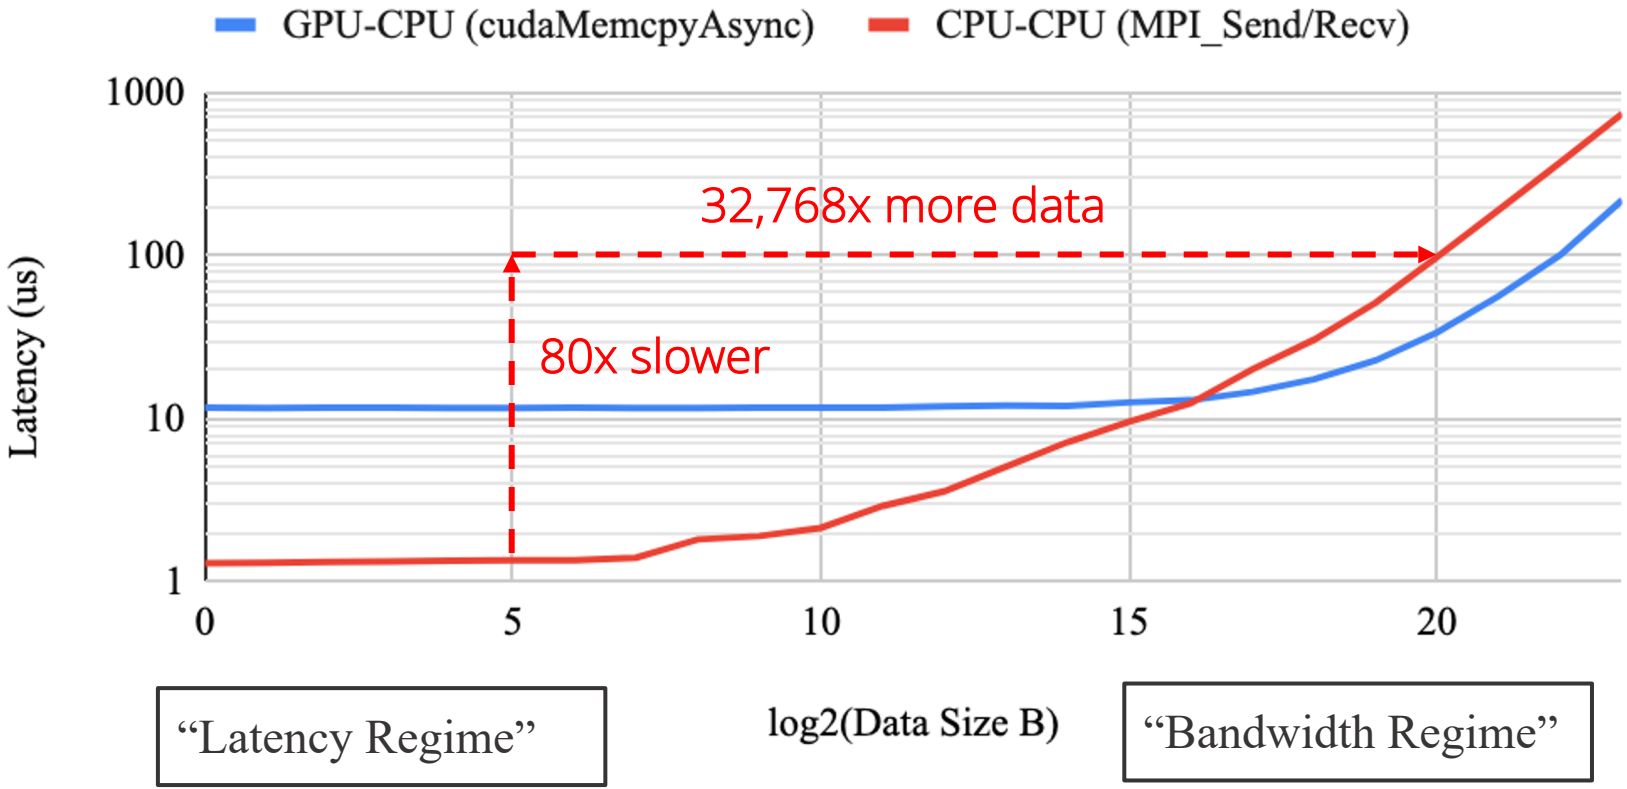

CUDA-Aware MPI + Packing

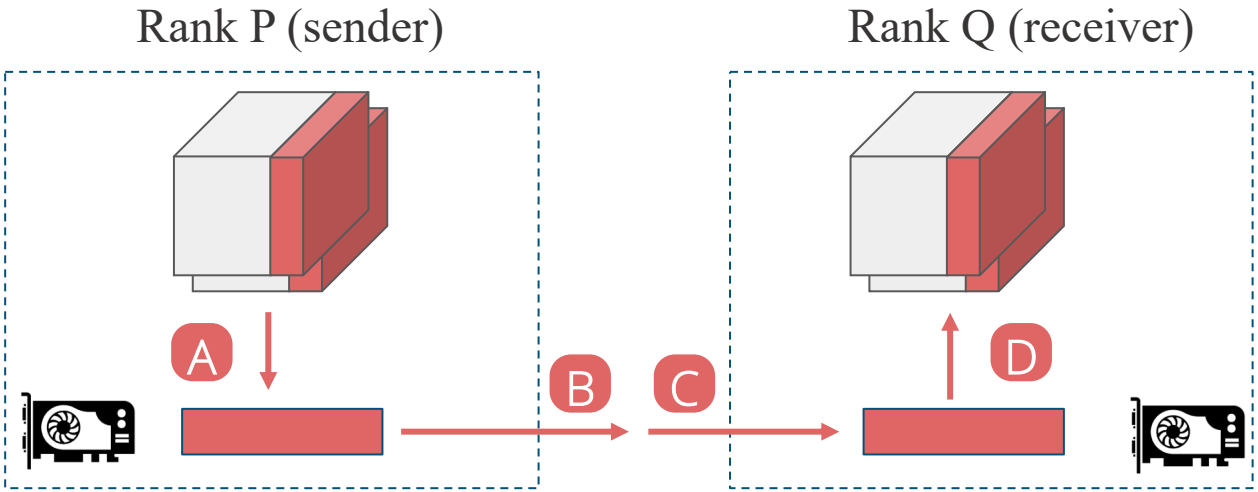

MPI Derived Datatypes

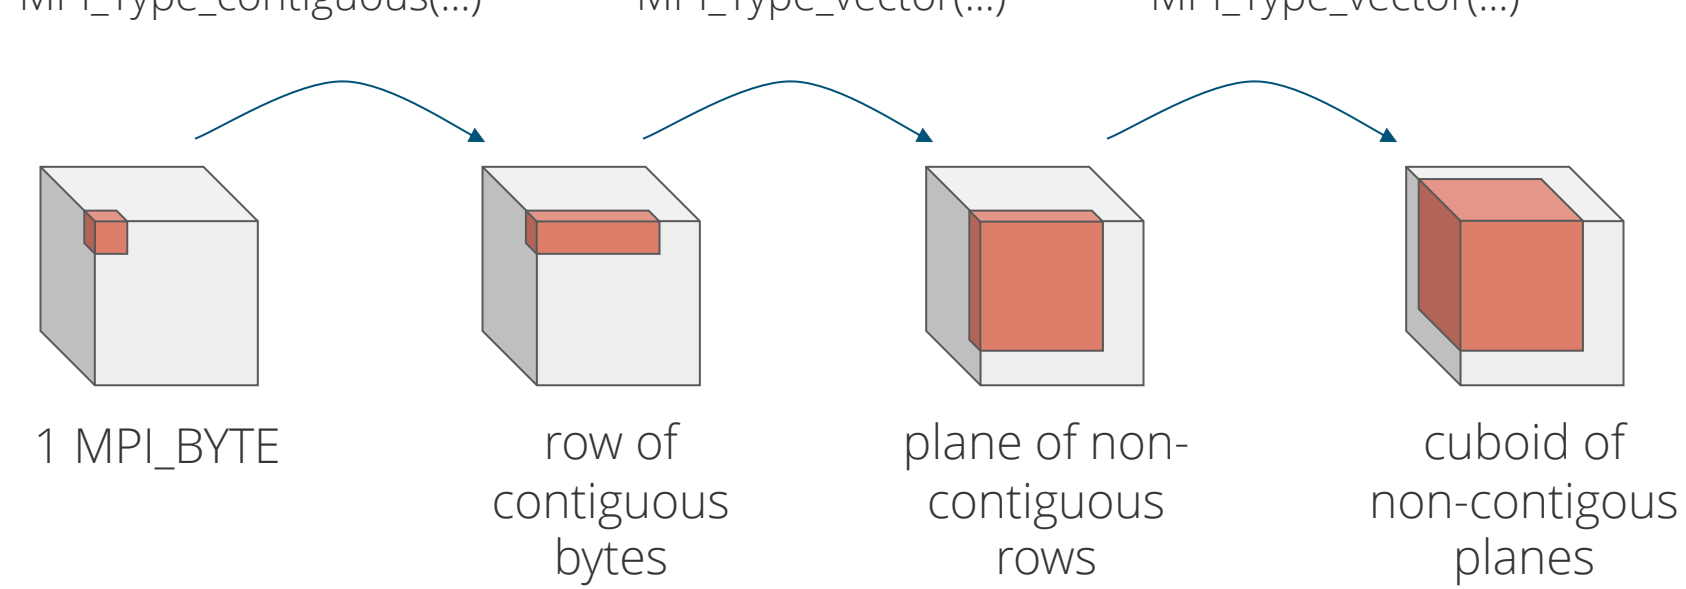

CUDA-Aware MPI + MPI Derived Datatypes

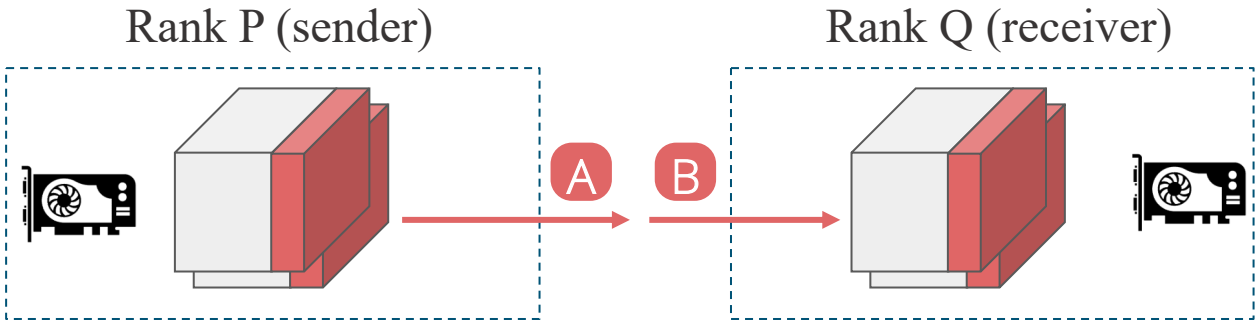

Setup (once):

```
MPI_Type...
MPI_Type...
...
MPI_Type_commit
```

```
MPI_Type...
MPI_Type...
...
MPI_Type_commit
```

Each halo exchange:

A `MPI_Isend`

B `MPI_Irecv`

“alltoallv” Halo Exchange

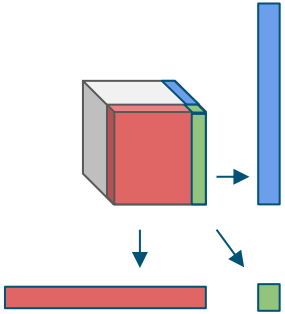

MPI_Packs

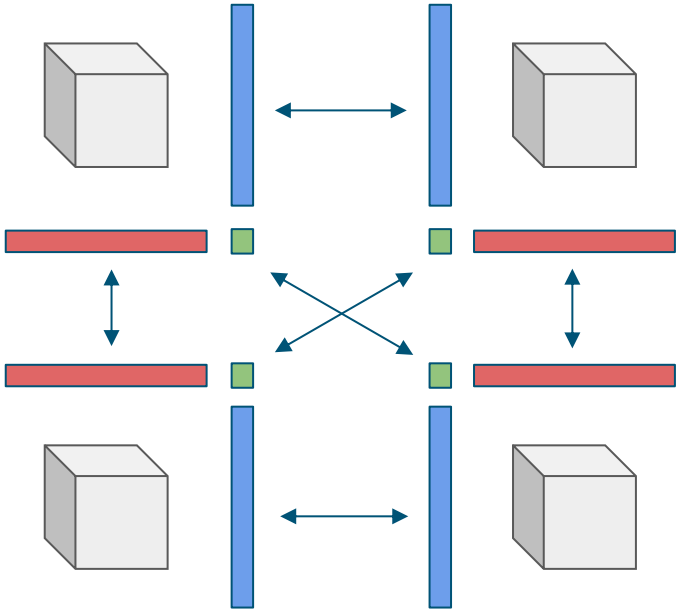

MPI_Neighbor_alltoallv

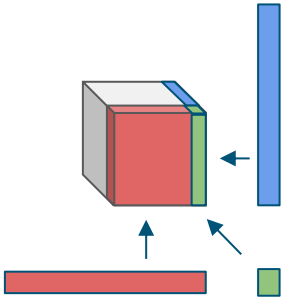

MPI_Unpacks

“alltoallv” with MPI derived types

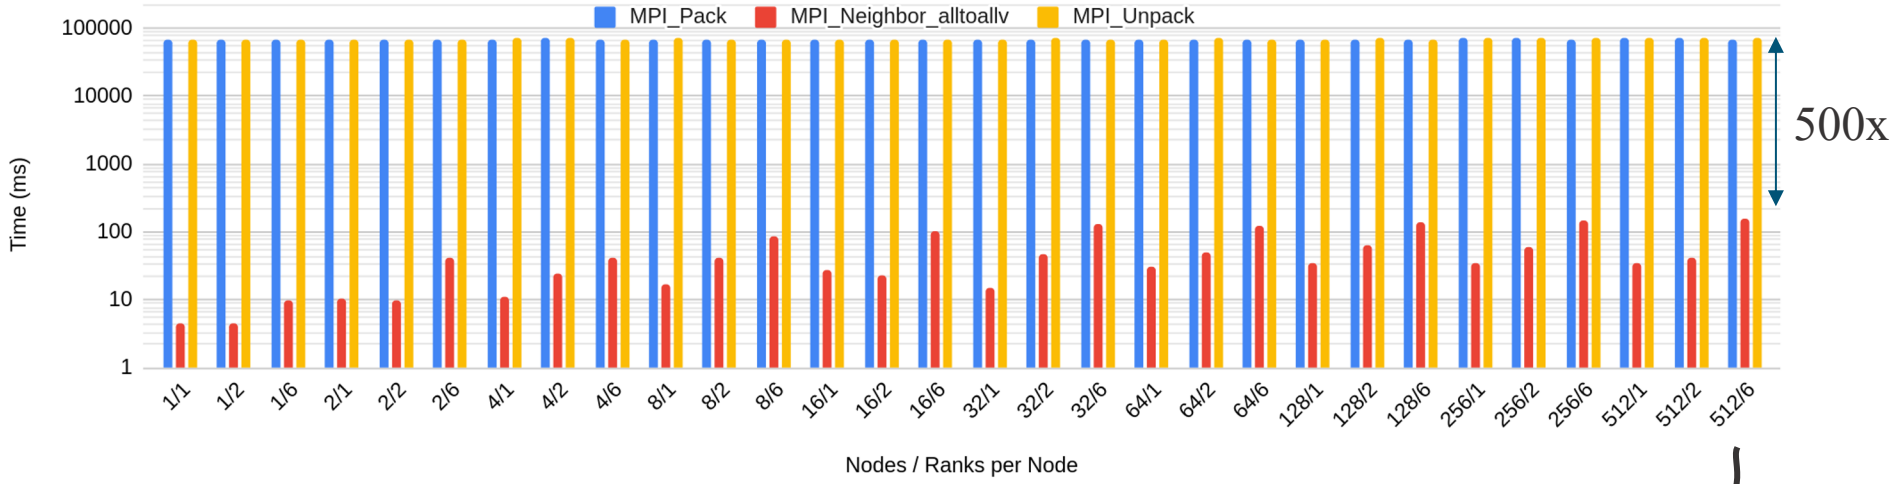

- MPI_Neighbor_alltoallv = ~500 MB/s/rank
- MPI_Pack / MPI_Unpack = **~1 MB/s**

MPI_Send (Spectrum MPI)

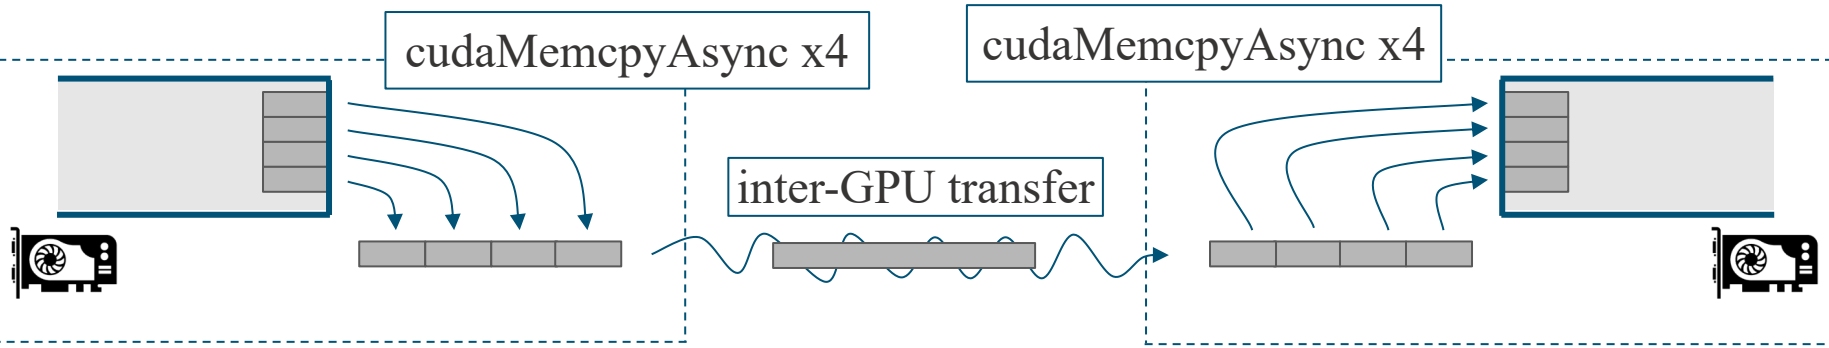

❌ “Latency Regime”

✅ “Bandwidth Regime”

❌ “Latency Regime”

Better MPI_Send (common foundation)

- GPU kernels to pack non-contiguous data

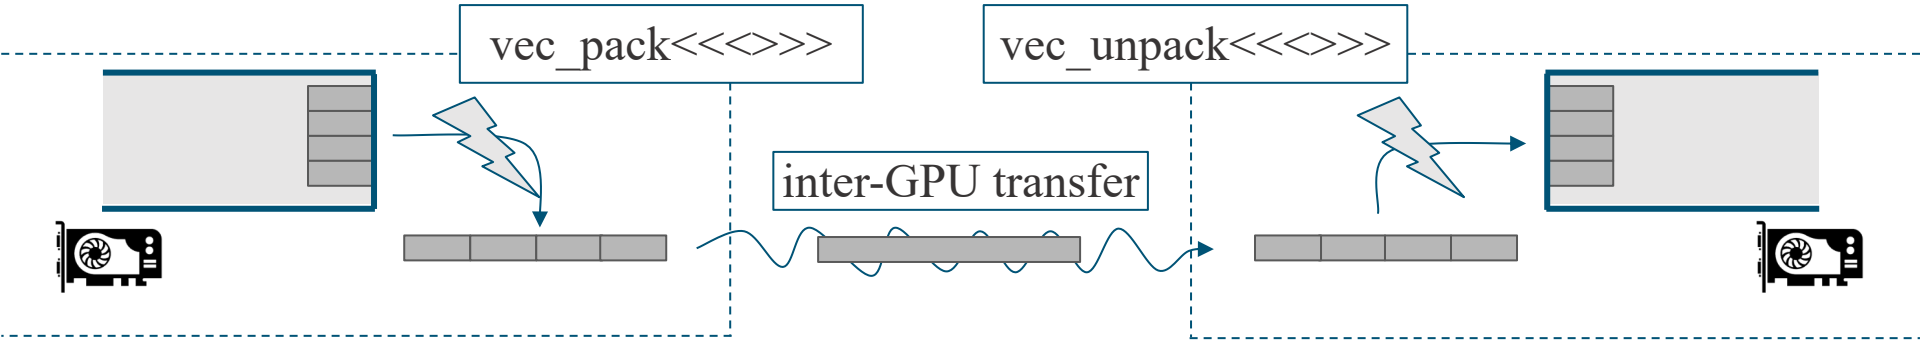

✅ “Bandwidth Regime”

✅ “Bandwidth Regime”

✅ “Bandwidth Regime”

MPI Derived Datatype Equivalence

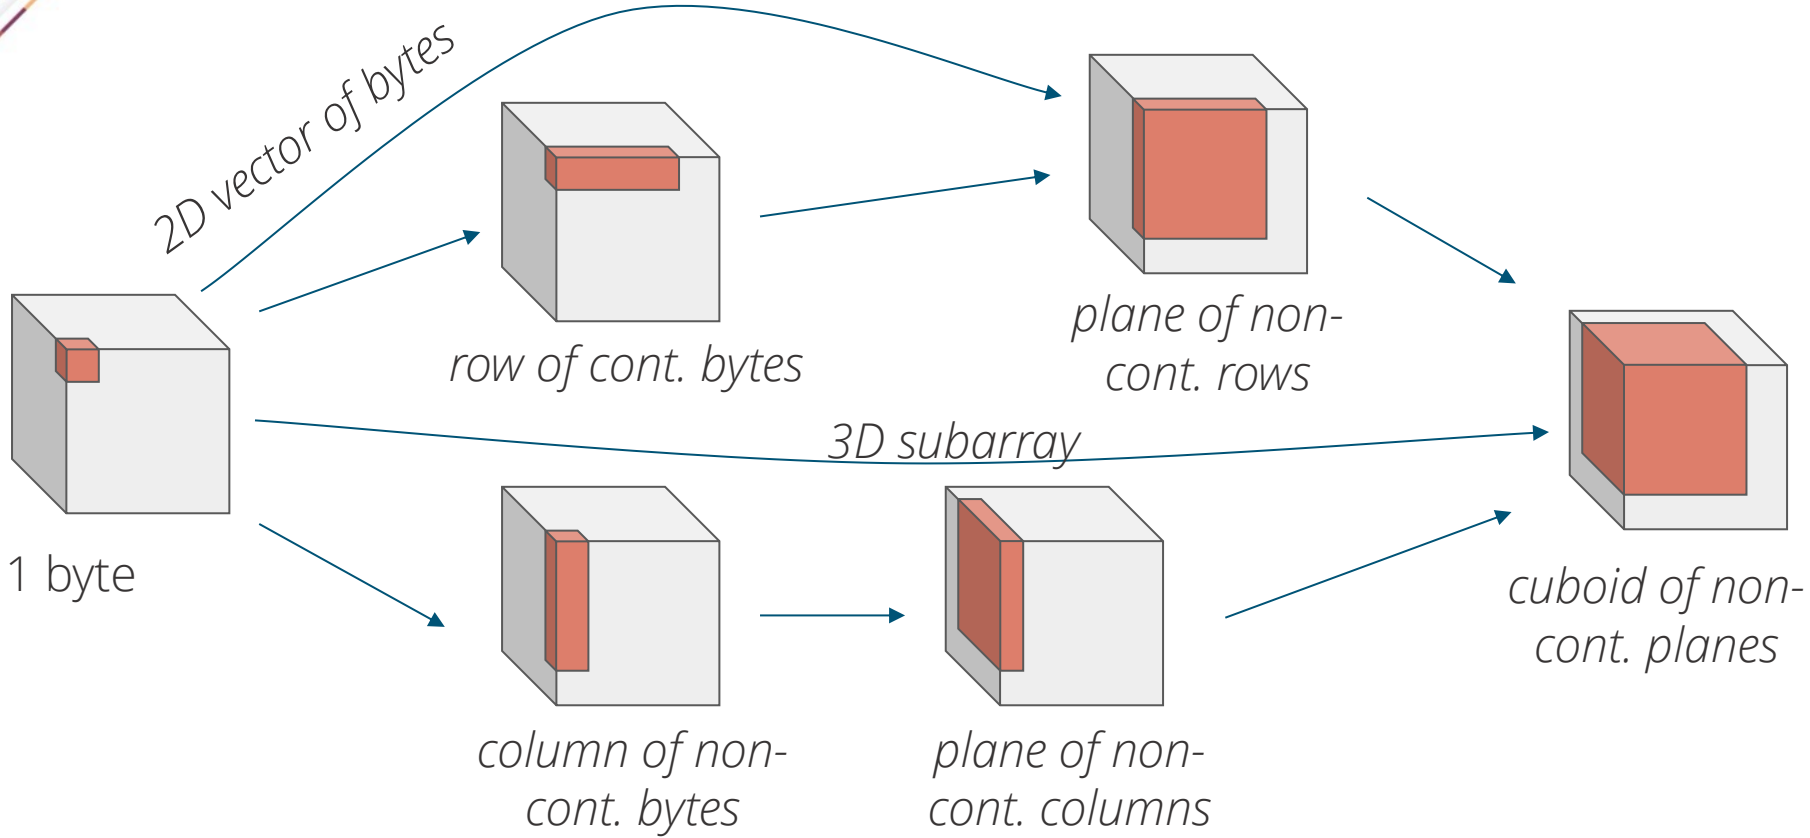

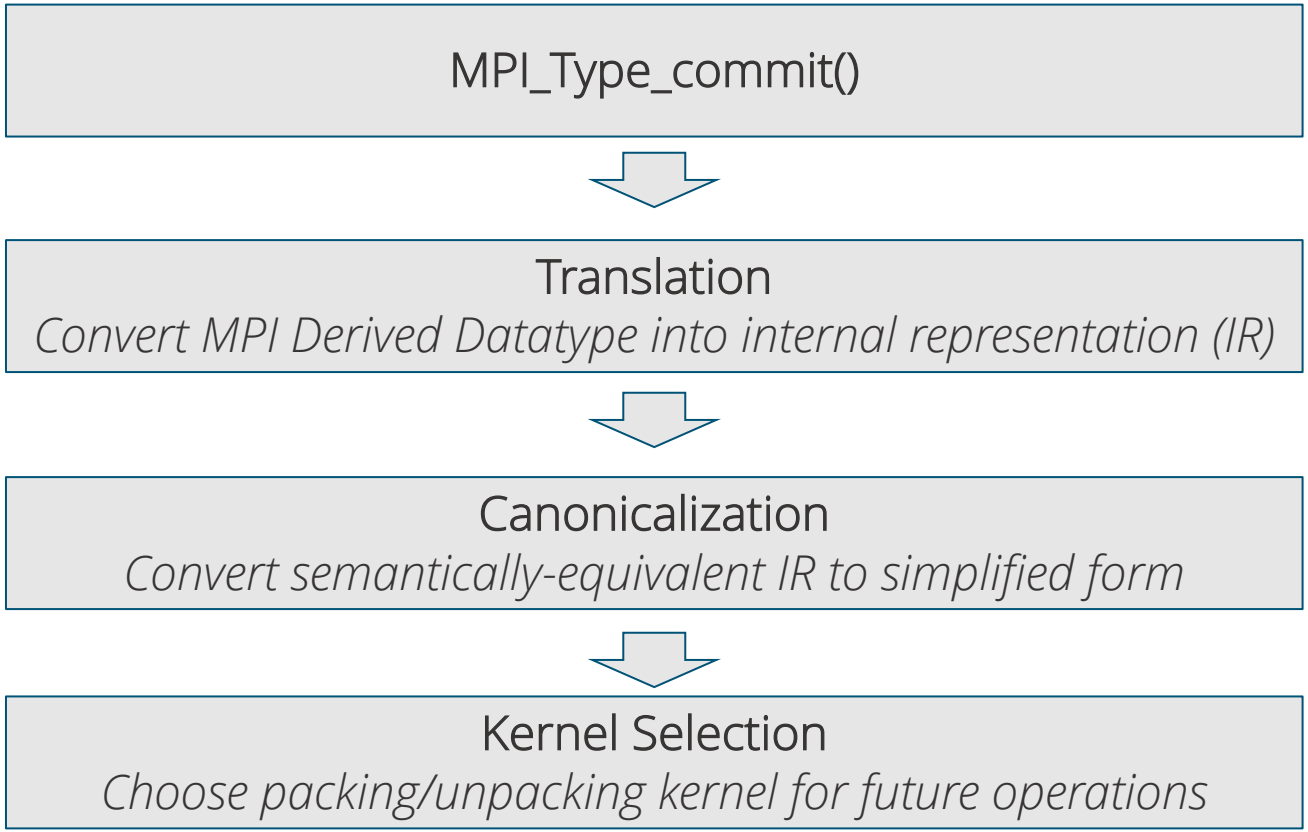
MPI_Type_commit()
Translation
Convert MPI Derived Datatype into internal representation (IR)
Canonicalization
Convert semantically-equivalent IR to simplified form
Kernel Selection
Choose packing/unpacking kernel for future operations

IR - “Internal Representation”

```
StreamData {
  integer offset; // offset (B) of the first element
  integer stride; // pitch (B) between element
  integer count;  // number of elements
}

DenseData {
  integer offset; // offset (B) of the first byte
  integer extent; // number of bytes
}
```

Hierarchy of StreamData, rooted at DenseData

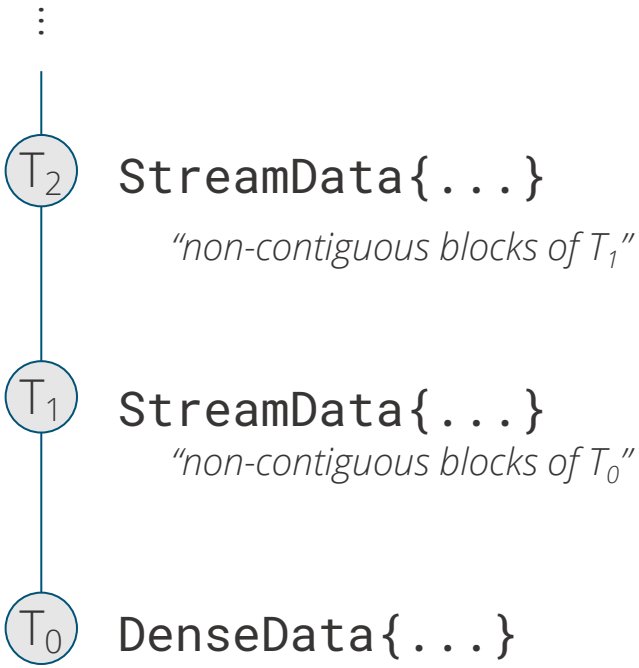

Example

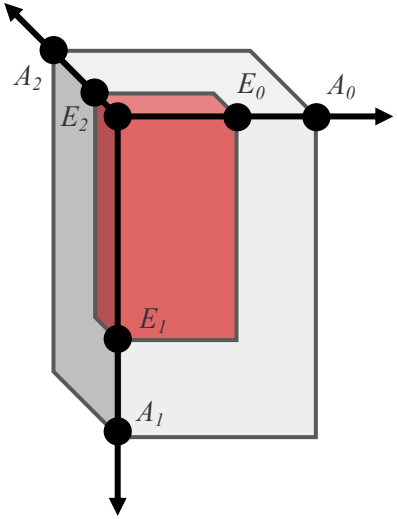

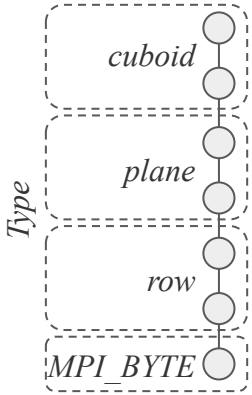


```
StreamData{offset:0, count:E2,   stride:A1*A0}
StreamData{offset:0, count:1,    stride:E1*A0}
StreamData{offset:0, count:E1,   stride:A0}
StreamData{offset:0, count:1,    stride:E0}
StreamData{offset:0, count:E0, stride:1}
StreamData{offset:0, count:1,    stride:1}
DenseData{offset:0,  extent: 1}
```

TEMPI Datatype Handling

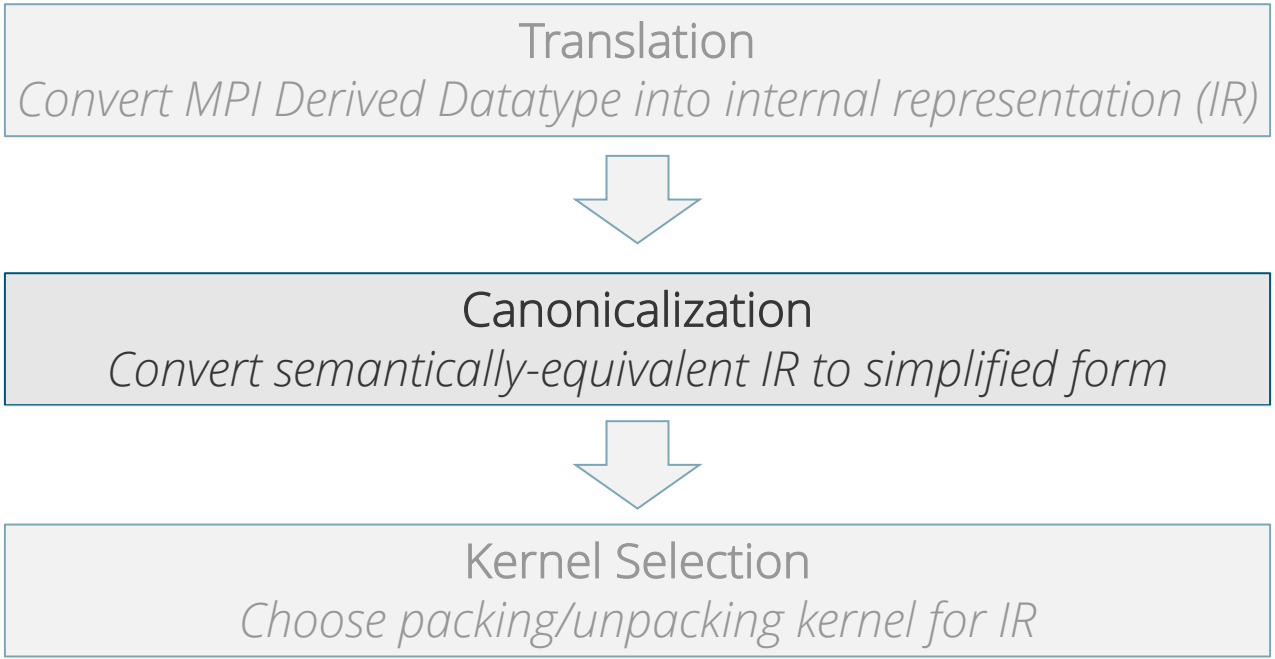

Canonicalization

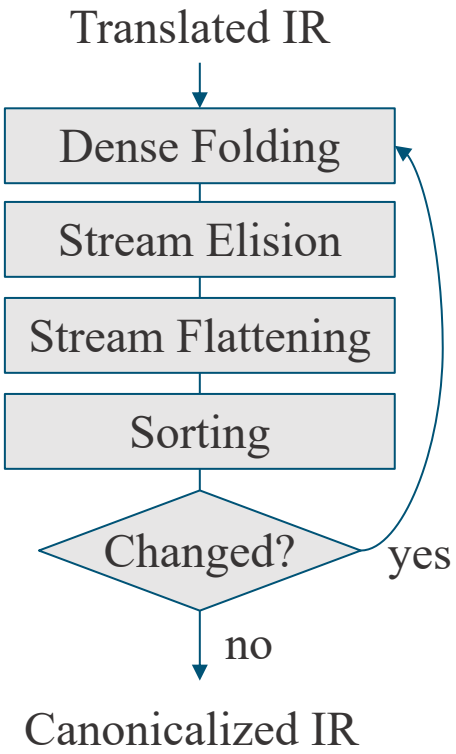

Canonicalization: Dense Folding

- `StreamData` is contiguous `DenseData` (parent of MPI named type)

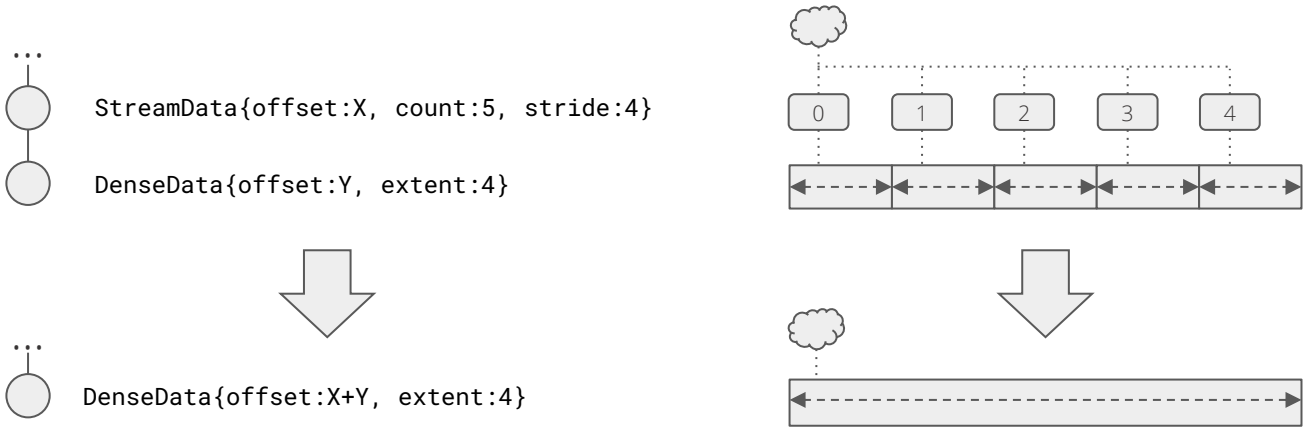

Canonicalization: Stream Elision

- Remove `StreamData` with `count = 1` (MPI Vector blocks commonly have one element)

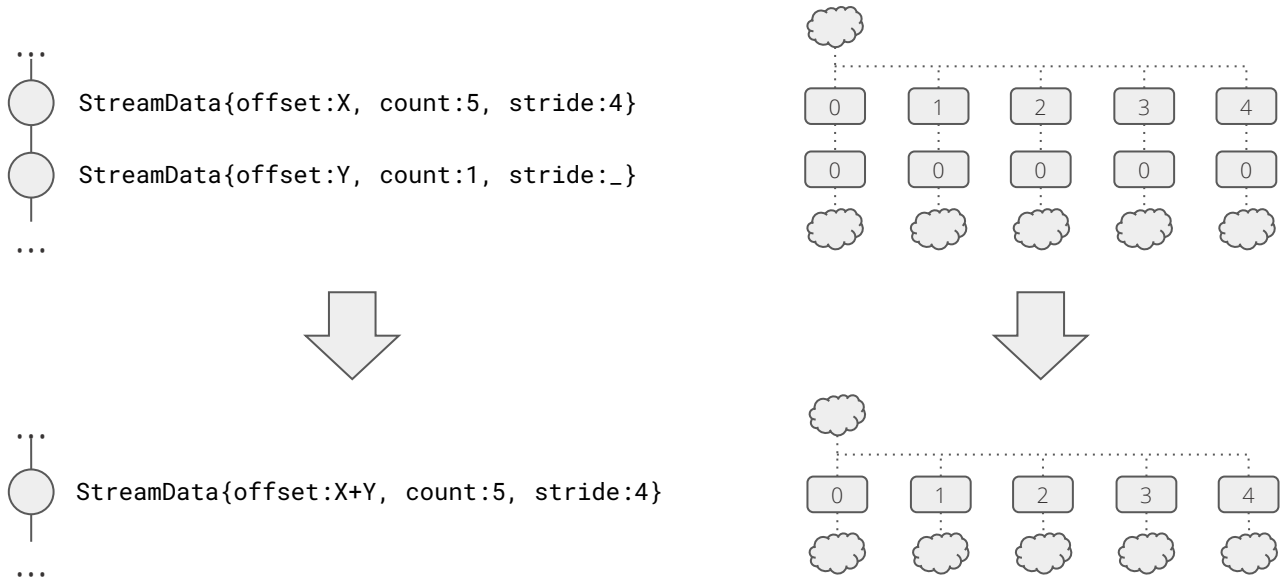

Canonicalization: Stream Flattening

- e.g. two vectors of three vs. one vector of six

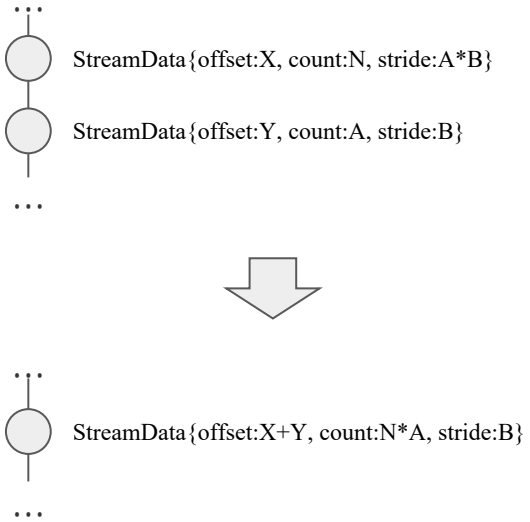

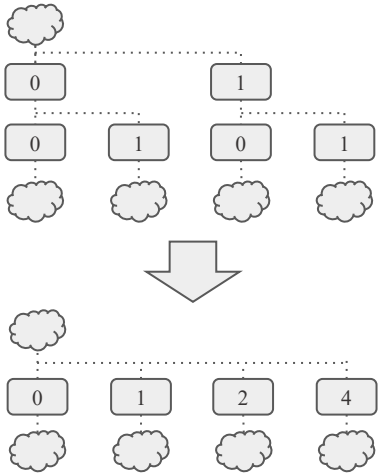

TEMPI Datatype Handling

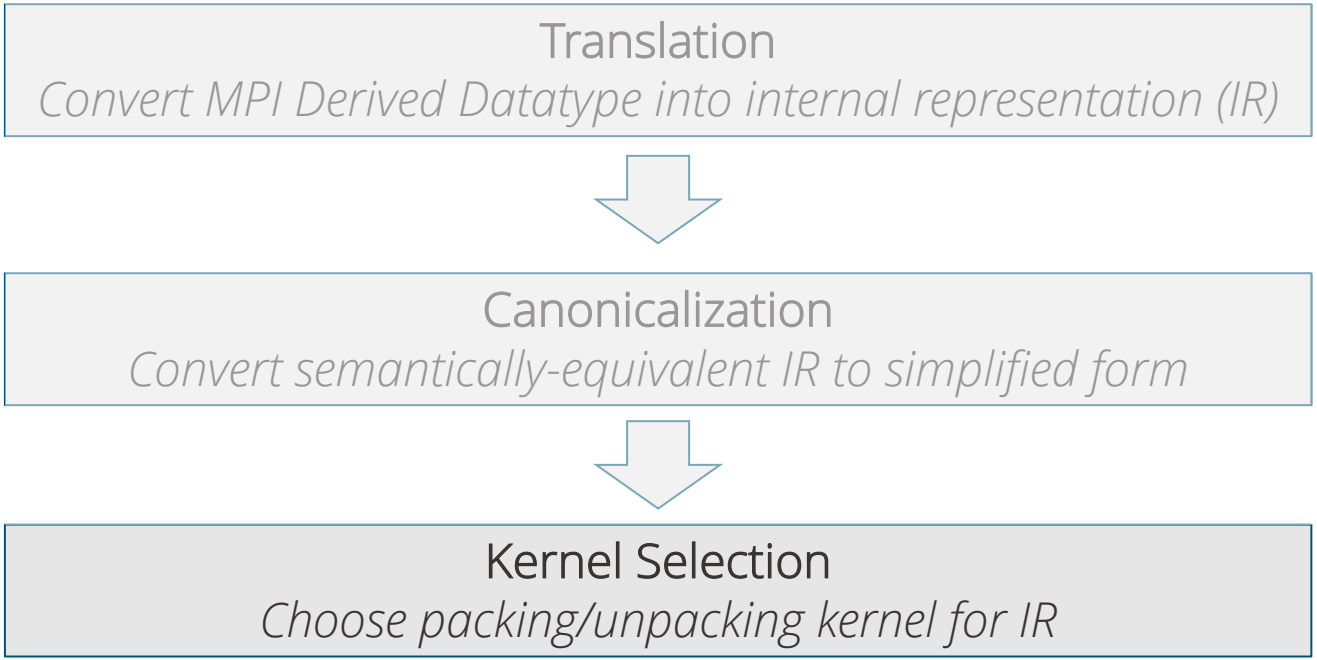

MPI_Type_commit Time

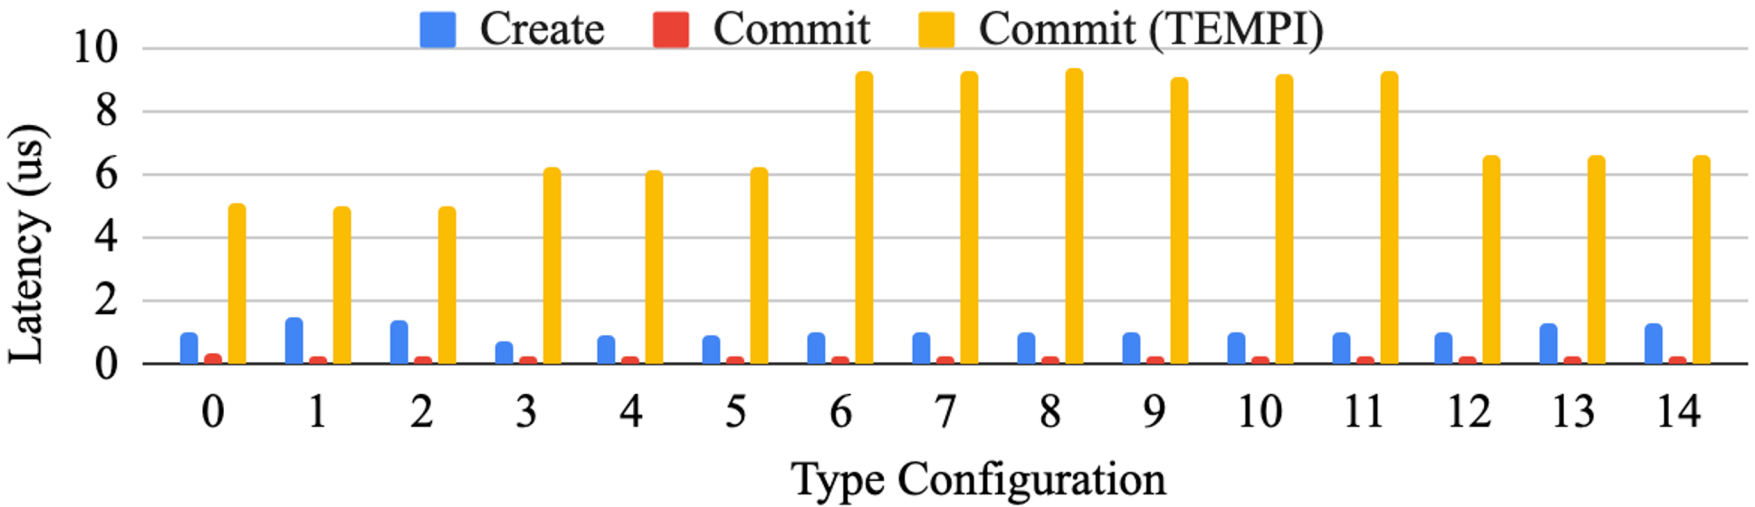

MPI_Pack Results

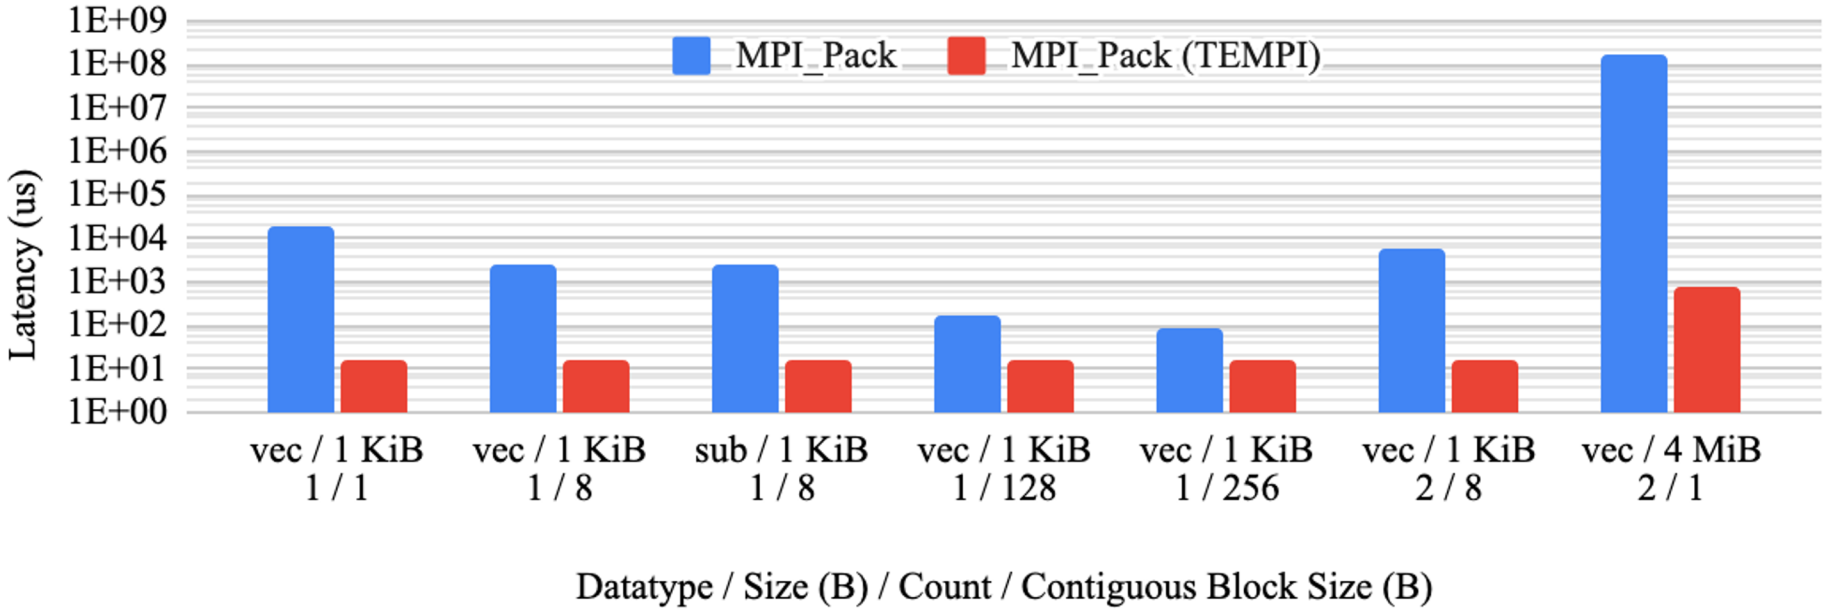

CUDA-Aware System MPI

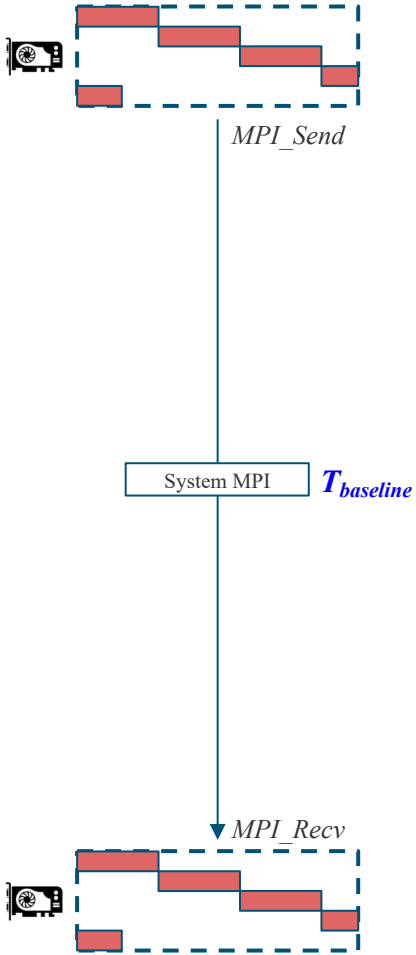

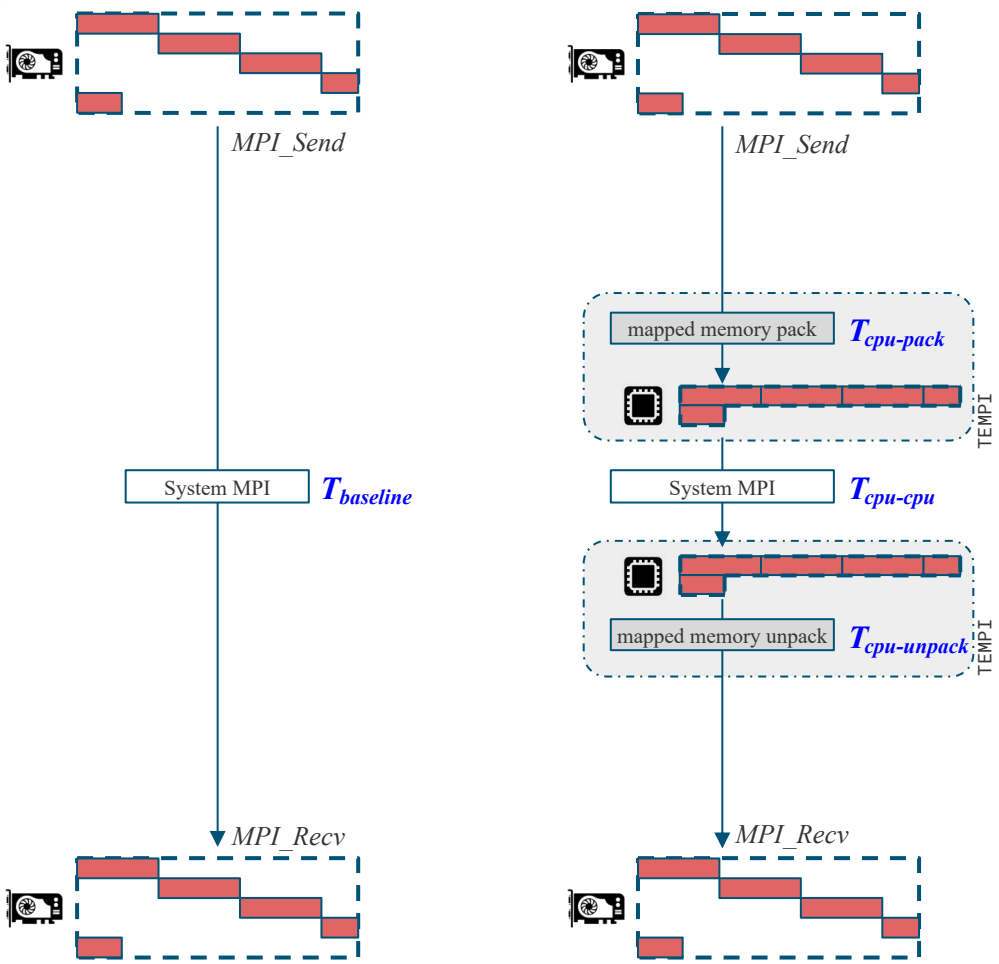
CUDA-Aware System MPI
TEMPI "One-shot"
MPI_Send
MPI_Send
mapped memory pack
$T_{cpu\text{-}pack}$
TEMPI
System MPI
$T_{baseline}$
System MPI
$T_{cpu\text{-}cpu}$
mapped memory unpack
$T_{cpu\text{-}unpack}$
TEMPI
MPI_Recv
MPI_Recv

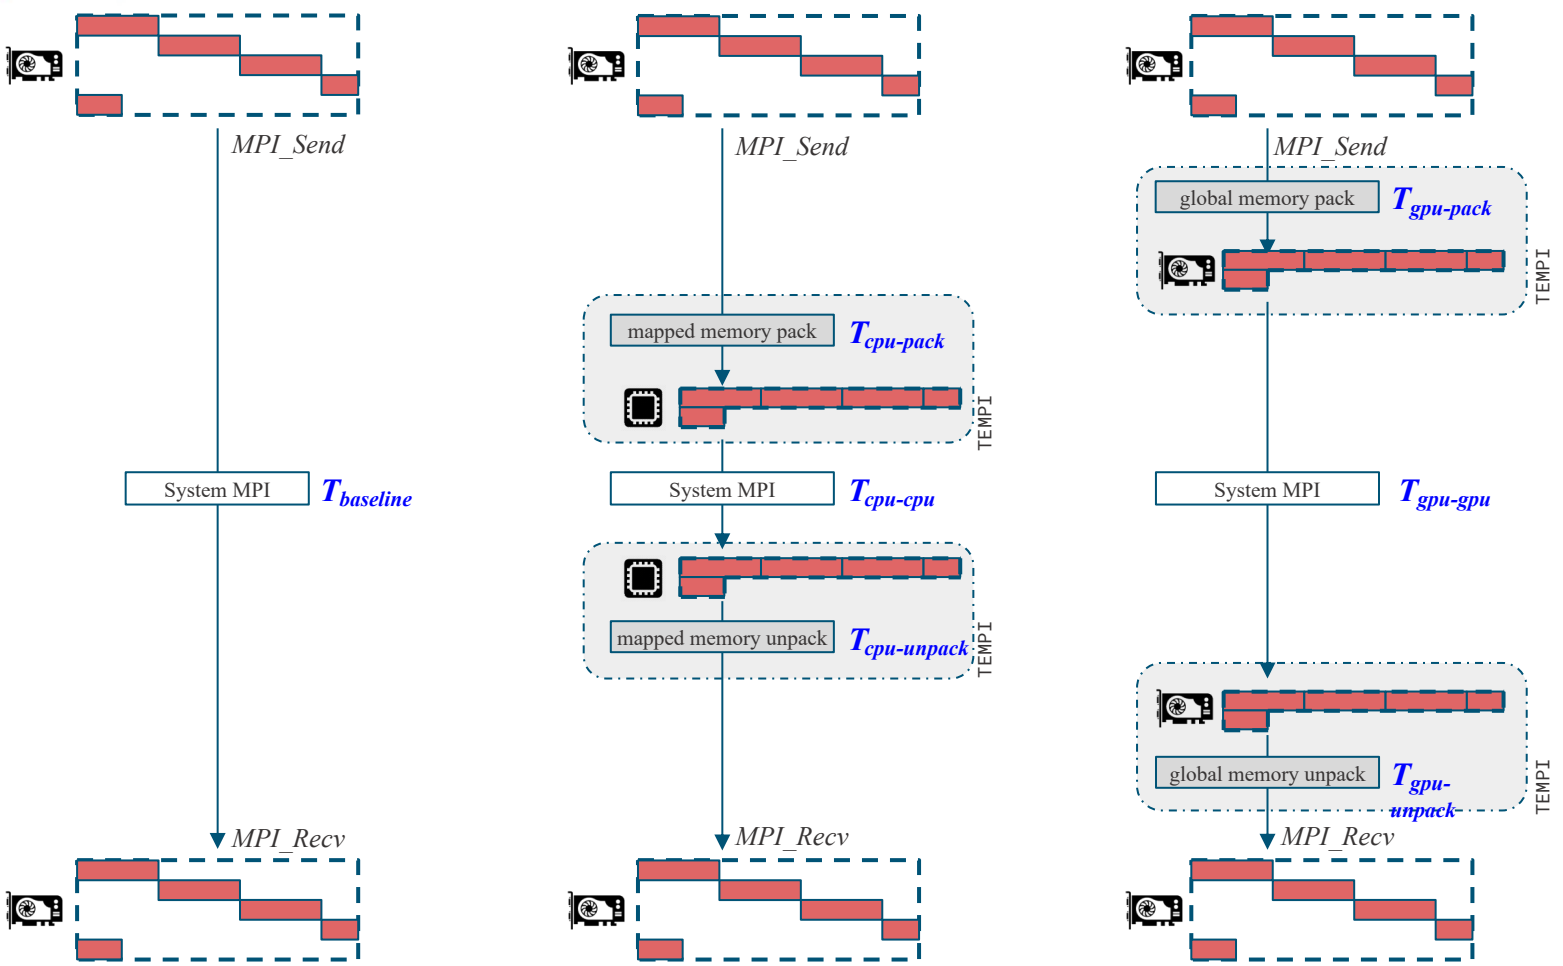
CUDA-Aware System MPI
MPI_Send
System MPI
$T_{baseline}$
MPI_Recv
TEMPI "One-shot"
MPI_Send
mapped memory pack
$T_{cpu-pack}$
TEMPI
System MPI
$T_{cpu-cpu}$
mapped memory unpack
$T_{cpu-unpack}$
TEMPI
MPI_Recv
TEMPI "Device"
MPI_Send
global memory pack
$T_{gpu-pack}$
TEMPI
System MPI
$T_{gpu-gpu}$
global memory unpack
$T_{gpu-unpack}$
TEMPI
MPI_Recv

Performance Modeling

$$T_{device} = T_{gpu-pack} + T_{gpu-gpu} + T_{gpu-unpack}$$

$$T_{oneshot} = T_{host-pack} + T_{cpu-cpu} + T_{host-unpack}$$

$$T_{staged} = T_{gpu-pack} + T_{d2h} + T_{cpu-cpu} + T_{h2d} + T_{gpu-unpack}$$

Contiguous Transfer Measurements

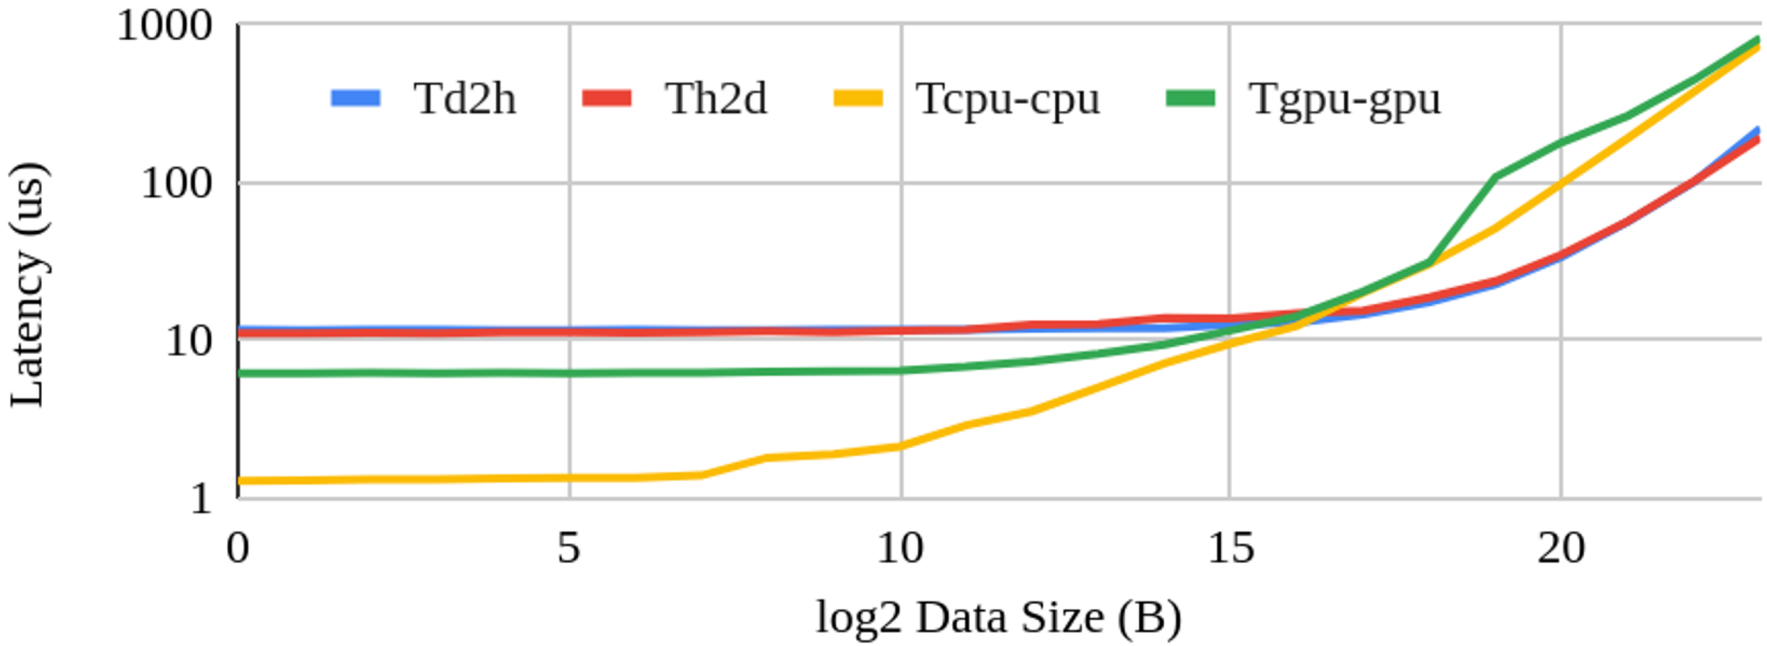

T depends on # bytes transferred

One-shot and Device Pack and Unpack

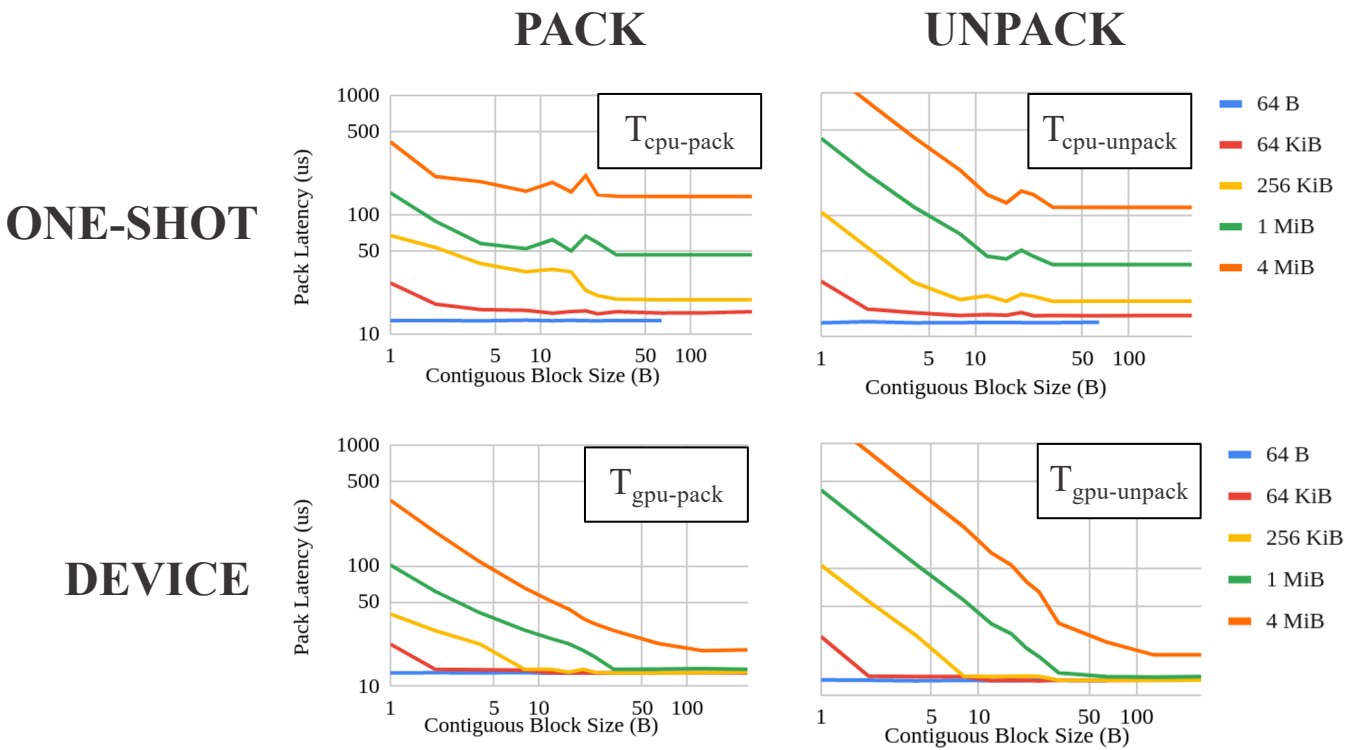

T depends on object size and block size

MPI_Send Microbenchmark

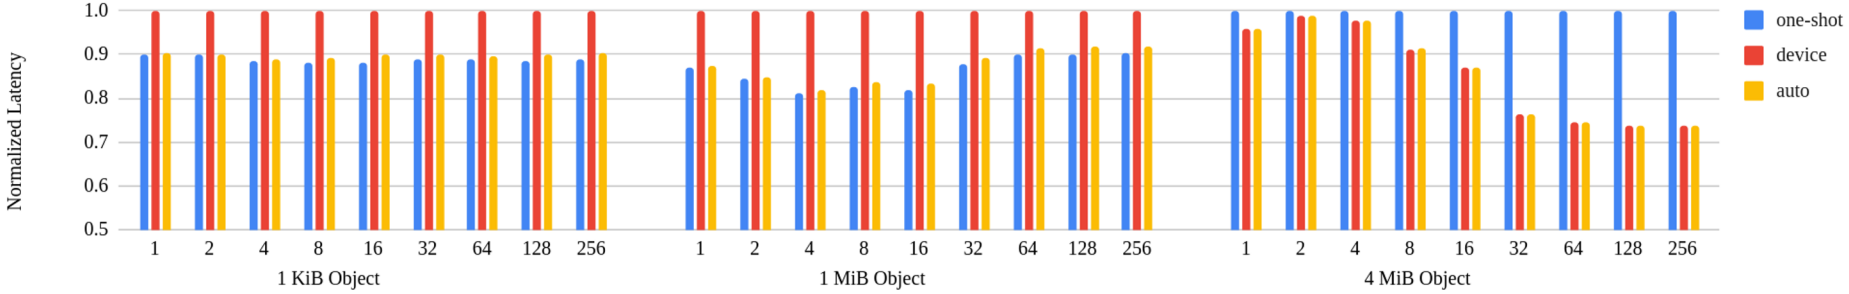

Minimal runtime performance-modeling overhead
Performance model reliably chooses faster method

MPI_Send / MPI_Recv

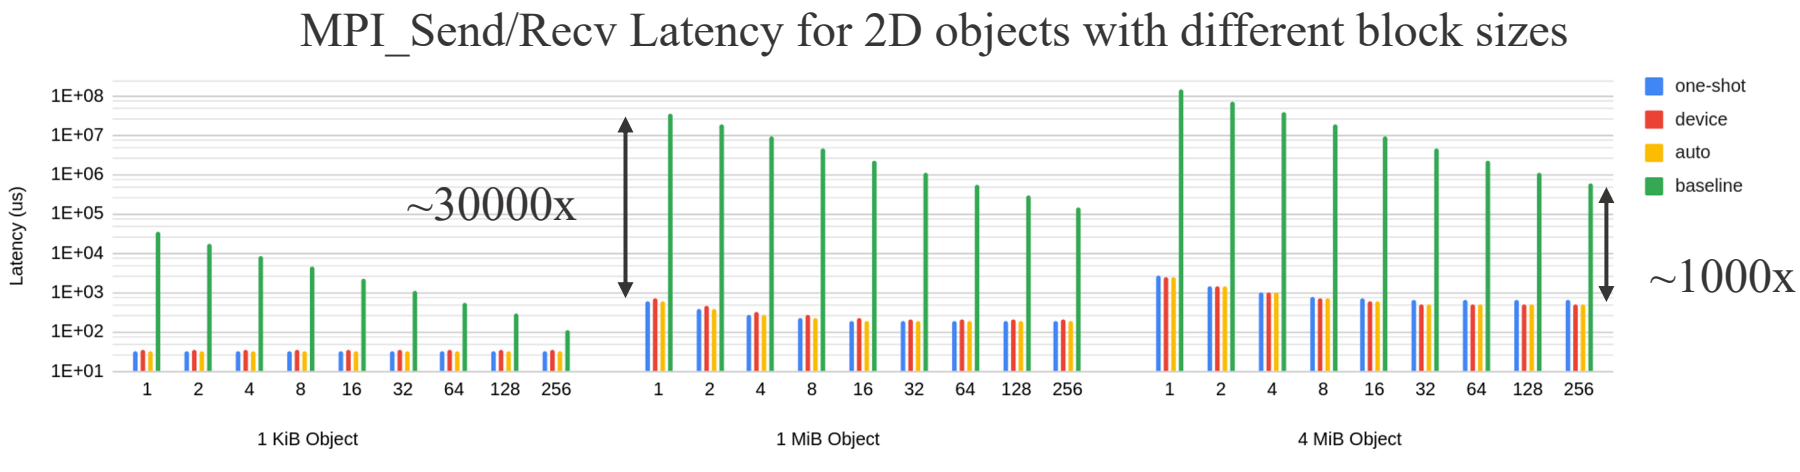

Large latency improvement from datatype handling
Small additional improvement from automatic method selection

Halo Exchange

With TEMPI, non-contiguous packing no longer dominates

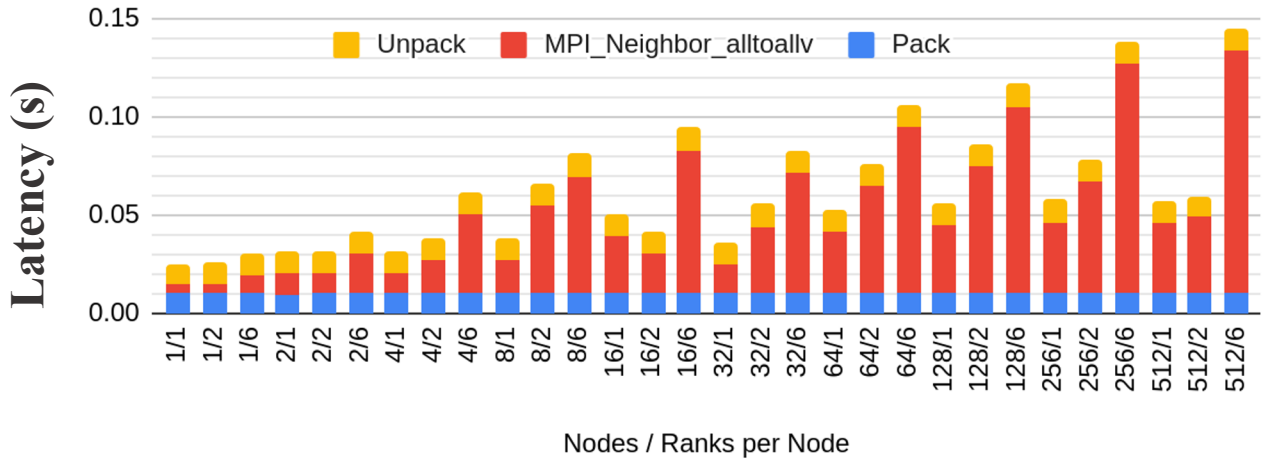

Large speedup in 3D halo exchange

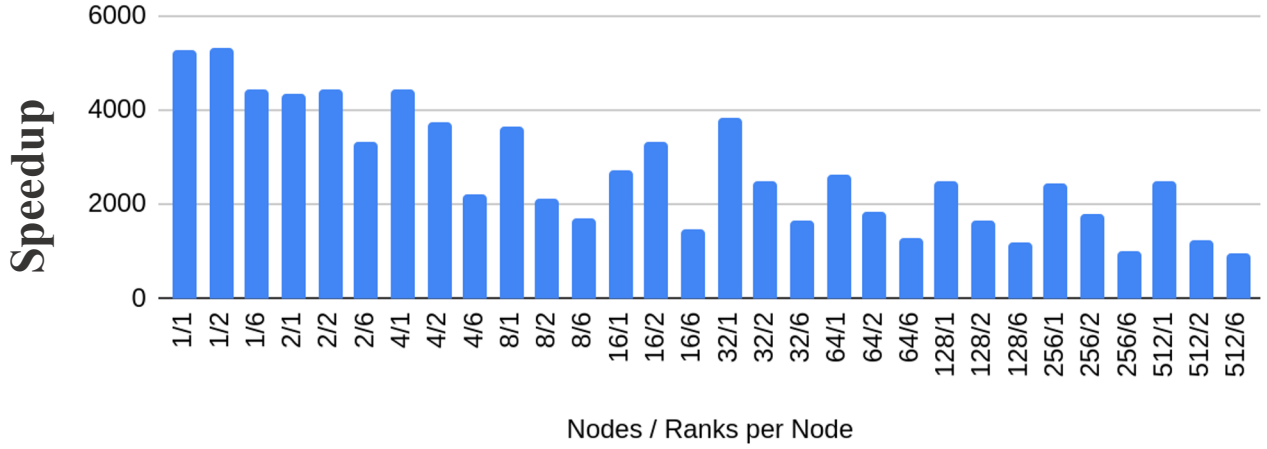

A Practical Challenge: Deploying MPI experiments on a Supercomputer

- Large-scale systems are tightly controlled
  - Can’t just make whatever changes you want
- Usually one MPI (or maybe two) are deployed on the system
  - Rarely bugfixed (if ever)
  - Even more rarely are new features added
- Difficult or impossible to make experimental modifications

- MPI has a well-defined standard
  - Take advantage of this + how OS loads libraries to inject modifications

TEMPI's Architecture: Interposer Library

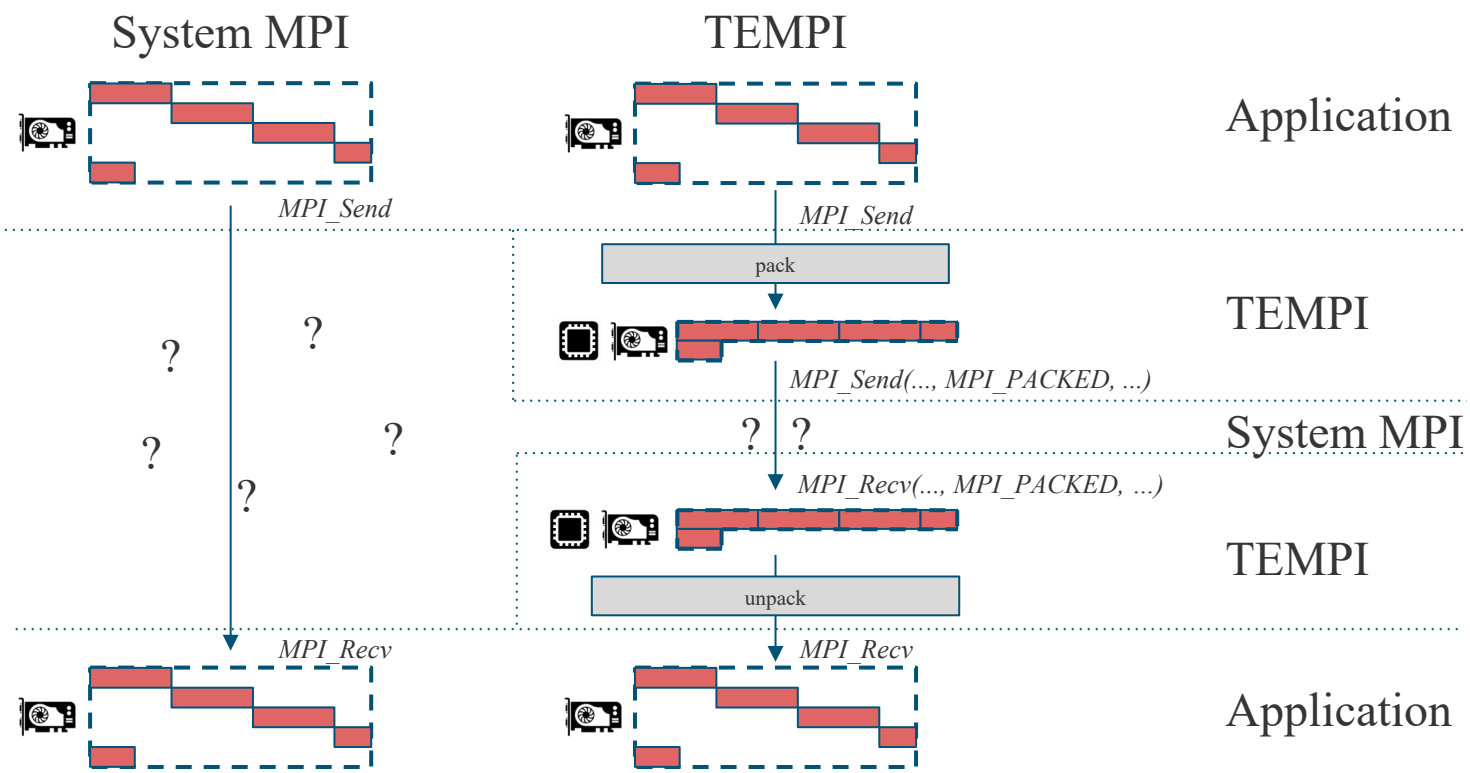

app.c

```
#include <mpi.h>

int main(int argc, char **argv) {
  MPI_Init(&argc, &argv);
}
```

1

```
gcc app.c -o app  \
  -I /mpi/include \
  -L /mpi/lib     \
  -l mpi
```

2

app

```
...
callq  7260
<MPI_Init@plt>
...
```

4

8

system MPI

mpi.h

```
int MPI_Init(...);
```

3

libmpi.so

```
...
0x86fb0 W MPI_Init
...
```

5

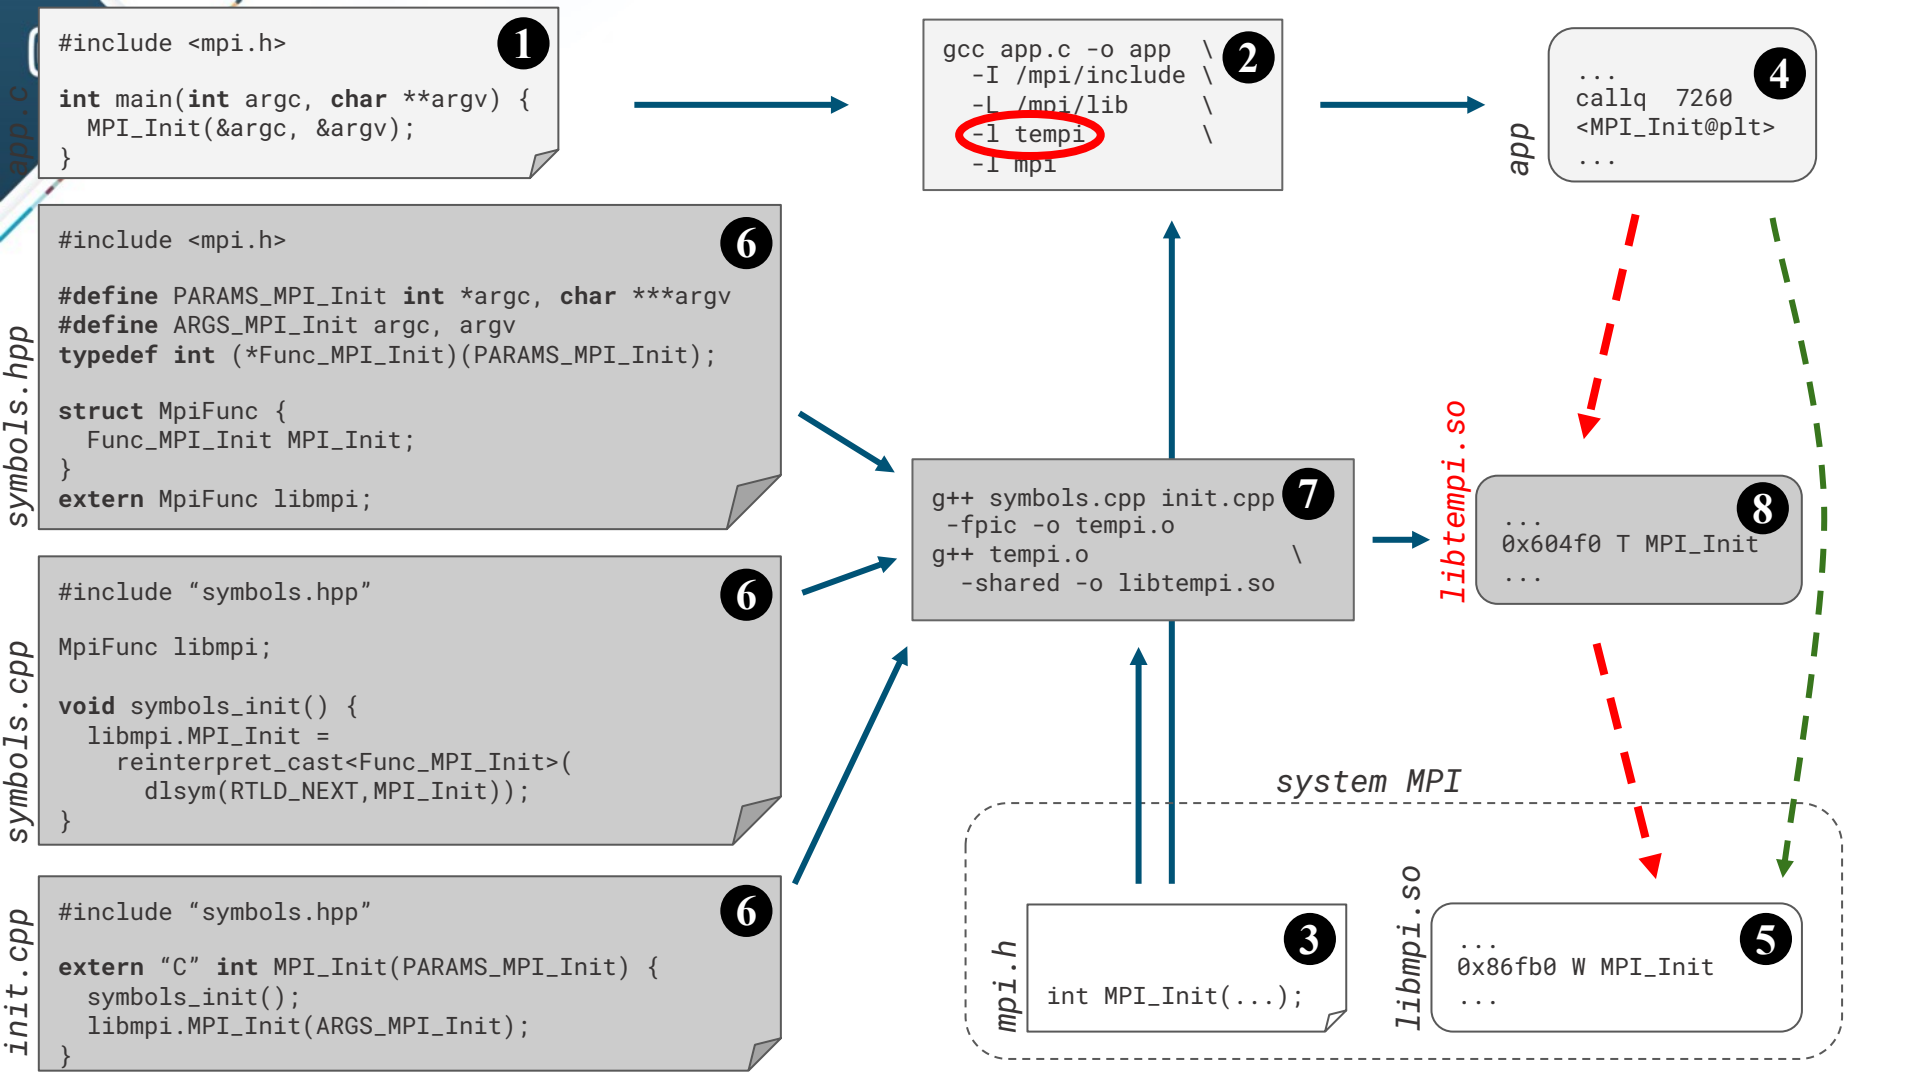
app.c
1
#include <mpi.h>
int main(int argc, char **argv) {
MPI_Init(&argc, &argv);
}
gcc app.c -o app \
-I /mpi/include \
-L /mpi/lib \
-l tempi \
-l mpi
2
app
4
...
callq 7260
<MPI_Init@plt>
...
symbols.hpp
6
#include <mpi.h>
#define PARAMS_MPI_Init int *argc, char ***argv
#define ARGS_MPI_Init argc, argv
typedef int (*Func_MPI_Init)(PARAMS_MPI_Init);
struct MpiFunc {
Func_MPI_Init MPI_Init;
}
extern MpiFunc libmpi;
symbols.cpp
6
#include "symbols.hpp"
MpiFunc libmpi;
void symbols_init() {
libmpi.MPI_Init =
reinterpret_cast<Func_MPI_Init>(
dlsym(RTLD_NEXT,MPI_Init));
}
init.cpp
6
#include "symbols.hpp"
extern "C" int MPI_Init(PARAMS_MPI_Init) {
symbols_init();
libmpi.MPI_Init(ARGS_MPI_Init);
}
g++ symbols.cpp init.cpp \
-fpic -o tempi.o
g++ tempi.o \
-shared -o libtempi.so
7
libtempi.so
8
...
0x604f0 T MPI_Init
...
system MPI
mpi.h
3
int MPI_Init(...);
libmpi.so
5
...
0x86fb0 W MPI_Init
...

Conclusions & Future Work

- Canonicalization approach works
  - 👍 Speedup for unmodified applications on OLCF summit
  - 👍 Any strided datatype
- Simple performance model to select GPU data transfer method
  - 👍 speedup
  - 👎 < 2x speedup
- Interposed library approach nice for experiments & prototype deployment
  - 👍 Easy to use without system privileges
  - 👎 Limited integration with existing MPI
- OpenMPI, MPICH, MVAPICH have other strategies
  - Use if available - covers some or all of the same problems
  - Comparison of MPI datatype performance would be useful to community
- github.com/cwpearson/tempi

Acknowledgements

This work is supported by IBM-ILLINOIS Center for Cognitive Computing Systems Research (C3SR) - a research collaboration as part of the IBM AI Horizon Network.

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy Contract No. DE-AC05-00OR22725.

This work utilizes resources supported by the National Science Foundation's Major Research Instrumentation program, grant #1725729, as well as the University of Illinois at Urbana-Champaign

Thank you

cwpears@sandia.gov

github.com/cwpearson/tempi

This work was completed prior to joining Sandia National Labs.