Chizu: A Framework to Enable Topology Aware Task Mapping

Harshitha Menon*, Abhinav Bhatele, Adam Kunen, Katie Lewis, Peter Robinson, Ulrike Meier Yang, Laxmikant V. Kale*

Center for Applied Scientific Computing & *University of Illinois at Urbana-Champaign

Network communication can become a bottleneck when scaling applications to large parallel machines. Topology aware task mapping can improve performance of parallel applications significantly by reducing the congestion on the network. We have developed a generic framework, called Chizu, which facilitates mapping of tasks. It uses the communication graph of the application and the network topology of the machine to optimize network communication. Chizu can exploit various graph partitioning and mapping libraries to perform mapping. We present our preliminary work on studying the communication and mapping of applications using Chizu.

Motivation

- Performance of communication bound applications is affected by latency and bandwidth
- Mapping tries to minimize network communication
- Wanted a generic framework to use and test existing and new mapping algorithms
- Wanted a simple library to support plug and play of different mapping strategies

Chizu

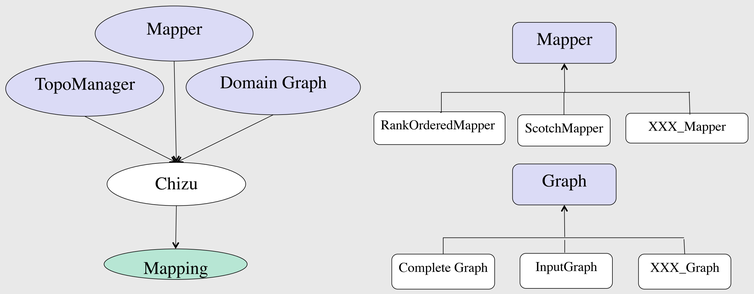

Framework to facilitate mapping of tasks.

- Can exploit many mapping algorithms
- Many applications can use this common library
- Can be used as a library or standalone binary
- The mapping produced can be used with srun:

srun --launcher-opts='--mapping map.txt' pgm

```
#include "chizu.h"

 {
Graph* dg; // Domain graph
Graph* ag; // Architecture graph
Mapper* m;
size_t* mapping;

dg = CreateGraphFromFile("foo");
ag = CreateGraphFromFile("bar");
m = CreateMapper(SCOTCH);
m->Map(m, dg, ag, mapping);
}
```

AMG Benchmark

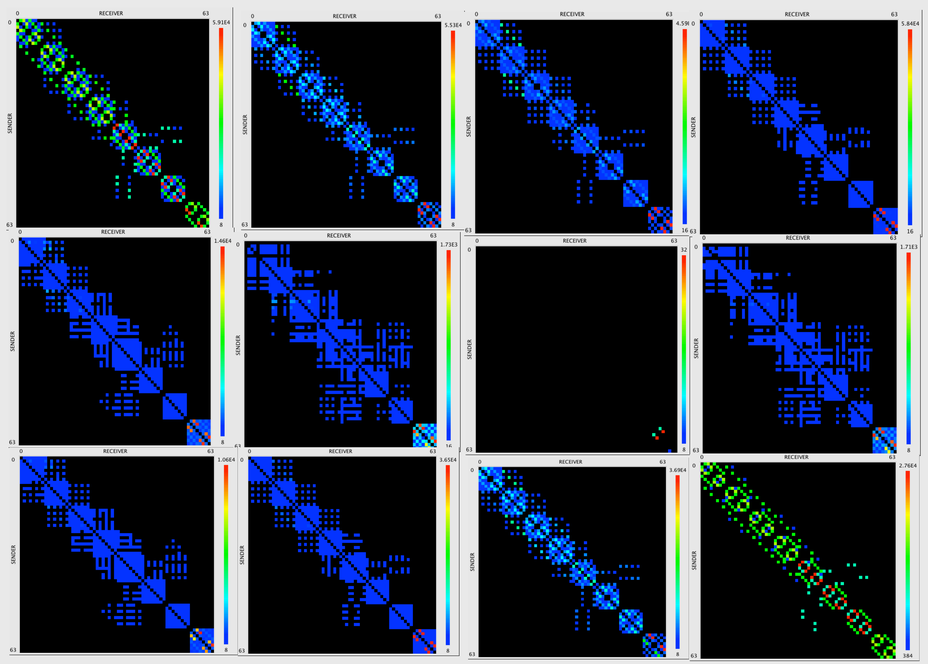

Communication pattern at each level of AMG solve

AMG2013 is a parallel algebraic multigrid solver for linear systems derived from BoomerAMG solver developed at LLNL. AMG has a multi-level communication pattern where the grid is coarsened at each level.

Domains	Average Hops	
	Default Mapping	Scotch Mapping
8192	2.0	1.2

Mapping results for AMG

Application 1

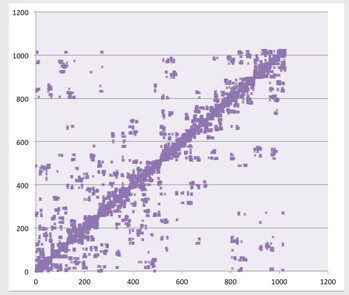

Domains	Average Hops	
	Default Mapping	Scotch Mapping
8192	2.2	1.9
16384	2.5	2.2
32768	2.7	2.1
65536	3.1	2.2

Mapping results

Summary

- Implemented a library, Chizu, which facilitates mapping of tasks given communication graph
- Chizu can be used as a library or standalone binary
- Chizu supports easy plug and play of different mapping strategies
- Studied communication pattern of AMG
- Used Chizu library on AMG

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. This work was funded by the Laboratory Directed Research and Development Program at LLNL under project tracking code 13-ERD-055 (LLNL-POST-641961).