SAND2021-10941C

Sandia National Laboratories

MiniMod: A Modular Miniapplication Benchmarking Framework for HPC

W. Pepper Marts

Matthew G. F. Dosanjh, Scott Levy,
Whit Schonbein, Ryan E. Grant, Patrick Bridges

U.S. DEPARTMENT OF ENERGY NNSA

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

MiniMod: Mission Justification

There is a wealth of new HPC middleware and app design

Yet HPC communication is dominated by MPI

- It is the conservative bet for performance
- Largely unchanged since MPI 2.0

Why are we not leveraging purported advancements?

- Robust exploration
- Demonstration of value
- You cannot devote enough app developer time to find out

MiniMod: How Do We Demonstrate Impact?

Proxy applications are the most popular solution

- Feasible for individuals and small teams
- Cover most application categories
- Demonstrably good evaluator

Proxies enable research, but there is room for improvement

- Comprehensive evaluation still requires multiple proxies
- Requires broad domain knowledge base

MiniMod: Evaluation Workflow Solution

We had several areas we wanted to improve

- Reduce repeated work
- Better cross comparability of results
- Utility for both application and middleware developers

MiniMod: a modular proxy application framework

- Runtime Configurable
- Compartmentalized Effort
- Low Overhead

MiniMod: Novel Benefits

Port once, test against dozens of configurations

- Leverage existing optimized work

Explore modules outside of the norm

- Remote memory access
- Fine-grained communication
- All selectable at runtime

MiniMod: Modular Architecture

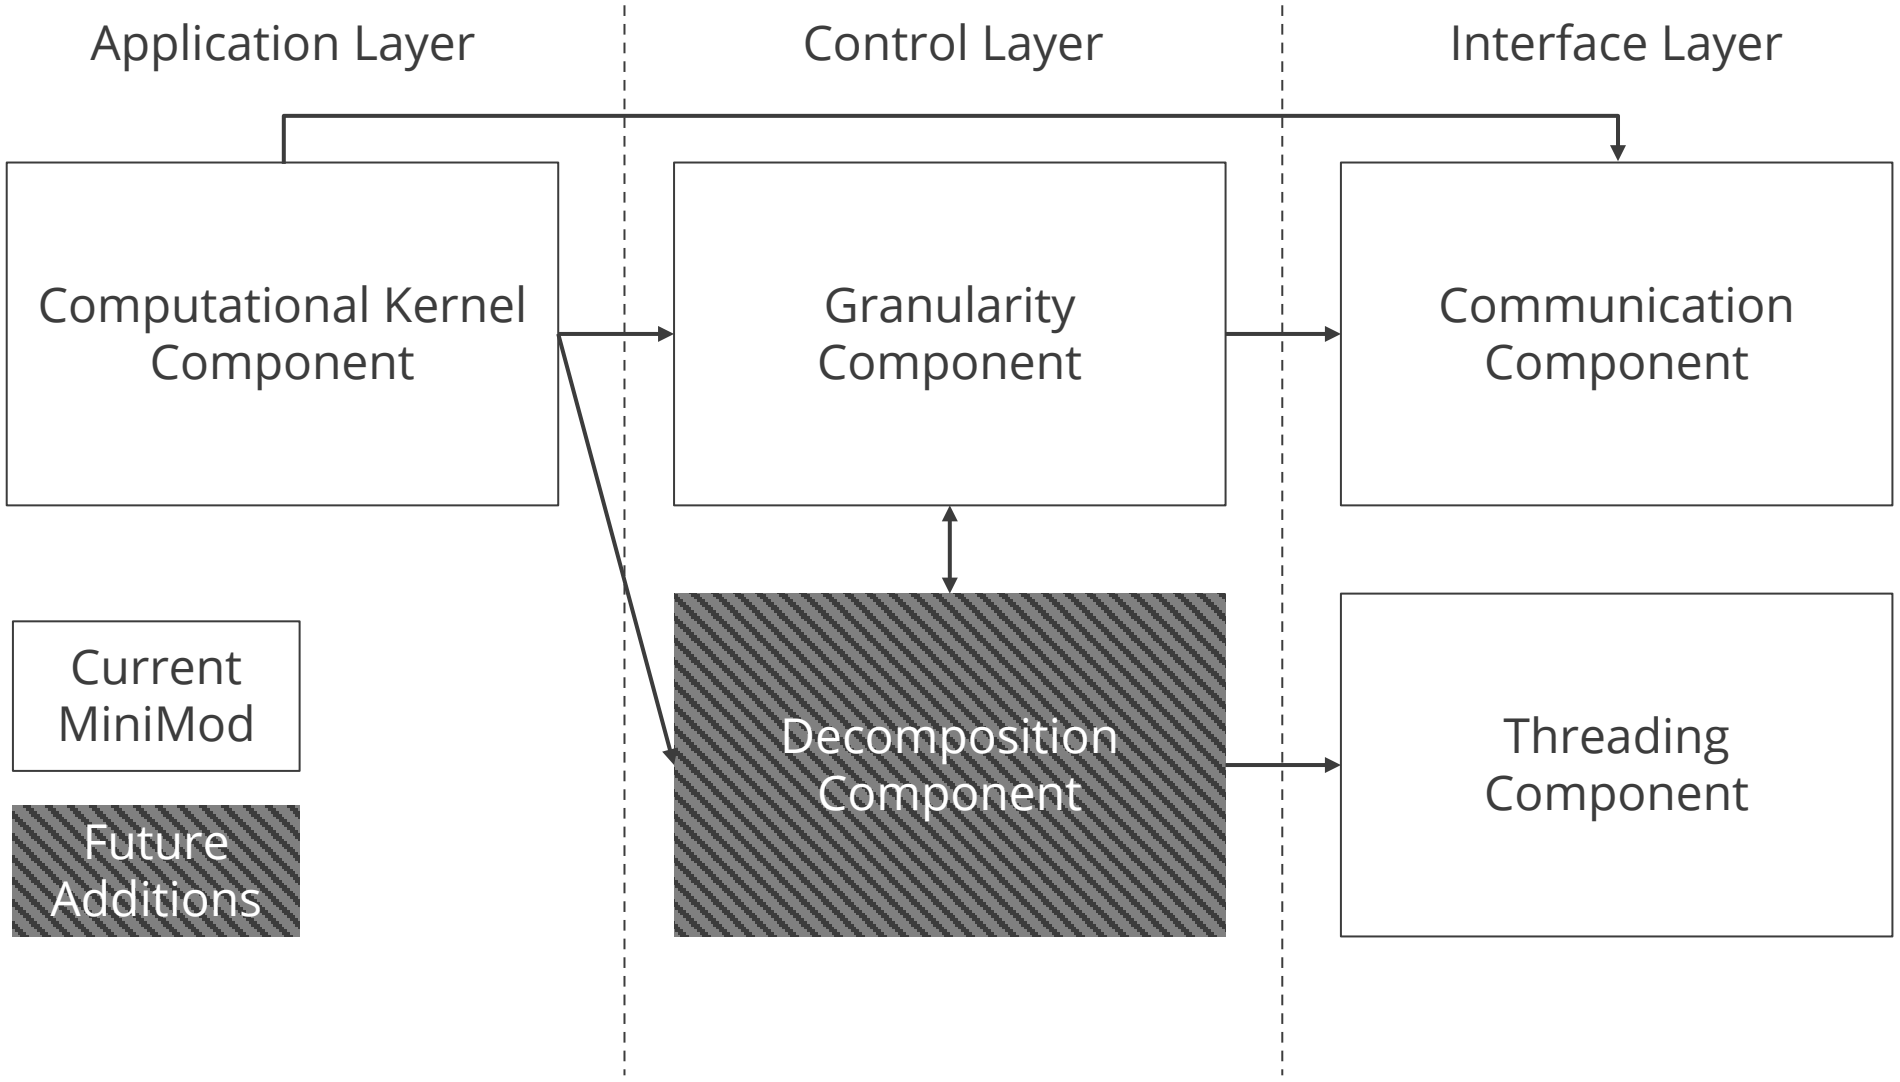

MiniMod: Calling Structure

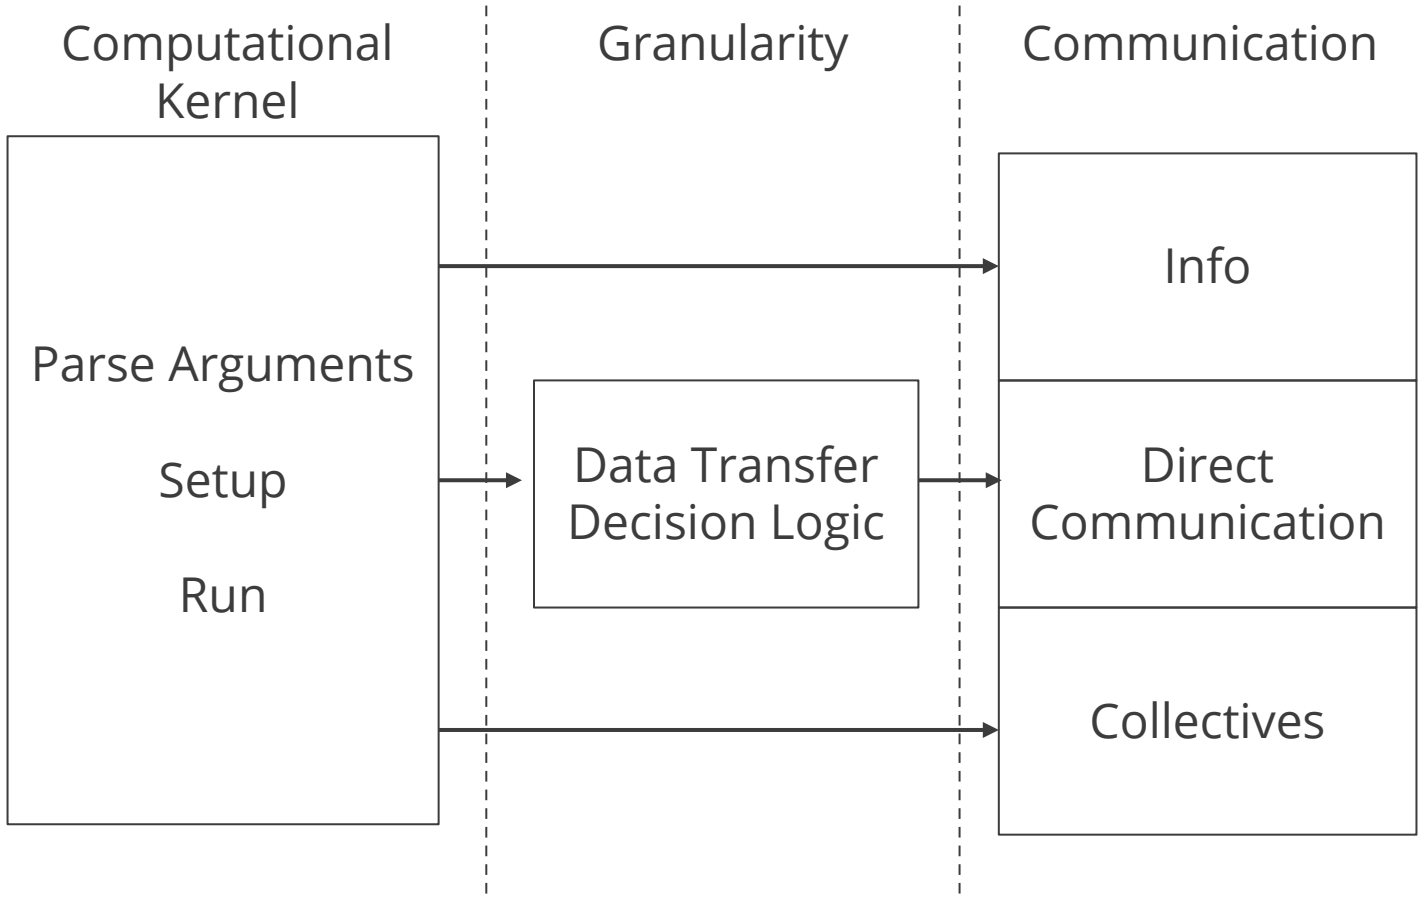

MiniMod: Module Selection

Layer	Component	Modules
Application	Computational Kernel	Micro-benchmarks Heat Diffusion MiniFE
Control	Granularity	Bulk synchronous Fine-grained
Interface	Communication	MPI two-sided MPI RMA Open SHMEM
	Threading	Pthreads OpenMP

MiniMod: Code Example

MiniMod Code Hooks:

- Initialize the pattern
- Mark the data ready
- Define the iteration

```
mm_init(mm_req);
for (int i = 0; i < iters; i++){
    mm_start(mm_req);
    for (int x = 0; x < elems; x++){
        // Compute
        data[x] = update(x)
        // Pack
        send_buffer[x] = data[x];
        mm_ready(mm_req, x);
    }
    mm_end(mm_req);
}
```

MiniMod: Validation and Experimentation

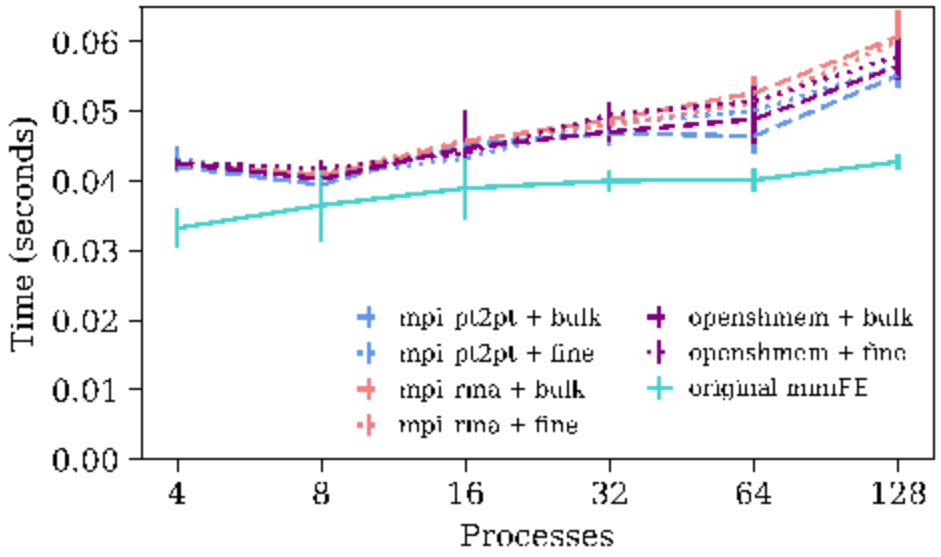

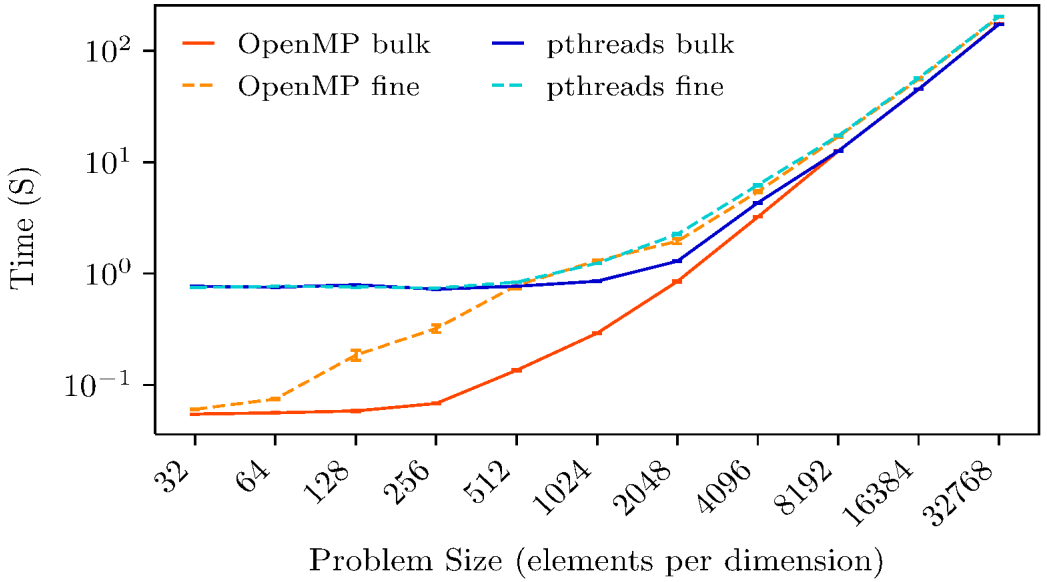

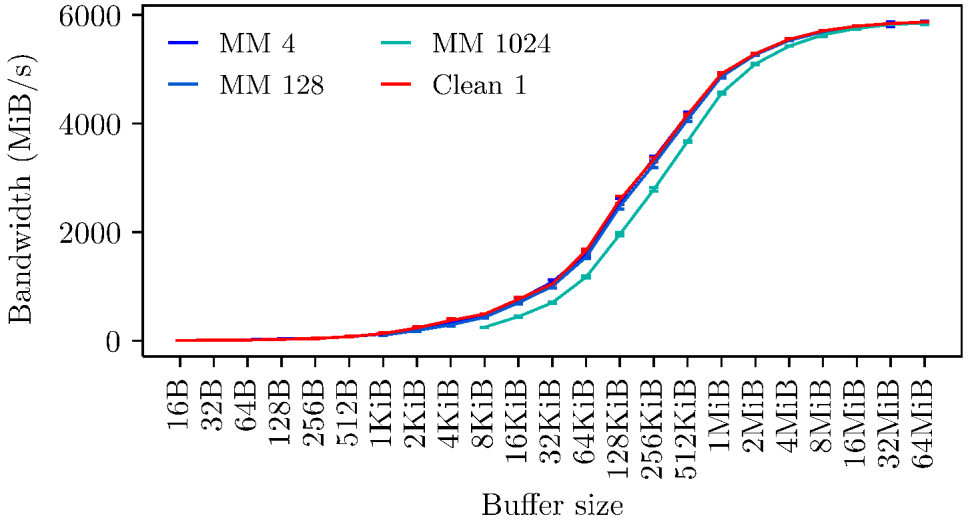

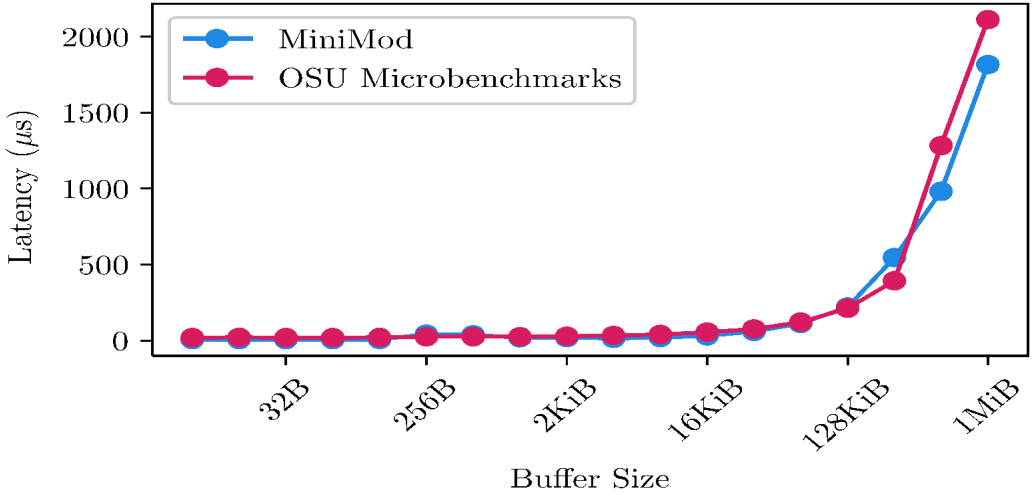

MiniMod: Future Work

When and how to leverage new communication methods

- Fine grained communication
- MPI partitioned communication
- Application traits and thresholds for benefit

Explore novel threading and control methodologies

- Software offloading
- Task based threading models

MiniMod: Questions?

Email: wmarts@sandia.gov