SAND2015-10655 C

An Asynchronous Many-Task Implementation of In-Situ Statistical Analysis using Legion

Philippe PÉBAY & Janine BENNETT
Sandia National Laboratories

(with special thanks to the Legion team @ Stanford)

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under Contract DE-AC04-94AL85000.

Motivation

- Design, develop, and evaluate unified data-driven approach for programming applications and *in situ* analysis and visualization;
- Study interplay between data-centric programming model requirements at extreme-scale;
- Assess impact of those requirements on design, capabilities, and implementation details for both applications and supporting *in situ* infrastructure.

Scope of this work

- Describe our first effort, using Legion, to:
  - support data-centric *in situ* data analysis,
  - provide foundation for efficient workflow handling.
- Report on:
  - proof-of-concept implementation,
  - current challenges with Legion.

MiniAero

- Compressible Navier-Stokes, 3 D, unstructured mesh, finite volume, explicit CFD mini-application;
- Implements 1 D shock tube problem set up in 3D domain discretized in rectilinear finite volume cells;
- MiniAero/Legion implementation used as basis for this study: surrogate for "real" scientific application.

Titan Statistics

- Set of C++ classes, part of VTK
- Parallel implementation using MPI.
- Learn/Derive/Assess model

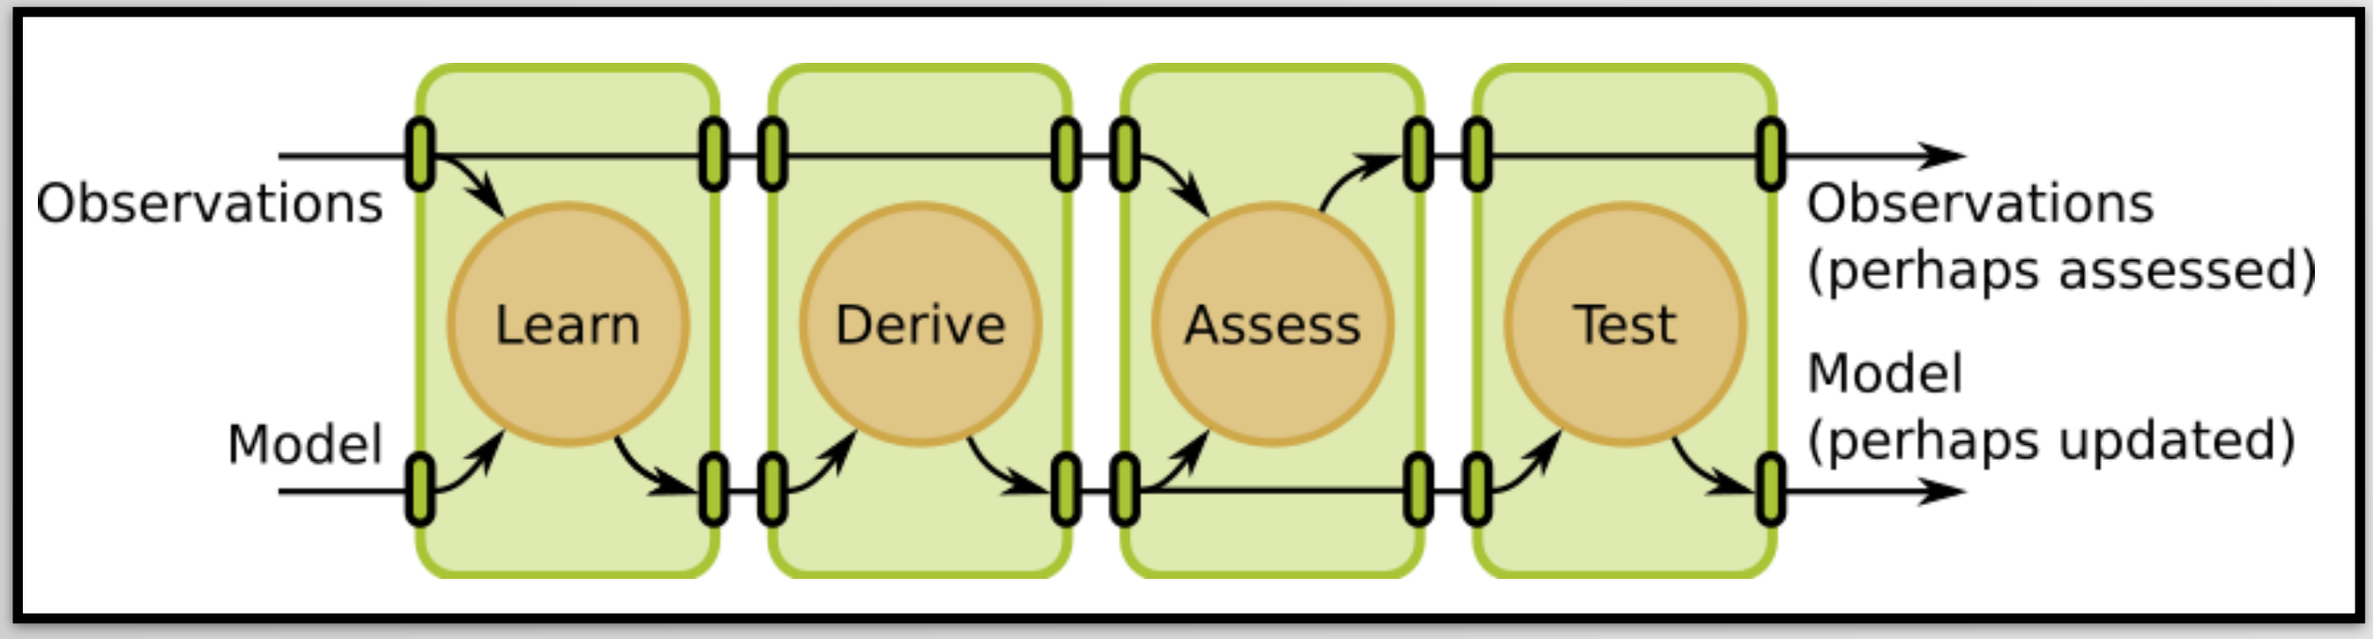

SPMD Implementation

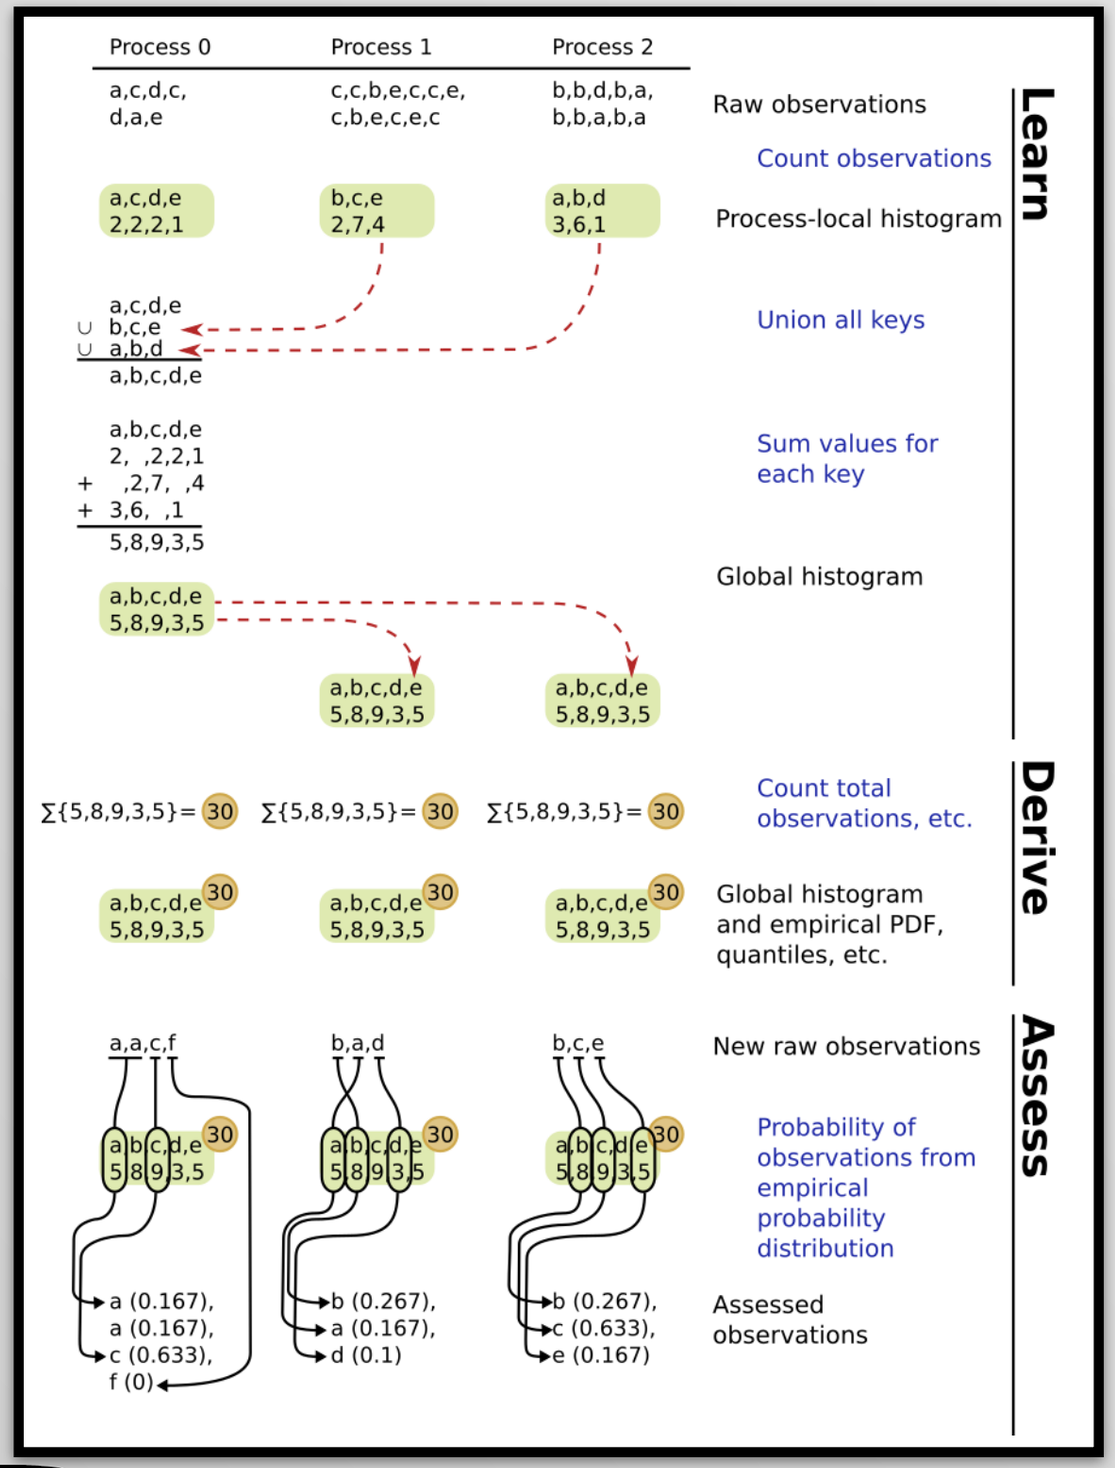

Scales optimally with up to $O(10^5)$ processes for:

- moment-based statistics,
- quanta-based statistics when quantization is adapted.

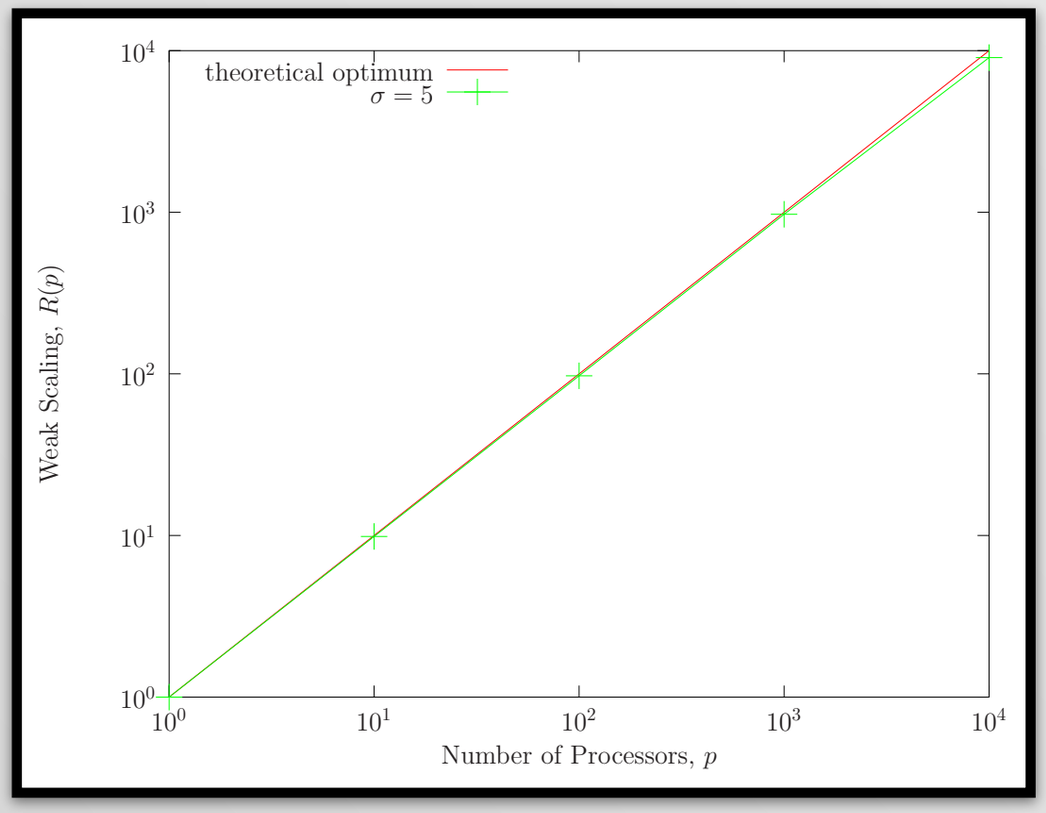

P. PÉBAY & J. BENNETT

AMT Approach

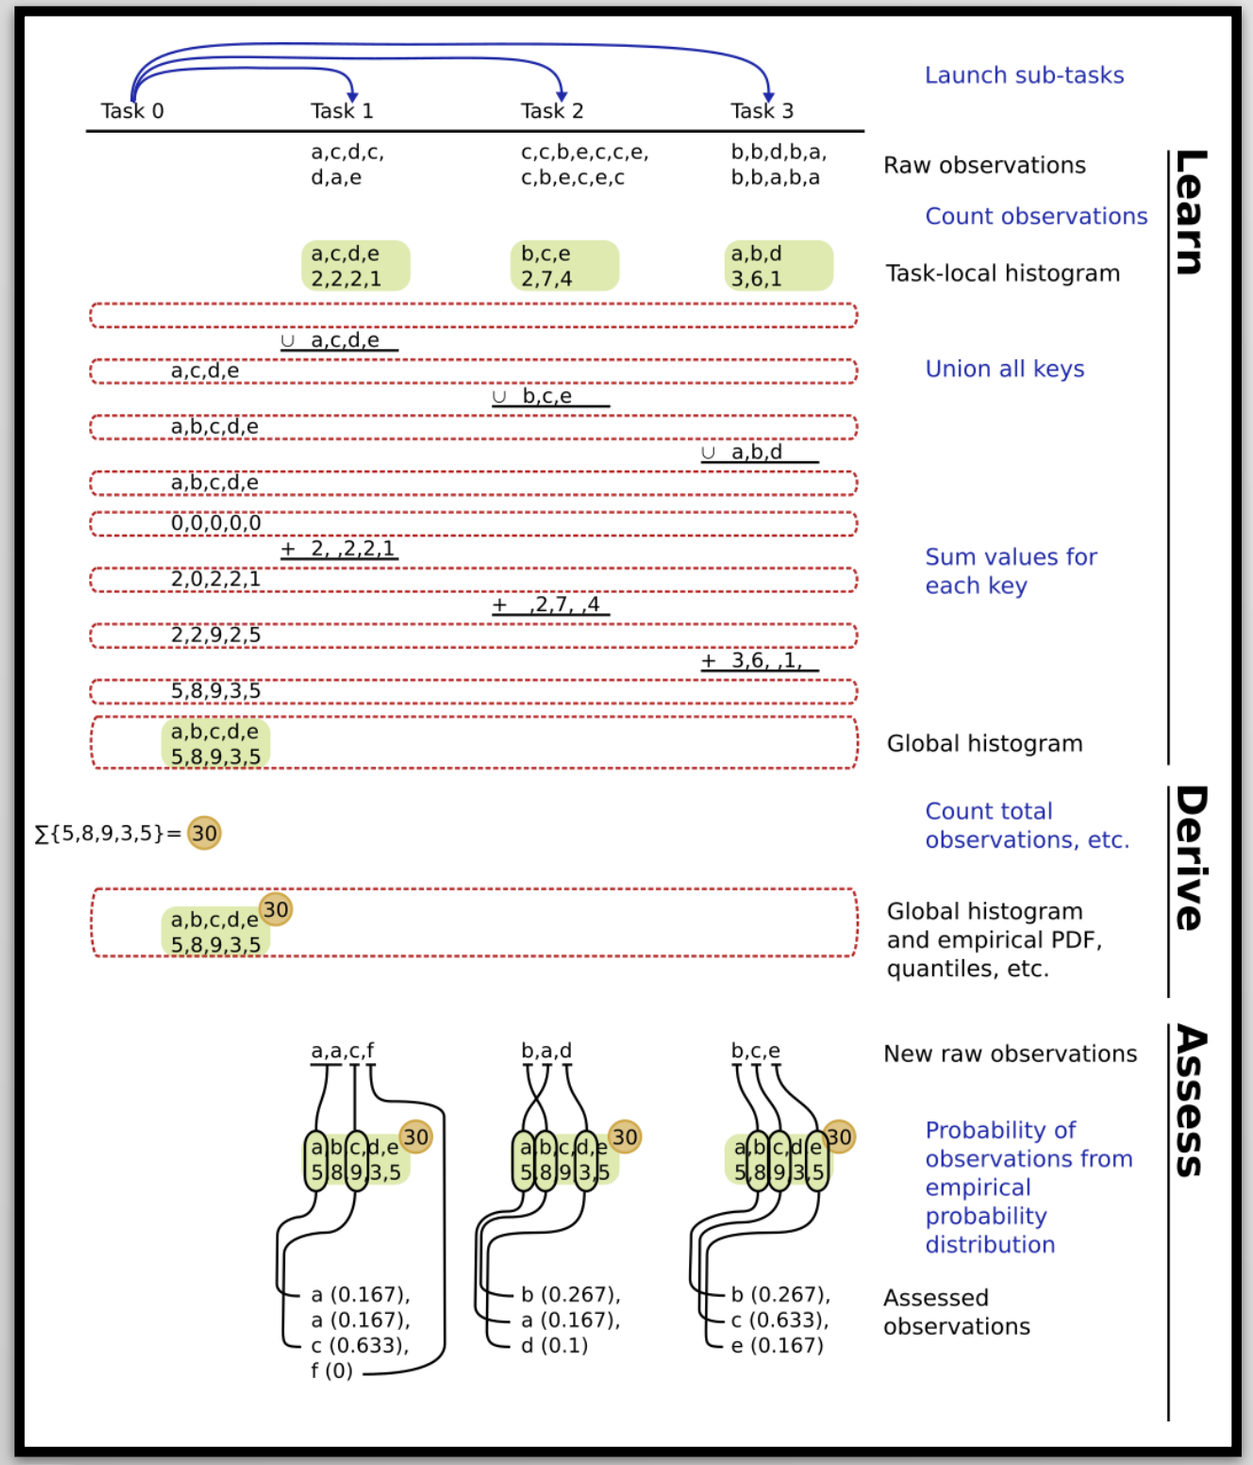

- Retain L/D/A work flow;
- Parallel Assess processes become tasks
- Replace MPI collectives with *aggregation region;*
- Implement with Legion regions (logical/physical);
- Derive and Assess are single tasks launched from top-level driver.

Implementation details

- Changes to MiniAero/Legion limited to a few files:
  - `TaskIDs.h`
  - `main.cc`
  - `Interface.h/cc`
  - `toplevel.cc`
  - `Mapper.cc`
- Most changes (aggregation regions) in `toplevel.cc`
- Total ~ 80 lines of code

Algorithm correctness

- "3D SOD" test case with 128x4x4 grid
- 2-core Intel Core i7
- Shared-memory LLRTS

Statistic	Reference	1 task	4 tasks	64 tasks	Octave
Cardinality	2048	2048	2048	2048	2048
Minimum	0.1039304	0.1039304	0.1039304	0.1039304	0.1039304
Maximum	0.8314390	0.8314390	0.8314390	0.8314390	0.8314390
Mean	0.4676847	0.4676847	0.4676847	0.4676847	0.4676847
Std. dev.	0.3638431	0.3638431	0.3638431	0.3638431	0.3638431
Skewness	0	-3.2766E-16	0	3.6015E-17	-1.9708E-13
Kurtosis	-2.0009763	-2.0009763	-2.0009763	-2.0009763	-2.0000000

Algorithm Scalability

- "3D SOD" test case with 128x4x4 grid
- 2-core Intel Core i7
- Shared-memory LLRTS

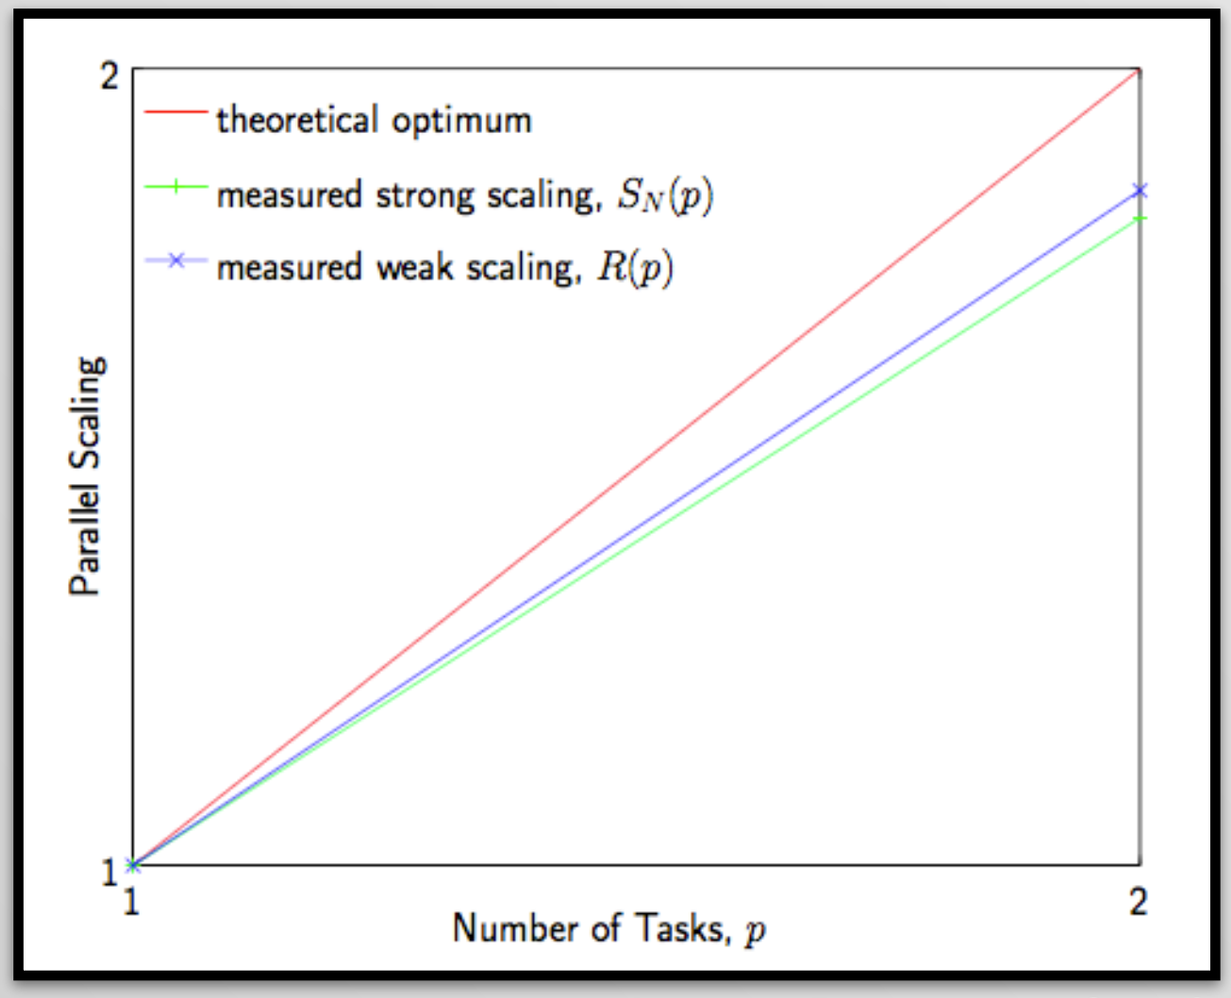

Main Findings

- Proof-of-concept approach for AMT *in situ* analysis appears valid; additional tests under way with several test clusters.
- Algorithm correctness in theory — but in practice there are some issues to be resolved.
- Sub-optimal scaling, confirmed with larger studies. This requires further investigation.
- Legion API improvements to be discussed.