# Large-Scale Solution of Constrained Systems via Monolithic Multigrid

Luke Olson
Department of Computer Science
University of Illinois at Urbana-Champaign

2017.03.02

## 1 Project Information

**Project title** Large-Scale Solution of Constrained Systems via Monolithic Multigrid

**Name and institution of principal investigator** Luke Olson, Department of Computer Science at the University of Illinois at Urbana-Champaign, *email:* `lukeo@illinois.edu` *web:* `http://lukeo.cs.illinois.edu`

**Names and affiliations of co-PIs and collaborators** Amanda Bienz, Department of Computer Science, University of Illinois at Urbana-Champaign; Andrew Reisner, Department of Computer Science, University of Illinois at Urbana-Champaign; Lukas Spies, Coordinated Science Laboratory, University of Illinois at Urbana-Champaign;

**Corresponding author name and contact information (email)** Luke Olson, `lukeo@illinois.edu`

## 2 Executive summary (150 words)

Physical problems with constraints arise in many fields of computational science and engineering, including fluid and solid mechanics, and the modeling of liquid crystals and ferromagnetic materials. A key component in the simulation of these problems is the solution of a linear system. Although algorithmic development for these problems has advanced substantially in recent years, the scalable performance of linear solvers has not been fully realized. The focus of this project is to bridge the gap between scalability and recent algorithmic advancements. A focus of this project was to analyze implementations at scale and seek algorithms that fully utilize the potential of the machine.

## 3 Description of research activities and results

### 3.1 Key Challenges: description of the science/engineering problem being addressed

In previous work we have focus on sparse matrix problems of the form

$$Ax = b, \tag{1}$$

where the data patterns are highly irregular. Here, we focus the work on *structured* computations, where $A$ may be described through logical stencils yielding predictable memory access and communication patterns. In Figure 1, a key challenge is observed: the lack of linear scaling when problems sizes (and core counts) are increased; here we would expect constant time in a perfect solver (ie, a flat line).

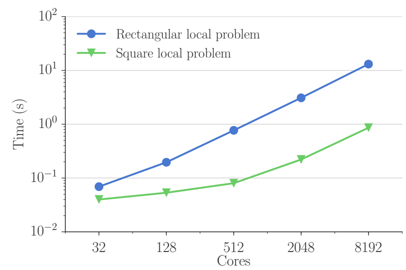


Figure 1: Total solution time of a structured, elliptic problem using robust, variational multigrid.

### 3.2 Why it Matters: description of the potential impact of solving this research problem, and, if appropriate, the educational outcomes. For exploratory allocations, indicate whether the results will be used to substantiate a future proposal.

Structured matrix computations are found in a range of simulations and represent a key operation and potential bottleneck for very large problems. Efficient, scalable methods are critical for this method, particularly as machines become more heterogeneous, uprooting traditional approaches and exposing new challenges.

The results due to this allocation show high potential and will likely contribute to a variety of computational problems. The allocation on Blue Waters has contributed to the basics of halo exchanges on GPU architectures, the design of optimal structured solvers at scale, and potential node-level abstractions that can be built into the software's parallel communication package to reduce overall communication loads.

### 3.3 Why Blue Waters: explanation of why you need the unique scale and attributes of Blue Waters, both the system and the staff, to address these challenges; if relevant, provide an assessment of code(s) performance on Blue Waters.

One challenge in structure matrix computations is the utilization of both CPUs and accelerators, such as GPUs. Blue Waters was instrumental in the algorithm development as experimental tests could be executed at large scales, exposing limitations of the method and the performance models.

### 3.4 Accomplishments: explanation of results you obtained

There were two main accomplishments of this work:

**Scalable variational multigrid solver** An optimal approach to coarsening was developed to address the challenges in Figure 1. Here, a performance model was used to guide the algorithm in finding the optimal redistribution of data (see Figure 2).

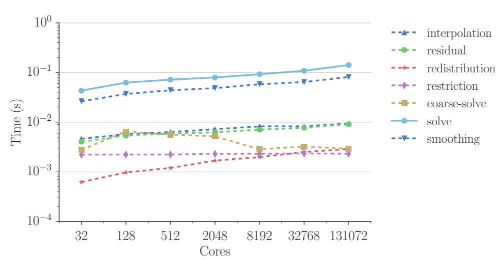


Figure 2: Scaling on Blue Waters.

**Heterogeneous Halo Exchange** The key kernel in a structured matrix computation is that of a halo exchange (Figure 3). With an efficient exchange layer, structured computations can be highly scalable, utilizing the available compute resources (both CPUs and GPUs). See Figure 4 for a comparison with a CPU-only, highly tuned package, MSG.

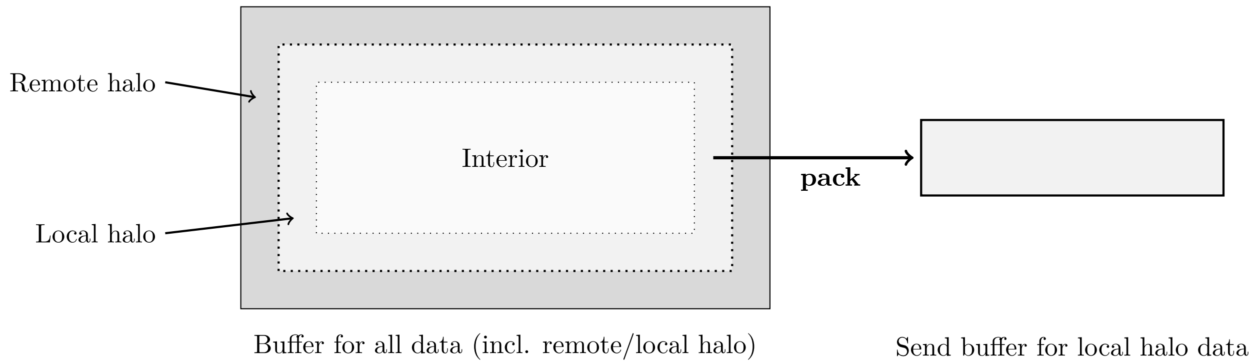


Figure 3: Packing local halo data into dedicated send buffer

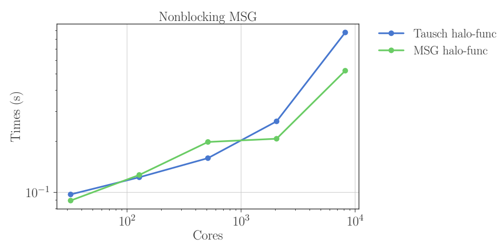


Figure 4: Performance of a halo exchange.

### 3.5 Next Generation Work: If relevant, state what you hope to accomplish on a next-generation Track-1 system in the 2019–2020 timeframe.

One extension of this work is to fully integrate the advances with emerging solver packages.

## 4 List of publications, data sets associated with this work

- *Scaling Structured Multigrid to 500K+ Cores through Coarse-Grid Redistribution*, by David Moulton, Luke N. Olson, and Andrew Reisner, in review.