TN

CI

SCALING ELLIPTIC SOLVERS VIA DATA REDISTRIBUTION

Allocation: Exploratory/60 Knh
PI: Luke Olson[1]
Collaborators: Amanda Bienz[1], Andrew Reisner[1], Lukas Spies[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Elliptic partial differential equations play a central role in many scientific simulations. Yet, solvers for such methods are limited in efficiency, particularly at high core counts. Structured multilevel solvers are especially effective for this class of problems; however, parallelism is limited with traditional methods. In this work, data are redistributed in order to extend the scaling limits of the method.

RESEARCH CHALLENGE

Structured multilevel methods are a highly effective and common tool in solving elliptic partial differential equations across a range of applications areas. Yet, scalability of these methods is limited, particularly at high core counts. The focus of this work is on the development of a multilevel method that redistributes data on coarse levels in the solver in order to optimally extend the parallel scalability of the solver.

METHODS & CODES

Structured multilevel solvers use a series of successively coarser grids to approximate the error in the original, fine-level problem. As grids coarsen, the amount of local work is significantly reduced, leaving high levels of communication per degree of freedom in the problem. Moreover, the coarsest levels often result in only a few (or one) degree of freedom per core, necessitating a gathering of the problem to one or a subset of processors.

In this work, Blue Waters is used to test a predictive performance model [1] for scaling the solver to 100,000 cores, while the baseline solver is limited to 4,000 cores. Central to the approach is the agglomeration of data as presented in Fig. 1. Here, a method is used to:

- Agglomerate processors into groups of coarse tasks;
- Gather or collect data for processors in each block; and
- Continue cycling redundantly with redistributed data.

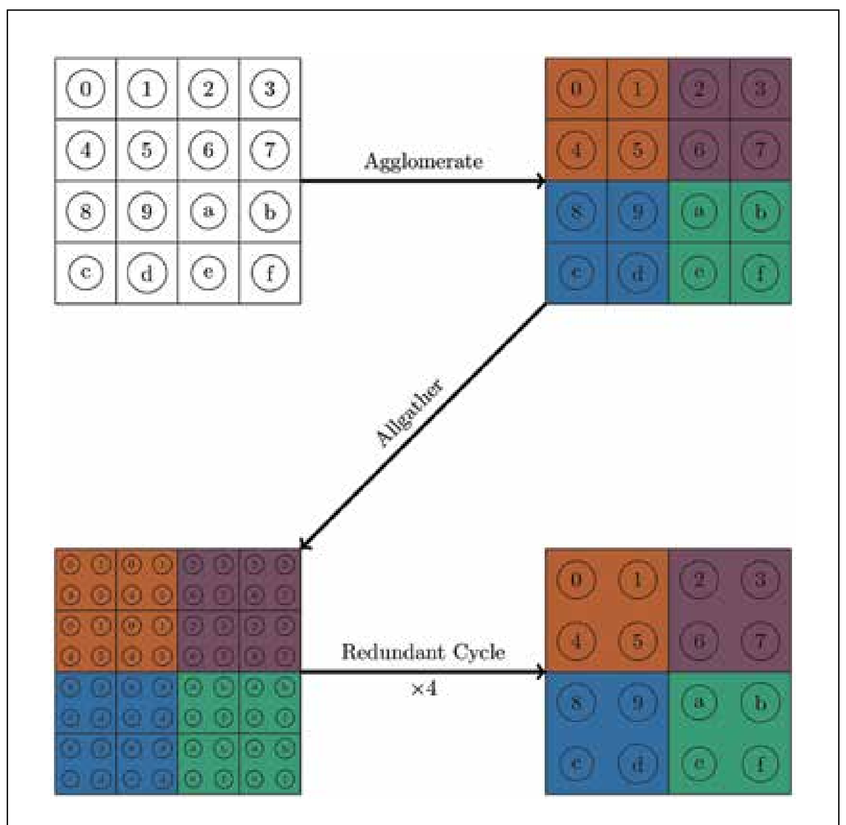

Figure 1: A 4 × 4 processor grid is redundantly reduced to a 2 × 2 processor grid.

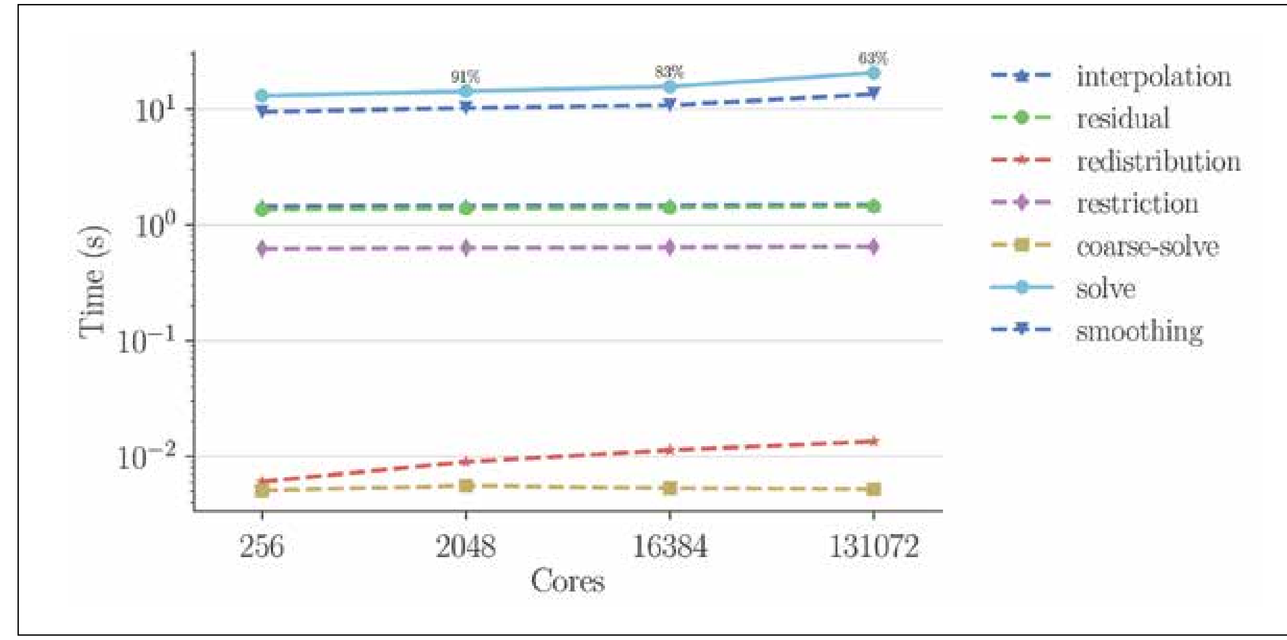

Figure 2: Weak scaling of a 3D elliptic problem on Blue Waters with a per-core size of 5M+ and a total problem size of over 660 billion points at the largest scale.

There are a large number of possibilities for redistributing data. In this work a search algorithm is used to identify the optimal redistribution based on a parallel performance model. As a result, the solver is capable of scaling to very large core counts with little overhead, as shown in Fig. 2. The code used and tested for this work is the Cedar Framework (https://github.com/cedar-framework/cedar).

RESULTS & IMPACT

There are several outcomes of this work. The first is the testing of a new scalable solver outlined in [1]. The method uses predictive performance models to guide the redistribution of data for optimal network traffic, leading to a robust and efficient method for a range of structured elliptic problems. Another key element of this work was the development of an efficient halo exchange suitable for heterogeneous systems. Finally, the modeling aspects of this work led to new insights on parallel performance. Together, these steps have contributed to fundamental advances in robust structured solvers.

WHY BLUE WATERS

Blue Waters was instrumental in developing accurate performance models and in testing the scalability of the methods. The allocation helped to extend the method and the code to new scales.

PUBLICATIONS & DATA SETS

Reisner, A., L. Olson, and D. Moulton, Scaling Structured Multigrid to 500K+ Cores through Coarse-Grid Redistribution. *SIAM Journal on Scientific Computing*, to appear (2018).