TN

CI

SCALABLE LINE AND PLANE SOLVERS

Allocation: Illinois/75 Knh
PI: Luke Olson[1]
Co-PI: Amanda Bienz[1]
Collaborator: Andrew Reisner[1]

[1]University of Illinois at Urbana–Champaign

EXECUTIVE SUMMARY

Structured elliptic solvers play an important role in a range of applications, from plasmas to solid mechanics. However, their scalability is limited owing to the global nature of the problem. Multigrid methods have been effective at solving this class of problems; however, structured multigrid solvers rely on sweeps over lines and planes in the mesh in order to yield an effective method, which severely limits scalability. Through this work, the research team has developed scalable kernels for line and plane relaxation methods.

RESEARCH CHALLENGE

Sparse matrix problems arising from elliptic partial differential equations present a significant computational challenge in many applications at scale. Iterative approaches such as multigrid preconditioned conjugate gradient methods have proven to be cost effective. Yet, in situations with high mesh anisotropy or variable problem coefficients, multigrid solvers require the use of more robust components. Point relaxation techniques such as the Jacobi method are insufficient, thus requiring the use of line and plane forms of these methods.

METHODS & CODES

Line relaxation (and in 3D, plane relaxation) plays an important role in preconditioning structured elliptic problems. Fig. 1 highlights the importance of line relaxation in comparison to a standard weighted Jacobi (pointwise) relaxation solver in multigrid for the case of an annulus with moderate 10:1 stretching of the grid. There is a significant reduction in total iterations, yet each pass of line relaxation requires a distributed tridiagonal solver.

The key bottleneck in a tridiagonal (or banded solver) is the limited work per processor. Fig. 2 highlights the *O(p)* effect, with *p* cores increasing in one dimension. A straightforward implementation of the tridiagonal solver results in a linear relationship. This work focuses on a multilevel form of the problem where a new tridiagonal problem is formed in the communication overlap, resulting in a *O(log p)* dependence and significantly reducing cost.

The current approach centers on two interrelated aspects to further advance the solver. First is the development of a scalable and accelerator-aware halo exchange library, called *Tausch* [1], which can be used in a variety of structured codes. Second is the extension of the structured solver *Cedar* [2] to accelerators, such as the XK portion of Blue Waters.

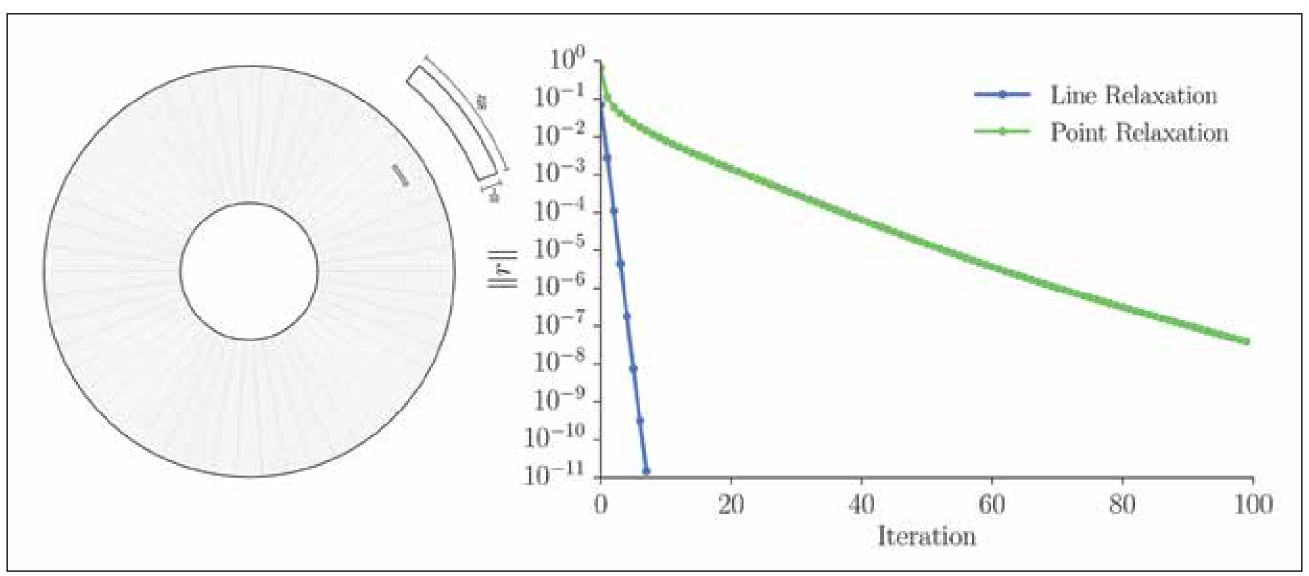

Figure 1: Point and line relaxation convergence for a Poisson problem on an annulus.

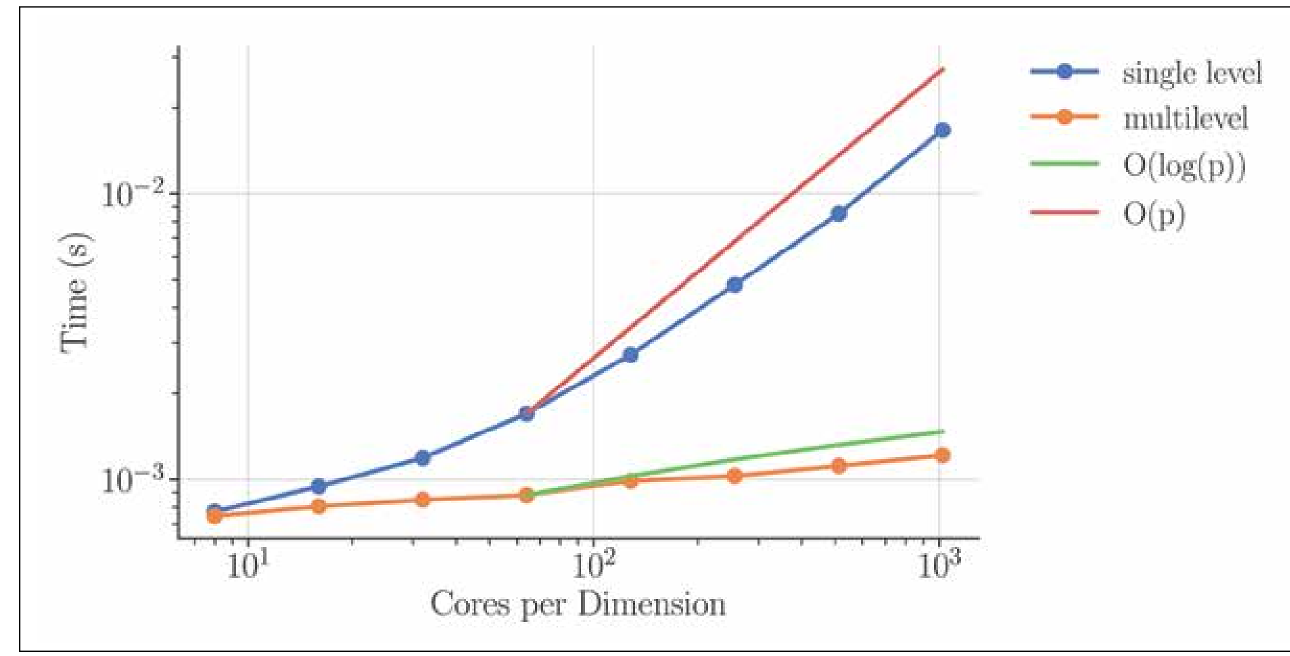

Figure 2: Scalability of sequential line relaxation versus multilevel line relaxation as the core count (*p*) increases.

RESULTS & IMPACT

This work has directly contributed to the algorithms, testing, and development of the *Tausch* [1], *Cedar* [2], and *RAPtor* [3] software packages. Each of these targets general use in a wide range of scientific applications.

WHY BLUE WATERS

Blue Waters has played a key role in testing structured solvers at scale, for example in [4]. Access to large core counts and consistent testing has been instrumental in developing scalable algorithms and accurate performance models that form the core of this work.