NUMERICAL LINEAR ALGEBRA WITH APPLICATIONS
Numer. Linear Algebra Appl. 2014; **21**:275–296
Published online in Wiley Online Library (wileyonlinelibrary.com). DOI: 10.1002/nla.1928

Reducing communication in algebraic multigrid using additive variants

Panayot S. Vassilevski and Ulrike Meier Yang*,†

Lawrence Livermore National Laboratory, 7000 East Ave, L-560, Livermore, CA, 94550

SUMMARY

Algebraic multigrid (AMG) has proven to be an effective scalable solver on many high performance computers; however, its increasing communication complexity on coarser levels has shown to seriously impact its performance on computers with high communication cost. Additive AMG variants provide not only increased parallelism as well as decreased numbers of messages per cycle but also generally exhibit slower convergence. We present various new additive variants with convergence rates that are significantly improved compared to the classical additive algebraic multigrid method and investigate their potential for decreased communication, and improved communication-computation overlap, features that are essential for good performance on future exascale architectures. Published 2014. This article is a US Government work and is in the public domain in the USA.

Received 28 May 2013; Revised 14 November 2013; Accepted 19 December 2013

KEY WORDS: multiplicative multigrid; additive implementation; parallelism; reduced communication

1. INTRODUCTION

Algebraic multigrid (AMG) is a popular solver for large-scale scientific computing and an essential component of many simulation codes. AMG has shown to be extremely efficient on distributed memory architectures [1, 2]. However, with single-core speeds plateauing, future increases in computing performance have to rely on increased concurrency provided by the architecture, leading to potentially billions of cores or threads. Applications have to match this increased level of concurrency to exploit the performance potential and hence face additional communication requirements. Future systems will also be subjected to strict power limitations for overall system power, and data movement, which includes communication, is responsible for a majority of the power consumed in a system. Therefore, to address these challenges, increased parallelism and reduced availability of power, and to successfully exploit future architectures, it is crucial to develop algorithms with reduced communication. Such an algorithm will consume less power, and its performance will be less affected by a reduction in power.

Algebraic multigrid obtains its optimal computation complexity by using smaller coarse grid problems to approximate the solution of the original fine grid problem. For traditional, matrix-based AMG with Galerkin or variational coarsening, which we consider here, the number of nonzeros per row for the coarse grid operators grows, and with it the number of neighbor processes. The communication complexity increases significantly, leading to a large number of messages. Contention at these levels can lead to a significant decrease in performance and scalability on current architectures with slower networks [3, 4] and is expected to be a bottleneck for future exascale machines. To counter the high communication complexities at the coarse levels, new variants with reduced

*Correspondence to: Ulrike Meier Yang, Lawrence Livermore National Laboratory, 7000 East Ave, L-561, Livermore, CA, 94550.

†E-mail: umyang@llnl.gov