

Multigrid Methods — An Overview


Luke Olson
University of Illinois at Urbana-Champaign
Nelder Fellow – Imperial College

W May 06 2020, 4–6pm BST

F May 08 2020, 4–6pm BST

M May 11 2020, 4–6pm BST

F May 15 2020, 4–6pm BST



Goals

The focus of this lecture series is on the fundamentals of multigrid methods and is open to both practitioners and methods developers in mathematics, computer science, and the applied sciences. The lecture series will cover the basics of multigrid methods in both a geometric and algebraic setting, introduce some key concepts in the theoretical treatment of these methods, and highlight their use in a parallel setting.

Suggested prerequisites: strong background in linear algebraic and some exposure to linear solvers; basic Python experience (for examples, but not required); general numerical analysis; introductory knowledge of partial differential equations.

Motivation

The solution of sparse linear systems is a significant computational challenge in many simulations. Often these arise from the discretization of partial differential equations, although the emergence of complex analytics in data science application has renewed interest in robust solvers. The focus of this lecture series is on so-called multigrid solvers, both theory and practice.

Multigrid methods have a rich history and continue to play a major role as the primary solver for many applications. From its inception in the 70s and through much of the 80s, the focus of much of the multigrid development was on *geometric* multigrid, a form that relies on inherent structure in the problem. The basic framework of an *algebraic* approach took form in the late 80s and through the 90s, culminating in range of new methods in the early 2000s for challenging (mostly elliptic) problems. In recent years we have observed significant advances in both theory and in application to a vast number of problems, including non-symmetric problems, systems of partial differential equations, graph problems, etc. To this end, one objective of this lecture series is to provide a survey of the multigrid methodology as well as modern ‘best practices’.

Outline

A goal of the proposed lecture series is to build a working knowledge of multigrid methods, both geometric and algebraic. In addition, an objective is to bring practitioners, both applications researchers and theoreticians, up to speed with the state of the art in the field. As such, we will cover the basics of projection methods and geometric-based multigrid, the basics of the algebraic form of multigrid, best practices in developing principled solvers for a range of problems, theoretical aspects of algebraic multigrid, and the implications of these approaches in a parallel setting.

Lecture #1: The Basics of Multigrid This lecture will cover the basic principles of geometric multigrid, from relaxation methods with a *smoothing* property to the introduction of a two-level method. Convergence properties will be outlined and extensions to multiple dimensions will be covered.

Lecture #2: Toward Algebraic Multigrid This lecture will begin with an overview of limitations of basic geometric multigrid. Operator based methods, such as BoxMG will be introduced, and application to nonlinear problems will be explored. The concept of *algebraically smooth* will also be used to motivate upcoming lectures.

Lecture #3: CF and Smoothed Aggregation Algebraic Multigrid An overview of the two basic forms of algebraic multigrid will be given. Classical, splitting-based methods along with aggregation-based approaches will be highlighted in detail, with several examples given in context.

Lecture #4: Theoretical Aspects of Algebraic Multigrid Multigrid convergence relies on both effective smoothing and robust coarse-grid correction. In this lecture, we give an overview of the underlying theory that supports algebraic multigrid and highlight several areas in which the theory is expanding.