

Multiscale computation: from fast solvers to systematic upscaling

A. Brandt*

The Weizmann Institute of Science, Rehovot 76100, Israel

Abstract

Most fundamental problems in physics, chemistry and engineering involve computation too hard even for future supercomputers, if conventional mathematical approaches are used. The reason is always a product of several complexity factors associated with the wide range of space and time scales characteristic to such problems. Each of these complexity factors can in principle be removed by various multiscale algorithms, i.e., employing separate processing at each scale of the problem, combined with interscale iterative interactions. Starting from multigrid fast solvers for discretized partial differential equations and from renormalization group methods in theoretical physics, the multiscale computational methodology has recently been extended to many other areas and new types of problems: linear and highly nonlinear, deterministic and stochastic, discrete and continuous, with particles and macromolecules, graphs and images.

Keywords: Multiscale algorithms; Multigrid methods; Renormalization group methods; Upscaling

Despite their dizzying speed, modern supercomputers are still incapable of handling many of the most vital scientific problems. This is primarily due to *the scale gap*, which exists between the microscopic scale at which physical laws are given and the much larger scale of phenomena we wish to understand.

This gap implies, first of all, a huge number of *variables* (e.g., atoms or gridpoints or pixels), and possibly even a much larger number of *interactions* (e.g., one force between every pair of atoms). Moreover, computers simulate physical systems by moving *one variable at a time*; as a result, each such move must be extremely small, since a larger move would have to take into account all the motions that should in parallel be performed by all other variables. Such a computer simulation is particularly incapable of moving the system across large-scale *energy barriers*, which can each be crossed only by a large, and unknown, coherent motion of very many variables.

This type of computational obstacles makes it impossible, for example, to calculate the properties of nature's building blocks (elementary particles, atomic nuclei, etc.) from a certain known underlying theory — and thereby to confirm the theory itself through comparison with measurements. Likewise, such obstacles are the main bottleneck in the drive to *computerize chemistry and materials science*: to replace expensive experiments with computer

simulations, yielding much more detailed understanding of molecular structures and interactions, creating the ability to design materials and processes, with enormous potential benefits for medicine, biotechnology, agriculture, material sciences, industrial processing, etc. Similar scale-born slowness factors and barriers, multiplying each other, plague many other engineering and scientific endeavors. All would be greatly facilitated if unlimited computing power were available — or if much better algorithms could be devised.

Just building ever faster machines will not do, in fact. With current computational methods the needed amount of computer processing often increases too steeply with the rise in problem size, so that in many important cases no conceivable computer will be adequate. Completely new mathematical approaches are needed.

Past studies have demonstrated that scale-born complexities can generally be effectively overcome, or drastically reduced, by *multiscale* (“multilevel”, “multigrid”, “multi-resolution”, etc.) algorithms.

Indeed, any many-variable problem defined in the physical space can have an approximate description at any given length scale of that space: a continuum problem can be discretized at any given resolution; average motions of a many-particle system can be represented at any given characteristic length; etc. The multiscale algorithm recursively constructs a *sequence* of such descriptions at increasingly larger (coarser) scales, and combines local processing (relaxation of equations, Monte Carlo simula-

* Tel.: +972 (8) 934-2345; Fax: +972 (8) 934-2945;
E-mail: achi@wisdom.weizmann.ac.il

tion of statistical relations, etc.) at each scale with various inter-scale interactions. Typically, the evolving solution (or the simulated equilibrium) on each scale recursively dictates the *equations* (or the Hamiltonian) on coarser scales while supplying large-scale corrections to the *solutions* (or configurations) on finer scales. In this way large-scale changes are effectively calculated on coarse grids, based on information previously gathered from finer grids.

As a result of such multilevel interactions, the fine scales of the problem can be employed very sparingly (just a couple of relaxation sweeps, for example), sometimes only at special small regions (local refinements just near singularities) or only in small representative “windows” (where larger-scale equations can be derived and then employed everywhere). Moreover, the inter-scale interactions can eliminate all kinds of scale-associated difficulties, such as: slow convergence (in minimization processes, PDE solvers, etc.); critical slowing down (in statistical physics); ill-posedness (e.g., of inverse problems); large-scale attraction basin traps (in global optimization and statistical simulations); conflicts between small-scale and large-scale representations (e.g., in wave problems); numerousness of long-range interactions (in many body problems or integral equations); numerousness of long-range (non-local) eigenfunctions (e.g., in quantum chemistry); the need to produce many fine-level solutions (e.g., in optimal control) or very many fine-level independent samples (in statistical physics); etc. Also, the evolving large-scale equations bring out the large-scale dynamics, or the macroscopic equations, of the physical system, which is often the very objective of the entire calculation.

Since the local processing (relaxation, etc.) in each scale can be done in parallel at all parts of the domain (e.g., at all cells of a given lattice), the multiscale algorithms, based on such processing, are ideal for implementation on massively parallel computers. Indeed, many problems cannot be solved efficiently by such computers without employing a multiscale procedure. For example, to fully parallelize a time-dependent calculation (i.e., to compute for earlier and later times simultaneously), a multiscale (multigrid) algorithm must be used.

Often, a combination of several multiscale approaches can benefit one particular problem in *many* different ways. Also, certain problems can benefit from multiscaling because this is the best way to *formulate* the problem, or some parts of it.

The main sources to the systematic multiscale approach are two major developments that took place in two different fields: multigrid (MG) methods in the field of applied mathematics, and renormalization group (RG) methods in the field of theoretical physics.

Multigrid methods were first developed as fast solvers for discretized elliptic partial differential equations (PDEs), based on two processes: (1) classical *relaxation schemes*, which are generally slow to converge but fast to smooth

the error function; (2) approximating the smooth error on a *coarser grid*, with a small coarsening ratio. The recursive combination of these two processes yields *linear complexity*, i.e., solution work proportional to the number of variables (unknowns) in the system. In a research stretched over many years, the range of applicability of these methods has steadily grown, to cover nearly all the major types of linear and nonlinear large systems of equations appearing in sciences and engineering. This has been accomplished by extending the concept of “smoothness” in various ways, finally to stand generally for any poorly locally determined solution component, and by correspondingly diversifying the types of coarse representations, to include for instance grid-free, multiple-coarsening and non-deterministic cases.

To obtain even further generality, there have however been two basic reasons to go beyond these multigrid methods. First, in all of them, a *correction* to a fixed current fine-grid configuration (a fine-level approximate solution or a fine-level statistical sample, depending on the type of problem) is represented as an interpolation from a coarser-level configuration. This is too restrictive for *highly nonlinear cases* (including cases of discrete-state and Lie-algebra variables), where configurations cannot be decomposed into weakly-interacting local and non-local parts (or non-smooth and smooth components).

Another basic reason to depart from correction-based multigrid methods is that for many systems, a linear complexity is not good enough, since the number of variables is huge. Such systems on the other hand are typically *highly repetitive*, in the sense that the same small set of equations (or Hamiltonian terms) keep repeating itself throughout the physical domain. This opens the way to the possibility of having, at the coarse level too, a small set of equations that are valid everywhere, and that can be derived from fine-level processing conducted only in some small representative “windows”, thus avoiding the need to operate with the huge number of fine-level variables. This is of course impossible to do with correction-based coarsening methods, whose coarse-level equations are not universal but depend on the local current fine-level approximation.

These two basic reasons point in fact in the same direction. Instead of relaxing the given system of equations so as to obtain a smooth error that can be approximated on a coarse level, one can use coarse level variables that are little sensitive to relaxation (representing chosen *averages* rather than individual values) and that represent the *full* solution rather than the correction to any given current approximation. Such coarse variables can be chosen (using a general criterion described in [1]) so that coarse-level equations can be derived just by local processing. We will use the term “*upscaling*” for this type of direct (full-solution) transition from a fine level to a coarser one. Its validity is very general, extending even to those highly nonlinear cases where all scales interact with each other so

strongly that correction-based multileveling is inapplicable.

In fact, upscaling, under the name “renormalization”, was first introduced into exactly those systems where all scales interact most strongly. These are the highly nonlinear systems of statistical mechanics at the critical temperature of phase transition. The *renormalization group* (RG) method was developed contemporaneously with, but independently of the multigrid method, its chief purpose having been to investigate the behavior of such critical systems at the limit of very large scales. It has thus focused on analyzing, theoretically and computationally, the fixed point of the *group* of successive renormalization steps, and various universal asymptotic power laws associated with it. Its computational efficiency remained however limited, mainly due to the lack of a systematic advance of the reverse transition, from coarser levels back to finer ones, which is needed either for accelerating the fine-level simulations and/or for confining them to small representative windows.

Realizing the complementary advantages of RG and MG, a new combined paradigm has emerged, first in the framework of critical statistical-mechanics systems, where it was named *Renormalization Multigrid* (RMG). We use the term *systematic upscaling* for the extensions of the RMG blueprint to other systems, statistical as well as deterministic.

Multilevel computation has evolved into a discipline by itself, having its own internal development, gradually increasing our understanding of the many types of multiscale interaction, their modes of operation and domains of application. Various underlying relations and algorithmic ideas are carried back and forth between widely varying types of problems.

Many recent developments are surveyed in [1]. The reported areas include: top-efficiency multigrid methods in fluid dynamics; inverse PDE problems and data assimilation; feedback optimal control; PDE solvers on unbounded domains and on adaptable grids; wave/ray methods for highly indefinite equations; rigorous quantitative analysis of multigrid; many-eigenfunction problems and ab-initio quantum chemistry; fast evaluation of integral transforms on adaptive grids; multigrid Dirac solvers; fast inverse-matrix and determinant calculations and updates; multiscale Monte-Carlo methods in statistical physics, including the renormalization multigrid (RMG) methods; molecular mechanics (including fast force summation, fast macromolecular energy minimization, and Monte-Carlo methods at equilibrium, both for macromolecules and for large ensembles of small molecules); combination of small-scale equilibrium with large-scale dynamics; image processing (edge detection and picture segmentation); tomography (medical imaging and radar reconstruction); efficient, general and highly accurate algebraic multigrid (AMG) and numerical homogenization schemes; fast practical graph algorithms; data clustering; and multiscale approaches to global optimization.

References

- [1] Brandt A. Multiscale scientific computation: review 2001. In: Barth TJ, Chan TF, Haimes R (eds.), *Multiscale and Multiresolution Methods*. Heidelberg: Springer Verlag, 2001; pp. 1–96. Appears also as “survey paper” at <http://www.wisdom.weizmann.ac.il/~achi>.