S0010-4655(05)00194-3/FLA AID:2927 Vol.●●●(●●●)
COMPHY:m3 v 1.36 Prn:29/03/2005; 8:48

cpc2927

[DTD5] P.1 (1-4)
by:Diana p. 1

Available online at www.sciencedirect.com
SCIENCE DIRECT®

Computer Physics Communications

ELSEVIER

Computer Physics Communications ●●● (●●●●) ●●●–●●●

www.elsevier.com/locate/cpc

Multiscale solvers and systematic upscaling in computational physics

A. Brandt

The Weizmann Institute of Science, Rehovot 76100, Israel

Abstract

Multiscale algorithms can overcome the scale-born bottlenecks that plague most computations in physics. These algorithms employ separate processing at each scale of the physical space, combined with interscale iterative interactions, in ways which use finer scales very sparingly. Having been developed first and well known as multigrid solvers for partial differential equations, highly efficient multiscale techniques have more recently been developed for many other types of computational tasks, including: inverse PDE problems; highly indefinite (e.g., standing wave) equations; Dirac equations in disordered gauge fields; fast computation and updating of large determinants (as needed in QCD); fast integral transforms; integral equations; astrophysics; molecular dynamics of macromolecules and fluids; many-atom electronic structures; global and discrete-state optimization; practical graph problems; image segmentation and recognition; tomography (medical imaging); fast Monte-Carlo sampling in statistical physics; and general, systematic methods of upscaling (accurate numerical derivation of large-scale equations from microscopic laws).

1. The scale gap

Despite their dizzying speed, modern supercomputers are still incapable of handling many 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 picture elements), and possibly even a much larger number of interactions. Moreover, computers simulate physical systems by moving few variables at a time; 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 coherent motion of very many variables.

This type of computational obstacles makes it impossible, for example, to calculate properties of elementary particles, atomic nuclei, etc., or to *computerize chemistry and materials science*, so as to enable the

E-mail address: achi@weizmann.ac.il (A. Brandt).

0010-4655/$ – see front matter © 2005 Published by Elsevier B.V.
doi:10.1016/j.cpc.2005.03.097

S0010-4655(05)00194-3/FLA AID:2927 Vol.●●●(●●●)
COMPHY:m3 v 1.36 Prn:29/03/2005; 8:48
cpc2927
[DTD5] P.2(1-4)
by:Diana p. 2

design of materials, drugs and processes, with enormous potential benefits for medicine, biotechnology, nanotechnology, agriculture, materials science, industrial processing, etc. With current common methods the amount of processing often increases so steeply with problem size, that even much faster computers will not do.

2. Multigrid and renormalization

Past studies have demonstrated that scale-born slowness can be overcome by multiscale algorithms. Such algorithms have first been developed in the form of fast *multigrid solvers* for discretized PDEs [1,2, 10–12]. These solvers are 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* (typically having twice the meshsize), by solving there equations which are derived from the PDE and from the fine-grid residuals; the solution of these coarse-grid equations is obtained by using recursively the same two processes. As a result, large scale changes are effectively calculated on correspondingly coarse grids, based on information gathered from finer grids. Such multigrid solvers yield *linear complexity* (solution work proportional to the number of unknowns).

In many years of research, the range of applicability of these methods has steadily grown, to cover most 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 solvers, called *algebraic multigrid* (AMG; see [5,6,13]), non-deterministic statistical mechanics problems [7–9,16] and multiple coarse-level representations [3].

It has been shown that the inter-scale interactions can indeed 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); conflicts between small-scale and large-scale representations (e.g., in wave problems, bridging the gap between wave equations and geometrical optics); numerousness of long-range interactions (in many body problems or integral equations); the need to produce many fine-level solutions (e.g., in optimal control) or very many fine-level independent samples (in statistical physics); etc. New AMG methods (BAMG—see [18]) can fast solve highly disordered systems, such as Dirac equations on critical gauge fields or PDE systems discretized on unstructured grids. Since the local processing (relaxation, etc.) in each scale can be done in parallel at all points of the domain, the multiscale algorithms, based on such processing, proved ideal for implementation on massively parallel computers.

To obtain even further generality, there emerged however two basic reasons to go much beyond these multigrid methods. First, they do not perform well for *highly nonlinear cases*, where configurations cannot be decomposed into weakly-interacting local and non-local parts. Second, for many systems, even 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 governing 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 governing equations that are valid everywhere, and that can be derived from fine-level processing conducted only in some small representative "windows" (see below).

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 should use coarse level variables that are little sensitive to relaxation (e.g., representing chosen *averages*, rather than a subset of individual fine-level values) and that represent the *full* solution rather than the correction to any given current approximation. Such coarse variables can be chosen (as described below) so that the coarse-level equations can be derived just by local processing. We use the term "*upscaling*" for this type of direct (full-solution) transition from a fine level to a coarser one. Such a transition is valid even in 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

S0010-4655(05)00194-3/FLA AID:2927 Vol.●●●(●●●)
COMPHY:m3 v 1.36 Prn:29/03/2005; 8:48
cpc2927
[DTD5] P.3(1-4)
by:Diana p. 3

where all scales interact most strongly: the highly nonlinear systems of statistical mechanics at the critical temperature of phase transition. The *renormalization group* (RG) method (see, e.g., [14,15]) 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. The RG method 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. Little has been done to upscale systems without a fixed point, which is the more prevalent situation in most practical problems. Also, the RG computational efficiency remained limited, due to the lack of a systematic coarse-to-fine transition, which is needed either for accelerating the fine-level simulations (as in multigrid solvers) and/or for confining them to small representative windows (as described below).

Realizing the complementary advantages and drawbacks of RG and multigrid, the following new combined paradigm has emerged.

3. Systematic upscaling (SU): An outline

Local equations and interactions. Computationally we deal only with discrete systems; their n variables $u_1, u_2, \ldots, u_n$ will typically be either values of discretized functions (grid values, or finite elements, etc.), or locations of particles. An equation or interaction is called *local* if it involves only O(1) neighboring variables. For simplicity of discussion we describe SU first for systems of local equations (including energy minimization with local interactions) or local interactions at equilibrium. Long range interactions are discussed later. *Extensions exist to dynamic and non-equilibrium systems*.

Coarsening. Similar to multigrid, SU is based on two processes: The usual local processing (relaxation in deterministic problems, Monte Carlo (MC) in stochastic ones) and repeated coarsening, creating increasingly coarser descriptions of the same physical system. At each coarsening stage, to each fine-level configuration $u = (u_1, \ldots, u_n)$ one defines (using the criterion below) a unique coarse-level configuration, denoted $u^c = (u^c_1, \ldots, u^c_m)$, which is a vector with a reduced number of variables; typically $.1n < m < .6n$.

Interpolation. To any given coarse configuration $U = (U_1, \ldots, U_m)$, there are of course many fine-level configurations u *compatible* with U (i.e., such that $u^c = U$). The interpolation (transition from U to a specific fine configuration u) is created by *compatible Monte Carlo* (CMC) (or compatible relaxation, in the deterministic case), i.e., by the local processing, restricted to configurations compatible with U. The interpolation is completed once the CMC has reached its equilibrium (or the compatible relaxation has converged).

The general coarsening criterion. The fine-to-coarse transformation is said to be *adequate* if (and to the extent that) the CMC equilibrates fast (or the compatible relaxation converges fast) independently of the system size n. For example, in the 2D Ising model at critical temperature, for the 2×2 majority blocking, the CMC autocorrelation time is very close to 1 [4].

A major problem in coarsening any system is to find a suitable set of coarse variables. The above criterion gives a general and very effective tool for developing such a set. The adequacy of that set implies essentially local dependence of the coarse variables, and hence the feasibility to construct, just by local processing, a set of *operational rules* (e.g., equations, or a Hamiltonian-like functional) that will govern simulations at the coarse level.

In highly repetitive systems (defined above), the local processing need not be done everywhere: the coarse-level equations can iteratively be derived by comparing coarse-level with fine level simulations, where the latter are performed only in some relatively small *windows* (subdomains, on the boundaries of which the fine level is kept compatible with the coarse level).

Thus, the fine level simulations supply the operational rules for the next coarser level, while the latter supplies the windows for those simulations. Iterating back and forth between all the levels quickly settles into a self-consistent multilevel compatibility. If the coarsening ration n/m is not large, no slowdown should occur, and at each level the computations need extend only over a collection of small windows, whose number depends on the diversity of local situations, not on the size of the problem.

S0010-4655(05)00194-3/FLA AID:2927 Vol.●●●(●●●)
COMPHY:m3 v 1.36 Prn:29/03/2005; 8:48
cpc2927
[DTD5] P.4 (1-4)
by:Diana p. 4

Long range interactions (e.g., between electrostatic charges) can each be decomposed into a smooth interaction and a local one ("smooth" and "local" being meant on the scale of the next coarse level; all familiar physical interactions can be decomposed this way [17]). The smooth part can directly be transferred to the coarse level (e.g., represented by aggregated charges), while the local part is transferred, together with all other local interactions, using the fine/coarse iterations described above.

Determinism and stochasticity. A stochastic system at the fine level often yields a deterministic system at large enough scales, and vice versa. The coarsening approaches described above can accommodate such transitions.

Reported examples of SU-type systems include an Ising spin model [4], a simple polymer and simple fluids (see §14 in [18] and references therein and also [19]). Several other systems are under development.

Acknowledgement

The research is supported by the Israel Science Foundation grant 295/01 and by the US Air Force Office of Scientific Research, contract F33615-03-D5408.

References

[1] A. Brandt, Multi-level adaptive solutions to boundary value problems, Math. Comp. 31 (1977) 333–390.
[2] A. Brandt, Guide to multigrid development, in: W. Hackbusch, U. Trottenberg (Eds.), Multigrid Methods, Springer-Verlag, 1982, pp. 220–312.
[3] A. Brandt, I. Livshits, Wave-ray multigrid method for standing wave equations, Electronic Trans. Num. An. 6 (1997) 162–181.
[4] A. Brandt, D. Ron, Renormalization multigrid (RMG): Statistically optimal renormalization group flow and coarse-to-fine Monte Carlo acceleration, Gauss Center Report WI/GC–11, June 1999, J. Stat. Phys. 102 (2001) 231–257.
[5] A. Brandt, S. McCormick, J. Ruge, Algebraic multigrid (AMG) for automatic multigrid solution with application to geodetic computations, Institute for Computational Studies, POB 1852, Fort Collins, Colorado, 1982.
[6] A. Brandt, Algebraic multigrid theory: The symmetric case, Preliminary Proc. Int. Multigrid Conf., Copper Mountain, Colorado, April 6–8, 1983; Appl. Math. Comp. 19 (1986) 23–56.
[7] A. Brandt, D. Ron, D.J. Amit, Multi-level approaches to discrete-state and stochastic problems, in: W. Hackbusch, U. Trottenberg (Eds.), Multigrid Methods, II, Springer-Verlag, 1986, pp. 66–99.
[8] A. Brandt, Multigrid methods in lattice field computations, Nucl. Phys. B (Proc. Suppl.) 26 (1992) 137–180.
[9] A. Brandt, M. Galun, D. Ron, Optimal multigrid algorithms for calculating thermodynamic limits, J. Stat. Phys. 74 (1994) 313–348.
[10] W. Hackbusch, Multigrid Methods and Applications, Springer, Berlin, 1985.
[11] U. Trottenberg, C.W. Oosterlee, A. Schüller, Multigrid, Academic Press, London, 2000.
[12] W.L. Briggs, V.E. Henson, S.F. McCormick, A Multigrid Tutorial, 2nd Ed., SIAM, 2000.
[13] J. Ruge, K. Stüben, Algebraic multigrid, in: S.F. McCormick (Ed.), Multigrid Methods, SIAM, Philadelphia, 1987, pp. 73–130.
[14] K.G. Wilson, The renormalization group and critical phenomena, Rev. Mod. Phys. 55 (1983) 583–600 (1982, Nobel Prize Lecture).
[15] M.E. Fisher, Renormalization group theory: Its basis and formulation in statistical physics, Rev. Mod. Phys. 70 (2) (1998) 653–681.
[16] J. Goodman, A.D. Sokal, Multigrid Monte Carlo methods for lattice field theories, Phys. Rev. Lett. 56 (1986) 1015–1018.
[17] A. Brandt, Multilevel computations of integral transforms and particle interactions with oscillatory kernels, Comp. Phys. Comm. 65 (1991) 24–38.
[18] A. Brandt, Multiscale scientific computation: review 2001, in: T.J. Barth, T.F. Chan, R. Haimes (Eds.), Multiscale and Multiresolution Methods: Theory and Applications, Springer-Verlag, Heidelberg, 2001, pp. 1–96. Available in www.wisdom.weizmann.ac.il/~achi.
[19] Multiscale computation of polypeptide backbone, cond-mat/0312185.