Think Globally, Act Locally: An Introduction to Domain-based Parallelism and Problem Decomposition Methods

David E. Keyes [*]   Yousef Saad [†]   Donald G. Truhlar [‡]

"Think globally; act locally." This bumper sticker maxim has a lot to say to practitioners of contemporary high performance computing. It is increasingly incumbent on computational scientists to respect the data access hierarchies that accompany the large memories required by applications programs. These hierarchies are imposed, ultimately, by the finite size of data storage media and the finite speed of light, but their presence is asserted more immediately by the hardware and software overheads of system protocols for the delivery of data. From the frame of reference of any given processing element, an approximate cost function can be constructed for the minimum time required to access a memory element that is any given logical or physical distance away. Such cost functions typically consist of plateaus separated by sharp discontinuities that correspond to software latencies where some boundary of the hierarchy, such as a cache size or a local memory size, is crossed. The ratio of times required to access remote and local data varies from 10 to 10^5 in typical architectures, the latter being characteristic of network cluster computing. An underlying motivation for the development of problem decomposition algorithms is that these discontinuities should explicitly be respected by user applications. If users cannot afford to treat memory as "flat" in large problems, then neither can they afford to treat all nonzero data dependencies on an equal footing. Consequently, algorithms must adapt to architecture, guided by knowledge of the relative strengths of different couplings from the underlying physics. Ironically, such forced adaptation sometimes results not in compromise, but in the discovery of intrinsically better methods for flat memory environments, as well.

Steady-state natural and human-engineered systems are often zero-sum

[*]Department of Computer Science, Old Dominion University, Norfolk, VA 23529-0162 and Institute for Computer Applications in Science and Engineering, NASA Langley Research Center, Hampton, VA 23681-0001. *email:* keyes@icase.edu.

[†]Department of Computer Science, University of Minnesota, Minneapolis, MN 55455-0154. *email:* saad@cs.umn.edu.

[‡]Minnesota Supercomputer Institute, 1200 Washington Ave. S., Minneapolis, MN 55415. *email:* af12101@sc.msc.edu.

networks in which the overall distribution of a quantity to be determined is conserved. The conservation principle holds over any size control volume, from the smallest scales requiring resolution up to the global domain. Somewhere between these extremes are the scales at which the latencies of the memory hierarchy are asserted. This suggests a multilevel discretization of the conservation laws, with coarse-grained interactions between "basins" of fast memory (thinking globally, but on a small problem) and with fine-grained interactions within them (acting locally, on the scales of the resolution required). Algorithms exploiting multilevel discretization have evolved naturally and somewhat independently in a variety of applications, both continuous (e.g., conservation of energy in a conducting body) and discrete (e.g., conservation of current in a network of electronic components). It is an objective of this volume to promote cross-fertilization of such applications by identifying analogous features between them.

It may be assumed without loss of generality that the challenges of writing algorithms for large-scale problems on hierarchical memory systems occur for physical systems that are irreducible in the matrix theoretic or group theoretic sense of the term. Each degree of freedom depends upon all of the others; no degrees of freedom may be removed and solved for exactly in isolation. For irreducibly coupled physical systems with *arbitrary* interactions between the components, there is not necessarily any benefit to a decomposition of the unknowns of the problem into sets that are proximate (in space) or strongly coupled (by dynamics) and a mapping into the global memory in a way that preserves their proximity or strong coupling. However, the interactions in the systems studied herein decay with an appropriate "distance" (in physical or basis function index space) sufficiently rapidly that remote interactions may be lumped or even ignored in certain phases of the solution process.

There is a history of applying both direct and iterative methods to such problems. Direct methods involve the construction by explicit condensation of lower-dimensional systems for degrees of freedom that act as separators. In the literature of differential equations, this is the Poincare-Steklov operator; in linear algebra, it is the Schur complement; in physics, it is the optical potential. The simplest iterative methods involve cycling between the subdomains whose unknown boundary data are updated by neighbors and may generically be called Schwarz methods. Many modern approaches combine direct and iterative aspects in the form of preconditioned Krylov methods.

The trade-offs involved in deciding what couplings may be lumped or ignored, with what consequences in terms of convergence rate or accuracy, and with what benefits in terms of mapping the computation to the memory hierarchy, constitute one of the main themes of this volume. A key concept in this regard is the selection of a reduced basis in which to represent the solution of a large-dimensional problem. This is an explicit choice in some

cases (as in a wave expansion), automated but still explicitly identifiable in some others (as in a Krylov method), and implicit in yet others (as in a multilevel or multipole method). In several chapters of this volume, the authors have brought out the benefits that accrue from selecting a good basis. These benefits range from getting any handle on the problem at all, to making a quantifiable asymptotic complexity reduction relative to a full-dimensional method, to identifying "reusable" bases for recurring computational tasks. A "good" basis is usually physically motivated (or problem-fitted), hierarchical, or orthogonal, and such good bases permit the solution process to be separated into distinct parts. A physically motivated or problem-fitted basis separates components of the result into dominant parts that may be suggested by some physical approximation and subdominant parts to patch in for more accuracy. A hierarchical basis separates components of the solution by their scales of variation. Expansion in an orthogonal basis provides another way to separate the components of the solution. Of course, these three attributes of a good basis are not mutually exclusive.

A signature of the choice of basis visible in some of the chapters is an expression of a key resolvent operator, or an approximation thereto, by a sum containing triple products of operators consisting of the inverse of a different-dimensional operator in the middle, with "rectangular" operators on either side that map between spaces of different dimensions. For instance, a Schur complement contains such triple products in which the middle term may be of higher dimension than the terms of the sum itself. A Schwarz preconditioner contains such triple products in which the middle term is of lower dimension. The "rectangular" operators can even be infinite dimensional in the long direction. In the chapters describing quantum chemistry applications, these triple products are sometimes expressed in bra and ket notation, while in the chapters originating from a problem in the continuum, linear algebraic expressions may be found.

Several other themes arise that transcend disciplinary barriers and are common within subsets of the chapters. These include:

1. opportunities to bring a physical understanding of the continuous problem into the discretization or the decomposition, particularly in the selection of partitions in problems in which the decay metric is anisotropic;

2. multiple discretizations of the same problem (e.g., on different scales, or to different orders of accuracy);

3. trade-offs in linear and nonlinear convergence rates that are mediated by a time-like parameter that stabilizes the nonlinear iteration while accelerating the linear iteration (by steepening the algebraic decay rate of the interactions at the same implicit time level), at the price

of requiring many such time steps;

4. opportunities for reuse of computational results from one iteration on related problems in subsequent iterations;

5. opportunities for and experience with parallel implementations.

In the rest of this introductory chapter we discuss several examples of of problem decomposition methods more specifically, each of which is the subject of one of the following chapters.

Xiao-Chuan Cai presents the classical Schwarz domain decomposition approach for the solution of elliptic and parabolic problems with operators that are dominated by the self-adjoint second-order terms, but need not be either self-adjoint or even definite. With a fixed geometric overlap between neighboring subdomains, and with a single coarse-grid problem involving approximately one degree of freedom per subdomain as part of the preconditioner at each Krylov iteration, an iteration count bound that is asymptotically independent of both the resolution of the problem and the number of subdomains can be achieved. The coarse-grid solution being critical, recent work examines how to obtain the coarse-grid operator in the context of irregular grids and decompositions.

Alfio Quarteroni describes domain decomposition methods for hyperbolic problems, in which characteristics play an essential role in selecting partitions and imposing interfacial boundary conditions. Scalar convection problems and systems of conservation laws are addressed, with applications from acoustics and elasticity. The author considers three examples of wave equations describing convective, acoustic, and elastic waves. He illustrates how these problems can be reformulated in the framework of a decomposition of the spatial domain and devises algorithms based on subdomain iterations. Finally, he addresses the interaction of time-differencing and space decomposition.

Petter Bjørstad and Terje Kårstad's contribution on two-phase immiscible, incompressible flow in oil reservoir simulation spans the subject matter of both of the first two chapters with an operator splitting that separately exploits the hyperbolic and elliptic features of the governing system of PDEs. The hyperbolic part of the problem is solved by a modified method of characteristics. Of particular interest is the resulting conflict between the optimal parallel mappings of the two split subproblems. In spite of the compromise, this chapter makes a strong case for the practicality of high-granularity parallel solutions to problems of real-world complexity. In particular, the resulting computational problems involve up to 16,384 subdomains (with one-element-wide overlap at their boundaries) and a coarse space. The solution is achieved via data parallel implementation with one subdomain per processor, approximate subdomain solvers, and a multigrid approach on the coarse grid.

V. Venkatakrishnan presents parallel solution techniques for the highly nonsymmetric Jacobian systems that arise when the convectively dominated Navier-Stokes equations are discretized on unstructured grids and solved by Newton's method. For these multicomponent problems, a coarse-grid operator leading to an optimal convergence rate is not known; nevertheless, a coarse system derived from agglomeration proves effective. The equations are solved by a preconditioned iterative method with a block diagonal preconditioner corresponding to a fixed sparsity pattern and involving a factorization within each processor subject to homogeneous Dirichlet boundary conditions. Such boundary conditions become more and more accurate as the outer Newton iteration progresses. Partitioning, node ordering, and the accuracy with which subdomain problems should be solved for most efficient solution of the overall steady-state problem are addressed. An implicit scheme for unstructured grids is demonstrated that requires fewer iterations for a given nonlinear residual reduction than the best single-grid method.

Dana Knoll and co-authors extend Krylov-Schwarz domain decomposition methods without a coarse-grid operator to nonlinear problems. The edge plasma fluid equations are a highly nonlinear system of two-dimensional convection-diffusion-reaction equations that describe the boundary layer in a Tokamak fusion reactor. There are six or more components with complicated interactions through composition-dependent transport coefficients and source/sink terms. A matrix-free version of Newton's method exploits the Krylov nature of the solver (in which the action of the Jacobian is probed only through matrix-vector products) to avoid forming the actual Jacobian of the nonlinear system, except for diagonal blocks used only in preconditioning and updated infrequently. Matrix-free methods depend critically upon numerical scaling since they approximate matrix-vector products through a truncated Taylor series. The implications for the robustness of various Krylov solvers are explored.

William Gropp and Barry Smith present an implementation philosophy and a publicly available implementation in portable parallel software of a variety of preconditioned Krylov algorithms for domain decomposition, in which the notion of subdomain is generalized to the block partitioning of a sparse matrix. The emphasis is on performance of such solvers on a variety of distributed memory architectures in the limit of large problem size, and the resulting trade-offs in convergence rate and parallel efficiency.

Andrew Lumsdaine and Mark Reichelt discuss the spatio-temporal simulation of semiconductor devices via accelerated versions of the waveform relaxation method, a classical method for systems of temporally varying ordinary differential equations. In contrast to conventional parabolic treatments, in which space parallelism only is sought at each time level, the entire space-time cylinder is partitioned for parallel processing purposes. Time, being causal the initial value problems under consideration here,

invites a special windowing treatment.

Graham Horton applies two-level and multilevel discretizations beyond the realm of PDEs to steady-state Markov chains, which arise, for instance, in queuing theory, and in the performance analysis of networks. Of particular interest is the derivation of a coarse-grid correction scheme that never violates the feasibility range of bounded variables, in this case probabilities. The resulting scheme is equivalent to a conventional multigrid method but with nonlinear (solution-dependent) intergrid transfer operators. Simple queuing networks with highly anisotropic coefficients, for which the novel multilevel method is particularly effective, are seen to have the same algebraic structure as convectively dominated transport equations.

Charbel Farhat also focuses on the coarse level of a multilevel preconditioner, from a parallel efficiency point of view and in the context of multicomponent problems of structural mechanics. The practically important problems of multiple right-hand sides in engineering analyses and how to amortize for multiple right-hand sides in the context of iterative methods are also addressed. Of particular interest are the extensions of domain decomposition methods for "nearby" systems that arise in design problems, time-dependent problems, and eigenvalue problems. Scalable results are demonstrated for structural mechanics problems.

Francois-Xavier Roux presents the dual Schur complement method of domain decomposition with application to nonlinear elasticity problems, and shows the dual to be preferable from a spectral convergence theory point of view. Along with Farhat, he addresses reuse of previous right-hand side work in reconjugation and extends to nonlinear cases in which the matrix also changes. Parallel implementation on distributed-memory parallel machines is discussed.

Roland Glowinski and co-authors show how domain decomposition and domain embedding techniques, seemingly complementary techniques for making irregular geometry amenable to acceleration by fast solvers, may be merged in the solution of both elliptic and time-dependent problems. This approach is based on using an auxiliary domain with a simple shape that contains the actual domain with a more complicated shape.

Jacob White and co-authors exploit the fast multipole and fast Fourier transform methods in the context of a boundary element discretization of electrostatic potential problems. Boundary element formulations lead to dense matrix operators of sufficient diagonal dominance and superior conditioning that rapid convergence of Krylov methods can be obtained without complex preconditioners; however, the matrix-vector multiply is dense, and hence expensive. The fast multipole method applies the action of the underlying operator without forming it explicitly, resulting in order-of-magnitude reductions in asymptotic complexity while guaranteeing an arbitrary given accuracy in the result. The techniques are applicable to a wide variety of engineering applications based on $1/r^2$ interactions.

The remaining chapters illustrate how problems and solutions analogous to those in mechanics applications in the preceding chapters also arise in quantum mechanics. In modern quantum mechanics, one works in basis function space rather than physical space, but the space is still structured into subsets that are strongly coupled within and weakly coupled between. Although the various quantum mechanical problems discussed have significant differences, there are recurring themes such as basis set contraction, which occurs in one way or another in all of these chapters.

The chapters of Ellen Stechel and Hans-Joachim Werner are concerned with large-scale electronic structure problems, which involve elliptic eigenvalue problems of very large dimension. Contraction occurs at several levels in electronic structure problems. Stechel includes an overview of recent attempts to reach the ultimate scaling limit whereby the computational effort scales linearly in the number of particles or dimensions. Some of the techniques employed are very similar to the work described by White. Werner reviews modern numerical methods for the treatment of electron correlation effects, including the internally contracted configuration interaction method in which sets of physically related many-body basis functions are treated as a single degree of freedom to reduce the size of the variational space. He also discusses the vectorization and parallelization strategies that are required to make the resulting algorithms efficient, including techniques for iterative solution of large matrix eigenproblems, solution of nonlinear equations in multiconfiguration self-consistent-field and coupled-cluster approaches, and the use of direct inversion on an iterative subspace. Problems of vectorization, parallelism, input/output bottlenecks, and limited memory are addressed, and the I/O bottleneck is addressed by disk striping. This provides an example of parallelism in communication that seems less widely discussed than parallelism associated with multiple processors.

Zlatko Bačić and Georges Jolicard and John Killingbeck discuss the vibrational eigenvalue problem in quantum mechanics. Bacic introduces the discrete variable representation (DVR), in which the analogies between function spaces and physical spaces are very clear, and he presents DVR-based divide-and-conquer computational strategies for reducing the dimensionality of the Hamiltonian matrix. Jolicard and Killingbeck discuss the wave operator theory as a tool to define active spaces and simplified dynamics in large quantum spaces. They present a partitioning integration method for solving the Schroedinger equation based on projections in reduced active spaces. For the Floquet treatment of photodissociation experiments, the choice of the relevant subspaces and construction of the effective Hamiltonians are carried out using the Bloch wave operator techniques. Recursive methods for the solution of the basic equations associated with these operators, based on Jacobi, Gauss-Seidel, and variational schemes are given.

David Schwenke and Donald Truhlar discuss large-scale problems in quantum mechanical scattering theory. In quantum mechanical scattering

theory the basis functions may be delocalized, and they are typically grouped
in sets associated with channels. At the highest level, associated with distortion potential blocks, Schwenke and Truhlar explicitly couple those channels
which physical arguments indicate are the most strongly interacting. At the
intermediate level, they can perform a sequence of calculations increasing
in complexity, optimizing the (contracted) basis functions at each step. At
the lowest level, they discuss replacing a class of weakly coupled channels
with a phenomenological optical potential. The optical potential idea can
also be introduced using a different kind of motivation for the partitioning,
as a way to reduce the computational effort by partitioning the energy-independent parts of the problem from the energy-dependent parts. The
resulting "folded" formulation has interesting computational analogies to
domain decomposition although it is accomplished in basis function space
rather than physical space. Finally the partitioning based on strength of
coupling can be re-exploited by solving the coupled equations iteratively
with preconditioners blocked by the same physical considerations as were
employed to block the distortion potentials.

The work summarized above underscores the importance in large problems of informing the solution process directly with the physics being modeled and with the architecture for which the computation is destined, and
portrays the tension between concentrating operations locally and taking
strategic account of remote information that dominates parallel algorithm
development today and for the foreseeable future.

DOMAIN-BASED PARALLELISM AND PROBLEM DECOMPOSITION METHODS IN COMPUTATIONAL SCIENCE AND ENGINEERING
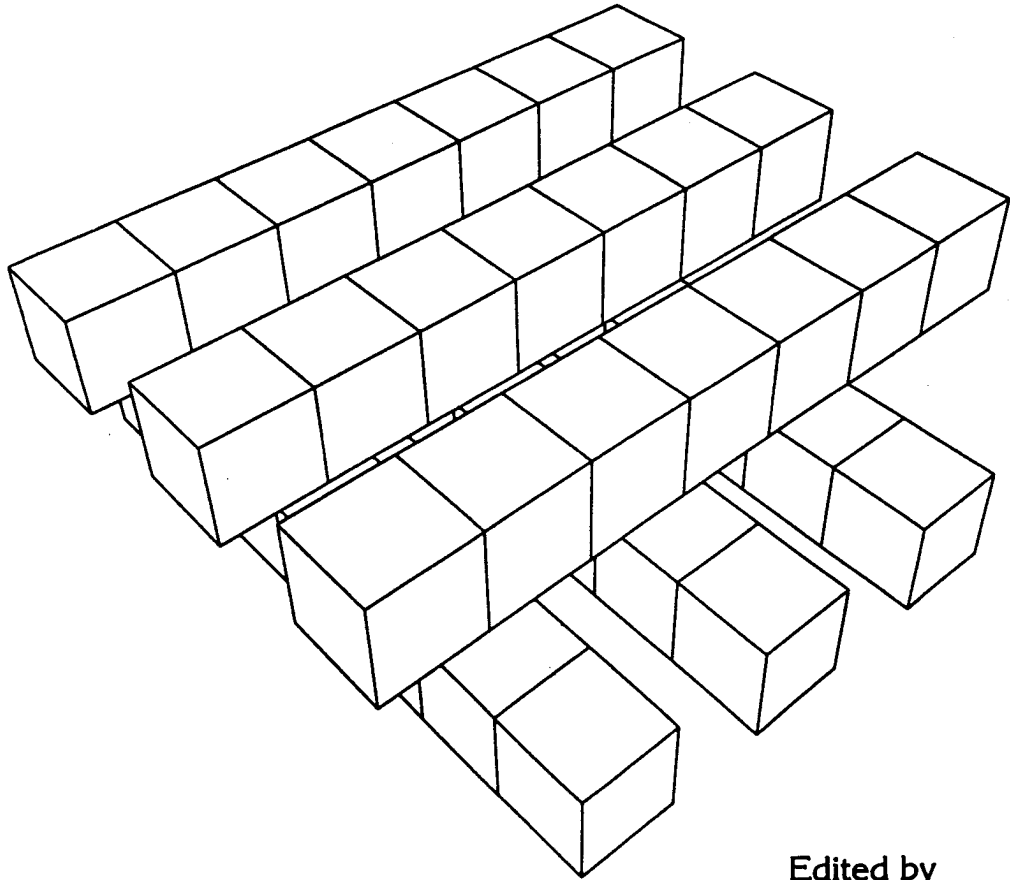
Edited by

□□□

David E. Keyes
Old Dominion University and ICASE
NASA Langley Research Center

□□□

Youcef Saad
University of Minnesota

□□□

Donald G. Truhlar
Minnesota Supercomputer Institute