1 Prior NSF Supported Work by PIs

1.1 Results from prior NSF support – J. Bernholc

Bernholc was supported by grant "Electronic Properties and Growth of Semiconductors," DMR 9408437, 7/1/94-6/30/97, \$225,000, and a two-year Creativity Extension 7/1/97-6/30/99, \$186,000. The main emphasis of this grant was on the initial stages of growth and doping of semiconductor materials and on the development of new theoretical methods for electronic structure calculations. During the five-year reporting period this research has led to 38 publications, 44 invited talks at conferences and 37 contributed presentations. The grant provided support for five graduate students and partial support for several postdoctorals. The PI has also served on Organizing and Program Committees of 9 conferences directly related to the grant, and was a Finalist for a Computerworld Smithsonian Award in the Science category in 1997.

Of particular interest to the present proposal is the development of a new method for computing quantum-mechanical properties of solid state and molecular systems. This method uses a real-space mesh to represent the wave functions, the charge density, and the ionic pseudopotentials. The density-functional equations are approximated using a *generalized* eigenvalue form:

$$\mathbf{H}_{mehr}[\psi_n] = \frac{1}{2}\mathbf{A}_{mehr}[\psi_n] + \mathbf{B}_{mehr}[V_{eff}\psi_n] = \epsilon_n \mathbf{B}_{mehr}[\psi_n],$$

where $\mathbf{A}_{mehr}$ and $\mathbf{B}_{mehr}$ are the components of the *Mehrstellen* discretization [44]. It uses a weighted sum of the wave function and potential values to improve the accuracy of the discretization of the *entire* differential equation, not just the kinetic energy operator. The short-ranged representation of the Hamiltonian leads to an efficient domain-decomposition-based implementation on massively parallel computers. Indeed, the calculations show a linear decrease in execution time up to the maximum number of SGI/Cray T3D and T3E processors that we could access (1024).

The details of the current multigrid implementation, which include new compact discretization schemes in real space for systems with cubic, orthorhombic, and hexagonal symmetries, and multilevel algorithms for the iterative solution of density functional equations are described in Ref. [28, 29, 12]. This methodology was tested and applied to a large number of systems. These include the C_{60} molecule, diamond, Si and GaN supercells, and *ab initio* MD simulations. Further applications include a simulation of surface melting of Si, properties of Si, GaN, InP surfaces, $In_xGa_{1-x}N$ quantum wells, and initial studies of a large amyloid β peptide ($C_{146}O_{45}N_{42}H_{210}$) implicated in Alzheimer's disease. The accuracy of the method was also tested by direct comparison with plane-wave calculations and Car-Parrinello dynamics [30], when possible, and the results were in excellent agreement in all cases.

1.2 Results from prior NSF support – T. Chan

Chan has a current NSF ACR contract ASC-9720257 on multilevel methods which will expire on 8/31/2000. In the current contract period, we worked on the design of efficient parallel multilevel algorithms for solving large systems of PDEs arising in CFD and other large scale scientific applications, such as the image restoration problem described later in this proposal. We give some highlights.

1. We have developed efficient agglomeration based algebraic methods for elliptic problems on unstructured grids [42]. These algorithms possess near optimal computational efficiency and yet are fairly easy to use due to their algebraic nature.

2. We also developed a unifying energy minimizing framework for deriving robust interpolation operators in algebraic multigrid methods [95]. The resulting method is robust with respect to anisotropy, jumps and oscillations in the PDE coefficients, unstructured grids and even integral operators.

3. We developed efficient non-overlapping domain decomposition algorithms for CFD problems [11].

4. We wrote an invited survey article on robust multigrid [41] that will appear in a special millenium issue of a major journal.

5. We developed some modified multigrid algorithms for the differential-convolution problems arising in image restoration [33].

6. We developed a multilevel spectral partitioning method for unstructured grids which is based on the algebraic multigrid technology we developed for elliptic solvers on these grids. We showed that spectral partitioning can be obtained at multigrid efficiency [38].

7. We developed domain decomposition algorithms for solving inverse elliptic problems, with potential applications to oil reservoir simulations [92].

In addition, Chan has supervised two PhD thesis in this research area: Susie Go (1998) and Justin W.L. Wan (1998). Go is working at HP Labs in Palo Alto and Wan is a Forsythe Fellow at the Computer Science Dept at Stanford and is a co-PI of this proposal. Chan has given invited plenary talks at the 10th GAMM Multigrid Workshop in 1998 and the European Multigrid Conference in 1999. He also gave a plenary talk on the image processing work at the Dundee Numerical Analysis Conference in 1999.

1.3 Results from prior NSF support – J. Xu

Xu has been continuously supported by National Science Foundation since 1989. Currently ongoing grants include: "*Theory and Application of Numerical Methods for Partial Differential Equations*" under DMS-9706949 for the period 8/1/1997-7/31/2000 in the amount of $150,000 and "*Parallel Multilevel PDE Solvers on Unstructured Meshes*" under ACI-9800244 in the amount of $90,000 for the period 7/15/1998-6/30/2000.

Xu has done a lot of work on numerical methods for partial differential equations and especially on multigrid and domain decomposition methods. Below is a brief summary of some results obtained mainly in the last *three* years. More detailed information on Xu's research may be found in his web page: `http://www.math.psu.edu/xu/`.

Continuing with relevant earlier works, [15, 14, 10, 102], Chan, Xu and Zikatanov [37, 35] developed new multigrid methods for unstructured grids based on agglomeration techniques (using graph level-set) and is now in the process of completing a *Fast Poisson Solver* for solving the Poisson equation discretized on general unstructured grids in both two and three dimensions. Continuing his early works on two-grid techniques (see [101, 100, 103]), Xu and Zhou developed a very efficient two-grid method for eigenvalue problems in [107] and a number of new local and parallel discretization and adaptive finite element algorithms in [105, 106]. Using partition-of-unity, Huang and Xu [68] devised a simple and effective finite element method for general non-matching (with or without overlapping) grids. Xu and Zikatanov [104, 108] developed an elegant monotone finite element scheme, called EAFE, for convection-diffusion equations in any spatial dimensions. For the solution of these monotone schemes, Wang and Xu [98] devised a cross-wind-strip block method based on block Gauss-Seidel iterations and obtained its rate of convergence. Kim, Xu and Zikatanov [75, 76] developed some efficent and robust iterative methods for convection diffusion equations by combining various techniques such as graph matching for defining special coarse subspaces and preconditioned GMRES with monotone scheme preconditioning.

In recent years, Xu also made a lot of efforts for developing efficient multigrid algorithms and numerical codes for several engineering applications including *subsurface flow simulation* (with Prof. G. Yeh, Civil Engineering, Penn State), *powder compaction simulation* (with Prof. V. Puri, Agricultural and Biological Engineering) and *electroochemical power device simulation* (with Prof. C. Wang, Mechanical Engineering). In each of these projects, the new code we developed based on our more advanced iterative methods dramatically outperformed the old codes.

Two graduate students, Hwanho Kim and Hongrong Qin, have been supported as research assistants. Kim became a postdoctoral fellow at Duke university last year. Dr. Ludmil Zikatanov and Dr. Mohammed Aiffa have also been *partially* supported as postdoctoral fellows.

2 Background and Motivation

The paradigm of multilevel algorithms is one of the most fundamental and powerful enabling technologies for high performance computing in support of computational science, which is an important component of information technology research. The reason is easy to understand. One of the most common reasons for an application domain to require high performance computing in the first place is due to the presence of a large number of degrees of freedom, coupled globally, leading to a large spectrum of scales in the problem which must be dealt with efficiently in any computational algorithm. A well known example is in computational material science, which is one of the main foci of our proposal. As is well-known, molecules, liquids, solids and all other macroscopic matter bind and interact due to Coulomb forces. On a *large* scale, the properties of matter can be described by the laws of classical physics and a continuum description suffices. At a smaller, *microscopic* scale, quantum-mechanical effects become essential and one needs to solve the Schrödinger equation for electrons that interact both among themselves and with the atomic nuclei. However, the properties on the *small* scale, e.g., the strength of bonding or the change in the electrical properties due to doping, will clearly affect those on the macroscopic scale. In a number of cases, these microscopic effects, once measured or computed, can be described in terms of a few macroscopic parameters, such as compressibility or carrier density, and macroscopic modeling can continue. On the other hand, macrocosm also affects the microcosm – the applied electric voltage can supply or deplete electrons – or a large strain can cause fracture originating in some well-defined weak spots.

The foregoing example provides an admittedly very simple illustration of a physical *multiscale* problem, where *back and forth* feedback is essential. Obviously, such problems also exist in the *temporal* domain, where small time-scale behavior affects large-scale motion, etc. In computational materials physics, the ultimate goal of any theory would be to develop a general scheme able to describe the properties of any material at any length-scale, and determine its time evolution. While a full achievement of this goal is some time off, several developments in mathematics, physics, and chemistry promise rapid progress in the near future.

A naive computational approach to such problems will involve an unnecessarily inefficient complete interaction of all degrees of freedom, leading to a computational complexity that grows much faster than linear in the degrees of freedom, thus severely limiting the size of problems a given computer system can handle. Such problems are ubiquitous in the computational sciences, from many particle systems in chemistry, plasma physics, astronomy, to multiscale modelling in material science and meterology, and to other application domains which are particularly relevant to ITR, such as in data mining, image processing, VLSI placement, device and circuit simulation. In our proposal, we will focus on a carefully selected subset of problems from this wide spectrum of problems.

What makes these problems solvable at all is that not all the possible interactions are equally strong — often the *strength* of interaction decays rapidly with some sort of distance, physical or logical. Algorithms which judiously exploit this fact can lead to a complexity that is closer to the optimal linearity. Almost invariably, such algorithms are *multilevel* in nature, proceeding in a hierarchical way, at each step using only efficient local interactions, and achieve effective global interaction *implicitly* through coarse representations of the original problem. Examples include the multigrid and domain decomposition algorithms for solving partial differential equations (PDEs) such as problems in computational fluid dynamics (CFD), the fast multipole method (FMM) and related tree algorithms for many particle problems, and multilevel algorithms for graph partitioning. In fact, it has been argued [40] that computer systems themselves should also be built in a hierarchical and multilevel way, to more naturally accomodate these algorithms which in turn adapt naturally to the "physics" of the problems they are solving. This is indeed turning out to be true, in view of the hierarchical memory systems employed in most high performance computing architectures. Together, multilevel and hierarchical algorithms and architectures is the only way to achieve true scalability — the "holy grail" of high performance computing. Our proposal will also address this issue.

Without doubt, multilevel algorithms have achieved spectacular success in many traditional scientific computing application domains. Yet we believe that there are still many application domains where multilevel methods have not yet played a critical role but potentially can have a lot to offer. These include image processing, data mining and material science, areas which we will address in more details later in this proposal.

Irrespective of the particular applications, there is a common unifying basic philosophy of multilevel methods: the creation of multiple levels of discrete representation of the same data and the same operators, transfers between these levels, and some kind of processing on each level. Thus what one learns from one domain often turns out to be very useful in another. This argues for an approach which, in addition to focusing on a small number of specific problems, also looks at similarities, generalizations and extensions in other problems. This cross-fertilization approach is the one taken in this proposal. For example, the vast knowledge base developed over the years for multilevel methods for solving PDEs can be applied to recently developed variational and PDE based image processing. In turn, by making the close analogy between image segmentation and clustering in data mining, classical multilevel methodologies can be brought to bear on this new class of problems. Another example is the recently developed *grid based* electronic density function theory opens up the possibility of using classical multilevel techniques to be applied efficiently to computationa material science problems. Conversely, what we learn in the new problem domains will surely bring new insights and understandings which can in turn be applied to classical PDE problems. For a broad survey of the variety of applications of the multiscale methodology to scientific computations, see the recent survey by Brandt [21].

The aim of our consortium of applied mathematicians, computer scientists, and application domain experts is to apply lessons learned and software constructed in one successful multilevel application in other, quite different, applications, and to uncover new problems and develop new potentials for multilevel algorithms in the process. We build from a foundation of multilevel expertise in our respective application domains, but with a shared vision of both the power and the generality of the multilevel solution framework. Together with our domestic and international collaborators, our team encompasses several major research centers on multilevel technologies, each with its own particular domain of expertise and applications. Among our team, there are already several existing collaborations but this proposal offers the opportunity to bring together our collective expertise for the first time. We expect that our common research will lead to advances in our individual disciplines as by-products of the work – and for this reason we have been able to line up sponsorship from US agencies outside of the NSF (e.g. LLNL) and from major international centers of multilevel research activity in Israel and Germany – but the principal added value of our collaboration is in its cross-disciplinary work.

3 Multilevel Methods in Non-linear Diffusion and Fluid Flow

There are literally countless practically interesting problems that we believe multigrid methods can make a huge difference for their more effective numerical simulations (see Brandt [22]). We choose to focus our attention on two problems, one new application that is important in wireless communication and the other is a well-known problem that is of fundamental practical and theoretical importance. The first problem is to apply multigrid method for the modeling and simulation for batteries and the second problem is to develop more efficient methods for Navier-Stokes equations. Battery modeling and simulation is, for example, apparently very important to wireless communication – one important component of information technology. This is a rather new area for multigrid application. Multilevel methods will be used in both the modeling and the simulation. Euler and Navier-Stokes equations are, however, a very well studied subject, but existing multigrid methods are still not very efficient for these equations in general and there are still some fundamental difficulties that need to be resolved. The resolution of these difficulties not only are important for fluid mechanics problems but will certainly have a great impact in many other applications as well.

3.1 Battery modeling and simulations for wireless communication

The design of batteries, electrochemical power sources, is an important part of many leading industries and inseparable part of our everyday life. Electrochemical batteries are used everywhere and in particular they are one of the most crucial components in *wireless communications*. Optimization of a battery design necessarily involves a large amount of time and experimental effort. Electrochemical power sources such as lead-acid, nickel-cadmium(Ni-Cd), Nickel-metal hydride(Ni-MH), and lithium batteries, as well as various fuel cells, are widely used in consumer applications and electric vehicles. These and future applications place an ever-increasing demand for developing more advanced power sources with higher energy density, higher power density, and longer life. Mathematical modeling and simulation are indispensable in this development process, because a cell model, once validated experimentally, can be used to identify cell-limiting mechanisms and forecast cell performance for design, scale-up, and optimization.

A typical battery can be modeled with a system of time-dependent nonlinear system (see Wang, Gu and Liaw [97] and [51, 49, 43, 65]). This system comes with highly oscillatory coefficients because of the porous electrode used inside the battery. Some volume averaging techniques are used in the aforementioned literature to derive effective coefficients for the porous electrode which consists of a mixture of solid and liquid materials. We propose to use mathematical homogenization techniques (see [45, 8, 77, 78, 79, 83]) and some multilevel upscaling techniques. In recent years, there are a number of attempts using multlevel upscaling techniques for obtaining effective equations. One idea is to use multigrid hierarchical basis ([82])

or wavelets functions ([7, 48]) for discretizing the original PDEs and then use certain approximate Schur complement to obtain an upscaling stiffness matrix. Another idea, first proposed in [9] and then extended to multidimensions in [66, 67], is to use some local harmonic (with respect to the underlying differential operators) functions as basis functions to obtain a coarse scale system. Recently in [69] we combine the aforementioned approach with the partition of unity technique and obtain encouraging results. We believe the technique in [69] has a lot of potential and we plan to have more careful investigation on this method and apply it to the modelling and simulations of battery systems.

After the homogenization or upscaling is properly addressed, we still need to develop efficient techniques for solving the resulting systems of nonlinear partial differential equations. This system consists of a number of coupled highly nonlinear elliptic and parabolic equations. Because of the nonlinear coupling, the design of effective multigrid method is not that easy. We plan to have a systematic study of this problem and develop actual multigrid code for solving this problem. In our preliminary study (see Wang, Wu, Xu and Zikatanov [96]) of this problem, we use Newton's method for linearization and then use a preconditioned GMRES method to solve the linearized system using block Gauss-Seidel method as a preconditioner and multigrid method to solve each block system. This approach gives a lot of improvement over the existing code in some simple 2D models that we have already tested. But there is still a long way to go for us to develop optimal multigrid algorithms together with actual working codes to simulate more general 3D models.

3.2 Steady-state fluid dynamics

Despite a lot of research efforts, many important partial differential equations can still not be solved by multigrid method with its full efficiency. One important example is the general Euler and Navier Stokes equations. The leading method has been based on multi-stage pseudo-time-stepping relaxation scheme [72, 71], but this type of algorithm is still very slow because some intermediate "characteristic components" can not adequately be reduced by coarse grids. The same problem also exists for other approaches. This is a subject that needs more extensive research efforts for more effective applications of multigrid methods.

While we would like to take the great challenge of getting multigrid method to work with its full efficiency for general Navier Stokes equations, as the first step, the relevant full potential equation will be our main focus of attention. Actually this equation possesses many major difficults in the more general case. Once we have an effective nonlinear multigrid solver for the full potential equation, we can then incorporate it as one of the relaxations steps in an outer multigrid solver for the entire Euler system and finally for the Navier-Stokes equations.

The full potential equation can actually be solved effectively by multigrid method in both the deep subsonic (uniformly elliptic) and the deep supersonic (uniform hyperbolic) case. the most difficult situation is near sonic regime (say, the Mach number is between .7 and 1.3). To attack this problem, some work has been done for linearized equation. For example, in [24, 23, 46, 47]), efficient multigrid methods have been developed based on a picewise semi-coarsening and rules for adding artificial dissipation at the coarser levels. The next task ahead is to extend the solver from these linear models to a general solver for the nonlinear full-potential operator in the entire transonic regime. In this investigation, we intend to explore the use of local mesh refinement techniques and more careful discretization schemes near the regions where the equation type changes.

4 Multilevel methods in computational materials science

4.1 Introduction to computational materials science

On the microscopic, quantum level, the properties of electrons and nuclei are governed by the Schrödinger equation. This equation is particularly difficult to solve, due to the fermion (half-integer spin) nature of the electrons: the many-electron wave function must be antisymmetric under any permutation of the electronic coordinates. Fortunately, for many problems in condensed matter physics it is sufficient to consider only the *ground state* of the material. In this case, *density functional theory*, whose development led to half of the 1998 Nobel prize in chemistry, has shown that focusing only on the electron density as opposed to each individual wave function leads to much simpler equations, which can be solved for hundreds and potentially thousands of atoms. Furthermore, the density functional equations are much easier to adapt to multiscale formalism than the many-electron Schrödinger equation.

There have been substantial developments in the methodology of computations as well. *Pseudopotentials* allow for the calculations to be restricted to only valence electrons, by simulating the nuclei and the chemically inactive core electrons with *effective core* potentials. Although not strictly necessary from the computational point of view, pseudopotentials significantly simplify the programming effort and often lead to substantial savings in computer time as well. Methods of quantum molecular dynamics, in which atoms move according to the Newton's equations but the forces are computed from density functional theory, have been developed and applied to a variety of systems. Finally, the early difficulties in using grids and grid-based methods in electronic structure calculations have been recently overcome, enabling applications to systems containing hundreds of atoms. This is important, because multiscale approaches are most easily realized in a grid-based approach, where it is possible to rigorously derive the equations appropriate to each level of resolution.

Significant progress has also been made in the development of simpler methods, such as classical molecular dynamics, where 2-, 3-and 4-body classical interatomic potentials provide a sufficiently accurate description of molecular interactions to enable studies of large-scale phenomena, and tight-binding methods, where the quantum hamiltonian is parameterized. These methods enable the studies of much larger systems and longer time scales, but are of course not as accurate or reliable as the *ab initio* or first-principles approaches. One of the main challenges of multilevel approaches to materials science is to develop appropriate coarsening procedures that would combine a multilevel description of the quantum differential equations with simpler parameterized approaches, and ultimately with classical and continuum physics.

4.2 Objectives and expected significance

The present proposal aims at developing new, vastly more efficient computational methods for multilevel simulations of materials properties. At the most advanced, *ab initio* level, the quantum solutions do not require any empirical input, but are based solely on the solutions of Schrödinger equation. They have true predictive power and are applicable to any material or chemical system, such as semiconductors, metals, chemical reactions, and biomolecules. However, large computational resources are required, and current applications are usually limited to several hundred "independent" atoms. (Calculations in solid state physics use periodic boundary conditions; therefore, the proper measure of computational complexity is the number of independent atoms in the periodic cell.) The methodological advances described in the proposal would enable calculations for much larger systems, significantly reduce overhead associated with treating "difficult" atoms, surfaces, and clusters, and seamlessly couple to simpler methods, thereby providing a natural framework for handling the multitude of length- and time-scales associated with complex physical systems.

There are many areas of physics, chemistry, materials science and molecular biology that would be impacted by the successful completion of the proposed work. In all of these fields, there is an increasing emphasis on understanding the relevant phenomena on a microscopic, quantum-mechanical level, while also predicting the mesoscopic behavior. For example, the continuing shrinkage of electronic devices places increasingly stringent requirements on "perfection" at the atomic scale, but nanodevices will likely include several thousand atoms. For new materials, improvements in growth and doping methodologies are a major concern, since the quality of bulk material is not sufficient for advanced applications. In chemistry and biochemistry, understanding reaction mechanisms and predicting complex structures would have a major impact on devising new industrial processes, unraveling the mechanisms of biological transformations and ultimately stimulating the development of new drugs.

The proposed research would result in scalable algorithms for multilevel methods, that would be portable across different parallel platforms, including workstation clusters. (The current algorithms, developed specifically for the high-bandwidth Cray T3D and T3E, do not port well to other architectures.)

4.3 Adaptable multilevel discretization

In the density functional method, one solves an effective Schrödinger equation (in Rydbergs):

$$\left[-\nabla^2 + V_N(\mathbf{r}) + \int \frac{2\rho(\mathbf{r}')d\mathbf{r}'}{|\mathbf{r}-\mathbf{r}'|} + \mu_{xc}(\mathbf{r})\right]\psi_i(\mathbf{r}) = \epsilon_i\psi_i(\mathbf{r}), i = 1, \dots N,$$

for the N electrons in the system. The first term represents the kinetic energy, the second the electrostatic potential due to the nuclei, the third the classical electron-electron repulsion potential, and the fourth the so-called exchange and correlation potential, describing the average effects of exchange and correlation derived from the theory of the electron gas. The last term depends only on the electron density $\rho(\mathbf{r}) = \sum |\psi_i(\mathbf{r})|^2$ and its gradient. For reviews of the applications of density functional methods see [88, 90].

The real-space discretization (unlike a plane-wave approach) allows a very substantial saving in the number of degrees of freedom needed to attain a given accuracy, by employing various forms of *local adaptation*. Each of these forms can be employed both in the main discretization of the density functional equations and in the discretization of the auxiliary Poisson equations needed to calculate the effective potential. A basic form of local discretization adaptation is the *grid adaptation*: employing fine grids in subdomains where the solution is rough, and coarse grids where it is smooth. In particular, this will lead to large savings in the regions far from nuclei, e.g., in studies of molecules or clusters.

Another form of local discretization adaptation exploits the known shape of a solution singularity or near-singularity to subtract it out. For example, the fact that, in the neighborhood of a nucleus, a molecular orbital ψ approximately assumes the shape of a corresponding atomic orbital ψ_0 and can be used as follows. Instead of ψ, one discretizes on the grid the smoother function $\psi_1 = \psi - \alpha\varphi\psi_0$, where φ is a smooth localizer (e.g., a Gaussian with a width of several meshsizes centered at the nucleus) and α is an *unknown* coefficient. The (linear) dependence of α on local grid values of ψ_1, and the resulting modification to the local discrete equations can be calculated once for all. With such proper choice of α, the function ψ_1 is much smoother than ψ, hence the meshsize used for its discretization can be much larger. Moreover, ψ_1 coincides with ψ several meshsizes away from the nucleus, thus retaining there the simpler discretization and avoiding explicit global interactions of different atomic orbitals.

4.4 The multigrid solver

For a given effective potential, a discretized density functional eigenfunction can be calculated by a multigrid solver in $O(N_g)$ computer operations, where N_g is the number of degrees of freedom (e.g., the number of gridpoints). A discretized Poisson equation, needed to calculate the Hartree potential and some forms of the exchange correlation potentials, can be solved to the required accuracy by a multigrid solver in about $40N_g$ computer operations. Essentially the same efficiency can be attained even when local adaptations are used.

This efficiency represents a modest advantage compared with plane-wave approaches, where the cost of a solver is at least $O(N_g \log N_g)$, assuming that the number of (conjugate gradient, for example) iterations is bounded independently of N_g. However, the grid operations are inherently local, facilitating effective parallelization through domain decomposition, while plane wave approaches must rely on FFTs, which require global operations. Moreover, plane wave methods cannot use local adaptations.

The density functional equations are inherently nonlinear, since both the electron-electron electrostatic potential V_{el} and the exchange-correlation potential depend on the computed eigenfunctions $\psi_i(r)$. This nonlinearity requires a number of applications of the linear solver, which can be quite large even when various acceleration schemes, such as conjugate gradients, are used. Although multigrid methods provide automatic acceleration and preconditioning at all length scales, the application of a linear solver still requires 10-15 iterations for the high level of convergence required (7 digits or more). However, a nonlinear "full multigrid" solver has been developed in other contexts. Its FAS version, can directly solve nonlinear problems as efficiently as solving their linear counterparts (see, e.g., [16, 17]). Such an approach can handle in an efficient way the problems related to the so-called "charge sloshing" (i.e., oscillations in electronic density appearing during self-consistent iterations) and can potentially substantially reduce the total number of iterations.

A one-shot solver for the density-functional equations will thus be developed. It is, in fact, quite straightforward: Using FAS multigrid, all it means is that the Poisson equation for V_{el} is relaxed, level by level, along with the relaxation of the Kohn-Sham eigen-equations. (Technically, it is easy to show that, at any given scale, in the extended PDE system comprised of the N_e eigen-equations plus the Poisson equation $-\Delta V_{el} = \rho$, the dependence of V_{el} on each ψ_i is not "principal", hence its enforcement in the relaxation processes can be retarded.) The total cost of this one-shot solver is only a fraction more than that of a multigrid solver for the problem with a *fixed* potential.

When a large number N_0 of orbitals is solved for, the total work becomes dominated by the additional $O(N_0^2 N_g)$ operations needed for the orthogonalization step. Furthermore, the FAS algorithm requires the solutions of the full density-functional equations on coarse grids, where orthogonalization becomes unstable for large systems. For this reason, the only application of FAS in this context has studied few-atom systems. We will investigate novel algorithms that would drastically reduce or eliminate the need for orthogonalization through hierarchical organization of orbitals and modifications of the original equations on coarse levels. Furthermore, we will implement additional acceleration by *recombining iterants*,which is what the conjugate-gradient method does. Since multigrid preconditioners are much better than those that can be used by the plane-wave approach, only few iterants will need to be recombined, diminishing the motivation to organize the recombinations in a conjugate-gradient-like procedure. Most importantly, using the FAS multigrid version, these recombinations can often be done *at coarse levels* of the solvers, which will cost only little in terms of memory and computing time.

As mentioned above, the multigrid solvers, being inherently local at each scale, can efficiently use a very high degree of parallel processing. This is true both for the main density-functional solver and for the auxiliary Poisson solvers employed in calculating the potentials [29]. Moreover, the relaxation processes of these different solvers can in fact run in parallel to each other.

4.5 Accelerated quantum molecular dynamics

Whereas in solving the electronic eigenproblem the nuclei are assumed fixed (the Born-Oppenheimer approximation), one actually needs to find the nuclear positions for which the electronic energy together with the inter-nuclear potential yield the minimal total energy. This external optimization would normally be done iteratively, requiring solving the electronic eigenproblem many times. Again, a one-shot multigrid solver + optimizer can and should be developed, incorporating suitable nuclei moves at each of the levels of the electronic solver. A model study reported in [19] has shown the feasibility of this approach and the exact multigrid techniques required for its full efficiency. (The main device is to calculate, along with each fine-to-coarse defect correction, also the coefficients of its linear dependence on the coordinates of neighboring nuclei, if any.)

The same multigrid techniques can also be used in *ab initio* molecular-dynamics calculations, where *implicit* nuclear positioning (using in the electronic equations the nuclear locations at the *end* of the time step) may be desired, in order to stabilize large time steps.

An interesting question that we plan to explore is the relation between the equations formed at the coarsest levels of such nuclear-positioning multigrid solver (due to the recursive fine-to-coarse defect corrections) and the inter-atomic (semi-classical) Hamiltonian that one would want to construct for the studied material. The coarse level description will aid in constructing a semi-classical parameterization that could be used in some parts of a large system. Conversely, the simpler molecular dynamics and tight-binding methods can be used in a "multilevel" fashion to estimate the initial atomic positions and propagate the atoms along a trajectory in-between the expensive quantum-mechanical steps. (The latter algorithm requires an efficient solver, such as described above, otherwise the additional iterations to converge the wave functions at the new positions will negate any gains.)

4.6 Classical molecular dynamics

For simulating very large molecular systems, a simplified form for the total potential energy E is derived, depending only on the n atomic positions $r = (r_1, r_2, \ldots, r_n)$, the electrons being implicit. Explicit approximations to $E(r)$ as a sum of various few-atom interactions are derived by computational experience compared with finer-scale quantum calculations such as described above, and with molecular measurement data (crystal structure geometries, vibrational spectroscopy, heats of formation, etc.). The most common few-atom interactions are of the following two kinds: (1) The bond interactions between chemically-bonded atoms, including three types: length (distance) interaction between 2 atoms, angle interaction between 3 atoms and torsion interaction between 4 atoms. (2) Non-bond interactions, including the short-range Lennard-Jones and hydrogen-bond terms and the long-range Coulombic potential. The aim of the calculations is usually either *dynamics* (calculating trajectories $r(t)$ which satisfy Newton's law $-\nabla E(r) = M\ddot{r}$, where M is the diagonal matrix of masses), or *equilibrium statistics* (average properties under the probability distribution $P(r) \sim \exp(-E(r)/k_B T)$), where k_B is the Boltzmann constant and T is the absolute temperature). The computing cost of current algorithms rises very steeply with problem size, restricting the modeling efforts to relatively small molecular ensembles and to time intervals many orders of magnitude smaller than needed. This steep rise in cost can be radically reduced by combining several types of multiscale approaches, such as the following. 1. *Fast summation of electrostatic forces.* A linear-complexity method described in [18] is based on decomposing each electrostatic interaction into the sum of a local interaction and a smooth interaction, and the summation of all the smooth interactions through coarser grids. The advantage over other linear-complexity approaches (e.g., the multiple multipole method) is in yielding the kind of multiscale description of force fields needed for the efficient multiscaling of dynamics and equilibrium calculations, where the charges are *moving*. 2. *Multiscale Monte-Carlo methods.* To calculate equilibrium statistics, an atom-by-atom Monte-Carlo process is usually performed. This process is usually extremely long, because of two basic reasons. (1) Due to the local nature of the Monte-Carlo process, only very slowly it affects large-scale conformational features, hence extremely many Monte-Carlo sweeps are needed to produce each new, statistically independent configuration. (2) Many such independent samples are needed to average out the deviation observed at each of them. For some simple lattice problems, *renormalization multigrid* (RMG) Monte-Carlo algorithms were developed which overcome *both* these complexity causes, by constructing a sequence of increasingly coarser descriptions of the simulated system. The Hamiltonian of each coarser level is derived by extensive, but only local, Monte Carlo simulations, iteratively adjusting it so as to fit coarse-level with fine-level simulations. Such iterations need be done only once for all similarly-structured molecular neighborhoods. The coarse levels thus constructed are then used both for accelerating the finest-level Monte Carlo, and for inexpensive calculations of large-scale statistics [27, 26]. Based on the same approach, the development of the following two types of multiscale Monte-Carlo is proposed. (i) *Multiscale Monte-Carlo for macromolecules.* The variables of each coarser level here are virtual atoms. Each such "atom" stands for the average location of several next-finer-level atoms. Preliminary tests show that the resulting coarse Hamiltonian, with vastly-reduced number of degrees of freedom and allowing much larger simulation steps, can very accurately reproduce large-scale statistics. (ii) *Multiscale Monte Carlo for fluids.* The efficient equilibrium simulation of gases and liquids at the atomic level, needed for the derivation of their large-scale behavior and macroscopic equations, is a central problem in scientific computation. Direct fine-scale MC simulations tend to be extremely inefficient due to the very slow change of various kinds of *clusters*, such as positional clusters and dipolar-alignment cluster, typical to water and other fluids at various scales. Here, coarse levels are constructed which are no longer atomistic; they are defined each on a uniform *lattice*. The values defined at each lattice point stand for average fluid properties in a lattice box around that point. The probability distribution of these values depend on similar values at neighboring-point, as specified in detail by a suitable "*conditional-probability (CP) table*". The CP table for each level is calculated by extensive, but only local, simulations at the next finer level. Our research objective is to develop these approaches and demonstrate their linear complexity in dynamics and equilibrium calculations. Moreover, the long term aim is to blend statistical approaches in the small with deterministic dynamics in the large, and to derive macroscopic equations at increasingly larger scales, leading eventually to very efficient continuum-level processing.

5 Multilevel Methods in Image Processing and Data Mining

Images are very efficient carriers of information, especially for human perception, and thus image processing (IP) is an fundamental problem domain for ITR. IP tasks range from image denoising and deblurring, to compression, segmentation and object identification. Each is computationally intensive, because even a static image contains a large number of pixels, and most tasks involve global interaction among pixels in some fashion, thus making IP an ideal domain to apply multilevel methods.

An important problem in data mining is data clustering. By drawing the obvious analogy between image segmentation and data clustering, there exists the possibility of developing efficient multilevel methods for data clustering based on ideas and concepts developed for multilevel methods for image processing and other traditional PDE applications.

In some sense, traditionally multilevel methods have already played an *implicit* role in IP, ranging from the use of FFTs and wavelet transforms to the development of super-resolution algorithms for obtaining high resolution images from several lower resolution ones. However, multilevel methodologies have traditionally not played an *explicit* role in IP.

We believe this situation will change in the near future because of recent developments in variational and PDE based IP [6, 84, 93, 89, 5, 1]. The PDEs are designed to possess certain desirable geometrical properties and part of the motivation is a more systematic approach to restoring images with sharp edges, as well as for image segmentation and active contours. Many of these new image models are variational in nature — the restored image, segmentation or active contours are characterized as minimizers of appropriately defined "energy" functionals. The Euler-Lagrange equations give rise to nonlinear PDEs but often it is important to exploit the objective function itself in both the modelling and the computational procedures.

The variational and PDE formulations allow multilevel methodologies to be applied to IP tasks in a natural and efficient way. The underlying continuous objective function and PDEs lend themselves naturally to a multiscale approach and one can borrow from the vast knowledge of multilevel methodology for traditional PDE problems. On the other hand, the PDE formulations call for new computational techniques which are different from the traditional frequency domain and algebraic approaches. Among the computational difficulties are the highly nonlinear and singular nature of the PDEs, and the presence of discontinuities (edges) in the solution. These special features of IP call for special adaption of standard multilevel techniques. There is also the need to develop new computational tools (such as the level set method described below) to properly handle the intrinsic geometric properties captured by these PDE models. We are proposing to develop these techniques.

In what follows, we outline a few specific IP problems for which we think multilevel methods can play a significant role.

5.1 Multilevel Methods for Differential-Convolution Problems in Image Restoration

In image restoration problems which involves deblurring, one often encounters the need to solve systems where the coefficient matrix corresponds to a sum of a differential and an integral operator. The latter corresponds to the blur (often a convolution) and the former a regularization (often a linearization of a highly nonlinear operator, often referred to as an anisotropic diffusion operator [89]). In our earlier work [33], we have carried out a preliminary study of adapting multigrid methods for this kind of problems. The usual multigrid algorithms do not work well for this class of problems because the small eigenvalues of the blurring operator have corresponding oscillatory, rather than smooth, eigenvectors. This is completely opposite to the situation for the more elliptic regularization operator, and renders the spectral behaviour of the combined operator difficult to handle for standard multigrid algorithms.

Several techniques have been used in the literature for these problems. One is the product-preconditioner approach of Vogel [94] which approximates the inverse of the sum by a product of inverses of shifted versions of the parts, making use of efficient preconditioners for the individual operators; e.g. multigrid for the elliptic one and FFT for the blur. Another approach is the cosine transform preconditioner of Chan, Chan and Wong [34], which computes the optimal preconditioner diagonizable by the cosine transform (corresponding to the natural Neumann boundary condition). A full two level additive Schwarz method has also been proposed by Hanke and Vogel [64]. While all of the above approaches have achieve some success, none is completely satisfactory.

A more fundamental approach which we would like to investigate is a fully nonlinear multilevel optimization procedure for minimizing the variational objective function directly. Ideas from existing multilevel optimization techniques [21] can be borrowed but again special features of the IP problem will require development of new techniques. We have already mentioned the difficulties for finding efficient multigrid smoothers for differential-convolution operators. In addition, the existence of many discontinuities in the solution will require special interpolation procedures to preserve their location and to avoid spurious oscillations (related to the famous Gibb's phenomenon). The coarsening procedure also needs careful study. We believe that the highly nonlinear and irregular coefficients of the elliptic operators will require a homogenization type approach for obtaining coarse scale representations. We will investigate whether recently developed techniques, such as those developed by Brandt [20], can be applied to this problem.

5.2 Multiscale Image Segmentation

A basic task in pattern recognition is the decomposition of a given image into meaningful segments. The criteria for blocking two picture elements into the same segment include similarity in color, intensity, absence of separating edges, etc. Quantitatively, these can be expressed in terms of *coupling* coefficients between neighboring pixels. Still, it is not uniquely defined how to derive the segments from the coupling coefficients. Multiscale approaches can play several fundamental roles in this task. We shall describe an approach first proposed in the recent work of Sharon, Brandt and Basri[91].

First, by identifying the coupling coefficients as conductances of an electrical network connecting the pixels, one can use algebraic multigrid coarsening techniques [25] to derive a coarse approximating network. If the coupling between a coarse pixel and its neighbors are weak, then the fine pixels corresponding to this coarse pixel can be identified as an image segment. One can view this procedure as an *iterated weighted aggregation* method. The use of weights is an important component of this approach and allow fractions of the same pixel to be associated with different coarser scale pixels, expressing the *likelihood* of pixels to belong together. These likelihoods can accumulate onto coaser levels, automatically reinforcing each other when appropriate. This weighted approach is fundamentally different from most segmentation algorithms which locally decides where two pixels (or segments) should merge or not, such as the well-known Mumford-Shah approach [85] which balances between strong coupling within a segment with the size of the boundary of a segment.

This algorithm in fact yields a *hierarchical and scaled segmentation*, in which smaller segments can be identified within segments, and scales of segments can be prescribed. Moreover, coupling between coarser pixels need not be dictated solely by fine scale couplings, but instead can use additional internal statistics of the emerging segments, such as center of mass, diameter, orientation, texture, etc., and use these statistics to introduce new important segmentation quantities, which are invisible at the fine scale. As a result, such multiscale methods in image processing allow both fast computations (with time linear in the total number of fine scale pixels) and the definition of large-scale segmentation determinants which are invisible to purely local processes.

We propose to enhance this multiscale segmentation algorithm in several ways: adding more coarse-level statistics, incorporating coarse edge information by making use of multiscale edge detection processes, the possbility of incorporating Mumford-Shah-like functionals at coarser levels to define the coupling coefficients (rather than merging of segments), etc.

5.3 Multilevel Methods for Level Set Based Active Contours

The basic idea in active contour models or snakes is to evolve a curve, subject to constraints from a given image, in order to detect objects in that image. For instance, starting with a curve around the object to be detected, the curve moves in the direction of its interior normal, and has to stop on the boundary of the object.

We have recently developed a new active contour model which improves on existing models [36, 39]. There are two key new features of our model. First, it does not require the object to be detected to have sharp edges. Blurred objects, in fact even objects without edges (e.g. a cluster of points), can be detected. Second, by using a level set [87] implementation, our active contour can detect multiple objects starting from one curve. We are proposing here to develop an efficient multilevel version of this active contour model.

Our model is in fact a combination between classical active contour models [73, 31, 32] and the Mumford-Shah functional for segmentation [85]. We minimize the sum of a regularizing term for the curve C (the length) and the differences in norm between the given image I and the averages of I inside and outside the curve C. Basically, this second term is zero (therefore the minimizer) if the curve is on the boundary of the object, and it is positive if the curve is outside or inside. More precisely, our model is:

$$\inf_{C,c_1,c_2} J(C,c_1,c_2) = \mu \text{length}(C) + \lambda \int_{inside\ C} |I - c_1|^2 + \lambda \int_{outside\ C} |I - c_2|^2,$$

where c_1 and c_2 are constants (representing the averages of I inside C and outside C).

To implement our model, we use the level set formulation of Osher and Sethian [87, 109]. The level set function $\phi : I\!R^2 \rightarrow I\!R$ is a Lipschitz function and the evolving curve will be represented by the zero-level set of ϕ. The level set formulation of our model is:

$$\inf_{\phi, c_1, c_2} F(\phi, c_1, c_2) = \mu \text{length}\{\phi = 0\} + \lambda \Big(\int_{\phi \geq 0} |I - c_1|^2 + \int_{\phi < 0} |I - c_2|^2 \Big),$$

where μ, λ are positive parameters. From here, it is not difficult to derive the Euler-Lagrange equations and the associated evolution problem for ϕ, parameterizing the descent direction by time:

$$\begin{cases} \frac{\partial \phi}{\partial t} = \delta(\phi) \Big[\mu \text{div} \Big(\frac{\nabla \phi}{|\nabla \phi|} \Big) - \lambda \Big((I - c_1)^2 - (I - c_2)^2 \Big) \Big] \text{ in } \Omega, \\ \phi(t, x, y) = \phi_0(x, y) \text{ in } \Omega, \quad \frac{\delta(\phi)}{|\nabla \phi|} \frac{\partial \phi}{\partial n} = 0 \text{ on } \partial \Omega. \end{cases}$$

An example of what our model can do is shown in Fig. 1, where clusters of points, including those "shielded" by other clusters, are identified.

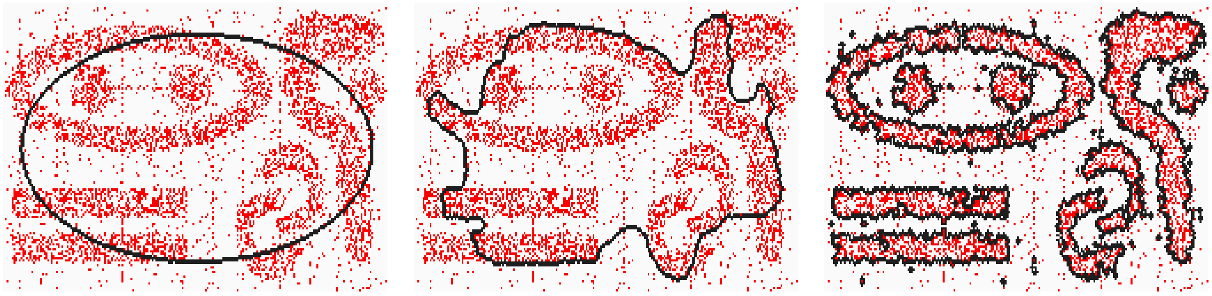

Figure 1: Clusters discovered by Chan-Vese level set based active contour contour model. Example from George Karypis, Eui-Hong Han, Vipin Kumar, "Chameleon: Hierarchical Clustering Using Dynamic Modeling", Computer, August 99, pages 68-75.

The explicit time marching scheme above can be slow. We propose to speed up the procedure by multilevel techniques. Generally, applying multilevel techniques to curve evolution is not so straightforward because the curve is usually parametrized as a one dimensional object in 2D. However, by using a level set formulation, this is made much easier. Essentially, the coarsening of the curve can be achieved *implicitly* by coarsening the 2D level set function, which is smooth by construction. The given image also has to be coarsened. Instead of standard coarsening based on centered averaging (i.e. down sampling), one should use a strategy which preserves the edge locations as much as possible. The coarse scale model is then trivially given by the same variational formulation, but with a coarser image I_c. The time marching scheme can still be used as a smoother. We believe that this multilevel approach can speed up the active contour procedure significantly, especially for the detection of "blocky" images which can be well-represented on the coarse scale. The challenge is how to handle the coarsening when small objects are present on the fine scale which cannot be represented on the coarse scale. For classical multigrid methods, this has been studied in some details by leading researchers such as Hackbusch and Yserentant and some of their techniques may be applicable here. We believe that the level set formulation may actually obviate the need for these more sophisticated method as a level-set representation of a contour can easily and smoothly go through topological changes, thus picking up the fine scale only objects easily.

5.4 Multilevel Method for Clustering Problems in Data Mining

Automatic data clustering is a central problem in data mining [70, 13]. The task is to group a set of data to maximize the similarity within the clusters and minimize the similarity between different clusters. This is a very active research area and there have been many algorithms proposed. The most well-known of these are the "k-median" and "k-mean" algorithms which iteratively determines the centers of the clusters by averaging over the current data points associated with each cluster, followed by recomputing the "nearest" data points associated with the new cluster centers. Multilevel methodology have been successfully applied to data clustering problems. For example, the **Chameleon** algorithm by Karypis et al [52] uses a hierarchical graph based algorithm which dynamically makes clustering decisions based on both interconnectivity and proximity information about the data.

Here we would like to exploit the obvious connection between image segmentation and clustering and propose several novel approaches which apply ideas from the former area (many discussed earlier in this section) to problems from the latter.

One direction is to explore the use of the aforementoned weighted hierarchical aggregration method of Sharon-Brandt-Basri (SBB) within the context of a Chameleon-like algorithm. The potential advantage is that the use of weights will allow detection of clusters which cannot be seen on finer levels. The basic data structure used in Chameleon is a graph and the algebraic multigrid based coarsening procedures used in SBB can be easily (and has been) extended to such structures. Moreover, in addition to interconnectivity and proximity, other problem-specific features and statistics of emerging clusters can be incorporated at coarser levels.

Another direction, admittedly more speculative, is to extend the level set based active contour models described earlier to data clustering. Clearly, from the results presented in Fig. 1 and the corresponding results in [52], our active contour methods can already find clusters *in low dimensions* as well as any existing clustering algorithms. Thus this algorithm can already be applied, at least in principle, to applications with low dimensional data. The fundamental barrier, however, is the extension to higher dimensions. We'd like to have an algorithm that runs at time linear in the number of data points. Unfortunately, the level set method, at least as implemented for images, requires at each iteration a pass over an underlying grid, which would be prohibitively expensive for higher dimensions. In fact, what we need to develop is a *gridless* level set methods. Meeting this challenge won't be easy but at least it seems possible. Ideas such as the use of gridless approximation schemes such as radial basis functions come to mind. But we believe a more fruitful attach would be to borrow the concepts, but not the algorithms, from the image active contours and set up an analogous model for clustering. For example, the Chan-Vese active contour basically defines an object in an image based on how well the object intensity can be approximated by a constant. An analogous definition for a cluster would be a sort of uniformity of data distribution within the cluster. Another example is the level set function value at a pixel basically denotes the proximity to the object boundary. Can such a concept be defined for data clusters? We certainly think so.

6 Performance Modeling and Improvement of Multilevel Codes

Multilevel codes, like other codes designed to solve sparse problems in production environments at national laboratories, typically achieve less than 10% of processor peak performance. While such figures are never proudly published, 1% of processor peak is probably not rare. (Careful studies of our own of a production NASA unstructured grid CFD code on 14 different processor/memory system combinations for the major commercial scientific processors with vendor or Linux compilers showed a range of 1.3% to 8.3%. [?].) On the same systems, the LINPACK benchmark achieves up to 80% of peak processor rate.

This gap in performance between what is achieved on regular (dense) computations and on sparse irregular computations has been called the "sparse matrix gap" [?]. The multilevel algorithms that we design (or redesign) in this proposal will target a much higher standard of performance on modern microprocessor architectures, such as the 25% of the SGI Origin at which we made the above-mentioned NASA code execute after some cache-reuse and register-reuse data structure transformations and algorithmic reorderings [?].

According to SGI-Cray's webpage on the T3E-900 (http://www.sgi.com/t3e/performance.html), the average number of floating point operations required to cover an access to main memory is 252 — about 25 times worse than on the CDC 7600, a machine introduced three decades ago. Proponents of the Hybrid Technology Multithreaded (HTMT) petaflops computer (http://htmt.caltech.edu) estimate that main memory latency will grow to 10,000 processor cycles under the earliest versions of their design, and increase thereafter. Ironically, one encounters the same phenomenon at the "low end" of high performance computing — beowulf clusters — for remote memory latencies. Given this dwarfing of the cost of computation by the cost of data transfer, application programs of the future may be dominated by the expression of data motion rather than operations. This daunting task is made more approachable by the identification of data motion templates, such as those listed below.

For computations with data access patterns that are regular and periodic in time (as in, for example, an iterative PDE computation on a static grid, even if the grid is spatially irregular), these latencies do not by themselves pose a serious threat to the feasibility of aggressive exploitation of hierarchical distributed memory machines. Latency can be covered, given sufficient bandwidth and proper data structure organization, by block transfers, prefetching, multithreading, or some combination. These solutions presume concurrency well beyond the process level – which is easily available in PDEs – and high bandwidth. (For problems with unpredictable data-dependent data access patterns high latency is far more serious, and perhaps fatal.)

Even if a latency problem is amenable to treatment, the required bandwidth cannot be taken for granted. The IBM F50 node (four 604e processors sharing a local memory) used in an early version of the ASCI Blue Pacific machine has barely enough memory bandwidth to keep *one* of the four processors fed during a streaming computation such as a vector scaling — even with the scaling operand fixed in a register. Since multilevel PDE computations touch $\mathcal{O}(n)$ data in $\mathcal{O}(n)$ operations, it is clear that memory bandwidth is at least as critical an issue as latency in present and future architectures. We therefore do not want to waste what little bandwidth there is.

Petaflops-scale computers, for which optimal multilevel algorithms are more important than ever before – since the data sets will be Petawords – will be useful only if efforts are made to bridge the gap between architects and algorithm/code designers. To this end, barriers to performance must be understood, workarounds identified, and templates constructed that embody the workarounds for generic actions involving data motions. Designers must converge on a minimal set of generic actions whose requirements make good targets for architects and low level library writers.

A general-purpose unstructured multigrid-preconditioned implicit simulation of a PDE system, may involve the following:

- *vertex loop*: concurrent at point level, no communication (typically state equation updates or vector BLAS)
- *edge loop*: concurrent at colored edge level, near neighbor communication (typically stencil updates or flux evaluations)
- *recurrence*: concurrent at subdomain or wavefront level, near neighbor communication (typically relaxation sweeps or backsolves)
- *global reduction*: concurrent at point level, but synchronizing (typically used in orthogonalization inner products for convergence acceleration and for convergence norms)
- *multilevel transfer*: concurrent at point level, possibly involving search heuristics (typically interpolation or agglomeration)

We will study and publish discussions of the hierarchical workingset structure of our multilevel codes (moving successively outward from points, through stencils, grid patches, sets of related grid patches, full subdomains of data, etc.) and characterize the relative scaling of workingset size in terms of overall problem size and process granularity. The overall focus of our proposal is *not* on flop/s or architectural issues, but our algorithmic development will be influenced by these important considerations.

There are (at least) two important strategies for the reuse of data in registers, caches, and local memories (before overwriting and refreshing): reuse of blocks of matrix coefficients on multiple vectors [**?**] and reuse of grid vectors over multiple iteration steps [**?**] We will investigate both.

In multilevel methods, there are two extreme data structure paradigms from the multilevel viewpoint (in addition to paradigms on grid function interleaving and spatial ordering, which are becoming well understood for tensor product and unstructured grids, and hybrids of these single grid systems). One extreme stores each point just once (modulo replication of parallel ghost cells) meaining that the coarser grid elements are not contiguous in the data structures. Another extreme replicates storage for each level. This uses (a small multiple) greater memory and involves extra copying, but elements within a level are stored contiguously, for superior spatial locality. Again, we will investigate both.

It is important to emphasize that barriers to performance may vary from one code segment to another. The NASA CFD code had three different barriers active in different regimes. During local "flux balance" phases that occupied approximately half of the execution time, the highly optimized code was limited by the number of load and store units on most architectures. Parallel scalability was embarrassing and memory bandwidth was sufficient. During the sparse linear algebra phases associated with calculating nonlinear corrections to the solution, which occupied approximately the other half of the time, memory bandwidth was limiting. Parallel scalability was adequate in the "fixed-memory-per-node" limit. During the orthgonalization phase, which grew to occupy about 10% of execution time on thousands of processors, the global synchronizations became the bottleneck. It is important to note that there are trade-offs between these phases that can be exploited to shift the balance of work away from what is least supported by the hardware. For instance, one can choose accelerators that orthogonalize less per matrix-vector application, or one can tune tolerances to do more matrix-vector applications per flux evaluation.

In multilevel algorithms there are *many* such relevant trade-offs between phases, such as: construction of coarsened operators, mapping functions between different levels of representation, smoothing on a level, and recalculating fine level operators (in a nonlinear problem). Choices of coarsening ratio, number of levels, number of pre- and post-smoothing steps, types of grid transfers, and types of coarsenings are among those that shift the balance of work between these of phases, which stress the memory and communication systems in different ways. Following purely mathematical guidelines to minimize the arithmetic complexity of a multilevel method may violate the minimization of data motion complexity. Systematic and comprehensive mapping of these trade-offs would involve detailed hardware monitoring and complex memory system models, which are beyond the scope of this relatively small proposal. Our interest is in building simple models of the data motion requirements, which would enable us to choose between data structure alternatives and algorithmic tunings in a memory-aware manner, and thus enhance the relevance of the algorithmic accomplishments of this project.

7 Education and Human Resources

Our budget only include support for one postdoc and 2 graduate students, with very little else. Thus the main use of the proposed funding is to train young researchers in this important interdisciplinary area. It is particularly difficult to train students in this area because it requires in depth knowledge of a wide spectrum of knowledge from computational mathematics, computationa science, high performance computing and the application area. We believe the team approach that we have built is a stimulating environment to train young researchers in this area.

We have asked for travel money to enable the postdoc, and possibly also the graduate students, to visit the various team members to facilitate the research collaboration. We have also asked for a small amount of money to support an annual group workshop in which the postdoc and graduate students will participate, possibly also with outside participants.

The PI (Chan) runs a weekly research seminar on image processing in which 10-15 faculty (e.g. Osher), postdocs and students participate. Together with Bjorn Engquist, he ran a similar research seminar on multilevel methods for several years and this will be resumed this year. One of the PIs (Xu) participated actively in this seminar when he spent a one year sabbattical at UCLA several years ago. One of the student participants, Justin Wan, was supported by LLNL's ISCR (where Keyes is Director) during his final dissertation year. Stan Osher also runs a similar seminar on level set methods in which many of the same people also participate. Thus together we have at UCLA a lively research environment for the topics covered in this proposal.

Our research seminars also have a diverse body of participants. Of Chan's current PhD students and postdocs, 4 are women, one of whom African American. We'll continue to broaden opportunities enabling the participation of the widest body of participants in our research project.

8 International Special Information

We have put together an international research team, consisting of the 3 US based PI's (Bernholc, Chan and Xu), collaborators from four national laboratories (Tim Barth, NASA Ames; Rob Falgout, David Keyes, Carol Woodward, LLNL; Dana Knoll, LANL; David Sidlikover, NASA Langley), and two major centers of multilevel computation research (Michael Griebel, Group Griebel, Univ. Bonn, Germany; Achi Brandt, Gauss Center, Weizmann Institute, Israel).

8.1 Prof. Michael Griebel, Institut für Angewandte Mathematik, Division for Scientific Computing and Numerical Simulation, Univ. Bonn

Activities in Computational Material Science. In Bonn, the group of M. Griebel is working now for some years on computational methods for material sciene and nanotechnologie. The focus of interest is threefold: First, there is work on the quantum mechanical level. To this end the Schrödinger equation is solved for the hydrogen and helium atoms under strong magnetic and electric fields, see [53, 54], by means of the sparse grid combination technique. This approach combines the solution of the Schrödinger equation on various scales of discretizations via a sort of multivariate extrapolation method. Its work count scales like $O(N(\log N)^{d-1})$ where d denotes the dimensionality of the problem, i.e. $d = 3$ for hydrogen and $d = 6$ for helium. Then, we have experience in the parallelization of codes for ab initio calculations of matter, i.e. the Hartree-Fock method [74] and the density functional theory, see [86]. The main focus was here on metallic clusters so far. Presently, in two theses [63, 99], we are implementing a parallel multilevel generalization of the Car-Parinello method and of the DFT. Third, on the molecular level, there is ongoing research with various variants of MD algorithms. Here, we implemented a fully adaptive version of the fast multipole method (adptive in h and p) for general long range potentials. It runs in parallel on for example the Cray T3E and shows good scaleup properties up to 1200 processors [4, 81]. Furthermore, based on GLOBUS, we coupled two parallel computers (T3E 512 proc.) and Parnass2 (144 proc.) and run on this parallel metacomputer molecular dynamics simulations involving complicated bond-order potentials [3]. Here the short range parts of the potentials are treated on one machine and the long range parts on the other. Besides we developed a code which allows for the coupling of the equation for the electrons (DFT-approximation) and the evolution equation of the nuclei (MD) via the Hellman-Feynman theorem. Here the movement of the nuclei is computed by a MD calculation and the computation of the current electron density distribution and thus the actual potential energy is achieved by the DFT. Furthermore there is work on the numerical simulation of nanotubes based on carbon and boron nitride. We computed the configuration of the geometry of caps for nanotubes by DFT and did simulations for extremely long nanotubes with up to 100 million atoms with MD involving Brenners potential, see [2]. Actual research is focussed on multiscale time integration to overcome the barrier of about 500 femtoseconds which is persistent in many MD calculatios.

To this end our expertise in parallel multigrid methods for PDEs comes in handy. Here there is theoretical work on the convergence of additive and multiplicative methods [61, 55, 56]. Also the parallelization of multigrid methods was dealt with in [59, 57, 62]. Furthermore new adaptive multilevel discretizations based on sparse grids [58, 80, 60] and tensorproducts of wavelets [50] had been studied.

Details for the international collaboration: In 2001, we plan in Bonn a GAMM Seminar on Multilevel Computation in the Material Science. This will be a follow up of the 10th GAMM Multigrid Conference in Bonn from 1998. We hope to win the collaborators in this project as invited speakers. Furtermore, in collaboration with LLNL (D. Keyes) we started an intensive exchange of PHD students and postdocs. This year, there will be six researchers staying in Livermore for one to two month each. We hope to be able to extend this collaboration also to UCLA (T.Chan), Penn State (J. Xu) and NCSU (J. Bernholc).

8.2 Prof. Achi Brandt, Gauss Minerva Center for Scientific Computation, Weizmann Institute

The Carl F. Gauss Minerva Center for Scientific Computation (The Gauss Center) resides in the Faculty of Mathematics and Computer Science in the Weizmann Institute of Science, Rehovot, Israel. Its objective is to act as a catalyst for the development of new fundamental computational approaches in physics, chemistry, applied mathematics and engineering, introducing, in particular, advanced multiscale (multilevel, multiresolution) methods.

The work at the Gauss Center treats multilevel computation as 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, interacting with many fields of application. The multiscale methodology is developed in a very systematic way, through sequences of model problems in which each type of computational difficulty is first isolated and studied in separation from others; then the learned techniques are combined to treat increasingly more complicated problems.

Current projects at the Gauss Center include multigrid solvers for partial differential equations (fluid dynamics, Dirac equations, quantum chemistry and others), inverse and optimal control problems, renormalization multigrid methods in statistical physics (including QCD calculations of elementary particles), fast integral transforms and integro-differential-equation solvers, multilevel molecular mechanics, advanced algorithms of medical imaging and multiscale approaches to image processing (edge detection, segmentation, picture coarsening, etc.).

The collaboration with the Gauss Center will thus benefit all parts of the current proposal. It will make it possible to draw upon the vast knowledge and experience accumulated at the Center, by mutual visits of senior staff, by sending our students for periods of study and collaboration at the Center, and by holding summer courses directed by Prof. Brandt.

For the Gauss Center this collaboration is an excellent opportunity to enhance its mission, through interactions with groups having similar objectives, learning and teaching new problems and new computational methods.

References

[1] Special issue on partial differential equations and geometry-driven diffusion in image processing and analysis. *IEEE Trans. in Image Processing*, 7, 1998.

[2] M. Griebel A. Caglar. On the numerical simulation of fullerene nanotubes: $c_{100.000.000}$ and beyond. In *Proceedings of NIC Workshop on Molecular Dynamics on Parallel Computers, J"ulich Jan. 1999.* World Scientific, 1999.

[3] M. Griebel A. Caglar and G. Zumbusch. Projekt Multiskalenmoleküldynamik im Gigbabit Testbed West. Bericht zum Gigabit Testbed West, Jülich, 2000.

[4] A. Caglar, M. Griebel, M.A. Schweitzer, G. Zumbusch. Dynamic load-balancing of hierarchical tree algorithms on a cluster of multiprocessor pcs and on the Cray T3E. In H. Fränznick H.W. Meuer, editor, *Proceedings 14th Supercomputer Conference, Mannheim 10.-12. Juni 1999.* Mateo Publisher, 1999. also as Report 27, SFB 256, Institut für Angewandte Mathematik, Universität Bonn, 1999.

[5] L. Alvarez, P.L. Lions, and J.M. Morel. Image selective smoothing and edge detection by nonlinear diffusion (ii). *SIAM J. Numer. Anal.*, 29:845–866, 1992.

[6] L. Alvarez and J. Morel. Formalization and computational aspects of image analysis. *Acta Numerica,*, pages 1–59, 1994.

[7] U. Andersson, B. Engquist, G. Ledfelt, and O. Runborg. Wavelet-based subgrid modeling: 1. Principles and scalar equations. Technical Report 97-13.

[8] M. Avellaneda and G. Milton. Optimal bounds on the effective bulk modulus of polycrystals. *SIAM J. Appl. Math.*, 49(3):824–837, 1989.

[9] I. Babuška, G. Caloz, and J. Osborn. Special finite element methods for a class of second order elliptic problems with rough coefficients. *SIAM J. Num. Anal.*, 31(4):945–981, 1994.

[10] R. E. Bank and J. Xu. An algorithm for coarsening unstructured meshes. *Numer. Math.*, 73:1–36, 1996.

[11] T. Barth, T.F. Chan, and W.P. Tang. A parallel non-overlapping domain decomposition algorithm for compressible fluid flow problems on triangulated domains. *AMS Contemporary Mathematics*, 218:23, 1998.

[12] J. Bernholc, J.-Y. Yi, and D.J. Sullivan. Structural transitions in metal clusters. *Faraday Society Discuss. Chem. Soc.*, 92:217, 1991.

[13] P.S. Bradley, U.M. Fayyad, and O.L. Mangasarian. Mathematical programming for data mining: Formulations and challenges. Technical Report MSR-TR-98-04, Univ. of Wisconsin, July 1998.

[14] J. H. Bramble, J. E. Pasciak, J. Wang, and J. Xu. Convergence estimates for multigrid algorithms without regularity assumptions. *Math. Comp.*, 57:23–45, 1991.

[15] J. H. Bramble, J. E. Pasciak, and J. Xu. The analysis of multigrid algorithms with nonnested spaces or noninherited quadratic forms. *Math. Comp.*, 56:1–34, 1991.

[16] A. Brandt. Multilevel adaptative solutions to boundary-value problems. *Math. Comp.*, 31(138):333, 1977.

[17] A. Brandt. Guide to multigrid development. In W. Hackbusch and U. Trottenberg, editors, *Multigrid Methods*, pages 220–312. Springer-Verlag, 1982.

[18] A. Brandt. Multilevel computations of integral transforms and particle interactions with oscillatory kernels. *Comput. Phys. Comm.*, 65:24–38, 1991.

[19] A. Brandt. The gauss center research in multiscale scientific computation. *Electronic Trans. Num. An.*, 6:1–34, 1997.

[20] A. Brandt. General highly accurate algebraic coarsening schemes. May 1999. Submitted to Elec. Trans. Num. Anal.

[21] A. Brandt. Multiscale scientific computation: six year research summary. *Tech Rpt, Gauss Minerva Center, Weizmann Inst., Israel*, June 1999.

[22] A. Brandt. Multiscale scientific computation: Six year research summary. Technical report, Weizmann Institute of Science, Rehovot, Israel, 1999.

[23] A. Brandt and B. Diskin. Multigrid solvers for non-aligned sonic flow. to appear in SIAM J. Sci. Comp., 1999.

[24] A. Brandt and B. Diskin. Multigrid solvers for non-aligned sonic flow: the constant coefficient case. *Computer & fluids*, 28:511–549, 1999.

[25] A. Brandt, S. McCormick, and J. Ruge. Algebraic multigrid (amg) for automatic multigrid soluton with application to geodetic computations. *Tech. Rpt., Inst. for Computational Studies, Fort Collins, CO*, 1982.

[26] A. Brandt and D. Ron. Renormalization multigrid (rmg): Statistically optimal renormalization group flow and coarse-to-fine monte carlo acceleration. *J. Stat. Phys., submitted.*

[27] A. Brandt and D. Ron. Statistically optimal renormalization group flow and coarse-to-fine monte carlo acceleration. *Gauss Center Report*, WI/GC-11:June, 1999.

[28] E.L. Briggs, D.J. Sullivan, and J. Bernholc. Large scale electronic structure calculations with multigrid acceleration. *Phys. Rev. B*, 52(8):5471–5474, 1995.

[29] E.L. Briggs, D.J. Sullivan, and J. Bernholc. Real-space multigrid-based approach to large-scale electronic structure calculations. *Phys. Rev. B*, 54(20):14362–14375, 1996.

[30] R. Car and M. Parrinello. Unified approach for molecular dynamics and density-functional theory. *Phys. Rev. Lett.*, 55(22):2471, 1985.

[31] V. Caselles, F. Catté, T. Coll, and F. Dibos. A geometric model for active contours in image processing. *Numerische Mathematik*, 66:1–31, 1993.

[32] V. Caselles, R. Kimmel, and G. Sapiro. On geodesic active contours. *Int. J. of Comp. Vision*, 22(1):61–79, 1997.

[33] R. Chan, T.F. Chan, and W.L. Wan. Multigrid for differential-convolution problems arising from image processing. *Scientific Computing, G. Golub, S. Liu, F. Luk and R. Plemmons (eds.), Springer, Singapore*, pages 58–72, December 1997.

[34] R. Chan, T.F. Chan, and C.K. Wong. Cosine transform based preconditioners for total variation deblurring. *IEEE Trans. Image Proc.*, 8:1472–1478, October 1999.

[35] T. Chan, J. Xu, and L. Zikatanov. A multigrid method on agglomerated coarse spaces. SIAM J. Sci. Comp., (submitted), 1999.

[36] T. F. Chan and L. Vese. Active contours without edges. 1998.

[37] T. F. Chan, J. Xu, and L. Zikatanov. An agglomeration multigrid method for unstructured grids. In *10-th international conference on Domain Decomposition methods*, volume 218 of *Contemporary Mathematics*, pages 67–81. American Mathematical Society, 1998.

[38] T.F. Chan, S. Go, and J. Zou. Multilevel spectral partitioning of unstructured grids. *Proceedings of 11th Domain Decomposition Conference, Greenwich, UK. DD Press*, July 1999.

[39] T.F. Chan, B. Sandberg, and L. Vese. Active contours without edges for vector-valued images. October 1999.

[40] Tony F. Chan. Hierarchical algorithms and architectures for parallel scientific computing. *Proceedings of the ACM Int'l Conf. Supercomputing, Amsterdam*, pages 318–329, June 1990.

[41] Tony F. Chan and W.L. Wan. Robust multigrid methods for elliptic linear systems. *UCLA Math CAM Rpt 99-30. To appear in State of the Art of Numerical Analysis, J. Comp. & App. Math.*, 2000.

[42] Tony F. Chan, J. Xu, and L. Zikatanov. An agglomeration multigrid method for unstructured grids. *AMS Contemp. Math.*, 218:67, 1998.

[43] Y. Chen and J. W. Evans. Thermal analysis of lithium-ion batteries. *J. Electrochem. Soc.*, 143:2708, 1996.

[44] L. Collatz. *The Numerical Treatment of Differential Equations.* Springer, Berlin, 1966.

[45] C. Collins, M. Luskin, and J. Riordan. Computational results for a two-dimensional model of crystalline microstructure. In *Microstructure and phase transition*, pages 51–56. Springer, New York, 1993.

[46] B. Diskin. Multigrid with conditional coarsening for the non-aligned sonic flow:. *Electron. Trans. Numer. Anal.*, 6, 1997.

[47] B. Diskin. *Efficient Multigrid Solvers for the Linearized Transonic Full Potential Equation.* PhD thesis, Feinberg Graduate School, Weizmann Institute of Science, June 1998.

[48] M. Dorobantu and B. Engquist. Wavelet-based numerical homogenization. Technical Report 96-3, UCLA, CAM report, February 1996.

[49] M. Doyle, J. Newman, A. S. Gozdz, C. N. Schmutz, and J.-M. Tarascon. Comparison of modeling predictions with experimental data from plastic lithium ion cells. *J. Electrochem. Soc.*, 143:1890, 1996.

[50] F. Koster, M. Griebel, N. Kevlahan, M. Farge, K. Schneider. Towards an adaptive wavelet-based 3D Navier-Stokes solver. In E.H. Hirschel, editor, *Numerical flow simulation I, Notes on Numerical Fluid Mechanics, Vol. 66*, pages 339–364. Vieweg-Verlag, Braunschweig, 1998.

[51] T. F. Fuller, M. Doyle, and J. Newman. Simulation and optimization of the dual lithium ion insertion cell. *J. Electrochem. Soc.*, 141:1, 1994.

[52] E.H. Han G. Karypis and V. Kumar. Chameleon: Hierarchical clustering using dynamic modeling. *Computer*, pages 68–75, August 1999.

[53] J. Garcke. *Berechnung von Eigenwerten der stationären Schrödingergleichung mit der Kombinationstechnik.* thesis, Institut für Angewandte Mathematik, Universität Bonn, 1998.

[54] J. Garcke and M. Griebel. On the computation of the eigenproblems of hydrogen and helium in strong magnetic and electric fields with the sparse grid combination technique. J. Comp. Physics, submitted, 2000.

[55] M. Griebel. Multilevel algorithms considered as iterative methods on semidefinite systems. *SIAM Int. J. Sci. Stat. Comput.*, 15(3), 1994.

[56] M. Griebel. *Multilevelmethoden als Iterationsverfahren über Erzeugendensystemen, Teubner Skripten zur Numerik.* Teubner Verlag, Stuttgart, 1994.

[57] M. Griebel. Parallel domain-oriented multilevel methods. *SIAM Journal on Scientific Computing*, 16(5):1105–1125, 1995.

[58] M. Griebel. Adaptive sparse grid multilevel methods for elliptic PDEs based on finite differences. *Computing*, 61(2):151–180, 1998.

[59] M. Griebel and T. Neunhoeffer. Parallel point- and domain-oriented multilevel methods for elliptic PDE's on workstation networks. *J. Comp. Appl. Math.*, 66:267–268, 1996.

[60] M. Griebel and P. Oswald. On additive Schwarz preconditioners for sparse grid discretization. *Numer. Math.*, 66(4):449–464, 1994.

[61] M. Griebel and P. Oswald. On the abstract theory of additive and multiplicative Schwarz algorithms. *Numer. Math.*, 70:163–180, 1995.

[62] M. Griebel and G. Zumbusch. Parallel multigrid in an adaptive PDE solver based on hashing and space-filling curves. *Parallel Computing*, 25:827–843, 1999.

[63] J. Hamaekers. *Numerische Umsetzung und Parallelisierung des Verfahrens von Car-Parinello.* thesis, Institut für Angewandte Mathematik, Universität Bonn, 2000.

[64] M. Hanke and C.R. Vogel. Two-level preconditioners for regularized inverse problems i: Theory. *Numerische Mathematik*, 1999.

[65] J.-S. Hong, H. Maleki, S. Al Hallaj, L. Redey, and J.R. Selman. Electrochemical-calorimetric studies of lithium-ion cells. *J. Electrochem. Soc.*, 145:1489, 1998.

[66] T. Y. Hou and X.-H. Wu. A multiscale finite element method for elliptic problems in composite materials and porous media. *J. Comp. Physics*, 134:168–189, 1997.

[67] T. Y. Hou, X.-H. Wu, and Z. Cai. Convergence of a multiscale finite element method for elliptic problems with rapidly oscillating coefficients. *Math. Comp.*, 68(227):913–943, 1999.

[68] Y. Huang and J. Xu. A partition-of-unity approach to non-matching grid finite element methods. In preparation and near completion, 1999.

[69] Y. Huang and J. Xu. A partition-of-unity finite element method for elliptic problems with highly oscillating coefficients. In *Proceedings for the Workshop on Scientific Computing 99*, Hong Kong, June 1999.

[70] A. K. Jain and R. C. Dubes. *Algorithms for Clustering Data.* Prentice-Hall, Inc. Englewood Cliffs, NJ, 1988.

[71] T. Jameson. Solution of the Euler equations by a multigrid method. *Appl. Math. Comput.*, 13:327–356, 1983.

[72] T. Jameson, W. Schmidt, and E. Turkel. Numerical solution of the Euler equations by finite volume method using Runge-Kutta, time stepping schemes. *AIAA*, 1295, 1981.

[73] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *Int. J. of Comp. Vision*, 1988.

[74] T. Kelterbach. *Zum Hartree-Fock-Verfahren und seiner numerischen Umsetzung.* thesis, Institut für Angewandte Mathematik, Universität Bonn, 1998.

[75] H. Kim, J. Xu, and L. Zikatanov. A multigrid method based on graph matching for convection diffusion equations. SIAM J. Sci. Comp., (submitted), 1999.

[76] H. Kim, J. Xu, and L. Zikatanov. Uniformly convergent multigrid methods for convection diffusion problems without any constraint on coarse grids. SIAM J. Sci. Comp., (submitted), 1999.

[77] R. Kohn and G. Strang. Optimal design and relaxation of variational problems. I. *Comm. Pure Appl. Math.*, 39(1):113–137, 1986.

[78] R. Kohn and G. Strang. Optimal design and relaxation of variational problems. II. *Comm. Pure Appl. Math.*, 39(2):139–182, 1986.

[79] R. Kohn and G. Strang. Optimal design and relaxation of variational problems. III. *Comm. Pure Appl. Math.*, 39(3):353–377, 1986.

[80] C. Zenger und S. Zimmer M. Griebel. Multilevel Gauss-Seidel-algorithms for full and sparse grid problems. *Computing*, 49:127–148, 1993.

[81] G. Zumbusch M. Griebel. Parallel multigrid in an adaptive PDE solver based on hashing and space-filling curves. *Parallel Computing*, 25:827–843, 1999.

[82] S. Knapek M. Griebel. *A multigrid-homogenization method*, volume 59 of *Modeling and Computation in Environmental Sciences*, pages 187–202. Vieweg-Verlag, Braunschweig, 1997.

[83] G. W. Milton. Bounds on the elastic and transport properties of two-component composites. *J. Mech. Phys. Solids*, 30(3):177–191, 1982.

[84] J. M. Morel and S. Solimini. *Variational methods for image segmentation*. Birkhauser, Boston, 1995.

[85] D. Mumford and J. Shah. Optimal approximation by piecewise smooth functions and associated variational problems. *Comm. Pure Appl. Math.*, 42:577–685, 1989.

[86] M. Griebel N. Rösch, S. Krüger and C. Zenger. Quantenchemie auf Parallelrechnern, Zur Perspektive der Dichtefunktionaltheorie. In *Proceedings der BMWF-Tagung HPSC95, Aachen, 1996*.

[87] S. Osher and J. Sethian. Fronts propagating with curvature dependent speed, algorithms based on a hamilton-jacobi formulation. *J. Comp. Phys.*, 79:12–49, 1988.

[88] M.C. Payne, M.P. Teter, D.C. Allan, T.A. Arias, and J.D. Joannopoulos. Iterative minimization techniques for ab initio total-energy calculations: molecular dynamics and conjugate gradients. *Rev. Mod. Phys.*, 64(4):1045, 1992.

[89] P. Perona and J. Malik. Scale space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Mach. Intelligence*, 12(7):629–639, July 1990.

[90] W. E. Pickett. Pseudopotential methods in condensed matter applications. *Comput. Phys. Rep.*, 9:115–198, 1989.

[91] E. Sharon, A. Brandt, and R. Basri. Fast multiscale image segmentation. September 1999.

[92] X.C. Tai, J. Froyen, M.S. Espedal, and T.F. Chan. Overlapping domain decomposition and multigrid methods for inverse problems. *AMS Contemporary Mathematics*, 218:523, 1998.

[93] Bart M. ter Haar Romeny. *Geometry-Driven Diffusion in Computer Vision*. Kluwer Academic Publishers, 1994.

[94] C. Vogel and M. Oman. A fast, robust algorithm for total variation based reconstruction of noisy, blurred images. *IEEE Trans. Image Proc.*, 7:813–824, 1998.

[95] W.L. Wan, T.F. Chan, and B. Smith. Energy minimizing interpolation for robust multigrid. *SIAM J. Sci. Comp., to appear*, 1999.

[96] C. Wang, J. Wu, J. Xu, and L. Zikatanov. Newton-Krylov-Multigrid method for lithium battery simulation. In preparation, 2000.

[97] C.Y. Wang, W.B. Gu, and B.Y. Liaw. Micro-macroscopic coupled modeling of batteries and fuel cells i. model development. *J. Electrochem. Soc.*, 145:3408, 1998.

[98] F. Wang and J. Xu. A cross-wind strip block iterative method for convection-dominated problems. *SIAM J. Sci. Comp.*, 21(2):646–665, 1999.

[99] R. Wildenhues. *Ein paralleles Multilevel-Verfahren für die Dichtefunktionaltheorie*. thesis, Institut für Angewandte Mathematik, Universität Bonn, 2000.

[100] J. Xu. A novel two-grid method for semilinear elliptic equations. *SIAM J. Sci. Comput.*, 15(1):231–237, 1994.

[101] J. Xu. Some two-grid finite element methods. In A. Quarteroni, J. Périaux, Yu. A. Kuznetsov, and O. B. Widlund, editors, *Domain Decomposition Methods in Science and Engineering: The Sixth International Conference on Domain Decomposition*, volume 157 of *Contemporary Mathematics*, pages 79–87, Providence, Rhode Island, 1994. American Mathematical Society.

[102] J. Xu. The auxiliary space method and optimal multigrid preconditioning techniques for unstructured meshes. *Computing*, 56:215–235, 1996.

[103] J. Xu. Two-grid discretization techniques for linear and nonlinear PDEs. *SIAM J. of Numer. Anal.*, 33(5):1759–1777, 1996.

[104] J. Xu. The EAFE scheme and CWS method for convection dominated problems. In P. Bjørstad, M. Espedal, and D. Keyes, editors, *Proceedings of the Nineth International Symposium on Domain Decomposition Methods in Science and Engineering*. John Wiley & Sons Ltd, 1997.

[105] J. Xu and A. Zhou. Some local and parallel properties of finite element discretizations. In *Proceedings of the Eleventh International Conference on Domain Decomposition Methods*, 1998. (to appear).

[106] J. Xu and A. Zhou. Local and parallel finite element algorithms based on two-grid discretizations. Math. Comp., (to appear), 1999.

[107] J. Xu and A. Zhou. A two-grid discretization scheme for eigenvalue problems. Math. Comp., (to appear), 1999.

[108] J. Xu and L. Zikatanov. A monotone finite element scheme for convection diffusion equations. *Math. Comp.*, 68:1429–1446, 1999.

[109] H.K. Zhao, T. Chan, B. Merriman, and S. Osher. A variational level set approach to multiphase motion. *J. Comp. Phys.*, 127:179–195, 1996.