Advances in Materials Science: SGI, Accelrys and Intel Work to Extend the Reach of Computer Simulation in Materials Research

Daniel P. Stevens DMD, Ph.D., Eric Renner, B.S., Roberto Gomperts, Ph.D., SGI; Tim Mattson, Ph.D., Intel; George Fitzgerald, Ph.D., Accelrys

Table of Contents

Materials Science: Challenges and Opportunities

To remain competitive in today's research environment, leading materials scientists must rely on the best available methodologies and tools to develop the next generation of innovative materials. Computer simulations are increasingly used to guide experiments, help interpret experimental results, and even to replace laboratory tests altogether. Materials Science, however, is at a crossroads. Realistic simulations that take into account effects at the atomic and molecular level are becoming more and more important in achieving an understanding of the fundamental nature of materials on the macro-, meso- and nano-scales.

Detailed molecular simulations, however, rarely scale economically to the numbers of atoms necessary to realistically simulate, understand and predict the properties of materials. As a result, rational material design is still in its earliest stages. Researchers are struggling to create realistic models that can be solved in a reasonable amount of time using current algorithms and computing technology. Common engineering simulations that are highly useful at the macroscale—such as finite element analysis (FEA) and computational fluid dynamics (CFD)—break down at the atomic scale. Bridging this gap is essential to continued advances.

Over the past ten to fifteen years a wide variety of codes have been developed to study a broad range of materials and molecular systems. These programs are capable of accurately predicting structural, thermochemical, electronic, optical and other properties. These tools are being used to increase understanding of catalysis, to analyze behavior at surfaces and interfaces and to attack a myriad of other important problems. Scaling such codes to accommodate

more complex systems has the potential to revolutionize materials science by rationalizing materials development and changing the way that scientists and engineers think about and approach problems. Using existing quantum mechanical-based methods, the time to run a calculation increases anywhere from N^3 to N^7 (where N is the number of atoms in the model). Employing approximations can reduce this to almost N^2, or even to order-N in some cases. Normally, however, doubling the size of a model increases the time to solution by at least 8-fold.

Improving code performance to effectively tackle large atom counts with high precision requires efficient, parallelized codes running on highly scalable multi-processor computing hardware. The memory required by a calculation is also driven by the number of atoms as well as the desired level of precision. For some calculations, available memory can be a bigger limiting factor than time to solution.

Increasing the Parallelism of Critical Materials Science Codes

The Accelrys Materials Studio® product suite provides advanced simulations of structural, energetic, optical and magnetic properties for chemicals and materials R&D. Materials Studio includes a large number of important codes from a wide variety of sources including leading academic institutions and research organizations worldwide, as well as internally developed programs. The Accelrys R&D staff works to incorporate additional features, further optimize the algorithms, improve scalability, and integrate these codes into the Materials Studio framework for ease-of-use. SGI and Intel have been working closely with Accelrys to increase and improve the parallel efficiency of many of these important programs. *"Our close collaboration with Intel and SGI is not only helping us improve our codes today, but giving us greater insight into how computing hardware will evolve,"* says Keith Glassford, Director, Materials Science Marketing at Accelrys. *"This allows us to better understand what we should be doing now to ensure that programs keep scaling in the future."*

Significant speedup has been achieved by improving parallelism within programs and eliminating computing bottlenecks. According to Roberto Gomperts, Principal Scientist at SGI, *"Accelrys scientists and their collaborators have done a great job developing and improving algorithms to make them as efficient as possible. SGI is helping improve the parallelism of the codes for execution on multi-processor systems to decrease the time to solution for large problems. For instance, we've enabled some of the features of the SGI Altix architecture, such as globally-addressable shared memory, to provide better scaling at high processor counts."*

The Message Passing Interface, MPI, is a standard API that is widely used for the creation of parallel programs. The MPICH library used by the Accelrys codes did not fully utilize the architectural features of Altix, which limited scalability to a modest number of processors. Beyond 8 processors communication starts to become a serious bottleneck decreasing scaling efficiency. As a result, 24 processors was the practical scaling limit for most of these codes. *"By simply substituting our optimized library for Altix we are able to scale to 128 processors and beyond with exceptional efficiency on some codes,"* Gomperts continues. *"In some cases, people were applying 16 processors to a problem and only getting a 3X speedup versus a single processor. In such a scenario the cost of computation is high and the incentive to apply multi-processing is correspondingly low. By comparison, we've achieved an over 10X speedup with 16 processors running the same simulation. The globally shared memory of the Altix architecture also makes it possible to tackle very large memory problems that would be beyond the reach of typical cluster architectures."*

The benefit obtainable from additional processors is directly dependent on the amount of parallelism in a program as governed by Amdahl's Law. Smaller problems with fewer atoms spend a smaller percentage of time in parallel execution than larger problems. *"The lesson is that you can get the most from your computing resources by matching the number of processors to the class of problem,"* says Gomperts. *"Running a small problem on 128 processors might be wasting valuable compute resources that could be better allocated to work on multiple*

problems in parallel for greater total throughput. On the other hand, improved parallelism can bring large problems that would formerly have been considered grand challenge problems into the realm of routine computation."

Dramatic Speed Up for Real World Problems

A few cases based on real customer examples illustrate the improvements that can be achieved in scalability and performance. Both test cases use quantum mechanical codes based on density functional theory (DFT). These algorithms have been proven to provide high accuracy at a low cost in terms of computational time. Both are available as part of Accelrys Materials Studio.

CASTEP

The first example uses the *ab initio* program CASTEP. This code has proven to be particularly accurate at predicting a variety of electronic, optical and structural properties of solids, interfaces and surfaces, and is therefore widely used by electronics companies, by organizations doing R&D on optical materials, and for alloy development in the automotive and aerospace industries. The pseudopotential plane-wave technology underlying CASTEP is well validated, with hundreds of scientific publications written each year showing new applications of the code.

Innovene (formerly BP-Amoco Chemicals, now a wholly-owned subsidiary of BP), is one of the five largest petrochemical companies in the world, and increasingly depends on computer simulation to understand complex catalytic processes. When scientists at Innovene needed to better understand the aging of the catalyst vanadyl pyrophosphate, SGI and Accelrys helped find a solution. This catalyst is widely used in the conversion of butane to maleic anhydride, a chemical intermediate used in the synthesis of a variety of important agricultural chemicals, dyes and pharmaceuticals. Over time, vanadyl pyrophosphate loses its catalytic ability through the loss of phosphorus. Although the catalyst can be reconstituted, Innovene wants to better understand this process in hopes of improving long term catalyst performance. Because there are easily measurable and characteristic changes in the optical properties of the catalyst associated with degradation, CASTEP was a perfect code for this study. The idea was to begin with the unit cell of vanadium pyrophosphate and model the removal of a single phosphorus unit.

Unfortunately, the 104-atom unit cell of the catalyst was much larger than the 40 to 50 atoms typically simulated on the 32-bit Linux cluster at Innovene. Such simulations already take several days to complete; doubling the number of atoms in this case increases the time to solution by a factor of 8X and dramatically increases the amount of memory required. Both of these factors create an intractable problem given the current hardware. For Innovene and many other cluster-users, the longer the time to completion, the greater the chances a job will be interrupted due to system failure, maintenance needs, or memory limitations.

SGI and Accelrys undertook this study using an Altix 3700 Bx2 with 1.6 GHz Itanium 2 processors with 9MB of cache per processor. It quickly became clear that the problem could not be run at all with less than 64GB of memory, putting it out of the reach of most or all standard cluster nodes. To determine the scalability of CASTEP running on Altix for this problem, SGI looked at the time required to complete each self-consistent field (SCF) cycle. Each geometry optimization step in the simulation requires computing the energy plus the gradient of the energy to find the minimum. The energy is computed using an iterative, SCF procedure, and each geometry optimization step requires about 20 SCF steps.

The results are summarized in the following chart and table:

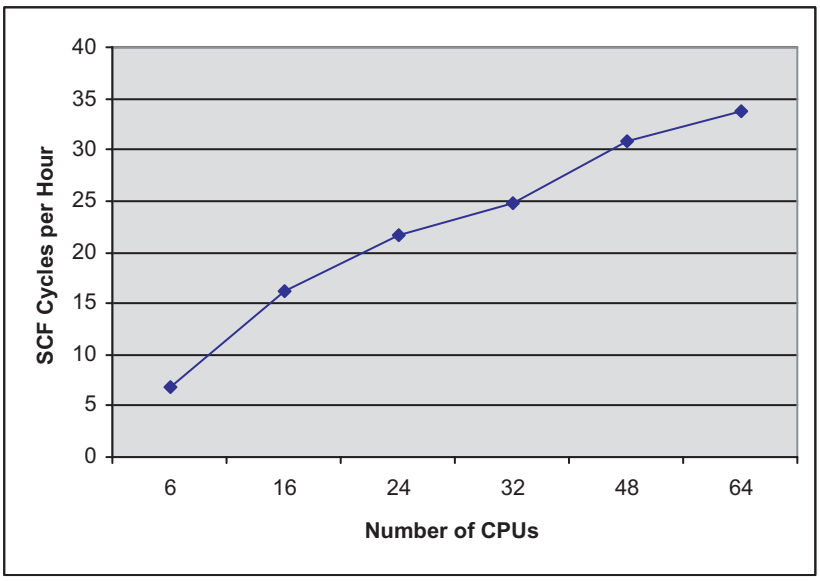

Number of Processors	SCF (Cycles per Hour)
6	6.92
16	16.18
24	21.56
32	24.88
48	30.87
64	33.74

Using 32-processors, the full calculation on the Altix takes about 2.5 days versus the greater than two weeks it would have taken on Innovene's cluster had the memory been sufficient to run the problem. For this problem it is clear that the benefits obtainable from scaling beyond 32 processors diminish, suggesting that for optimal efficiency, this simulation should be run on no more than 32-48 processors and additional resources are best allocated to run other jobs.

The optical properties calculated as a result of the simulation correlate well with experimental results, giving Innovene greater understanding of what is happening to the unit cell as the catalyst ages. Scientists at Innovene would like to be able to do more of this type of modeling in the future to understand more complex problems, tackling unit cells as large as 300 atoms. According to Gerry Zajac, Innovene Research Associate, *"The ability to routinely do these types of large calculations in a reasonable timeframe would create a lot of opportunities for us to look at systems that we couldn't previously approach. Right now we have to break up problems into little pieces and use a time-consuming, mosaic approach. With this level of ability we could do complete studies of catalytic reactions including side reactions, intermediates and by-products, and have results in a matter of months rather than years. We could also look at systems that previously could only be studied experimentally."*

Unit cell of the catalyst vanadyl pyrophosphate studied at Innovene.

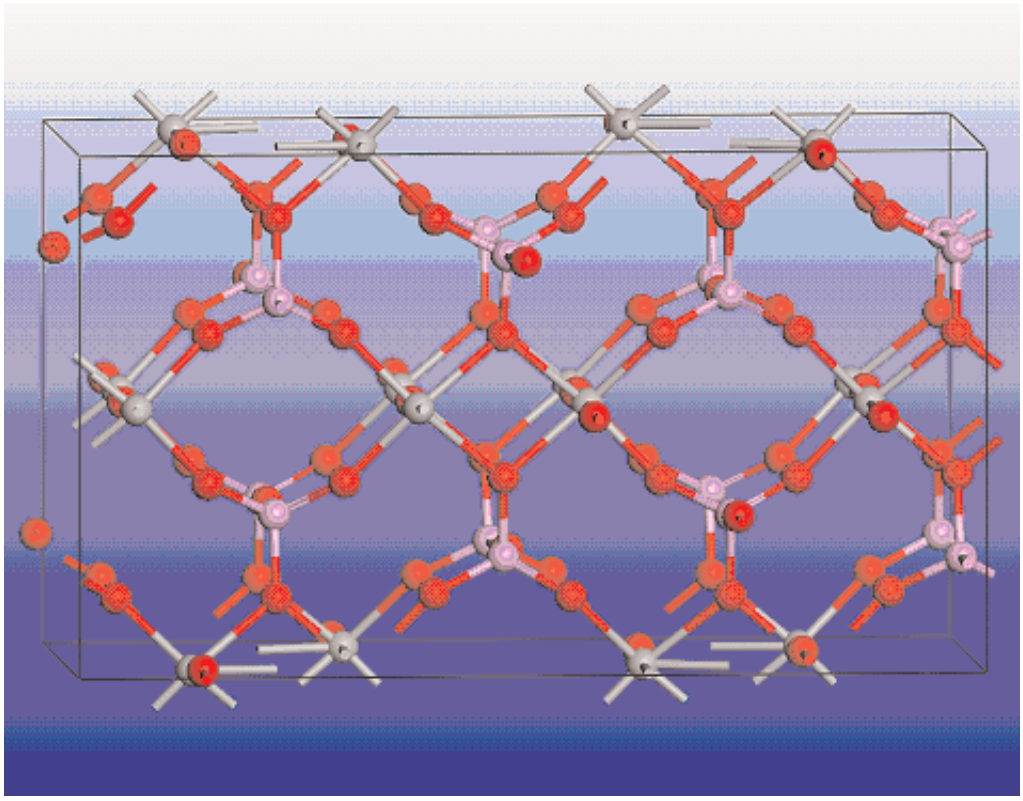

Unit cell of the catalyst after deactivation through loss of a phosphorus unit.

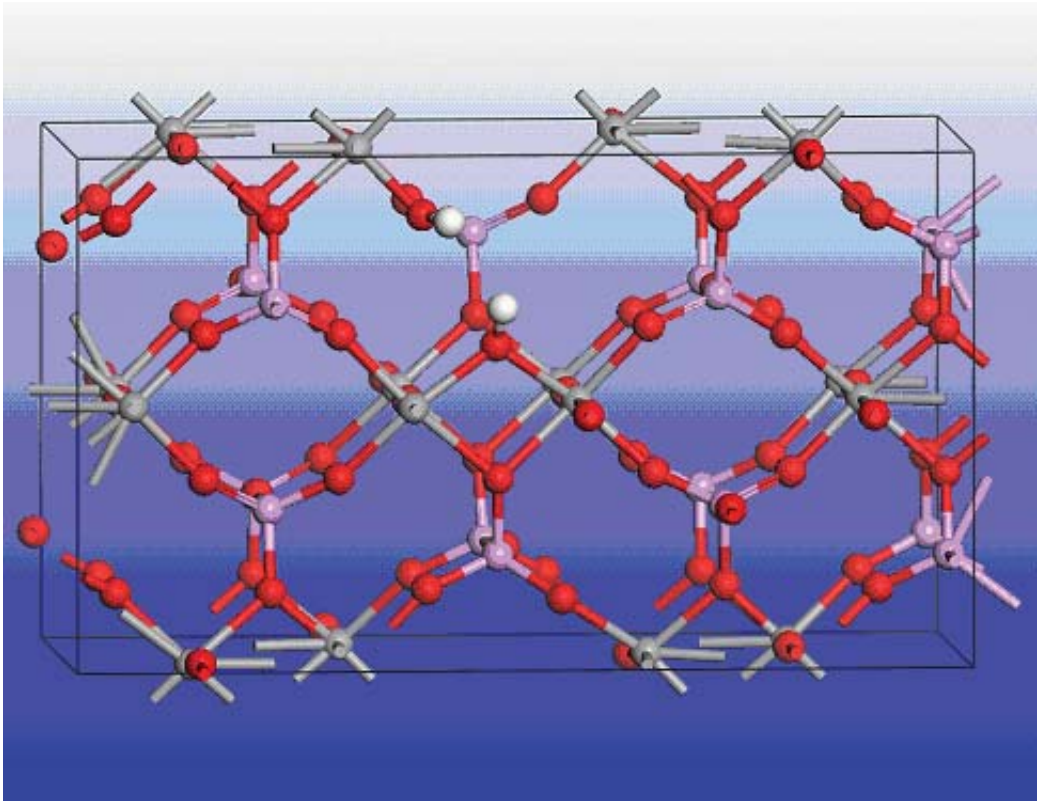

DMol³

The second example uses another popular quantum mechanical code, DMol³. In contrast to plane-waves, this program constructs molecular orbitals from combinations of atomic orbital basis functions, making it very intuitive for use by chemists. The code is ideal for the study of chemical reactions, especially on surfaces or within microporous materials. Like CASTEP, DMol³ provides highly accurate predictions of thermochemical, electronic and optical properties. The code is widely used in the petrochemical, fine and specialty chemical industries.

Nitrotoluene is an important chemical precursor for agricultural chemicals, explosives and pharmaceuticals. In all these applications, the para- is the desired orientation. The traditional method of preparation results in more or less equal amounts of ortho-, and para-forms of nitrotoluene—an inefficient process requiring an additional separation step—plus it creates a significant amount of waste acid and other byproducts that are difficult to treat.

An alternative approach combines nitric acid and acetic anhydride in beta zeolite to yield acetyl nitrate. Adding toluene to this mixture produces nitrotoluene—with 70% in the desired para-orientation. The only byproduct of the reaction is environmentally benign acetic acid (vinegar).

Accelrys scientists are using computer simulation to better understand the mechanism of this reaction in beta zeolite and the strong preference for the para-orientation—a study that is extremely well suited to the strengths of DMol³. Each calculation consumes about a week on 8 processors (4 nodes) using a 32-bit Linux cluster (dual 1.2 GHz Intel® Pentium® III processors, 512KB of cache, with 1GB of memory per node, 100Mb/sec Ethernet cluster interconnect). The average time for a Single Point Energy Calculation is about 57 minutes.

For comparison, the same simulation was run on 8-processors of an SGI Altix 3700 Bx2 with 1.6GHz Itanium 2 processors. Each Single Point Energy Calculation completes in about 14 minutes

Transition state of toluene nitration in the pores of beta zeolite. The nitration reaction is aided by an acetic acid molecule at the Bronsted acid site.

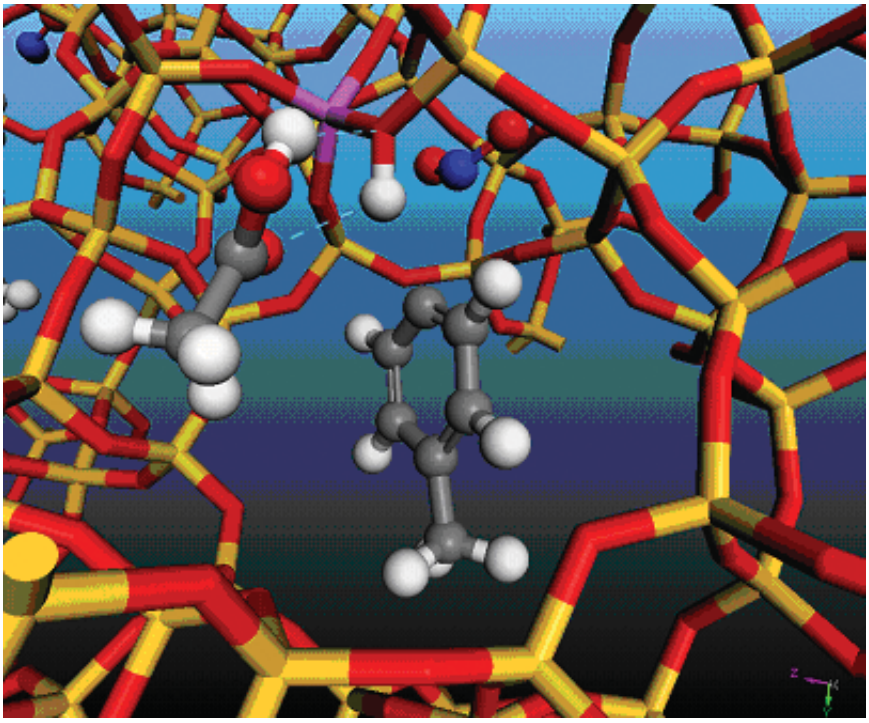

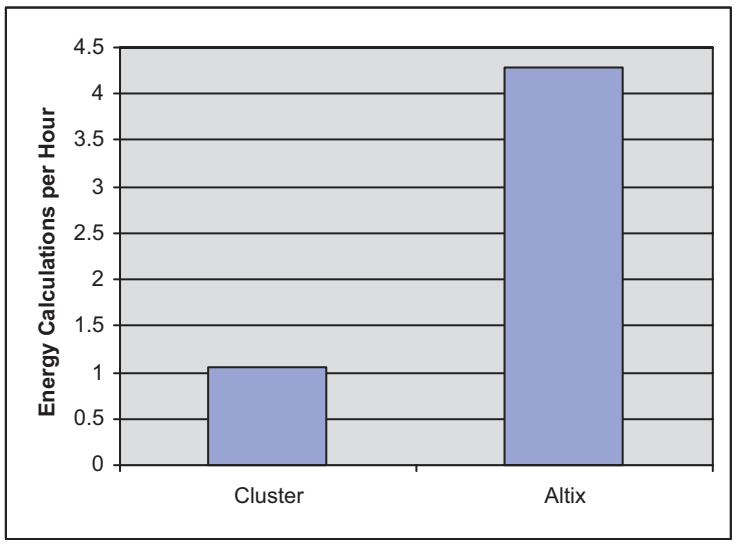

for a 4X improvement in overall job execution. On an 8-processor Altix system the entire job can be completed in 1-2 days versus the 6-7 days required on the Pentium III-based Linux cluster. The speedup allows Accelrys scientists to complete 4X the work in the same amount of time to better elucidate reaction mechanisms in complicated zeolite systems.

SGI® Altix® and Intel® Itanium® 2 Processors

The speedup achieved with the codes in the above examples versus that seen on other multi-processor systems is due in large part to the advantages of the SGI Altix platform and Intel Itanium 2 microprocessors. Accelrys' Materials Science codes benefit significantly from the inherent parallelism of Itanium 2 combined with the global shared memory of SGI Altix.

The Itanium 2 microprocessor is based on the Explicitly Parallel Instruction Computing or EPIC architecture. Since the instruction set is explicitly parallel, the architecture attains high levels of parallelism inside the processor (so called "instruction level parallelism"). The Intel Itanium 2 compiler takes directives from the programmer to create highly parallelized machine code that takes the fullest advantage of the processor's multiple execution units and massive on-dye resources. The Itanium 2 is able to process 6 instructions per clock cycle versus 3 instructions per clock cycle for the Intel Xeon chip. Itanium 2 actually significantly outperforms Xeon chips running at twice the clock rate when executing CASTEP.

SGI Altix family of servers

Leveraging the EPIC architecture of the Itanium 2 processor, SGI Altix adds an advanced parallel system architecture, immense internal resources, and exceptional bandwidth to memory and storage systems, to deliver industry-leading performance for scientific computing applications. *"The Itanium 2 processor's EPIC architecture is perfect for materials science simulations using codes such as CASTEP and DMol³"* said Tim Mattson Ph.D., Parallel Computing Evangelist at Intel. "*The Altix interconnect lets applications scale up to large numbers of Itanium 2 processors. At the same time, programming is dramatically simplified due to the large shared address space provided*

by SGI's NUMAflex architecture. This significantly lowers the barrier for institutional programmers and ISVs that want to port to the platform."

In addition to the raw microprocessor performance of the Itanium 2 processor, 64-bit Altix offers scientists a virtually unlimited memory space coupled with an ultra-high speed communication fabric that connects the system's microprocessors, memory, graphics processing units (GPUs) and other peripherals to eliminate the bottlenecks that can limit scalability. This capability enables research teams to manage CPU and memory resources more effectively so the most urgent jobs complete more quickly. The ability to dynamically allocate resources enables scientists to be more flexible and hence make more effective use of the computing resources at their disposal for maximum productivity.

Conclusion

The results obtained in these tests illustrate the power of the SGI Altix platform for elucidating critical materials science investigations. Altix provides the scalability and memory capacity to accommodate problems that are beyond the reach of other systems. SGI, Accelrys, and Intel are committed to work to further improve the scalability of important materials science codes.

"The speedup achieved with Itanium 2 and the Altix platform makes our simulation codes applicable to bigger systems than ever before," says Keith Glassford. *"This brings us closer to bridging the gap between the engineering world and the molecular world with far-reaching implications. As more powerful tools enable engineers to adopt a more comprehensive view they will naturally transition to a more exploratory approach to their work that will increase the rate of innovation."*

3852 [08.2005] J15007