SANDIA REPORT

SAND97-0770 • UC-405

Unlimited Release

Printed April 1997

Phase Behavior and Molecular Scale Dynamics of Confined Fluids Using New Atomistic Simulation Techniques: CRADA #1140 Close-Out Report

RECEIVED

MAY 27 1997

OSTI

Grant S. Heffelfinger, Brian K. Peterson, Martin E. Lewitt

Prepared by

Sandia National Laboratories

Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000.

MASTER

Approved for public release; distribution is unlimited.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Prices available from (615) 576-8401, FTS 626-8401

Available to the public from
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Rd
Springfield, VA 22161

NTIS price codes
Printed copy: A03
Microfiche copy: A01

DISCLAIMER

**Portions of this document may be illegible
in electronic image products. Images are
produced from the best available original
document.**

SAND97-0770
Unlimited Release
Printed April 1977

Distribution
Category UC-405

Phase Behavior and Molecular Scale Dynamics of Confined Fluids using New Atomistic Simulation Techniques: CRADA #1140 Close-Out Report

Grant S. Heffelfinger
Manager
Computational Materials Science Department
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-1111

Brian K. Peterson (formerly with Mobil R&D)
Senior Principal Computational Chemist
Air Products and Chemicals, Inc.
7201 Hamilton Blvd.
Allentown, PA 18195-1501

Martin E. Lewitt (formerly with nCUBE)
Technical Consultant
Hewlett Packard
P.O. Box 729
Sandia Park, NM 87047-0729

Abstract

The Mobil-nCUBE-Sandia CRADA #1140 was aimed at developing the tools and capabilities to model fluid phase behavior and flow in pores--ultimately leading to a better understanding of these phenomena. In this CRADA close-out report the collaboration's original goals are revisited, the course of the CRADA is traced, and the technical accomplishments of the CRADA are presented.

Contents

Summary	vi
---------------	----

Phase Behavior and Molecular Scale Dynamics of Confined Fluids using New Atomistic Simulation Techniques: CRADA #1140 Close-Out Report

Introduction.....	1
Background and History of the CRADA.....	1
Description of the Revised CRADA.....	2
The Course of the CRADA.....	3
Technical Accomplishments.....	3
The Grand Canonical Monte Carlo Code and Spatial Parallel Algorithm.....	4
Validation of the PGCMC Algorithm.....	8
Performance and Optimization of the PGCMC Code.....	8
Scaling Characteristics of PGCMC.....	11
Extension to Fluid-Pore Systems and Realistic Model Development.....	14
Impact of the CRADA On Other Sandia Projects and Benefits to DOE.....	17
References.....	18
Appendix--Quarterly Report Dated 9/95.....	A-1

Figures

Figure 1.The Spatial Monte Carlo Decomposition.....	5
Figure 2.Subdomain Sizing.....	6
Figure 3.Parallel Energy Calculation for the Spatial Monte Carlo Decomposition. .	9

Tables

Table 1.A Comparison Between PGCMC and NVT MD for State Point A.....	10
Table 2.State Points	11
Table 3.Scaling Results for the Spatial Decomposition (PGCMC)	12
Table 4.Speedup and Efficiency Results for PGCMC for State Point A	13

Summary

The Mobil-nCUBE-Sandia CRADA #1140, originally entitled, "Large-Scale Modeling of Synthetic Lubricants and Elastohydrodynamic Contacts," was renegotiated to become "Phase Behavior and Molecular Scale Dynamics of Confined Fluids using New Atomistic Simulation Techniques," roughly 6 months after its execution. The revised CRADA was aimed at developing materials modeling methods and companion massively parallel algorithms to enable the modeling of fluid phase behavior and flow in pores. The technical challenges of the proposed work included developing new atomistic simulation methods and companion massively parallel algorithms, software optimization, developing realistic fundamental models for the systems of interest, and applying the software tools to the specific problems involved with economically removing oil locked in mesoporous rock.

The technical accomplishments of the CRADA include 1) developing a parallel algorithm and grand canonical Monte Carlo code (**PGCMC**) for simple homogeneous bulk systems, 2) extending the **PGCMC** code to apply to more complex fluid-pore systems while adding the capability to determine necessary thermodynamic quantities such as the pressure tensor, and 3) developing realistic models of the fluid-pore system components. Other planned work was interrupted by restructuring at Mobil Research and Development Corporation as well as career changes for all three of the collaborators (Brian Peterson, Mobil, Martin Lewitt, nCUBE, and Grant Heffelfinger, Sandia).

The Mobil-nCUBE-Sandia CRADA was synergistic with existing activities and core competencies at Sandia by design. It drew upon Sandia's computational science and computational facilities and expertise as well as Sandia's petroleum exploration and production research sponsored by DOE Energy Research. The parallel algorithm employed in the **PGCMC** code found immediate use in an important simulation tool developed at Sandia, Dual Control Volume Grand Canonical Molecular Dynamics (DCV-GCMD), and was incorporated into the DCV-GCMD code, LADERA, immediately after its development. The DCV-GCMD method enables the dynamic simulation of a system experiencing a steady-state chemical potential gradient and has found application to several important DOE Energy Research and Defense Program problems including diffusion in zeolites and silicalites, diffusion of drugs through biomembranes, and diffusion-limited oxidation of polymers. The latter has been identified as an important materials aging problem for the enduring stockpile due to its effect on the elastomeric materials used in o-rings resulting in the loss of hermetic seal. Other new algorithms developed in the course of this work included extending the **PGCMC** code to a wide variety of physical systems and fluid types. The Harasima definition of the pressure tensor was modified, the spatial decomposition of the GCMC algorithm was extended to inhomogeneous systems, and arbitrary types of particles and potentials were added, increasing the scope of physical systems able to be modeled. All of these were achieved in a parallel scheme without significant change to the scaling properties of the final program.

Phase Behavior and Molecular Scale Dynamics of Confined Fluids using New Atomistic Simulation Techniques: CRADA #1140 Close-Out Report

Introduction

The Mobil-nCUBE-Sandia CRADA #1140, was originally entitled "Large-Scale Modeling of Synthetic Lubricants and Elastohydrodynamic Contacts." Roughly six months after it's execution the scope of work was renegotiated at the request of Mobil Research & Development and the title of the modified CRADA became "Phase Behavior and Molecular Scale Dynamics of Confined Fluids using New Atomistic Simulation Techniques." The new CRADA was aimed at developing the tools and capabilities to model fluid phase behavior and flow in pores--ultimately leading to a better understanding of these phenomena. While excellent work was carried out on these subjects during the course of the CRADA, the collaboration was fraught with problems throughout it's 2.5 year course, largely due to ongoing changes in Mobil Research & Development Corporation. In this CRADA close-out report the original goals of the CRADA are revisited, the course of the CRADA is traced, and the technical accomplishments of the CRADA are presented.

Background and History of the CRADA

The CRADA was executed June 1993 with its original title, "Large-Scale Modeling of Synthetic Lubricants and Elastohydrodynamic Contacts." The work was aimed at enhancing the ability to predict and understand the thermodynamic and transport properties of fluids while developing new algorithms for parallel materials modeling codes. In particular, the main goal of the work was to improve the material properties of commercial lubricants by studying the relationships between molecular structure and rheological properties as well as between rheological properties and lubricant performance in elastohydrodynamic contacts. These phenomena were to be studied via molecular dynamics simulations and large-scale fluid dynamics simulations, respectively. Another area of mutual interest was the porting and optimization of molecular and fluid dynamics simulation codes for parallel architectures.

By October 1993 Mobil Research & Development Corporation was in the middle of significant restructuring leading to a reduction in personnel of 20% in Mobil R&D Engineering and a shift in Mobil R&D's research emphasis from the processing of oil

(downstream operations) to the exploration for and production of oil (upstream operations) and ultimately a request for a change in the technical directions of the CRADA. As the materials modeling techniques and the scale of the computational requirements of the new CRADA work were largely unchanged, the tasks to be performed by all three partners, Mobil, nCUBE, and Sandia, were largely unaffected. Thus the CRADA renegotiation was accomplished quite easily and work began in earnest on the first of the two major thrusts of the CRADA.

Description of the Revised CRADA

The revised CRADA, "Phase Behavior and Molecular Scale Dynamics of Confined Fluids using New Atomistic Simulation Techniques," was aimed at developing materials modeling methods and companion massively parallel algorithms to enable the modeling of fluid phase behavior and flow in pores. As stated in the revised Joint Statement of Work, "As known petroleum reserves shrink, and production from existing fields become economically unattractive, US oil companies increasingly need technological advantages to explore and produce new petroleum resources. MOBIL's effectiveness in this regard depends on its ability to understand and model, at the microscale, the fluid distribution and flow in pores." Most microscopic models of confined fluids are limited by the fact that they use macroscopic continuum level theories and constitutive relationships for pores ranging in size from mesopores (>20 Angstroms) to macropores (>500 Angstroms). These relationships are known to fail for small pores yet the exact nature of this breakdown as well as at what pore size it occurs is unknown. Therefore, models and methods for these systems were needed to determine where the existing macroscopic models fail and to predict fluid behavior in these regions. Atomistic computer simulation was the most promising approach, however, the size and complexity of the computer simulations required massively parallel processing.

The diverse requirements and technical challenges of the revised CRADA included: 1) developing new atomistic simulation methods, 2) developing companion massively parallel algorithms for and software optimization of these new simulation methods, 3) developing realistic fundamental models for the systems of interest, and 4) applying of the software tools to the specific problems and questions involved with economically removing oil locked in mesoporous rock. Sandia's contribution to the CRADA drew upon its experience and expertise in computer simulation technology, especially developing codes for the massively parallel computers. Mobil contributed realistic models for the fluid-pore systems and experimental data. nCUBE's interest, helping to establish massively parallel supercomputers as an atomistic simulation resource through software development was reflected by their contribution to porting serial applications to MIMD massively parallel supercomputers and optimizing such applications for parallel architectures.

The revised CRADA, focused on fluids confined in mesopores, had two thrusts: equilibrium and transport behavior.

Equilibrium Behavior: Capillary Pressure and Phase Behavior for Fluids-Pore Systems

Because the region of validity of Laplace's equation, which predicts capillary pressure for fluid-pore systems, is not only complex but also poorly understood, one goal of the CRADA was to clarify the region of validity of Laplace's equation and to develop a suitable improved/modified version of this equation. Grand canonical molecular dynamics and Monte Carlo simulations were determined to be the most suitable tools for this work and thus a focus of the resulting work involved developing a new massively parallel grand canonical Monte Carlo (GCMC) code suitable for addressing these questions.

Transport Behavior: Validity of Darcy's Law and Constitutive Relation Development

This second goal of the CRADA was to realistically model transport in multiphase systems confined in pores of various sizes and to use the resulting code or codes to elucidate the region of validity of Darcy's Law and develop suitable improvements/modifications to this equation.

The Course of the CRADA

Work on the first project (Task 1) began July 1993 after the revised CRADA was signed and progressed steadily until early 1995. At that point Mobil R&D began another restructuring process that ultimately led to a further reduction in personnel of 30% and the end of Mobil's Central Research Laboratory (CRL), a 30 year-old institution with nearly 150 employees before the restructuring process began. Work on the CRADA throughout 1995 was steady and remained focused on Task 1 of the revised CRADA. However, by November of 1995 Brian Peterson, the primary Mobil employee working on the CRADA, had left the CRL and moved to the Mobil's new Strategic Research Center (SRC) located near at the Paulsboro refinery in southern New Jersey. Meanwhile, Martin Lewitt, the primary nCUBE employee on the CRADA, left nCUBE in August 1995 and in November 1995, Grant Heffelfinger, the primary Sandia employee working on the CRADA assumed new job responsibilities as the acting department manager of the Parallel Computing Sciences Department. At that point, work ceased on the CRADA and it was subsequently terminated in the spring of 1996. By May of 1996, Brian Peterson had left Mobil R&D Corporation.

Technical Accomplishments

Because the three major personnel disruptions discussed above coincided not only with each other but also with the planned initiation of the CRADA's Task 2--modeling transport behavior in fluid-pore systems, to begin January 1996, work never commenced on Task 2 of the revised CRADA. The technical accomplishments of the CRADA, while significant, are therefore confined to Task 1--modeling equilibrium behavior in fluid-pore systems. This is borne out by the CRADA quarterly report dated September 1995 which is included in the

Appendix. Therefore, the contents of this close-out report will center on the milestones and accomplishments associated with Task 1 of the CRADA: 1) developing a parallel algorithm and grand canonical Monte Carlo (GCMC) code for simple, homogeneous bulk systems, 2) extending the parallel GCMC code to apply to more complex fluid-pore systems while adding the capability to determine necessary thermodynamic quantities such as the pressure tensor, and 3) developing realistic models of the fluid-pore system components.

The Grand Canonical Monte Carlo Code and Spatial Parallel Algorithm

The major effort associated with Task 1 was focused on developing a new massively parallel algorithm and code for grand canonical Monte Carlo simulation. While the code, **PGCMC**, and spatial parallel algorithm were first presented in November, 1994¹, the details of both are thoroughly discussed in a paper² which compared the spatial decomposition to alternative parallel decompositions. Therefore, what follows is an introduction and brief description of the spatial GCMC decomposition; the interested reader can find the details in the Journal of Computational Chemistry paper².

While molecular dynamics (MD) has proven relatively easy to adapt to parallel architectures due to the fact that the bulk of the computational work of an MD simulation, the calculation of the forces and new positions of the atoms, is performed for each atom in the system at the same time, parallel Monte Carlo (MC) simulations, on the other hand, have been more difficult to achieve due to the fact that atoms in a MC simulation are moved one at a time. While the bulk of the computational work in a MC simulation, the energy calculation of a proposed move, can be shared across processors, given that many MC simulations employ short-ranged potentials this approach is inefficient because only a small subset of the system's atoms interact with any one atom at any given point in the simulation. Therefore, any efficient massively parallel Monte Carlo algorithm for systems with short-ranged interactions must involve simultaneous moves on different processors.

The spatial GCMC decomposition relies on the fact that for short-ranged fluids, such as the cut and shifted Lennard-Jones potential relevant for the fluid-pore models proposed by Mobil, atoms separated by a greater distance than the reach of the potential act independently, thus different processors can work concurrently in regions of the same system which are sufficiently far apart. While the algorithm was developed for Monte Carlo simulations in the grand canonical ensemble (constant volume, V , chemical potential, μ , and temperature, T), it is equally well-suited for application to NVT Monte Carlo (constant number of atoms, N , volume, V , and temperature, T).

To take advantage of the fact that for short-ranged interaction potentials, processors can work simultaneously in the same volume as long as the regions in which they operate are separated by a distance further than the reach of the potential, the spatial decomposition requires that the system volume be divided into "domains," with one domain assigned to

each processor. For example, for a system composed of 64 domains: 2, 4, and 8 domains long in the x , y , and z directions, respectively, the **PGCMC** code employs 64 processors. These domains are further subdivided into 8 "subdomains". An 8 domain subsystem and its corresponding 64 subdomains (8 per domain) is shown in Figure 1. By assigning each processor to work in corresponding subdomains at the same time and choosing the subdomain dimension sufficiently large, each processor can perform the three types of GCMC moves (displacement, creation, and destruction) independently of its neighboring processors.

Figure 1. The Spatial Monte Carlo Decomposition.

A subsystem with 8 domains and 64 subdomains. Subdomain 0 in each domain is marked to show that the 8 processors responsible for the domains 0-7 of this subvolume are working concurrently in subdomain 0.

Subdomain Sizing

The simulation begins with each processor attempting a displacement followed by an attempted creation or destruction in subdomain 0 of the domain for which it is responsible (note the highlighted subdomain's 0 in Figure 1). We can see from Figure 1 that the distance between two subdomain 0's in different domains (and thus on different processors) is the edge length of a single subdomain. Thus, the events occurring in two subdomain 0's in neighboring domains should be independent if the subdomains are sized to be at least r^c (the range of the potential) on edge. However, it is possible that two simultaneous displacements in neighboring domains might occur such that the two atoms move within range of the interaction. For example, consider the example shown in Figure 2. If a proposed displacement of an atom near the edge of subdomain 0 in domain 0 moved the atom the maximum distance possible, Δr_{max} , toward subdomain 0 in domain 1 it would land in subdomain 1 in domain 0. Meanwhile if a simultaneous proposed displacement of an atom near the edge of subdomain 0 in domain 1 placed it Δr_{max} deep in subdomain 1 in domain 0, contention between the two proposed moves would occur unless the subdomains were sized greater than $r^c +$ twice the maximum possible displacement on edge. Thus, to remove the possibility of interprocessor contention yet obtain maximum speedup, the

subdomains should be sized exactly $r^c + 2\Delta r_{max}$ on edge. For example, if values of $\Delta r_{max} = 0.15\sigma$ and $r^c = 2.5\sigma$, the subdomains should be sized 2.8σ on edge.

Figure 2. Subdomain Sizing.

An example of two atoms in different domains (0 and 1) and thus on different processors, simultaneously moving the maximum possible distance, Δr_{max} , toward each other yet still remaining beyond the reach of the potential due to the proper sizing of the subdomains: $r^c + 2\Delta r_{max}$ or greater on edge.

Energy Calculation for Proposed Moves

The largest computational task for each type of move is the calculation of the energy change for a proposed move. The energy calculation is carried out in parallel as follows. Each processor, working in corresponding subdomains, chooses a set of test coordinates for which the system energy is required. For the case of an creation, the test coordinates represent the proposed atom's location; for a destruction, the coordinates will be for an atom proposed for destruction; for a displacement, the test coordinates will be either an atom's original position or its proposed new position. Each processor first calculates the contribution to its test coordinate's energy due to the atoms residing in the same subdomain as the test coordinate, followed by the energy contribution due to the atoms in the other 7 subdomains in the processor's domain. Thus, if all processors are working in subdomain 2 of their domains, the processor responsible for domain 5 in Figure 1 will first calculate the contribution to the subdomain 2 test coordinates due to the atoms in its subdomain 2,

followed by those in its subdomains 0, 1, 3, 4, 5, 6, and 7. Each processor then sends its test coordinates to its 7 "set I" neighbors, those possessing a subdomain which borders the subdomain containing the test coordinates, and receives the test coordinates from each of its "set II" neighbors, those for which it possesses a subdomain which borders the neighbor's subdomain which contains the neighbor's test coordinates. After calculating the energy contribution for all of its atoms in the appropriate subdomains and for the test coordinates of all set II neighbors, each processor sends this information to the set II neighbors and receives the corresponding information from its set I neighbors. Each processor is now in possession of the energy of its test coordinates *due to the entire system* and proceeds to accept or reject the proposed displacement, creation, or destruction.

The 7 set I neighboring processors responsible for subdomains which border subdomain S in domain D can be broken down into three groups based on the nature of their border with subdomain S and domain D (see Figure 3). Three of the processors are responsible for domains which share a face with domain D. Each of these three faces accounts for 4 subdomains, making 12 subdomains that border S and are in domains that share a face with S's domain D. Another 3 processors are responsible for domains which share an edge with domain D. Each edge accounts for 2 subdomains, giving another 6 subdomains that border S and are in domains that share an edge with S's domain D. Finally, 1 processor is responsible for a subdomain which shares a corner with S, bringing the total number of subdomains neighboring S and on the 7 set I processors to 19. Counting the 7 subdomains which border subdomain S and lie in the same domain D and are thus on the same processor, a grand total of 26 subdomains which border subdomain S is obtained, 19 on 7 set I processors and 7 on the same processor. The symmetry of this arrangement results in each processor needing the energy contributions of atoms which lie in subdomains on 7 other processors for each of its **own** subdomains (set I processors), as well as each processor providing the energy contributions for the atoms in its domain to 7 other processors (set II processors). Two 8 by 7 arrays are used to keep track of these two sets of neighbors for the eight subdomains.

Monte Carlo Moves: Creations, Destructions, and Displacements

When a processor accepts a proposed creation in one of the subdomains within the domain for which it is responsible, the x , y , and z coordinates of the atom are simply added to the arrays containing the coordinates of all atoms in that domain. Similarly, for a proposed destruction, the coordinates of the destroyed atom are removed from these arrays. Accepted displacements that occur within a domain are simply handled like an accepted creation (the atom's new position) followed by a destruction (the atom's old position). Since the arrays containing the coordinates are local to each processor, and thus to each domain, displacements that change subdomains while remaining in the same domain can be handled entirely by the processor responsible for a domain. However, occasionally a displacement is accepted for an atom that changes domains and therefore processors. This is handled as follows. When each processor proposes a new set of coordinates for an atom in one of its subdomains (keeping in mind that all processors are working in the correspondingly numbered subdomain within their domain), it checks to see if the new position is outside

its domain. Similarly, when each processor calculates the energy contribution of the atoms in its domain to the proposed displacement coordinates of its 7 set II neighbors, it checks each set of coordinates to determine if they lie within its domain. After the energy has been calculated for all proposed displacements, each processor determines if its own displacement is accepted or rejected. If a processor had determined that its displacement placed the atom outside its domain, it then sends a message to the processor responsible for the destination domain telling whether or not the displacement was accepted. Meanwhile, if a processor determined that the new coordinates of a neighboring processor's displacement landed the atom in its domain, it waits for this message from the neighboring processor regarding whether or not the displacement was accepted. If the displacement was rejected the simulation proceeds unchanged, if it was accepted, the processor responsible for the destination domain "takes possession" of the atom simply by adding its coordinates to the arrays containing the coordinates for the atoms in its domain (exactly like a creation) and the processor responsible for the departure domain treats the displacement like a destruction, removing the original coordinates of the atom from its coordinate arrays.

Validation of the PGCMC Algorithm

The spatial parallel GCMC program, **PGCMC**, was validated for several state points by comparing with an *NVT* molecular dynamics program. This was accomplished by using **PGCMC** to obtain the number of atoms, N , for a given set of values for μ , V , and T . Using T and values of N and V which yielded a corresponding density, $\rho = N/V$, in the *NVT* molecular dynamics program, the corresponding chemical potential, μ was determined. This value of the state point's chemical potential was then compared to that used as input to **PGCMC**. The results for an example state point (state point A) are shown in Table 1. For this example, the *NVT* MD simulation employed 2048 atoms and statistics were

accumulated for 80,000 MD timesteps of $\Delta t \sqrt{\frac{\epsilon}{m\sigma^2}} = 0.01$, after being warmed-up from a

lattice for 20,000 MD timesteps. The chemical potential was calculated using Widom's method³ every 10 timesteps. More discussion about the MD technique can be found in Allen and Tildesley⁴.

Performance and Optimization of the PGCMC Code

While **PGCMC** was also implemented on Sandia's 1872 processor Intel Paragon, most of the development and optimization work was carried on Sandia's 1024 processor nCUBE 2 massively parallel computer. The nCUBE2 is a hypercube topology, distributed memory, multiple instruction stream, multiple data stream (MIMD) computer. The individual nCUBE processor is a single chip VLSI implementation of nCUBE's proprietary instruction set architecture, integrating both communications and memory control. The 28 DMA channels used to implement the hypercube interconnect can operate independently in parallel with each other and in parallel with computation, limited ultimately only by their

contention for shared memory bandwidth. The channels implement "wormhole" routing which reduces the performance impact of communicating to non-nearest neighbor processors.

Figure 3. Parallel Energy Calculation for the Spatial Monte Carlo Decomposition.

a) A subvolume with domain D which contains subdomain S, and the neighboring 7 domains which are contained on subdomain S's 7 set I processors. (b) An exploded view of the domains and subdomains of the 7 set I processors for subdomain S. Subdomain faces which are shared with subdomain S are designated with a bold square, subdomain edges which are shared with subdomain S are designated with a bold line, and the corner which subdomain S shares with a subdomain on one of the 7 set I processors is designated with bold sphere.

Table 1. A Comparison Between PGCMC and NVT MD for State Point A

State Variable	PGCMC (* denotes an independent variable)	NVT MD (* denotes an independent variable)
$\frac{\mu'}{kT}$	-3.0*	-2.97
$\rho\sigma^3 = \frac{N}{V}$	0.278	0.278*
N	50,050.78	2048*
V	179,830.784*	7359.41*
$\frac{kT}{\varepsilon}$	3.0*	3.0*

$$\frac{\mu'}{kT} = \frac{\mu}{kT} + 3\ln\left(\frac{\sigma}{\Lambda}\right) \quad (\varepsilon \text{ and } \sigma \text{ from the Lennard-Jones model}^2).$$

Programs for the nCUBE 2 are implemented in FORTRAN, C or C++. All co-ordination or co-operation between processors, i.e., parallel processing is performed via explicit message passing calls. Message writes are non-blocking because the data is copied to a communication buffer and control is returned to the program. Message reads are blocking.

PGCMC is implemented in FORTRAN. Performance was optimized by attempting to overlap communications both with other communications and with computation. The parallel transfer capabilities of the DMA channels were exploited by issuing several of the non-blocking writes in a row and are handled by the operating system transparently to the user. By issuing the message writes as soon as possible and delaying the message reads with computation as long as possible, communication may sometimes be completely hidden behind computation. This technique was exploited in the energy calculations by first writing the test coordinates to all seven of the set I neighbors. Then the energy contributions of the processors own subdomains were calculated. The requests from the 7 set II processors were then read and worked upon, and results written back one at a time. Finally

the energy contributions of the other processors were read. This effectively used the DMA channels in parallel and hid the message writing of the test coordinates behind computation.

PGCMC uses nCUBE supplied routines for embedding a 3-D mesh in the hypercube in such a way that all communication between neighboring domains along the axis of the mesh and across periodic boundaries were on processors that were also directly connected nearest neighbors in the hypercube. This resulted in all communications to processors whose subdomains share a face being "1-hop" or nearest neighbor. All communications to processors whose subdomains shared just edges were 2-hop and to processors sharing just a corner were 3-hop. This minimized the potential contention for channels.

Table 2. State Points

State Point	$\frac{\mu'}{kT}$	$\frac{kT}{\epsilon}$	Resulting Density $\rho\sigma^3$
A	-3.0	3.0	0.28
B	-2.0	1.0	0.73

Scaling Characteristics of PGCMC

In order to investigate the scaling characteristics of **PGCMC**, timings of varying system sizes and number of processors for the state points in Table 2 were obtained. Each system, regardless of size, was "warmed-up" and the density checked to be sure the system was fully equilibrated. The number of "warm-up" steps necessary was a function of density. The system sizes and scaling results for the spatial decompositions for both state points, is shown in Table 3.

From the results in Table 3 we can see that the spatial GCMC algorithm scales nearly exactly with system size for both state points. However, we can see from Table 3 that the time per step (C_{step}) of the spatial decomposition does increase slightly with the number of processors. This is due simply to the fact that the interprocessor communication necessary to gather the thermodynamic data on processor 0 for print-out to the data file (done with logarithmic "fold" routines⁵), a very small part of the overall simulation time, increases with increasing numbers of processors.

Table 3. Scaling Results for the Spatial Decomposition (PGCMC)

Number of Processors P	System Dimensions			System Volume	State Point A		State Point B	
	$\frac{x}{\sigma}$	$\frac{y}{\sigma}$	$\frac{z}{\sigma}$	$\frac{V}{\sigma^3}$	Ave. No. of Atoms <<N>>	Elapsed Time sec/step	Ave. No of Atoms <<N>>	Elapsed Time sec/step
8	11.2	11.2	11.2	1404.928	392	0.041		
16	22.4	11.2	11.2	2809.856	783	0.042		
32	22.4	22.4	11.2	5619.712	1,566	0.043	4126	0.071
64	22.4	22.4	22.4	11239.424	3,131	0.043	8231	0.072
128	44.8	22.4	22.4	22478.848	6,254	0.044	16492	0.072
256	44.8	44.8	22.4	44957.696	12,519	0.044	32910	0.072
512	44.8	44.8	44.8	89915.392	25,043	0.045		
1024	89.6	44.8	44.8	179830.784	50,047	0.045		

Once the time per step, C_{step} , is known, it can be multiplied by a measure of the number of steps to reach a given value of the standard error to yield the amount of elapsed time to reach that value of the standard error, C_{SE} . For the purposes of demonstrating how the overall efficiency of the spatial decomposition scales with system size (and therefore numbers of processors) we have determined the speedup, S ,

$$S(P) = \frac{C_{SE}(P = 8)}{C_{SE}(P)}, \quad (1)$$

and the overall efficiency, E ,

$$E(P) = \frac{S}{P/8} \times 100\%, \quad (2)$$

as a function of the number of processors (divided by 8) for **PGCMC** at state point A (Table 4).

Table 4. Speedup and Efficiency Results for PGCMC for State Point A

Number of Processors P	Number of GCMC Steps	Number of Raw Data Points N_{pts}	Bin Size Used to Refine Raw Data b	Number of Raw Data Points to a Standard Error of 0.0008, $N_{pts}^{SE=0.0008}$	Elapsed Time to a Standard Error of 0.0008 (hours)	Speed-Up (eqn 1) S	Efficiency (eqn 2) E
8	3200K	6400K	12500	5862.5K	33.4	1.0	100%
16	2400K	4800K	12000	1792K	10.5	3.2	159%
32	2400K	4800K	12000	1088K	6.5	5.1	129%
64	2400K	4800K	8000	760K	4.5	7.4	93%
128	1920K	3840K	8000	288K	1.8	18.6	116%
256	960K	1920K	8000	248K	1.5	22.3	70%
512	400K	800K	8000	80K	0.50	66.8	104%
1024	480K	960K	8000	40K	0.25	133.6	104%

Note that the definitions of both the speedup and the efficiency are based on a comparison between results for P processors and 8 processors, as opposed to a single processor ($P = 1$), hence the use of $P/8$. This is due to the fact that **PGCMC** runs on a minimum of 8 processors. Furthermore, we can see from this figure that for some system sizes (numbers of processors) the **PGCMC** algorithm actually experiences superlinear speedup and thus efficiencies greater than 100%. These results are due to scatter in the data and thus the statistical nature of the evaluation of the speedup and efficiency.

Finally, while the spatial grand canonical Monte Carlo algorithm scales with system size and is thus well suited for studying large open systems of atoms with short-ranged interactions, for example fluids in mesopores, the overall efficiency of the another decomposition (the "time" decomposition) was shown to be superior to that of the spatial decomposition for small systems². However, for the work of the CRADA, the spatial decomposition is preferred because of the need for larger systems; the advantages in efficiency of the time decomposition were shown to be outweighed by its warm-up time and memory requirements for large systems².

Extension to Fluid-Pore Systems and Realistic Model Development

To design and implement a simulation geometry which is commensurate with the spatial parallel Monte Carlo algorithm and which would provide scientific results useful in interpreting physical oil/gas/water behavior in meso- and micro-porous rocks, several challenges were addressed:

1. The extension of the spatial parallel Monte Carlo algorithm to inhomogeneous systems.
2. The addition of "external fields".
3. The design and selection of a model pore geometry and interatomic potential parameters.
4. The design and implementation of a method to calculate the Pressure Tensor in the fluid.

Inhomogeneous Systems

The basic spatial parallel Monte Carlo algorithm was designed and tested on homogenous, isotropic systems. When implemented in certain systems with solid boundaries or with two fluid phases, it became apparent that correct solutions were not being obtained. An investigation led to the fact that the original scheme for choosing particles on which Monte Carlo moves or destructions would be attempted had failed whenever inhomogeneous fluid states occurred near subdomain boundaries. This was due to the selection being done with a uniform probability over the particles which were present in a given subdomain: If there were 10 particles in a subdomain, each might be selected 10% of the time for attempted destruction, while if there were 100 particles in a neighboring subdomain, each would only be selected about 1% of the time. This, of course, did not occur in homogeneous systems where, on average, the number of each type of particle was the same in each subdomain. The result for inhomogeneous systems was a biased Monte Carlo algorithm that did not produce the Markov chain necessary for creating averages in the desired Grand Canonical ensemble.

The problem was solved by selecting particles on the basis of the number of particles in the most populated subdomain in the system. If the subdomain containing the largest number of particles contained 200 particles, and the program was selecting a particle to destroy in a subdomain with only 10 particles, each particle would be selected with probability $10/200 = 0.05$, while 95% of the time, no destruction at all would be attempted in that subdomain. This solution entailed adding more communication overhead to the program, but the penalty was minor and a workable algorithm was found which produced correct results for both homogeneous and inhomogeneous systems.

External Fields

In order to be able to simulate and test a variety of types of pore system, we chose to implement a system whereby the walls of pores could be modeled either by individual atoms or by external fields. The individual atoms were further enabled to be either fixed or moveable and to interact with other atoms through any desired potential function.

One way to model a pore wall with this technique would be to simply build it up out of atoms, essentially no different than the fluid atoms, and then either fix those atoms in space or give them potentials such that they remained an essentially rigid wall. A more efficient way is to add a type of particle with a non-central potential interaction with the fluid atoms. The energy due to the interaction of this particle would be calculated in the normal way, whenever the particle was within a specified cut-off distance from the fluid particles, but the potential might not depend on this distance but rather some other parameter like the distance from a nearby wall. The "wall" particle simply represents the external field and produces the interaction energy in the normal course of the calculation for each fluid particle. The wall particles were placed at the center of each domain; thereby eliminating communication between processors during calculations of complicated 3D external fields. This was accomplished without undue memory burdens since the potential was calculated like any other potential in the system.

The result was a program for which it was extremely easy and quick to add or modify different pore shapes. For instance, we were able to quickly test the program by adding a cylindrical-pore potential and reproducing published results for one-component, two-phase equilibria. In addition, we reproduced published results for bulk two-component, two-phase systems.

Pore Geometry and Potential Parameters

The goals of the project could only be met by looking at different types of pore geometry; some which were efficiently implemented and easy to analyze and some of which were more faithful to the physical systems of interest. The cylindrical pore model mentioned above was selected for testing and benchmarking purposes. One of the problems with this system is that some processors are left idling since the spatial parallel Monte Carlo algorithm depends on Cartesian geometry and parallelepiped domains. Mapping a cylindrical pore into a square-sided box can leave much space outside of the pore and therefore the processors responsible for the unused areas are performing no useful work during part of the Monte Carlo cycle. To address this efficiency problem and also to ease analysis, we decided on a "wedge" pore system.

Consider the letter "X" in a box where the corners of the figure are coincident with the corners of a perspective view of the entire space to be modeled. The simulation system continues into the paper so that the X is drawn-out in a third dimension. The X can represent the walls of 4 wedge-shaped regions; left, right, top, and bottom. When combined with periodic boundary conditions, PBC, only two pores are modeled; the left/right pore and the

top/bottom pore. Particles moving out the right-hand side of the figure move back into the left-hand side and similarly for the top and bottom. If the walls of the system are made thin enough so that they don't take up much space, but are thick enough so that the two pore-systems don't interact (thicker than the fluid-fluid cut-off distance) we achieve a system where almost all of the processors covering the space are responsible for domains that are filled with fluid particles and which therefore spend most of their time doing useful calculations with reasonable load-balance. If the X is inscribed inside a rectangle, the two pores are geometrically identical (and which when simulated provide useful redundancy and improved statistics over just one pore) and each contains corners with two different angles. This is important since one of the chief aims of the project was to map out the dependence of capillary pressure on the radius of the fluid-fluid meniscus in the pore: the LaPlace equation.

We implemented the wedge geometry with fixed wall particles and were able to begin simulations on two- and three-phase, two-component fluids in these pores. The correct qualitative trends were obtained for the dependence of the system on the chemical potentials of the two components, on the potential parameters, and on the temperature. Due to the termination of the CRADA, no further work was done.

Pressure Tensor

To measure capillary pressure (the macroscopic parameter most useful and widely measured in physical systems), the pressure tensor of the confined multiphase fluid system was needed. The pressure tensor is not unambiguously defined for microscopic systems and we used this ambiguity to come up with a definition and an algorithm that would enable us to calculate the desired quantities. The most intuitive definition of the pressure tensor, due to Irving and Kirkwood, considers forces between particles to act along the vector directly between them. This definition, while intuitive and the most commonly used, is difficult to implement in the spatial decomposition of the GCMC algorithm. The pressure tensor contributions from particles within domains, say, A & B, may have to be allocated to a region of space on a third domain, C, because the interparticle vector may partially lie in each of the three domains. This would mean that communication would be needed with a processor which contains neither of the particles; a significant departure from the communication patterns used in the rest of the code. To overcome this difficulty we made two choices; to use an alternative definition of the pressure tensor and to have multiple processors responsible for certain "overlapping" subdomains.

The Harasima definition of the pressure tensor⁶ considers contributions to be allocated along a path which runs at right angles with respect to a surface (a fluid/fluid surface in previous work). We chose to use a Harasima-like definition for the entire space of our systems even though we were not going to choose a surface-plane to anchor the definition; we calculated pressure-tensor components along the three Cartesian coordinates. The three segments making up an interparticle path were first along x from particle 1's x -position to particle 2's x -position, then along y , and then along z . These segments were

much easier to deal with in the context of the square decomposition of space in the algorithm.

In addition to a new definition of pressure tensor, we expanded the range of space that processors were responsible for when calculating pressure tensor components. Each domain was associated with an extended domain which included one subdomain's width on either side (in all three dimensions) of the domain of interest. Pressure tensor contributions in these regions were summed by different processors and the final results summed in a communication which was necessary only at the end of a simulation run. The method was implemented and tested and for uniform systems and the correct bulk pressure was found.

Impact of the CRADA On Other Sandia Projects and Benefits to DOE

The Mobil-nCUBE-Sandia CRADA was synergistic with existing activities and core competencies at Sandia by design. It drew upon Sandia's computational science and computational facilities as well as Sandia's petroleum exploration and production research sponsored by DOE Energy Research. While the **PGCMC** code is not involved in any specific DOE supported work at Sandia at this moment, its parallel algorithm found immediate use in an important simulation tool, Dual Control Volume Grand Canonical Molecular Dynamics (DCV-GCMD), and was incorporated into the DCV-GCMD code, LADERA, immediately after its development.

The DCV-GCMD method⁷, as implemented in LADERA, enables the dynamic simulation of a system experiencing a steady-state chemical potential gradient. It thus enables the determination of the diffusivity in the presence of a chemical potential gradient, D_i^x , of each species i in a multicomponent mixture. This is accomplished by inserting two chemical potential control volumes in a standard NVT molecular dynamics simulation and using GCMC-like insertions and deletions in the control volumes to maintain constant chemical potential of each species. If the chemical potential of a given species is maintained at different values in the two control volumes, steady state chemical potential and density gradients are established. By measuring the gradient of the density (or concentration) profile, $(d\rho_i(x))/(dx)$, and flux, J_i^x , the diffusivity, D_i^x , of component i can be calculated via Fick's law, generally written as:

$$J_i^x = -D_i^x \frac{d}{dx} \rho_i(x). \quad (3)$$

Because the flux is calculated by counting the number of atoms of species j which move through a flux plane between the two control volumes, quality statistics are obtained only

if the system's cross sectional area is large. This dictates the need for a large system and thus the use of massively parallel hardware.

LADERA, the parallel implementation of the DCV-GCMD method, was developed with Laboratory Directed Research & Development (LDRD) funding and has subsequently found application to several important DOE Energy Research and Defense Program problems. The code has been applied to diffusion in zeolites and silicalites^{8,9} as part of an ER funded effort to develop new ceramic membranes and separations technology. An additional LDRD was also established to extend the method to bonded systems. This will enable the application of the DCV-GCMD method to problems such as the diffusion of drugs through biomembranes and diffusion-limited oxidation of polymers. The latter has been identified as one of the most important materials aging problem for the enduring stockpiles due to its effect on the elastomeric materials used in o-rings.

Finally, several new algorithms were developed in the course of this work. The **PGCMC** method was extended and rendered easier to adapt to a wide variety of physical systems and fluid types. The Harasima definition of the pressure tensor was modified, the spatial decomposition of the GCMC algorithm was extended to inhomogeneous systems, and arbitrary types of particles and potentials were added, increasing the scope of physical systems able to be modeled. All of these were achieved in a parallel scheme without significant change to the scaling properties of the final program.

References

- ¹Heffelfinger, G.S. and Lewitt, M.E., "A Spatial Decomposition for Massively Parallel Grand Canonical Monte Carlo Computer Simulation," American Institute of Chemical Engineers (AIChE) Annual Meeting, San Francisco CA, 13-18 November, 1994.
- ²Heffelfinger, G.S., and Lewitt, Martin E., "A Comparison Between Two Massively Parallel Algorithm for Monte Carlo Computer Simulation--An Investigation in the Grand Canonical Ensemble," *Journal of Computational Chemistry*, **17**, 250, (1996).
- ³Widom, B., *Journal of Chemical Physics*, **39**, 2808 (1963).
- ⁴Allen, M.P. and Tildesley, D.J., *Computer Simulations of Liquids* (Clarendon Press, Oxford, 1987).
- ⁵Fox, G., Johnson, M., Lyzenga, G., Otto, S., Salmon, J., and Walker, D., *Solving Problems on Concurrent Processors, Volume I* (Prentice Hall, Englewood Cliffs, New Jersey, 1988).
- ⁶Harasima, A., *J. Phys. Soc. Japan*, **8**, 343 (1953).
- ⁷Heffelfinger, G.S., van Swol, F., "Diffusion in Lennard-Jones Fluids Using Dual Control Volume Grand Canonical Molecular Dynamics Simulation (DCV-GCMD)," *J. Chem. Phys.* **100**, 1 (1994).

⁸Pohl, P.I., Heffelfinger, G.S., Smith, D.M., "Molecular Dynamics Computer Simulation of Gas Permeation in Thin Silicalite Membranes," in preparation, (1995).

⁹Heffelfinger, G.S., Pohl, P.I., and Frink, L.J.D., "Molecular Dynamics Computer Simulations of Diffusion in Porous Silicates," Mat. Res. Soc. Symp. Proc., **366**, 225 (1995).

APPENDIX

Quarterly Report Dated 9/95

CRADA SURVEY FOR QUARTER

4 /1995

19-Sep-95

To:

From: Grant S Heffelfinger, 1421, MS-1111

Subject: SANDIA QUARTERLY PROJECT REPORT SURVEY

I have attached the relevant parts on all 3 of the ascii text Quarterly Report forms:

Base CRADA Information

Project Highlights: We are about to present our first scientific results at the Annual Meeting of the American Society of Chemical Engineers in Miami, FL. in November.

Briefly describe the progress of your project (modify or replace highlights from last quarter): Our progress since last quarter has been minimal.

Project Deviations: There are 2 primary reasons for the deviations in the project. First, Mobil Research and Development Corporation has been convulsing for the last 6 months. They have totally shaken down the entire company, including a name change (Mobil Technology Corporation) and liquidated the Central Research Laboratory (where our collaborators were located). This has had a profound effect on the project although we are still working hard and beginning to get back on track. Secondly, we've had some difficulties implementing the needed pressure tensor in the parallel GCMC code. We should have this finished in a few weeks.

Briefly describe any cost, schedule or performance deviations. This narrative must match the progress indicated by the completed tasks and milestones (modify or replace last quarter's deviations): We are about 2 months behind on Task 4 now but should have it done by the conference in November. This does not present a problem to the overall CRADA schedule but rather will just give Mobil a little less time to apply the code with our assistance. In the meantime, we've started thinking about Task 2, due Jan 1st.

Select the statement which best describes partner participation (if participation is not in kind with that agreed to in the Project briefly explain):

_____ Partner In Kind participation IS consistent with that agreed to in the CRADA.

X Partner In Kind participation IS NOT consistent with that agreed to in the CRADA.

Comments: This has been the situation for the last few months while Mobil has undergone an extensive restructuring. However, things seem to be settling down now and we are making good progress again.

TASKS AND MILESTONES SURVEY FOR QUARTER 4/ 1995

19-Sep-95

ALL SURVEYS MUST BE RETURNED 10/2/95

TASKS AND MILESTONES: NO CHANGE x CHANGED

CRADA: 1140.0

TASKS 1a and 2a: Identify appropriate models for wetting and nonwetting components in an ideal pore.

DUE: 8/1/93 START: 6/93 COMPLETE: 8/1/93

TASK 1b: Develop Grand Canonical Monte Carlo (GCMC) codes for bulk and confined fluid systems.

DUE: 1/1/94 START: COMPLETE: 1/1/94

TASK 1c: Develop and implement parallel algorithms for bulk and confined fluid GCMC codes.

DUE: 10/1/94 START: COMPLETE: 9/1/94

TASK 1d: Calculate the surface tension for the bulk binary system while concurrently determining the capillary pressure for the confined binary fluid at varying chemical potentials.

DUE: 7/1/95 START: COMPLETE:

TASK 2b: Investigate possible methods and algorithms for modelling pressure induced flow in fine capillaries for high density fluids. Once appropriate methods and algorithms have been identified, develop massively parallel code for the algorithms.

DUE: 1/1/96 START: COMPLETE:

TASK 1e: Using these results, establish the region of validity for the Laplace equation for capillary pressure and refine the Laplace equation for the capillary pressure and refine the Laplace equation for microscopic pores.

DUE: 1/1/96

START:

COMPLETE:

TASK 2c: Investigate region of validity of Darcy's Law and establish reasons for failure.
Develop new/refined Darcy-type models.

DUE: 7/1/96

START:

COMPLETE: