



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-TR-746794

High Performance Parallel Processing (HPPP) Shallow Junction Devices Final Report CRADA No. TC-824-94B2

T. Diaz delaRubia, T. Keeler

February 23, 2018

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed 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 or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

DISCLAIMER

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

High Performance Parallel Processing (HPPP)

Shallow Junction Devices

Final Report

CRADA No. TC-824-94B2

Date: September 15, 1998

Revision: 2

A. Parties

The project is a relationship between the Lawrence Livermore National Laboratory (LLNL) and Cray Research Inc. (CRI).

University of California
Lawrence Livermore National Laboratory
PO Box 808, L-795
Livermore, CA 94550

Cray Research Inc. (CRI)
655 Lone Oak Drive
Eagan, MN 55121

B. Project Scope

Massively parallel computers, such as the Cray T3D, have historically supported resource sharing solely with space sharing. In that method, multiple problems are solved by executing them on distinct processors. This project developed a dynamic time- and space-sharing scheduler to achieve greater interactivity and throughput than could be achieved with space-sharing alone. Cray Research Inc. (CRI) and LLNL worked together on the design, testing, and review aspects of this project. There were separate software deliverables. CRI implemented a general purpose scheduling system as per the design specifications. LLNL ported the local gang scheduler software to the LLNL Cray T3D. In this approach, processors are allocated simultaneously to all components of a parallel program (in a "gang"). Program execution is preempted as needed to provide for interactivity. Programs are also relocated to different processors as needed to efficiently pack the computer's torus of processors.

In phase one, CRI developed an interface specification after discussions with LLNL for system level software supporting a time- and space-sharing environment on the LLNL T3D. The two parties also discussed interface specifications for external control tools (such as scheduling policy tools, system administration tools) and applications programs. CRI assumed responsibility for the writing and implementation of all the necessary system software in this phase.

In phase two, CRI implemented job-rolling on the Cray T3D, a mechanism for preempting a program, saving its state to disk, and later restoring its state to memory for continued execution. LLNL ported its gang scheduler to the LLNL T3D utilizing the CRI interface implemented in phases one and two.

Binary collision Monte Carlo codes, interfaced to the MD kernel, were developed by AT&T and were optimized with CRI on the T3D. Interatomic potentials for silicon-silicon and silicon-dopant interactions based on the tight binding approximation were developed by LLNL as well. A tight binding parallel molecular dynamics code was implemented on the T3D by LLNL.

During phase three, the functionality and effectiveness of the LLNL gang scheduler was assessed to provide input to CRI time- and space-sharing efforts. CRI will utilize this information in the development of general schedulers suitable for other sites and future architectures.

The full ion implantation simulator was tested by AT&T and LLNL, and the results implemented the processing of representative device structures by AT&T.

All phases of this project were completed on time and all deliverables were met without significant changes to the original statement of work.

C. Technical

Cray Research was able to provide a highly effective job-rolling mechanism. This mechanism also provided the flexibility of restoring the program's state to different processors than were initially utilized.

LLNL completed the initial installation of the gang scheduler in March of 1996. The next few months were spent perfecting the scheduling algorithms and tuning. Early simulations indicated that interactivity could be improved dramatically, and the computer's saturation point could also be increased. Since LLNL's Cray T3D was originally configured for a high level of interactivity with moderate throughput, we were able to reconfigure the computer for dramatically higher throughput. The gang scheduler was able to provide even greater interactivity while utilization was increased from the 30 percent range to the 90 percent range — a phenomenal level of throughput for a massively parallel computer.

Issues outside of the scope of this CRADA, but of interest for further study include: the comparison of job-paging rather than rolling, gang-scheduling on distributed memory computer architectures, and gang-scheduling on computers architectures in which greater flexibility in processor assignment exists.

The primary technical objective of this CRADA was to develop a physics-based simulator for ion implantation modeling of semiconductor processing. Towards that end, LLNL and AT&T developed molecular dynamics and Monte Carlo kernels that were implemented with the assistance of CRI on the T3D machine. In

addition, accurate ab-initio interparticle potentials for silicon-silicon and silicon-dopant interactions were developed and implemented into the codes.

The milestones in this program are highlighted below:

<u>Task Description</u>	<u>Month</u>
<u>LLNL/AT&T/CRI Common Milestones</u>	
- port classical molecular dynamics (MD) simulation code to T3D massively parallel platform	0- 6
- optimize classical molecular dynamics (MD) simulation code for T3D massively parallel platform	0-12
- develop analysis and visualization codes for data manipulation and parameter extraction	0-12
- apply computer codes in direct simulation	12-36
- develop new Energetic Beam Solid Interaction Simulator (EBSIS) on T3D. This code will incorporate results of the classical and quantum mechanical MD simulations	12-36
- documentation for simulation codes developed for the T3D	30-36
<u>LLNL/CRI Milestones</u>	
- develop and implement ab-initio and quantum chemical tight binding MD methods on T3D.	6-30
- develop ab-initio interparticle potentials for silicon-silicon and silicon-dopant interactions.	6-30
<u>AT&T Milestones</u>	
- test results of process modeling codes for representative device structures	12-36

D. Expected Economic Impact

This advanced atomistic modeling approach represents a paradigm shift. Application of the simulators will reduce device development time and cost.

E. Partner Contribution

All deliverables have been met. CRI developed a number of documents and a highly effective job-rolling mechanism. The LLNL gang scheduler has proven to be highly effective at improving system throughput and interactivity. In parallel with the development of LLNL's gang scheduler, CRI developed a gang scheduler based upon the Dynamic Job Manager, which has been distributed by CRI. Some U.S.

government agencies have installed LLNL's gang scheduler on their Cray T3D computers.

F. Documents/Reference List

1) Reports

"Improved Utilization and Responsiveness with Gang Scheduling," Dror G. Feitelson and Morris A. Jette, *Job Scheduling Strategies for Parallel Processing Workshop*, (publication pending) April 1 997.

"The Gang Scheduler—Timesharing on a Cray T3D," Morris Jette, David Storch, and Emily Yim, *Cray User Group Meeting*, March 1996.

Substantial documentation available on the internet at "<http://www-1c.llnl.gov/dctg/gang>".

Proceedings of the second International Conference on Computer Simulation of Radiation Effects in Solids. Edited by T. Diaz de la Rubia, George H. Gilmer and M. Caturla. *Nuclear Instruments and Methods in Physics Research B*, Vol 102, 1995.

T. Diaz de la Rubia and G.H. Gilmer, "Structural transformations and Defect Production in Ion Implanted Silicon: A Molecular Dynamics Simulation Study". *Phys. Rev. Lett.* **74**, 2507 (1995)

G.H. Gilmer, T. Diaz de la Rubia, D.M. Stock, and M. Jaraiz. "Diffusion and Interaction of Point Defects in Silicon: Molecular Dynamics Simulations" *Nucl. Instr. Meth. B* **102**, 247 (1995)

M. Jaraiz, G.H. Gilmer, J.M. Poate, and T. Diaz de la Rubia. "Atomistic Calculations of Ion Implantation in Si: Point Defect and Transient Enhanced Diffusion Phenomena" *Appl. Phys. Lett.* **68**, 409 (1996).

J. Zhu, T. Diaz de la Rubia, L. Yang G.H. Gilmer and C. Mailhot, "Ab initio Pseudopotential Calculations of B Diffusion and Pairing in Silicon", *Phys. Rev. B* **54** (7), 4741 (1996).

M.J. Caturla, T. Diaz de la Rubia, L. Marques, and G.H. Gilmer: "Ion Beam Processing of Silicon at keV Energies" *Phys. Rev. B* **54**, 16683 (1996).

L. Marques, M.J. Caturla, T. Diaz de la Rubia and G.H. Gilmer. "Ion Beam Induced Recrystallization of Amorphous Silicon: A Molecular Dynamics Study" *J. Appl. Phys.* **80**, 6160 (1996).

2) Intellectual Property**i) Subject Inventions**

Art. I: "Subject Invention" means any invention of The Regents or Participant conceived or first actually reduced to practice in the performance of work under this CRADA.

Art. XIV: The Parties agree to disclose to each other each and every Subject Invention, which may be patentable or otherwise protectable under the Patent Act.

LLNL Sole Subject Inventions: None

AT&T Sole Subject Inventions: None disclosed

Joint Subject Inventions: None

ii) Computer Software:

CRADA Article XIII requires that for all copyrighted computer software produced in the performance of this CRADA, the Party owning the copyright will provide the source code, an expanded abstract, and the object code and the minimum support documentation needed by a competent user to understand and use the software to DOE's Energy Science and Technology Software Center.

Computer software developed by LLNL:

Previously developed gang scheduler software was ported to the LLNL Cray T3D.

Computer software developed by AT&T:

Binary collision Monte Carlo codes

Computer software developed jointly by LLNL and AT&T:

Molecular Dynamics and Monte Carlo kernels

Copyrighted computer software:

AT&T did not disclose its assertion of copyright on any software developed under this CRADA.

iii) Licensing activity:

Appendix C, Intellectual Property Agreement, provides that to the extent that The Regents obtains title or authority to license Intellectual Property first arising or produced under this CRADA, The Regents shall negotiate in good faith with AT&T Bell Laboratories for a license to rights in such Intellectual Property, on reasonable commercial terms, for the term of this CRADA plus a period of not more than six months after the completion or termination of this CRADA.

AT&T Bell Laboratories has not indicated an interest in licensing any LLNL Intellectual Property developed under this CRADA.

G. Acknowledgment

Participant's signature of the final report indicates the following:

- 1) The Participant has reviewed the final report and concurs with the statements made therein.
- 2) The Participant agrees that any modifications or changes from the initial proposal were discussed and agreed to during the term of the project.
- 3) The Participant certifies that:
 - a) all reports either completed or in process are listed;
 - b) all subject inventions attributable to the project have been disclosed or are included on a list attached to this report; and
 - c) appropriate measures have been taken to protect intellectual property attributable to this project.
- 4) The Participant certifies that if tangible personal property was exchanged during the agreement, all has either been returned to the initial custodian or transferred permanently.
- 5) The Participant certifies that proprietary information has been returned or destroyed by LLNL.

Tim Keeler
Cray Research Inc.

Date

 Tomas Diaz-De La Rubia
Lawrence Livermore National Laboratory

Date

Attachment I - Final Abstract
Attachment II - Project Accomplishments Summary
Attachment III -Final Quarterly Report

High Performance Parallel Processing (HPPP)

Shallow Junction Devices

Final Abstract
Attachment I
CRADA No. TC-824-94B2

Date: August 26, 1998

Revision: 1

Massively parallel computers, such as the Cray T3D, have historically supported resource sharing solely with space sharing. In that method, multiple problems are solved by executing them on distinct processors. This project developed a dynamic time- and space-sharing scheduler to achieve greater interactivity and throughput than could be achieved with space-sharing alone. CRI and LLNL worked together on the design, testing, and review aspects of this project.

This project was a multi-partner CRADA with nine industrial partners. During this portion of the CRADA, AT&T and Cray enabled predictive modeling of shallow junction formation for next generation semiconductor device fabrication.

There were separate software deliverables. CRI implemented a general purpose scheduling system as per the design specifications. LLNL ported the local gang scheduler software to the LLNL Cray T3D. In this approach, processors are allocated simultaneously to all components of a parallel program (in a "gang"). Program execution is preempted as needed to provide for interactivity. Programs are also relocated to different processors as needed to efficiently pack the computer's torus of processors.

High Performance Parallel Processing (HPPP)

Shallow Junction Devices

Project Accomplishments Summary (Attachment II)

CRADA No. TC-824-94B2

Date: August 26, 1998

Revision: 1

A. Parties

The project is a relationship between the Lawrence Livermore National Laboratory (LLNL) and Cray Research Inc. (CRI).

University of California
Lawrence Livermore National Laboratory
PO Box 808, L-795
Livermore, CA 94550

Cray Research Inc. (CRI)
655 Lone Oak Drive
Eagan, MN 55121

B. Background

Massively parallel computers, such as the Cray T3D, have historically supported resource sharing with space sharing only. In that method, multiple problems are solved by executing them on distinct processors, and once a problem begins execution, it cannot be interrupted. When a heavy load exists on such a computer, initialization of jobs may be substantially delayed until other jobs terminate. This paradigm makes interactivity and/or throughput inherently poor. LLNL had experience developing a gang scheduler for the BBN TC2000 massively parallel computer, which successfully addressed these problems.

C. Description

This project developed a dynamic time- and space-sharing scheduler to achieve greater interactivity and throughput than could be achieved with space-sharing alone. The LLNL Cray T3D was initially configured for a high level of interactivity during working hours, frequently resulting in very poor throughput. When a very heavy interactive load existed, very long delays could be experienced for job initialization. We were also forced to severely limit job size and run times to limit the adverse impact a single job could have. Our objective was to develop a dynamic time- and space-sharing scheduler to resolve all of these problems, providing excellent throughput, excellent interactivity, and permitting the execution of larger and longer running jobs.

Cray Research Inc. and LLNL worked together on the design, testing, and review aspects of this project. There were separate software deliverables for this project. CRI implemented a general purpose mechanism for the preemption of jobs, saving the

job's state to disk, and later reloading its state into memory for continued execution. LLNL ported the local gang scheduler software to the LLNL Cray T3D. In this approach, processors are allocated simultaneously to all components of a parallel program (in a "gang"). Program execution is preempted as needed to provide for interactivity. Programs are also relocated to different processors as needed to efficiently pack the computer's torus of processors.

LLNL completed the initial installation of the gang scheduler in March of 1996. Since LLNL's Cray T3D was originally configured for a high level of interactivity with moderate throughput, we were able to reconfigure the computer for dramatically higher throughput. The gang scheduler was able to provide even greater interactivity while utilization was increased from the 30 percent range to the 90 percent range — a phenomenal level of throughput for a massively parallel computer.

LLNL and CRI were responsible for developing the parallel molecular dynamics codes and porting them to the T3D. AT&T implemented the Monte Carlo interface to the molecular dynamics algorithm to study the collisional aspects of the problem and tested the process modeling codes for representative device structures.

D. Expected Economic Impact

The utilization of the LLNL Cray T3D increased by a factor of three. With computers of this size costing tens of millions of dollars, the benefit to taxpayers has been enormous. The LLNL gang scheduler has been provided to other U. S. government agencies with Cray T3D computers. Like LLNL, these sites are expected to enjoy a substantial increase in throughput and interactivity in their Cray T3D supercomputers with minimal expense for the installation of this software.

CRI has also developed its own dynamic time- and space-sharing scheduler, the Dynamic Job Manager (DJM). The development of the DJM has benefited from much of this CRADA's work and has been installed by CRI on a number of their Cray T3D computers leading to significant performance improvements.

This project was a multi-partner CRADA with nine industrial partners. During this portion of the CRADA, AT&T and Cray enabled predictive modeling of shallow junction formation for next generation semiconductor device fabrication.

E. Benefits to DOE

The Cray T3D located at LLNL was able to achieve a dramatic improvement in throughput and interactivity. In the months following the introduction of the gang scheduler, average monthly utilization increased from the 30 percent range to the 90 percent range while interactivity also increased. LLNL was able to triple the effective use of its Cray T3D without additional hardware cost and with minimal software development costs. Problems using up to 40 hours and 64 processors are executed on a routine basis. Problems have been executed with up to 256 processors (the entire LLNL T3D) and 40 hours of execution time. Problems of this size could not realistically be solved formerly.

F. Industry Area

This project demonstrated that time- and space-sharing of massively parallel computers is not only possible, but highly advantageous. This work benefits a wide number of industries in which massively parallel computers are used, including: aircraft, automotive, medical, etc. In part due to our success, several other companies are currently pursuing gang schedulers for their parallel computers

G. Project Status

This project is completed.

H. LLNL Point of Contact for Project Information

Tomas Diaz De La Rubia

Lawrence Livermore National Laboratory

PO Box 808, L-353

Livermore, CA 94550

925-422-6714

FAX: 925-422-7300

Company Size and Point(s) of Contact

AT&T has 128,000 employees. Primary contact for this project is G. H. Gilmer.

Dr. George H. Gilmer

Lucent Technologies, Bell Laboratories

Murray Hill, NJ 07974

908-582-5547

Cray Research has 2,100 employees. Primary contact for this project is Kevin Lind.

Mr. Tim Keeler

Cray Research Inc.

655 Lone Oak Drive

Eagan, MN 55121

612-452-6650

FAX: 612-683-5310

J. Project Examples

There are no project examples.

K. Intellectual Property

i) Subject Inventions

Art. I: "Subject Invention" means any invention of The Regents or Participant conceived or first actually reduced to practice in the performance of work under this CRADA.

Art. XIV: The Parties agree to disclose to each other each and every Subject Invention, which may be patentable or otherwise protectable under the Patent Act.

LLNL Sole Subject Inventions: None

Cray Research Sole Subject Inventions: None disclosed

Joint Subject Inventions: None

ii) Computer Software:

CRADA Article XIII requires that for all copyrighted computer software produced in the performance of this CRADA, the Party owning the copyright will provide the source code, an expanded abstract, and the object code and the minimum support documentation needed by a competent user to understand and use the software to DOE's Energy Science and Technology Software Center.

Computer software developed by LLNL:

Previously developed gang scheduler software was ported to the LLNL Cray T3D.

Computer software developed by AT&T:

Binary collision Monte Carlo codes

Computer software developed jointly by LLNL and AT&T

Molecular Dynamics and Monte Carlo kernels

Copyrighted computer software:

AT&T Corporation did not disclose its assertion of copyright on any software developed under this CRADA.

iii) Licensing activity:

Appendix C, Intellectual Property Agreement, provides that to the extent that The Regents obtains title or authority to license Intellectual Property first arising or produced under this CRADA, The Regents shall negotiate in good faith with AT&T Bell Laboratories for a license to rights in such Intellectual Property, on reasonable commercial terms, for the term of this CRADA plus a period of not more than six months after the completion or termination of this CRADA.

Cray Research Inc. has not indicated an interest in licensing any LLNL Intellectual Property developed under this CRADA.

L. Release of Information

I certify that all information contained in this report is accurate and releasable to the best of my knowledge.

Karena McKinley
Karena McKinley, Director
Industrial Partnerships and Commercialization

11/15/08

Release of Information

I have reviewed the attached Project Accomplishment Summary prepared by Lawrence Livermore National Laboratory and agree that the information about our CRADA may be released for external distribution.

Tim Keeler
Cray Research, Inc.

Date

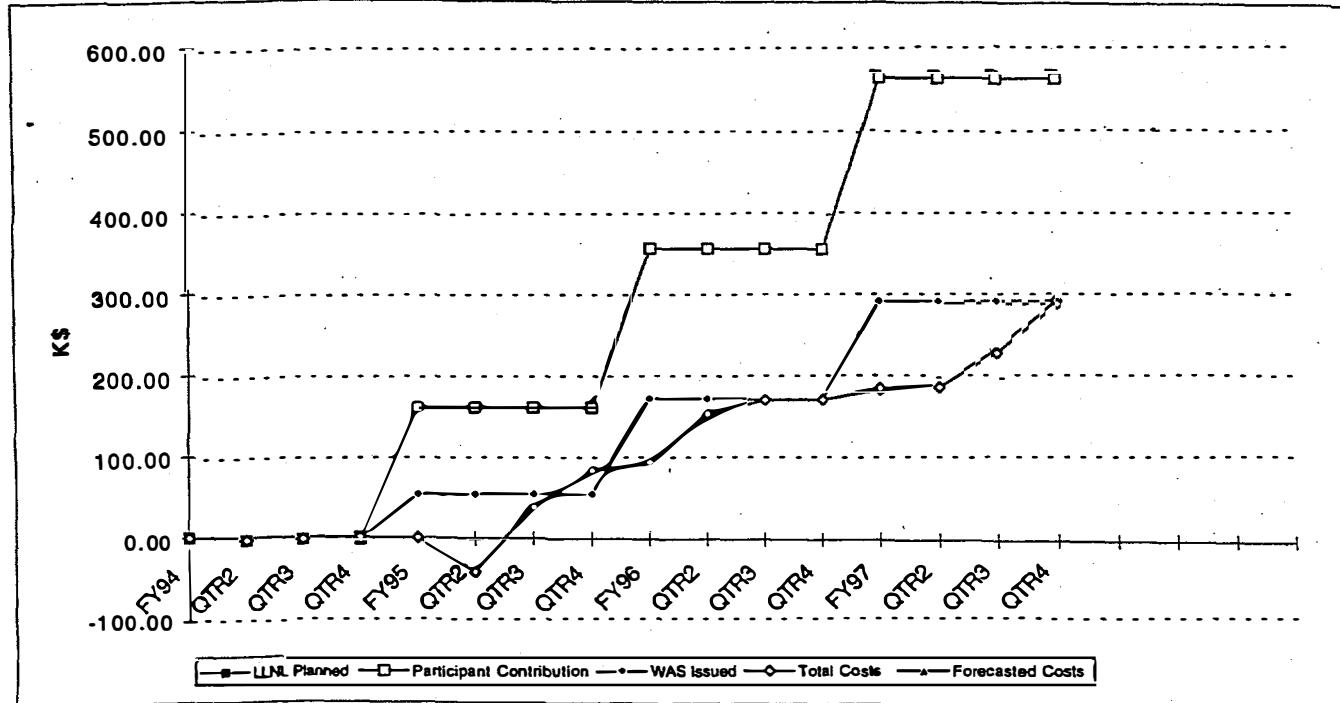
Lawrence Livermore National Laboratory

Title: HPPP Shallow Junction Devices
 Participant: AT&T
 DOE TTI No.: 94-MULT-003-XX-1
 CRADA No.: TC-0824-94 (B)
 Account Numbers: 4745-81, 91
 Accounts Closed: N/A

Reporting Period: 07/01/96 - 09/30/97
 Date CRADA Executed: 1/1/95
 DOE Approval Date: 12/21/94
 Scheduled Ending Date: 12/31/97
 Project Completion Date: N/A
 B & R Code (S): DP0301, YN01000

Approved Funding Profile (\$K)

	FY94	FY95	FY96	FY97	FYOUT	Total
LLNL Planned	0	162	195	209	0	566
Participant In-Kind	0	162	195	209	0	566
Participant Funds-ir	0	0	0	0	0	0
WAS DP0301	0	55	118	118	0	291
LDRD Funds	0	107	0	0	0	107
Total Costs	0	85	88	118	0	291



DP0301	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	July	Aug	Sep	FYTD
FY94	0	0	0	0	0	0	0	0	0	0	0	0	0
FY95	0	0	0	22	25	-89	30	21	31	17	23	6	85
FY96	-44	32	21	31	13	16	26	-14	6	1	0	0	88
FY97	0	12	0	0	3	0	11	0	31	31	15	14	118
FYOUT	0	0	0	0	0	0	0	0	0	0	0	0	0

291

YN01000	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	July	Aug	Sep	FYTD
FY94	0	0	0	0	0	0	0	0	0	0	0	0	0
FY95	0	0	0	0	0	0	0	0	0	0	0	0	0
FY96	0	0	0	0	0	0	0	0	0	0	0	0	0
FY97	0	0	0	0	0	0	0	0	0	0	0	0	0
FYOUT	0	0	0	0	0	0	0	0	0	0	0	0	0

0

STAFF w/phone:

Lab PI: Tomas Diaz de la Rubia (510) 422-6714
 Resource Manager: Steve Stinson (510) 423-2888
 DOE OAK: Jerry Scheinberg (510) 637-1653

Participant: George Gilner (908) 582-5547

DOE HQ: Alex Larzelere (202) 586-1101

Lawrence Livermore National Laboratory

Reporting Period : 07/01/96 - 09/30/97
DOE TTI No.: 94-MULT-003-XX-1
CRADA No.: TC-0824-94 (B)

Page 2

Milestones and Deliverables:

List the complete set of milestones for all phases of the CRADA. Continue on a separate page if necessary.
Report any changes from the original CRADA or previous quarterly report on the CRADA Change Form.

Completion Date:
Scheduled Actual

1
2
3
4
5
6
7
8

Verification of participants' in-kind contribution was made in accordance with LLNL policy. Explain basis of verification:

Please initial: YES NO

List any subject inventions by either party (include IL# for LLNL inventions), additional background intellectual property, patents applied for, software copyrights, publications, awards, licenses granted or reportable economic impacts

Accomplishments

Describe Technical/Non-Technical lessons learned (address and be specific about milestones, participant contributions)
Summarize causes/justification of deviations from original scope of work. Continue on a separate page if necessary.

Reviewed by CRADA project Program Manager:

Date:

Reviewed by Karena McKinley, Director, LLNL/IP&C:
Direct questions regarding this Report to IP&C Resource Manager, Carol Asher, at (925) 422-7618

Date:

Karena McKinley Date: 11/15/00