REPORT DOCUMENTATION PAGE

Form Approved OMB NO. 0704-0188

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)
03-02-2020	Final Report	12-Sep-2018 - 11-Sep-2019

4. TITLE AND SUBTITLE	
Final Report: Heterogeneous Computer System for Code Development, Large-Scale Simulations, and Data Post-Processing	5a. CONTRACT NUMBER W911NF-18-1-0454
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER 611103
6. AUTHORS	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES	8. PERFORMING ORGANIZATION REPORT NUMBER
New Mexico State University PO Box 30002, MSC OGC Anderson Hall E1200, Espina and Stewart Streets Las Cruces, NM 88003 -8002	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES)	10. SPONSOR/MONITOR'S ACRONYM(S) ARO
U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211	11. SPONSOR/MONITOR'S REPORT NUMBER(S) 72039-EG-RIP.1

12. DISTRIBUTION AVAILIBILITY STATEMENT

Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

The views, opinions and/or findings contained in this report are those of the author(s) and should not contrued as an official Department of the Army position, policy or decision, unless so designated by other documentation.

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Andreas Gross
UU	UU	UU	UU		19b. TELEPHONE NUMBER 575-646-6179

Standard Form 298 (Rev 8/98)
Prescribed by ANSI Std. Z39.18

RPPR Final Report
as of 12-Mar-2020

Agency Code:

Proposal Number: 72039EGRIP **Agreement Number: W911NF-18-1-0454**

INVESTIGATOR(S):

Name: Ph.D. Andreas Gross
Email: agross@nmsu.edu
Phone Number: 5756466179
Principal: Y

Name: Diana V. Dugas
Email: dugasdvt@nmsu.edu
Phone Number: 5756466498
Principal: N

Name: Carlo M. Henderson
Email: matt117@nmsu.edu
Phone Number: 5756465077
Principal: N

Organization: **New Mexico State University**
Address: PO Box 30002, MSC OGC, Las Cruces, NM 880038002
Country: USA
DUNS Number: 173851965 EIN: 856000401

Report Date: 11-Dec-2019 Date Received: 03-Feb-2020
Final Report for Period Beginning 12-Sep-2018 and Ending 11-Sep-2019
Title: Heterogeneous Computer System for Code Development, Large-Scale Simulations, and Data Post-Processing
Begin Performance Period: 12-Sep-2018 **End Performance Period:** 11-Sep-2019
Report Term: 0-Other
Submitted By: Ph.D. Andreas Gross Email: agross@nmsu.edu
Phone: (575) 646-6179
Distribution Statement: 1-Approved for public release; distribution is unlimited.

STEM Degrees: 0 **STEM Participants:** 5

Major Goals: The major goals of the project were to
(1) decide on a heterogeneous compute architecture that
(1.a) mirrors those at the DoD HPC centers to allow for the development and testing of code that can be ported to very-large-scale heterogeneous HPC systems
(1.b) provides considerable local compute power for the CFD lab and other researchers
(1.c) allows for compute intensive data post-processing
(1.d) offers a state-of-the-art environment for training students in all aspects of supercomputing
(2) to purchase and install the system
(3) to make the system operational (software installs etc.) and to demonstrate that it meets the requirements.

Accomplishments: All goals were met.
Together with the NMSU Information & Communication Technologies (ICT) Cyber Infrastructure Architect Team, the DURIP system was designed, purchased, and installed.
User data were collected over several months that demonstrate that the system is well utilized.
The system has been used for the training of five graduate students.

Training Opportunities: Since 2019 the new system has been used for the training of five graduate students:
One Ph.D. student is using the system for developing OpenCL-based subroutines for accelerating the compressible CFD code. OpenCL is a framework for writing programs that execute across heterogeneous platforms consisting of central processing units (CPUs), graphics processing units (GPUs), digital signal processors (DSPs), field-programmable gate arrays (FPGAs) and other processors or hardware accelerators.
Two M.S. and two Ph.D. students are using the system for CFD simulations. The M.S. students which don't have access to DoD HPC systems are also learning how to submit and run compute jobs on HPC-like computer systems and how to post-process data.

Results Dissemination: Nothing to Report

Honors and Awards: Nothing to Report

Protocol Activity Status:

Technology Transfer: Nothing to Report

PARTICIPANTS:

Participant Type: PD/PI
Participant: Andreas Gross
Person Months Worked: 1.00 **Funding Support:**
Project Contribution:
International Collaboration:
International Travel:
National Academy Member: N
Other Collaborators:

Participant Type: Co PD/PI
Participant: Diana Dugas
Person Months Worked: 1.00 **Funding Support:**
Project Contribution:
International Collaboration:
International Travel:
National Academy Member: N
Other Collaborators:

Final Report

Heterogeneous Computer System for Code Development, Large-Scale Simulations, and Data Post-Processing

U.S. Army Research Office
Contract Number: W911NF-18-1-0454

Program manager: Dr. Matthew Munson

A. Gross
Mechanical and Aerospace Engineering Department
and
D. V. Dugas
Information & Communication Technologies

New Mexico State University
Las Cruces, NM 88003

January 31, 2020

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*
01/31/2020

2. REPORT TYPE
Final Report

3. DATES COVERED *(From - To)*
10/12/18-10/11/19

4. TITLE AND SUBTITLE

Heterogeneous Computer System for Code Development, Large-Scale Simulations, and Data Post-Processing

5a. CONTRACT NUMBER
W911NF-18-1-0454

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

6. AUTHOR(S)

A. Gross, D. V. Dugas

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

New Mexico State University
1050 Stewart Street
Las Cruces, NM 88003

8. PERFORMING ORGANIZATION REPORT NUMBER
N/A

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Dr. Matthew Munson
Army Research Office
800 Park Office Drive
Suite 4229
Research Triangle Park, NC 27709

10. SPONSOR/MONITOR'S ACRONYM(S)
ARO

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT

DISTRIBUTION STATEMENT A: Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT
The combination of CPUs, Graphics Processing Units (GPUs) and other purpose-designed hardware into heterogeneous systems enables affordable high performance computing. Novel simulation codes and algorithms have to be developed to unleash the compute power inherent in these new hardware architectures. A heterogeneous system was purchased and installed at New Mexico State University. This system allows researchers and graduate students to develop highly effective and efficient codes for heterogeneous systems. Equally important, the system significantly enhances the local supercomputing research infrastructure. The system also constitutes an ideal education and training environment for graduate students in all aspects of modern supercomputing.

15. SUBJECT TERMS
High performance computing; Heterogeneous computer system

16. SECURITY CLASSIFICATION OF: a. REPORT	b. ABSTRACT	c. THIS PAGE	17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
					Andreas Gross
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UU	5	**19b. TELEPHONE NUMBER** *(include area code)* 575-646-6179

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

Contents

1. Summary 2

2. Introduction 2

3. Computational Fluid Dynamics Lab as Motivator for New System 3

4. Description of New System 4

5. Student Training 4

6. User Data for 2019 5

7. Conclusions 5

1. Summary

Accurate first-principles based simulations of flows of engineering interest require grid resolutions that challenge existing Central Processing Unit (CPU) based computer systems. The combination of CPUs, Graphics Processing Units (GPUs) and other purpose-designed hardware into heterogeneous systems results in a leap in compute power over older, purely CPU-based systems. Novel simulation codes and algorithms have to be developed to unleash the compute power inherent in these new hardware architectures for cutting-edge research.

A heterogeneous system was purchased and installed at New Mexico State University (NMSU). This system allows researchers and graduate students at NMSU to develop highly effective and efficient codes for heterogeneous systems that will greatly benefit our funded research. Equally important, the system significantly enhances the supercomputing research infrastructure at NMSU. The system allows us to develop codes and algorithms that exploit the massive parallelism of the latest heterogeneous systems at the Department of Defense (DoD) high performance computing (HPC) centers. The system also constitutes an ideal education and training environment for our graduate students in all aspects of modern supercomputing.

2. Introduction

Computer codes are typically developed locally and then transferred to High Performance Computing (HPC) centers for massive simulation runs. Recently a shift occurred at these centers towards heterogeneous systems that are comprised of Central Processing Units (CPUs) and Single Instruction Multiple Data (SIMD) devices such as Graphics Processing Units (GPU) and/or Intel Xeon Phi coprocessors. Although early supercomputers such as the Cray 1 were SIMD devices, the industry moved away from the concept in the 90s, and, up until recently, HPC was dominated by shared or distributed memory CPU systems. The Message Passing Interface (MPI) standard was developed for making efficient use of large CPU systems. Recently the trend is reversing and large heterogeneous systems are going online that inherit traits of the older SIMD systems. This paradigm shift was motivated by the abundant availability of relatively inexpensive GPUs and new customized SIMD chips for supercomputing applications. For example, many of the systems on the TOP500 supercomputer list, such as the Sunway TaihuLight (Wuxi, China), Tianhe-2 (Guangzhou, China), and Titan (Oak Ridge), make extensive use of some form of non-CPU based accelerators (64 SIMD vector processors per CPU for the TaihuLight, 48,000 Phi coprocessors for the Tianhe-2, and 18,688 GPUs for the Titan). Such massively parallel architectures (which inspired the creation of words such as "exascale" or "pervasive computing") would allow for much larger cases and faster turnaround times. This makes the development of multi-threaded codes indispensable for the full and efficient utilization of such systems.

Graphical Processing Units are becoming increasingly important for HPC. Their exceptional performance compared to CPUs can be attributed to the fact that GPUs are capable of executing thousands of threads concurrently. Currently, NVIDIA GPUs outperform Intel CPUs on floating point performance and memory bandwidth by a factor of ten or more. GPUs are designed for vector processing which refers to the simultaneous alteration of multiple data sets based on a common set of instructions. Therefore, the architecture of GPUs can be classified as SIMD. Compute-intensive tasks of an application can be grouped into a common instruction set such that each GPU core works on different data but executes the same instruction set. The amount of parallelism is breathtaking. For example, the NVIDIA Tesla P100 has 3,584 compute cores for a maximum of 5.3TeraFLOPS of double precision floating point operations. The amount of fast memory available to each thread is relatively small. The Intel Xeon Phi chip can be seen as a middle ground between GPUs and CPUs. Also a SIMD, it is less restrictive in terms of instruction set and memory use. The most powerful Phi chip, the Xeon Phi 7290F has 72 cores (288 threads) delivering up to 3.5TeraFLOPS of double precision floating point operations.

3. Computational Fluid Dynamics Lab as Motivator for New System

Computational fluid dynamics (CFD) is one area of research that relies heavily on high performance computing. In CFD the governing equations are solved in a discrete manner on a computational grid. For direct numerical simulations (DNS) all scales of fluid motion up to the smallest resolved scales have to be resolved which is often only possible on the latest and largest HPC systems. Similar to DNS, in large-eddy simulations (LES) a large part of the unsteady fluid motion is resolved. However, the smaller scale motion, which is more stochastic, is modeled. For applications of engineering interest, the grid resolution requirement is typically still immense and CFD again relies heavily on HPC.

Most of the CFD research at NMSU centers on implicit LES (ILES) which have gained considerable popularity in recent years. Similar to subgrid stress models in traditional LES, in ILES the numerical diffusion of the discretization removes energy at the unresolved scales. In ILES, the accuracy of the numerical scheme and the grid resolution determine how much of the turbulence kinetic energy (TKE) spectrum is captured. Larger simulations require more computational resources. This drives the development and advancement of CFD codes that make the most efficient use of new computer systems and numerical methods.

A large number of our simulations are carried out with a compressible CPU-based CFD code that was developed in-house with funding from the Department of Defence (DoD). The formal order of accuracy of the convective terms is up to 9th-order. More accurate schemes are hardly practical and, therefore, any increase in Reynolds number, Re, typically requires a proportional increase in grid resolution (Chapman shows that the number of grid points for LES scales with $Re^{9/5}$). The numerical dissipation of the underlying weighted-essentially non-oscillatory discretization removes enough energy at the grid scales for the simulations to qualify as ILES. Compared to other simulation strategies that employ turbulence modeling, ILES are more accurate but also more computationally expensive. By running simulations on up to 2,048 cores at the DoD HPC centers, we are able to resolve a large fraction of the wave number spectrum (such that the grid cutoff falls in the inertial subrange) for Reynolds numbers (e.g., based on chord) of up to approximately 300,000. Larger Reynolds numbers require simulations to be carried out on more cores and/or to make efficient use of SIMD chips which are increasingly seen on the latest DoD HPC systems. In its present form, our CFD code cannot take advantage of SIMD devices which form the backbone of the most recent large-scale HPC systems.

Keeping all this in mind, together with the NMSU Information & Communication Technologies (ICT) Cyber Infrastructure Architect Team, requirements for the new DURIP system were formulated. It was decided that the new system shall

- Offer a heterogeneous compute architecture that mirrors those at the DoD HPC centers to allow for the development and testing of code that can be ported to very-large-scale heterogeneous HPC systems
- Provide considerable local compute power for the CFD lab and other researchers
- Allow for compute intensive data post-processing
- Offer a state-of-the-art environment for training students in all aspects of supercomputing

4. Description of New System

Quotes were obtained from different vendors, among them Lenovo and Dell. The most competitive system, which was offered by Dell, was purchased. The cost was $199,851.08 and did thus exactly match the total amount of the contract.

The system was modeled after modern DoD HPC systems which typically offer a mix of compute nodes with typically two GPUs per CPU and pure CPU compute nodes. This architecture was chosen so that code that is developed locally at NMSU can later be ported to DoD HPC systems. In the following a brief outline of the system is provided.

The system has five mixed GPU/CPU nodes and 10 CPU nodes. The nodes are connected with an Infiniband switch. No Phi chips were included since it was learned that the Phi chip was to be discontinued.

Each one of the five GPU/CPU node houses:

- Two Xeon Gold 5117 with 14 cores each
- Two Tesla P100 GPUs with 3584 cores and 16GB RAM each
- 192GB of RAM
- One 480GB hard drive

Each one of the 10 CPU node houses:

- Two Xeon Gold 5117 with 14 cores each
- 192GB of RAM
- One 480GB hard drive

The total number of GPU cores is $5 \times 2 \times 3584 = 35{,}840$.
The total number of CPU cores is $5 \times 2 \times 14 + 10 \times 2 \times 14 = 420$.
The total amount of RAM is 2.88TB and the total hardware storage space is 7.2TB.

The operating system is Centos 7. The system is housed in an air-conditioned server room at the Information & Communication Technologies (ICT) center and accessed through the slurm job scheduler. The system maintenance is done by ICT staff.

5. Student Training

Within the CFD lab, since 2019 the new system has been used for the training of five graduate students:

- One Ph.D. student is using the system for developing OpenCL-based subroutines for accelerating the compressible CFD code. OpenCL is a framework for writing programs that execute across heterogeneous platforms consisting of central processing units (CPUs), graphics processing units (GPUs), digital signal processors (DSPs), field-programmable gate arrays (FPGAs) and other processors or hardware accelerators.
- Two M.S. and two Ph.D. students are using the system for CFD simulations. The M.S. students which don't have access to DoD HPC systems are also learning how to submit and run compute jobs on HPC-like computer systems and how to post-process data.

6. User Data for 2019

The reporting tool called Open XDMoD was employed to monitor the system usage from March 2019 till now for the CFD laboratory. The CFD laboratory alone used in excess of 600,000 compute hours since March 2019 (Fig. 1). The CFD lab is given priority on the system. Other NMSU entities can utilize the system as well through a backfill queue. Compute hours used by these entities add to the number stated above. For example, the total number of CPU hours since 10/15/2019 is 429,148; 339,029 of which were used by the CFD lab, 84,779 were used by the Open Science Grid, and 5,340 were used by other researchers at NMSU via the backfill queue.

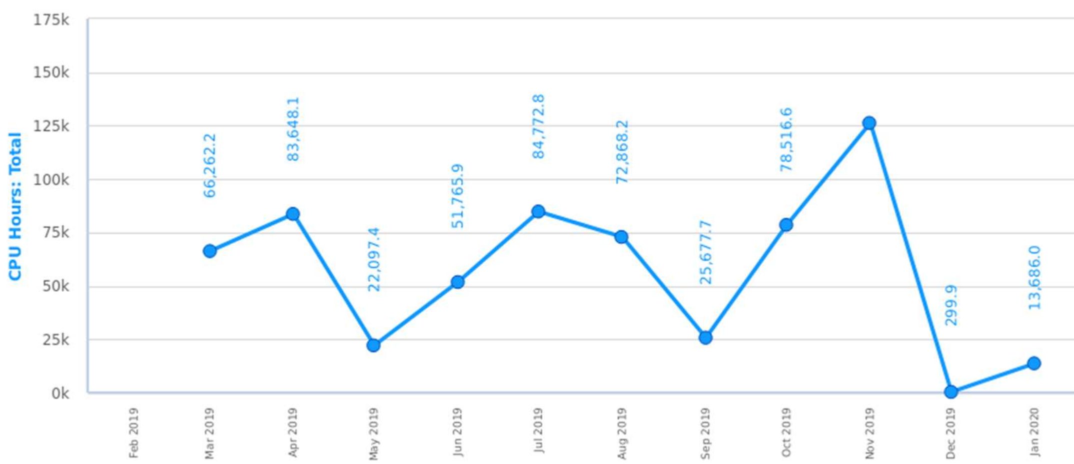

Fig. 1: System Usage.

7. Conclusions

The purchased system mirrors the latest system architectures at the DoD HPC centers. This makes the system suitable for the development of efficient codes and algorithms for heterogeneous systems. The system also provides a much needed and substantial upgrade of the local computer infrastructure that would not have been possible without DURIP support. It permits us for the first time ever to carry out production runs locally, which dramatically benefits our DoD funded research, but also offloads some strain from the DoD HPC centers. Finally, the system provides an excellent environment for student training in the use of state-of-the-art HPC systems so that valuable compute time at the DoD HPC centers is used for research and not training or code optimization.