# REPORT DOCUMENTATION PAGE

Form Approved OMB NO. 0704-0188



| 1. REPORT DATE (DD-MM-YYYY) | 2. REPORT TYPE | 3. DATES COVERED (From - To) |
|---|---|---|
| 01-11-2016 | Final Report | 1-Aug-2015 - 31-Jul-2016 |

| 4. TITLE AND SUBTITLE | 5a. CONTRACT NUMBER |
|---|---|
| Final Report: A Transitional Computational Platform to Migrate Parachute Simulation from Workstation to HPC | W911NF-15-1-0403 |
| | 5b. GRANT NUMBER |
| | 5c. PROGRAM ELEMENT NUMBER<br>611103 |
| 6. AUTHORS<br>Xiaolin Li | 5d. PROJECT NUMBER |
| | 5e. TASK NUMBER |
| | 5f. WORK UNIT NUMBER |

| 7. PERFORMING ORGANIZATION NAMES AND ADDRESSES | 8. PERFORMING ORGANIZATION REPORT NUMBER |
|---|---|
| Research Foundation of SUNY at Stony Brc<br>W-5510 Melville Library<br><br>Stony Brook, NY 11794 -3362 | |

| 9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) | 10. SPONSOR/MONITOR'S ACRONYM(S) |
|---|---|
| U.S. Army Research Office<br>P.O. Box 12211<br>Research Triangle Park, NC 27709-2211 | ARO |
| | 11. SPONSOR/MONITOR'S REPORT NUMBER(S)<br>66858-MA-RIP.1 |

12. DISTRIBUTION AVAILIBILITY STATEMENT

Approved for Public Release; Distribution 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

Supported by the DURIP grant W911NF-15-1-0403 for 'A Transitional Computational Platform to Migrate Parachute Simulation from Workstation to HPC', we purchased a parallel cluster from the Advanced Cluster Technologies, Inc.. This parallel cluster is named 'Intruder', after a type of sports parachute. It consists of one head node, 21 computing nodes (20 CPU (Central Processing Unit) nodes

15. SUBJECT TERMS

parallel cluster, strong and weak scaling, front tracking, HPC

| 16. SECURITY CLASSIFICATION OF: | | | 17. LIMITATION OF ABSTRACT | 15. NUMBER OF PAGES | 19a. NAME OF RESPONSIBLE PERSON |
|---|---|---|---|---|---|
| a. REPORT | b. ABSTRACT | c. THIS PAGE | | | Xiaolin Li |
| UU | UU | UU | UU | | 19b. TELEPHONE NUMBER<br>631-632-8354 |

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

## Report Title

Final Report: A Transitional Computational Platform to Migrate Parachute Simulation from Workstation to HPC

## ABSTRACT

Supported by the DURIP grant W911NF-15-1-0403 for 'A Transitional Computational Platform to Migrate Parachute Simulation from Workstation to HPC', we purchased a parallel cluster from the Advanced Cluster Technologies, Inc.. This parallel cluster is named 'Intruder', after a type of sports parachute. It consists of one head node, 21 computing nodes (20 CPU (Central Processing Unit) nodes and 1 GPU (Graphic Processing Unit) node), connected with 56Gb/s InfiniBand and 1000MB/s Ethernet. Each node was populated with dual Eight-Core Intel E5-2630v3 'Haswell' 2.4GHz processors with different size of RAM (Random-access Memory) and Storage. The head node and compute node have 32GB of RAM for each, and the GPU node has 128GB of RAM. A 32TB of network file system using RAID6 (Redundant Array of Independent Disks) is installed in the head node, and is shared with other nodes. Each node also has a clone of the operation system on its local disk: 2TB SSD on head node, 1TB SATA drive on compute nodes, and 240GB SSD on GPU node. The parallel computing can be further accelerated by including the GPU node, which contains seven NVIDIA Tesla K40 GPUs with 12GB of RAM for each device.

---

## Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

### (a) Papers published in peer-reviewed journals (N/A for none)

Received Paper

**TOTAL:**

**Number of Papers published in peer-reviewed journals:**

---

### (b) Papers published in non-peer-reviewed journals (N/A for none)

Received Paper

**TOTAL:**

**Number of Papers published in non peer-reviewed journals:**

## (c) Presentations

**Number of Presentations:** 0.00

## Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

| <u>Received</u> | <u>Paper</u> |
| --- | --- |

**TOTAL:**

**Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):**

## Peer-Reviewed Conference Proceeding publications (other than abstracts):

| <u>Received</u> | <u>Paper</u> |
| --- | --- |

**TOTAL:**

**Number of Peer-Reviewed Conference Proceeding publications (other than abstracts):**

## (d) Manuscripts

| <u>Received</u> | <u>Paper</u> |
| --- | --- |

**TOTAL:**

**Number of Manuscripts:**

## Books

<u>Received</u> <u>Book</u>

**TOTAL:**

<u>Received</u> <u>Book Chapter</u>

**TOTAL:**

## Patents Submitted

## Patents Awarded

## Awards

## Graduate Students

| <u>NAME</u> | <u>PERCENT_SUPPORTED</u> |
|---|---|
| **FTE Equivalent:** | |
| **Total Number:** | |

## Names of Post Doctorates

| <u>NAME</u> | <u>PERCENT_SUPPORTED</u> |
|---|---|
| **FTE Equivalent:** | |
| **Total Number:** | |

## Names of Faculty Supported

NAME PERCENT_SUPPORTED

**FTE Equivalent:**
**Total Number:**

## Names of Under Graduate students supported

NAME PERCENT_SUPPORTED

**FTE Equivalent:**
**Total Number:**

## Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period: ...... 0.00

The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:...... 0.00

The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:...... 0.00

Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):...... 0.00

Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:...... 0.00

The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense ...... 0.00

The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields:...... 0.00

## Names of Personnel receiving masters degrees

NAME

**Total Number:**

## Names of personnel receiving PHDs

NAME

**Total Number:**

## Names of other research staff

NAME PERCENT_SUPPORTED

**FTE Equivalent:**
**Total Number:**

## Sub Contractors (DD882)

# Inventions (DD882)

# Scientific Progress

Technical Report on the Transitional 'Intruder' Cluster

Supported by the DURIP grant W911NF-15-1-0403 for 'A Transitional Computational Platform to Migrate Parachute Simulation from Workstation to HPC', we purchased a parallel cluster from the Advanced Cluster Technologies, Inc. The parallel cluster is named 'Intruder' a type of sports parachute. It consists of one head node, 21 computing nodes (20 CPU (Central Processing Unit) nodes and 1 GPU (Graphic Processing Unit) node), connected with 56Gb/s InfiniBand and 1000MB/s Ethernet. Each node was populated with dual Eight-Core Intel E5-2630v3 'Haswell' 2.4GHz processors with different size of RAM (Random-access Memory) and Storage. The head node and compute node have 32GB of RAM for each, and the GPU node has 128GB of RAM. A 32TB of network file system using RAID6 (Redundant Array of Independent Disks) is installed in the head node, and is shared with other nodes. Each node also has a clone of the operation system on its local disk: 2TB SSD on head node, 1TB SATA drive on compute nodes, and 240GB SSD on GPU node. The parallel computing can be further accelerated by including the GPU node, which contains seven NVIDIA Tesla K40 GPUs with 12GB of RAM for each device. A hybrid computation utilizing both CPU and GPU has been performed on this platform. The detailed description of the hardware is summarized in Table 1 of the attached pdf file.

The software environment consists of the CentOS 6.7 operating system; the Intel OpenMPI compiler 1.8 for MPI-based applications; Sun Grid Engine 2011.11p1 as the job scheduler, PETSc3.6 with benchmark examples; FronTier with computational fluid dynamics applications. Benchmark tests have been performed within this environment, and the performance are compared with the workstation and Cray supercomputer.

We have carried out several experiments to investigate the impact of the computing platform on our main application: FronTier, which relies on a few parallel computing API (such as MPI, CUDA) and external packages (such as PETSc, HYPRE). In order to distinguish their impact of these external packages on FronTier, we created a few independent programs by including PETSc, MPI and CUDA separately. In the first test, a two dimensional Poisson equation was solved with PETSc using the KSP solver and HYPRE as the preconditioner. To test the strong scaling, we fix the domain size to be 4096X4096 while consecutively double the number of processors from 1 to 1024. The scaling results are summarized in Table 2 and illustrated by Figure 1 of the attached pdf file which and showed that the HPC has a wider range of linear scaling than "Intruder" cluster and workstation. The result also suggests that the speedup becomes slower when the machine is nearly fully occupied. We interpret this fact as an indicator of the limited bandwidth of the main memory which is substantially slower than the speed of the modern CPU. The second experiment is to test the weak scaling of FronTier by solving 2-D Riemann problem with MPI library. The base size for each processor is set to be 100X100, 200X200, 400X400 and 800X800. The efficiency is measured by T_1/T_N, where N is the number of processors occupied. The results are displayed in Figure 2 and Table 3 in the attached pdf file, which implies that an efficiency over 50% can be achieved up to 256 cores when the base size is smaller than 400X400 while the efficiency decreases to 20% when the base size is 800X800$. The third experiment tests the GPU acceleration on the cloth simulator. The GPU code is implemented with CUDA library, which is a parallel computing platform created by NVIDIA. We compare the computational time of solving spring model for different parachute types using or without using GPU device and calculate the speedup. As shown in Table 3 of the attached pdf file, using GPU device can achieve at least 16 times and up to 21 times speedup for cloth simulation.

Due to the price reduction since we wrote the proposal, there is a fraction of fund left after the purchase of the cluster. We used the remaining fund to buy two workstations and three desktops to be used by graduate students in the research project, high school students in the HSAP program, and undergraduate students who are interested in our research.

## Technology Transfer

**Recipient: Joseph Myers**
**Army Research Office**

**ARO–DURIP: W911NF-15-1-0403**

**Award Report**

**A Transitional Computational Platform to Migrate Parachute Simulation from Workstation to HPC**

**Principal Investigator: Xiaolin Li**
**University at Stony Brook**

**Reporting Period:**
**September 1, 2015 – August 31, 2016**

**Recipient:**
**Research Foundation**
**University at Stony Brook**
**Stony Brook, NY 11794-3366**

**Unexpended Funds: $0.00 (0.0%)**
**Exceeds 10% of available funds:** No

# Technical Report on the Transitional "Intruder" Cluster

November 1, 2016

Supported by the DURIP grant W911NF-15-1-0403 for "A Transitional Computational Platform to Migrate Parachute Simulation from Workstation to HPC", we purchased a parallel cluster from the Advanced Cluster Technologies, Inc. The parallel cluster is named "Intruder", after a type of sports parachute. It consists of one head node, 21 computing nodes (20 CPU (Central Processing Unit) nodes and 1 GPU (Graphic Processing Unit) node), connected with $56Gb/s$ InfiniBand and $1000MB/s$ Ethernet. Each node was populated with dual Eight-Core Intel $E5-2630v3$ "Haswell" $2.4GHz$ processors with different size of RAM (Random-access Memory) and Storage. The head node and compute node have $32GB$ of RAM for each, and the GPU node has $128GB$ of RAM. A $32TB$ of network file system using RAID6 (Redundant Array of Independent Disks) is installed in the head node, and is shared with other nodes. Each node also has a clone of the operation system on its local disk: $2TB$ SSD on head node, $1TB$ SATA drive on compute nodes, and $240GB$ SSD on GPU node. The parallel computing can be further accelerated by including the GPU node, which contains seven NVIDIA Tesla K40 GPUs with $12GB$ of RAM for each device. A hybrid computation utilizing both CPU and GPU has been performed on this platform. The detailed description of the hardware is summarized in Table Table 1.

| Node Type | CPU | RAM | Storage for data | Storage for OS | GPU |
|---|---|---|---|---|---|
| Head node | 2x8-core Intel E5-2630v3 | 32GB | 32TB (RAID 6) | 2TB (SSD) | none |
| CPU node | 2x8-core Intel E5-2630v3 | 32GB | none | 1TB (SATA) | none |
| GPU node | 2x8-core Intel E5-2630v3 | 128GB | none | 240GB (SSD) | 7 Tesla K40 |

Table 1: Summary of the hardware

The software environment consists of the CentOS 6.7 operating system; the Intel OpenMPI compiler 1.8 for MPI-based applications; Sun Grid Engine $2011.11p1$ as the job scheduler, PETSc3.6 with benchmark examples; FronTier with computational fluid dynamics applications. Benchmark tests have been performed within this environment, and the performance are compared with the workstation and Cray supercomputer.

We have carried out several experiments to investigate the impact of the computing platform on our main application: FronTier, which relies on a few parallel computing API (such as MPI, CUDA) and external packages (such as PETSc, HYPRE). In order to distinguish their impact of these external packages on FronTier, we created a few independent programs by including PETSc, MPI and CUDA separately. In the first test, a two dimensional Poisson equation was solved with PETSc using the KSP solver and HYPRE as the preconditioner. To test the strong scaling, we fix the domain size to be $4096 \times 4096$ while consecutively double the number of processors from 1 to 1024. The scaling results are summarized in Table 2 and illustrated by Figure 1, and showed that the HPC has a wider range of linear scaling than "Intruder" cluster and workstation. The result also suggests that the speedup becomes slower when the machine is nearly fully occupied. We interpret this fact as an indicator of the limited bandwidth of the main memory which is substantially slower than the speed of the modern CPU. The second experiment is to test the weak scaling of FronTier by solving 2-D Riemann problem with MPI library. The base size for each processor is set to be $100 \times 100$, $200 \times 200$, $400 \times 400$ and $800 \times 800$. The efficiency is measured by $T_1/T_N$, where $N$ is the number of processors occupied. The results are displayed in Figure 2 and Table 3 which implies that an efficiency over 50% can be achieved up to 256 cores when the base size is smaller than $400 \times 400$ while the efficiency decreases to 20% when the base size is $800 \times 800$. The third experiment tests the GPU acceleration on the cloth simulator. The GPU code is implemented with CUDA library, which is a parallel computing platform created by NVIDIA. We compare the computational time of solving spring model for different parachute types using or without using GPU device and calculate the speedup. As shown in Table 4, using GPU device can achieve at least 16 times and up to 21 times speedup for cloth simulation.

Due to the price reduction since we wrote the proposal, there is a fraction of fund left after the purchase of the cluster. We used the remaining fund to buy two workstations and three desktops to be used by graduate students in the research project, high school students in the HSAP program, and undergraduate students who are interested in our research.

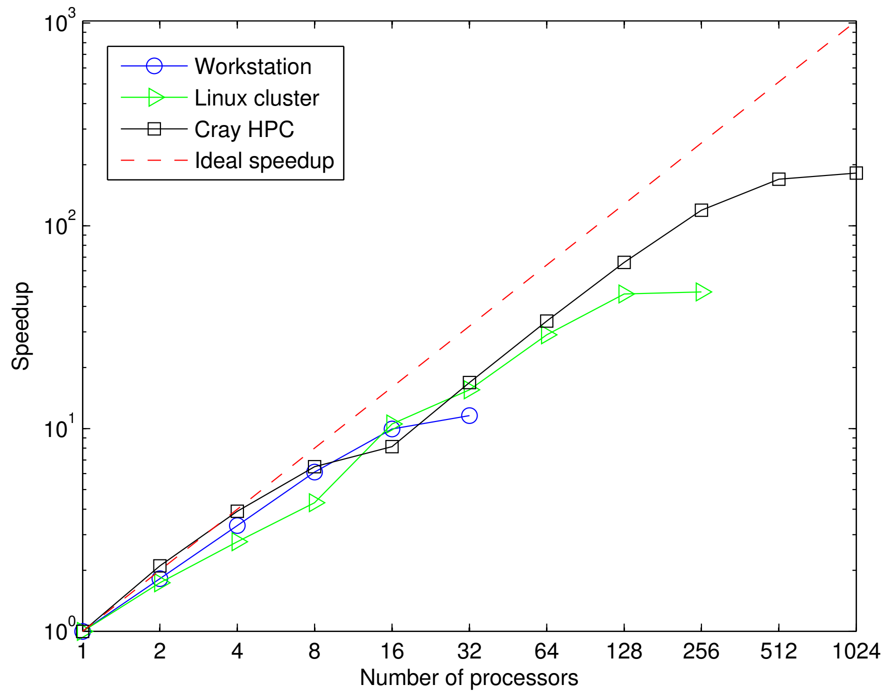


Figure 1: The figure displays the speedup of linear solver solving the 2-D Poisson equation with domain size $(4096 \times 4096)$ on different machine: workstation, linux cluster and Cray supercomputer.

| Number of processors | 1 | 2 | 4 | 8 | 16 | 32 | 64 | 128 | 256 | 512 | 1024 |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Time on workstation (s) | 136.80 | 88.61 | 48.45 | 27.71 | 16.99 | 14.59 | - | - | - | - | - |
| Time on linux cluster (s) | 129.10 | 87.59 | 54.86 | 37.06 | 15.14 | 10.28 | 5.77 | 3.62 | 3.70 | - | - |
| Time on HPC (s) | 78.06 | 56.63 | 31.53 | 20.90 | 17.67 | 9.98 | 5.34 | 3.05 | 1.082 | 1.31 | 1.22 |

Table 2: Computational time summary of strong scaling test on workstation, Linux cluster and CRAY HPC by solving the 2D Poisson equation.

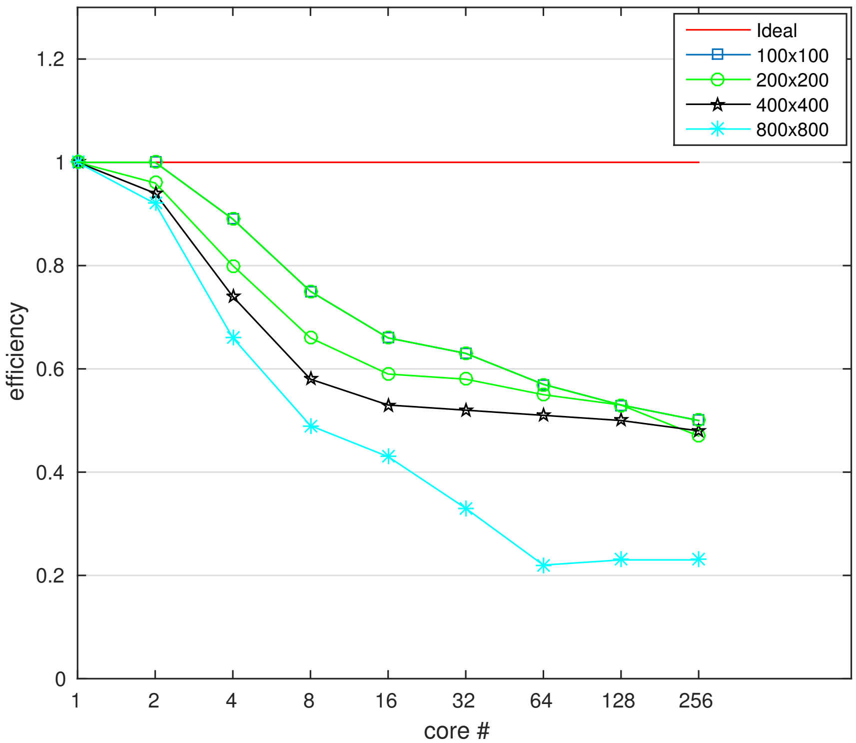


Figure 2: This figure shows the weak scaling test on 2d Riemann problem. The efficiency is calculated with $T_1/T_N$, where $N$ is the number of processors used. The results implies that 50% efficiency has been achieved when mesh size per processor is smaller than $400 \times 400$.

| Size per proc | num of procs | 1 | 2 | 4 | 8 | 16 | 32 | 64 | 128 | 256 |
|---|---|---|---|---|---|---|---|---|---|---|
| $100 \times 100$ | time (s) | 40 | 40 | 45 | 53 | 61 | 64 | 70 | 75 | 80 |
| | efficiency | 1.00 | 1.00 | 0.89 | 0.75 | 0.66 | 0.63 | 0.57 | 0.53 | 0.50 |
| $200 \times 200$ | time (s) | 85 | 89 | 106 | 129 | 145 | 149 | 154 | 159 | 181 |
| | efficiency | 1.00 | 0.96 | 0.80 | 0.66 | 0.59 | 0.58 | 0.55 | 0.53 | 0.47 |
| $400 \times 400$ | time (s) | 401 | 427 | 544 | 695 | 762 | 773 | 783 | 809 | 838 |
| | efficiency | 1.00 | 0.94 | 0.74 | 0.58 | 0.53 | 0.52 | 0.51 | 0.50 | 0.46 |
| $800 \times 800$ | time (s) | 2081 | 2263 | 3161 | 4244 | 4839 | 6334 | 9282 | 8865 | 9114 |
| | efficiency | 1.00 | 0.92 | 0.66 | 0.49 | 0.43 | 0.33 | 0.22 | 0.23 | 0.23 |

Table 3: Weak scaling test on 2d Riemann problem by gradually increase the mesh size per processor through $100 \times 100$, $200 \times 200$, $400 \times 400$, $800 \times 800$.

| Parachute type | CPU/GPU | Time(s) | Avg time per step(s) | Speedup |
|---|---|---|---|---|
| C9 | CPU | 2805.85 | 3.39 | 1.00 |
| | GPU | 131.90 | 0.16 | 21.2 |
| G11 | CPU | 5101.47 | 5.41 | 1.00 |
| | GPU | 243.18 | 0.26 | 20.81 |
| Intruder | CPU | 1252.65 | 2.00 | 1.00 |
| | GPU | 69.67 | 0.11 | 18.18 |
| T10 | CPU | 5540.02 | 5.99 | 1.00 |
| | GPU | 282.74 | 0.36 | 16.64 |
| T11 | CPU | 6791.9 | 5.12 | 1.00 |
| | GPU | 352.07 | 0.29 | 17.66 |

Table 4: A comparison of computational time between different parachute type on CPU or GPU. The speedup is calculated based on the computing time by CPU.