Active Project (2016 - 2018)

Accelerating Memory-Access-Limited HPC Applications via Novel Fast Data Compression, Phase II Project

SBIR/STTR Programs | Space Technology Mission Directorate (STMD)

ABSTRACT

A fast-paced continual increase on the ratio of CPU to memory speed feeds an exponentially growing limitation for extracting performance from HPC systems. Breaking this memory wall is one of the most important challenges that the HPC community faces today. In Phase I we introduced aggressive innovations enable the injection of unprecedented acceleration into vast classes of memory-access-bound HPC codes via ultra-fast software-based data compression. Groundbreaking speedup on a fully functional NPBCG prototype was delivered to NASA, thus validating the tremendous potential of our approach. The proposed approach is based on a revolutionary theory of compression spearheaded by Accelogic (Compressive Computing), which is able to provide enormous compressive gains for the typical floating point data of HPC applications. In Phase II we will build on our success with the NPBCG benchmark, and move boldly into tackling the acceleration of a real-life high-profile code, namely NASA's Cart3D, improving its performance by a paradigm-shifting 2x to 4x end-to-end wall-clock time acceleration by the end of Phase II. Our firm has accumulated crucial know-how and has synthesized its expertise into a powerful industrial-quality process for software acceleration that will be used to ensure success on completing Phase II objectives. In Phase II we also plan on injecting a second NASA code with basic Compressive Computing techniques, and providing it with base levels of acceleration of ~1.3-2x. We will choose this second code from a pool of high-profile codes that have already signed up as early adopters for this project: FUN3D, USM3D, Enzo, and WRF. The work on a second NASA code will also serve as the ultimate field test of the broadness and ease-of-infusion of the proposed technology. We have secured complementary funds in the amount of $500,000 to increase resources and ensure that the proposed Phase II proposed will be successfully accomplished.

Table of Contents

Technology Maturity

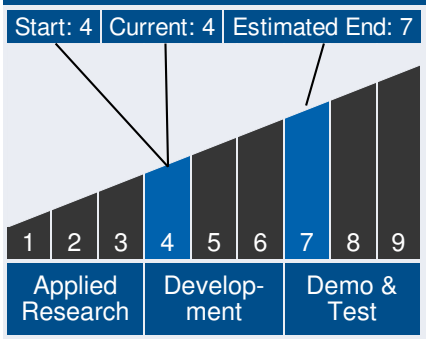

Management Team

Program Executives:

- Joseph Grant
- Laguduva Kubendran

Program Manager:

- Carlos Torrez

Continued on following page.

Accelerating Memory-Access-Limited HPC Applications via Novel Fast Data Compression, Phase II Project

ANTICIPATED BENEFITS

To NASA funded missions:

Potential NASA Commercial Applications: The impact of the proposed technology spans most areas of importance to NASA's scientific missions, including: aerospace, weather forecasting, cosmology, combustion, climate research, and chemistry, among others. To this date, five of the Top NASA HPC applications have enlisted as partners of the project to become early adopters of the technology. This fact speaks clearly about the interest that the NASA community has shown on the potential uses and benefits of infusing the knowledge generated from this project into NASA. Furthermore, once the technology is fully operational, it will benefit tens of thousands of users, who will see substantially increased performance in their regular, day-to-day runs, as well as in their massive, supercomputer-based production runs. One of the lead developers of NASA's Top Codes mentions that this technology "can be considered critical in achieving the next generation of so-called exascale software applications, [and] in turn, these efforts will enable scientific and engineering breakthroughs previously considered computationally intractable".

To the commercial space industry:

Potential Non-NASA Commercial Applications: The resulting technology will increase the efficiency of memory access in most modern computer architectures, thus directly enabling unprecedented speedups in memory-access-bound HPC applications. With a significant fraction of HPC codes belonging to this "memory-bound" category, numerous scientists, developers, researchers, and complete industries will benefit, in areas as varied as aerospace, climate research, molecular dynamics, chemistry, weather forecasting, energy, civil engineering, geophysics, and life sciences, among others.

Management Team *(cont.)*

Principal Investigator:

- Juan Gonzalez

Technology Areas

Primary Technology Area:
Modeling, Simulation, Information Technology and Processing (TA 11)
└ Computing (TA 11.1)
  └ Ground Computing (TA 11.1.2)

Accelerating Memory-Access-Limited HPC Applications via Novel Fast Data Compression, Phase II Project

SBIR/STTR Programs | Space Technology Mission Directorate (STMD)

U.S. WORK LOCATIONS AND KEY PARTNERS

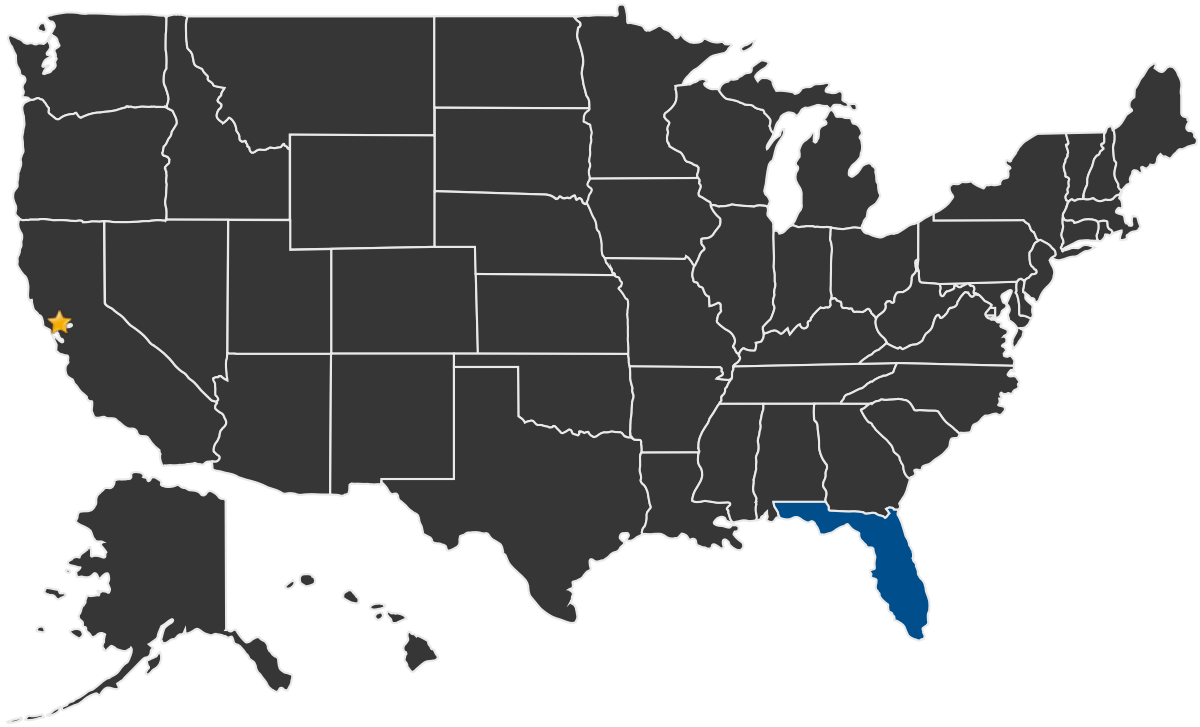

U.S. States With Work

Lead Center:
Ames Research Center

Other Organizations Performing Work:

- Accelogic, LLC (Weston, FL)

PROJECT LIBRARY

Presentations

- Briefing Chart
  - (http://techport.nasa.gov:80/file/23550)

IMAGE GALLERY

Accelerating Memory-Access-Limited HPC Applications via Novel Fast Data Compression, Phase II

DETAILS FOR TECHNOLOGY 1

Technology Title

Accelerating Memory-Access-Limited HPC Applications via Novel Fast Data Compression, Phase II

Potential Applications

The impact of the proposed technology spans most areas of importance to NASA's scientific missions, including: aerospace, weather forecasting, cosmology, combustion, climate research, and chemistry, among others. To this date, five of the Top NASA HPC applications have enlisted as partners of the project to become early adopters of the technology. This fact speaks clearly about the interest that the NASA community has shown on the potential uses and benefits of infusing the knowledge generated from this project into NASA. Furthermore, once the technology is fully operational, it will benefit tens of thousands of users, who will see substantially increased performance in their regular, day-to-day runs, as well as in their massive, supercomputer-based production runs. One of the lead developers of NASA's Top Codes mentions that this technology "can be considered critical in achieving the next generation of so-called exascale software applications, [and] in turn, these efforts will enable scientific and engineering breakthroughs previously considered computationally intractable".