Active Project (2014 - 2016)

Accelerating Communication-Intensive Applications via Novel Data Compression Techniques, Phase II Project

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

ABSTRACT

The traditional view of arithmetic operations dominating the computational cost of numerical algorithms has changed. As performance of new processors increases, we are moving into a new reality in which data movement is expensive and operations are becoming nearly free. In Phase I we discovered a new theory of data compression with unprecedented capability to reduce data movement in distributed HPC applications. Phase I results include: 1) the demonstration via proof-of-concept prototypes of lossless compressive gains of 4x-20x for NAS Parallel Benchmarks, and 2) the formulation of COPA, a Compression OPportunities Auditing & discovery process that facilitates the analysis of any given large-scale code with the purpose of optimizing it for performance via data movement reduction. Phase II targets the consolidation of the new theory through two major thrusts: 1) the development of software that facilitates integrating high-yield compression into HPC codes, and 2) the further infusion of the technology into NASA applications. Thrust 1 will produce a fully functional prototype of a software suite for the automatic/semi-automatic acceleration of HPC codes via compression. The prototype will incorporate ease-of-infusion features (e.g., MPI-enabled compression and decompression routines), as well as ergonomic features that allow the seamless integration of new compression modules into the suite. Thrust 2 will begin with the application of the Chapter 1 of COPA to five HPC codes of utmost importance to NASA, namely Cart3D, FUN3D, USM3D, Enzo, and WRF. Accelogic will collaborate with the developers of these codes to integrate the compression technology under three different levels of "integrability," namely automatic, semi-automatic, and manual. By the end of Phase II, at least one of these applications will undergo the complete COPA optimization process. Complementary Phase II/III funds for $1+ million have been secured to ensure successful commercialization.

Table of Contents

Technology Maturity

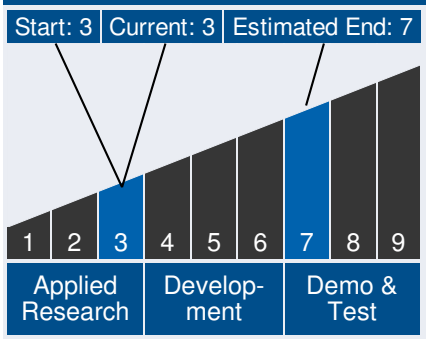

Management Team

Program Executives:

- Joseph Grant
- Laguduva Kubendran

Program Manager:

- Carlos Torrez

Continued on following page.

Accelerating Communication-Intensive Applications via Novel Data Compression Techniques, Phase II Project

ANTICIPATED BENEFITS

To NASA funded missions:

Potential NASA Commercial Applications: The impact of this technology on NASA is high and very broad. The fact that, in the short time span of the Phase I project, five of the Top NASA HPC applications have submitted letters of intent to Accelogic to become Phase II participants, speaks vastly about the significance of this technology for NASA. One of these participants mentions in his letter of intent that this technology "can be considered critical in achieving the next generation of so-called exascale software applications." Another participant mentions that a rough audit of their codes in production runs suggests that speedup factors on the order of 2-4 would be reasonable to expect across the spectrum of their user's cases. Such speedups would represent new opportunities for the code users to explore new larger science problems, faster and at a lower cost. The impact of this compression technology spans most areas of importance to NASA's scientific mission, including aerospace, weather forecasting, combustion, climate research, and chemistry, among many others.

To the commercial space industry:

Potential Non-NASA Commercial Applications: The landscape for non-NASA government and commercial applications is similar, with this compression technology having a potential strong impact on most fields for which HPC is crucial. This includes a vast array of Government and commercial fields such as: climate research, molecular dynamics, chemistry, defense, weather forecasting, energy, finance, economic forecasting, civil and environmental engineering, geophysics, life sciences, and semiconductors.

Management Team *(cont.)*

Project Manager:

- David Kao

Principal Investigator:

- Juan Gonzalez

Technology Areas

Primary Technology Area:
Modeling, Simulation, Information Technology and Processing (TA 11)
└ Information Processing (TA 11.4)
└ Intelligent Data Understanding (TA 11.4.2)
└ Intelligent Data Collection and Prioritization Toolset (TA 11.4.2.1)

Secondary Technology Area:
Modeling, Simulation, Information Technology and Processing (TA 11)
└ Modeling (TA 11.2)
└ Frameworks, Languages, Tools, and Standards (TA 11.2.5)

Additional Technology Areas:
Modeling, Simulation, Information Technology and Processing (TA 11)
└ Information Processing (TA 11.4)

Accelerating Communication-Intensive Applications via Novel Data Compression Techniques, Phase II Project

U.S. WORK LOCATIONS AND KEY PARTNERS

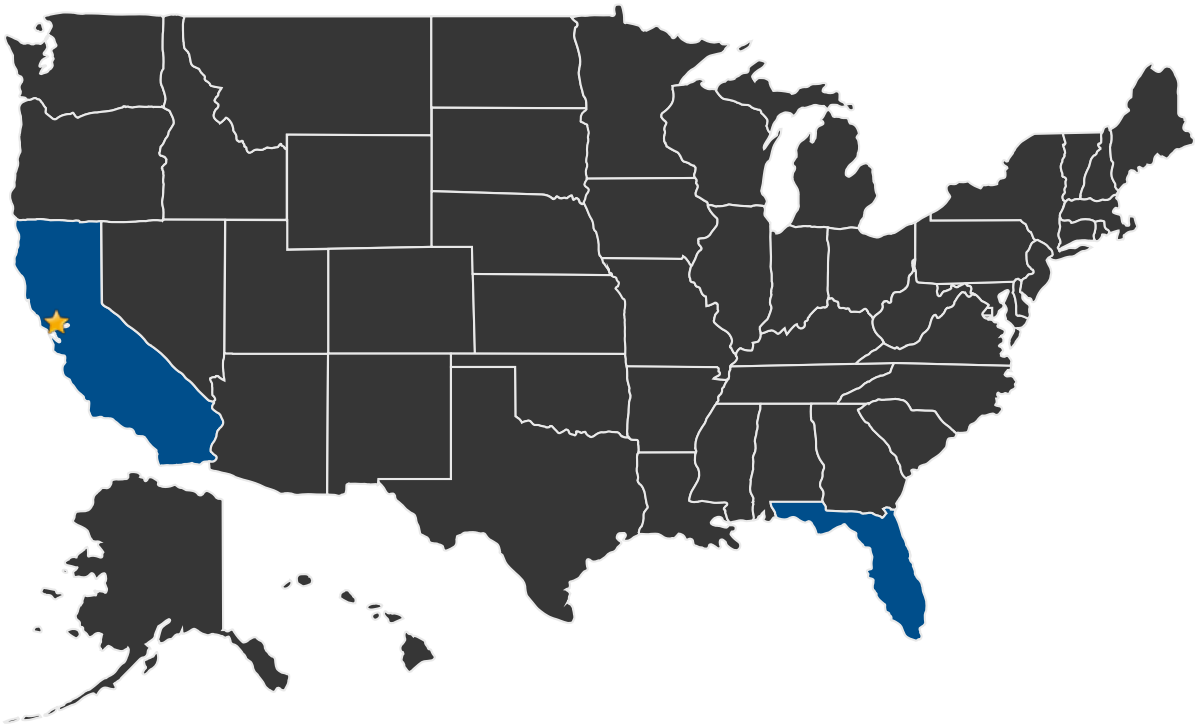

U.S. States With Work

Lead Center:
Ames Research Center

Other Organizations Performing Work:

- Accelogic, LLC (Weston, FL)

PROJECT LIBRARY

Presentations

- Briefing Chart
  - (http://techport.nasa.gov:80/file/23043)

IMAGE GALLERY

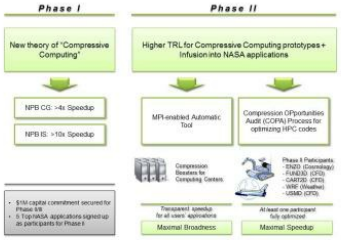

Accelerating Communication-Intensive Applications via Novel Data Compression Techniques, Phase II

DETAILS FOR TECHNOLOGY 1

Technology Title

Accelerating Communication-Intensive Applications via Novel Data Compression Techniques

Potential Applications

The impact of this technology on NASA is high and very broad. The fact that, in the short time span of the Phase I project, five of the Top NASA HPC applications have submitted letters of intent to Accelogic to become Phase II participants, speaks vastly about the significance of this technology for NASA. One of these participants mentions in his letter of intent that this technology "can be considered critical in achieving the next generation of so-called exascale software applications." Another participant mentions that a rough audit of their codes in production runs suggests that speedup factors on the order of 2-4 would be reasonable to expect across the spectrum of their user's cases. Such speedups would represent new opportunities for the code users to explore new larger science problems, faster and at a lower cost. The impact of this compression technology spans most areas of importance to NASA's scientific mission, including aerospace, weather forecasting, combustion, climate research, and chemistry, among many others.