ParaMEDIC

Parallel Metadata Environment for Distributed I/O and Computation

ParaMEDIC is a high-performance and portable framework to decouple computation and I/O in applications that require large quantities of both resources simultaneously. While there are multiple institutes that provide either lots of compute power or large storage resources, very few centers manage to meet both. Thus, computation has to be performed at one site and the entire generated data moved to another site for storage. This model is clearly inefficient, especially for worldwide distributed environments.

The ParaMEDIC Framework

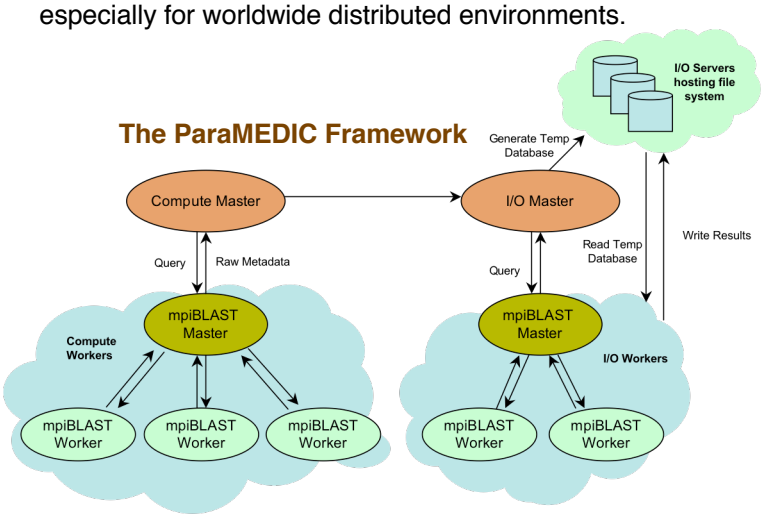

ParaMEDIC provides an easy-to-use plug-in framework that allows different applications to easily convert their output to orders-of-magnitude smaller metadata at the compute sites, transfer the metadata to the storage site, and re-convert the metadata to the actual output at the storage site, thus reducing the amount of data communicated substantially.

ParaMEDIC Architecture

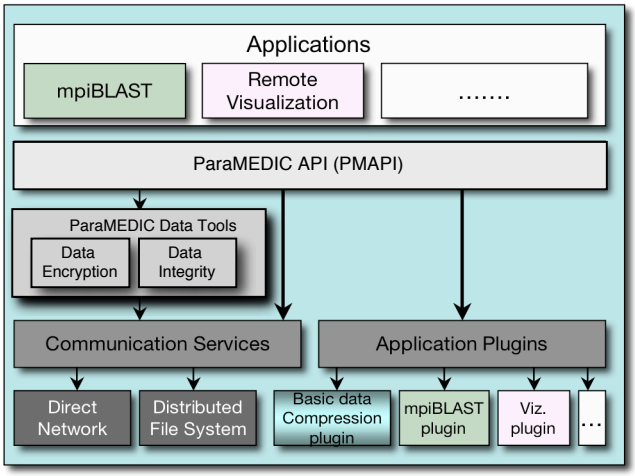

ParaMEDIC with mpiBLAST

mpiBLAST is a freely available, open-source parallelization of the popular genomic sequence search toolkit NCBI BLAST. While mpiBLAST has shown to achieve high performance in clusters with fast local file-systems, its I/O processing remains a concern for scalability in systems having limited I/O capabilities such as those using distributed file-systems spread across a wide-area network.

ParaMEDIC improves the performance of mpiBLAST by partitioning the worker processes into compute and I/O workers. Compute workers, instead of directly writing their output to the distributed file-system, plug into the ParaMEDIC framework that allows them to convert their output to orders-of-magnitude smaller meta-data, and send it to the I/O workers. I/O workers, which physically reside closer to the storage, process this meta-data to re-create the actual output.

Performance of ParaMEDIC vs. mpiBLAST on Teragrid infrastructure

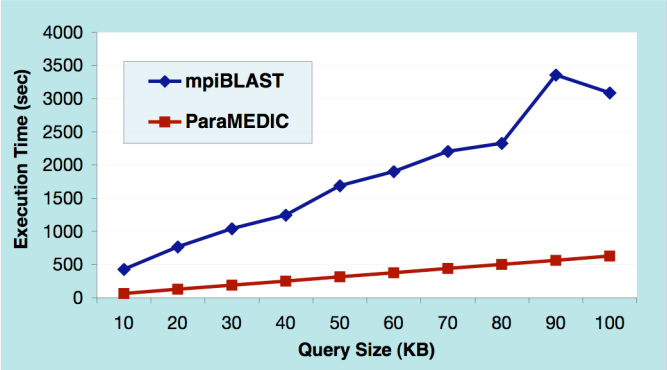

ParaMEDIC powered mpiBLAST has been shown to achieve close to a **five-fold speed** up on the Teragrid infrastructure, where the distributed file-system (GPFS) is hosted at the San Diego Supercomputer Center (SDSC), while the compute resources are distributed all over the U.S., including University of Chicago (Illinois), Purdue University (Indiana), Texas Advanced Computing Center (Texas), and others.

Argonne NATIONAL LABORATORY

VirginiaTech

ParaMEDIC

Parallel Metadata Environment for Distributed I/O and Computation

Storage Challenge Finalist at Supercomputing '07

Goals of the Storage Challenge

Sequence-search all the 567 microbial genomes (that have been completed to date) against each other in order to discover missing genes via mpiBLAST sequence-similarity computations.

Generate a complete genome sequence-similarity tree, based on the above results, in order to structure the sequence databases so as to enable pruning of the sequence-search space and, thus, accelerate the sequence-search process.

The enormous compute power and storage resource requirements for solving the above problems significantly reduced the potential supercomputing solutions that can be used. Consequently, we created a worldwide supercomputer that aggregates supercomputers and storage resources around the world.

Compute and Storage Requirements

Number of Processors utilized: 10,000+
Data generated: 1 petabyte

Storage Bandwidth Utilization

mpiBLAST (without ParaMEDIC): 2.4 to 80 Mbps
mpiBLAST (with ParaMEDIC): 5000 Mbps

Utilized close to 90% of the available storage bandwidth

Principal Investigators

Pavan Balaji, Argonne National Laboratory
(balaji@mcs.anl.gov)
Wu-chun Feng (a.k.a "Wu"), Virginia Tech
(feng@cs.vt.edu)

Partners

Tokyo Institute of Technology
SUN Microsystems
University of Chicago
Renaissance Computing Institute
Louisiana State University
San Diego Supercomputing Center

Storage Challenge Finalist Presentation

Argonne NATIONAL LABORATORY

Tuesday, Nov. 13th at 1:30pm
Where: A10/A11
Contacts: Pavan Balaji (ANL), Wu Feng (VT)

VirginiaTech

ParaMEDIC Demos

Showcasing ParaMEDIC with MPICH2, mpiBLAST and NetEffect 10-Gigabit Ethernet
Nov. 13th, 14th and 15th; 11:00 am -12 noon at ANL booth (#551)
All-day demos at Virginia Tech. booth (#2803)

MPICH2 NETEFFECT mpiBLAST