

Gordon

Fostering Computations on **NON-TRADITIONAL RESEARCH**

When SDSC's *Gordon* supercomputer debuted in early 2012, the system's architects envisioned that its innovative features—such as the first large-scale deployment of flash storage (300 terabytes) in a high-performance computer—would open the door to new areas of research.

"I view *Gordon* as a new kind of vessel, a ship that will take us on new voyages to make new discoveries in new areas of science," said Mike Norman, SDSC Director of *Gordon*'s principal investigator, prior to its launch.

Fast forward to 2013, and it's clear that *Gordon* already reached out to areas relatively new to advanced computing, such as political science, mathematical anthropology, finance, and even the cinematic arts.

Following are three examples of what might be considered "non-traditional" computational research assisted by *Gordon* during the past year.



SDSC's innovative *Gordon* supercomputer made its debut in 2012. (above)

Gordon installation in the SDSC Datacenter (left)



Large Scale Video Analytics

Virginia Kuhn, associate director of the Institute for Multimedia Literacy and associate professor in the School of Cinematic Arts at the University of Southern California (USC), has been using *Gordon* to search, index, and tag vast media archives in real time, applying a hybrid process of machine analytics and crowd-sourced tagging.

The Large Scale Video Analytics (LSVA) project is a collaboration among cinema scholars, digital humanists, and computational scientists from the IML (Institute for Multimedia Literacy); ICHASS (Institute for Computing in the Humanities, Arts and Social Sciences); and NSF's (National Science Foundation) XSEDE (eXtreme Science and Engineering Discovery Environment). The project is customizing the Medici content management system to apply various algorithms for image recognition and visualization into workflows that will allow real-time analysis of video.

"Contemporary culture is awash in moving images," said Kuhn, principal investigator for the research project. "There is more video uploaded to YouTube in a day than a single person can ever view in a lifetime. As such, one must ask what the implications are when it comes to the issues of identity, memory, history, or politics."

The LSVA project turned to *Gordon* for its extensive and easily accessed storage capacity, since video data collections can easily reach multiple terabytes in size. Work was initially performed using a dedicated *Gordon* I/O node, and later expanded to also include four dedicated compute nodes. "Persistent access to the flash storage in the I/O nodes has been critical for minimizing data access times, while allowing interactive analysis that was so important to this project," said Robert Sinkovits, director of SDSC's Scientific Applications Group.

The system enabled data subsets that were most heavily used to reside in areas that provided fast random access.

Predictive Analytics Using Distributed Computing

For operations on large data sets, the amount of time spent moving data between levels of a computer's storage and memory hierarchy often dwarfs the computation time. In such cases, it can be more efficient to move the software to the data rather than the traditional approach of moving data to the software.

Distributed computing frameworks such as Hadoop take advantage of this new paradigm. Data sets are divided into chunks that are then stored across the Hadoop cluster's data nodes using the Hadoop Distributed File System (HDFS), and the MapReduce engine is used to enable each worker node to process its own portion of the data set before the final results are aggregated by the master node. Applications are generally easier to develop in the MapReduce model, avoiding the need to directly manage parallelism using MPI, Pthreads, or other fairly low-level approaches.

Yoav Freund, a computer science and engineering professor at UC San Diego, specializes in machine learning, a relatively new field of research which bridges computer science and statistics. A recognized authority in Big Data analytics, Freund recently taught a graduate-level class in which students used a dedicated Hadoop cluster on *Gordon* to analyze data sets ranging in size from hundreds of megabytes to several terabytes. While student projects were diverse—analysis of temperature and airflow sensors on the UC San Diego campus, detection of failures in the Internet through the analysis of the communication between BGP (Border Gateway Protocol) nodes, and predicting medical costs associated with vehicle accidents—they had one need in common: the need for rapid turnaround of data analysis.

"Given the large data movement requirements and the need for very rapid turnaround, it would not have been feasible for the students to work through a standard batch queue," said Freund. "Having access to flash storage greatly reduced the time for random data access."

A *Gordon* I/O node and the corresponding 16 compute nodes were configured as a dedicated Hadoop cluster, with the HDFS mounted on the solid state drives (SSDs). To enable experimentation with Hadoop, SDSC also deployed MyHadoop, which allows users to temporarily create Hadoop instances through the regular batch scheduler.



Robert Sinkovits is the Director of SDSC's Scientific Applications Group and has responsibility for ensuring that data-intensive problems can make effective use of this innovative flash memory-based system.

Modeling Human Societies

Doug White, a professor of anthropology at UC Irvine, is interested in how societies, cultures, social roles, and historical agents of change evolve dynamically out of multiple networks of social action. These networks can be understood using graph theory, where individuals are represented as vertices and the relationships between them are represented using edges.

In the 1990s, White and his colleagues discovered that networks could be meaningfully analyzed by a new algorithm that derives from one of the basic theorems in graph theory: the Menger vertex-connectivity theorem announced in 1927.

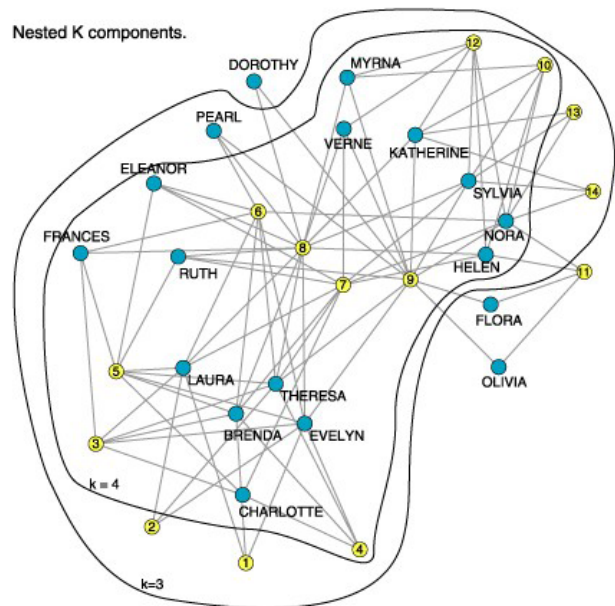
Computer scientists had previously coded Menger's edge-connectivity theorem, known as the Max-Flow (Ford-Fulkerson) network algorithm, to optimize the routing of traffic. But Menger's theorem also proved that vertex-connectivity, defined as the number of independent paths between a pair of vertices that share no vertices (other than the end points) in common, is equivalent to the number of vertices that need to be removed from the graph so that there is no longer a path between the two vertices.

Using Menger's theorem, it is possible to identify cohesive groups in which all members remain joined by paths as vertices are removed from the network. Finding the boundaries of cohesive subgroups, however, remain a challenge for computation, far more challenging than finding cliques of fully connected sets of vertices within a graph. White and his team are currently testing their recently developed software on synthetic networks constructed from well-known models, and plan to apply their methods to data sets derived from networks of co-authorship, benefits of marriage alliances, and splits in the scientific citation networks in periods of contention.

"The most computationally difficult step in the team's design of new high-performance software requires calculating the number of vertex-independent paths between all pairs of vertices in the network," said White.

Robert Sinkovits, director of SDSC's Scientific Applications Group, addressed scalability issues in the parallel algorithm and ported the application to *Gordon's* vSMP nodes following guidelines provided by ScaleMP. For larger graphs, such as a 2400 vertex network built using the Watts-Strogatz small-world model, near linear scalability was achieved when using all 256 cores (an estimated 243x speedup relative to one core).

"We're coordinating teams at UC Irvine, UC San Diego, the Santa Fe Institute, and Argonne National Laboratory so that after 130 years, we finally put comparative research on a solid footing of replicable scientific findings that are comparable from small to large and super-large datasets," said White.



K-cohesive levels of "southern women" neighborhood parties. K varies from 1 to 4 (connected to 4-connectivity). Image courtesy of James Moody and Douglas R. White