HP 9000 SERVERS FOR ADVANCED SCIENTIFIC RESEARCH

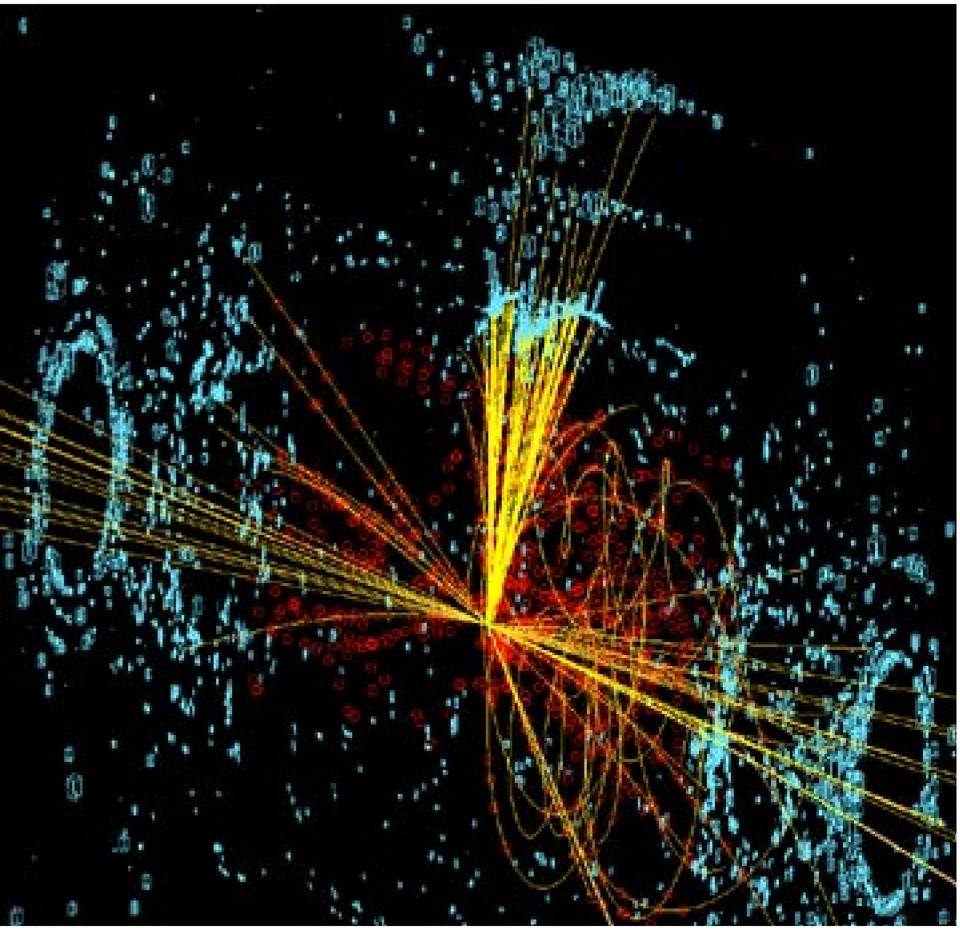

Caltech CACR Provides the Computing Power for Developing an Experimental Petabyte-Sized System of Networked Databases

Caltech's Center for Advanced Computing Research (CACR) was formed in 1995 to enhance Caltech's multidisciplinary activities in high-performance computing and to provide access to state-of-the-art computing resources that enable scientific engineering computations not previously possible.

"Our recent experience has shown that large systems such as the HP technical server can successfully be used as general-purpose computing resources for the high-energy physics community."

Harvey Newman
Professor of Physics, Caltech

The Challenge:

Provide the computer performance to gather, store, retrieve and analyze massive quantities of data and make the data available quickly to thousands of research physicists worldwide for distributed processing.

The Solution:

Use the 256-processor HP technical server at the Caltech Center for Advanced Computing Research (CACR) as the central computing engine for modeling a globally interconnected, object-oriented database network.

Impact:

In building a prototype for a petabyte-sized network for research scientists, Caltech CACR is also designing a prototype architecture that may serve the next generation of corporate computing.

As part of this mission, computational scientists and engineers at Caltech are working directly with Hewlett-Packard's computer scientists and developers to ensure that the computing power is available to meet the needs of the scientists and engineers on the front lines of advanced research.

The California Institute of Technology (Caltech) and the Jet Propulsion Laboratory (JPL) have an HP technical server at Caltech which is a 256-processor system running a single-system image. Designed to support advanced research applications, the system comprises sixteen 16-processor nodes, a distributed, shared-memory architecture, 64 GB of memory and more than one terabyte of disk storage. Its peak performance is 184 GFLOPS. The HP server's scalable parallel architecture groups processors, memory and I/O components into hypernodes for a balanced approach to scalability. The acquisition of the HP technical server was funded by the National Science Foundation (NSF), the United States Department of Energy (DoE) and JPL.

Scientists at CACR are using the HP system for a number of critical research applications, including a high-energy physics project being constructed to search for some of the last remaining unknown particles in the universe.

High-Energy Physics Search for the Holy Grail

Previous generations of experiments at the "high-energy frontier" of their day revealed the structure of the atom with its nucleus and orbiting electrons. Then the nucleus was discovered to be composed of protons and neutrons. In the 1960s, it was discovered that protons and neutrons are themselves made up of smaller particles called quarks. The heaviest quark, or "top quark," was discovered in 1995 by scientists using a particle accelerator in which protons and electrons collided at approximately two trillion electron-volts (TeV) of energy. ***[See Project Background Sidebar]***

Now scientists are hoping that the Large Hadron Collider (LHC), in which particles will collide at energies up to 14 TeV, will help them uncover another particle that exists in theory but has not been verified through experiments—the Higgs particle. This particle, named for a British scientist who hypothesized its existence in the 1960s, is believed to be responsible for giving other particles their mass. Called "the Holy Grail of particle physics," the Higgs particle is the key missing piece of our present understanding of the laws of nature, known as the Standard Model. Finding the Higgs—or proving that it does not exist—is the primary goal of experiments to be conducted at CERN, the European Laboratory for Particle Physics in Geneva, Switzerland, beginning in 2005.

Research Requirements for Massive Computing Power and Bandwidth

Collisions inside the LHC are expected to occur at the rate of about 800 million per second. Of these millions of events, only about 100—or .0000001 percent—are “interesting” to physicists exploring the frontier of the highest-collision energies and the smallest-distance scales. Part of the job of the detector is to identify and select the interesting events for recording, a process called “triggering.” As particles from interesting events travel through the detector, data about these particles will be recorded at the rate of approximately one megabyte per event. Data are thus coming out of the detector at 100 megabytes per second. The two largest experiments to be conducted at the LHC, Compact Muon Solenoid (CMS) and ATLAS, are each expected to generate a petabyte of data per year during their 20-year lifespans (1 petabyte = 10^9 megabytes, or 2^{50} bytes).

Raw event data at the LHC will amount to several petabytes each year. The data are already highly compressed when they emerge from the triggering process, and they must be stored in their entirety. From these data, reconstructions of physics “objects,” such as tracks, clusters and jets, will take place in near real time, requiring processing power of some

PROJECT BACKGROUND

In order to understand the fundamental nature of matter, scientists construct particle accelerators in which electrically charged particles such as protons, traveling at near the speed of light, collide at very high energies. These high-energy collisions, termed “events,” produce more particles, which fly away from the collision point in collimated sprays called “jets.” These jets pass through “detectors,” massive computerized instruments that record all data about a particle, data which may be used to determine properties such as its energy, momentum, mass and electrical charge. By analyzing these data, scientists can reconstruct the event that produced a given particle.

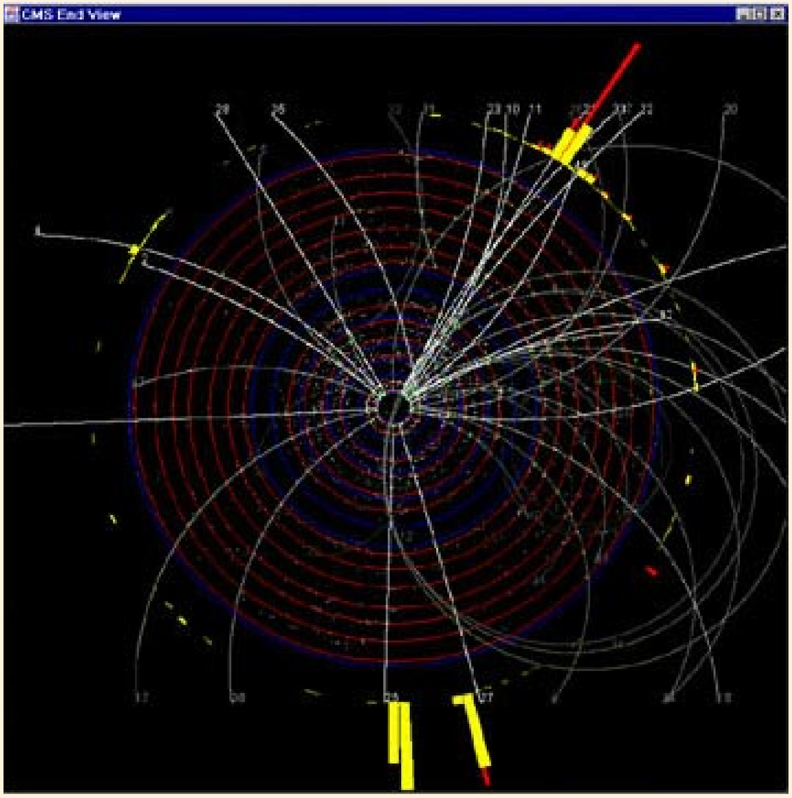

Figure 1: ***An end view of the CMS detector to be used at the LHC, showing the structure of the Tracker (red and blue rings) together with individual “hits” (green dots) and reconstructed particle tracks (white and gray lines with numbers). Particle tracks of lower momentum (gray) curl in the magnetic field from the large solenoid magnet (not shown) that surrounds the Tracker detector. The Electromagnetic Calorimeter crystal tower energies are shown in yellow, with reconstructed energy clusters in red. The CMS detector weighs 14,500 metric tons and is 22 meters long and 15 meters in diameter.***

CUSTOMER AT A GLANCE

The California Institute of Technology (Caltech), located in Pasadena, California, is one of the world's leading research centers. Caltech has 900 undergraduate and 1,100 graduate students, and a faculty of about 280 professorial members and nearly 500 research members. The Institute also manages NASA's Jet Propulsion Laboratory (JPL) and operates six other astronomical, seismological and marine biology facilities. Over the years, 26 Nobel Prizes have been awarded to faculty members and alumni of Caltech.

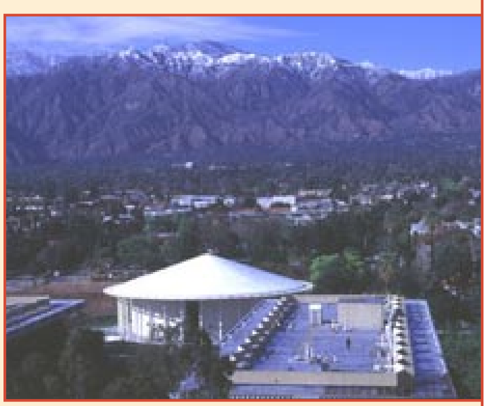

10 trillion instructions per second (TIPS), or 10^7 MIPS. Each reconstructed object will add about 100 k (10^7) bytes to the already prodigious data volume. By the time the LHC program is mature, CERN projects that the data volume will be about 100 petabytes, stored on both tape and disk.

Managing LHC data also requires distributed computing, and serving and processing data across a hierarchy of networks. The CMS and ATLAS experiments will each involve more than 1,700 physicists spread among 150 institutions in more than 30 countries. Physicists located at collaborating institutes will need the same level of access to the latest data as those located at CERN. This requirement will necessitate continuous transport of the data across the network, or rapid decisions on where to execute analysis queries (i.e., whether to move the data across the network to a local compute resource or move the query/application to the data and ship the results back to the physicist). The offline software (reconstruction, analysis, calibration, event display, etc.) will be developed throughout the lifetime of the experiments. As a result, many developers and users will be involved in each software component.

The search for the Higgs particle thus represents an unprecedented challenge for computation, data storage, access and networking technology.

HP Technical Server Answers Physics Data Management Challenge

To meet this challenge, Harvey Newman of Caltech and Julian Bunn of CERN approached CACR and Hewlett-Packard about a collaborative effort, a project called GIOD (pronounced GEE-ode), or Globally Interconnected Object Databases. The project will model the computing environment required for data analysis at the LHC. According to Julian Bunn, "We knew the Center for Advanced Computing Research at Caltech had recently acquired this very large 256-CPU HP technical server as part of a joint project with JPL. Added to that, CACR had a working production HPSS hierarchical storage management system. The integration of that system with the HP technical server, together with an object database, made a very good infrastructure for a project of this sort."

The CACR is particularly well positioned to look at the distributed processing issue, explained Harvey

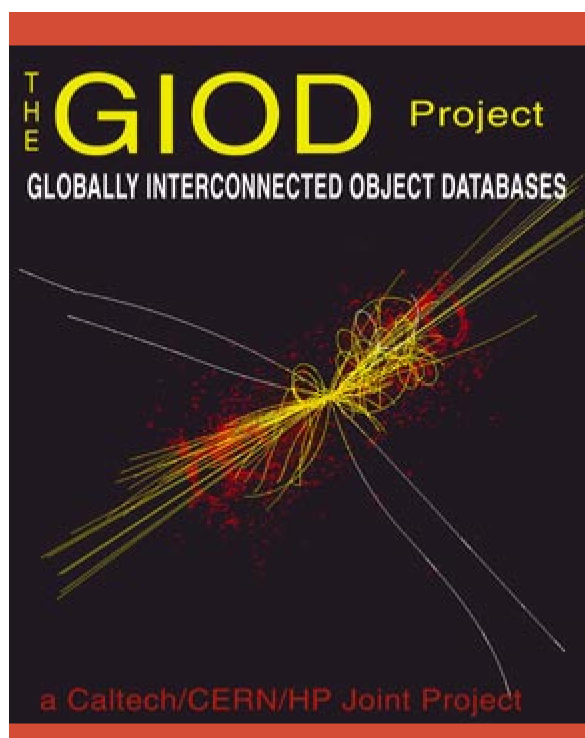

Figure 2

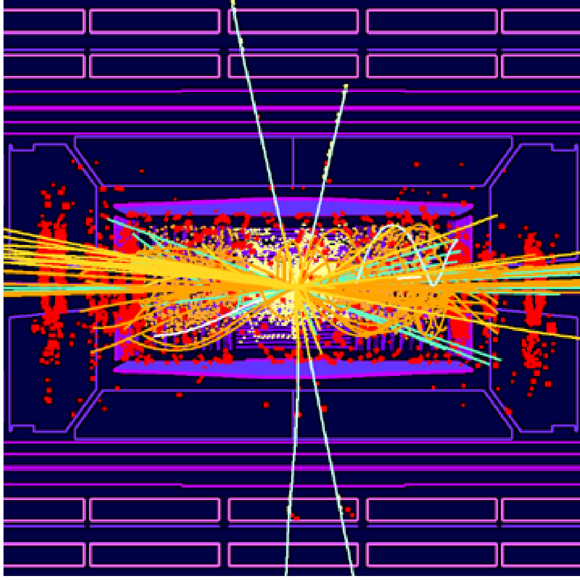

Figures 2 and 3: Events are generated and tracked through the detector.

Newman, professor of physics at Caltech and project co-leader. “We’ve been responsible for international networking for the high-energy physics community, and we had an emphasis on developing and managing networks. As an outcome of this history, we now manage the trans-Atlantic link from the United States to Europe for our collaboration, and we collaborate with CERN in managing and operating this link. That gives us a strong basis for collaborating with CERN on these database projects.”

Event data from both CMS and ATLAS experiments will be stored and distributed using Object Database Management Systems (ODBMS), coupled with hierarchical storage management systems. Project leaders Bunn and Newman agree that this is the only approach that can accommodate the complexity and sheer volume of the data, the geographical dispersion of the collaborating institutes, and the large number of participating physicists.

In addition, regional computing centers, interconnected by very fast network links, will serve data and resources to physicists located at nearby institutes. Typically, they will offer a replica of the event database, the master of which is populated at CERN directly from the data acquisition systems in the experiments.

The goal of the GIOD project is to construct a prototype LHC regional computing center. In the present model, the HP technical server sits at the center of a web of hardware and software that includes HPSS, a hierarchical storage management system; the Objectivity/DB Object Database Management System (ODBMS); and various high-speed local area and wide area network infrastructures.

HP Technical Server Has Power and Scalability to Meet the Most Demanding Requirements

One of the functions of the GIOD project will be to demonstrate the ability of commercially available computing systems such as the HP technical server to serve the high-energy physics community. So far the results have been most encouraging.

In one set of experiments, the HP technical server was used to carry out large-scale Monte Carlo simulations of the CMS detector. In a typical run, 64 simultaneous jobs were executed in a 64-CPU system partition, each of which generated 100 fully detailed simulations of the decay of the postulated Higgs particle into four muons. The 6,400 events that resulted from this run were generated and tracked through the detector in a record time of two hours. Bunn said that recent tests with 256 individual clients reading from and writing to an Objectivity database showed that “the system scales very nicely up to about 200 clients, and by simply increasing the number of I/O channels, we could go even higher.”

The HP technical server is also being used to run traditional high-energy physics code that is used for simulating particle collisions inside the detectors. “The HP server is very good at running very large numbers of simulations in production,” said Bunn. “We have generated about a million simulated events, and we are using those simulated events to populate terabyte-scale object databases.” The Objectivity ODBMS is then called upon to read in the simulated events and create objects in the database that correspond to the raw data, and then to reconstruct the data, which are then in turn put as objects into the object database.

At a review of the GIOD project, Caltech’s Newman said, “Access to the HP technical server . . . is making it possible to simulate the Higgs events along with the copious

Figure 3

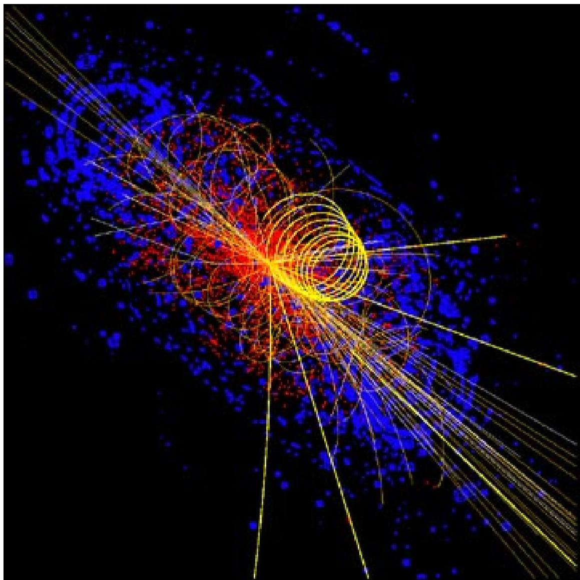

backgrounds from standard multijet events, with a fidelity and statistical precision that was previously inaccessible.”

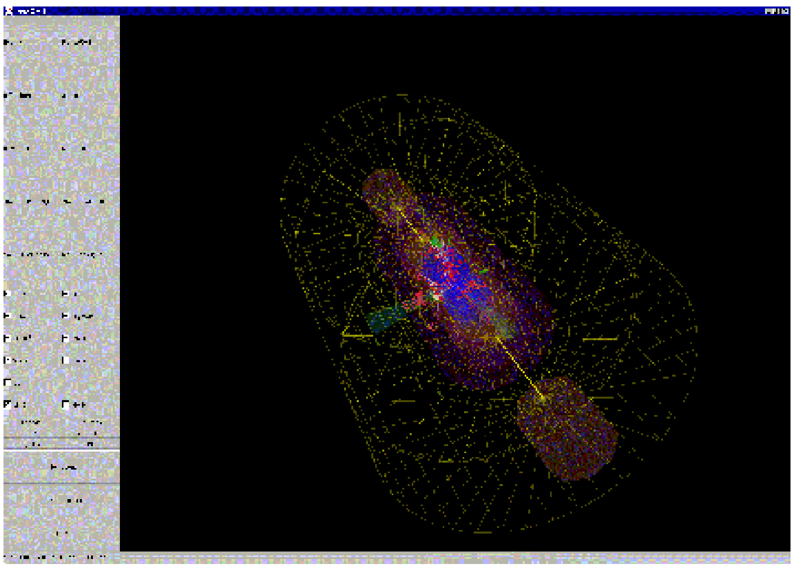

Figure 4

Project scientists have also developed a 3D applet for viewing the events, so that scientists can actually see the tracks and the clusters. Figures 4 and 5 are screen shots from the 3D event viewer applet, showing a 3D rendering of almost the whole of the CMS detector. The Inner Tracker detector is depicted by the innermost faint blue structure at the center of the detector. Digitizings (“hits”) are shown by red points in this detector. The white and gray lines are the reconstructed tracks. Immediately surrounding the Tracker is the Electromagnetic Calorimeter (ECAL), shown as a faint green tube, with shower deposits (from decaying photons and electrons) in the crystals appearing as green towers pointing away from the center of the whole detector. Surrounding the ECAL is the Hadronic Calorimeter, in faint purple, with shower deposits (from decaying hadrons like pions) represented as pink towers. The outermost layer shown is a cylinder representing the solenoid. Reconstructed “jets,” high momentum sprays of electromagnetic and hadronic energy, are shown as large towers pointing from the center of the detector and labeled “JET.”

The CMS event viewer “is a very useful tool for physicists,” said Bunn. It allows the user to remove or add all the subdetector detail piecewise and to rotate and zoom the picture. It will eventually allow the user to select individual objects in the display and interrogate the Object Database about their properties.

Meeting the Corporate Data Management Challenge

The GIOD project is helping to define the further development of object-oriented database management systems. Industry, academic and scientific research will also benefit from better network management systems, improved tools for storing and analyzing large amounts of complex data, and better tools for collaborative software development. “We are pleased to be participating in a project that will help define the next generation of corporate computing,” said Greg Astfalk, chief scientist, HP High Performance Systems Division.

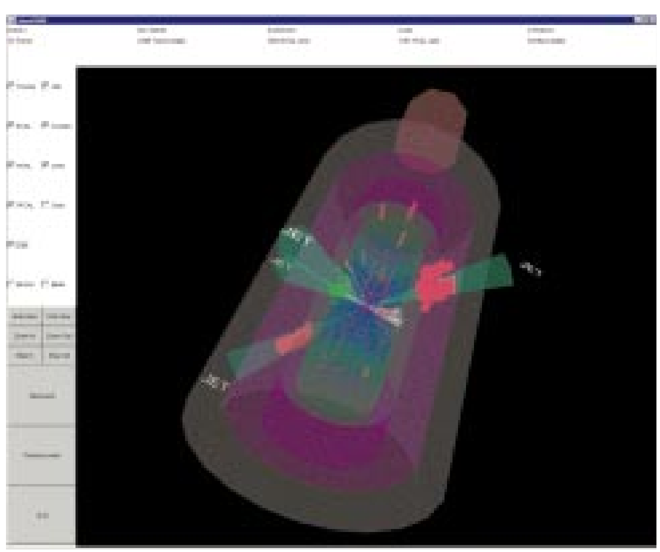

Figure 5

“High-energy physics always seems to be at the far leading edge of data storage and processing technology,” noted Newman. But this time the edge is not so far out. Current projections are that the volume of U.S. corporate data will exceed 500 petabytes within six years. Therefore, within a very few years, medium-sized to large corporations will face petabyte-sized data management challenges.

LHC will not be online until 2005. The GIOD project in its present form will end in 1999, but work on distributed object databases and networked computing systems will continue at CERN, at Caltech, elsewhere in Europe, and in the United States. This means that by the time corporations are faced with managing petabyte-sized distributed computing networks, the technology will be ready.

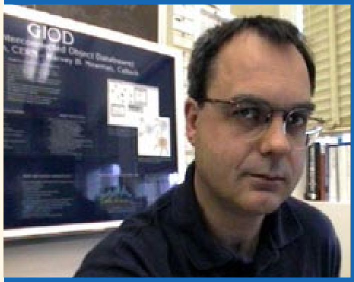

“We have generated about a million simulated events, and we are using those simulated events to populate terabyte-scale object databases.”

Julian Bunn
CERN

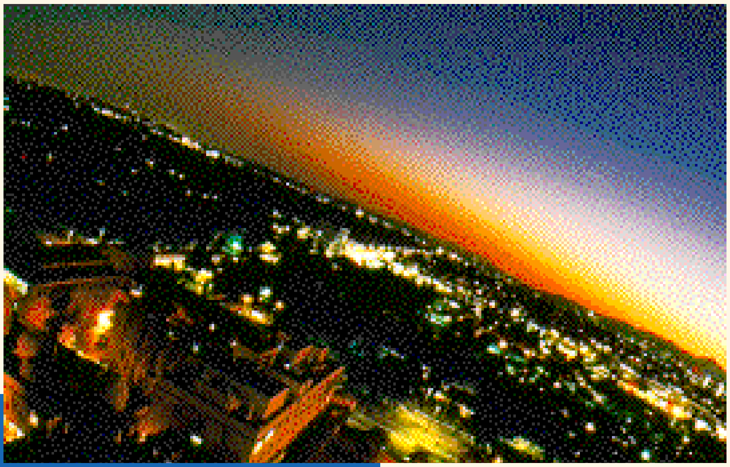

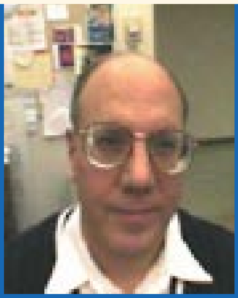

“Access to the HP technical server . . . is making it possible to simulate the Higgs events along with the copious backgrounds from standard multijet events, with a fidelity and statistical precision that was previously inaccessible.”

Harvey Newman, Ph.D.
Professor of Physics, Caltech

For the location of the nearest Hewlett-Packard sales office, call:

UNITED STATES OF AMERICA:
+1 800 637 7740

CANADA:
Hewlett-Packard Ltd.
5150 Spectrum Way
Mississauga, Ontario L4W 5G1
+1 905 206 4725

JAPAN:
Hewlett-Packard Japan, Ltd.
Korakuen Shinjuku Bldg.,
4-15-7
Nishi-Shinjuku, Shinjuku-ku,
Tokyo, 160 Japan
+81 3 5371 1342

LATIN AMERICA:
Hewlett-Packard
Latin American Region
Headquarters
Waterford Building, 9th Floor
5200 Blue Lagoon Drive
Miami, Florida 33126 USA
+1 305 267 4220
Refer to country phone numbers

AUSTRALIA/NEW ZEALAND:
Hewlett-Packard Australia Ltd.
31-41 Joseph Street
Blackburn, Victoria 3130
Australia (A.C.N. 004 394 763)
+61 3 9272 2895

ASIA PACIFIC:
Hewlett-Packard
Asia Pacific Ltd.
17-21/F, Shell Tower
Times Square
1 Matheson Street
Causeway Bay
Hong Kong
+8522 599 7777

EUROPE/AFRICA/MIDDLE EAST:
Hewlett-Packard S.A.
150, Route du Nant-d'Avril
CH-1217 Meyrin 2
Geneva, Switzerland
+41 22 780 81 11

EUROPEAN MULTICOUNTRY:
+41 22 780 81 11

MIDDLE EAST AND AFRICA:
+41 22 780 71 11

EUROPEAN HEADQUARTERS:
+41 22 780 81 81
Refer to country phone numbers

For direct country contact, call:

ARGENTINA:
+541 787 7145

AUSTRIA:
+43 1 25 000 0

BELGIUM AND LUXEMBOURG:
+32 2 778 31 11

BRAZIL:
+5511 7296 8000

CHILE:
+562 203 3233

DENMARK:
+45 45 99 10 00

EAST CENTRAL EUROPE, CIS, AND YUGOSLAVIA:
+43 1 25 000 0

FINLAND:
+358 0 887 21

FRANCE:
+33 1 69 82 60 60

GERMANY:
+49 7031 140

GREECE:
+30 1 689 644

HUNGARY:
+36 1 252 7300

ICELAND:
High Performance Systems hf.
+354 1 67 10 00

IRELAND:
+353 12 88 33 99

ISRAEL:
Computation and Measurement
Systems (CMS) Ltd.
+972 3 5380 333

ITALY:
+39 2 92122770

MEXICO:
+525 326 4600

NETHERLANDS:
+31 20 547 6911

NORWAY:
+47 22 7356 00

POLAND:
+48 22 608 77 00

PORTUGAL:
+351 1301 7343

RUSSIA AND THE CIS, EXCL. UKRAINE:
+7 095 923 5001

SLOVENIA:
+38 61 55 84 72

SOUTH AFRICA:
Hewlett-Packard
South Africa (Pty) Ltd.
+27 11 806 1000

SPAIN:
+34 1 631 1600

SWEDEN:
+46 8 444 2000

SWITZERLAND:
+411 735 7111

TURKEY:
+90 312 468 8770

UNITED KINGDOM:
+44 1344 369231

VENEZUELA:
+582 239 4133

HEWLETT® PACKARD

For more information, contact any of the HP worldwide sales offices or HP Channel Partners (in the U.S. call 1-800-637-7740) or look for HP on these World Wide Web sites:

- www.hp.com/go/techservers
- www.hp.com/go/workstations

For more information about Caltech, go to www.caltech.edu

Printed in USA 1/99

5968-3499E