

# CONTEXT

---

***In this issue:***

---

UNC-CH sharpens  
Hubble images — *page 4*

---

Intelligent manufacturing  
from Ford Alpha — *page 6*

---

Excellence at NASA Lewis — *page 8*

---

MAFIA — solving electromagnetic  
problems — *page 11*

---

Paris Observatory — *page 12*

---



## **C4/XA — newest C Series system announced**

Convex is setting a new standard in supercomputing performance with a new family of high-end supercomputers in the C Series product line. The C4/XA is capable of exceeding one billion floating point operations per second (GFLOPS) per CPU — a level formerly achieved only with expensive liquid-cooled systems. The C4/XA is compatible with previous C Series software and with a broad range of software applications optimized for the system's architecture.

*Continued on page 2*



**C4/XA...**

*continued from page 1*

Initial C4/XA customers include BMW Rolls-Royce in Dahlewitz, Germany, CEDEX in Madrid, Spain, and National Chiao Tung University in Taiwan. BMW Rolls-Royce will use the system for structural analysis and computational fluid dynamics. Mr. B. Lenk of BMW Rolls-Royce said, "The C4/XA high-end vector parallel system fits seamlessly into our UNIX-based client/server environment, and helps to shorten the design cycles for a new generation of fuel efficient aircraft engines."

At CEDEX (Centro de Estudios y Experimentacion de Obras Publicas), a public works research and experimentation agency of the MOPTA public works and transportation department, the C4/XA was selected as a central resource to do sophisticated computing problems in civil engineering. "The C4/XA will be used as a multipurpose host for performance critical applications in our distributed network of UNIX systems," explained Juan-Antonio Vera of CEDEX.

The C4/XA systems are ideal for scientific and technical users who have high compute and data management performance requirements, and provide a high-performance upgrade path for C2 and C3 users. For example, the National Chiao Tung University in Taiwan is upgrading its Convex C240 to a C4/XA system which will operate as a high performance compute server for the campus. The C4/XA system is available in four models ranging from one to four processors, and are field-upgradeable to meet growing needs.

The eXtended Architecture (XA) of the new C4 combines sophisticated chip technology

with an advanced vector subsystem design that incorporates six vector functional units per processor. "This architectural enhancement can yield more than a six-fold increase in processor performance over previous high-end C Series models," explained Steven J. Wallach, Convex senior vice president of technology.

"The new processor feature results in exceptional single processor performance, as indicated by the Linpack 1000x1000 benchmark results of 647 MFLOPS (64-bit), and convolution results of 1.4 GFLOPS (32-bit)," Wallach added. "In addition, C4/XA uses Convex's patented ASAP (Automatic Self-Allocating Processors) technology to deliver fine-grained parallelism without diminishing overall system throughput in multijob environments."

The C4/XA supports hardware-based IEEE floating point operations. This capability allows binary file compatibility with workstations and other RISC-based systems, significantly easing integration into heterogeneous client/server environments.

Vector processing systems, such as C4/XA, with applications optimized over the past decade, will continue to provide the highest performance on a specific class of compute intensive applications. Parallel processing systems, such as Convex's Exemplar systems, can deliver innovative solutions for a new class of parallel applications such as crash simulation.

So whatever your high-performance computing needs are, Convex provides the solution. ■



*The C4/XA system can exceed one billion GFLOPS per CPU—a level formerly achieved only with expensive liquid-cooled systems.*

# Exemplar — the right choice

"We've held out for a long time on our MPP [massively parallel processing] purchase because we wanted to have a productive parallel computer," says Dr. Danilo Zevrtanik, the director of Jozef Stefan Institute in Ljubljana, Slovenia. "Back in 1992, Convex showed us their plans and convinced us to wait for their product. Now I am convinced that we have made the right choice."

Jozef Stefan Institute (JSI) is one of the first Convex customers to install the new Convex Exemplar SPP system. Other customers to receive the first shipments include the National Center for Supercomputing Applications (NCSA); the Navy Command, Control and Ocean Surveillance Center (NCC/OSC); NKK, a Japanese steel manufacturer; and the University of Erlangen in Germany.

Introduced in March, Exemplar, the new line of multi-purpose parallel servers, combines desktop compatibility and productivity with supercomputing performance. The Exemplar systems are based on Hewlett-Packard Company's industry-leading PA-RISC technology and will be sold jointly by Convex and HP.

"Convex's Exemplar truly delivers the best of two worlds — the ease of use and applications availability of workstations and the performance of supercomputers," says Dr. Kazuo Kunioka, Managing Director of NKK. "Now the power and the applications to solve today's most complex problems are available to a range of mainstream business users as well as to the research community."

Besides being one of the first users of Exemplar, NKK has also participated in Convex's Preferred Partners Program. Customers who represented a mix of industries were involved in the development of Exemplar from its earliest days. By testing Exemplar system technologies, these customers provided critical feedback that was used to refine and perfect the product. Some of the other Preferred Partners included The Scripps Research Institute, and NCSA.

The Exemplar family includes two models: the SPP100/CD (Compact Design) workgroup server configured with two to 16 processors for up to 3.2 GFLOPS of peak performance and the SPP100/XA (eXtended Architecture) scalable system with eight to 128 processors for over 25 GFLOPS of peak performance.

"Convex has married the best attributes of supercomputers, MPP and the best of desktop workstations in the Exemplar servers," says Steven J. Wallach, Convex senior vice president of technology. "This is the same 'best of both worlds' approach we used to create the world's first affordable supercomputer a decade ago — combining the performance of supercomputers

with the ease of use of minicomputers. The Exemplar systems represent the first step in our teraflop initiative."


"Convex has set a new benchmark in the product development process," says Matt Blanton, Convex's vice president of advanced development. "It has been a case study in concurrent engineering, calling on the involvement and expertise of every area of Convex. And the contribution of our Preferred Partners Program gave us an invaluable, real-world evaluation every step of the way."

Blanton adds that Hewlett-Packard, too, was an ideal development partner. "Our strategic relationship with HP has been incredible," says Blanton. "It gets the highest marks possible for this kind of partnership."

The results of Convex's partnerships in developing SPP have been applauded by its initial users.

Larry Smarr, director of NCSA, says, "We see the addition of the Exemplar SPP to NCSA's metacomputing architecture as a crucial step in allowing corporate America to reap the benefits of parallel computing."

Dr. Claus-Uwe Linster of the University of Erlangen adds, "With these new scalable systems, the University can address all current needs today and easily expand to meet our needs in the future." ■



Sunlight reflects off Endeavour's aft windows and the shiny HST prior to its post servicing deployment near the end of the 11-day STS-61 mission. (Photo courtesy of NASA Space Telescope Science Institute.)

## UNC-CH sharpens Hubble images with Convex

*The following article is adapted from the Spring/Summer 1994 issue of OIT Review (ISSN #1071-3360), published by the Office of Information Technology of the University of North Carolina at Chapel Hill. Address all queries about the Review to the editor: (jim\_barnes@unc.edu). Adapted and reproduced with permission.*

When a space shuttle crew released the Hubble Space Telescope (HST) into its orbit 280 miles above the earth's surface on April 24, 1990, NASA and the world of astronomy eagerly awaited data from the most expensive scientific instrument ever sent into orbit.

The \$1.5 billion telescope, with its nearly eight-foot-diameter primary mirror, was capable of collecting images of deep space unattainable to ground-based

telescopes. From its orbit, the Hubble would not have its images filtered through the miles of atmosphere that shimmer before traditional ground-based instruments. Thus, astronomers were anticipating images of unprecedented clarity and depth, images that would help them in their search for distant stars and the age of the galaxy itself.

So on June 27, 1990, when NASA announced that a mirror flaw was distorting the images from the Hubble, thousands of scientists the world over joined in a collective groan. It seemed that this project, destined to spearhead new dimensions in deep-space viewing and measurement, would fall far short of its planned importance.

Dr. Bruce Carney, the Dean E. Smith Professor of Astronomy at UNC-CH, was one of the disappointed scientists. He was part of a team investigating the ages of the oldest stars throughout the galaxy. "We were given time on the Hubble's first cycle [of observation time] before we knew it wasn't going to work," Carney says. "So when we did learn — much to our horror — about the performance of the telescope, we realized our project was in jeopardy."

NASA traced the problem with the Hubble imaging system to a faulty measurement when the optics company set an instrument to grind the eight-foot-diameter mirror. From a discrepancy less than the width of a human hair, the mirror produced blurry images that made age-dating the stars all but impossible.

According to Carney, the age of a star may be inferred from measurement of its color and brightness. Although mathematical algorithms could be used to deblur the Hubble's images, the known algorithms omitted crucial data such as the relative brightness of the stars. "The images could be made to look better," explains Carney, "but we couldn't learn anything quantitative from them."

*Continued on page 15*

## UNC pilots SAS on Convex

Besides distinguished astrophysical research, the Convex C3840 system at the University of North Carolina at Chapel Hill also supports users on the SAS system.

SAS, a product of the SAS Institute of Cary, North Carolina, is an integrated suite of software products that provides enterprise-wide information delivery, including tools to access, manage, analyze and present user data.

Sally M. Muller, who is the SAS consultant at UNC, says, "SAS has been far and away the most heavily used applications software at UNC. It is used by students, faculty and staff on a wide range of computing platforms."

Until recently, the Convex system was the only publicly available processor that did not support SAS. Since many Convex users needed the data management and statistical features of the SAS System, UNC was intrigued to hear that SAS Institute and Convex were developing a version of the product for Convex systems. UNC offered to become a beta test site.

Muller and her colleague Kenneth A. Hardy reported to the 19th Annual SAS Users Group International Conference, "The physical

memory of the Convex C3840 greatly enhances the speed of calculations for large problems and can allow users to solve very large SAS problems without receiving SAS memory error messages. Another benefit is that users of some UNIX workstations and PCs at UNC can take advantage of the large memory and vector processing support on the Convex C3840 to undertake problems that would either take too many hours or not run at all on these smaller platforms. Also, the fast sequential I/O of the Convex system can speed up multistep SAS programs working with very large files."

While most Convex users have traditionally been involved in science and engineering, most of the SAS users are in the social sciences — geography, economics, sociology, business, education, etc. But some graduate students and faculty in these areas were candidates for running their SAS applications on the Convex system because their jobs were either highly memory- or compute-intensive. At this point, says Muller, there are about 20 "serious" users running their SAS applications on the Convex system. "These represent a whole new category of Convex users," says Muller. ■

# Intelligent manufacturing from Ford Alpha

Early in the 20th century, Ford Motor Company established itself as a leader in automotive manufacturing by letting human workers take advantage of the power and capability of machinery.

Now at the dawn of the 21st century, Ford is establishing its leadership again — this time by letting humans take advantage of the power and capability of high performance computing with a concept known as Intelligent Manufacturing.

Leading the charge for Ford is its Alpha Simultaneous Engineering Division, a 450-person organization in Dearborn, Michigan. Alpha was formed in the mid-1980s to help Ford respond to profound changes in the auto industry.


“Alpha is Ford’s global technology resource,” says David F. Hagen, general manager of Alpha. “We provide ‘first proof’ that an application works and provides value to the customer.”

Convex systems have been at the heart of Alpha’s intelligent manufacturing since the first C120 was installed in 1989. Currently, an eight-processor C3880 is in use, and Alpha is looking at moving to the newest Convex systems by year-end.

Dr. Howard Crabb is the manager of CAE Functions at Alpha. He explains that Alpha is the proving ground for the advanced technologies that will be implemented in other Ford divisions. “We work with other divisions to get their buy-in to change design or manufacturing

processes. We evaluate the effectiveness of advanced CAE tools and also determine the most effective methods for applying the tools to real-world engineering tasks,” says Crabb. “Alpha pilots a new process, then once it’s proven, helps implement the process into production in the divisions.”

The division uses cross-functional teams to focus on several key priorities that offer wide-ranging benefits to Ford. These include Computer-aided Engineering (CAE), Electrical/Electronic Systems, Manufacturing Flexibility, Lightweight Components, and Lightweight Body Structures.



*Computational Fluid Dynamics (CFD) analysis of internal airflow characteristics.*

Alpha's work spans the world of Ford: North American Automotive Operations (NAAO), International Automotive Operations (IAO), Automotive Components Group (ACG), and Jaguar.


The Convex C3880 was selected as the main platform because it was at the time the only system available with eight processors. It is used at the concept level of design to predict crash analysis, air-flow simulations and vehicle system and sub-system performance.

A key application on the Convex system is visual engineering. Designers and engineers can create realistic 3D models and animation of any application and technology. These can include components, subsystems and structures, test procedures, full vehicles — even entire factories — all in the computer.

Most of the modeling is done using the MSC/Aries, MCAE software which makes it possible for the design engineer to develop the baseline model, analyze its performance and make necessary refinements. The Convex system networks with all the workstations.

An important current goal at Ford is to bring new car programs into production 36 months or less from the time engineering begins. As a demonstration project to help accomplish this goal, Alpha engineers recently designed a new lower control arm for a current model production vehicle in 95 person-days, compared to the 229 person-days that would have been required using conventional engineering methods.

"Using MSC/Aries Solids made the concept development process rather simple," says Crabb. "One reason was that the design engineer could formalize



*Thermal distribution analysis of new brake rotor design.*


his ideas with 3D shaded-solid images. Another key factor was that complicated shapes could be quickly generated and modified using the MSC/Aries software. The resulting 3D images leave no doubt about what has been created. These images are clear enough to present to other engineers and less technically oriented managers and marketing personnel to get feedback on the concept."

The MSC/Aries software produces an input file for the MSC/NASTRAN finite element package. Twenty iterations of analyzing and changing the model were performed in just 80 days, compared to a conventional approach that would have required 196 days.

There have been other dramatic reductions in time and cost from visual engineering. "As just one example," notes Crabb, "we worked with the SN95 (1994 Mustang) team to reduce the number of instrument panel prototypes by 50 percent and to reduce the prototype build cycle by 30 weeks."

The MCAE technology is now being transferred to operating divisions so that front-line engineers can greatly reduce component design time requirements and reduce program lead-time to world-class levels and better, according to Crabb.

Just as Ford revolutionized the automotive industry in 1913 with the moving assembly line, it is revolutionizing the industry in 1994 with intelligent manufacturing. By giving workers the tools they need to make their work more productive and rewarding, Ford is continuing its industry leadership in quality products and processes. ■



*This computer graphic shows a thermal footprint for a four-jet, side-inlet generic Short Takeoff, Vertical Landing (STOVL) aircraft configuration. The blue areas represent the lowest temperatures, and the red indicate the highest. The four red dots indicate locations where the lift jets impinge on the ground. An outline of a 27903 type aircraft is shown for reference.*

## Excellence at NASA Lewis

The NASA Lewis Research Center in Cleveland, Ohio is a National Aeronautics and Space Administration (NASA) Center of Excellence for air breathing propulsion. In various areas of its research, Convex systems are integral to the Center's excellence.

There are two major Convex installations at NASA Lewis. One, a Convex C3220 DataSeries, is part of an overall Mass Storage Service. The other, a C220 system, is used as a compute server for a wide range of research.

Lewis is conducting research for NASA's High-Speed Research Program in the areas of combustor design and enabling propulsion materials. It is advancing technologies to support advanced short take-off and vertical landing aircraft, and it is managing the Advanced Communications Technology Satellite. Further, Lewis is playing a key role in NASA's program to enable a more effective access to low Earth orbit and geosynchronous orbit.

The more recent of the two Convex systems is a C3220 DataSeries system. Released for production in September 1993, the C3220 is the heart of the central mass storage service which also features 100 gigabytes of magnetic disk storage, plus one terabyte of StorageTek robotic storage and 650 Metrum Robotic Storage. The Convex system runs UniTree fileserver software.

That's a lot of data storage and file service capability. But it's crucial, according to NASA Lewis' mass storage project manager Steve Prahst, who notes that the Center's data storage requirements are growing at a rate of approximately one terabyte per year.

Aside from having huge data storage requirements, the researchers are extremely demanding for reliability. "It's simple," says Prahst. "Reliability is indispensable. We cannot lose our users' data. If we do it once, they don't come back."

Prahst says that the Center had an in-house developed mass storage system for many years, but decided to adopt a commercially available system that would be portable between computing platforms. "In our selection process, we pretty much homed in on UniTree," says Prahst, "and although we considered running it on other platforms, we were swayed that UniTree was proven and running in many Convex sites."

The scientists at NASA Lewis are using the Convex DataSeries system to support their simulations and testing of jet engine components.

The other Convex installation is a C220 system. About four years before the C3220 DataSeries was installed, NASA Lewis installed the C220 system for a wide variety of research. Following the Center's mission of researching jet propulsion, most of the Convex-based work has been computational fluid dynamics for fluid mechanics and analysis in turbo engines.

NASA Lewis, which is adjacent to the Cleveland airport, sometimes makes the Convex C220 system available for projects other than jet propulsion. For example, in September 1993, when NASA launched the Advanced Communications Technology Satellite (ACTS), the Convex system was used for real-time mission planning. Engineers calculated the correct target vectors for each potential opportunity to deploy the satellite.

At about the same time, a doctoral student in medicine at the Case Western Reserve University, Cleveland, Ohio, used the Convex system to simulate a human heart in order to analyze normal function and to locate any

cardiac defects. In this case, NASA Lewis was supporting the government grant awarded to the student for the research.

For its core projects of jet propulsion research as well as critical projects beyond this scope, NASA Lewis depends on the high performance and reliability of its Convex systems. And as the Center's simulated and experimental data continue to explode, NASA Lewis will also continue to rely on Convex to meet its mass storage needs. ■



NASA's Lewis Research Center is using computational fluid dynamics to obtain a full three-dimensional solution for the unsteady flow field of counterrotation propellers. This figure of a General Electric Unducted Fan shows pressure contours in a plane just downstream of both blade rows. The blue areas represent lowest pressure contours and the pink represent the highest.

# SPP binary compatibility: A major advantage

"The Exemplar series architecture matches perfectly with the ANSYS Revision 5.0 parallel design, including shared memory, multiple threads and large private cache," says Dr. John Swanson, president of Swanson Analysis Systems, Inc. "In addition, the magic words 'binary compatible' bring joy to the hearts of software developers."

Swanson Analysis Systems, Inc. is one of several leading independent software vendors that have endorsed Convex's environment and porting strategy. Others include: Hibbitt, Karlsson & Sorensen, Inc. (ABAQUS); Gaussian, Inc. (Gaussian 92); Fluid Dynamics International, Inc. (FIDAP); Computational Dynamics Limited (STAR-CD); Livermore Software Technology Corporation (LS-DYNA3D); Visual Numerics, Inc. (IMSL); the MacNeal-Schwendler Corporation (MSC/NASTRAN); Numerical Algorithms Group, Inc. (NAG); Fluent, Inc. (FLUENT); Engineering Systems International, S.A. (PAM-CRASH); and Mecalog, S.A. (RADIOSS-CRASH).

There are hundreds of thousands of Hewlett-Packard PA-RISC-based workstations and servers throughout the world, and thousands of software applications that run on the HP-UX operating system. With the announcement of Convex's Exemplar SPP systems, users of these systems and these applications can look forward to the performance increases of parallelized processing while maintaining a familiar user environment.

"These new PA-RISC-based products from Convex provide a high level of performance and scalability that our customers in the technical marketplace need for their highly parallel applications," says Gary B. Eichhorn, vice president and general manager of HP's Workstation Systems Group. Hewlett-Packard and Convex will jointly sell these new systems as part of the on-going relationship between the two companies.

Organizations trying to solve exceedingly large scientific or engineering problems can also look to the Exemplar SPP system for massively parallel processing capabilities that can be put to practical use.

Dr. Danilo Zevertanik of the Jozef Stefan Institute, one of the first customers to receive an Exemplar SPP system, believes that the system's software capabilities that are ready to use are a major advantage. "We wanted to have a productive parallel computer," says Zevertanik, "one

where our researchers wouldn't spend most of their time unraveling the complexities of parallel programming, instead of using computer power in support of their specific areas of research."

According to Rick Roloff, manufacturing market development manager for Convex, work is well under way to tune and parallelize key parts of several important codes for full SPP processing. Additionally, work is in process to validate numerous applications with HP-UX binary code on a single SPP processor.

The Convex Exemplar SPP system is unique in its ability to provide both the platform for parallel processing development required by Jozef Stefan Institute and other customers, as well as the binary compatibility with the HP-UX operating system that enables access to the broadest spectrum of applications in high-performance computing.

"NCSA believes the Convex Exemplar series will lead the way in applications-rich scalable computing," says Larry Smarr, director of the National Center for Supercomputing Applications (NCSA). "It represents the first massively parallel computer to provide true binary compatibility with widely used workstations." ■

# MAFIA — Solving electromagnetic problems on a Convex supercomputer

by Martin Witting, SICAN GmbH,  
and Ingolf Sessler, Convex  
Hamburg

The MAFIA (MAXwell Finite Integration Algorithm) software system for solving electromagnetic problems was developed over the last decade through a collaboration of DESY (Deutsches Electron-Synchrotron, Germany); KfA (Kernforschungsanstalt, Germany); LANL (Los Alamos National Laboratories, USA); and later on at the Technical University of Darmstadt in Germany under the guidance of Prof. T. Weiland. It is based on the Finite Integration Algorithm, by which Maxwell's equations in their integral form are directly "translated" on three-dimensional (or for problems with rotational symmetry, two-dimensional) discrete mesh in space, and (for time domain purposes) in time. By some conversions, almost similar to analytical techniques known from the theory of electromagnetic fields, this set of "Maxwell grid equations" is reduced for special cases and then solved numerically.

Therefore, MAFIA is capable of doing computations for a wide range of electromagnetic problems. For this purpose, MAFIA is divided into a preprocessor module for the generation of the discrete mesh and the input of the geometry, a set of solver modules, and a postprocessor module for the display of the results and further processing of the gained

field data. These solver modules in particular are:

- S** A scalar potential is introduced into the equations and a "Poisson grid equation" is derived, which is solved by an iterative SOR or a multigrid scheme.
- R, E** The eigenvalue problem for lossfree media without driving currents is solved by a subspace iteration.
- W3** The fields for time harmonic driving currents are calculated in the frequency domain. This is done by an iterative SOR or a conjugate gradient scheme or a mix of the two.
- T3, T2** Transient electromagnetic fields excited by incident waves or elementary dipole antennas are calculated by the so-called Yee-algorithm, an explicit integration of the time domain grid equations.
- TS3, TS2** The same as in T3 and T2, but also with solving the self-consistent equations of motion for moving charges taking Lorentz forces into account.

The MAFIA software is distributed as a binary code for the most common workstation computers.

At SICAN (Gesellschaft fuer Silizium-Anwendungen und CAD/CAT Niedersachsen mbH of Hannover, Germany) there was strong demand for more computational power, so the MAFIA solver modules were installed on a Convex C Series system, a C3880 with eight processors and four Gbytes of memory. To take full advantage of the vector and parallel capabilities of this machine, the code has been adapted and optimized. SICAN designs digital integrated circuits, mixed signal ICs, Multichip Modules, and ASICs.

The optimization procedure took great advantage of the Convex profiler (PROF) and performance analyzer (CXpa). The first one was used to identify the most time-consuming subroutines while the latter one allowed a loop-level analysis. Through this we could verify the effectiveness of vectorizations often based on manually added compiler directives. The next and more sophisticated step used the hints delivered by the loop-level analysis to make more or less subtle code modifications. These changes included addition statements to simplify array index calculations as well as the unification of formerly separated loop nests into a single one to reduce the loop instrumentation overhead. Finally

# The Paris Observatory searches the heavens

High in the Andes Mountains, two 3.6 meter telescopes at the European Southern Observatory (ESO) study galaxies, clusters and faint stars in the deep skies. A sister telescope on Mauna Kea in Hawaii also studies planet surfaces and comets and asteroids in the solar system.

Radio-telescopes of the French-German Institute of Radio-Astronomy in Nançay in the French Alps observe the cold objects of the cosmos such as forming stars or pulsars.

Stellar telescopes in Haute Provence and Pic du Midi monitor galaxies and planets, and solar telescopes in Pic du Midi and in Meudon study the flares and other atmospheric events of the sun.

And deep in outer space, telescopes of the Ulysses spacecraft — the first spacecraft orbiting out of the ecliptic plane of the sun and planets — survey the solar wind and plasma of the solar system.

The massive astrophysical data from all these telescopes are collected and processed by a Convex C3440 system at the Observatory of Paris in Meudon, France.

The Observatory, which operates under the authority of the Ministry of National Education, has core activities which include fundamental research in astronomy and astrophysics. It employs 200

researchers in all fields of astronomy, including solar physics, planetology, stellar, galactic, extragalactic, and cosmological physics.

The Observatory acquired its Convex system in April, 1992. "The vectorial calculating power and prodigious memory of the Convex system has made it possible for us to solve problems that we could not before," says astronomer Jean-Marie Malherbe, head of the Paris Observatory Computer Department. "The system's 32-bit architecture is perfectly adapted for image and data processing. And for modeling observations, theoretical studies and numerical simulations, the 64-bit architecture gives us the precision we require."

More than 100 users in nine departments are networked to the Convex system. Both FDDI and Ethernet networks connect the users' high-performance workstations with the Convex C3440.

With the assistance of the processing performance of the Convex system, the Paris Observatory is preparing for additional studies and projects. In 1998, the Very Large Telescope in Chile (the European Southern Observatory) will be one of the most powerful in the world. The Observatory is participating in the construction of this telescope which will have four mirrors, each with an 8-meter diameter.

Also, the Observatory is constructing a solar telescope (THEMIS) on the Canary Islands to record the solar vector magnetic field with high spatial resolution. This new instrument will be one of the most competitive, and will be operational in 1996.

As the Observatory of Paris searches the heavens, the Convex system helps the researchers find their answers in the stars. ■

## MAFIA...

*continued from page 11*

the powerful pregon analysis tool of the Convex performance analyzer was an excellent guideline to get the best available thread balance for the parallelized T3 version.

To get reliable results on performance gain, we used an example geometry including a mesh of approximately one million points (1217\*127\*60) and 6000 timesteps. ■

Version	Elapsed Time	CPU Time (secs) (User + System)	Acceleration (Compared with Scalar)
Scalar	43:22:11	140800	-----
Vector	03:11:41	11154	13
Parallel	00:45:41	13346	57

Time values are up; performance gains are rounded off.

# Beginning a new history of computing at University of Tübingen

The first computing machine installed at the University of Tübingen was not a Convex. The year of the first installation was 1623, and the name of the actual manufacturer of the machine was Wilhelm Schickard, a professor at the university and a friend of Johannes Kepler, the great astronomer.

However, in the early 1980s, the University of Tübingen, which is located about 50 km south of Stuttgart, the capital of the State of Baden-Württemberg, was one of the first Convex customers in Germany. Today, the university boasts a sophisticated computing center with two Convex systems at its heart.

One system, a six-processor C3860, is used as a compute server. According to system manager Dr. Werner Dilling, this system is primarily used for highly computational applications in the natural sciences and in the arts and humanities. "The users enjoy the performance of the C3860," says Dilling. "And they keep it heavily loaded 24 hours a day."

The second system, a C3220, is used as a fileserver. Equipped with 265 megabytes of main memory and 50 gigabytes of disk, it interfaces with a Metrum RS-600 robot providing 8.7 terabytes of tape storage. This is equivalent to 8 million 500-page books. The huge archive is used for workstation back-up and long-term data storage. It is the largest robotic archive in the State of Baden-Württemberg.

"The institutes use the C3220 for back-up and software archiving and distribution," says Dilling. "Some of their software programs are quite specialized and they don't need to use them often. But with the fileserver, they can access them quickly and easily, around the clock."

There are about 27,000 students at the University of Tübingen. The major schools are arts and humanities, economics and social science, natural science, and medicine. There are approximately 6,000 users with about 1,500 workstations and personal computers throughout the University of Tübingen campus. Sprawled over the city, the workstations and PCs access the computer center via a 100 Mb/s FDDI backbone.

About 90 percent of the use of the fileserver is automatic back-up. "Our users can put a large amount of their old files on the fileserver because they can retrieve them so easily," says Dilling.

"This also unloads the compute server. We have a more transparent access to the backed-up data. Before the fileserver system, we could only restore files manually — now it's all automated."

The C3220 and C3860 represent the third generation of Convex systems that University of Tübingen has used. And the university's new client-server concept represents a major change in the service philosophy of the computer center. From a single-vendor, proprietary computing environment, the computer center has moved into a UNIX-based, open systems environment. "Grandpa's computer center is dead," quipped Dr. Dietmar Kaletta, the director of the computer center.

The University of Tübingen is hoping for additional funding for enhancing the computer system. "Convex has some products we're very interested in," notes Dilling. "For example, we increased the disk capacity on our file server. We also have a great interest in the scalable parallel systems."

As a member of the Convex Scalable Computing Working Group (SCWG), the University of Tübingen has been active in advancing the state of the art in scalable algorithms and applications design. ■

# TeamQuest Baseline

"We thought," says Norb Thomes of TeamQuest, "that users of our Baseline product would pretty much just be system administrators. It turns out, though, that 'help desks' use it to answer callers' questions about what's going on with their jobs, and other administrative groups use Baseline when they're putting up new applications to mark their progress."

Thomes adds, "There's just a lot of use for a package that provides a very complete picture of system performance and can analyze innumerable, disjunct bits and pieces of information."

TeamQuest Corporation of Clear Lake, Iowa, recently announced a graphical performance analysis and capacity planning tool for Convex C Series machines running ConvexOS 10.x and 11.1.

Joe Rich, also of TeamQuest, adds, "Of course, the higher performance of the Convex C Series increases the performance of Baseline. But the real advantage of Baseline to Convex users is that it can maximize system performance and productivity, and help organizations get the most from their high-value computer investment."

Known as TeamQuest Baseline release 5.2, the product includes the capability to:

- monitor system performance in real time
- report past performance from a performance database
- report performance at the system level, workload level, or process level
- perform automated correlation analyses to find possible cause/effect relationships
- perform long-term trend analysis to plan for future system needs.

TeamQuest Baseline supports two types of performance reporting: real-time monitoring and historic performance reporting. Real-time monitors report current system performance data and are updated on a timed basis to reflect the most current data available. This provides the earliest warning of problems, allowing users to investigate bottlenecks as they occur. Historic reporting enables users to look back at data, find and examine performance and system usage trends, and predict future system needs.

All aspects of data collection and retention are user-defined, including the characterization of workloads. Performance data is consolidated around business functions, giving the administrator clear and concise reports showing exactly how each group of users utilizes the system.

TeamQuest Baseline also features automated correlation analysis. Performance peaks can be automatically checked against the other parameters in the database to greatly simplify the investigation of performance bottlenecks.

TeamQuest Baseline uses a client/server architecture. The client software runs on the desktop and handles all user interface tasks. The server software runs on the C Series systems and handles all data collection, data maintenance and correlation analysis. To generate a report, the client software formats a request and sends it to the server software for processing. Data packets are returned to the client and used to create line graphs, area charts, bar charts, scatter plots, pie charts and textual reports.

Multiple report windows from a variety of systems can be opened within the desktop, enabling users to examine facets of performance on a single system or compare the performance of more than one system with side-by-side reports. In addition, predefined report definitions are installed with the product. These definitions can be modified at any time and new reports can be added. Users can copy and paste report windows into other desktop applications. Reports can also be sent to hardcopy devices such as laser printers and plotters, and performance data can be exported to other applications.

Easier access to system performance and capacity management information will enhance decision making by both technical staff and management.

For more information on using TeamQuest Baseline, contact TeamQuest at (800)551-TEAM or (515) 357-2700 FAX (515) 357-2778 or email [info@teamquest.com](mailto:info@teamquest.com). ■

**Hubble images...**


*continued from page 5*

While NASA discussed a possible repair mission for 1993, other scientists tried to solve the telescope's problem by correcting the data already collected. Two other scientists at UNC-CH joined with Carney to attack the problem. They were Dr. James Coggins, associate professor of computer science, and Laura Kellar Fullton, a doctoral candidate in astrophysics.

Coggins became involved because he had done work with biomedical image restoration. While visiting the Goddard Space Flight Center one day, someone asked him whether he knew anything about optics and deblurring. "It was in fact the Hubble problem," says Coggins. So he reviewed his past work in restoration and "learned some more math," he notes.

"One day while I was at the Goddard Space Flight Center, I had the equations on the board and a diagram showing the basic iterative method for deblurring images," says Coggins. "I was looking at it, and I said, 'Wait a minute — that's not what you want to do.' I went over to the board and added something to the diagram and looked at it and said, 'Hey, that's kinda cool.'"

It turned out to be quite cool, because the simple twist — simple at least to an expert in the field — yielded a much more robust, faster and more accurate deblurring process that preserves the relative brightness information from the images.



*Astronaut Story Musgrave prepares to ride to the top of the Hubble telescope to install protective covers on the telescope magnetometers. Astronaut Jeff Joffman assisted Musgrave with final service tasks on the telescope, wrapping up 5 days of space walks. (Photo courtesy of NASA Space Telescope Science Institute.)*

Known as the Basic Iterative Reconvolution Algorithm, Coggins' algorithm deblurs images. "We take a well-informed guess as to what the true image ought to be," explains Coggins. "Then, essentially, we blur the guess. If the guess was right, the blurred guess is identical to the observed image. Otherwise, we make the guess better by adding to the guess the difference between the observed image and the blurred image. Then, iteratively, the deblurred image obtained at one level is fed back up as the correction to the guess at the previous level."

Coggins discussed his algorithm with Fullton, a member of Carney's research team. Both Fullton and Carney were impressed that the algorithm might indeed deblur images while retaining the relative brightness information essential to their own Hubble research.

Using the Convex C3840 system at UNC's Office of Information Technology, Fullton wrote computer code simulating Hubble images and, implementing Coggins' deblurring algorithm, discovered they were able

*Continued on page 16*

**Hubble images...***continued from page 15*

to measure the brightness of stars in the simulation to within one or two percent over a very large range of brightness. "So we're measuring 98 to 99 percent accuracy in the brightness of the stars, which is about as good as you're ever going to get," Fullton explains.

"When I started developing the code," says Fullton, "I was using Sun workstations. The problem was that I was working with 800x800 real images, and performing Fourier transforms. A workstation might work for one Fourier transform image, but for

40 or 50 images, I needed much more performance because it was incredibly time-consuming."

Fullton reports that she was pleased with the Convex system's performance. "The vector operations made everything go really well," she says. "I used C++, and the object oriented language let me treat the images as objects that I could store in a library and manipulate easily."

Coggins' algorithm has compared quite well in side-by-side tests against several other image restoration algorithms. Fullton perhaps puts it best when she says, "We are very confident in the algorithm as it stands now."

Fullton has had requests from all over the world for copies of her software program that contains the computer code for running the algorithm.

The UNC community takes pride in a problem-solving effort that linked three of its scientists in an effort to recover with ingenuity what was threatened by mechanical error. All three will continue their involvement in various phases of the Hubble project, but they share the satisfaction of their collaboration in the image restoration project. ■

---

Convex, the Convex logo ("C"), C Series, Exemplar, CXpa, C3800, C3400, SPP-UX, C4/XA, ASAP, SPP1000/XA and SPP1000/CD are trademarks of Convex Computer Corporation. Hewlett-Packard, HP-UX, and PA-RISC are trademarks of Hewlett-Packard Company. NASTRAN is a registered trademark of the National Aeronautics and Space Administration. MSC/NASTRAN is an enhanced, proprietary version developed and maintained by MacNeal-Schwendler Corporation. UNIX is a registered trademark of UNIX System Laboratories, Inc. SAS is a registered trademark of the SAS Institute, Inc. METRUM is a trademark of METRUM Information Storage. UniTree is a registered trademark of OpenVision Technologies. StorageTek is a trademark of Storage Technology Corporation. ANSYS is a registered trademark of Swanson Analysis Systems, Inc.

*Context* Editor  
Convex Computer Corporation  
P.O. Box 833851  
Richardson, TX 75083-3851

**Context Staff**

Editor  
Lillian R. Geary

Writer  
Emily Martin

*Context* is published quarterly by Convex Computer Corporation to provide news about Convex applications, technology, service and support, and other trends in the supercomputer marketplace. Subscriptions are free of charge to all Convex customers.

Although the material contained herein has been carefully reviewed, Convex does not warrant it to be free of errors or omissions.

Convex reserves the right to make changes in equipment design or components as progress in engineering or manufacturing may warrant.

Send address changes and subscription inquiries to *Context* Editor, P.O. Box 833851, Richardson, TX 75083-3851.

© 1994 Convex Computer Corporation