Powerful Supercomputing at Economical Prices

The University of Stuttgart takes advantage of the latest advances in computing power with the Cray XT5m supercomputer, offering researchers and scientists high-performance computing without the high-end cost

In academic institutions around the world, researchers strive to push the boundaries of scientific knowledge. While advancements in computing over the last fifty years have solved many computational scientific challenges, the ability of computing to keep up with the rapid changes in the scientific and engineering fields is constantly challenged. Scientists working in fields as diverse as physics, fluid dynamics, chemistry, materials sciences, biotechnology and mechanical engineering are constantly finding new challenges to explore, tackle and solve with high performance computing.

Staying at the forefront of computational science that is relevant to our daily lives takes continuously increasing computing power. It also takes collaboration with commercial research organizations. The University of Stuttgart in Germany is taking advantage of the latest advances in computing power and applying savvy insight from commercial research organizations to provide researchers and scientists with high-performance computing at the cost of traditional high-end computing.

As the site of the first installation of the Cray XT5m supercomputer, the University of Stuttgart's High Performance Computing Center (abbreviated HLRS) now benefits from a high-performance system that is powerful and affordable. The Cray XT5m system, which starts around $500,000, takes advantage of the hardware and software advancements of the Cray XT5 supercomputer, the basis of the petascale system currently in use at the U.S. Department of Energy's Oak Ridge National Laboratory.

> "The less time we spend figuring out how the system works, or programming specialized processors, or optimizing simulations, the more time we have to run simulations and the more productive we are."
>
> — Prof. Michael Resch
> Director of HLRS

The Cray XT5m supercomputer's deployment and design is bringing new flexibility and insight to both academic and corporate research efforts, allowing the University of Stuttgart to be on the cutting edge of not only midrange supercomputing, but also public/private sector partnerships. Supported by its state government and by local industrial corporations, HLRS is a crucial part of a triumvirate that is advancing scientific knowledge while making it available to a wider spectrum of participants.

How the HLRS Works

HLRS is one of three supercomputing centers in Germany. It participates in two key partnerships. One, known as High Performance Computing Center for Academia and Industry (abbreviated HWW), is a partnership of academia, government and business. The University owns 12.5 percent of the partnership, as does the state of Baden-Württemberg. Automobile manufacturer Porsche AG owns 10 percent, and German telecom company T Systems owns 40 percent. Other academic institutions, which have access to the system, own the remaining 25 percent. These groups have access to supercomputing cycles for scientific visualization, computational fluid dynamics, physics and other fields.

Academic researchers work through the University of Stuttgart to gain cycle time on its supercomputers (besides the Cray XT5m, it also has systems from IBM, NEC and Sun); the researchers account for about 70 percent of its use. Companies who want to collaborate by obtaining cycle time on the supercomputers work through HWW. While the University is responsible for operation, hosting, simulation and optimization activities, HWW handles administrative activities such as its security and accounting.

The other key partnership of HLRS is the Automotive Simulation Center/Stuttgart (ASCS), which was formed to focus on common technical or scientific problems in the automobile industry. The members of ASCS include Daimler, Porsche and Opel in the automotive industry; Karmann in the supply industry; engineering and scientific software developers such as INTES, ESI Group, DYNAmore, Abaqus and Altair; and hardware vendors such as Cray, IBM and NEC.

The ASCS is collaborating on tackling two key challenges facing automobile manufacturers around the world: improvements to the internal combustion engine and multidisciplinary optimization. The former includes research into issues such as fluid flow, chemical reactions and material science. The latter relates to correlating the multitude of design and engineering parameters that go into making a car as efficient as possible.

"You can optimize [a car's] fuel efficiency by making it lightweight and using materials to avoid as much weight as possible," explains Prof. Michael Resch, director of HLRS, whose degrees span disciplines including applied mathematics, computer science, economics and mechanical engineering. "But if you reduce the weight of the car, the vehicle is less stable. You need a lightweight structure that will consume less fuel but is as stable as or more stable than any other car." Designing for

aerodynamics may involve changing the shape of the car, which may increase production costs.

"The goal is to find the optimal shape at the optimal minimum weight for the optimum cost," says Resch. "That's multidisciplinary optimization, and it is a challenge because all these calculations should be done simultaneously. Without a system like the XT5m, we have to do these calculations for these mathematical models individually, which is inefficient and time-consuming." With the help of the XT5m supercomputer, which accommodates calculations for multiple elements simultaneously, engineers can more precisely determine parameters. The result: more efficient, safe and fuel-efficient vehicles.

Cray XT5m Helps Identify Advances in Computational Science

The process by which the ASCS partners work to develop these models is highly collaborative. Engineers at the industrial partners provide goals and requirements for a project, which the HLRS and other research partners use to create mathematical modeling and algorithms. The researchers then conduct experiments to validate the models. Once validated, the ISVs implement the code into their software. The software is then passed back to HLRS, which validates the software in a pre-production environment. The cycle is completed when the code is delivered to the industrial partners for their individual use.

One example of this relationship comes from ESI Group, a global supplier of digital simulation software for prototyping and manufacturing processes focusing on material sciences and physics. Its software is designed to simulate product behavior during testing and subsequently to improve the efficiency of the manufacturing processes to attain that performance. Using virtual rather than physical prototypes can save manufacturers millions of dollars.

ESI ported its PAM-CRASH software to the Cray XT5m supercomputer, taking advantage of up to 1,024 CPU cores to analyze millions of potential elements relating to materials in a collision scenario. Using the Cray system and its development tools, the ISV has been able to bring computation time down from hours to 25 minutes.

Ultimately, Resch says, this iterative relationship helps identify potential advances in numerical and computational methods and get those advances incorporated into the ISV's code. Previously, it would have been too expensive for any one of the participants to work on their part of the cycle, but by working in concert everyone benefits.

For industry, the collaboration with ASCS translates into two key metrics: improved design and shorter time to market for those designs. "Thanks to the XT5m, our industrial partners can do more design simulations in the same amount of time," says Resch. "Because the system is less expensive than other supercomputing systems, each of the simulations costs less to run. The sooner they get crucial design information, the better. Simulation is part of the design phase, and the earlier in the design process you can conduct complex simulations, the easier it is to avoid errors at the manufacturing phase. The more errors you avoid, the more money you save. For the automotive industry here in Stuttgart, that becomes a competitive advantage."

Cray XT5m Architecture Helps Lower TCO

In offering its cutting-edge program, HLRS relies on several key technical capabilities of the Cray XT5m. Foremost among these is its use of the "best-of-class" standard x86 processors. Because it uses AMD Opteron processors, the XT5m can provide industry leading price-performance in a fully upgradeable system. Supercomputing systems that use specialized vector processors do not derive the advantage of research and development that then gets applied across millions of processors; as a result, those systems are more expensive by default. "The ten supercomputing centers alone around the world cannot support the R&D costs necessary to improve vector processors. In the end this technology will have to become part of a global standard platform," Resch says.

Resch also cites the SeaStar interconnect technology, which interfaces with the compute and service nodes in the architecture. Cray designed the SeaStar interconnect technology to provide leading performance and reliability in large-scale, distributed-memory supercomputing systems, outscaling standard interconnects such as Infiniband for the largest supercomputing systems. Cray now offers a scalable, upgradeable and cost-optimized version of the interconnect in the XT5m. HLRS has used standard processors in high-performance clusters in the past, but Resch notes that their performance has been limited as applications scale to large core counts.

Another important facet of the Cray supercomputer for HLRS is its standard programming environment, which takes advantage of the Linux operating system; standard compilers for C, C++, Fortran and UPC (as well as Cray's own); debuggers; performance optimization and workload management tools; and parallel programming platforms.

"This is an important advantage," notes Resch. "Cray has a long history of development software for high-performance computing systems, and it's no different for the XT5m. I have to prepare, program and run the system. The less time we spend figuring out how the system works, or programming specialized processors, or optimizing simulations, the more time we have to run simulations and the more productive we are. To get the greatest efficiency, we need the combination of standard processors and standard architecture on the one hand, and we need the speed and programmability on the other."

Standard processors bring other advantages to research efforts. "Sometimes researchers might want to do simple testing in the initial stages of development," explains Resch. "Because it's the same standard processor architecture, they can do that testing on a smaller workstation. They don't have to request time on the supercomputer." Once the researchers have confirmed the viability of the tests, they can run large scale production simulations on the XT5m. It is necessary to apply parallelization techniques for the simulations to run efficiently on the supercomputer, but the process is much faster thanks to the commonality of the processors.

Resch is also excited by the XT5m supercomputer's capability for processor upgrades. "Whenever AMD comes up with a new processor for the volume market, we can upgrade our processors," he says. In addition, the XT5m is designed so that upgrades to processors or memory can be done in the field at a low cost without a lot of downtime, while preserving the overall infrastructure.

Similarly, when the SeaStar interconnect technology is replaced by the next-generation technology, code-named Gemini, that capability can be upgraded as well. Resch believes that the XT5m will serve HLRS for several years thanks to these upgrade paths, a capability that will significantly decrease the University's total cost of ownership.

Cray XT5m Puts Industry and Academia on the Cutting Edge

Deploying the Cray XT5m at the University of Stuttgart has resulted in both academic and business benefits. Resch reports that the supercomputing system is so popular, every available hour and cycle has been reserved around the clock. While researchers may have to wait two weeks for access, they have greater access to more supercomputing cycles at a lower cost than ever before. This helps those researchers tackle more questions and get feedback sooner.

The benefit for industry is clear as well. Without having to invest themselves in a supercomputer, the industrial organizations in Stuttgart—which include global manufacturers such as Bosch, Karcher and Stihl—have access to supercomputing cycles. This helps them tackle the design and engineering of 21st-century products that take into account the greatest possible efficiencies in terms of materials, weight and energy. Being able to create more efficient products and get them to market faster means greater industrial efficiency and cost savings all around.

Resch notes that when industrial users need access to the supercomputing system, particularly for a just-in-time simulation, the HLRS will sometimes give industry a priority spot. The industrial users are then subject to a surcharge for that real-time usage, the money from which HLRS uses to upgrade the systems. "When we have the money to upgrade systems every two years rather than every four years, this means our researchers can be more competitive within academia," says Resch. HLRS can improve its systems without having to rely solely on the government or the University to fund improvements. "It amounts to a real boon for the University, both from a financial standpoint and for our reputation as a place for research," says Resch.

Thanks to the Cray XT5m, HLRS has been able to create a symbiotic relationship between industry and academia that benefits both by improving efficiencies while keeping costs down.

About AMD

Advanced Micro Devices (NYSE: AMD) is an innovative technology company dedicated to collaborating with customers and technology partners to ignite the next generation of computing and graphics solutions at work, home and play.

For more information, visit http://www.amd.com

AMD

The future is fusion