SUCCESS STORY | DRESDEN UNIVERSITY OF TECHNOLOGY

ADVANCING NEW DISCOVERIES WITH A GPU-POWERED SUPERCOMPUTER

NVIDIA® Tesla® K80 accelerators are opening new doors to scientific discovery for researchers at TU Dresden facing complex challenges and ever-growing data sets.

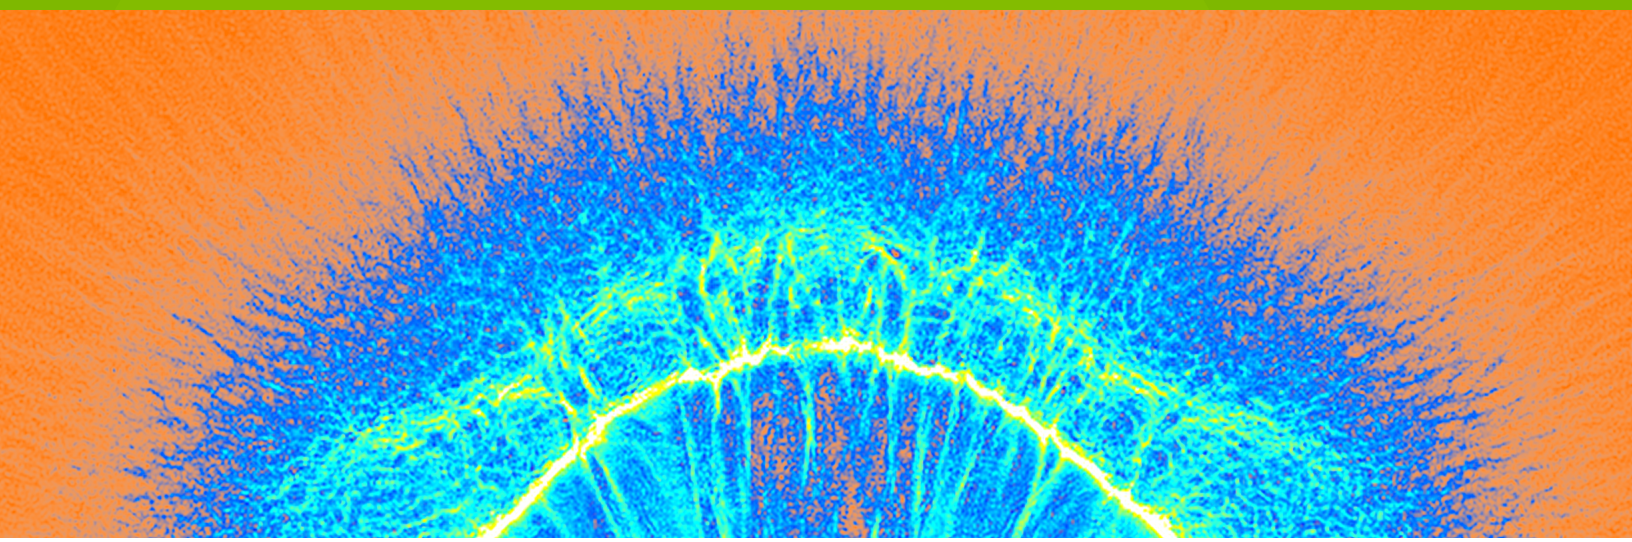

NVIDIA

Pioneering new medical treatments that target cancer more precisely, while helping preserve healthy tissue, is just one way that Tesla K80 GPUs are redefining what's possible.

AT A GLANCE

CUSTOMER PROFILE

Customer: Dresden University of Technology

Industry: Higher Education

Locations: Dresden, Germany

Size: About 45,000 students and staff

System: #364 Supercomputer in July 2015 Top500 List, #9 on the Green500 List (GPU partition of the installation)

SUMMARY

- A leading German university is conducting complex research across diverse fields.
- The existing CPU system did not meet the needs of numerous researchers, who depend on GPUs for timely simulations and visualisations.
- 128 NVIDIA Tesla K80 GPUs in 64 nodes were added to the data centre.
- Usage of the GPU partition of the new supercomputer installation hit 80% on its first day and continued to rise thereafter.

The Dresden University of Technology (TU Dresden) is one of the largest universities in Germany, with about 37,000 students supported by more than 500 professors, plus 7,200 staff members. TU Dresden traces its roots back to 1828, making it one of the oldest universities in Germany. Its commitment to furthering scientific research and engineering, as well as humanities and social sciences, earned it the title of "University of Excellence" in 2012.

CHALLENGE

The traditional CPU-only approach to computing presents a major challenge for researchers across multiple disciplines. They have big problems to solve, but not enough compute horsepower to solve them with CPUs. For example, Dr. Michael Bussmann, from Helmholtz-Zentrum Dresden-Rossendorf, is researching innovative techniques for treating cancer with high-power lasers that minimise the effect on healthy tissue. With complex simulation and in-situ visualisation requirements, his research depends on GPUs to deliver meaningful results within a usable time-frame—a task that would be almost impossible using CPUs alone.

From an IT perspective, data centre usage is a key metric for measuring the success of its investments. Four months after deployment, TU Dresden's CPU-only system gradually reached 50% usage—a sign that users were still searching for ways to transfer their challenges onto the system.

"Our data centre included over 800 traditional servers, but even these were not powerful enough to support some of the advanced research taking place here," said Dr. Guido Juckeland, IT-Architect and secretary of the GCoE at TU Dresden. "We soon decided to experiment with accelerated computing and turned to NVIDIA for help."

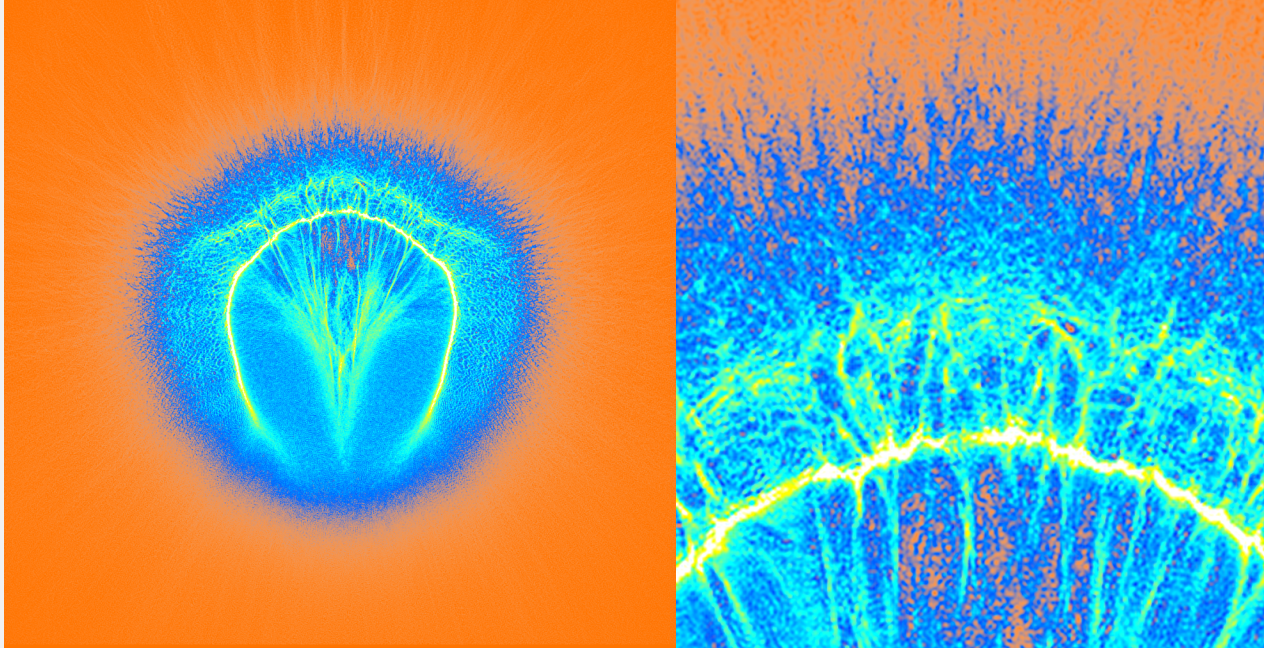

Right: Ions in a micron-sized spherical target after a high-power laser has hit. Source: TU Dresden.

REASONS FOR TESLA K80

1. There are 370+ GPU-accelerated applications.
2. A single Tesla K80-powered server replaces up to eight CPU-only servers.
3. Higher throughput leads to more discoveries.

SOLUTION

The first wave of GPUs at TU Dresden consisted of 88 NVIDIA Tesla K20 accelerators deployed in 44 servers. To Dr. Juckeland's surprise, usage hit more than 80% from the first day the system became available for users. Demand from new and existing GPU users, coupled with more than 20 GPU-accelerated applications, continued to drive usage even higher thereafter.

"Giving our users access to multiple high-end NVIDIA Tesla GPUs is allowing our researchers to make discoveries that would be next to impossible with CPU-based systems," Juckeland continued.

> "Even high-end workstations and laptops don't have the GPU power needed for demanding computation, which is where the accelerated servers come in."

The NVIDIA Tesla K20 GPUs quickly attracted new users across diverse fields, including deep learning, which further increased usage. Demand was so high that the IT department decided to convert part of a planned CPU purchase to 64 nodes equipped with dual Tesla K80 GPUs.

RESULT

By expanding their supercomputers with Tesla K80 accelerators, researchers are accelerating discoveries and even tackling new problems that would simply not be possible with CPU-based systems. For example:

Dr. Bussmann's research on radiation cancer therapy using laser-driven ion beams at Helmholtz-Zentrum Dresden-Rossendorf can simulate nearly 250 million particles on a single Tesla K80. Accessing all 256 GPUs allows highly accurate in-situ visualisation and simulation that a CPU system couldn't match. The performance boost provided by the Tesla platform is helping the research team find ways to target cancer cells better, while minimising impact on surrounding tissues.

The Max Planck Institute of Molecular Cell Biology and Genetics (MPI CBG) exploits massively parallel applications on a day-to-day basis. Dr. Gene Myers and his research team apply GPUs in everything from electromagnetic wave propagation simulations and real-time microscope control and visualisation, to image de-noising and multi-view reconstruction of 3D optical imagery.

The MPI CBG team around Pavel Tomancak uses Tesla K80 GPUs to speed the time-to-solution of applications running intensive Fast Fourier Transformations and multi-angle 3D image deconvolutions by 6x to 20x compared to traditional CPU methods.

GPUs are essential for research using deep learning. Professor Carsten Rother's group from the Computer Vision Lab has developed semantic understanding of the surroundings of an autonomously driving car that recognises pedestrians, other cars, streets, and signs. The training of convolution neural networks is a daunting task that would not be feasible in the very agile development environment that the group is participating in when only using CPUs.

Below: Given a sequence of images, the goal is to recover a rich, detailed representation of the 3D world, ranging from physical to semantic aspects. Source: Computer Vision Lab Dresden.

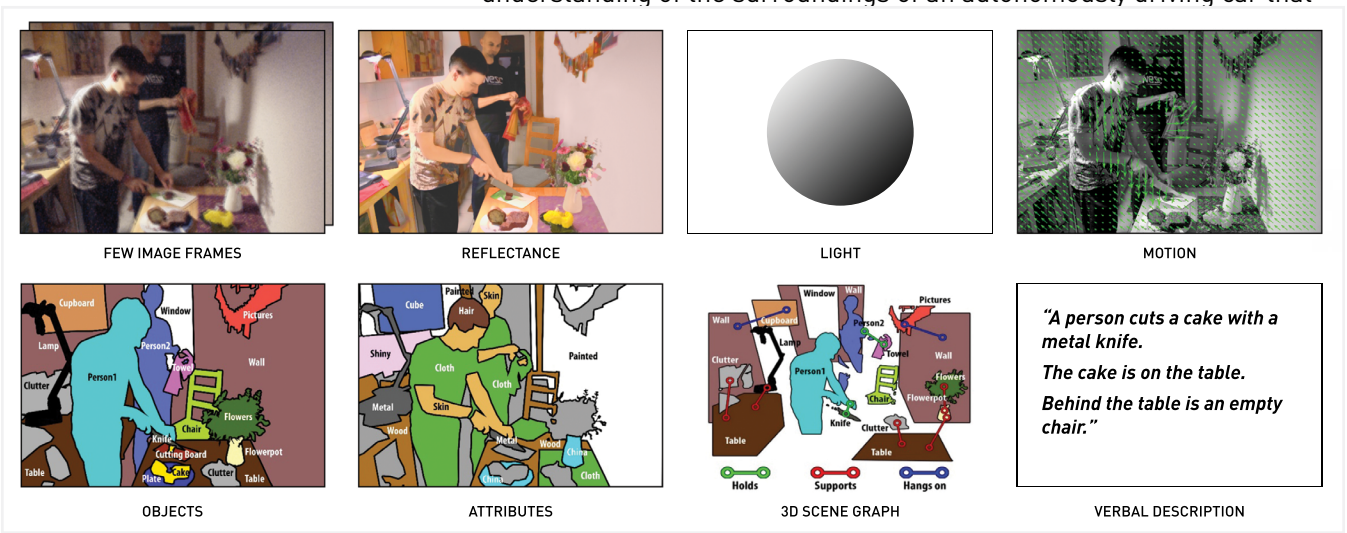

The 64 NVIDIA Tesla K80-powered servers provide roughly the computing power of 1,400 CPU-based servers. But the real benefit of GPU acceleration is dramatically faster turnaround times for research. Not all applications can scale across so many servers; however, they can all take full advantage of the available GPUs.

"The ever-growing use of GPUs at MPI CBG and in biological data analysis, simulation and image reconstruction allows scientists to study more organisms and follow their development from embryo to adult on a cellular level at much higher resolutions than previously possible," said Juckeland.

"Adding the NVIDIA Tesla K80 GPUs to our data centre will allow us to manage steadily rising data rates as automated processing becomes the norm and as new algorithms yielding higher resolutions demand ever more processing power."

To find out more about NVIDIA Tesla, visit **www.nvidia.com/tesla**

JOIN US ONLINE

blogs.nvidia.com

@GPUComputing

linkedin.com/company/nvidia

Google.com/+NVIDIA

NVIDIA