Advanced Monitoring Techniques for Data Centers Using Virtual Reality

By Jürgen P. Schulze

In project GreenLight, researchers at the California Institute for Telecommunications and Information Technology (Calit2) have built an instrument to measure the energy usage of computing systems under real-world conditions. This instrument consists of clusters of CPUs and GPUs, storage systems, and a 10-Gigabit Ethernet network infrastructure. The power consumption of each component can be measured separately, and includes temperature sensors in various components, to study the energy flow through the system. The goal of this project is to learn how computer hardware and software can be made more energy efficient. This research is intended to have an impact on the power consumption of data centers such as those supporting financial, business, communication, or entertainment applications. This article reports on findings and experiences relevant to the movie industry.

INTRODUCTION

With fast-growing costs for energy and the social and political urge to decrease carbon emissions, the reduction of the electric energy consumption has become a priority in data centers. Large data centers support financial, business, communication, entertainment, and mobile applications worldwide. As of 2006, the electricity use related to servers and data centers nationwide was estimated at approximately 1.5% of total U.S. electricity consumption, equivalent to 5.8 million U.S. households.[1] As of 2009, the computing industry contributes 2% to the global CO_2 emissions, similar to the worldwide aviation industry, with a predicted 6% increase per year.[2] Of this large amount of energy consumed by data centers, approximately half is used for IT equipment; the rest is used for cooling and other non-computation-related components.

In the past few years, the computer industry developed a variety of technologies to reduce power consumption and thus IT's carbon footprint. New low-power CPUs have been designed to reduce the energy needed for servers. Virtualization is now intensively used for server consolidation to reduce the number of servers required. Also, the architecture of the buildings hosting data centers has been redesigned, leading to cutting-edge solutions, such as container-based data centers.[3] This new type of data center is engineered to optimize space and cooling, achieving great power usage effectiveness (PUE) values. Microsoft and Google recently reported a PUE of 1.2 for their container-based data centers, compared to an average PUE of 1.9 to 2.0 for other data centers around the world.[4]

PROJECT GREENLIGHT

In project GreenLight at Calit2,[5] researchers have built an instrument to measure the energy usage of computing systems under real-world conditions, with the ultimate goal of getting computer designers and users in the scientific community to re-think the way they do their jobs from an energy efficiency perspective. The researchers at Calit2 have recently completed the initial GreenLight instrument, which consists of a Sun Modular Datacenter (Sun MD) container, filled with a variety of computer and graphics servers, storage systems, and network switches. The container houses eight racks capable of supporting up to 40 servers each. All devices are equipped with per outlet power sensors using Avocent PM3000 network monitorable power distribution units (PDUs) and per-server Intelligent Platform Management Interfaces (IPMI) to provide accurate power measurements for each device in realtime. These measurements include details of consumption for computation and communication, as well as for handling air flow, heat exchanger dissipation, chilled water flow, and related thermal management. Energy data, both realtime and historical, are being collected on a server and available on the Internet to assist energy efficiency research in areas such as computer architecture, software architecture, visualization, media content storage and distribution (a typical usage scenario in the movie industry), energy monitoring and modeling, and domain science applications in biology, geology, and bioengineering.

Conserving energy, reducing "carbon footprints," and deploying sustainable or renewable energy sources are increasing priorities in public and private institutions, across all sectors of the economy, both in the U.S. and abroad. Project GreenLight provides the tools to enable the research and understanding of energy conservation in information and communications technology (ICT), one of the world's largest growth industries and leading energy consumer. Given the pervasiveness of ICT in every industry today and its central role in the creation of new industries tomorrow, energy efficiency improvements in this sector will have worldwide and far-reaching impact. Project GreenLight is now providing tools and data to help researchers and industry make such improvements and is providing users with information on the energy use of various types of computing clusters.

THE VIRTUAL DATA CENTER

To help power efficiency researchers understand power consumption and temperature distribution in the Sun MD container, a 3D model of the Sun MD was developed to visualize these measurements spatially. The 3D model, depicted in **Fig. 1** is based on a CAD model of the data center from Sun, to which the computer systems Calit2 was added within the GreenLight project, such as servers, network switches and storage systems. The 3D model is a complete replica of the container, which also allows the user to open the doors, enter the container, and pull out the computer racks, all by directly interacting with the visual components.

Figure 1. 3D model of SUN Mobile Data Center.

Our virtual reality (VR) application connects via web services to a server that collects and stores the realtime information from the sensors in the container and retrieves in realtime the status of the container. **Figures 2 and 3** show how the instrumented components in the container are displayed in different colors, depending on the type of measurement selected from a 3D menu. For instance, when looking at temperature data, the nodes may be depicted in red to indicate a high temperature, and green when they are running cool. This gives the user a quick overview of the state of the devices in the container. Once the data has been loaded from the data bank it will be updated in synchronization with the data bank.

The interactive functionality was created using COVISE (Collaborative Visualization and Simulation Environment),[6] a software framework for immersive visualization systems that is used in all of the virtual reality spaces at Calit2. The purpose of these immersive virtual reality systems is to provide more data to the user as a result of the increased resolution and realism than a single monitor would be capable of providing. The user interacts with the 3D environment by using a 3D tracking system that tracks the user's head as well as the input device (a 3D joystick). With the joystick, the user can select individual devices in the container to display live information about them and their status.

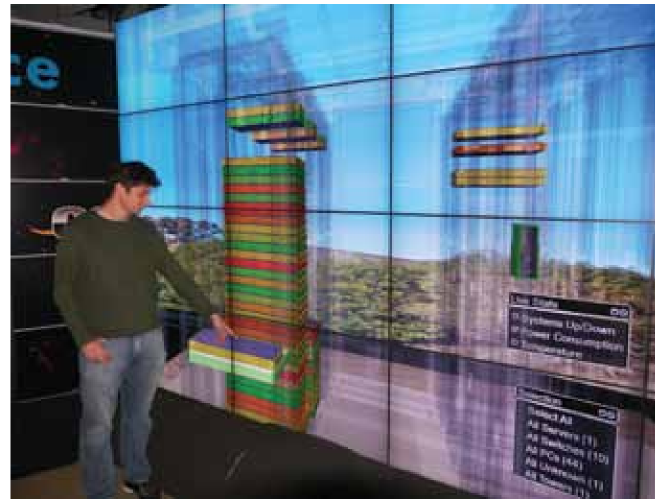

Figure 2. X-ray view of the GreenLight Instrument on AESOP display wall.

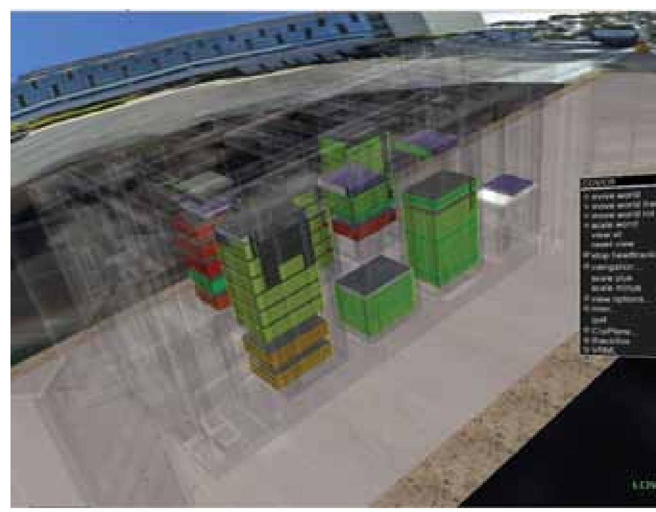

Figure 3. The GreenLight Instrument fully equipped showing the color coded power consumption of each component.

The 3D visualization software has been installed in Calit2's StarCAVE, the NexCAVE, and on the AESOP wall. The StarCAVE[7] is a 5-walled, 34 megapixel, head-tracked 3D environment. The NexCAVE is an LCD-display-based immersive environment with exceptional image quality and 10 megapixels. The AESOP wall is a 4 x 4 tiled display wall, consisting of narrow bezel LCD displays. It cannot show a 3D image in stereo, but is more accessible to the GreenLight researchers. The goal of the AESOP installation is to provide a permanent portal into the GreenLight Instrument, which the participating researchers can approach at any time and intuitively discuss the collected data. **Figures 2 and 4** show examples of the application running on the AESOP display wall.

Both the AESOP wall and the StarCAVE have been partially equipped with power consumption sensors to provide information about the power consumption of these cluster-based display systems.

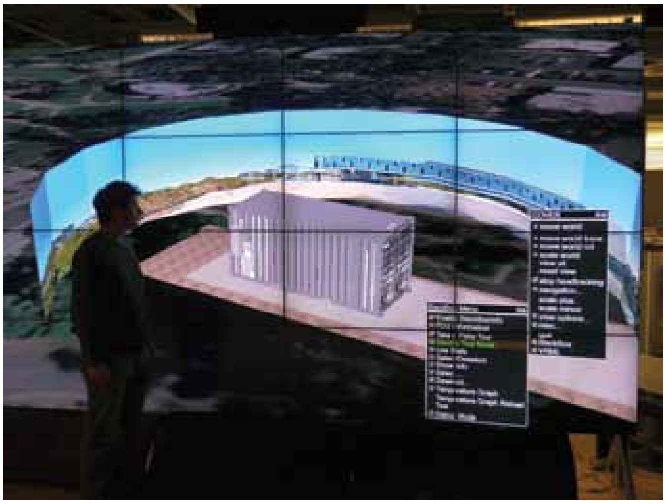

Figure 4. Virtual GreenLight Instrument on the AESOP display wall at Calit2.

The 3D application allows a user to virtually visit the data center, without having to physically go to the container. This saves on transportation cost and time, but also has further reaching consequences—by allowing technicians and researchers to view the status of the container without having to go there, the measurements will not be interrupted by routine maintenance jobs. For instance, every time the container door is physically opened, cold air escapes and has to be regenerated, which affects the power measurements and internal airflow dynamics.

The main purpose of the virtual container is to allow scientists who measure the power consumption and performance of their algorithms, to view these measurements spatially. Given the large number of sensors and their three-dimensional distribution in the container, this is impractical with a web interface, which Calit2 has also created for the GreenLight project,[8] and which provides a quick and easy way to see the measurements, albeit without spatial information. The virtual environment may give clues on whether one hardware component affects another by, for instance, blowing warm air into the other machine's air intake, in which case that machine might run warmer than it would without the hot machine. These side effects can be much more easily detected in a spatial model of the hardware components than in the web interface.

The 3D visualization application can also be used for educational purposes or to reduce maintenance times, by indicating to technicians the containers that experience problems and where the faulty components in the container are located.

RESULTS

The GreenLight project and the visualization of the acquired data has resulted in findings in different areas of computer science.

Data Visualization

One result of our research on how to effectively display the data from the GreenLight Instrument is about the display of this data in the 3D model. Whereas we first implemented a method to display the data as 2D graphs, which we displayed in the 3D environment, we found that this method made it hard to associate the graph with a certain component in the container, even if the component was highlighted. Also, only a limited number of graphs can be displayed at the same time. Therefore, we decided to put the information about the devices directly on their visual representation by coloring their surface with an appropriate color scheme. This approach makes it harder to see small differences in the data because it is hard to distinguish small differences in color, however, it allows the state of all the hardware to be displayed in the container at once.

This visualization method of coloring the components in the container led to implementation of the x-ray mode. Originally, all of the geometry of the container and its components was opaque, so the user had to pull out the computer racks to see the components; this was cumbersome, and did not allow simultaneous viewing of all devices. Making the non-instrumented geometry in the container translucent, allowed for simultaneous viewing of all the instrumented geometry, without the need for moving racks. The user is allowed to switch between opaque and x-ray mode because it is sometimes useful to see the container the way it looks in reality, for example, technicians can train for where to find defective components and how to get to them.

During our research and development of the virtual data center, we found that the configuration of the components in the container changes quite frequently, mostly when new components are added and old ones get replaced. Previously, our human 3D modeler updated the 3D model with a CAD tool to represent the new configuration of devices, and in addition, the sensor data was remapped to reflect these changes. To allow a system administrator without 3D modeling skills to make these changes, and to make them much faster, a database was implemented to describe the configuration of the devices in the container. This approach allows devices to be added, moved, and removed quickly and easily without the need for a 3D modeling tool. The 3D modeling tool is still needed when new types of devices are added to the container for which a 3D model has not yet been created.

Image Rendering

Another example of a result from this project is a study that compares central processing unit (CPU) and graphic processing unit (GPU)-based rendering of a computer generated animation, which is very similar to the type of rendering done by special effects houses in their data centers. In this project, the frames for an animation of fractal geometry, a 3D Julia set (named after French mathematician Gaston Julia), were generated by a GPU cluster in the Sun MD. This is an extreme case of a computationally intensive problem in floating point space where every pixel is rendered after a couple of hundred iterative calculations. It is also an embarrassingly parallel problem that ports well to various parallel architectures. Rendering a 4K animation of approximately 5,000 frames takes more than 20 days on a fast serial processor. This code was run on a 12-node GPU cluster in the Sun MD, each node of which has dual quad-core Intel Xeon E5440 processors with 8 GB of RAM and dual GTX 295 nVidia cards. The result was then compared to the algorithm running on the same cluster, but without using the GPUs and running on the CPUs only. It turns out that the GPU solution outperformed the pure CPU solution by a factor of 30, while consuming only 8% of the energy.

Power Consumption of a VR System

Our final example of a result from the GreenLight project is the power consumption of a large virtual reality system—in this case the StarCAVE. With its multitude of computers, projectors, displays, and tracking devices it is not obvious which elements and operations contribute most to energy consumption. We need to understand its power consumption characteristics to optimize energy efficiency during its usage. The power consumption of the StarCAVE was tested in six typical modes, including (1) with the projectors off, only the PCs running; (2) with the projectors on but no graphics application running; (3) with the graphics application (OpenCOVER) running, but not displaying data; (4) displaying static data, not moving; (5) displaying static data while moving; (6) moving and using a very CPU, GPU and a network intensive application.

Results of the measurements showed that the projectors consume approximately 320 watts, regardless of what data they display, and they consume 9 watts in standby. The PCs consume approximately 295 watts without a graphics application running, and only about 3% more with an application running, regardless of the type of application or how intensely it is being used. The latter result was surprising, because a greater difference in power consumption was expected, once the GPUs became active.

CONCLUSION

Many of the computer and visualization components used in Project GreenLight are the same as those used by the movie industry, whether in movie theaters (projectors, graphics computers for playback) or computer animation houses (CPU and GPU render farms, storage systems, network components). Furthermore, with the massive shift in the industry from analog film to digital data storage, power consumption of IT components and data center management play an increasingly important role in the budgets of production companies. The past and future results of the GreenLight project can inform and guide the movie industry on their path toward complete digitalization, while keeping the related cost and carbon footprint at a minimum. Many of the approaches reported in this paper are applicable to small and large data centers in the industry.

FUTURE WORK

In the future, we plan to use the virtual environment not only to view the state of our data center, but also to actively control it. This will require the addition of a data path back to the data center, along with access control mechanisms, but it will be a very intuitive, yet powerful way to administer a data center. We also plan to install additional sensors to be able to obtain an accurate spatial map of the temperature distribution in the Sun MD. This will help optimize the spatial arrangement of the IT devices in the container to minimize the HVAC requirements.

ACKNOWLEDGMENTS

This publication is based in part on work supported by Award No. US 2008-107, made by King Abdullah University of Science and Technology (KAUST). The author would like to thank Dr. Claudiu Farcas for the statistical data on worldwide energy consumption.

A contribution received May 2010.

REFERENCES

1. (EPA), U.S. Environmental Protection Agency, *Report to Congress on Server and Data Center Energy Efficiency*. Aug. 2007.
2. Computing's Carbon Footprint Gets Bigger, *The New York Times*, Jan. 2009.
3. Oracle. *Sun Modular Data Center*. [Online] [Cited: 05 25, 2010.] http://www.sun.com/service/sunmd/.
4. Michael Manos, "Out of the Box Paradox—Manifested," [Online] [Cited: 05 23, 2010.] http://loosebolts.wordpress.com/2008/10/20/out-of-the-box-paradox-manifested-aka-chicago-area-data-center-begins-its-journey/.
5. Project GreenLight Home Page. [Online] [Cited: 05 23, 2010.] http://greenlight.calit2.net.
6. D. Rantzau, K. Frank, U. Lang, D. Rainer, and U. Woessner, *COVISE in the CUBE: An Environment for Analyzing Large and Complex Simulation Data.*. s.l. : Proceedings of IPTW 98, 1998.
7. T.A. DeFanti, G. Dawe, D.J. Sandin, J.P. Schulze, P. Otto, J. Girado, F. Kuester, L. Smarr, and R. Rao, *TheStarCAVE, a third-generation CAVE and virtual reality OptIPortal*. s.l. : Elsevier, Future Generation Computer Systems 25(2), 2009.
8. Claudiu Farcas, Filippo Seracini. GLIMPSE Home Page. [Online] [Cited: 05 23, 2010.] http://glimpse.calit2.net/.

Jürgen P. Schulze is a Research Scientist at the California Institute for Telecommunications and Information Technology, and a lecturer in the computer science department at the University of California San Diego. His research interests include scientific visualization in virtual environments, human-computer interaction, realtime volume rendering, and graphics algorithms on programmable graphics hardware. He holds an M.S. degree from the University of Massachusetts and a Ph.D. from the University of Stuttgart, Germany. After his graduation he spent two years as a post-doctoral researcher in the Computer Science Department at Brown University.