

Friday, Nov. 10, 2017

Cleary Alumni and Friends Center
Stzrelczyk Great Hall

- 11 a.m. **TECHNICAL SYMPOSIUM**
In-Data Center
Performance Analysis
of a Tensor Processing
Unit
- 4 p.m. **KEYNOTE**
Hints for Special-Purpose
Computer Design

Check-in and refreshments 30 minutes before each lecture.
All events are free and open to the public.

PARKING OPTIONS:

- Visitors can purchase a half (\$3) or full day (\$5) permit from Parking Services, located at 605 17th St. North, in the parking ramp.
- Visitors can also use the new pay stations, any commuter lot, as well as the first level of the parking ramp. (pay-by-phone app also available with these stations)
- Parking information, parking map and a link to purchase permits can be found online at www.uwlax.edu/parking.

For further information contact:

Steve Senger, Ph.D.
Department of Computer Science
University of Wisconsin-La Crosse
1725 State St.
La Crosse, WI 54601
608.785.6805
email: compsci@uwlax.edu

www.cs.uwlax.edu

UNIVERSITY of WISCONSIN
LA CROSSE


COMPUTER SCIENCE DEPARTMENT
221 Wing Technology Center | 1725 State St. | La Crosse, WI 54601 USA
www.cs.uwlax.edu

UNIVERSITY OF WISCONSIN-LA CROSSE

Distinguished Lecture Series in Computer Science

Friday, Nov. 10, 2017

UW-La Crosse Campus




Cliff Young, Ph.D.

Co-sponsored by the

University of Wisconsin-La Crosse Foundation Inc.

Department of Computer Science

College of Science and Health



Cliff Young, Ph.D.

is a member of the Google Brain team, whose mission is to develop deep learning technologies and deploy them throughout Google. He is one of the designers of Google's Tensor Processing Unit (TPU), which is used in production applications including Search, Maps, Photos, and Translate. TPUs also powered AlphaGo's historic 4-1 victory over Go champion Lee Sedol. Before joining Google, Cliff worked at D. E. Shaw Research, building special-purpose supercomputers for molecular dynamics, and at Bell Labs. Research areas include Hardware and Architecture, Machine Intelligence and Software Systems.

UNIVERSITY of WISCONSIN
LA CROSSE

Distinguished Lecture Series in Computer Science

LECTURE TOPICS

SYMPOSIUM IN-DATA CENTER PERFORMANCE ANALYSIS OF A TENSOR PROCESSING UNIT

With the ending of Moore's Law, many computer architects believe that major improvements in cost-energy performance must now come from domain-specific hardware. The Tensor Processing Unit (TPU), deployed in Google datacenters since 2015, is a custom chip that accelerates deep neural networks (DNNs). We compare the TPU to contemporary server-class CPUs and GPUs deployed in the same datacenters. Our benchmark workload, written using the high-level TensorFlow framework, uses production DNN applications that represent 95% of our datacenters' DNN demand. The TPU is an order of magnitude faster than contemporary CPUs and GPUs and its relative performance per Watt is even larger. The TPU's deterministic execution model turns out to be a better match to the response-time requirement of our DNN applications than are the time-varying optimizations of CPUs and GPUs (caches, out-of-order execution, multithreading, multiprocessing, prefetching, ...) that help average throughput more than guaranteed latency. The lack of such features also helps explain why despite having myriad arithmetic units and a big memory, the TPU is relatively small and low power.

KEYNOTE HINTS FOR SPECIAL-PURPOSE COMPUTER DESIGN

A new age of special-purpose computers is beginning. Both supply and demand are forcing computer architects and system designers to explore domain-specific architectures. Our exponential supply of transistors, Moore's Law, is at long last slowing down. At the same time new applications (some driven by deep learning) are generating new demands for effective computation. Machines that focus on a single application offer opportunities to tailor hardware and software structures to algorithms, and perhaps surprisingly, also offer opportunities for the converse, adapting algorithms to build a better system overall. This talk takes its inspiration from Butler Lampson's 1983 'Hints for Computer System Design', and focuses on guidelines, rules of thumb, and experiences relevant to these new kinds of machines.

The University of Wisconsin-La Crosse Distinguished Lecture Series in Computer Science is funded by private gifts to the UW-La Crosse Foundation Inc. and through support from the Department of Computer Science and the College of Science and Health. The purpose of the series is to bring to La Crosse each year a computer scientist whose significant accomplishments can inspire and enrich the careers of students, faculty and the computer community in general.

The Department of Computer Science at UW-La Crosse is the second oldest in the state, behind Madison. Our program was founded in 1968. The department was created as the result of efforts by Jack Storlie, a chemistry professor at the time, who could see that computing would have broad applications in many fields. It has always been a goal of the department to provide students with a strong foundation in software development and the broadest possible opportunity to study the range of sub-disciplines in computer science. The department believes that this maximizes the employment opportunities for our students and well prepares them for a career of innovation in a rapidly evolving discipline.

Currently the department consists of 13 faculty. It offers a B.S. in computer science, Master of Software Engineering (MSE), dual degree five year B.S./MSE degree track and emphasis in Computer Engineering Technology in collaboration with Western Technical College. The department faculty and students are active in research, regularly publish in peer-review journals and give presentations at conferences. It also sponsors a student chapter of the Association for Computing Machinery (ACM), a recently organized Women in Computer Science (WiCS) group and a chapter of the honorary computer science society, Upsilon Pi Epsilon.

For more information about the UWL Department of Computer Science, visit our website at www.cs.uwlax.edu.

www.cs.uwlax.edu