



THE UNIVERSITY OF CHICAGO

Computer Science Department

CERES UNSTOPPABLE SPEAKER SERIES

Rob Schreiber
Engineer
Cerebras Systems Inc.

July 31, 2019, Crerar 390 at 2:00pm

“New Hardware Technology to Accelerate AI”

Abstract:

Deep neural networks are an excellent market opportunity for the computer industry. Training them demands performance far beyond what CPUs can provide. The demand has grown far faster than the growth in chip capability that Moore's Law provides. GPUs do the job better than CPUs and for some time both training and inference have been a new and valuable market niche for the GPU. But the GPU is not optimized for neural networks. Now, new and better adapted architectures are appearing. But the growth in demand for performance continues, and consequently the time to train large networks remains high, hours to days or even more.

Cerebras Systems is a ventured-funded startup. We are designing a new system for training deep networks aimed at more performance than the fastest systems of today. I'll indicate the new approaches in both architecture and hardware technology that allow us to do that.



Bio:

Rob Schreiber is a Distinguished Engineer at Cerebras Systems, Inc., where he works on architecture and programming of systems for accelerated training of deep neural networks. Schreiber's research spans sequential and parallel algorithms for matrix computation, compiler optimization for parallel languages, and high performance computer design. With Moler and Gilbert, he developed the sparse matrix extension of Matlab. He created the NAS CG parallel benchmark. He was a designer of the High Performance Fortran language. Rob led the development at HP of the PICO system for synthesis of custom hardware accelerators. He has help pioneer the exploitation of photonic signaling in processors and networks. He is an ACM Fellow, a SIAM Fellow, and was awarded, in 2012, the Career Prize from the SIAM Activity Group in Supercomputing.

The talk is at 2:00pm, Wednesday, July 31, Crerar 390

Refreshments during talk, Crerar 390

Host: Yanjing Li

Contact: 773-702-3508

<http://cs.uchicago.edu/calendar>

