Facultad de Informática
Universidad Complutense de Madrid

ANUNCIO DE CONFERENCIA

Sparse computing architecture using nonvolatile main memory

D. Peter Hsu

Facultad de Informática

Sala de Grados - 19 de noviembre de 2025 - 11:00

Entrada libre hasta completar el aforo

Resumen:

Datacenter energy use are rising to unsustainable levels. Sparse computing avoids doing redundant or irrelevant computations and has the potential to reduce energy consumption, but it requires irregular memory references and control flow which are very inefficient on existing parallel architectures such as GPU. This talk proposes an architecture using emerging new non-volatile memory technologies such as MRAM as main memory.

We show simulation results of Sparse Matrix-Vector Multiplication (SpMV) suggesting high energy efficiency on this architecture. This is a work-in-progress talk and the results are very preliminary and limited. My objective is to introduce this topic and hopefully generate interest for further research

Sobre Peter Hsu:

Peter Hsu received his Ph.D. from University of Illinois Urbana-Champaige. He started work at IBM T. J. Watson Research Center on the 801 Project. He joined SGI in 1990 as architect of MIPS R8000 TFP microprocessor; the chip powered fifty TOP500 supercomputers in 1994.

Peter co-founded ArtX in 1997 to develop the Nintendo GameCube video game console. He joined Oracle Labs in 2011 as Architect and designed a fifty-thousand core parallel SQL database accelerator. Peter moved to Europe in 2018 and was visiting professor at EPFL University in Switzerland, then senior researcher at the Barcelona Supercomputing Center in Spain working on the European Processor Initiative (EPI) project. He is now an independent consultant living in Girona.