Inici > SORS: A Memory Model for RISC-V

SORS: A Memory Model for RISC-V
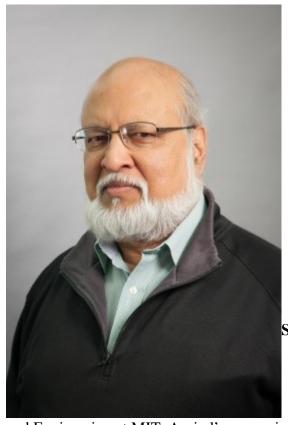
Objectives

Presentation can be viewed following the link

Abstract: Historically memory models for multiprocessors have not been designed deliberately but have just *emerged*. Practically every microarchitectural optimization, which is transparent in a single threaded setting, becomes programmatically visible in a multithreaded setting. This has created a cottage industry for masochists who try to classify and understand every nuance of permitted behaviors.

RISC-V offers us a unique opportunity to fix this historical wrong: specify a memory model first and make implementations conform to it. We will discuss the current proposals being debated for RISC-V memory model and point out the salient issues in this debate. These issues are atomic vs non-atomic memory systems, permitted instruction reorderings especially in the presence of dependencies, fences to control instruction reordering and ease of porting TSO programs to RISC-V. You will have a chance to express your opinions which I will report to the committee.

Short Bio: Arvind is the Johnson Professor of Computer Science and Engineering at MIT. Arvind's group, in collaboration with Motorola, built the Monsoon dataflow machines and its associated software in the late eighties. In 2000, Arvind started Sandburst which was sold to Broadcom in 2006. In 2003, Arvind co-founded Bluespec Inc., an EDA company to produce a set of tools for high-level synthesis.In 2001, Dr. R. S. Nikhil and Arvind published the book "Implicit parallel programming in pH". Arvind's current research focus is to enable rapid development of embedded systems. Arvind is a Fellow of IEEE and ACM, and a member of the National Academy of Engineering and the American Academy of Arts and Sciences.

Barcelona Supercomputing Center - Centro Nacional de Supercomputación