

January 24th, 2019

Synchronized Progress in Interconnection Networks (SPIN), A New Approach to Address Deadlocks

Paul V. Gratz

Texas A&M International University, USA

Abstract

One of the most fundamental design challenges in any interconnection network is that of routing deadlocks. A deadlock is a cyclic dependence between buffers that renders forward progress impossible. Deadlocks are a necessary evil and almost every on-chip/HPC network today avoids it either via routing restrictions across physical channels (Dally's Theory) or with at least one escape virtual channel (Duato's Theory). This ensures that a cyclic dependence between buffers is never created in the first place. Moreover, each solution is tied to a specific topology, requiring an updated policy if the topology were to change. Alternately, solutions have also been proposed to reserve certain resources (buffers) and allocate them only upon detection of a deadlock, thereby breaking the dependence chain and recovering from the deadlock. Unfortunately, all these approaches fundamentally lead to a loss in available bandwidth due to routing restrictions or buffer resource usage restrictions.

This presentation will challenge the theoretical notion of viewing deadlocks as a lack of routing resource (buffers) problem, the view that every solution addressing deadlocks to date is based upon. I will argue that a deadlock can in fact be considered as a lack of coordination between distributed entities. I will show that orchestrating a forward movement of every packet in the deadlocked ring at exactly the same time, can guarantee forward progress and eventually lead to deadlock resolution with a bounded number of rotations.



I will present a new technique based upon these observations, SPIN (Synchronized Progress in Interconnection Networks). I will illustrate the benefit of this approach by designing a novel, truly one-VC, fully- adaptive routing algorithm. I will show that, versus conventional deadlock avoidance techniques, this approach provides up to 80% higher throughput, 52% lower area and 50% lower power for an on- chip 64-core mesh, and up to 83% higher throughput, 53% lower area and 55% lower power for a off-chip 1024-node dragon-fly.

Short bio

Paul V. Gratz is an Associate Professor in the department of Electrical and Computer Engineering at Texas A&M University, currently visiting the University of Edinburgh on sabbatical. His research interests include efficient and reliable design in the context of high performance computer architecture, processor memory systems and on- chip interconnection networks. He received his B.S. and M.S. degrees in Electrical Engineering from The University of Florida in 1994 and 1997 respectively. From 1997 to 2002 he was a design engineer with Intel Corporation. He received his Ph.D. degree in Electrical and Computer Engineering from the University of Texas at Austin in 2008. His papers "Path Confidence based Lookahead Prefetching" and "B-Fetch: Branch Prediction Directed Prefetching for Chip- Multiprocessors" were nominated for best papers at MICRO '16 and MICRO '14 respectively. At

ASPLOS '09, Dr. Gratz received a best paper award for "An Evaluation of the TRIPS Computer System." In 2016 he received the "Distinguished Achievement Award in Teaching – College Level" from the Texas A&M Association of Former Students and in 2017 he received the "Excellence Award in Teaching, 2017" from the Texas A&M College of Engineering.