Advances in the use of Graph Databases to Aid in the Timing Analysis of the World's Fastest Microprocessors

Kerim Kalafala
IBM Systems, Poughkeepsie, NY

Bio:

Kerim Kalafala is a member of the IBM Academy of Technology, a Senior Technical Staff Member in the IBM Systems Group, and an IBM Master Inventor. His current role is lead architect of static timing and noise analysis software tools in IBM EDA. Kerim has received multiple prestigious IBM Research Division awards, an ACM/IEEE Technical Impact Award, as well as a best-paper award at the DAC, and was recognized for co-authoring a top-10 most cited paper in the 50 year history of DAC. He is an inventor on 44 issued patents worldwide and approximately 16 more pending. Kerim is a member of the executive board for the Rhinebeck Science Foundation, and volunteers extensively in his local community. Before joining IBM, Kerim received his undergraduate and graduate degrees in Computer and Systems Engineering from Rensselaer Polytechnic Institute.

Abstract:

When it comes to advanced technology, time-to-market is a critical factor. Processor manufacturers are under immense pressure to deliver new generations of chips that offer faster processing and greater reliability with higher efficiency, which means that each new generation is more complex, and has more (and smaller) components, with four times the number of internal connections. In this context, we will present our experiences using a commercially available graph database (Neo4J) on the Linux on Power platform, as a critical cog to aid in the timing analysis of next generation of IBM Power (P) and Z processors. Ingestion and query analysis on timing graphs consisting of hundreds of millions of nodes and billions of properties will be emphasized along with unique ways in which we are leveraging web-services hosted in a private cloud environment as part of our overall design system.

Sound interesting? Learn about how IBM is using persistent graph databases, hosted in an internal private cloud environment, to tackle the world's most complex computer design and engineering problems. In particular, we will showcase how graph databases are becoming a critical component to the timing analysis of the world's fastest microprocessors (the "brains" of every single computer designed by IBM, including those that sit at the heart of the IBM Watson Cognitive Computing platform) The complexity at hand includes analysis of ~10+ billion transistors to an accuracy tolerance of a trillionth of a second (roughly the time taken for light to travel a fraction of a millimeter). Throughout a microprocessor design project lifecycle, multi terabytes of data are produced, which must be efficiently analyzed and acted upon within ever shrinking time-to-market windows. Within this mountain of data, relationships are crucial entities (e.g., which circuits talk to each other? How does high-level logic design correspond to detailed physical implementation? What are the shortest/longest paths through a given network of electrical components?) - And therefore, we have found graph databases to be uniquely well suited to the analysis tasks at hand.