



# Department of Computer Science

Model-Centric Speculative Parallelization for Scalable Data Processing

Junqiao Qiu  
Michigan Technological University

Hosted by Pan Xu

**Date:** Friday, February 4, 2022

**Time:** 11:00 AM to 12:00 PM

**WebEx:** <https://njit.webex.com/njit/j.php?MTID=m112bce294f15c909a0e07b46b3bb21a9>

<http://cs.njit.edu/seminars>

## **Abstract:**

Exploiting parallelism is key to designing and implementing high-performance data processing on modern processors. However, there are many data processing routines running in serial, due to the sequential nature of their underlying computation models, such as finite-state machines (FSMs), a classic but inherently sequential computational model with applications in data decoding, parsing, and pattern matching.

In this talk, I will present techniques using speculation to “break” the inherent data dependencies, thus enabling scalable data-parallel processing. First, I will introduce a basic speculative parallelization scheme that breaks the state transition dependencies in FSM computations. Then, more interestingly, I will show how a broader range of applications, known as bitstream processing, can benefit from FSM-based speculative parallelization techniques. The key idea is to extract the variable bits that cause dependencies from programs and model their value-changing patterns with FSMs. With this approach, we demonstrate that a rich set of performance-critical bitstream kernels can be effectively parallelized, with up to linear speedups on parallel processors. Finally, I will also briefly discuss the major challenges and my forward-looking research ideas in designing effective speculative parallelization frameworks for critical data-intensive but inherently sequential computations.

## **Bio:**

Dr. Junqiao Qiu is an assistant professor in the Department of Computer Science at Michigan Technological University. He received his Ph.D. in the Computer Science and Engineering Department at University of California Riverside in 2020 under the supervision of Prof. Zhijia Zhao. His research interests are broadly in the area of parallel computing, compiler techniques, and systems. He is a recipient of UC DYP Award, UC Riverside Dean’s Distinguished Fellowship, and Best Paper Award at ASPLOS’20.