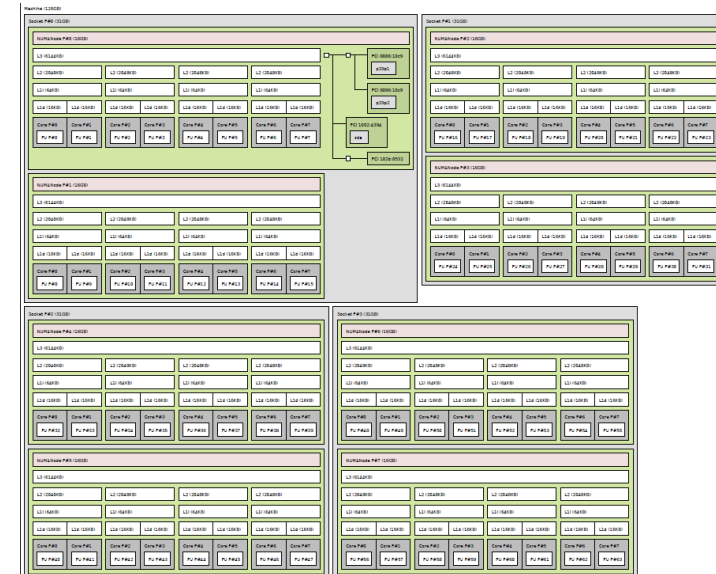
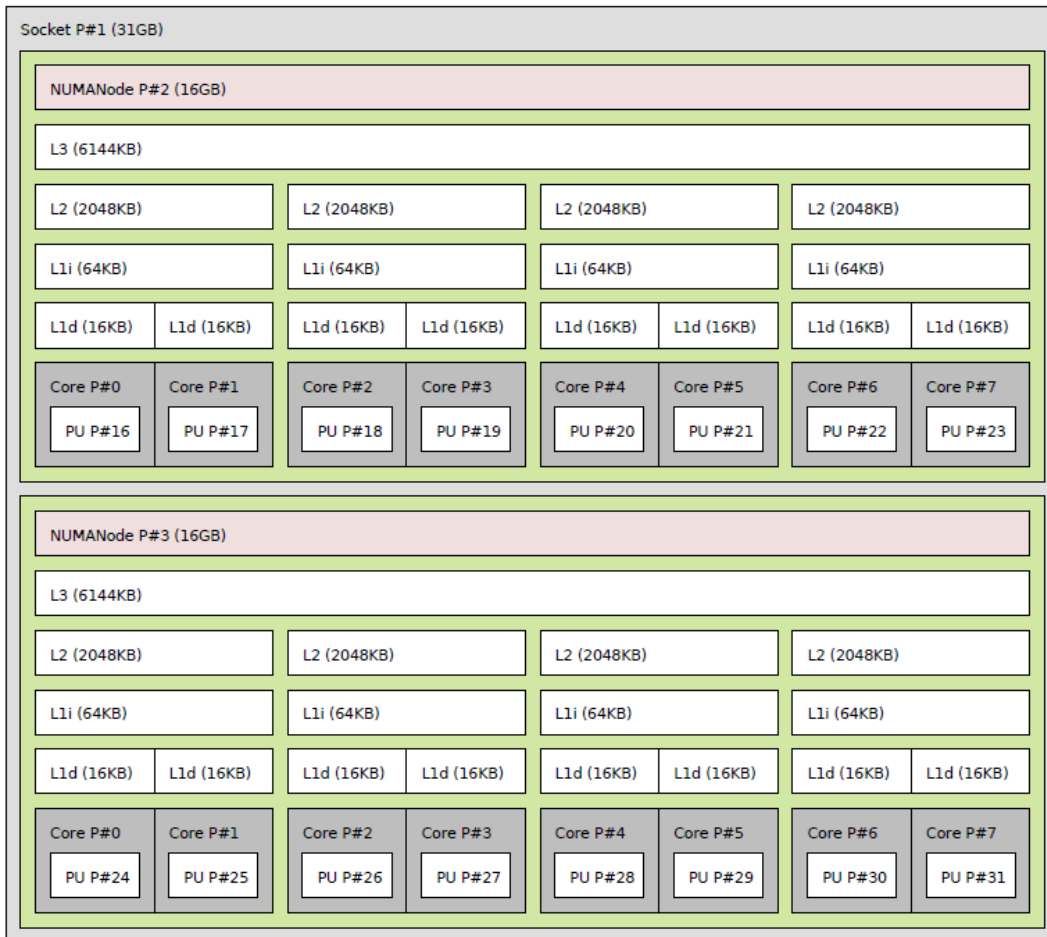


STM Scalability on NUMA architectures

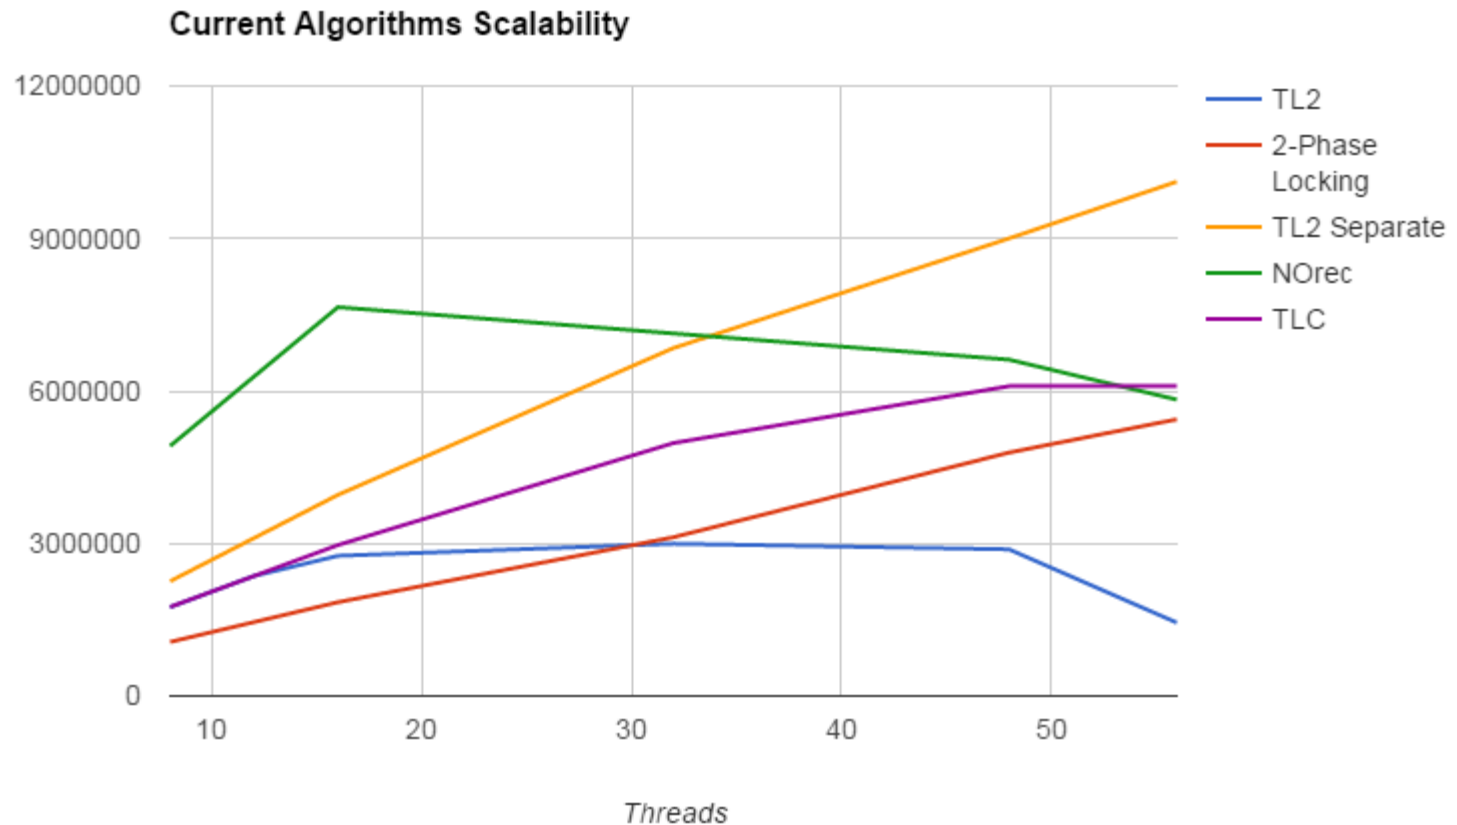
Mohamed Mohamedin, Roberto Palmieri, Binoy Ravindran

Systems Software Research Group
Virginia Tech


Example of NUMA architecture



Current STM Algorithms Do NOT Scale on NUMA




Everything high locality



Intuition of our proposal


- Differentiate between:
 - handling contention within a single socket
 - handling contention among sockets

In your socket:



Make use of current STM algorithms (e.g., TL2) by threads running inside the socket

Across sockets:



Use minimal communication.
Lower priority without starvation

Preliminary results

