# Distributed Reinforcement Learning for Power Limited Many-core System Performance Optimization

Zhuo Chen
Electrical and Computer Engineering
Carnegie Mellon University
Email: zhuoc1@andrew.cmu.edu

Diana Marculescu
Electrical and Computer Engineering
Carnegie Mellon University
Email: dianam@cmu.edu

***Abstract*—As power density emerges as the main constraint for many-core systems, controlling power consumption under the Thermal Design Power (TDP) while maximizing the performance becomes increasingly critical. To dynamically save power, Dynamic Voltage Frequency Scaling (DVFS) techniques have proved to be effective and are widely available commercially. In this paper, we present an On-line Distributed Reinforcement Learning (OD-RL) based DVFS control algorithm for many-core system performance improvement under power constraints. At the finer grain, a per-core Reinforcement Learning (RL) method is used to learn the optimal control policy of the Voltage/Frequency (VF) levels in a system model-free manner. At the coarser grain, an efficient global power budget reallocation algorithm is used to maximize the overall performance. The experiments show that compared to the state-of-the-art algorithms: 1) OD-RL produces up to 98% less budget overshoot, 2) up to 44.3x better throughput per over-the-budget energy and up to 23% higher energy efficiency, and 3) two orders of magnitude speedup over state-of-the-art techniques for systems with hundreds of cores.**