ECE382N.23: Embedded System Design and Modeling

Lecture 10 – ML-Based Mapping & Exploration

Andreas Gerstlauer
Electrical and Computer Engineering
University of Texas at Austin
gerstl@ece.utexas.edu

The University of Texas at Austin
Chandra Department of Electrical and Computer Engineering
Cockrell School of Engineering

Lecture 10: Outline

- **ML-based mapping formulations**
  - Reinforcement learning for mapping & exploration
- **Reinforcement learning**
  - Multi-arm bandits (MAB)
  - Monte-carlo tree search (MCTS)
  - Q-learning
  - Deep reinforcement learning
- **Bayesian optimization**
  - Classic Bayesian methods
  - DSE-specific extensions

Traditional Optimization

- **Traditional (meta-)heuristics**
  - Do not retain knowledge across runs/designs
  - E.g., evolutionary/genetic algorithms

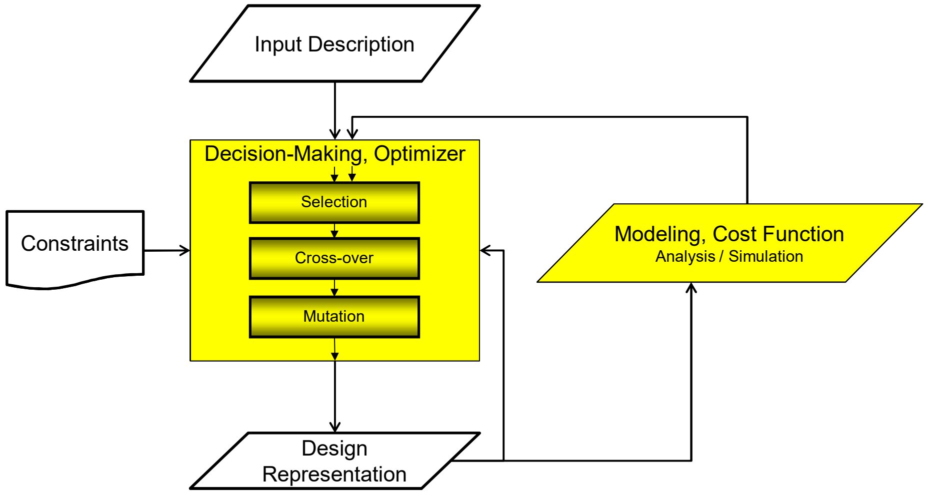

Machine Learning for Optimization

- **Supervised learning**
  - Limited set of decisions with offline training & separate inference
  - E.g. dynamic power management or online prefetcher selection

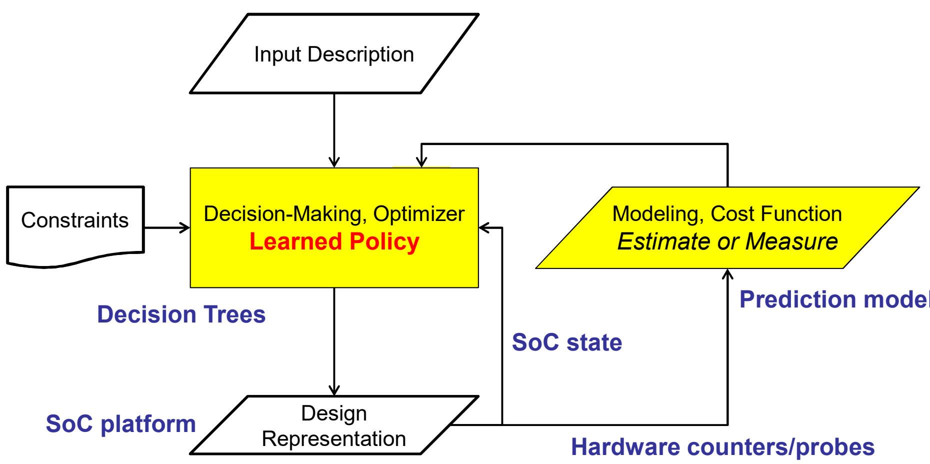

D. DiTomaso, *et al.*, "Machine Learning Enabled Power-Aware Network-on-Chip Design," *DATE*, 2017.
E. S. Alcorta, N. Yadwadkar, A. Gerstlauer, "Lightweight ML-based Runtime Prefetecher Selection on Many-core Platforms," *MLArchSys*, 2023.

Machine Learning for Optimization

- **Reinforcement learning (RL)**
  - Sequential, step-wise decision making
  - Online learning with exploration vs. exploitation
  - E.g. robotics, game play

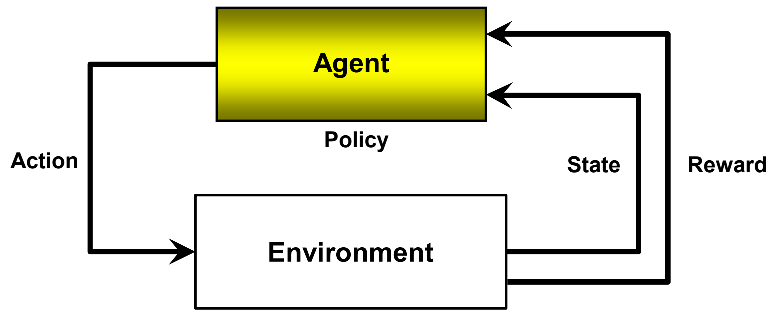

Reinforcement Learning (RL) for Design

- **Natural mapping of the design process**
  - E.g. physical design (floorplanning)

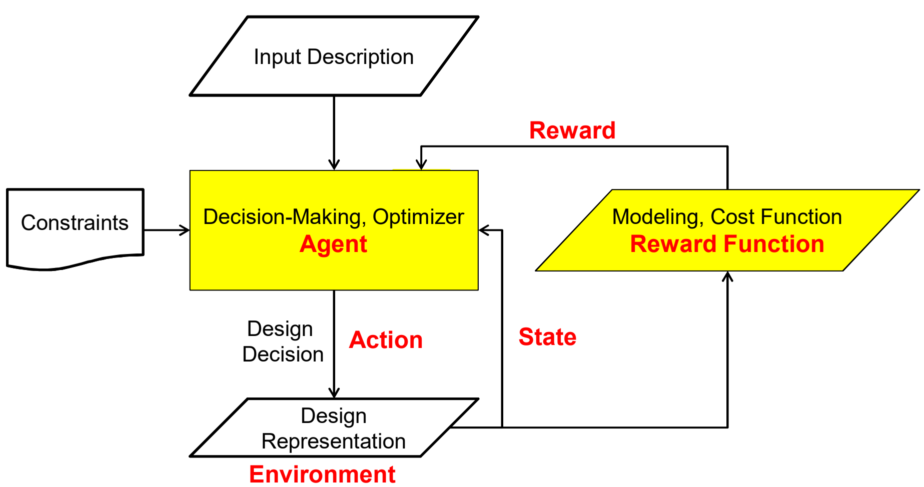

A. Mirhoseini, A. Goldie, M. Yazgan, *et al.*, "A Graph Placement Methodology for Fast Chip Design," *Nature* 594 (2021).

RL for System-Level Design

- **Constructive/incremental**
  - E.g. NoC allocation

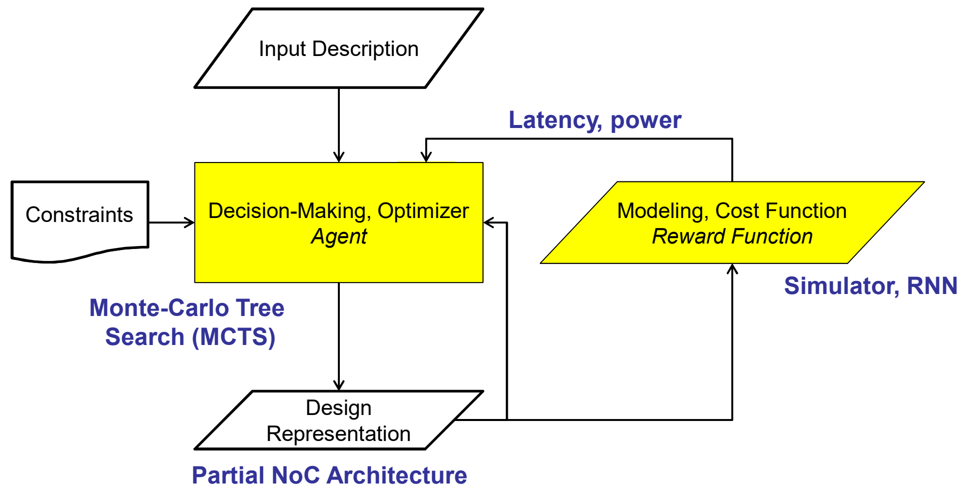

Y. Hu, *et al.*, "Wavefront-MCTS: Multi-objective Design Space Exploration of NoC Architectures based on Monte Carlo Tree Search," in *ICCAD*'18.
Y. Hu, *et al.*, "Machine Learning Approaches for Efficient Design Space Exploration of Application-Specific NoCs," *ACM TODAES*, 25(5), 2020.

RL for System-Level Design

- **Iterative**
  - E.g. online thermal-aware task (re-)mapping

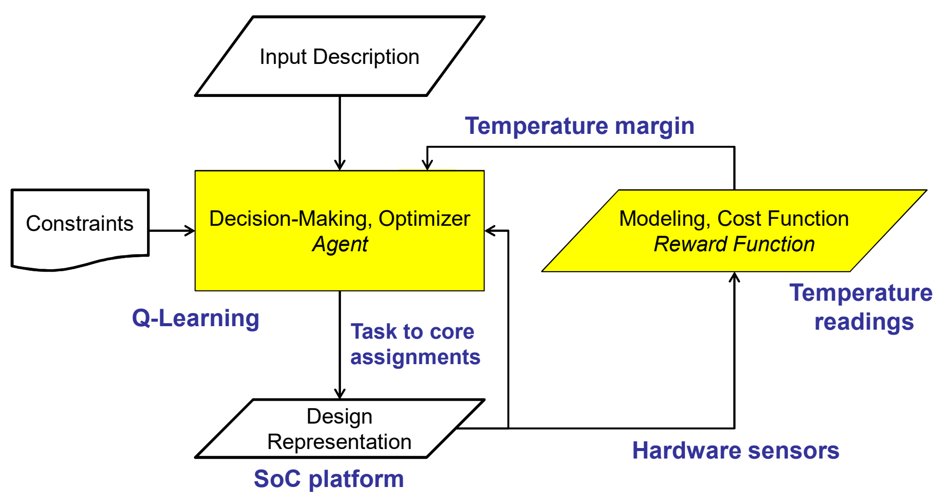

S. Lu, *et al.*, "Reinforcement Learning for Thermal-aware Many-core Task Allocation," in *GLSVLSI*, 2015.
U. Gupta, *et al.*, "A Deep Q-Learning Approach for Dynamic Management of Heterogeneous Processors," *CAL*, 18(1), 2019.

Learning-Based Optimization Challenges

- **Competitiveness of ML approaches**
  - Long history of very effective optimization heuristics
  - Manually encoding human knowledge and intuition
  - Not clear whether ML can compete, e.g. physical design [Markov'23]
- **Training data and overhead**
  - Limited high-quality designs to learn from
  - Online design space exploration
  - Requires fast design quality prediction models
- **Generalization of learned policies**
  - Beyond a single design