Catur
Learning the Norm, Shielding the Tail

Yibo Zhao, Tianyuan Wu, Hui Xue, Qi Chen, Zhenhua Han, Zikai Xu, Yuntai Chang, Rui Gao, Steve Deng, Ray Jui-Hao Chiang, Mingxia Li, Yuqing Yang, Cheng Tan, Fan Yang, Peng Cheng, Yongqiang Xiong, Lili Qiu, Lidong Zhou
Northeastern University, Microsoft Research

Motivation

NUMA placement is critical for cloud VM performance.

- Poor placement causes >30% performance degradation
- 100M+ VMs with diverse types, hardware, and evolving patterns
- Rule-based policies (Xen, KVM, Hyper-V) cannot generalize across the fleet

Key Idea

Catur: RL-based NUMA placement with anomaly shielding.

- **Learning the norm:** RL optimizes long-term average defect with reward shaping and robust action space
- **Shielding the tail:** Speculative state-graph traversal prevents performance anomalies
- Drift-aware continuous learning adapts to evolving workloads

Results

Evaluated on 100 million production VM traces.

Policy	Ticket Ratio	Avg. Defect
Xen	0.94%	1.11%
Nova-Pack	1.22%	1.33%
Tetris	1.28%	1.46%
E-PVM	1.22%	1.41%
Catur	**0.66%**	**0.73%**

Rule-based policies can't handle cloud-scale NUMA diversity. **RL** learns the norm; **shielding** protects the tail — 34–50% fewer defects on 100 million production VMs.