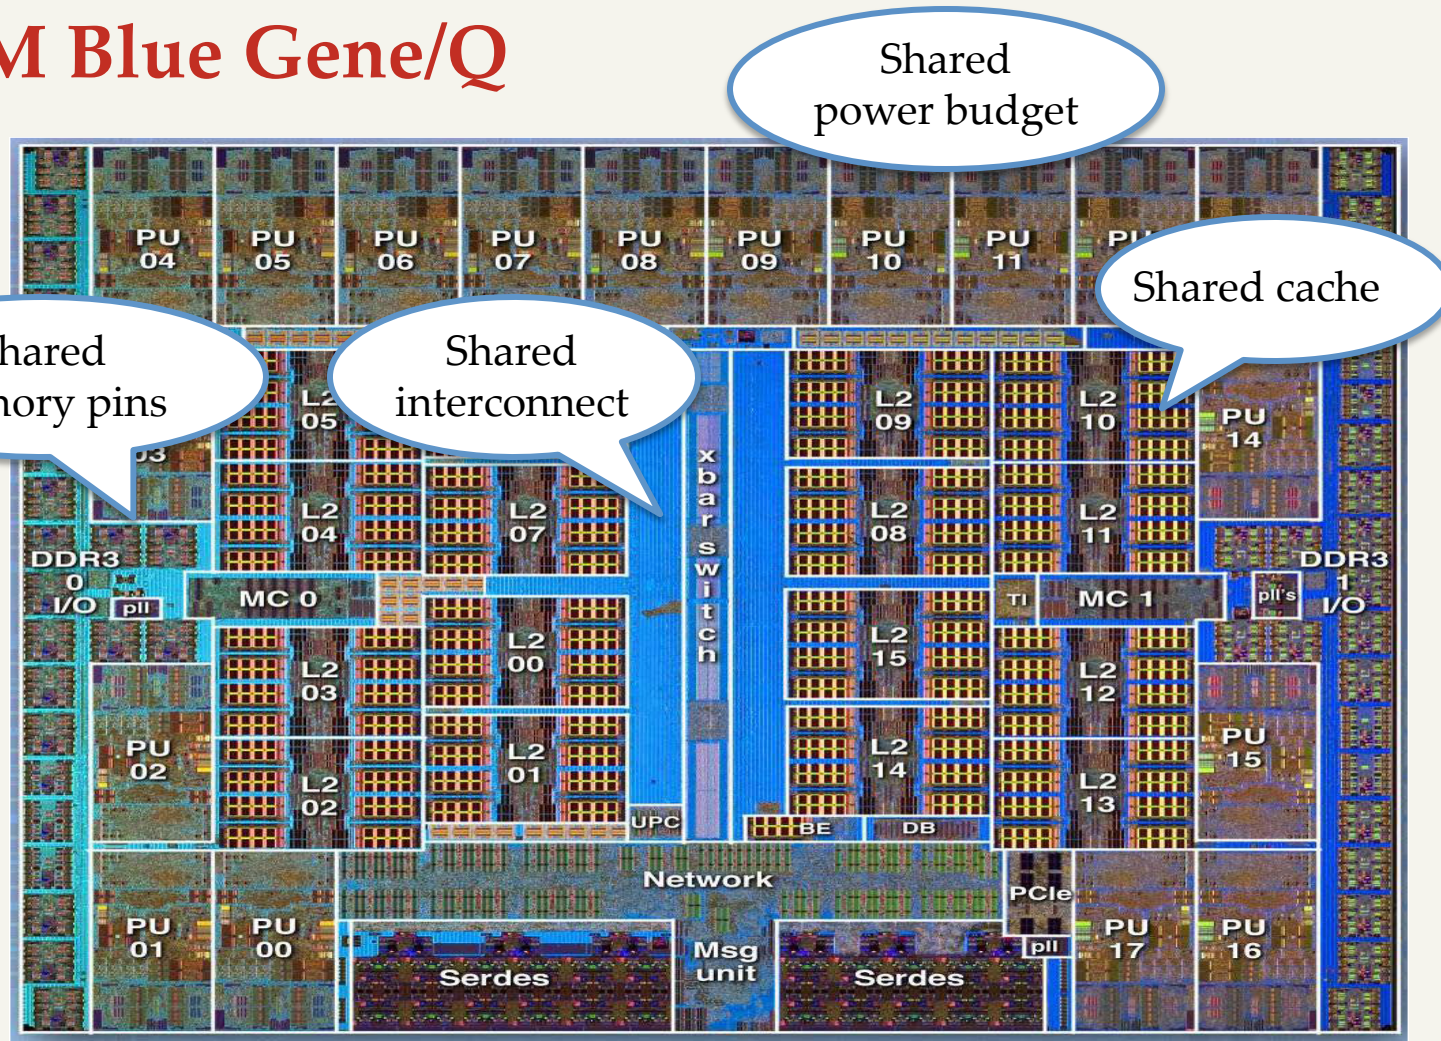




# SHARED ON-CHIP RESOURCES

## ■ IBM Blue Gene/Q




Source: IBM

REBUDGET

## ■ Global allocation space very large

- Exponential increase with no. of cores, no. of resources, granularity of resources
- Global hill-climbing unlikely to scale gracefully






- **Trade off global optimality with simplicity**
  - Global optimum very expensive
  - **Market: Relatively simple, distributed mechanism**
    - » (1) reasonably good; (2) potential for scalability



Source: Wikipedia

- **Effective market-based approach**
- **Market players (cores)**
  - Endowed w/ finite budgets to purchase resources
  - Goal: Maximize own utility—regardless of others
  - Naturally concurrent
- **Market maker**
  - Reconciles bids, posts new prices based on the supply
  - Very fast and simple
- **Converge to market equilibrium dynamically**
  - High performance: 21% average gain in throughput
  - Scalability: < 0.5% overhead for up to 128 cores





## ▪ Market side

- Prices are set in proportional to players' bids

$$price_j = \frac{\sum_i bid_{ij}}{total\_resource_j}$$

$$resource_{ij} = \frac{bid_{ij}}{price_j}$$

## ▪ Player side (cores)

- Model utility—resource relationship
- Bid optimally within its budget constraint to maximize its utility



- **Budget assignment depends on the definition of optimality**
  - **Fairness-oriented:** Give same budget to everyone (XChange-EqualBudget)
  - **Performance-oriented:** Assigning budgets in proportional to the performance gap between minimum and maximum allocation (XChange-Balanced)



## ▪ Optimality of market equilibrium?

- Pareto optimality  $\neq$  Global optimality
- Tragedy of commons
- Performance/fairness bounds?

## ▪ Performance vs. fairness trade-off?

- XChange suggests budget can function as “knob”
  - » Equal budget = “fairness-oriented”
  - » Utility-proportional budget = “performance-oriented”
- Control performance/fairness meaningfully?





- **Theoretic lower bounds with respect to global optimum**
  - System efficiency (performance)
  - Fairness
  
- **Budget assignment mechanism**
  - Builds on top of XChange
  - Control efficiency vs. fairness systematically
  - Evaluation: ReBudget  $\gg$  theoretic bound often



$$Eff = \sum_i U_i(r_i)$$

## ▪ Price of Anarchy (PoA)<sup>2</sup>

- A “superpower” determines the resource allocation that maximize overall efficiency:  $Eff(OPT)$
- PoA measures the cost of distributed market mechanism over global optimum:

$$PoA = \frac{Eff(Market)}{Eff(OPT)}$$

- PoA is lower bound: Worst-case system efficiency
  - » Market *at least as good as* PoA (alt., can be *as bad as* PoA)

[2] Papadimitriou, *Algorithms, games, and the Internet*, STOC 2004.



## ■ How “tight” can PoA bound be?

- The best-known theoretic bound of efficiency loss for a market under some specific conditions:<sup>3</sup> PoA is 75% of global optimum

## ■ How close to theoretic limit can XChange be?

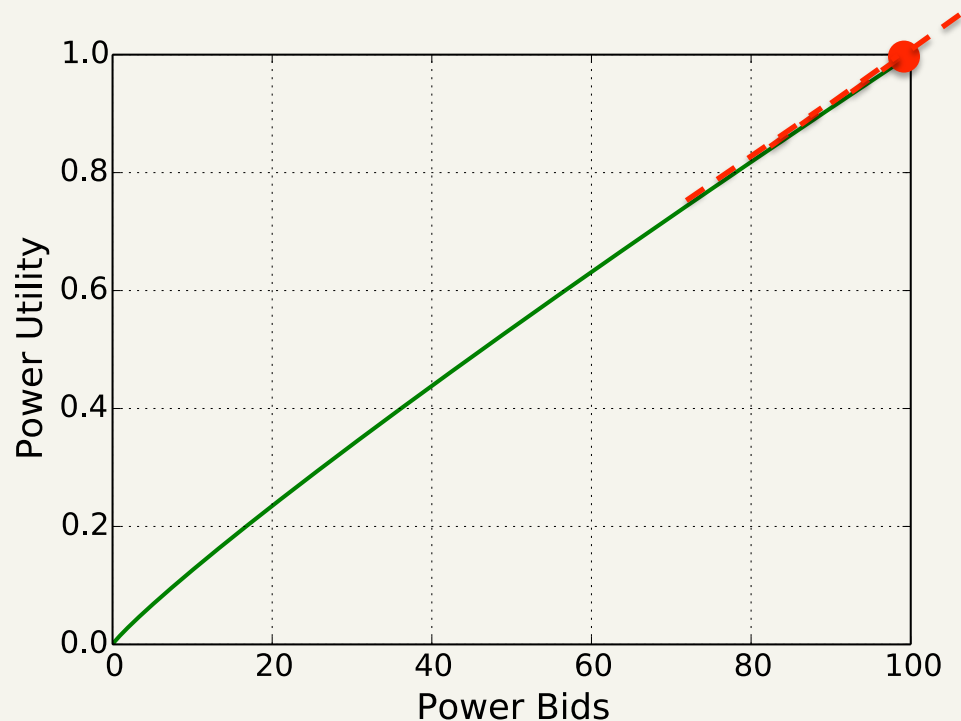

- Hypothesis: We can tighten PoA in XChange by adjusting players' budgets relative to each other

[3] Ramesh Johari and John N Tsitsiklis. Efficiency loss in a network resource allocation game, Mathematics of Operations Research, 29(3):407–435, 2004.



## ■ Market players (cores)

- Have finite budgets  $B_i$  to purchase resources in the system
- Try to maximize their own utility by bidding resources



$$\text{Utility} = \text{Utility}_{\text{cache}} + \text{Utility}_{\text{power}}$$



## ■ Player's incentive to adjust bids


- Intuition: if a player bids optimally, its “incentive” for different resources are the same
- Otherwise, a revision of its bids can improve its utility

## ■ Mathematical representation

- Assume utility functions are smooth and concave\*
- Define player  $i$ 's marginal utility (incentive) for resource  $j$ :  $\lambda_{ij} = \frac{\partial U_i}{\partial b_{ij}}$ , and incentive of player  $i$ :  $\lambda_i = \max_j \lambda_{ij}$
- Lagrange multiplier method:

$$\lambda_{ij} = \frac{\partial U_i}{\partial b_{ij}} \begin{cases} = \lambda_i & \text{if } b_{ij} > 0 \\ < \lambda_i & \text{if } b_{ij} = 0 \end{cases}$$

- Intuition: at equilibrium, player  $i$ 's incentive  $\lambda_i$  measures its utility improvement if it is given more money (i.e., assigned a larger budget)
  - Consider 2-player market:

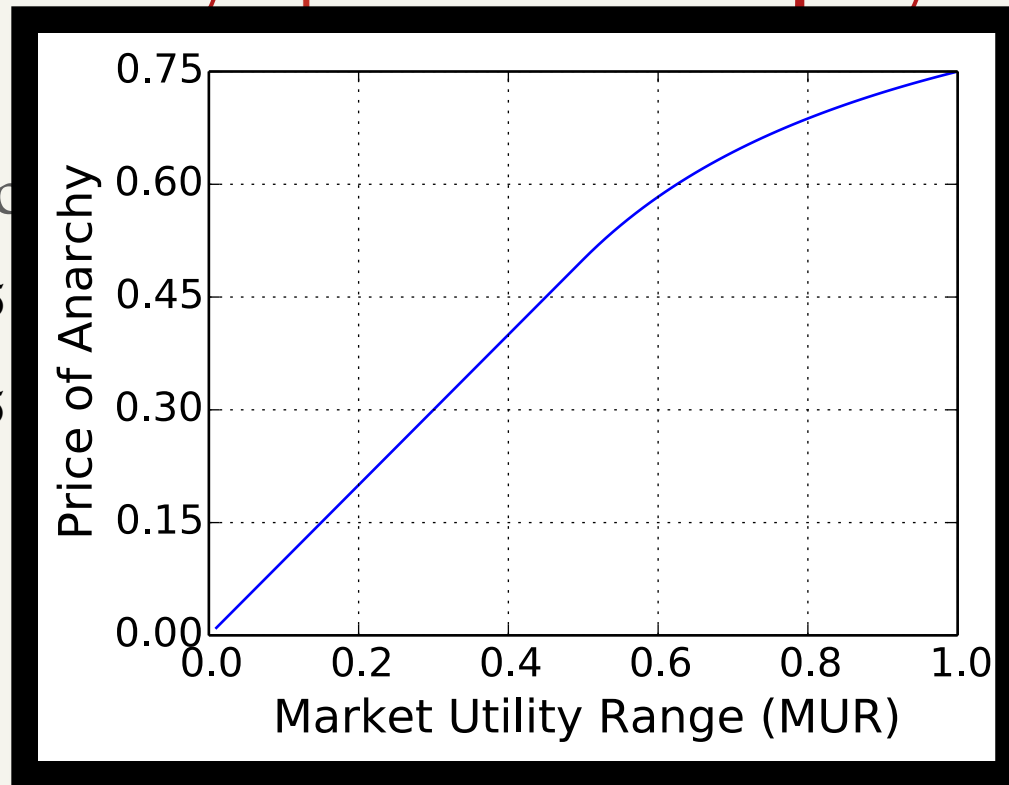


- **Market utility range (MUR):** maximum variation of marginal utility  $\lambda_i$  of all market players:

- Introduction

If  $MUR \geq 0.5$

If  $MUR < 0.5$



## ▪ Proof

[3] Ramesh Johari and John N Tsitsiklis. Efficiency loss in a network resource allocation game, Mathematics of Operations Research, 29(3):407–435, 2004.



## ■ System fairness: Envy-freeness

- Envy: how a player prefer others resources:  $EF_i(r) = \frac{U_i(r_i)}{\max_j U_i(r_j)}$
- Given allocation  $\mathbf{r} = (\mathbf{r}_1, \mathbf{r}_2, \dots)$ , envy-freeness is defined:  
$$EF(r) = \min_i EF_i = \min_{i,j} \frac{U_i(r_i)}{U_i(r_j)}$$
- C-approximate envy-free:  $EF(r) \geq C$






- Players' budget difference is a natural way to measure fairness in the market

- Equal budget
- Larger budget

- Define Market Budget Range (MBR) as the variation in budget

- Introduce Envy-freeness [4]:  
• The market is  $(2\sqrt{1+MBR})$  times more efficient than the MBR




[4] Li Zhang. The efficiency and fairness of a fixed budget resource allocation game. In *Automata, Languages and Programming*, pages 485–496. Springer, 2005.

- **Measuring MUR and MBR, we can quantify the loss in efficiency and fairness of market equilibrium**
- **Use MUR and MBR to guide budget re-assignment**
  - To improve PoA: punish the players whose incentives are low
  - Guarantee minimum fairness set by system admin

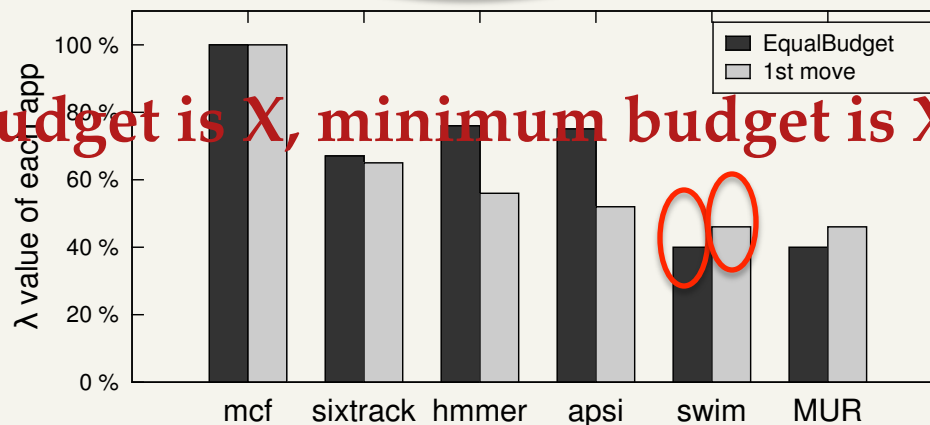


## ■ Given a fairness target

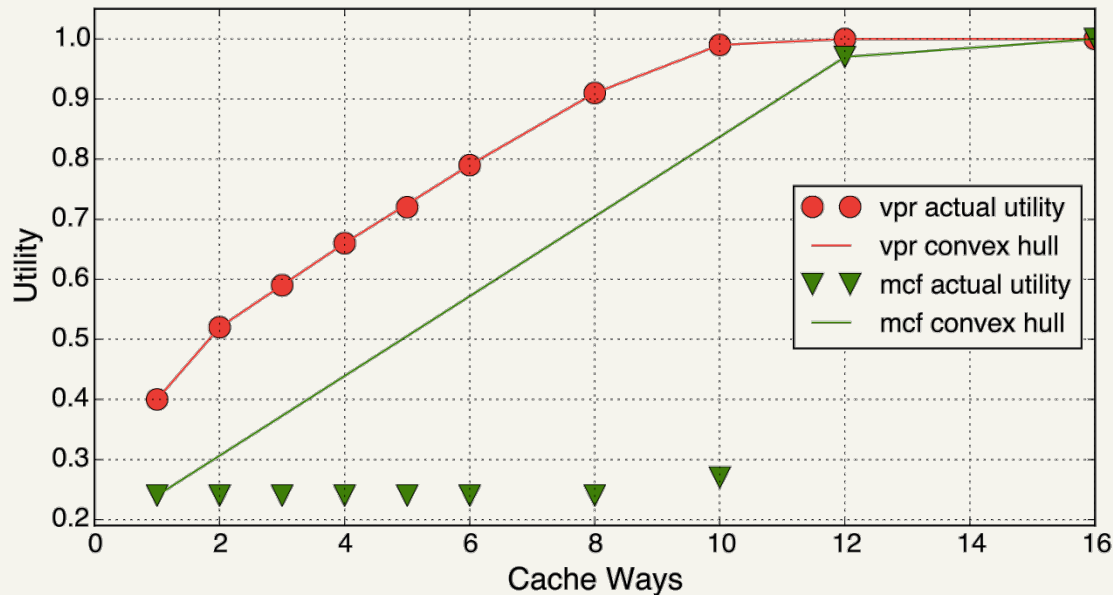
1. Compute MBR:  $EF = 2\sqrt{1 + MBR} - 2$  ; initialize “penalty”:  $(1 - MBR) \cdot \frac{X}{2}$
2. Use XChange to reach market equilibrium
3. Player  $i$  with incentive lower than 50% of the maximum incentive has to pay a “penalty” (equivalent to a budget cut)



## ■ Maximum Budget is $X$ , minimum budget is $X \cdot MBR$



- **Assumption: a player's utility function is non-decreasing, continuous, and concave**
  - Adopt Talus<sup>4</sup> and Futility Scaling<sup>5</sup>, which convexifies the cache behavior



[4] Beckmann and Sanchez, HPCA, 2015.

[5] Wang and Chen, MICRO 2014



## ■ Simulation setup

- 4 GHz 4-way OoO core, 32 kB i/d L1
- 8- and 64-core CMP
- 4MB (16 way) and 32MB (32 way) L2 Cache
- 10W per core as equal-share
- DDR3-1600 channels, 4 ranks ea., 8 banks/rank

## ■ Performance analysis

- Mix of SPEC2000 and SPEC2006 multi-programmed workloads



## ▪ **Analytical experiment**

- Profile 24 SPEC 2000 and SPEC 2006 applications
- Assume cache and power behavior are perfectly continuous and concave
- Create 240 bundles that fall into six categories: CPBN, CCPP, CPBB, BBNN, BBPN, and BBCN

## ▪ **Implement ReBudget in simulator**

- Utility function is modeled in the run-time
- ReBudget is triggered every 1ms to adapt to application phase changes

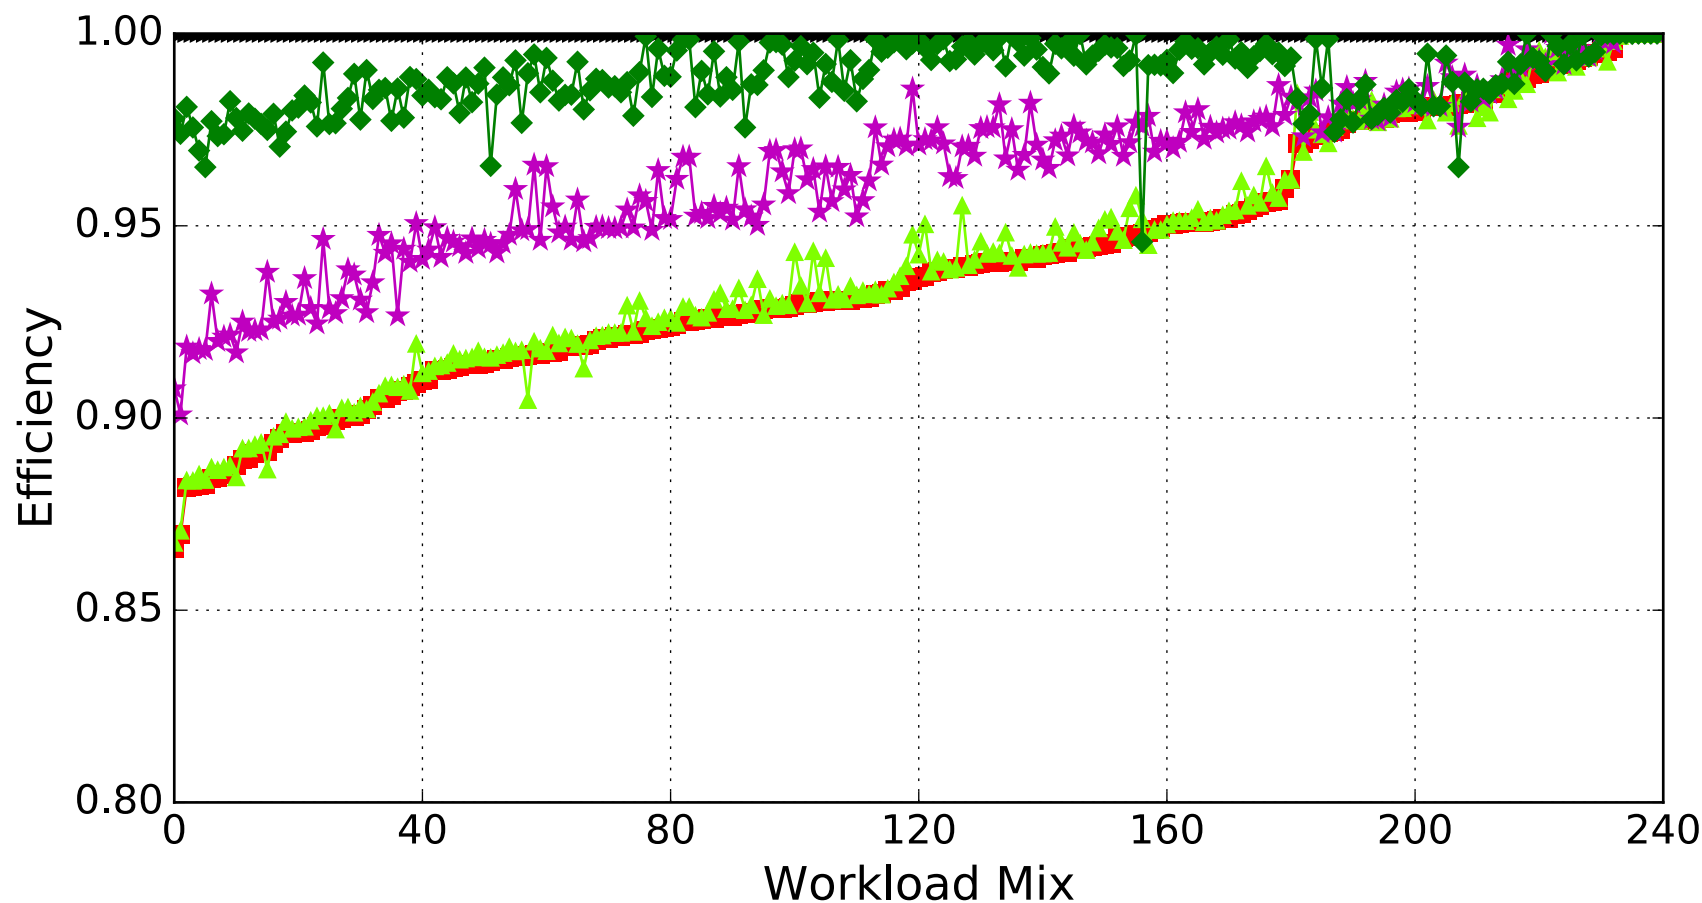


# EFFICIENCY

XChange,  
Wang and Martínez,  
HPCA 2015

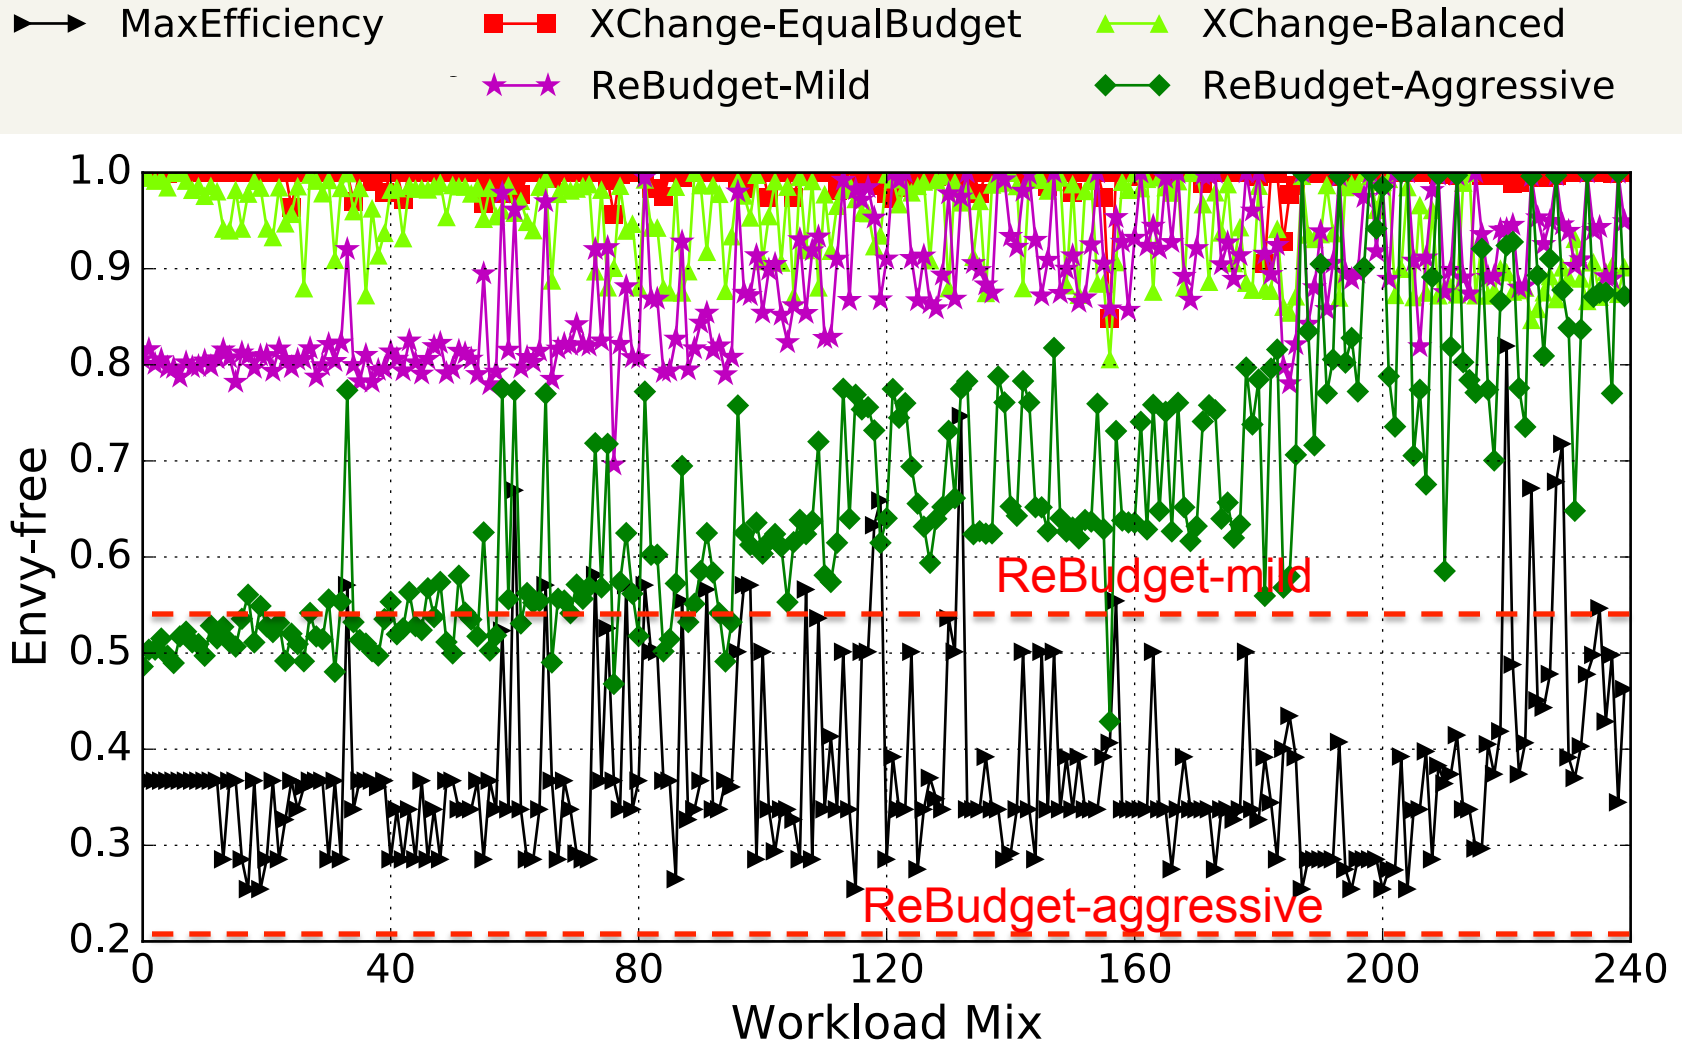
➡ MaxEfficiency

|   |                     |   |                     |
|---|---------------------|---|---------------------|
| ■ | XChange-EqualBudget | ▲ | XChange-Balanced    |
| ★ | ReBudget-Mild       | ◆ | ReBudget-Aggressive |




REBUDGET








# SIMULATION RESULTS: EFFICIENCY





- **Heuristic-based market-based approach can suffer low efficiency: tragedy of commons**
- **We provide theoretic guarantees on system efficiency and fairness for arbitrary budget assignment**
  - Introduce two metrics: market utility range (MUR) and market budget range (MBR) to quantify the loss of system throughput and fairness
- **ReBudget efficiently trade off system efficiency and fairness**
  - ReBudget-aggressive achieves 95% efficiency for all application bundles





# **ReBUDGET: TRADING OFF EFFICIENCY VS. FAIRNESS IN MARKET-BASED MULTICORE RESOURCE ALLOCATION**

Xiaodong Wang

José F. Martínez

Computer Systems Lab  
Cornell University

