Cost-aware scheduling for heterogeneous enterprise machines (CASH'EM)

Jennifer Burge, Partha Ranganathan, Janet Wiener

Duke University and HP Labs

Data center provider decisions

- Provisioning
  - Which machines to buy? Replace?
- Scheduling
  - Which customers to write contracts for?
  - Which machines to allocate to each customer and when?
    - Which machines to leave idle? Turn off?
    - What is risk that higher-paying customer will arrive?

Scheduling

- Problem: assign customer requests to machines
  - In time to meet customer requirements
  - Exact workload not known in advance
- Goal: maximize profit = value – cost
- Heterogeneous machines have variable costs
  - All prior work assumes uniform costs…
- Customers with different requirements (SLAs) have variable values and deadlines
- Handle overload and underload

Data center coordinator

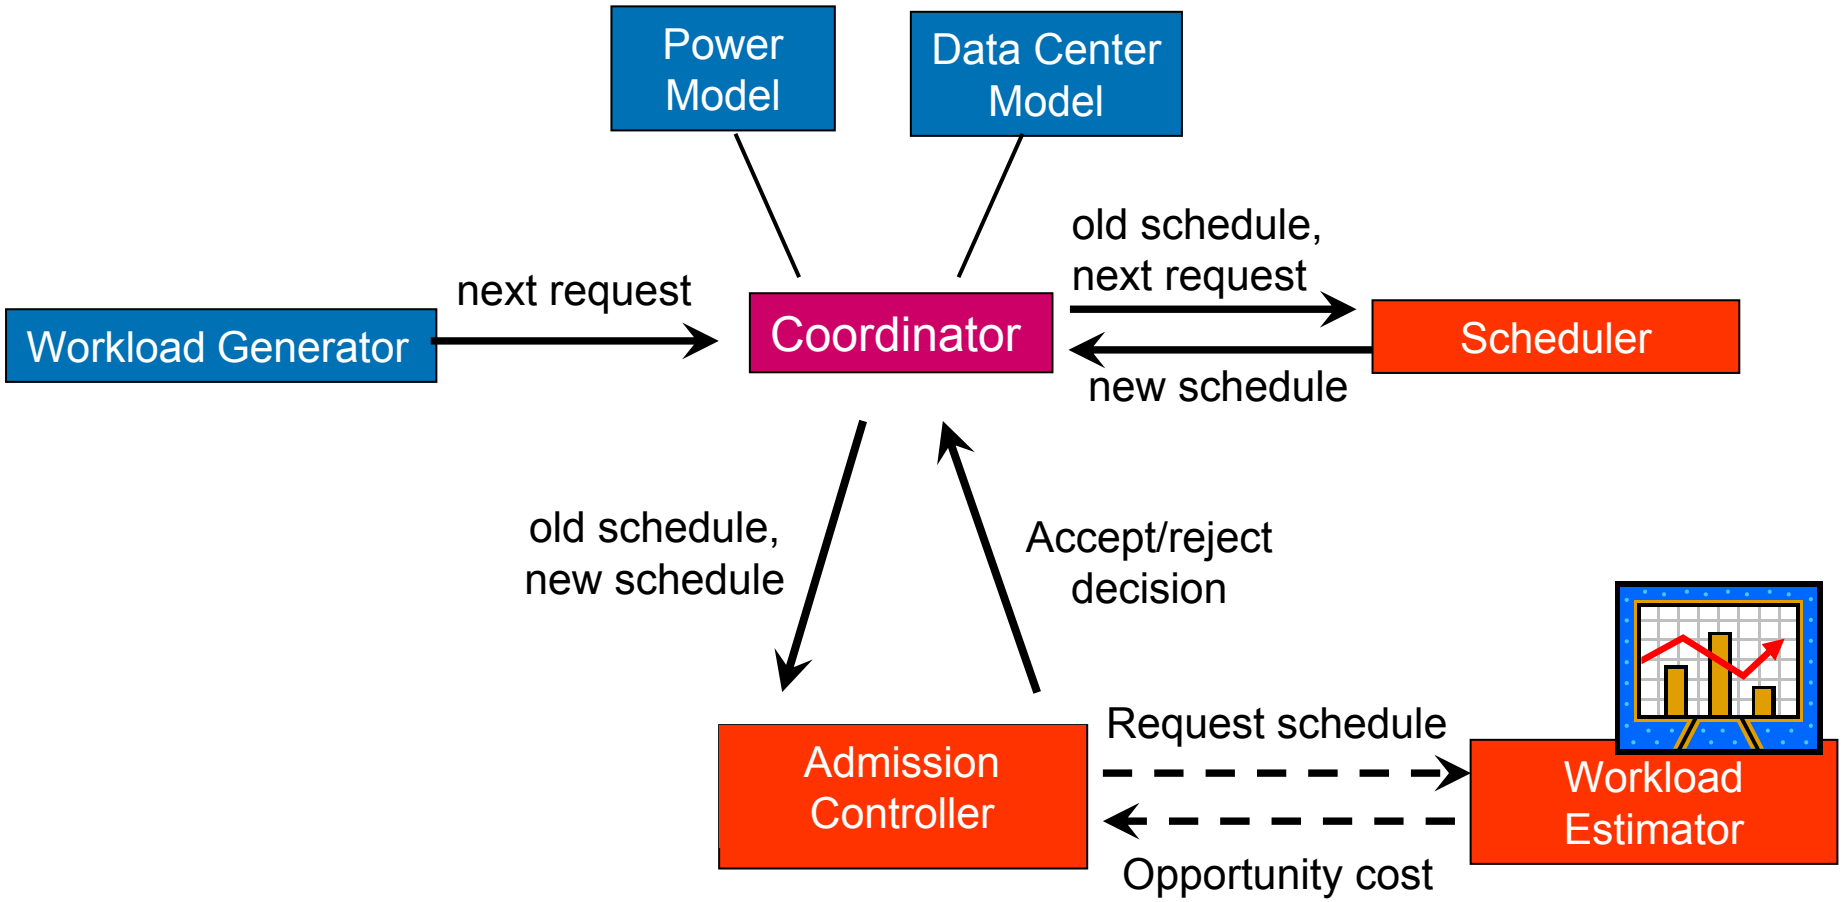

Scheduling algorithms

- FIFO
  - Each new request goes to next available machine
- FIFO-Opportunity
  - Look at first available machine of each type
  - Choose machine where request earns highest profit
    - Prefers cheaper machine if request will finish on time
- Profit-rate [AuYoung2006]
  - Prioritize requests by profit-rate
- “Best” (unachievable): profit upper bound
  - Offline, ignore deadlines

Admission control algorithms

- Higher
  - Compare schedule with and without new request
  - Admit request if schedule with request is more profitable
    - May pay penalties for rejection, cancellation of requests
- Risk
  - *"By scheduling this request, I risk missing better ones"*
  - Predict arrival rates and values of future requests
  - Calculate opportunity cost
    - Will more profitable request arrive soon?
    - Will it be rejected or delayed by this request?
    - Time farther in the future = higher opportunity cost
    - Fewer machines available = higher opportunity cost

Simulation data center

- Power Model

Machine class	100% util watts	Idle watts	Off watts	Time factor	Power-perf
BL	150	60	0	1.00	150
DL	200	100	25	1.25	250
BC	25	10	0	3.75	92

- Data center model
  - 100 machines
  - Equal amounts of each type
  - Mimics HP data center in Germany

Simulation workload

- 90 days at 10 minute granularity
- Requests arrive during simulation
  - 6 minute inter-arrival time, exponential rate
- 4 request value-rates
  - 1.2 to 4.4 cents/hour (each 25%)
  - Only 2 highest value-rates are profitable on all machines
- 2 request durations
  - 2.5 hours (50%), 24 hours (50%)
- 2 request patience levels
  - 1 week (most), 2 hours (10%)

Compare algorithms

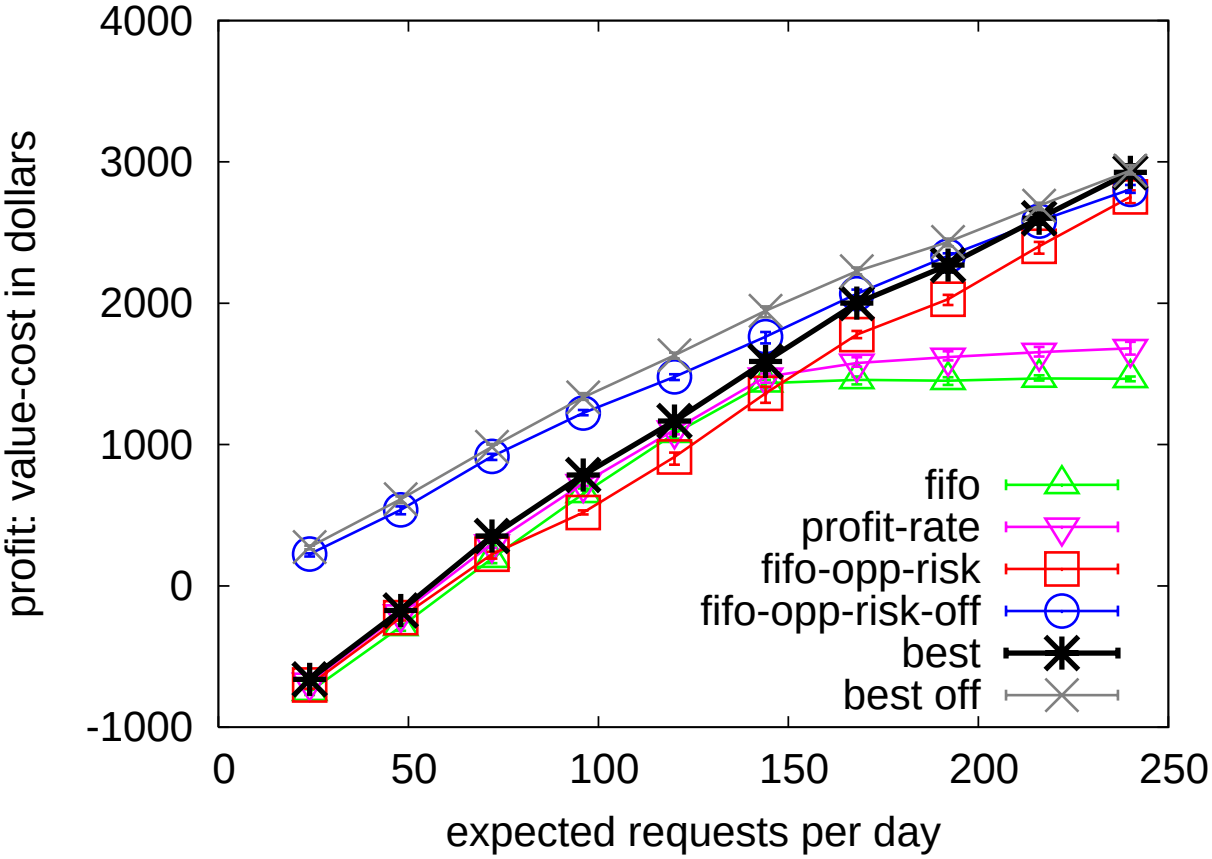

- At undersaturation, turn machines off
- At oversaturation, risk algorithms win

Fate of requests by duration

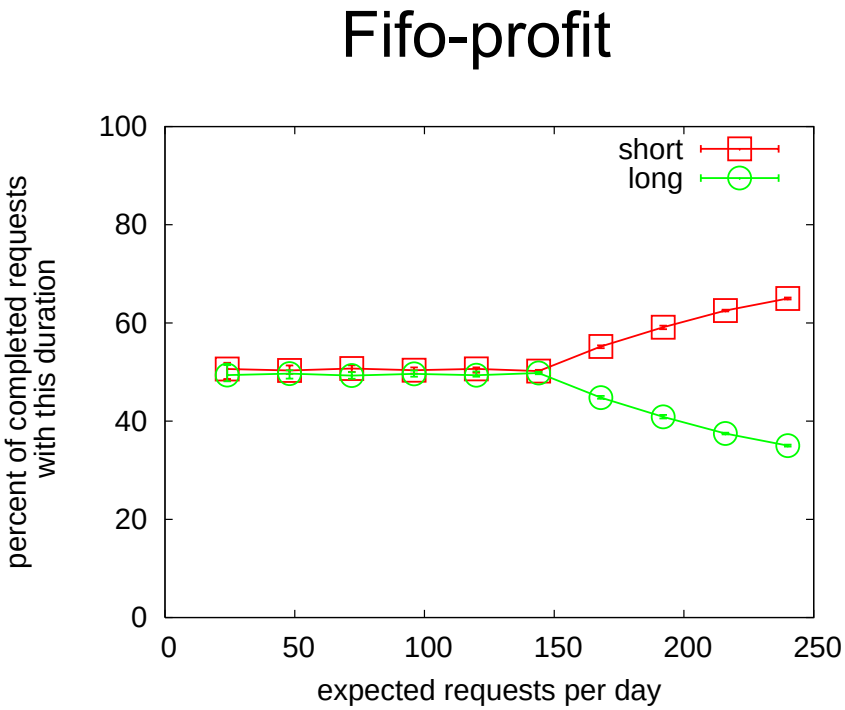

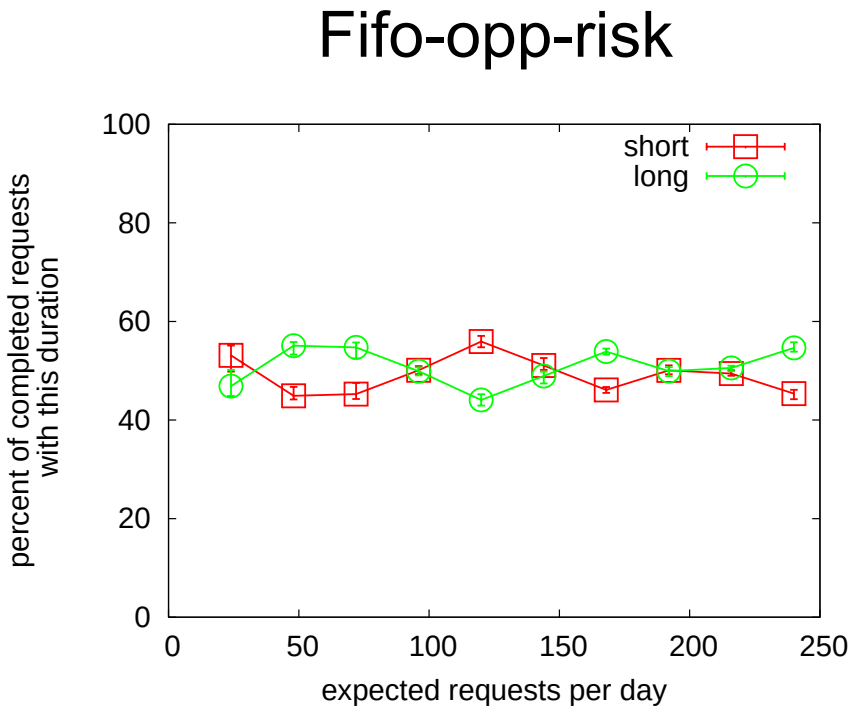

- Fifo-profit favors short requests on slower machines
- Fifo-opp-risk does not discriminate by duration
  - So it can discriminate by value

Fate of requests by value-rate

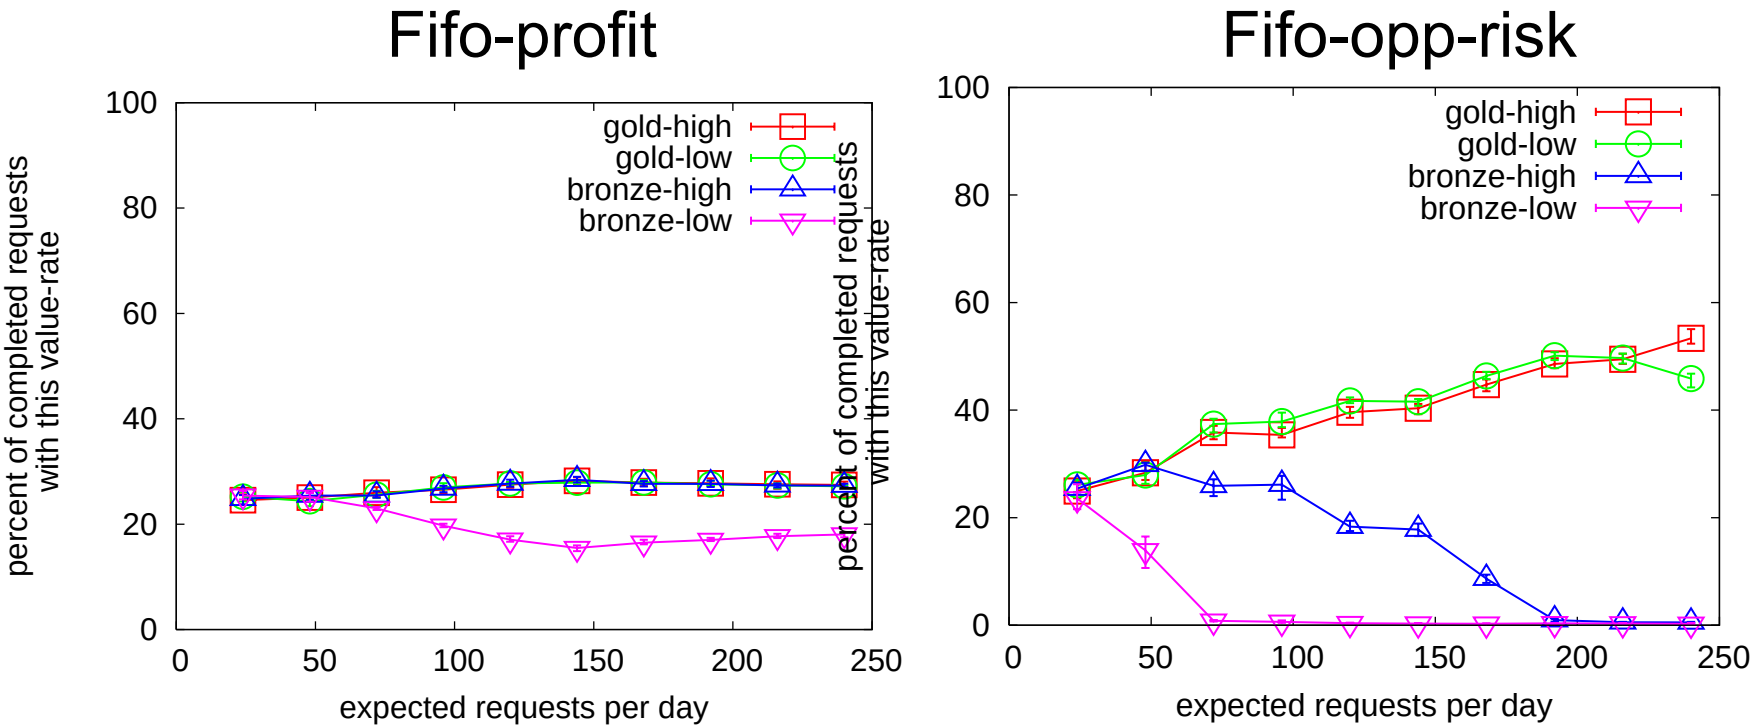

- Fifo-opp-risk prefers the highest-valued requests
  - It essentially saves space in schedule for them

Utilization of machines of different types

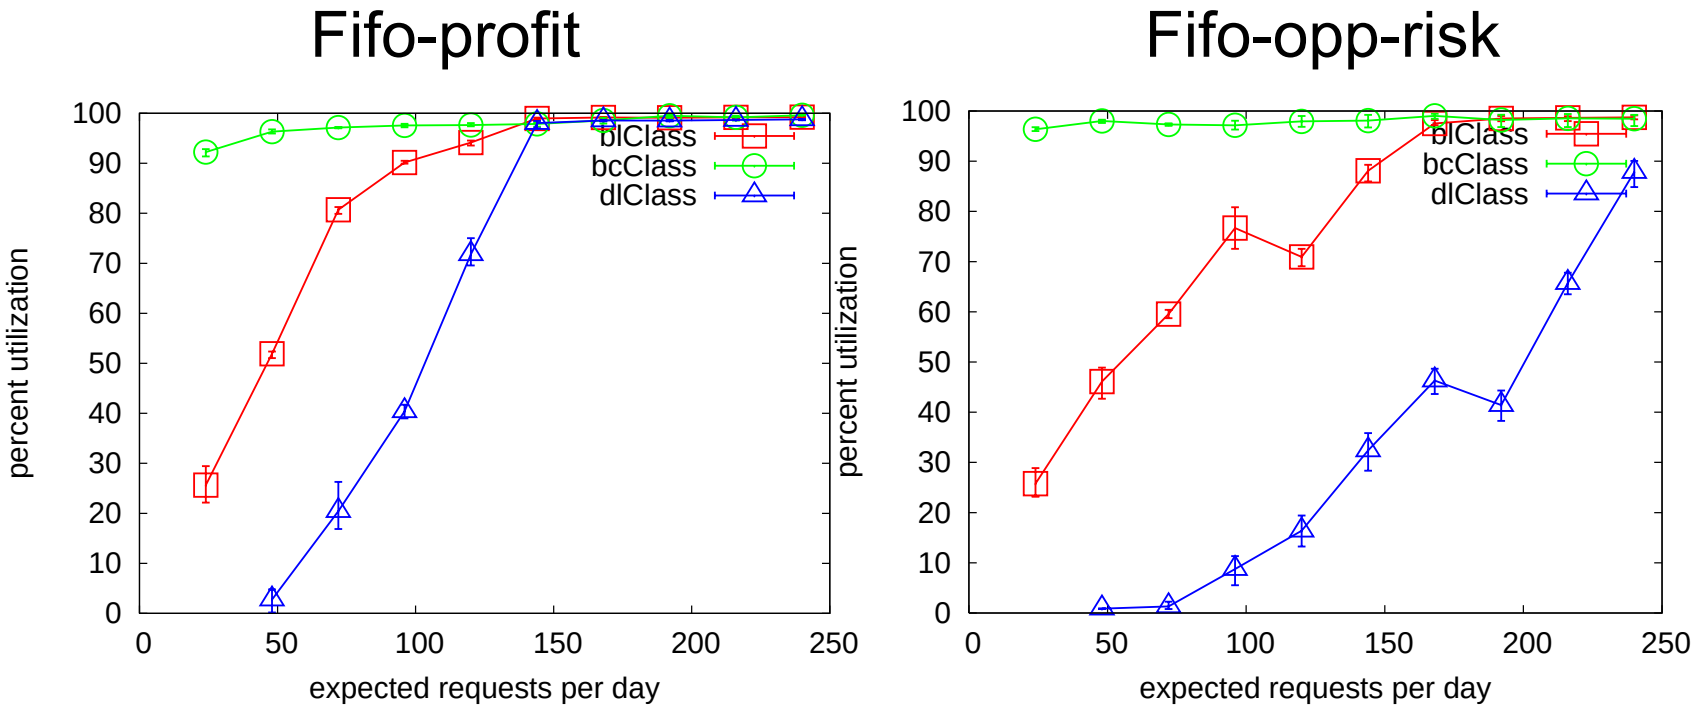

- Fifo-opp-risk requires more requests to saturate the data center
  - Only high-value requests make a profit on DL machines

Related work

- Profit-maximizing online scheduling algorithms
  - Millenium, RiskReward ignore cost
- Scheduling algorithms that predict risk
  - Disk scheduling, online knapsack
- Scheduling to conserve power
  - Use power-states
  - Turn machines off
  - Focus on reducing cost, no flexibility in value

Conclusions

- It matters which machines are assigned to which customer
  - Especially when data center is undersaturated
- Risk-based admission control is effective
  - Anticipate likelihood of higher-value requests in future
- Turning unused machines off saves $$
- Heterogeneity in machines is actually good
  - Can satisfy some customers with lower cost, lower power machines