List scheduling heuristics for virtual machine mapping in cloud systems

S. Nesmachnow
S. Iturriaga

B. Dorronsoro
E.-G. Talbi

P. Bouvry

Universidad de la República
Uruguay

INRIA/University of Lille
France

University of Luxembourg
Luxembourg

Outline

1. Motivations and Contributions
2. The Virtual Machine Mapping Problem
3. List Scheduling Algorithms
4. Experiments
5. Conclusions and Future Work

Outline

1. **Motivations and Contributions**
2. The Virtual Machine Mapping Problem
3. List Scheduling Algorithms
4. Experiments
5. Conclusions and Future Work

Motivations

- Cloud computing
  - Novel computing paradigm
  - Cloud service provisioning models: IaaS, PaaS, SaaS
  - VMs that can be easily allocated and deallocated
  - Elasticity, flexibility, seemingly infinity of resources, etc.
- Broker: intermediary entity between cloud providers and users
  - Finding the best deal
  - 'Cloudifying' applications
- New business model for cloud
  - Book reserved instances (RI) on a number of cloud providers
    - Low investment
  - Sublet them to its customer as on-demand resources
    - 20% cheaper than the price cloud providers offer
  - Overloaded situations: cloud bursting

Contributions

- Definition of the novel business model
- Formulation of the optimization problem that arises
- Resolution of the problem with 8 heuristics
- Generation of a novel benchmark
  - 400 instances
  - Diverse workloads and scenarios
  - Real pricing data (AWS and Azure)
  - Available: http://www.fing.edu.uy/inco/grupos/cecal/hpc/VMMP

Outline

1. Motivations and Contributions
2. **The Virtual Machine Mapping Problem**
3. List Scheduling Algorithms
4. Experiments
5. Conclusions and Future Work

The VMMP Problem

- Allocate all customers VMs requests into the available RIs

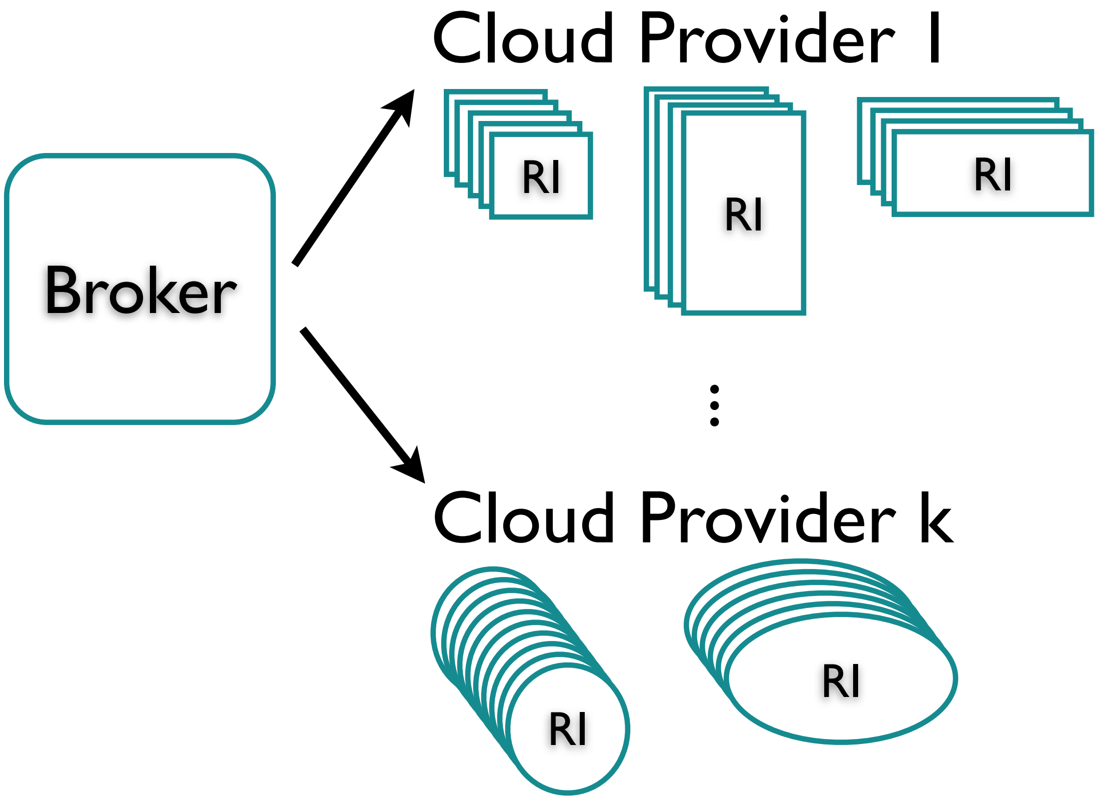

The VMMP Problem

- Allocate all customers VMs requests into the available RIs

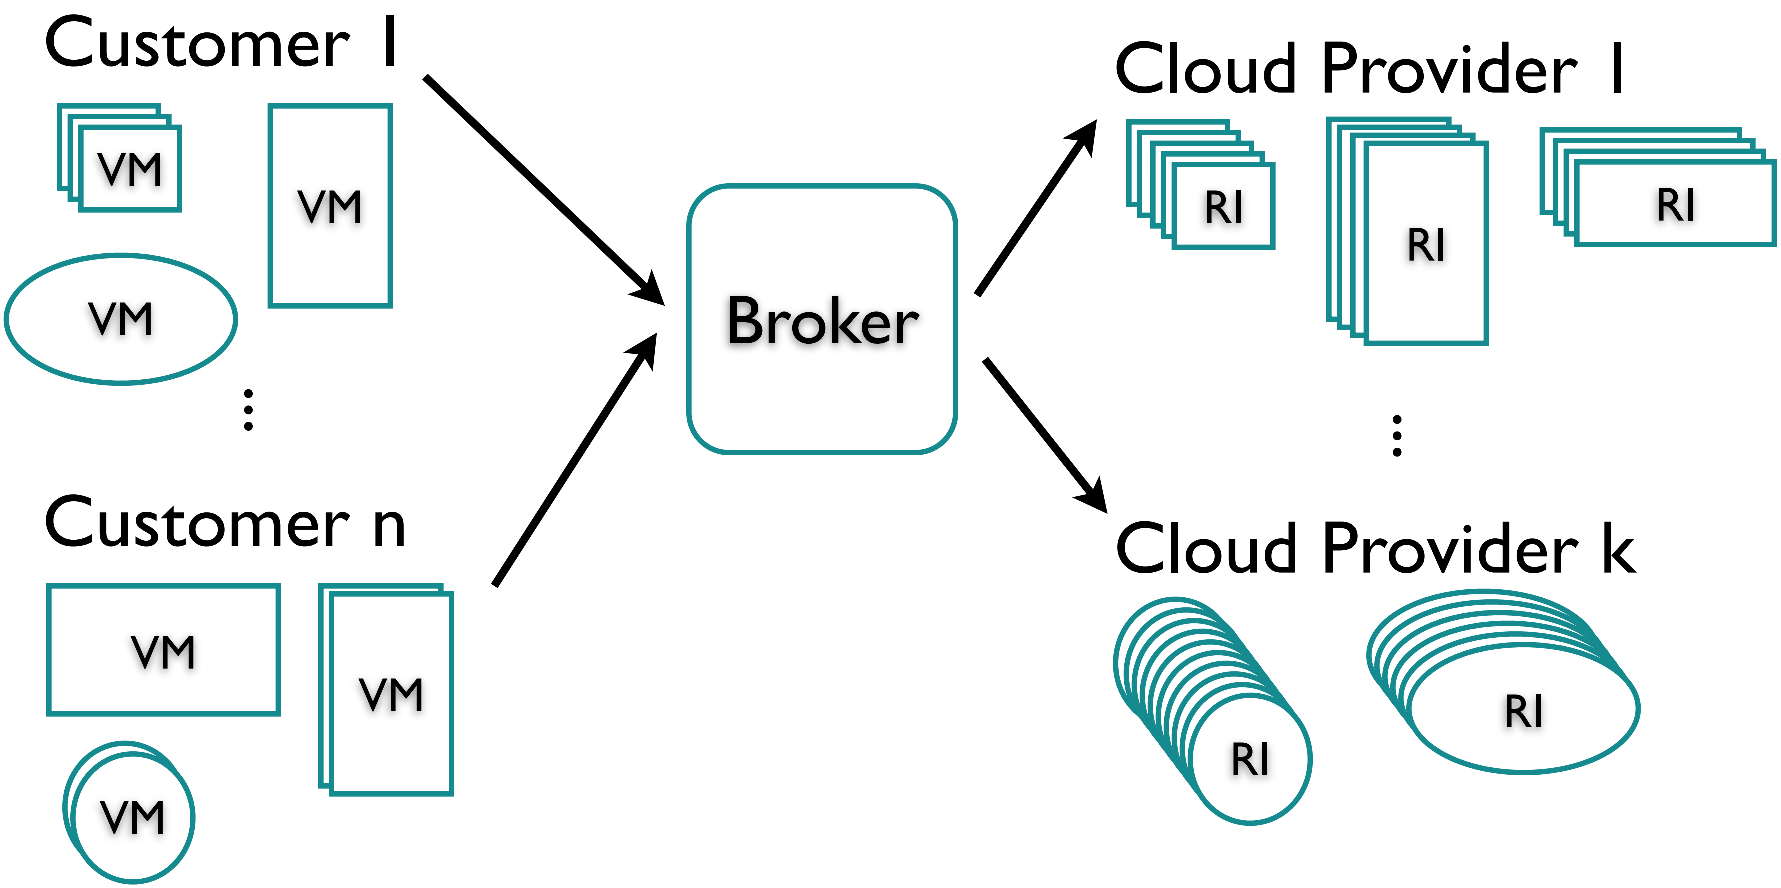

The VMMP Problem

$$\max \quad \sum_{j=1}^{j=m} \left(\sum_{i:f(v_i)=b_j} (p(BF(v_i)) - C(b_j)) \times T(v_i) \right) +$$

Profit

$$\sum_{h:ST(v_h)>D(v_h)} (p(BF(v_h)) - COD(BF(v_h))) \times T(v_h)$$

Cost of deadline violations handling

$$\text{subject to} \quad M(v_i) \leq M(b_j) \ P(v_i) \leq P(b_j)$$

$$S(v_i) \leq S(b_j), \ nc(v_i) \leq nc(b_j)$$

Outline

1. Motivations and Contributions
2. The Virtual Machine Mapping Problem
3. **List Scheduling Algorithms**
4. Experiments
5. Conclusions and Future Work

List Scheduling Algorithms

Res.

- Best fit resource (BFR): assigns every VM to its most suitable RI

Time

- Earliest finish time (EFT): VMs that finish earlier first
- Lower gap first (LGF): VMs with tightest deadlines first
- Shortest task first (STF):VMs with shortest execution time first
- Earliest deadline first (EDF):VMs with earliest deadlines first (arrival time is not taken into account)

Cost

- Cheapest instance (CI): VMs are assigned to the cheapest RI that can execute it, in a FIFO
- Max profit (MaxP): VMs that provide higher profit first

Time & Cost

- Shortest request to cheapest instance (SRCI): Shortest VMs are first assigned to the cheapest instance that can execute them

Outline

1. Motivations and Contributions
2. The Virtual Machine Mapping Problem
3. List Scheduling Algorithms
4. **Experiments**
5. Conclusions and Future Work

Experiments

- 20 workload files (information on VMs request)
  - Batches of 50, 100, 200, and 400 VMs
- 20 scenario files (information on available RIs)
  - 10, 20, 30, and 50 RIs (AWS and Azure data)
- 8 different kinds of VMs
- Available online
- Pricing: 20% cheaper than the cloud provider price

400 problem instances

#	VM id	provider	memory	storage	proc.	nc	price	C	COD
1	m1.small	Amazon	1.7 GB	160 GB	1.0 GHz	1	0.048	0.027	0.06
2	m1.medium	Amazon	3.75 GB	410 GB	2.0 GHz	2	0.096	0.054	0.12
3	A2.medium	Azure	3.5 GB	489 GB	1.6 GHz	2	0.096	0.09	0.12
4	m1.large	Amazon	7.5 GB	850 GB	2.0 GHz	4	0.192	0.108	0.24
5	m2.xlarge	Amazon	17.1 GB	420 GB	3.25 GHz	2	0.192	0.136	0.24
6	A3.large	Azure	7.0 GB	999 GB	1.6 GHz	4	0.328	0.18	0.41
7	c1.xlarge	Amazon	7.0 GB	1690 GB	2.5 GHz	8	0.384	0.316	0.48
8	A4.xlarge	Azure	14.0 GB	2039 GB	1.6 GHz	8	0.464	0.36	0.58

Results: Broker Profit

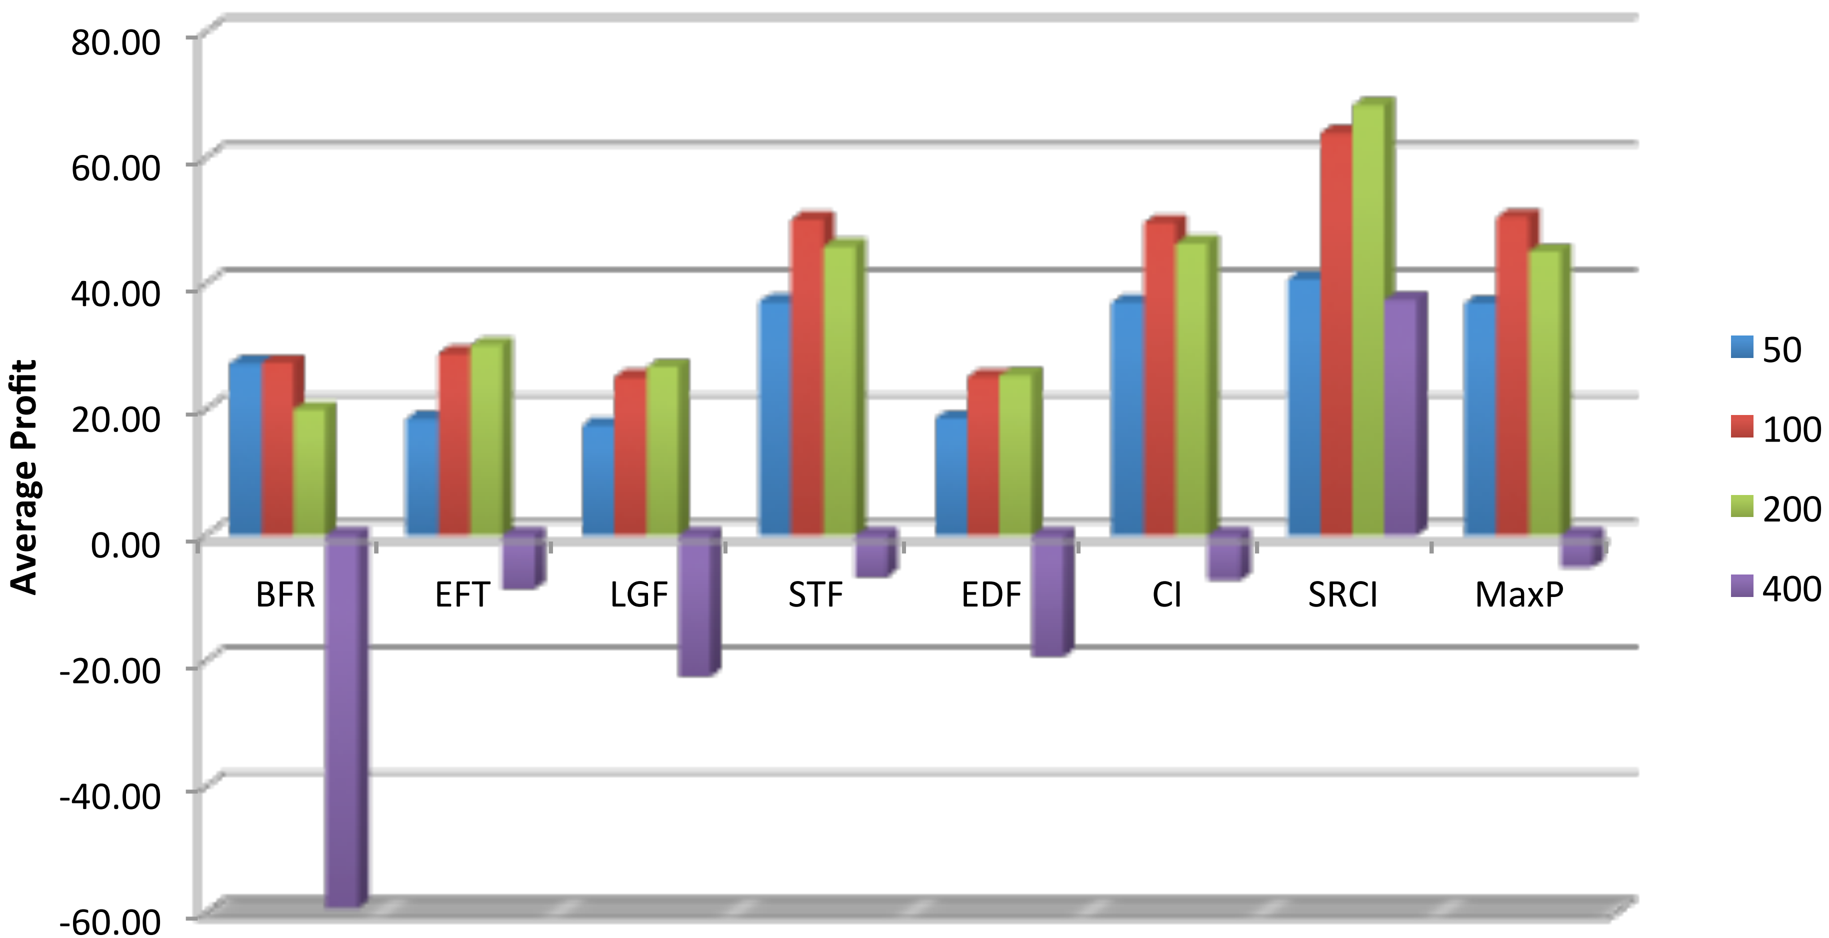

Results: Broker Profit

- Friedman test

		batch dimension (n)				overall
		50	**100**	**200**	**400**	
heuristic	**BFR**	5.35	6.34	6.50	7.59	6.44
	EFT	6.61	5.81	5.25	4.26	5.48
	LGF	6.76	6.44	5.82	5.47	6.12
	STF	2.99	3.36	3.77	4.20	3.58
	EDF	6.41	6.19	6.03	5.33	5.99
	CI	3.17	3.56	3.61	4.21	3.64
	SRCI	1.71	1.16	1.10	1.41	1.34
	MaxP	3.02	3.14	3.92	3.54	3.41

Results: Broker Profit

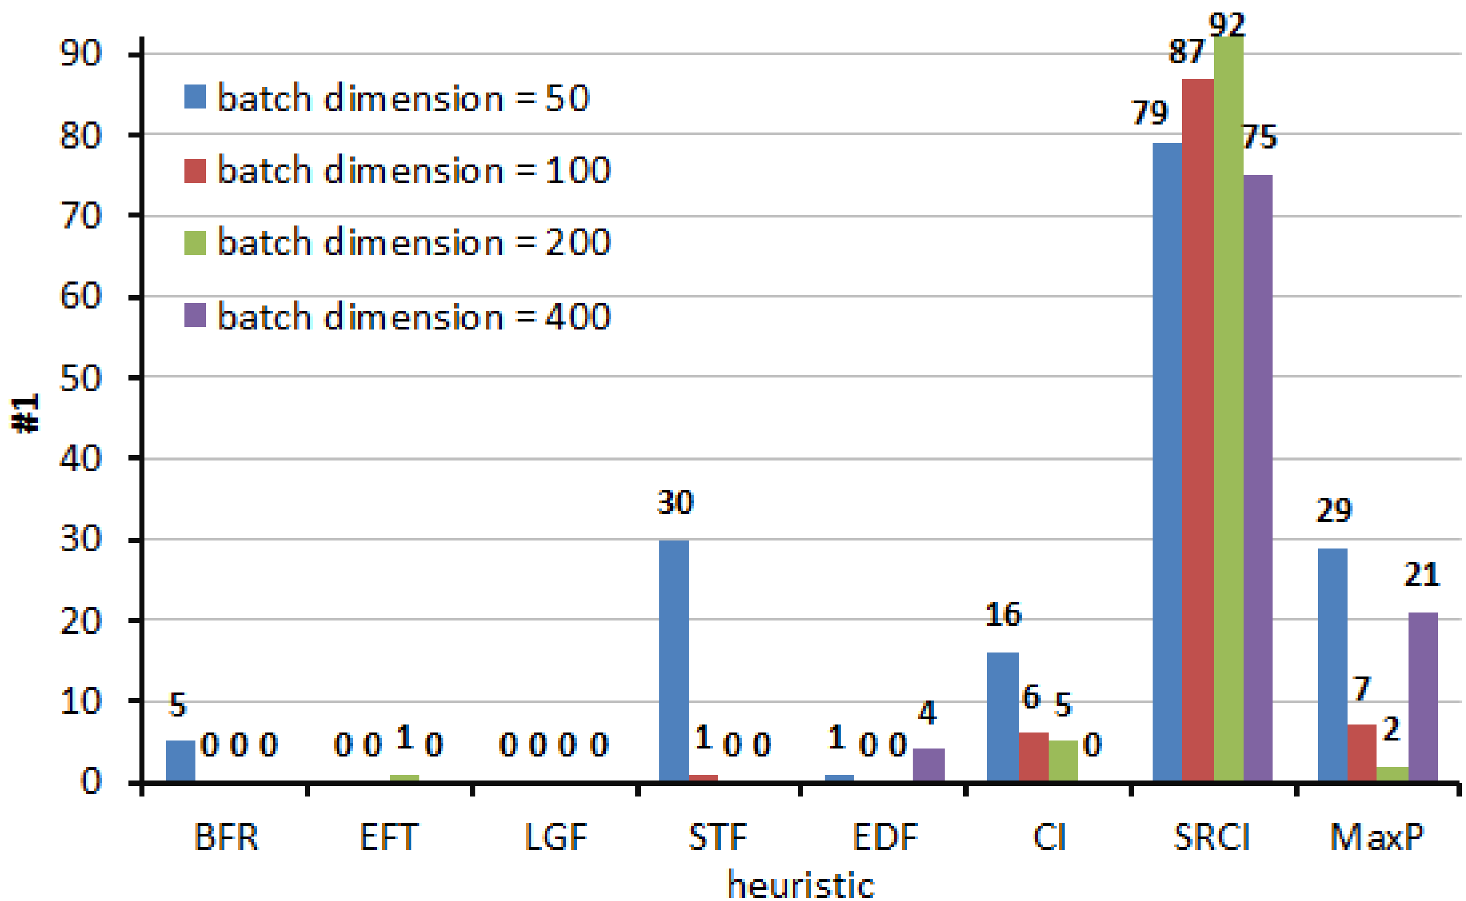

Results: Violated Requests

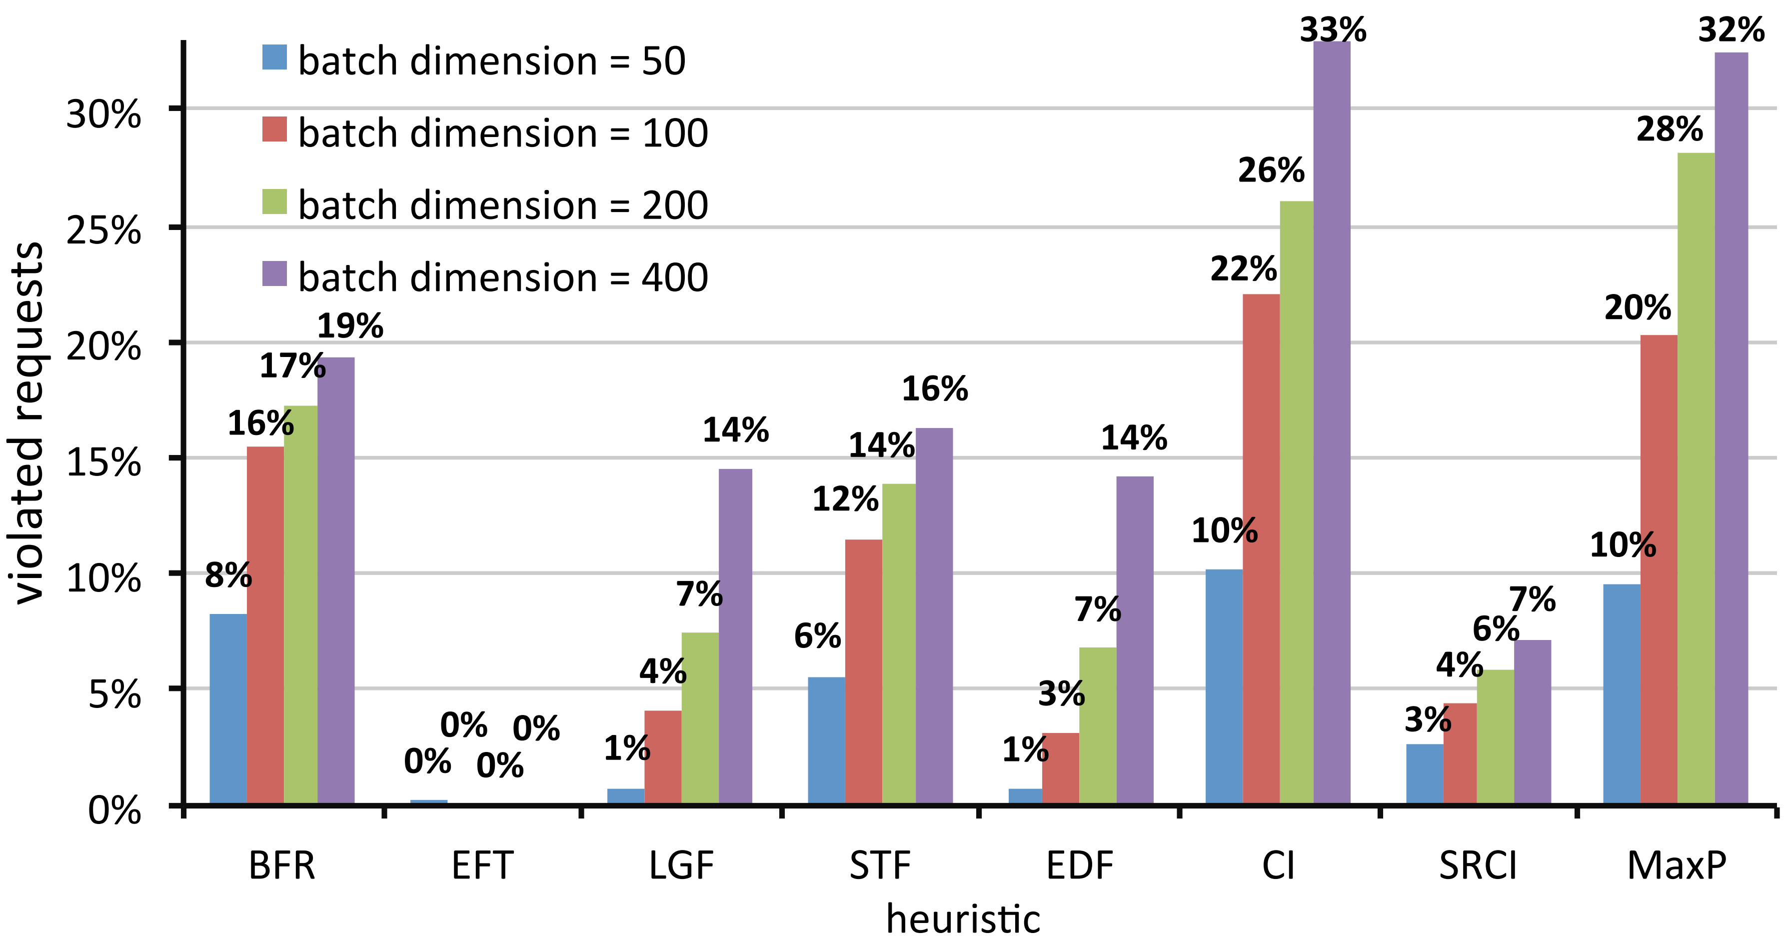

Outline

1. Motivations and Contributions
2. The Virtual Machine Mapping Problem
3. List Scheduling Algorithms
4. Experiments
5. **Conclusions and Future Work**

Conclusions & Future Work

- Novel cloud brokering model
  - Reserved nodes are sublet in an on-demand basis
  - Profitable: large difference between on-demand and reserved VMs cost
- VMMP: novel problem to plan the resources utilization
  - VMs requests must be mapped into RIs, maximizing profit
  - Constraint violations imply profit reduction
- Eight heuristics to solve the problem
  - SRCI outperformed the others
- Future work
  - Use of metaheuristics
  - Consider nesting technology
  - Development techniques for accurately managing the number and kind of RIs

Thank you.