School of
Information Systems

SMU
SINGAPORE MANAGEMENT UNIVERSITY

Probabilistic Inference Based Message-Passing For Resource Constrained DCOPs

Supriyo Ghosh, Akshat Kumar and Pradeep Varakantham

School of Information Systems
Singapore Management University

Decision Making in Multiagent Systems

- Examples
  - Sensor networks, Distributed meeting scheduling, multi-robot coordination

Goal: How do a set of agents decide the best alternative using only local coordination?

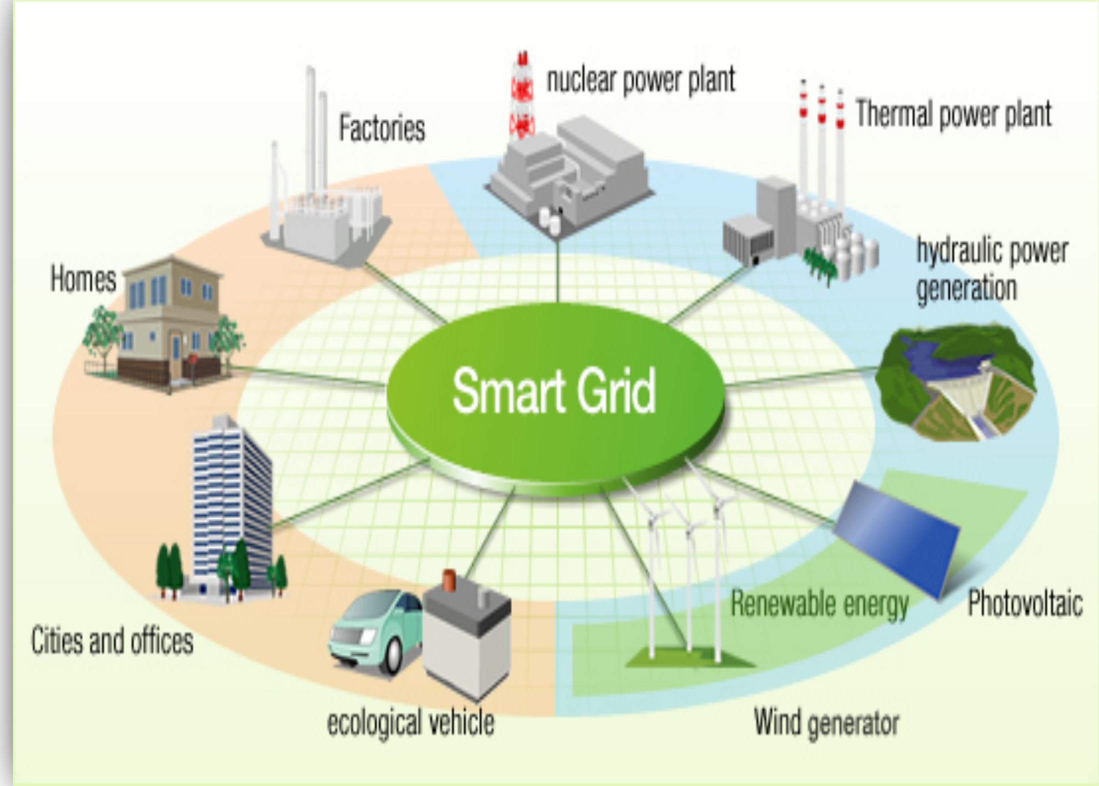

- Challenges
  - No central control/knowledge
  - Communication overhead
  - ***Shared resources***

Motivating Example

- Meeting scheduling with budget constraint
  - Two branch each has a limited travel budget
- Two meetings have to be scheduled, each having two options
- **Goal:** Schedule meeting such that total utility is maximised

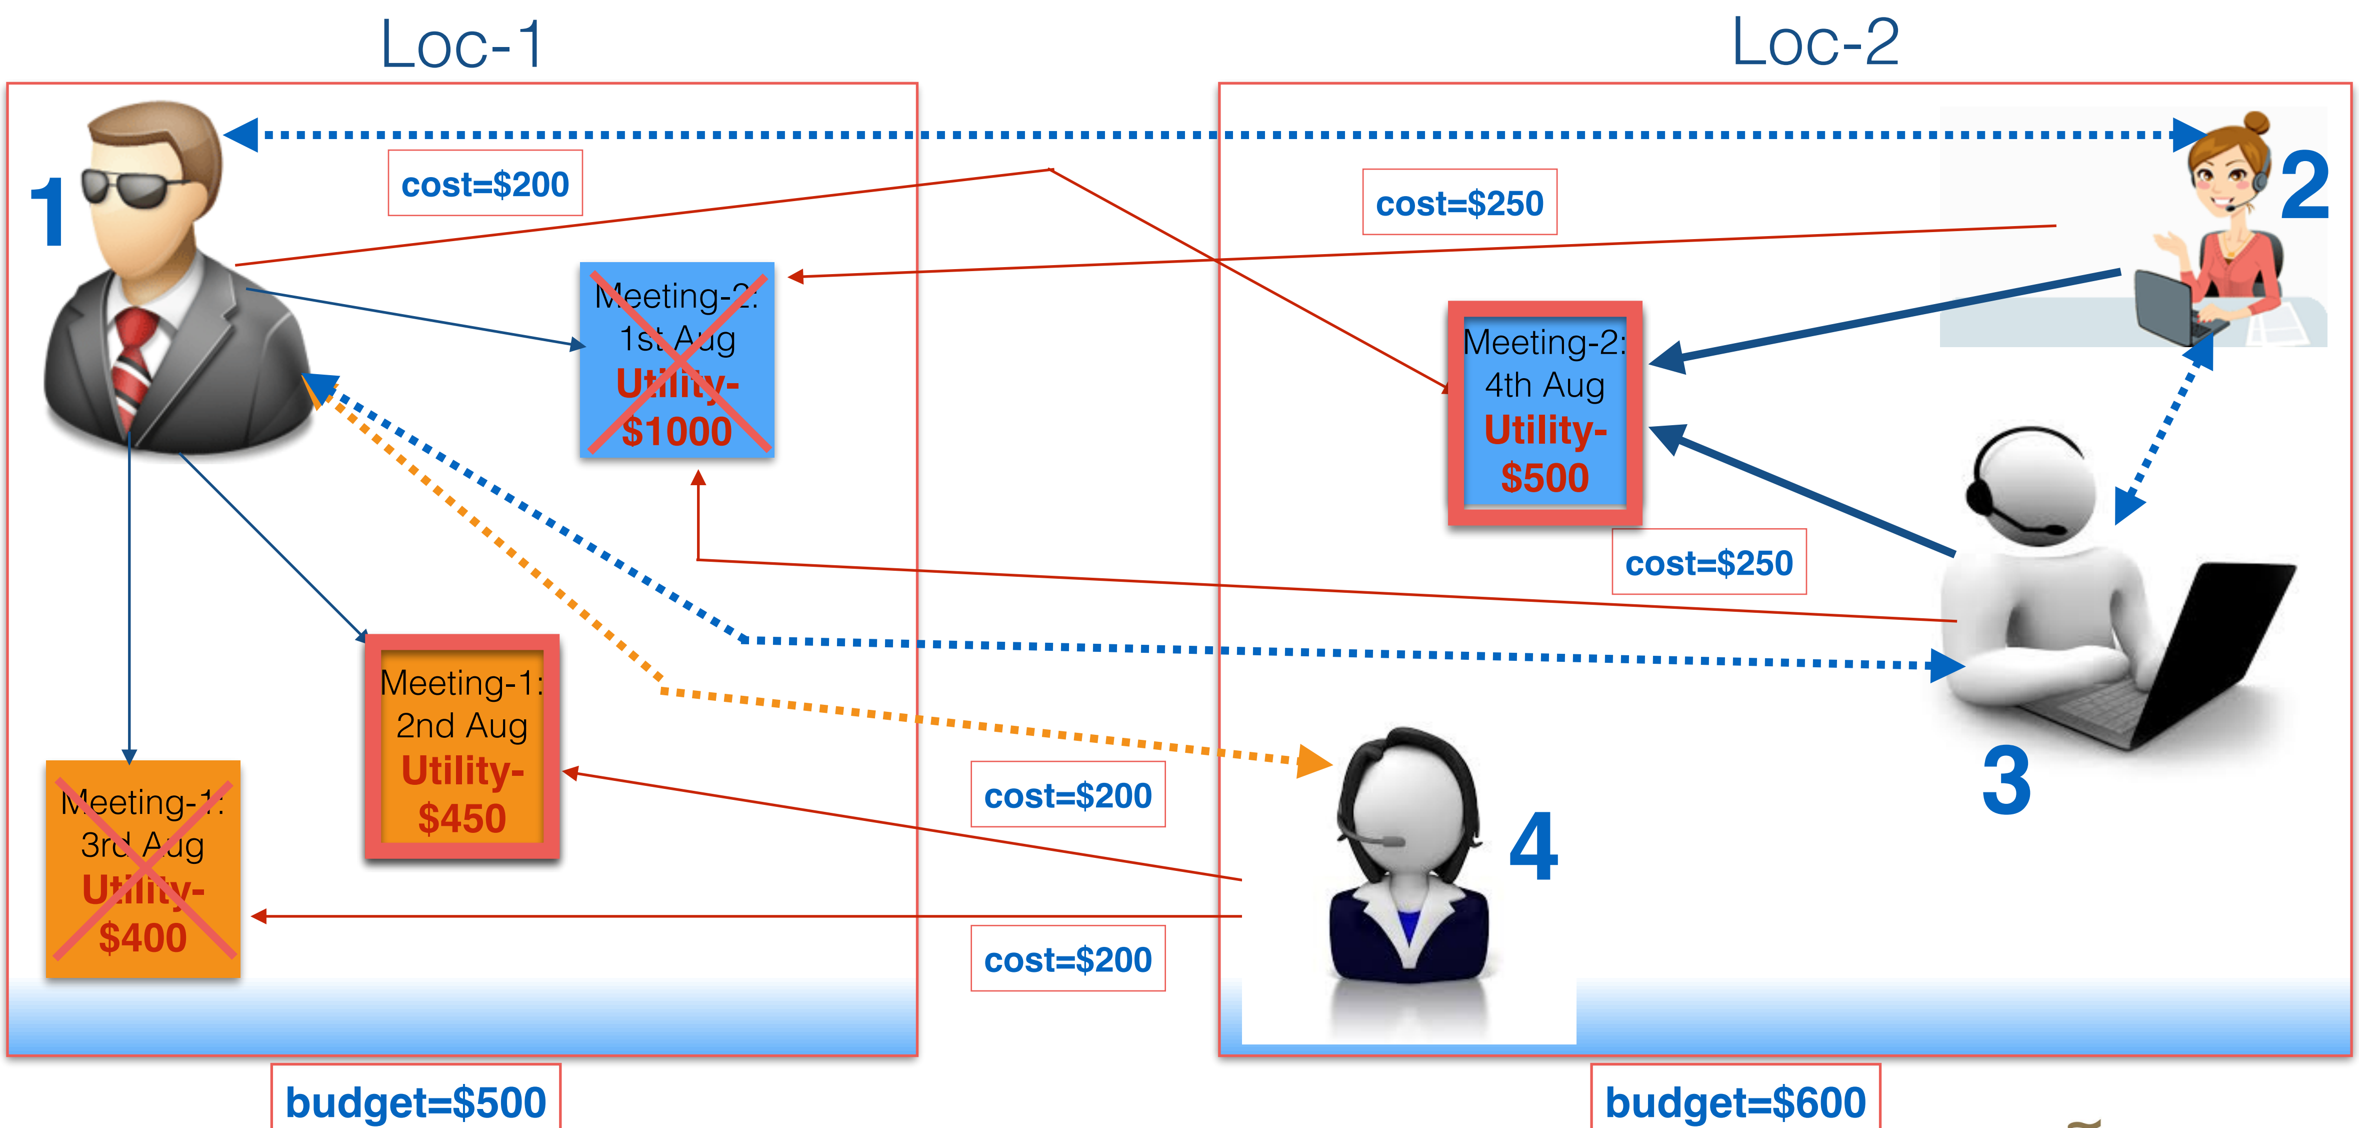

School of Information Systems

Resource Constrained DCOPs

- variables: $X = (x_1, ..., x_n)$, finite domains
- Constraints: $\theta_{ij}(x_i, x_j) : x_i \times x_j \rightarrow \Re$
- **Goal :** Find joint assignment s.t. $\max_x \sum_{(i,j)} \theta_{ij}(x_i, x_j)$

- A set $R = \{r_1, ..., r_m\}$ of shared resources
- Utilisation: $u_i(r, x_i) : R \times x_i \rightarrow \Re^+$
- Resource constraint:

$$\forall r \in R : \sum_{i \in V} u_i(r, x_i) \leq C(r)$$

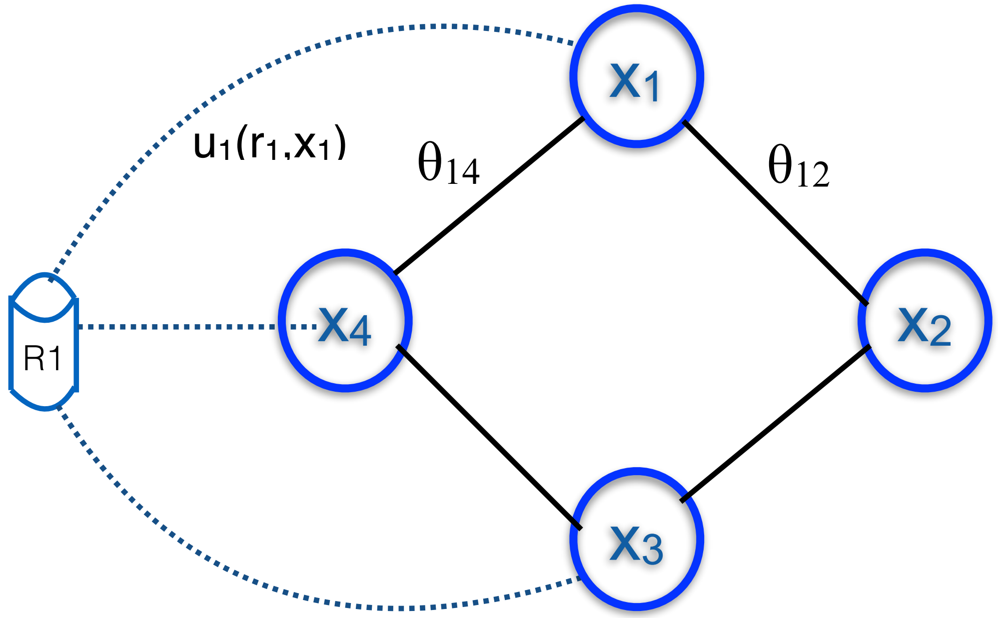

Related Work and Challenges

- Exact Algorithm:
  - [Bowring *et al.*, AAMAS-06]- Extends ADOPT with multiply constraints
  - [Matsui *et al.*, AAAI-08]- Introduces resources as a virtual variable.
  - Limited Scalability
- Approximate Algorithm:
  - No dedicated approximate solver for RC-DCOP
  - For tight resource constraints, approximate DCOP solvers fail to find even one feasible solution
  - Empirically true for max-sum
- Probabilistic Inference
  - [Kumar and Zilberstein, NIPS-10]- MAP estimation in graphical model
  - Extend the above to handle resource constraints

Our Contribution

Formulate QP for continuous relaxation of RC-DCOP

↓

QP formulation $\equiv$ Likelihood Maximization in a Bayes Net Mixture model

↓

Maximise the likelihood in the mixture of BN using Expectation Maximisation

↓

Develop Message-Passing based on EM

Continuous Relaxation of RC-DCOP

$$\max_{p=\{p_1,\dots,p_n\}} \sum_{(i,j)\in E} \underbrace{\sum_{x_i,x_j} p_i(x_i)p_j(x_j)\theta_{ij}(x_i,x_j)}_{\text{expected contribution by edge (i,j)}}$$ ← Maximize overall utility

$$\text{s.t.} \quad \sum_{x_i} p_i(x_i) = 1 \qquad \forall i \in V$$ ← Normalisation Constraint

$$\sum_{i\in Nb(r)} \underbrace{\sum_{x_i} p_i(x_i)u_i(r,x_i)}_{\text{consumption by agent i}} \leq C(r) \qquad \forall r \in R$$ ← Resource Constraint

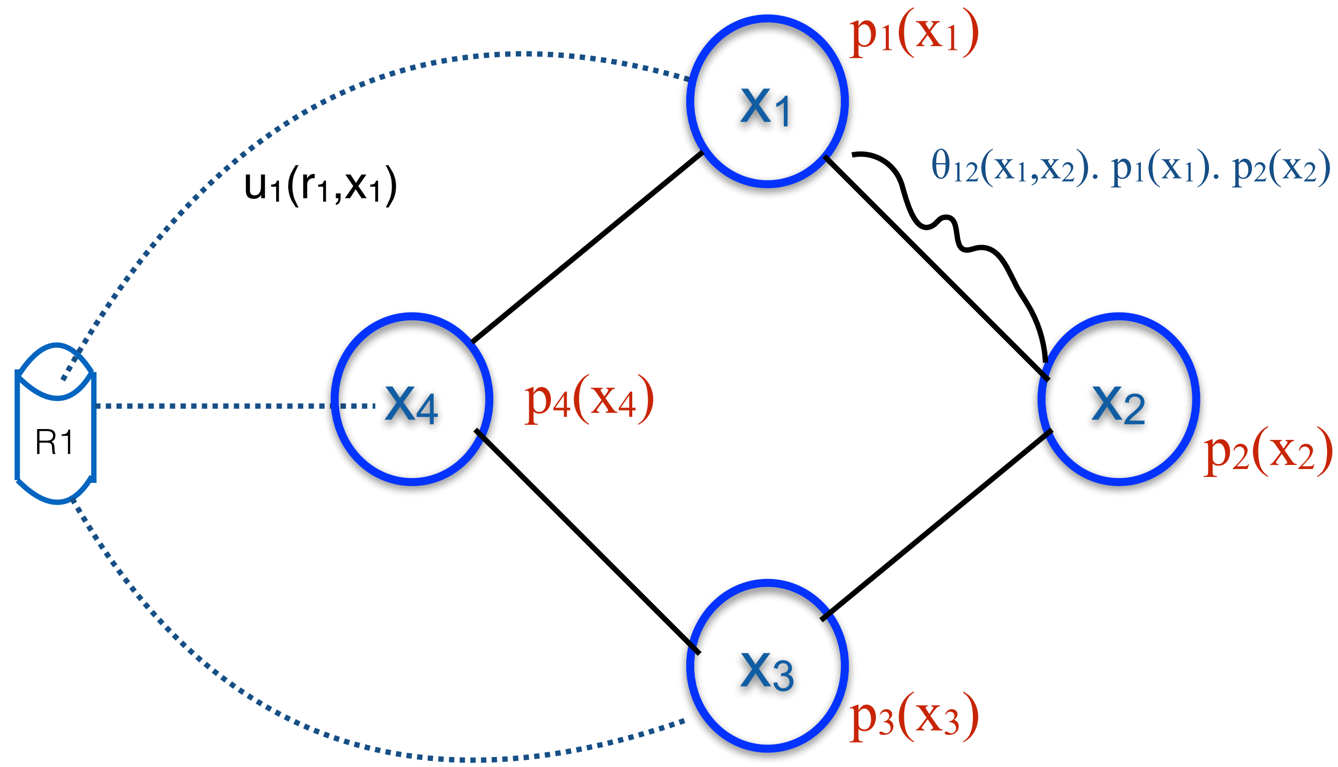

Our Contribution

Formulate QP for continuous relaxation of RC-DCOP

↓

QP formulation $\equiv$ Likelihood Maximization in a Bayes Net Mixture model

↓

Maximise the likelihood in the mixture of BN using Expectation Maximisation

↓

Develop Message-Passing based on EM

Conversion to mixture of Bayes Net

+ CPT for Bayes Net $l=1$:

$$P(\hat{\theta}=1|x_1,x_2)=\frac{\theta_{12}(x_1,x_2)-\theta_{min}}{\theta_{max}-\theta_{min}}$$

Theorem: Maximising the likelihood $L^p = P(\hat{\theta}=1;p)$ of observing the reward variable subject to resource constraint is equivalent to solving the QP relaxation of RC-DCOP

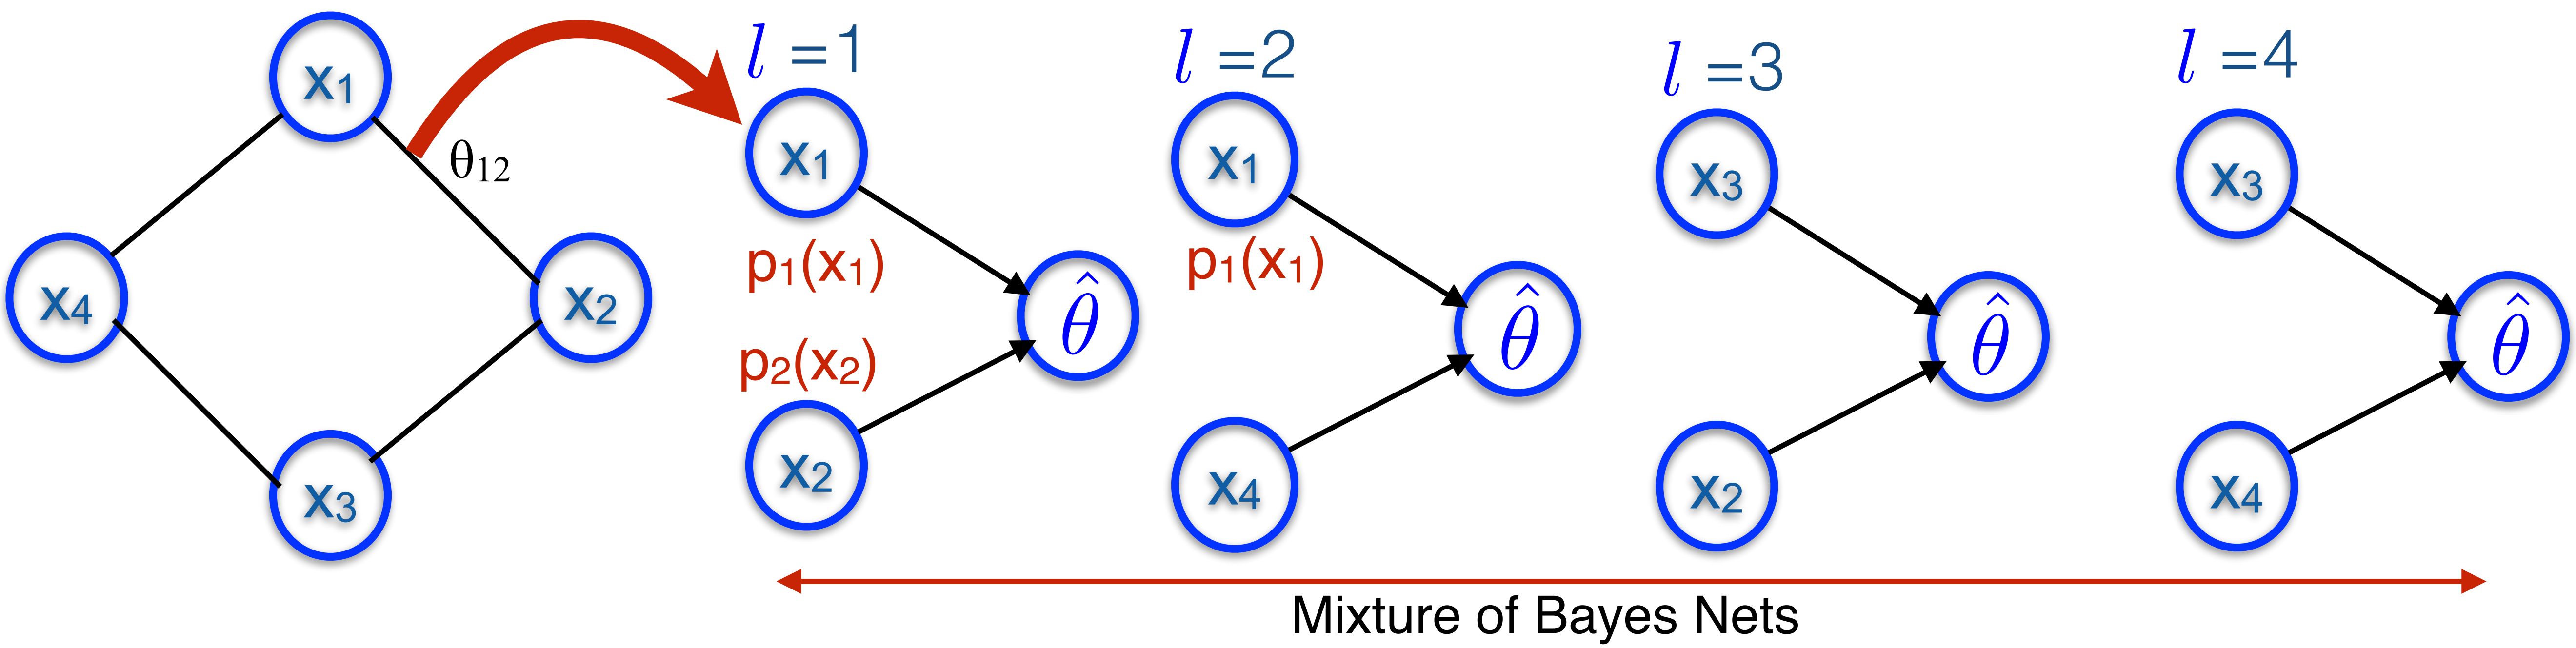

School of Information Systems

[Kumar and Zilberstein, 2010] Akshat Kumar and Shlomo Zilberstein. MAP estimation for graphical models by likelihood maxi- mization. In *NIPS*, pages 1180–1188, 2010.

Our Contribution

Formulate QP for continuous relaxation of RC-DCOP

↓

QP formulation $\equiv$ Likelihood Maximization in a Bayes Net Mixture model

↓

Maximise the likelihood in the mixture of BN using Expectation Maximisation

↓

Develop Message-Passing based on EM

EM for Mixture of BN

- Use EM [Dempster *et al.*,1977] to maximize the likelihood of $\hat{\theta} = 1$
- Hidden Variables: $x_{i,l}$
- Observed Variables: $\hat{\theta} = 1$
- Unknown Parameters: $p_i(x_i)$
- Parameters constraints includes the resource, normalisation constraints

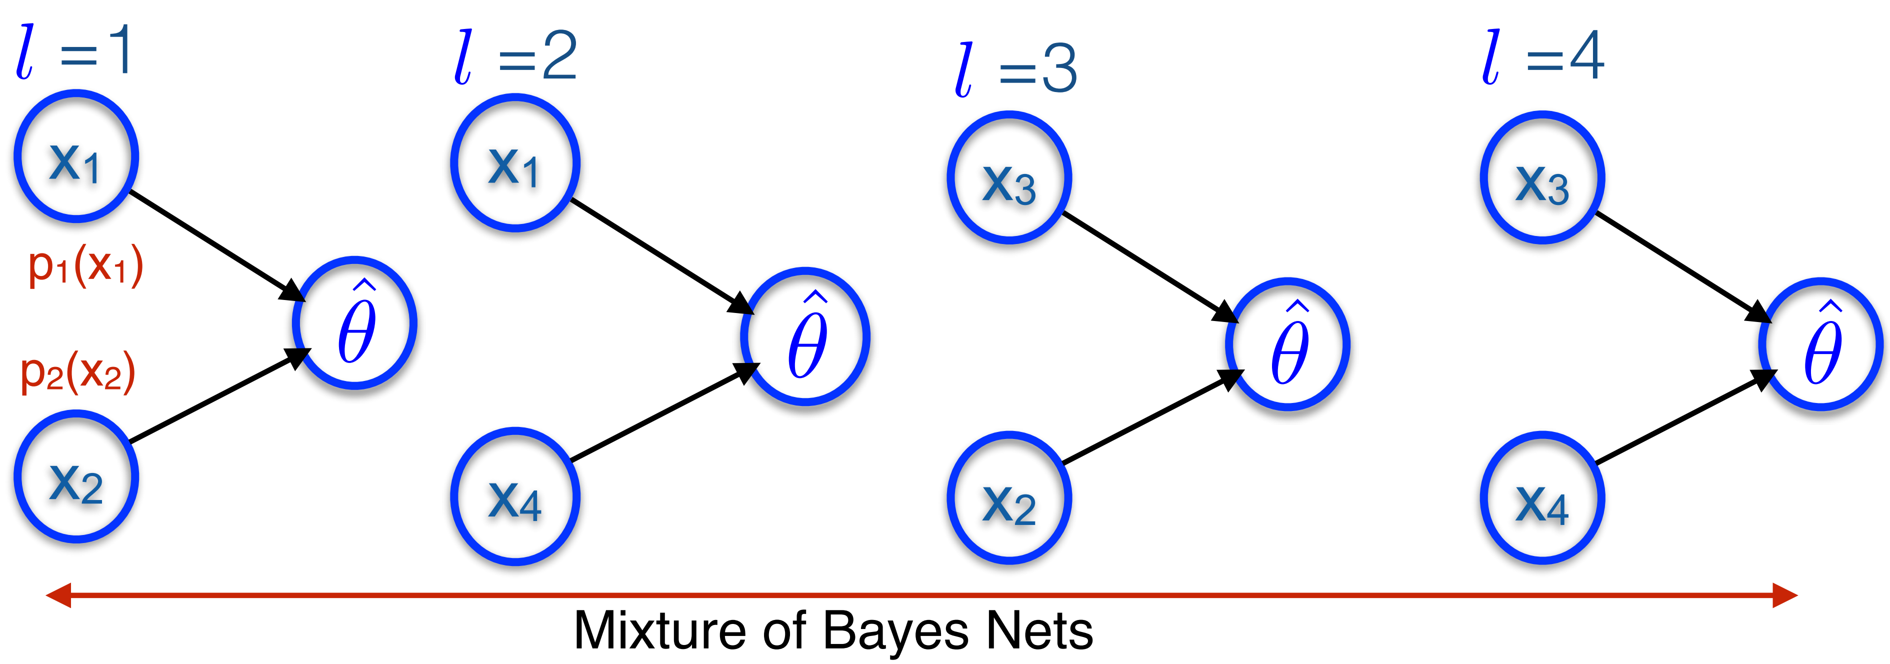

School of
Information Systems

SMU
SINGAPORE MANAGEMENT UNIVERSITY

Overview of EM for RC-DCOP

- EM is an iterative approach consisting of E & M-step
- E-step computes expectation over hidden variables
  - Implemented using message passing

- M step maximises expected log-likelihood
  - A convex optimization problem
  - No analytical solution
  - Solved by iteratively optimising the dual of this convex M-step problem (block coordinate descent strategy)
  - Implementable using message-passing

Our Contribution

Formulate QP for continuous relaxation of RC-DCOP

↓

QP formulation $\equiv$ Likelihood Maximization in a Bayes Net Mixture model

↓

Maximise the likelihood in the mixture of BN using Expectation Maximisation

↓

Develop Message-Passing based on EM

E-Step

Message:1(Agent→Agent): $\gamma_{i \to j}(x_j) \leftarrow \sum_{x_i} p_i(x_i)\hat{\theta}_{x_i x_j}$

Message:2 (Agent→Resource): $\delta_{i \to r}(x_i) \leftarrow p_i(x_i) \sum_{k \in Nb(i)} \gamma_{k \to i}(x_i)$

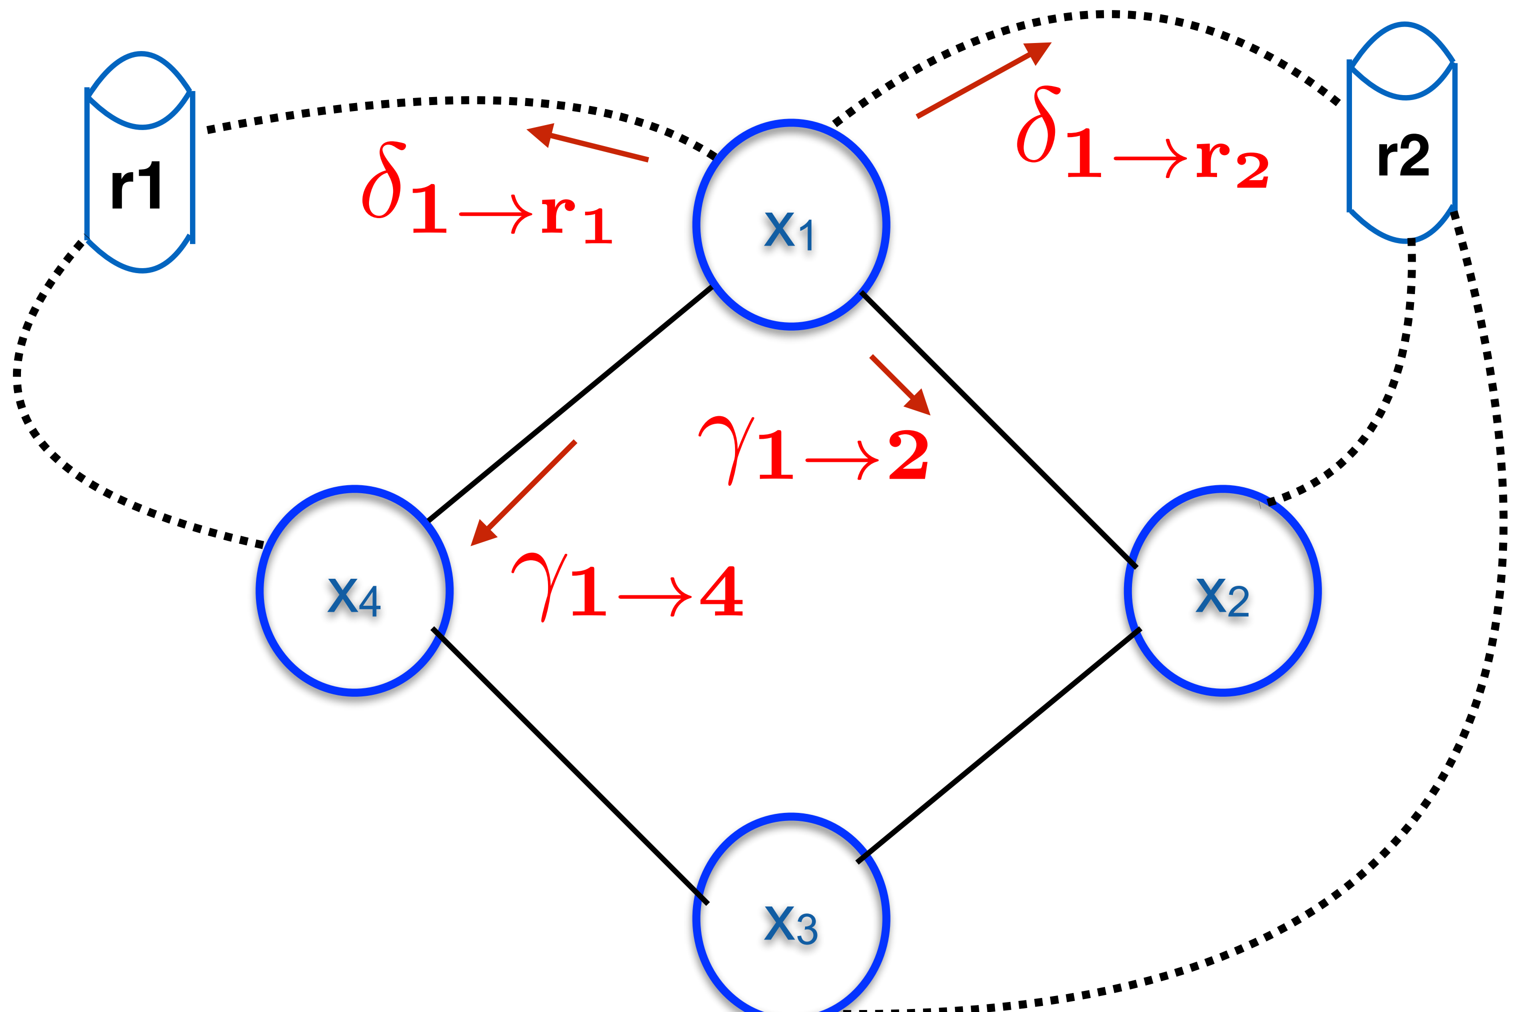

M-Step

Message:1(Resource→Agent): $\mu_{r\rightarrow i} \leftarrow max(0, \mu_r)$

Message:2(Agent→Resource): $\nu_{i\rightarrow \hat{r}}(x_i) \leftarrow \lambda_i + \sum_{r' \in Nr(i)\setminus \hat{r}} \mu_{r'\rightarrow i} u_i(r', x_i)$

Parameter Update: $p_i^*(x_i) \leftarrow \frac{p_i(x_i) \sum_{k\in Nb(i)} \gamma_{k\rightarrow i}(x_i)}{\lambda_i + \sum_{r\in Nr(i)} \mu_{r\rightarrow i} \cdot u_i(r, x_i)}$

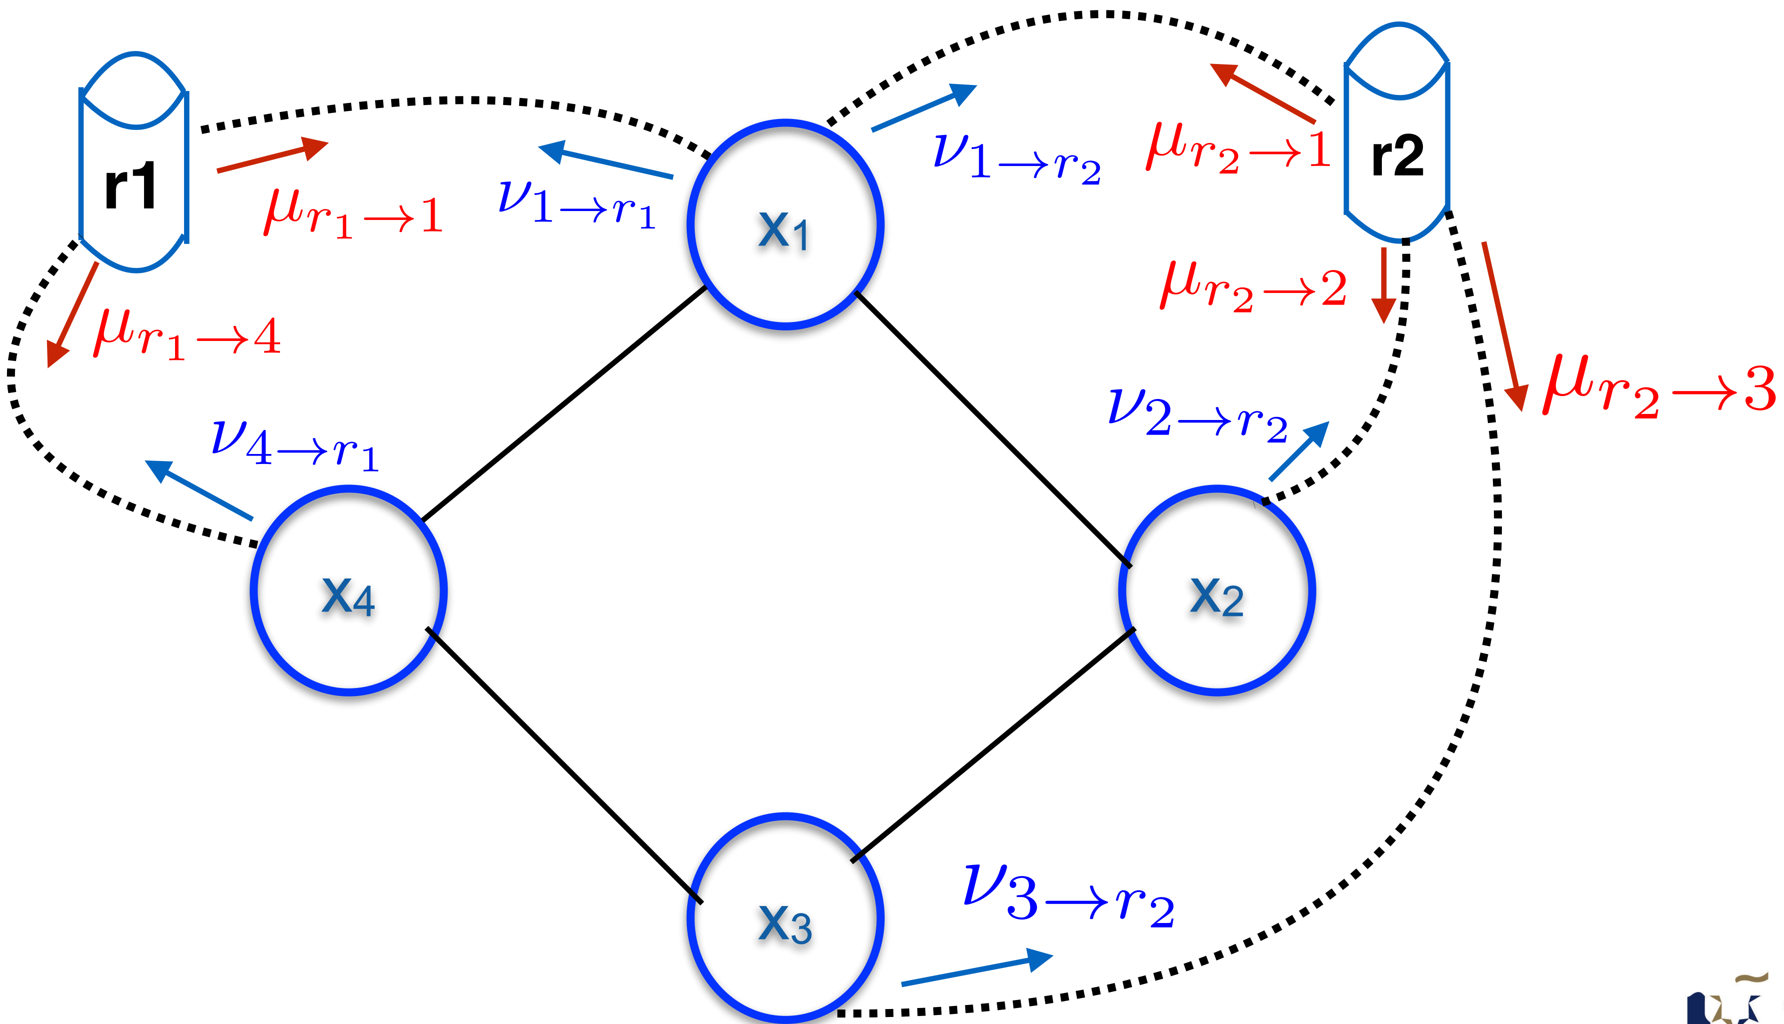

School of Information Systems

SMU SINGAPORE MANAGEMENT UNIVERSITY

Experimental Results

- Experiment on Two benchmarks
- Random Graphs(30 and 40 node):
  - Edge density is varied from 0.5 to 0.9
  - Random utility θ_{ij} between 1 to 10
  - Resource capacity is varied from 20%-60% of consumption
- Graph Colouring:
  - # of Nodes is varied from 20 to 50
  - Use same settings provided in [Farinelli *et al.*, AAMAS-08]
- Comparison Algorithm:
  - Toulbar2 [Allouche et al., INRA-10]
  - Max-Sum in Frodo implementation [L'eaut'e et al., IJCAI-09]

Experimental Results(2)

- Failure - No resource feasible solution found
- EM has deterministic outcome, single run
- Max-Sum run multiple times due to variable outcome

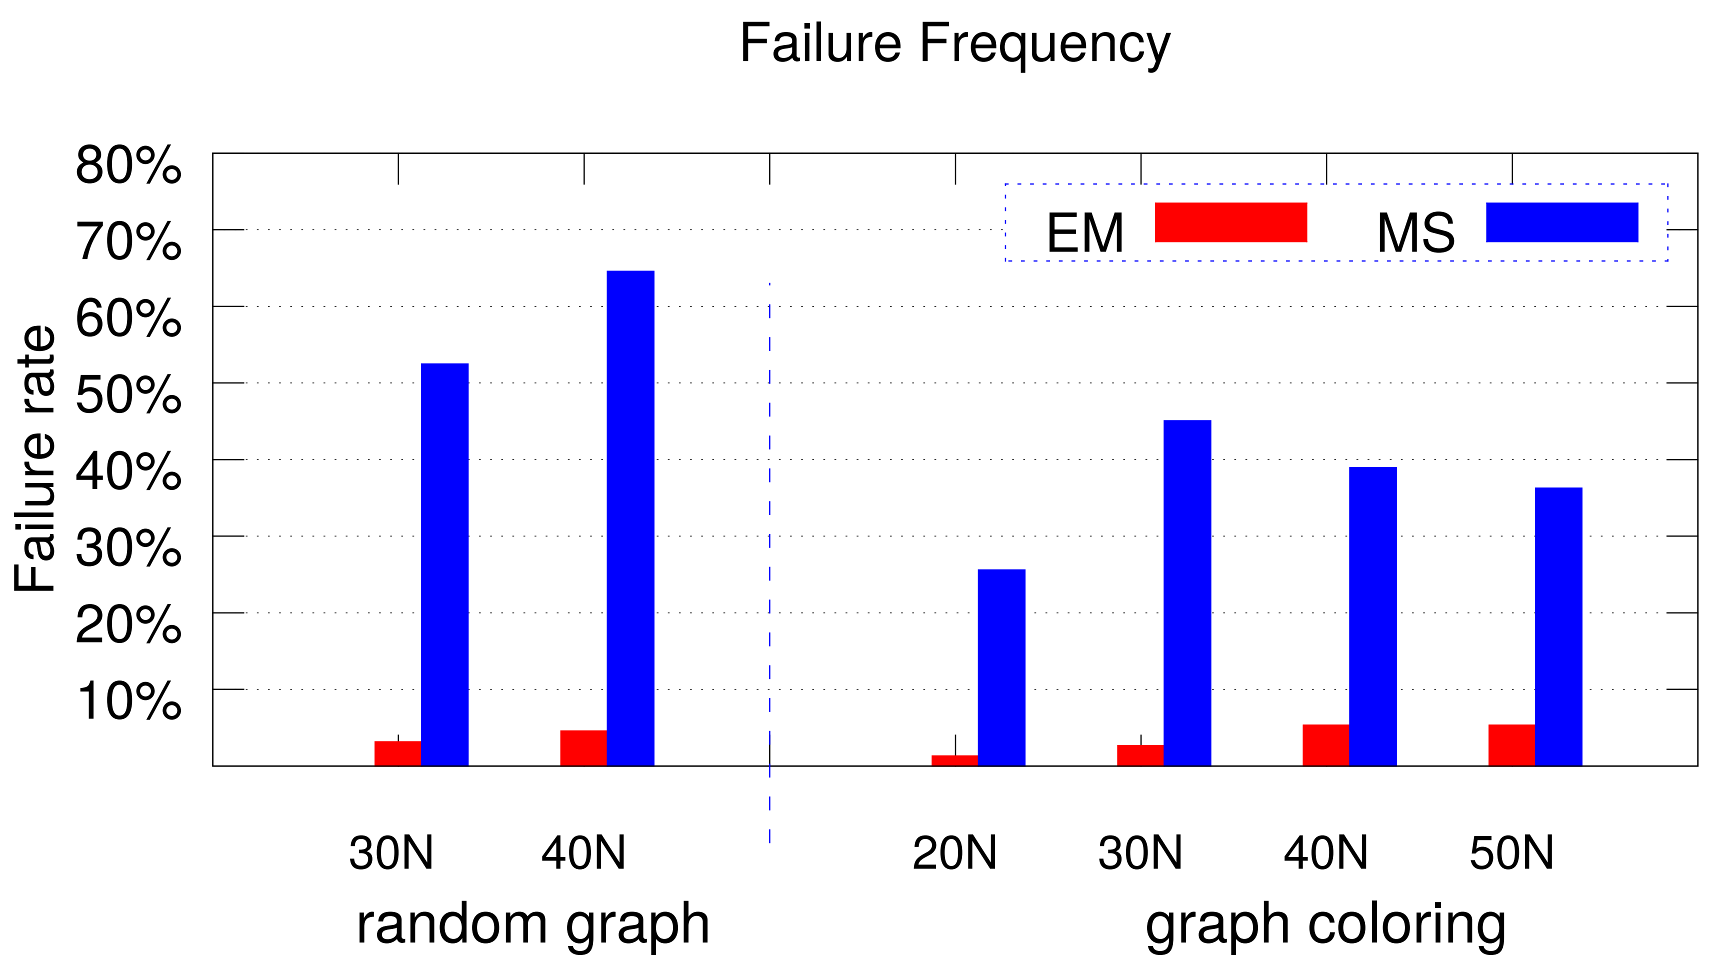

Experimental Results(3)

- EM outperforms toulbar2 as problem complexity increases.
- EM solution quality is noticeably better than Max-Sum.

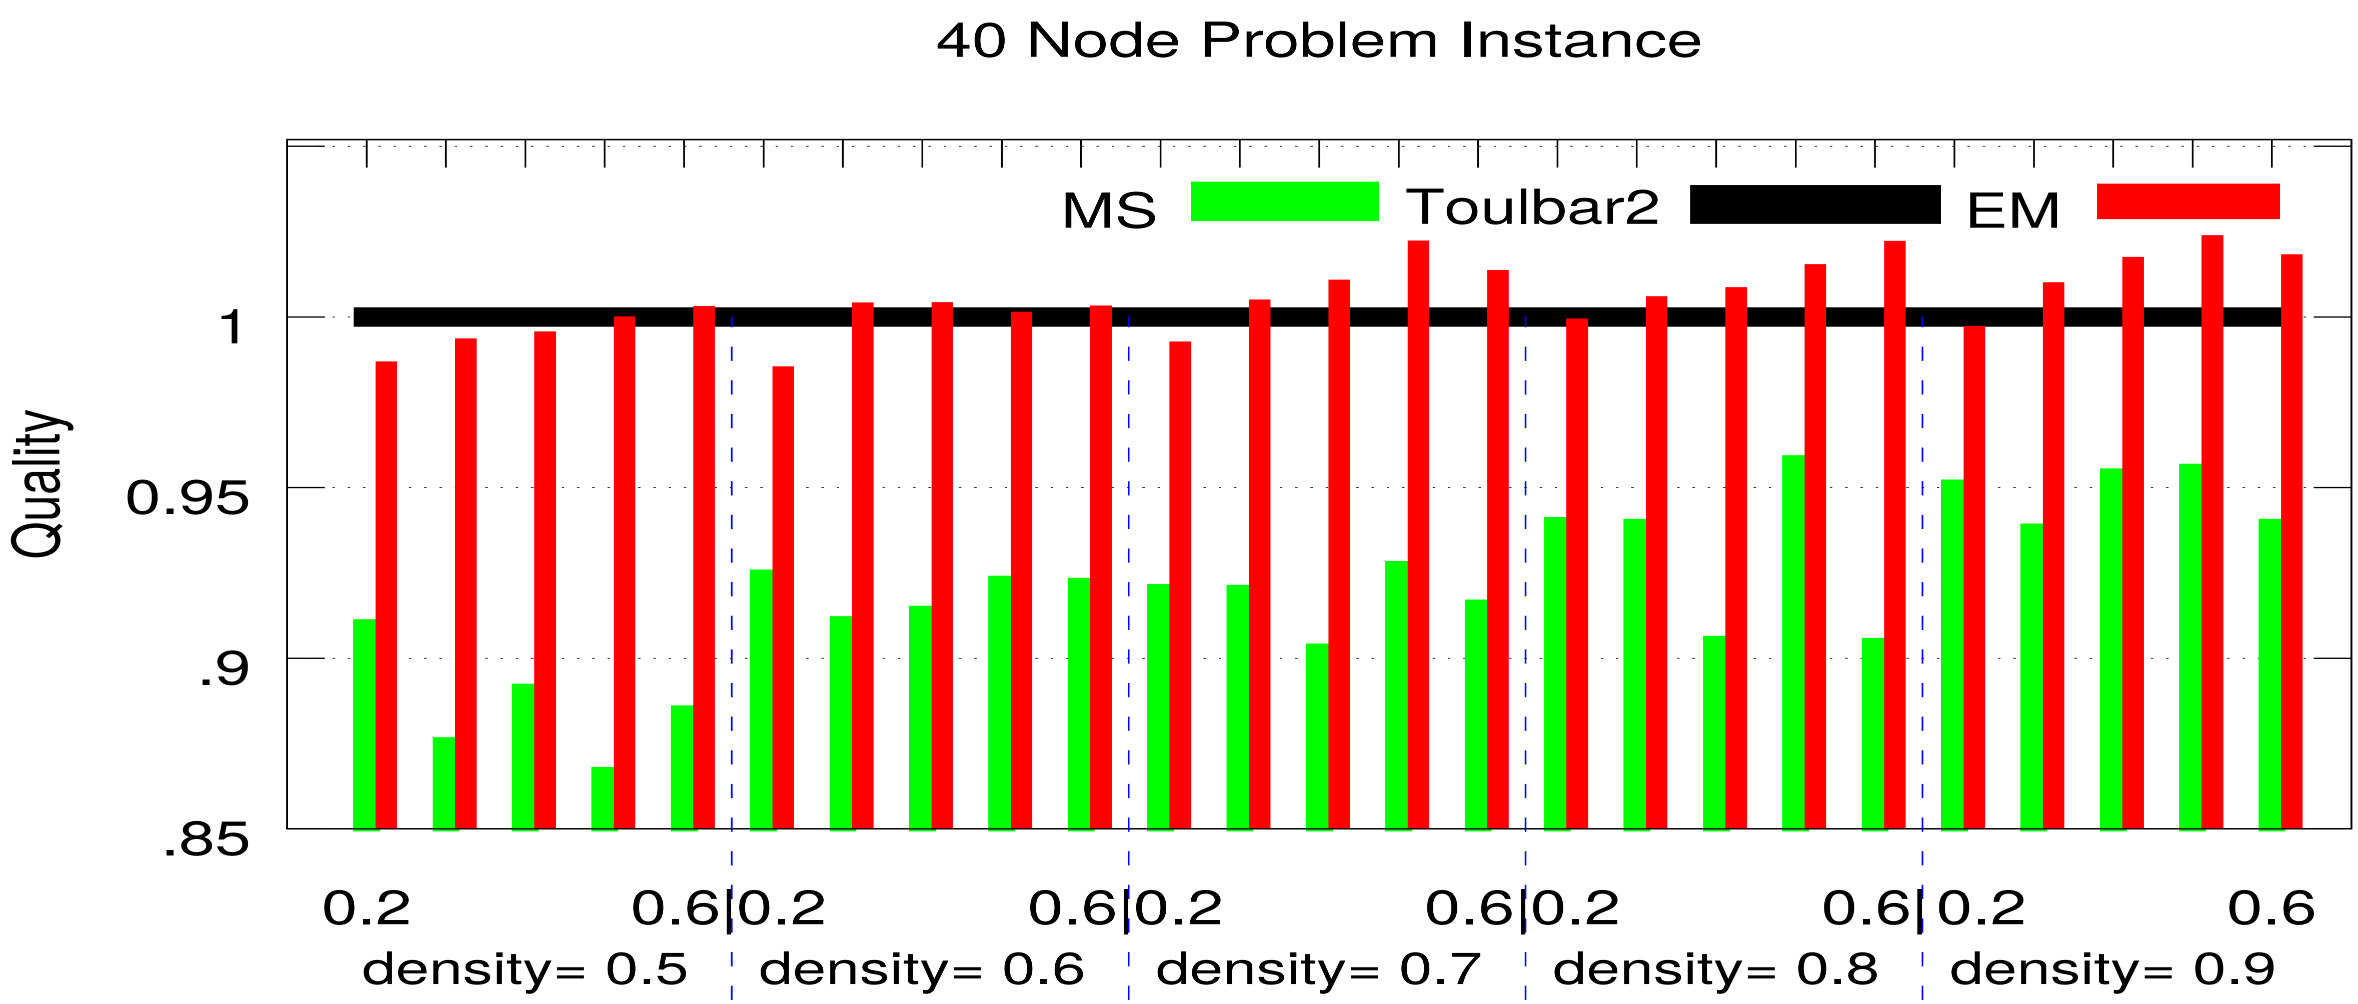

Experimental Results(4)

- EM provides near-optimal solution for graph colouring problems
  - Toulbar2 finds optimal solution
- Solution quality of EM is always better than Max-Sum for 30 node problems

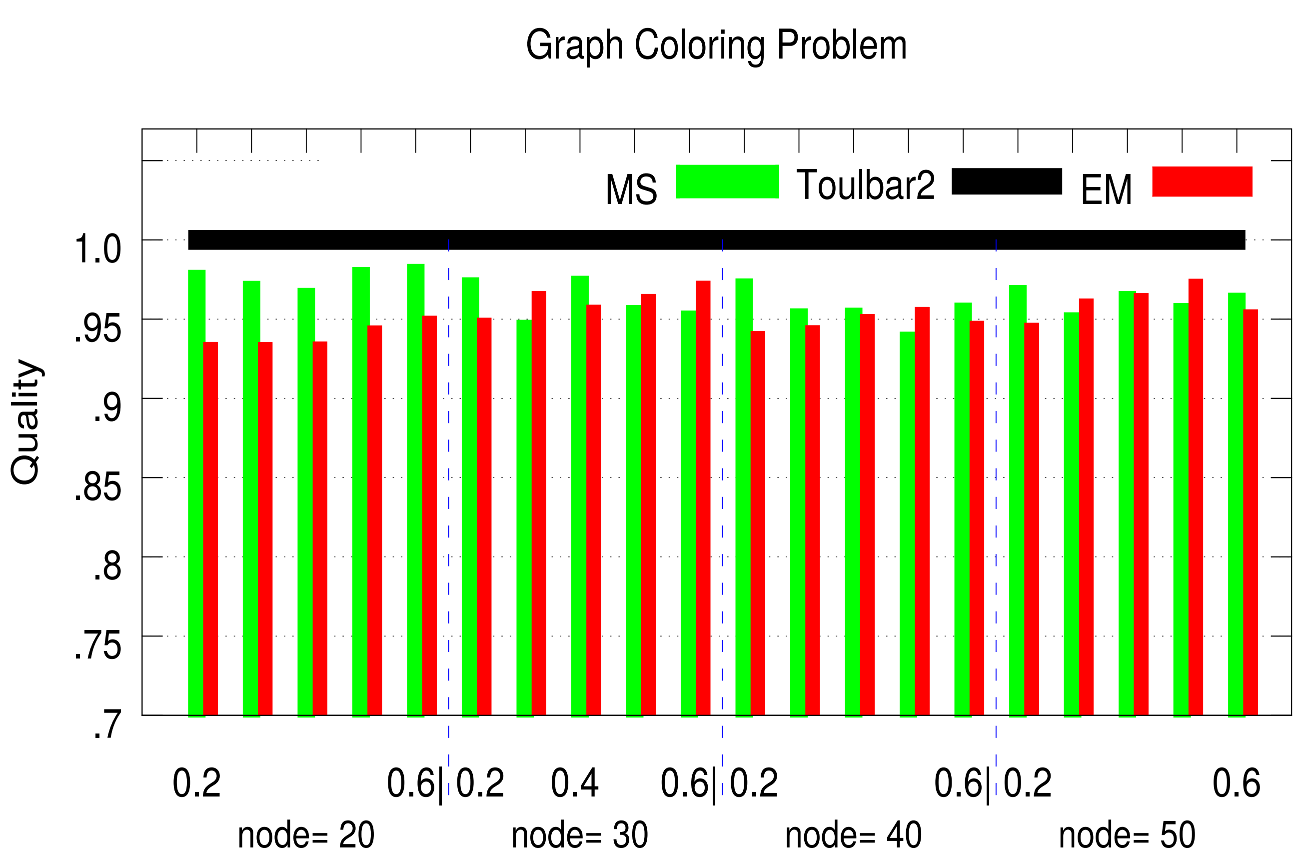

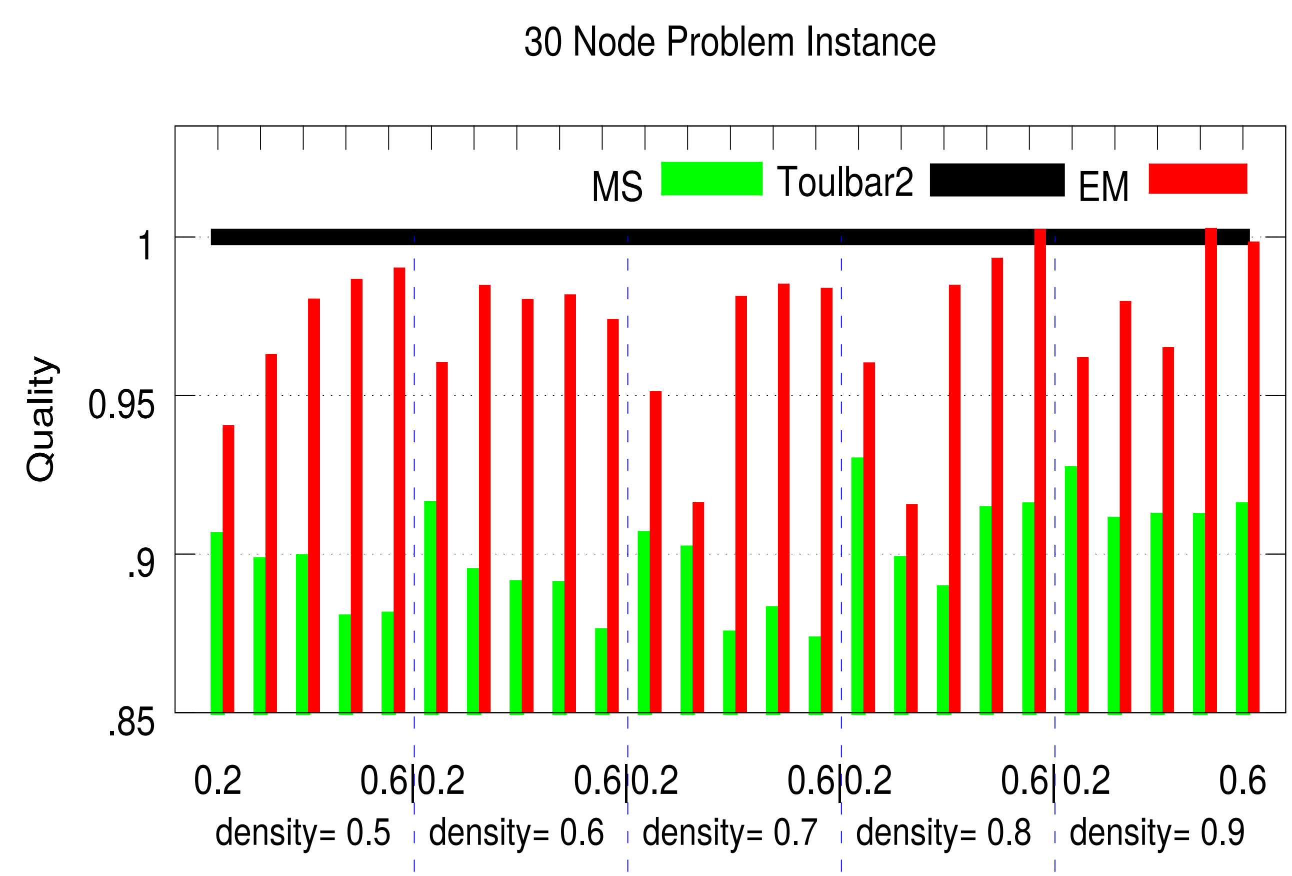

School of Information Systems

SMU SINGAPORE MANAGEMENT UNIVERSITY

Experimental Results(5)

- EM almost always achieves solution within 3 minutes.
- Although Max-Sum takes much lower time, its solution quality is worse

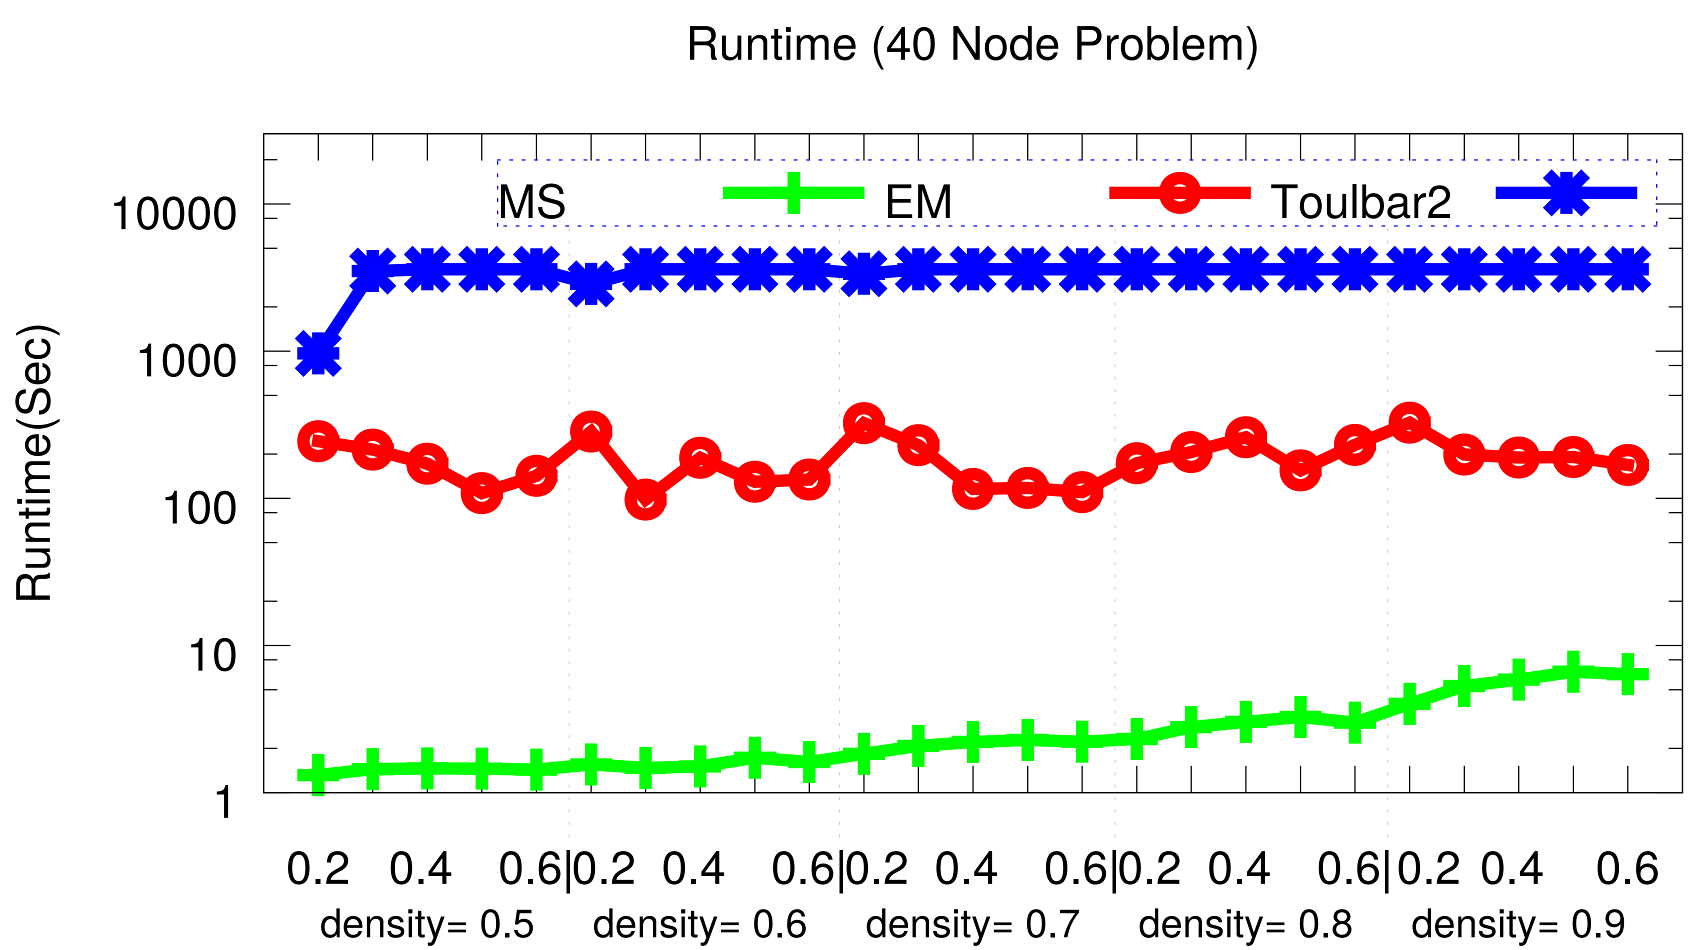

School of Information Systems

SMU Singapore Management University

Conclusion

- Present a promising class of approximate algorithm for RC-DCOP using probabilistic inference
- Solving RC-DCOP is equivalent to Likelihood Maximization
- Use machine learning technique for likelihood maximization
- Develop EM as message-passing algorithm for RC-DCOP
- EM has much lower failure rate than Max-Sum, provides good quality solution

Questions???

supriyog.2013@phdis.smu.edu.sg

Primal Extraction

- Extract integral solution from p^* using rounding technique from [Ravikumar & Laffety, 2006]
- If $p_i^*(x_i) \geq \delta$, set $p_i^{int}(x_i) = 1$ and $p_i^{int}(\hat{x}_i) = 0, \forall \hat{x}_i \setminus x_i$
- For each unlabelled node $i \in V$ find $\arg\max_{x_i} \sum_{j \in Nb(i)} \sum_{x_j} \theta_{ij}(x_i, x_j)$
- Label node with x_i that satisfy the resource constraints
- Iterate the process until convergence

*Pradeep Ravikumar and John Laf- ferty. Quadratic programming relaxations for metric labeling and Markov random field MAP estimation. In *ICML*, pages 737–744, 2006.

School of Information Systems

SMU SINGAPORE MANAGEMENT UNIVERSITY

Our Contributions

- Provide a new class of approximate algorithm for RC-DCOP
- Mapping of RC-DCOP to that of probabilistic inference in mixture of Bayes Nets
- Maximises the likelihood (LM) in Bayes Nets that is equivalent to solve the RC-DCOP
- Interpret LM as message-passing

Likelihood Maximisation in Mixture of BN

$$\underbrace{P(\hat{\theta}, x_{l_1}, x_{l_2} | l; p)}_{\text{Joint for Bayes net l}} = \underbrace{P(\hat{\theta} | x_{l_1}, x_{l_2}, l)}_{\hat{\theta}_{x_l}} p_{l_1}(x_{l_1}; p) p_{l_2}(x_{l_2}; p)$$

Likelihood for Bayes net l : $L_l^p = P(\hat{\theta} = 1 | l; p) = \sum_{x_l} P(\hat{\theta} = 1, x_{l_1}, x_{l_2} | l; p) = \sum_{x_l} \hat{\theta}_{x_l} p_{l_1}(x_{l_1}; p) p_{l_2}(x_{l_2}; p)$

Likelihood for complete mixture: $L^p = \sum_l P(l) L_l^p = \frac{1}{|E|} \sum_l \sum_{x_l} \hat{\theta}_{x_l} p_{l_1}(x_{l_1}; p) p_{l_2}(x_{l_2}; p)$

$$\underbrace{\sum_l \sum_{x_l} \theta_l(x_l) p_{l_1}(x_{l_1}; p) p_{l_2}(x_{l_2}; p)}_{\text{Objective of QP}} = |E|(\theta_{min} + (\theta_{max} - \theta_{min}) L^p)$$

→ QP formulation ≡ LM in BN mixture Model

School of Information Systems

SMU Singapore Management University

Optimisation problem in M-Step

$$\max_{p^*} \sum_{i \in V} \sum_{x_i} p_i(x_i) \log p_i^*(x_i) \underbrace{\sum_{j \in Nb(i)} \sum_{x_j} \hat{\theta}_{x_i x_j} p_j(x_j)}_{\text{Expected contribution for value } x_i}$$

$$\text{s.t.} \quad \sum_{x_i} p_i^*(x_i) = 1 \qquad \forall i \in V$$

$$\sum_{i \in Nb(r)} \sum_{x_i} p_i^*(x_i) u_i(r, x_i) \leq C(r) \qquad \forall r \in R$$

Expectation Maximization

Theorem 1: Maximising the following log-likelihood **Q(p, p*)** w.r.t. **p*** iteratively finds the optimal solution for DCOP. [kumar et al. 2011]

$$Q(p,p^*) \propto \sum_{i \in V} \sum_{x_i} p_i(x_i) \log p_i^*(x_i) \sum_{j \in Nb(i)} \sum_{x_j} \hat{\theta}_{x_i x_j} p_j(x_j)$$

$f_i(x_i)$

+ **Need to satisfy resource & normalisation constraints for RC-DCOP**

+ **Optimising the dual:**
  + Find the Lagrangian $L(p^*, \lambda, \mu)$ by using price variable λ, μ
  + Find the dual function: $q(\lambda, \mu) = \min_{p^*} L(p^*, \lambda, \mu)$
  + Parameter update: $p_i^*(x_i) = \dfrac{p_i(x_i) f_i(x_i)}{\lambda_i + \sum_r \mu_r \cdot u_i(x_i)}$

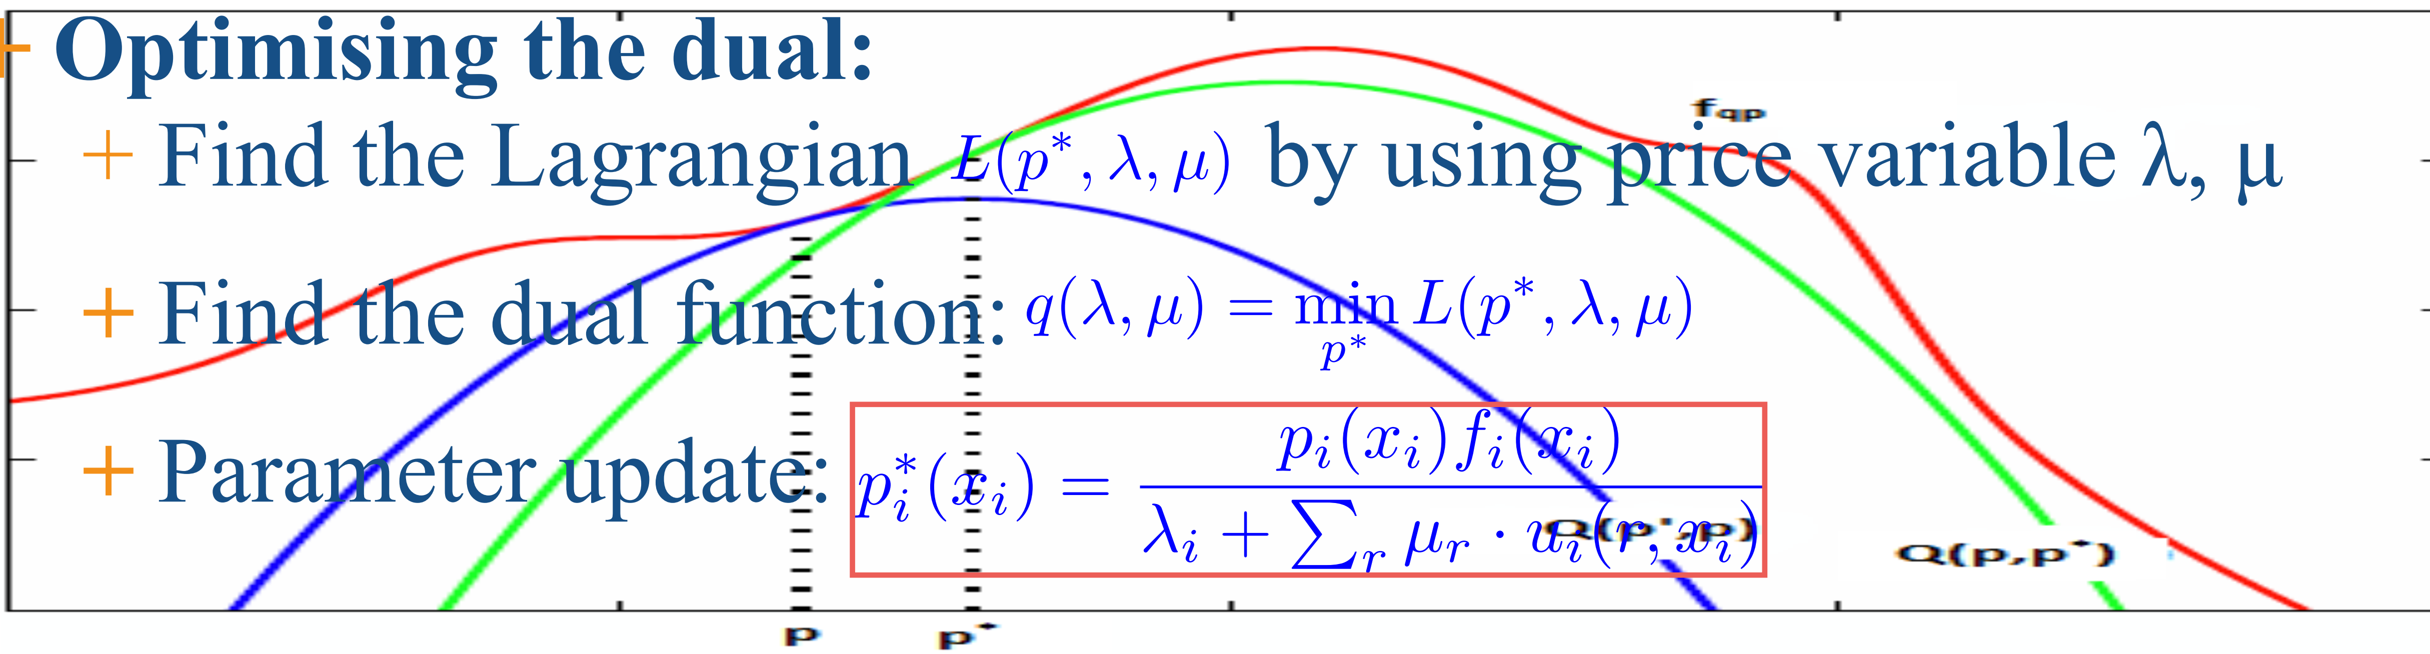

[Kumar *et al.*, 2011] Akshat Kumar, William Yeoh, and Shlomo Zilberstein. On message-passing, MAP estimation in graphical models and DCOPs. In *DCR Workshop*, pages 57–70, 2011.

M-Step

- **Block coordinate descent (BCD):**
  - Choose arbitrary variable λ_i, fix all other variables and optimise q(.) w.r.t. λ_i

$$\sum_{x_i} \frac{p_i(x_i) f_i(x_i)}{\lambda_i + \sum_r \mu_r \cdot u_i(r, x_i)} - 1 = 0$$

- **The largest root is the only feasible solution**

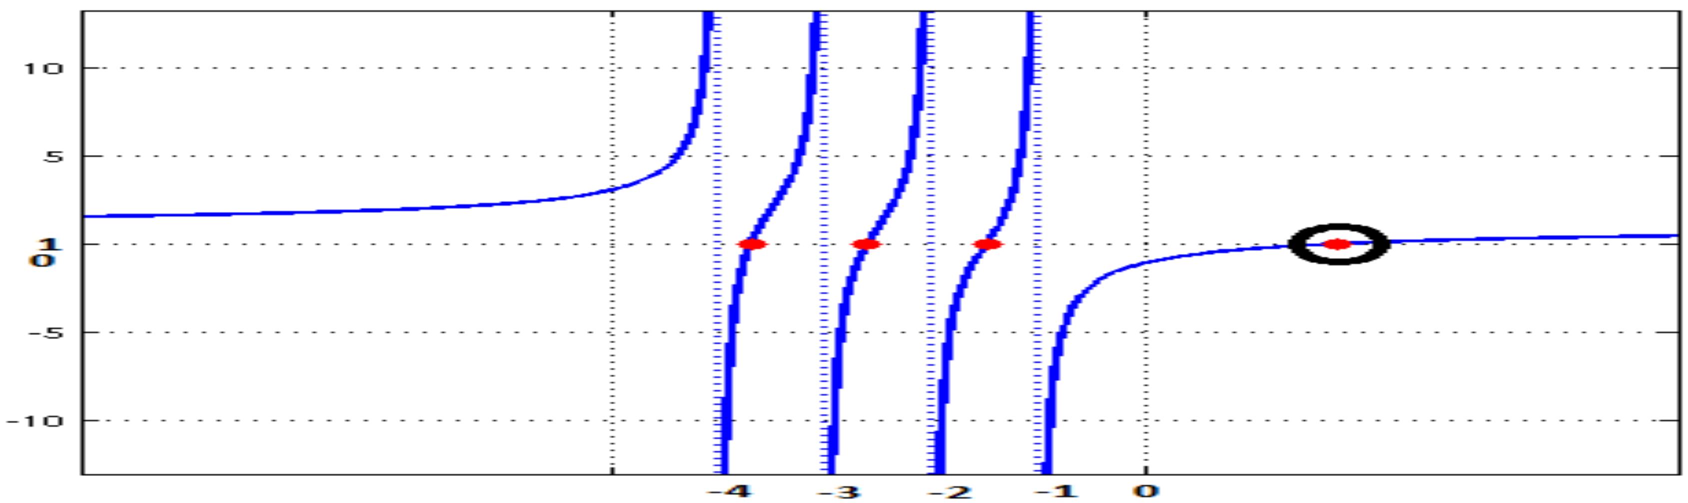

1/(x+1)+1/(x+2)+1/(x+3)+1/(x+4) -1=0

$$g(x) = \sum_{t=1}^{T} \frac{a_t}{x + b_t}$$

- **Minimise q(.) w.r.t. μ_r to find the value of price variable μ_r**
- **As Solution is uniquely determined, BCD is guaranteed to converge.**

Supplementary Slide

- Maximizes the likelihood in mixture of bayes net

$$\sum_{l} P(l) \underbrace{\sum_{x_l} P(\hat{\theta} = 1, x_{l_1}, x_{l_2} | l; p)}_{\text{likelihood for net l}} = \frac{1}{|E|} \underbrace{\sum_{l} \sum_{x_l} \hat{\theta}_{x_l} p_{l_1}(x_{l_1}; p) p_{l_2}(x_{l_2}; p)}_{\text{Objective of the QP for RC-DCOP}}$$

- Utility function for graph colouring problem

$$\theta_{ij}(x_i, x_j) = \begin{cases} 0 & if \quad x_i = x_j \\ 1 + \gamma_i(x_i) + \gamma_j(x_j) & Otherwise \end{cases}$$