ED-BP: Belief Propagation via Edge Deletion

UCLA Automated Reasoning Group

Arthur Choi, Adnan Darwiche,
Glen Lenker, Knot Pipatsrisawat

Last updated 07/15/2010: See the **speaker notes** on each slide for commentary (PowerPoint only).


Glen
Adnan
Knot
Arthur

ED-BP: Idea

loopy BP

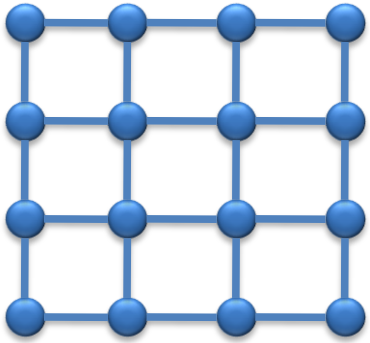

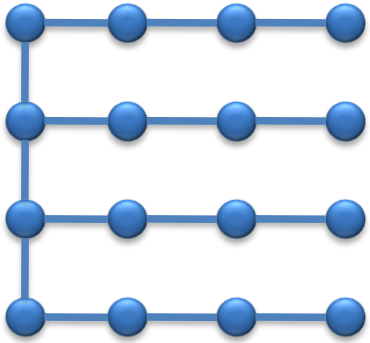

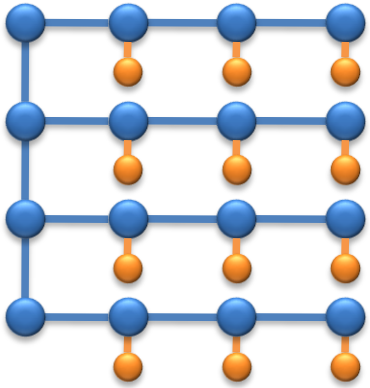

Delete Edges

Compensate

Characterizing Belief Propagation

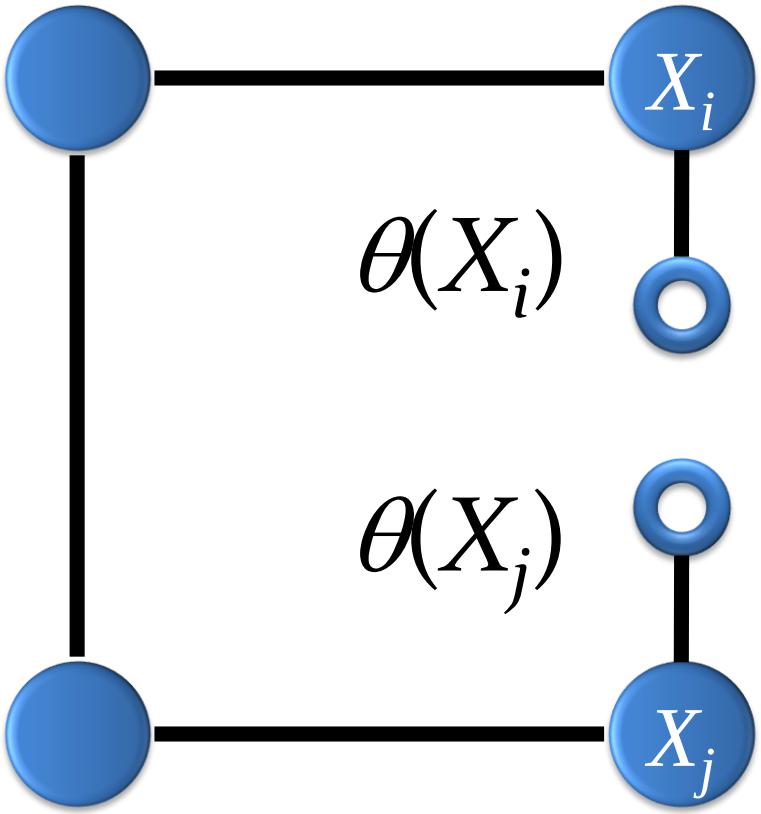

- ED-BP characterization:

$$\Pr(X_i = x) = \Pr(X_j = x) = \theta(X_i = x)\,\theta(X_j = x)/z_{ij}$$

- in a tree:
  - MAR: BP marginals
  - PR: Bethe
  - MPE: BP max-marginals

Edge Deletion

exact

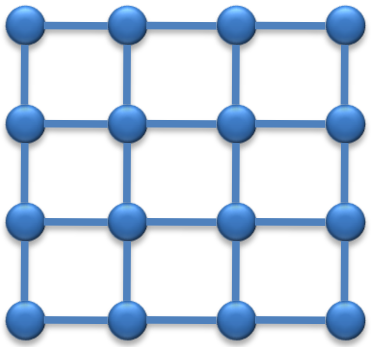

delete edges

?

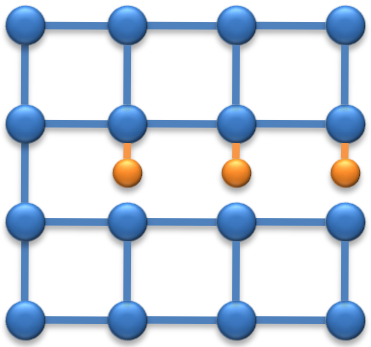

Edge *Recovery*

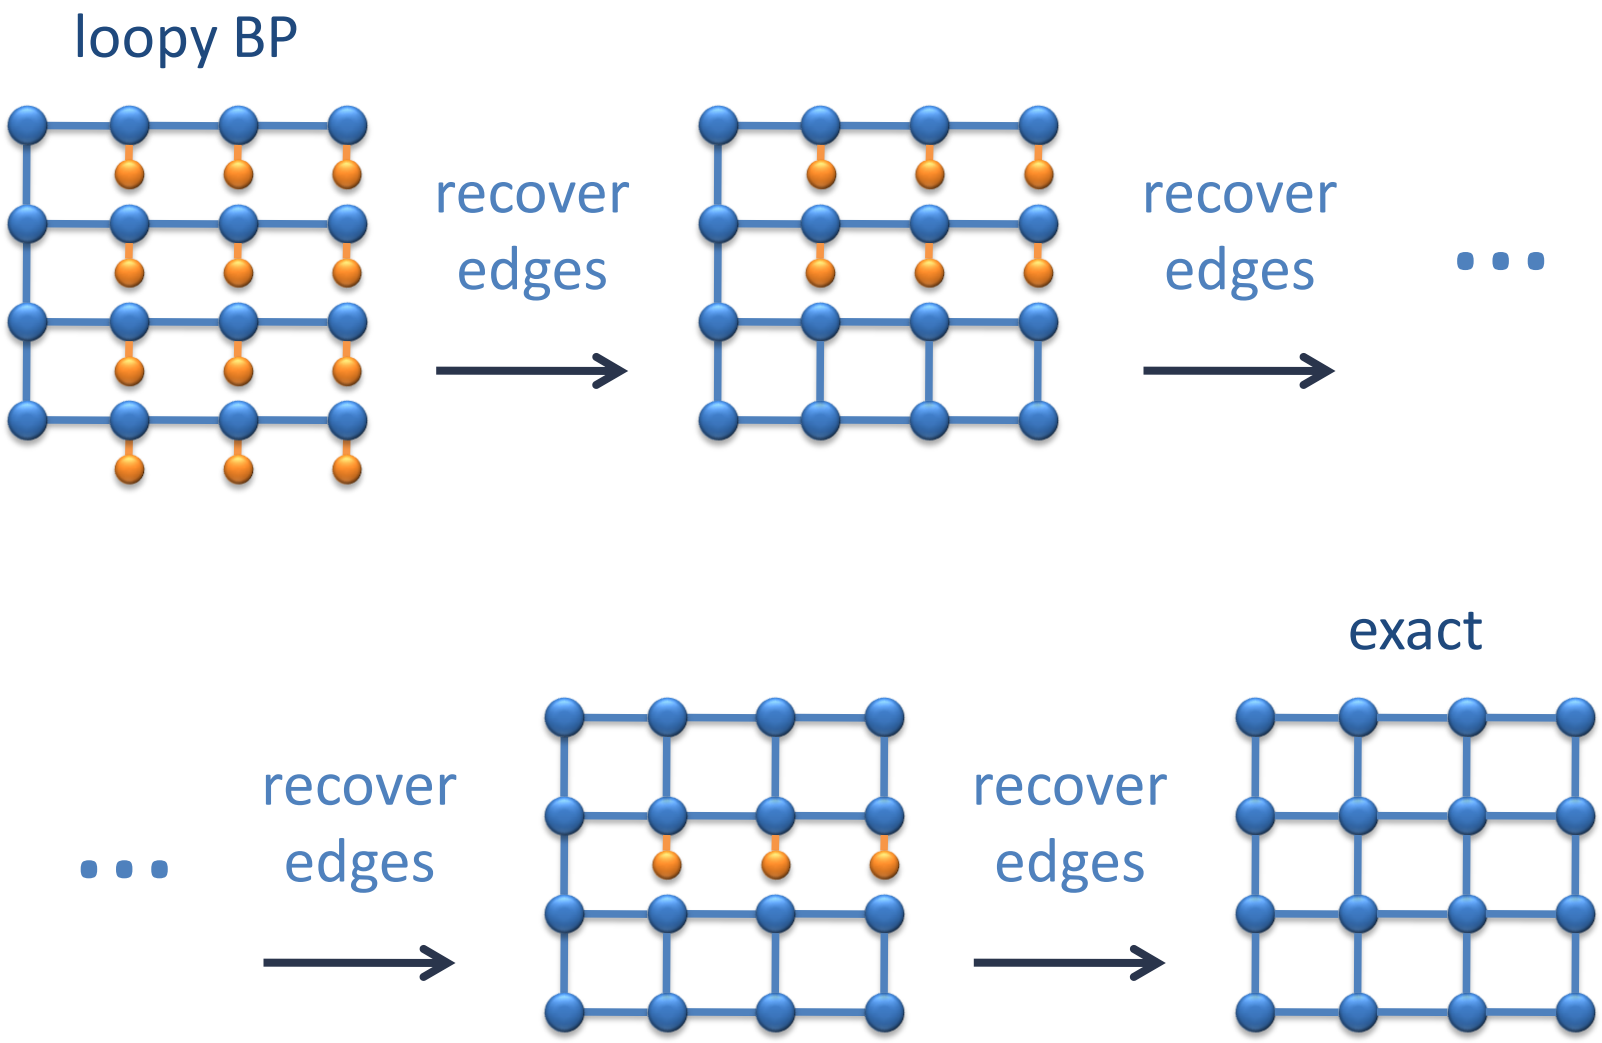

Edge Recovery: Old Idea

loopy BP

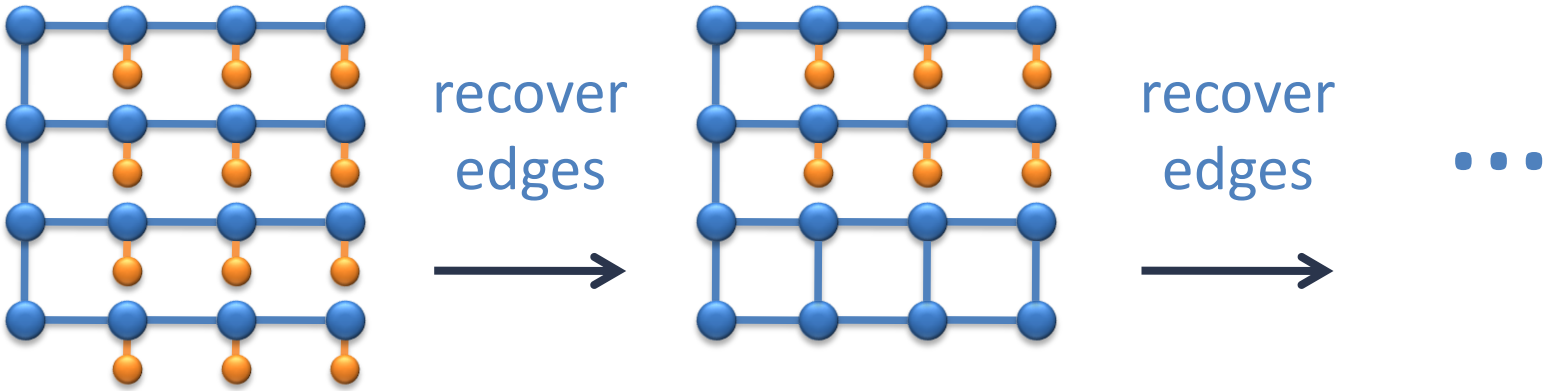

[CD06]: target quality: use mutual information

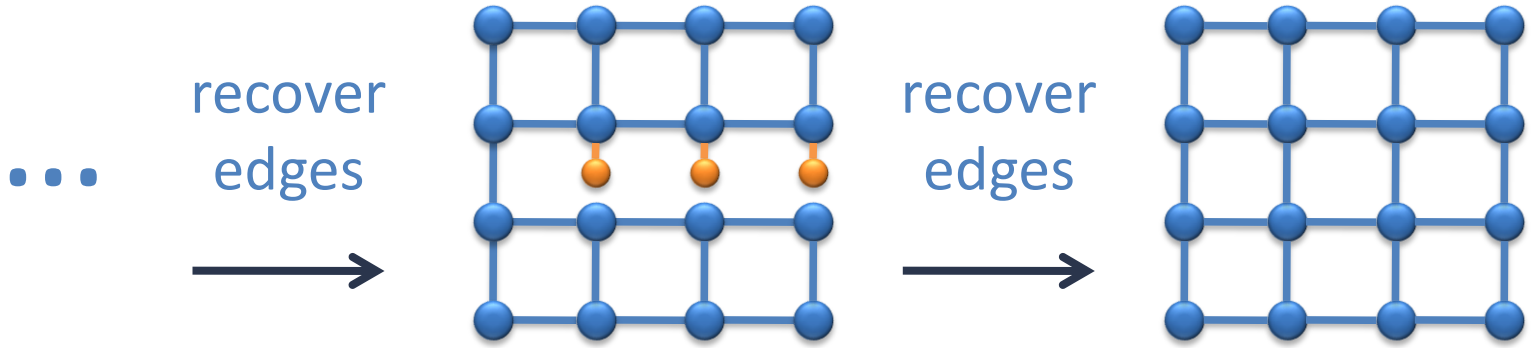

Edge Recovery: New Idea

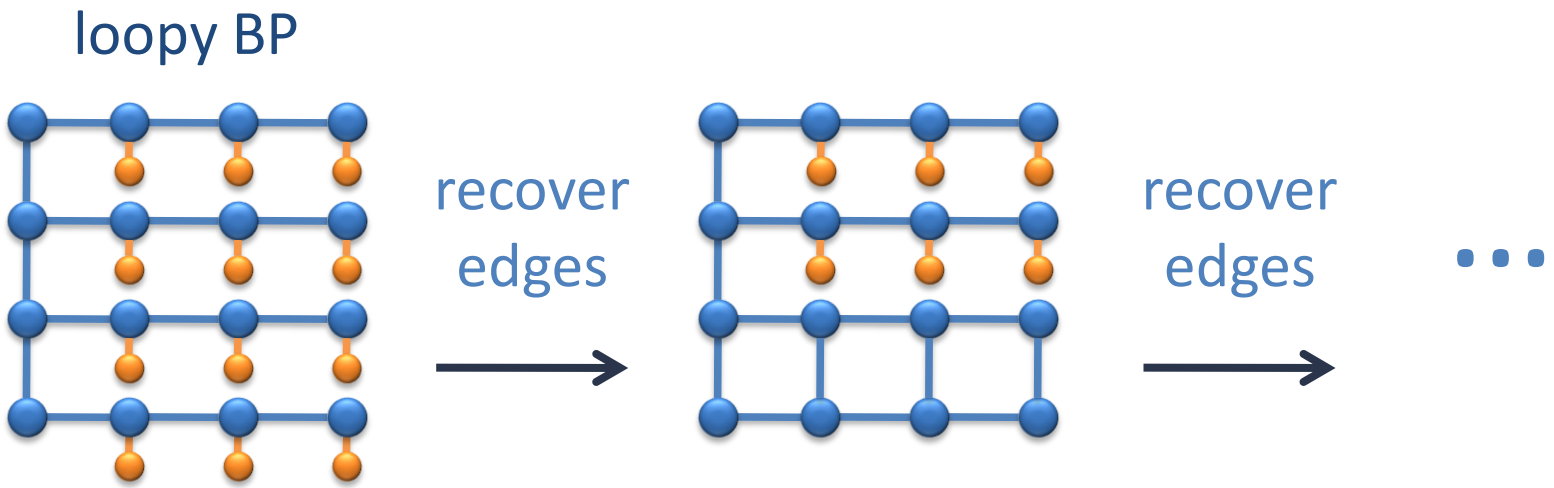

Challenge UAI-10: encourage convergence, *residual recovery*

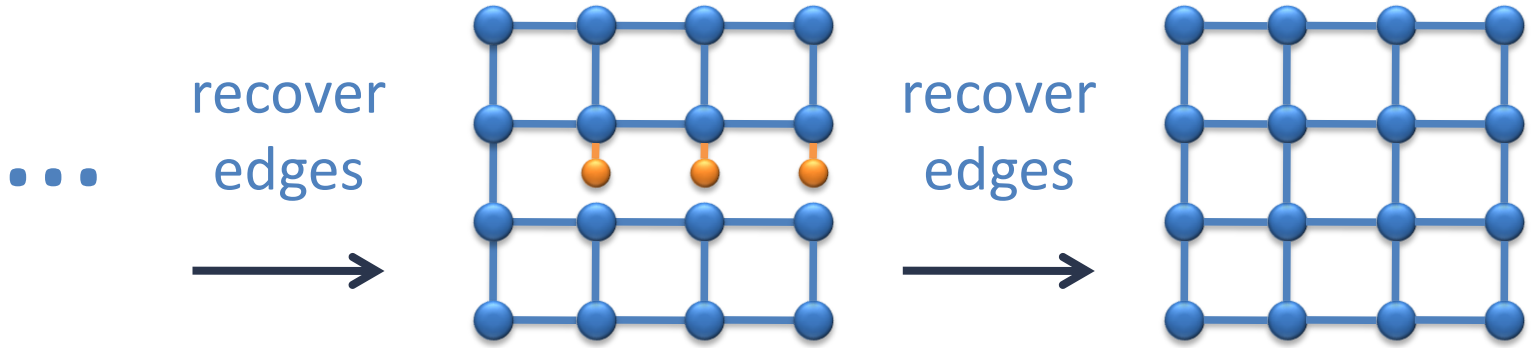

ED-BP: Residual Recovery

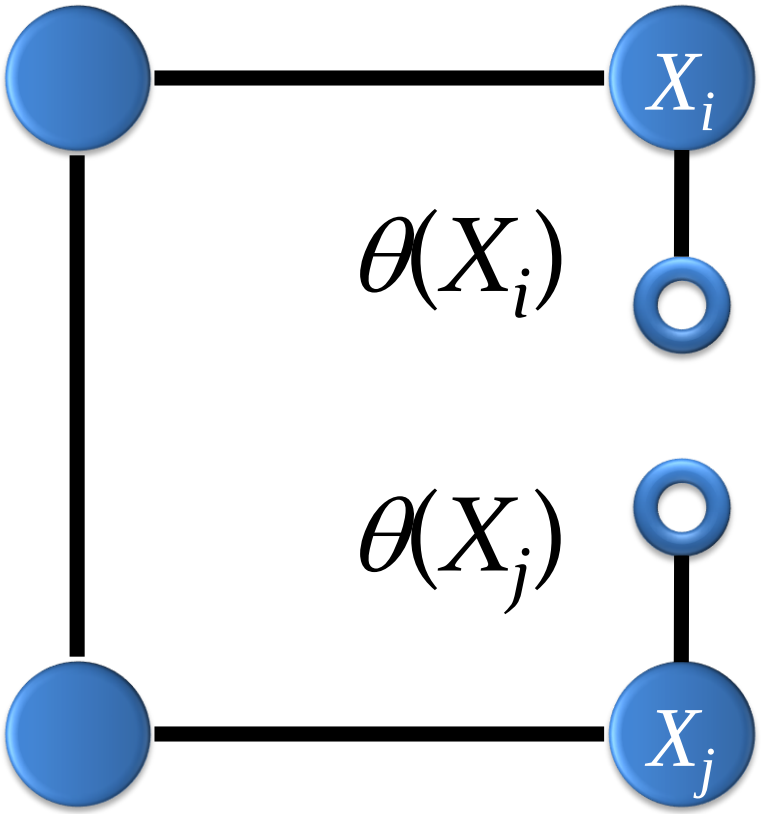

- Recover edges based on how close they are to convergence
- ED-BP characterization:

$$\Pr(X_i = x) = \Pr(X_j = x) = \theta(X_i)\theta(X_j)/z_{ij}$$

- Ongoing: try residuals as in residual BP

Exact Solvers

- Exact inference in a simplified network: ED-BP can use any black box inference engine
  - currently using vanilla Hugin and Shenoy-Shafer jointree algorithms
  - not currently using Ace, or other advanced inference engines …

Overall Results

PR Task: 20 Seconds

Solver	Score
edbr	1.7146
vgogate	2.1620
libDai	2.2775

MAR Task: 20 Seconds

Solver	Score
edbq	0.2390
libDai2	0.3064
vgogate	0.4409

Overall Results

PR Task: 20 Minutes

Solver	Score
vgogate	1.2610
edbp	1.3063
libDai	2.0707

MAR Task: 20 Minutes

Solver	Score
ijgp	0.1722
edbq	0.1742
libDai3	0.2810

Overall Results

PR Task: 1 Hour

Solver	Score
vgogate	1.2609
edbr	1.2699
libDai	2.0707

MAR Task: 1 Hour

Solver	Score
ijgp	0.1703
edbr	0.1753
libDai3	0.2639

Overall Results

PR Task: 1 Hour

Solver	Score
vgogate	1.2609
edbr	1.2699
libDai	2.0707

MAR Task: 1 Hour

Solver	Score
ijgp	0.1703
edbr	0.1753
libDai3	0.2639

Congratulations Vibhav

More Slides on the ED-BP Solver

ED-BP: The Solver

- Based on UAI'08 solver, new MPE version
- Numerous improvements
  - pre-processing
  - initial spanning tree
  - internal inference engine for *exact* reasoning
  - ***edge recovery***
    - led to biggest impact in performance

ED-BP: The Solver

- Pre-processing
  - lightweight
  - RSat: infer fixed values from network zero's
- Initial spanning tree
  - random spanning tree
  - max spanning tree (mutual information)
- Black box engine for exact inference
  - jointree algorithms: shenoy-shafer versus hugin
  - in the future: compilation to ACs (Ace)

References

Arthur Choi, Hei Chan, and Adnan Darwiche. On Bayesian Network Approximation by Edge Deletion. In *Proceedings of the 21st Conference on Uncertainty in Artificial Intelligence (UAI)*, 2005.

Arthur Choi and Adnan Darwiche. An Edge Deletion Semantics for Belief Propagation and its Practical Impact on Approximation Quality. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*, 2006.

Arthur Choi and Adnan Darwiche. A Variational Approach for Approximating Bayesian Networks by Edge Deletion. In *Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence (UAI)*, 2006.

Arthur Choi, Mark Chavira, and Adnan Darwiche. Node Splitting: A Scheme for Generating Upper Bounds in Bayesian Networks. In *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence (UAI)*, 2007.

Arthur Choi and Adnan Darwiche. Approximating the Partition Function by Deleting and then Correcting for Model Edges. In *Proceedings of the 24th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2008.

Arthur Choi and Adnan Darwiche. Focusing Generalizations of Belief Propagation on Targeted Queries. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence (AAAI)*, 2008.

Arthur Choi and Adnan Darwiche. Many-Pairs Mutual Information for Adding Structure to Belief Propagation Approximations. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence (AAAI)*, 2008.

Arthur Choi and Adnan Darwiche. Approximating MAP by Compensating for Structural Relaxations. In *Proceedings of the Twenty-Third Annual Conference on Neural Information Processing Systems (NIPS)*, 2009.

Arthur Choi, Trevor Standley, and Adnan Darwiche. Approximating Weighted Max-SAT Problems by Compensating for Relaxations. In *Proceedings of the 15th International Conference on Principles and Practice of Constraint Programming (CP)*, 2009.