

MOTIVATION

- Bayesian Networks (BNs) are popular models with a firm theoretical foundation and inherent interpretability. Exact inference on BNs is fundamental to many fields such as healthcare.
- Exact inference on BNs is NP-hard. The well-known junction tree (JT) algorithm has a high computational complexity.
- Prior parallelization strategies accelerating JT have efficiency issues and cannot adapt to various network structures.

PRELIMINARIES

- BNs define a set of random variables and their dependencies.
- **Exact inference on BNs** computes the posterior distribution of a subset of variables given observed values of other variables.
- **JT** first converts a BN into a junction tree with multiple *cliques* and *separators*, then conducts inference on the tree.

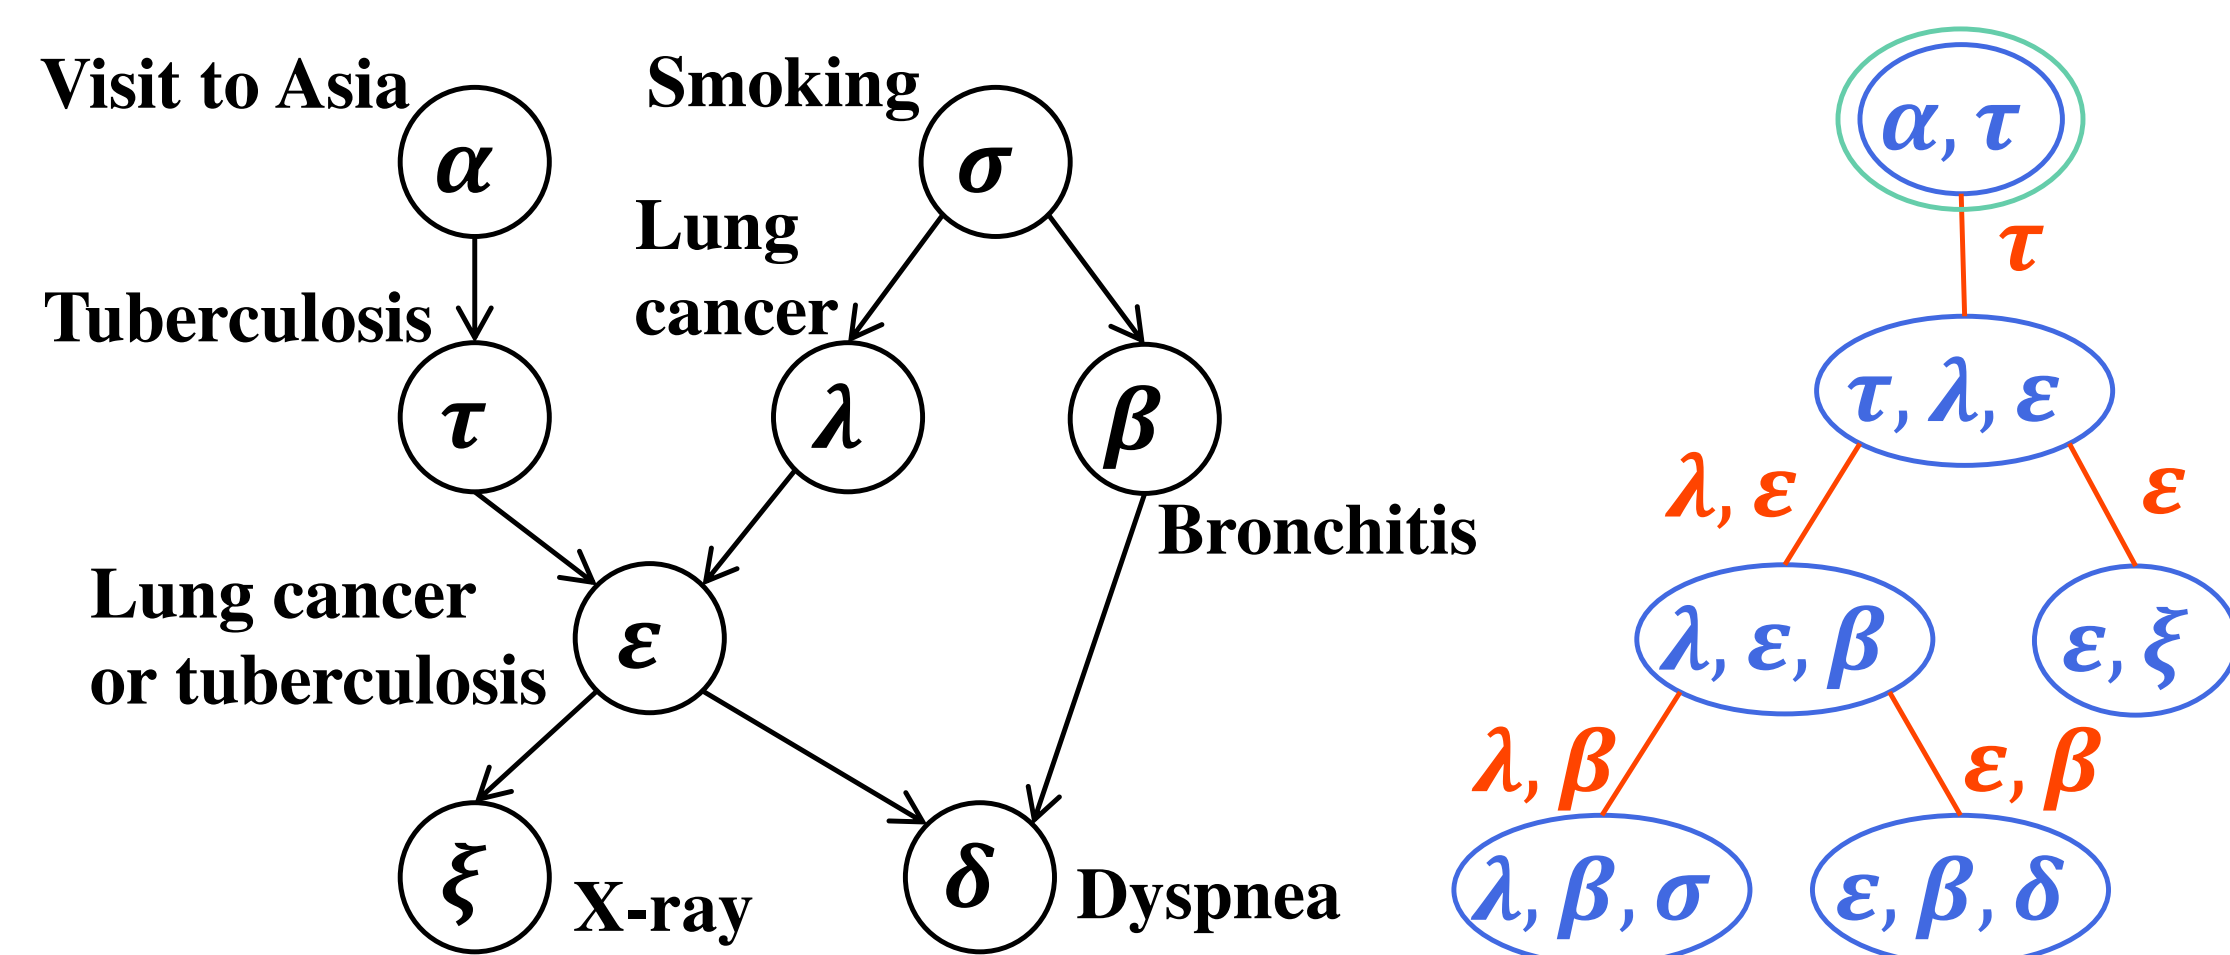


Figure 1: Network structure and the corresponding junction tree of the example BN namely “Asia”. Cliques are in blue, separators are in red, and the double circle indicates the root of the tree.

OVERVIEW OF OUR PROPOSED FAST-BNI

The proposed Fast-BNI is a parallel BN exact inference solution that accelerates the JT algorithm using hybrid inter- and intra-clique parallelism. Compared with prior works that accelerate JT using only one of the two granularities of parallelism, Fast-BNI has the following three main advantages.

- Workload balancing.
- Lower parallel overhead.
- Adaptability to various network structures.

HYBRID INTER- AND INTRA-CLIQUE PARALLELISM

- For the inter-clique parallelism, we develop a *traversal method* to exploit parallelization opportunities across the tree, and a *root selection strategy* to construct a more balanced tree with the minimal number of layers to reduce the total number of parallelization invocations.
- Each clique or separator contains a subset of random variables in the BN and maintains a potential table over these variables. For the intra-clique parallelism, we develop primitives for three dominant potential table operations, including potential table *marginalization*, *extension* and *reduction*.
- Fast-BNI utilizes a hybrid parallelism that closely integrates inter- and intra-clique parallelism. At the beginning of each layer, all the potential table entries corresponding to this layer are packed to constitute one of the parallel tasks. The tasks are then distributed to the parallel threads to perform concurrently.

EXPERIMENTAL RESULTS

We compared sequential and parallel implementations of Fast-BNI with the existing implementations, as shown in Table 1.

- The sequential implementation of Fast-BNI can be 1.2 to 13.1 times faster than UnBBayes.
- When comparing the parallel implementations, Fast-BNI-par can run 1.2 to 15.2 times faster than the counterparts.
- Fast-BNI has more advantages over existing implementations on larger networks.

Table 1: Comparison of Fast-BNI with other implementations, and speedup of Fast-BNI over each compared implementation.

BN	Sequential implementation			Parallel implementation						
	Execution time (sec)		Speedup	Execution time (sec)				Speedup		
	UnBBayes [1]	Fast-BNI-seq	UnBBayes	Dir. [2]	Prim. [3]	Elem. [4]	Fast-BNI-par	Dir.	Prim.	Elem.
Hailfinder	28.3	4.0	7.1	3.0	3.2	4.0	2.5	1.2	1.3	1.6
Pathfinder	319.2	68.9	4.6	40.5	23.6	27.8	11.1	3.6	2.1	2.5
Diabetes	90961	6944	13.1	3016	2311	3316	558.6	5.4	4.1	5.9
Pigs	43714	3729	11.7	3353	1068	2380	221.7	15.1	4.8	10.7
Munin2	3054	2643	1.2	1951	934.7	1638	241.7	8.1	3.9	6.8
Munin4	258194	34198	7.6	20364	10348	21398	3021	6.7	3.4	7.1

REFERENCES

- [1] R. Carvalho, KB. Laskey, P. Costa, M. Ladeira, L. Santos, and S. Matsumoto. UnBBayes: modeling uncertainty for plausible reasoning in the semantic web. *Semantic web*, pages 953–978, 2010.
- [2] AV. Kozlov and JP. Singh. A parallel Lauritzen-Spiegelhalter algorithm for probabilistic inference. In *Supercomputing*, pages 320–329, 1994.
- [3] Y. Xia and VK. Prasanna. Node level primitives for parallel exact inference. *SBAC-PAD*, pages 221–228, 2007.
- [4] L. Zheng. *Parallel Junction Tree Algorithm on GPU*. PhD thesis, Carnegie Mellon University, 2013.

ACKNOWLEDGEMENT

Professor Ajmal Mian is the recipient of an Australian Research Council Future Fellowship Award (project number FT210100268) funded by the Australian Government.

CONTACT INFORMATION

Name Jiantong Jiang
Email jiantong.jiang@research.uwa.edu.au
Fast-BNI Code <https://github.com/jjiantong/FastBN>