

Parallelizing Exact Inference in Bayesian Networks

High Performance Embedded Computing 2006

Vasanth Namasivayam, Animesh Pathak
and Viktor K. Prasanna*

{namasiva, animesh, prasanna}@usc.edu

University of Southern California

<http://ceng.usc.edu/~prasanna/>



Problem Definition

Bayesian Network (BN)

a popular tool for presenting joint probability distributions

Exact Inference on BN

The time complexity of inference is exponential with the BN's density, variable ranges and clique width.

$$O(nk^3 + n^2w + wr^wn + nr^w)$$

Objective: Scalable Parallel Exact Inference on an **arbitrary** BN

Input

An arbitrary BN
Evidence
(Observations)

Convert
Arbitrary BN
to Junction
Tree in
Parallel

Parallel Exact
Probabilistic
Inference
Algorithm

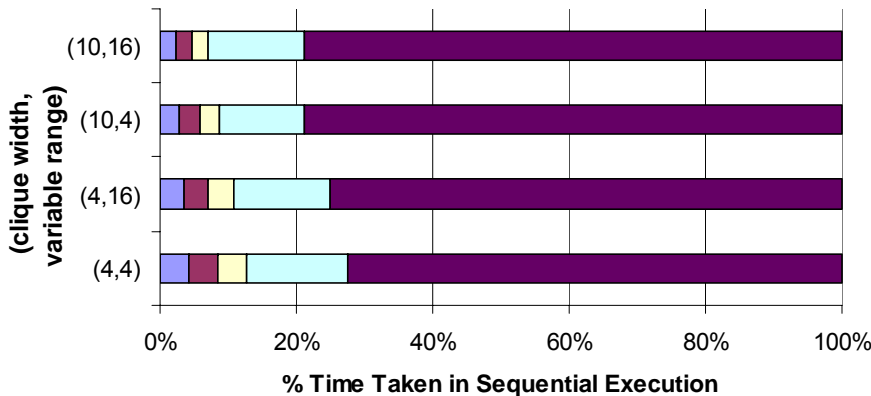
Output

Posterior
probability of
query variables

Probabilistic Network Library (PNL)


- Full function, graphical model library
- Open Source, BSD style license
- Created by Intel, part of Intel Open Source Library
- Written in C++
- Parallelized using OpenMP

Proportion of time taken in various stages




Performance of PNL

w.r.t. Bayesian Network type




w.r.t. (clique width, variable range)




Exact Inference on BN

Our Techniques Include



Parallel Junction
Tree Creation



Pointer
Jumping

Running time


$$O(nk^3/p + n^2w/p + wr^wn/p + nr^w \log n/p)$$

Scalable for $1 \leq p \leq n$

Currently exploring the implementation of our techniques on the runtime system and architecture developed as part of CEARCH(DARPA ACIP Program)

Performance

Balanced Tree



Random Tree

