

Abstract Submitted
for the MAR14 Meeting of
The American Physical Society

Quantum Inference on Bayesian Networks THEODORE YODER,
GUANG HAO LOW, ISAAC CHUANG, Massachusetts Institute of Technology —
Because quantum physics is naturally probabilistic, it seems reasonable to expect physical systems to describe probabilities and their evolution in a natural fashion. Here, we use quantum computation to speedup sampling from a graphical probability model, the Bayesian network. A specialization of this sampling problem is approximate Bayesian inference, where the distribution on query variables is sampled given the values e of evidence variables. Inference is a key part of modern machine learning and artificial intelligence tasks, but is known to be NP-hard. Classically, a single unbiased sample is obtained from a Bayesian network on n variables with at most m parents per node in time $\mathcal{O}(\frac{1}{P(e)})$, depending critically on $P(e)$, the probability the evidence might occur in the first place. However, by implementing a quantum version of rejection sampling, we obtain a square-root speedup, taking $\mathcal{O}(\frac{1}{\sqrt{P(e)}})$ time per sample. The speedup is the result of amplitude amplification, which is proving to be broadly applicable in sampling and machine learning tasks. In particular, we provide an explicit and efficient circuit construction that implements the algorithm without the need for oracle access.

Theodore Yoder
Massachusetts Institute of Technology

Date submitted: 15 Nov 2013

Electronic form version 1.4