Expressing and Verifying

Probabilistic Assertions

A new way to check properties of programs that behave statistically.

Approximate Computing

Compute: Allow random errors in operations to improve hardware efficiency.
Check: Output is likely to be high-quality even in the face of error.

Data Obfuscation for Privacy

Compute: Add random noise to private data to avoid divulging exact information.
Check: Obfuscated data is still useful in aggregate.

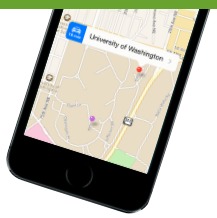

Mobile and Sensing

Compute: Draw conclusions from noisy sensor data.
Check: Conclusions are still useful to the user.

Probabilistic programs need **probabilistic assertions.**

`passert` `e, p, c`: Expression *e* is true with at least probability *p* at confidence level *c*.

distribution extraction

The concrete (ordinary) semantics are nondeterministic.

A symbolic semantics captures a program's probabilistic behavior.

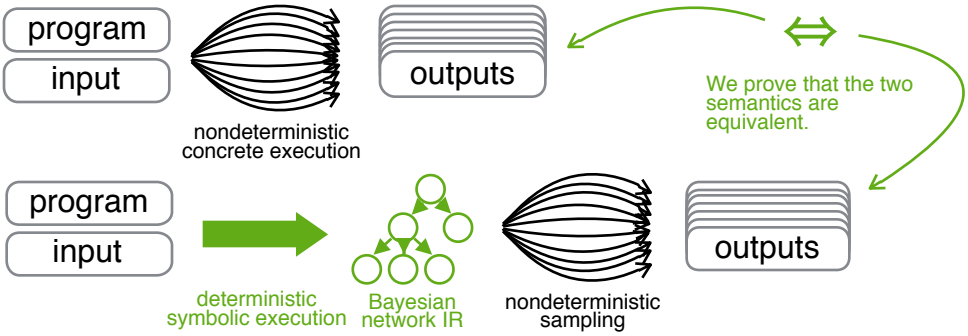

statistical optimizations

The expression dag output from distribution extraction is a Bayesian network, a representation of probability distributions that lets statistical properties act as optimizations.

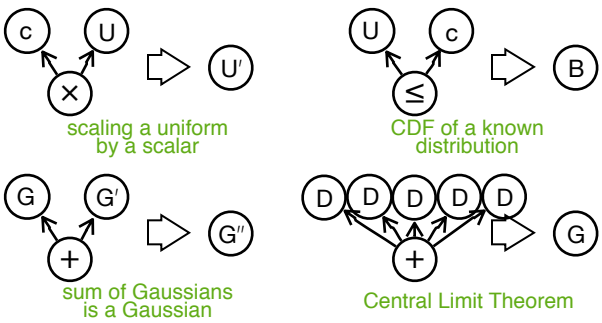

Simpler networks mean faster sampling.

verification

Optimizations collapse the network to a Bernoulli and we verify the `passert` exactly.

OR

Sample the network and perform a hypothesis test to get a statistical guarantee.

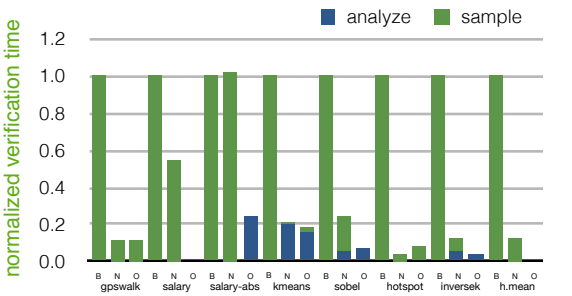

for each benchmark: baseline, no optimization, optimized

Our verifier checks `passerts` 24× faster than a naive checker on average.

sampa

Microsoft Research

Adrian Sampson
Pavel Panchekha
Todd Mytkowicz
Kathryn S. McKinley
Dan Grossman
Luis Ceze

contact:
asampson@cs.washington.edu
http://homes.cs.washington.edu/~asampson