

Research



Cite this article: Deka B, Birklykke AA, Duwe H, Mansinghka VK, Kumar R. 2014 Markov chain algorithms: a template for building future robust low-power systems. *Phil. Trans. R. Soc. A* **372**: 20130277.
<http://dx.doi.org/10.1098/rsta.2013.0277>

One contribution of 14 to a Theme Issue 'Stochastic modelling and energy-efficient computing for weather and climate prediction'.

Subject Areas:

electrical engineering

Keywords:

Markov chain, error tolerance, algorithmic fault tolerance

Author for correspondence:

Biplab Deka
e-mail: deka2@illinois.edu

Markov chain algorithms: a template for building future robust low-power systems

Biplab Deka¹, Alex A. Birklykke^{1,2}, Henry Duwe¹,
Vikash K. Mansinghka³ and Rakesh Kumar¹

¹Department of Electrical and Computer Engineering, University of Illinois at Urbana Champaign, Champaign, IL, USA

²The Technology Platforms Section, Aalborg University, Aalborg, Denmark

³Intelligence Initiative, Massachusetts Institute of Technology, Cambridge, MA, USA

Although computational systems are looking towards post CMOS devices in the pursuit of lower power, the expected inherent unreliability of such devices makes it difficult to design robust systems without additional power overheads for guaranteeing robustness. As such, algorithmic structures with inherent ability to tolerate computational errors are of significant interest. We propose to cast applications as stochastic algorithms based on Markov chains (MCs) as such algorithms are both sufficiently general and tolerant to transition errors. We show with four example applications—Boolean satisfiability, sorting, low-density parity-check decoding and clustering—how applications can be cast as MC algorithms. Using algorithmic fault injection techniques, we demonstrate the robustness of these implementations to transition errors with high error rates. Based on these results, we make a case for using MCs as an algorithmic template for future robust low-power systems.

1. Introduction

It is becoming increasingly clear that disruptive low-power solutions are needed to meet future power-performance goals. Such disruptive solutions are already being explored in the context of microprocessors. Examples include approaches based on trading robustness for lower power such as better-than-worst-case

design [1]. However, it is not clear whether such approaches scale well to post CMOS technologies where hardware fault rates are going to be significantly higher [2]. As such, a radical rethinking of system design might be needed to build robust systems in the fault-prone post CMOS era.

In this paper, we explore one such disruptive solution—robustifying applications by casting these as Markov chain (MC) algorithms. This approach is based on the insight that MC algorithms can produce acceptable results even in the face of certain errors (§3*b*). We specifically focus on the Las Vegas formulations of MC algorithms. In Las Vegas algorithms, the solution to a problem is embedded in the state of the MC and the solution is discovered when the MC visits a solution state. In this formulation, we observe that transition probability errors only affect the expected runtime of the algorithm while the algorithm still produces the correct solution. This formulation is not unlike the use of MCs in Markov chain Monte Carlo (MCMC) approaches that have found wide use in inference applications [3]. However, in this paper, we focus on applications that are generally not implemented as MCs and make a case for implementing them as MC applications for increased robustness.

Given the above insight, we propose to replace the rigid and sensitive control flow used in many conventional algorithms for solving deterministic problems with a Markov process that guides the execution towards the solution through a sequence of random steps. To leverage the robustness of such algorithms, we propose building systems that let faults in devices manifest themselves only as transition errors at the algorithmic level. Such an approach could form the basis of building systems that save power by using low-power, possibly unreliable, devices and allow only controlled errors; errors that the application can tolerate. Such systems may also be able to exploit inherent parallelism exhibited by such algorithms and provide overall energy benefits [4].

Clearly, the underlying assumption here is that it is possible to construct an MC that converges to the solution of a given problem. We show how such MCs can be constructed with four example applications: Boolean satisfiability (SAT), low-density parity-check (LDPC) decoding, sorting and clustering. In addition, we note that such implementations already exist for a wide range of problems such as the ones solved using stochastic local search strategies.

This paper makes the following contributions:

- We propose a novel algorithmic approach to building robust applications—by transforming applications into MC algorithms. We propose executing such algorithms on a system that limits the manifestation of hardware faults to transition errors in the MCs.
- We present a methodology for constructing MC algorithms for four problems—Boolean SAT, LDPC decoding, sorting and clustering. We argue that the methodology can be applied to a large class of applications.
- We quantify the robustness benefits of casting applications into MC algorithms using algorithmic error injections. We show that MC implementations are indeed significantly more robust than the traditional implementations of these applications. The MC applications were able to tolerate transition errors with rates as high as 10–20%. Even at high error rates, MC algorithms produced acceptable results in terms of runtime and output quality. Also, these algorithms showed gradual degradation in runtime or output quality with increasing error rates.

Our results on the robustness of MCs show that MCs may indeed be an attractive template for building future low-power systems. Such systems could trade off robustness for power benefits using one of several strategies, such as (i) voltage overscaling, (ii) use of low-power, possibly unreliable, post CMOS devices [2], (iii) eliminating design guardbands [1] or (iv) alleviating the need for on-chip redundancy techniques for robust execution. Such systems may also be able to leverage inherent parallelism exhibited by MCs.

2. Related work

The proposed approach to build robust applications by casting them into MC algorithms can be thought of as a novel approach for *application robustification* [5]. Sloan *et al.* [5] propose an approach for application robustification where an application with unknown correct output x^* is cast into a minimization function $f(x)$ whose minimum lies at x^* . The minimization function is then solved using an error-tolerant solver such as gradient descent or conjugate gradient. Our approach, on the other hand, casts an application into an MC where peak(s) in the limiting distribution correspond to the application output(s).

More generally, the proposed approach falls into the category of algorithm-based fault tolerance (ABFT) [6]. Most prior ABFT work is focused on error detection. Correction still relies on checkpointing and restart. Checkpoint-and-restart approaches may have prohibitive overhead at high fault rates. Some works do exist on algorithmic correction [6,7]. However, these approaches are specific to algebra applications.

A related body of work exists on probabilistic computing [8]. However, prior approaches have focused on techniques to program and execute probabilistic applications. This is the first work to the best of our knowledge that attempts to build robust general applications by casting them into probabilistic implementations via MC algorithms.

3. Background and motivation

In this section, we review some salient properties of MCs and discuss how applications can use MCs to find their solutions and why such MC algorithms would be robust to certain errors.

(a) Applications and Markov chains

Consider an application that has a set of possible solutions or states. One of those solutions is actually the correct solution or a goal state. For example, consider sorting. Any permutation of the inputs is a state and one such permutation is the goal state. Assuming that there exists an efficient mechanism to check if a particular solution is the correct solution, one approach to find the correct solution of the application is to generate samples from the state space and check to see if any of them is indeed the correct solution. First, consider the case when these samples are generated completely randomly (figure 1a). This scenario corresponds to generating samples from a uniform distribution over the states. For practical applications, the state space is generally large. So, if samples are generated completely randomly, it might take a large number of samples before we find the correct solution. Clearly, this scheme is not efficient.

An MC algorithm performs the above sampling more intelligently. An MC algorithm is an iterative algorithm which, in every iteration, produces a sample from the sample space of the application (figure 1c). These chains are constructed such that, for a given application, the distribution over states has a significant peak at the correct solution (or the goal state; figure 1b), i.e. the probability of generating the goal state as a sample is significantly higher than that of the other states. This suggests that if we use these samples to check for a solution, we will find the solution much faster than when using completely random sampling. Note that such a strategy can also be adopted for the case when the application can have more than one correct solution. In that case, the *steady-state distribution* over states will have multiple peaks, one at each of the solution states.

How does an MC algorithm ensure such a steady-state distribution over states? This is accomplished by guaranteeing specific *transition probabilities* between states. These transition probabilities dictate what state the MC will produce next given the state that it has just produced. When these transition probabilities are such that the MC has certain properties (*irreducibility and aperiodicity*), there is a unique steady-state distribution over states [9]. This provides applications, a mechanism, to ensure that the overall steady-state distribution over states has a specific

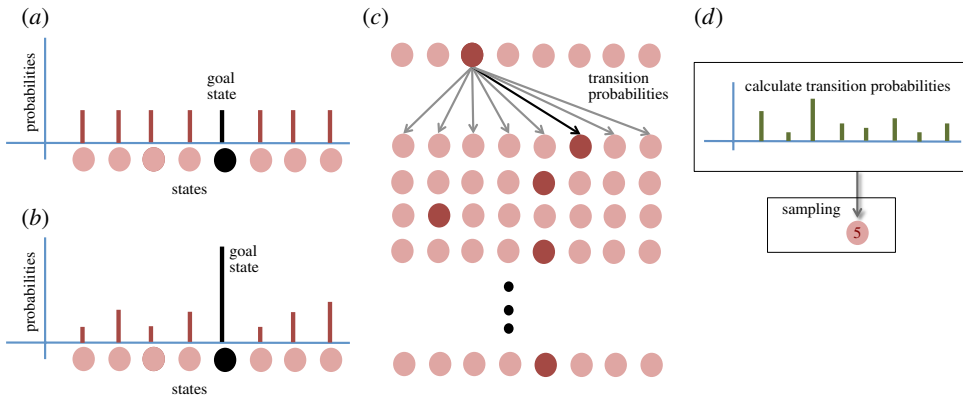


Figure 1. (a) Random sampling in the state space results in a uniform steady-state distribution over states. (b) MC sampling results in a steady-state distribution over states that has a peak at the goal state. (c) An MC produces a sample from the state space in each iteration. (d) In each iteration, the MC algorithm performs two computations: calculating the transition probability distribution and generating a sample from the distribution. (Online version in colour.)

structure (one with a peak at the goal state) simply by ensuring specific transition probabilities in each iteration of the algorithm.

The specific method to calculate the transition probabilities in each iteration depends on the application. However, for all applications, each iteration of the MC algorithm consists of two essential computations (figure 1d). First is the calculation of the transition probabilities. These probabilities depend on the current state and the application inputs. Second is drawing a sample for the next state from the transition probability distribution. This becomes the state that the MC produces in the next iteration. These computations in each iteration, when performed appropriately for each application, result in an MC that has a steady-state distribution over states with peaks at the correct solutions of the application. We present several examples of this process in §4.

(b) Robustness of Markov chain algorithms

In this section, we discuss why MCs are expected to be robust to certain kinds of errors. As described in §3a, an MC algorithm performs two operations in every iteration: calculating the transition probability distribution and sampling from this distribution (figure 1d). Given this understanding, we can think of the robustness of such algorithms at two different levels. First, *within an iteration*, if there are errors in calculating the transition probability distribution, it does not necessarily mean that the sample generated in that iteration would be different from the case in which there were no errors. This is the first level of error tolerance.

Second, even if errors did result in a different sample, when we look *across iterations* at the overall algorithm, the effect is that the steady-state distribution over states will be different. However, it is not essential for efficient execution to have an exact distribution as long as the altered distribution has a peak at the correct solution. For such an altered distribution, it is still possible to arrive at a solution in a reasonable amount of time. Owing to the above two levels of error tolerance, we expect MCs to produce acceptable outputs even in the presence of errors. These errors can be in transition probability calculations, in sampling from the transition probability distribution or other errors that lead to a transition error.

Errors may have an effect on runtime, however. The more the steady-state distribution is altered due to errors, the higher is the potential runtime. Errors could also degrade the output quality for certain applications. We discuss these effects for example applications in §6. Ultimately, what transition error rates are tolerable to an application would depend on what runtime and output quality are acceptable to that particular application.




Figure 2. An approach for converting applications to MC algorithms. An MC is constructed such that the steady-state distribution over states has peaks at the correct solutions. Samples produced by this MC are checked to see whether they are the correct solutions. (Online version in colour.)

(c) Generality of Markov chain algorithms

MCs have found wide use in the domain of statistical inference [10] and machine learning [3]. In addition, they are used to solve many other computational problems—especially hard problems—using stochastic local search methods. A few examples include the travelling salesman problem [11], the knapsack problem [12], VLSI placement, estimation of the matrix permanent [13] and constraint satisfaction problems [14]. As such, many applications already have MC implementations.

In addition, applications normally *not* considered as sampling may be implemented using MC algorithms by using the general methodology described in §3*a*. This involves coming up with a strategy to calculate transition probabilities in each iteration of the MC such that the steady-state transition probability over states has peaks at the correct solutions (figure 2). In §4, we describe how MC implementations for four example applications are made possible by employing this methodology.

4. Casting applications as Markov chain algorithms

In this section, we present a case study of four applications that are representative of a wider range of applications to demonstrate the general methodology of casting applications as MC algorithms. After discussing the details of each of these applications, we describe how MC algorithms are used to find solutions for these applications and also the output quality metrics for each application.

We chose two well-understood combinatorial applications: Boolean SAT and sorting. These applications are representative of very wide classes of algorithmic problems. SAT is the canonical \mathcal{NP} -complete problem [15]. Therefore, if we can implement SAT as an MC algorithm, we can obtain MC implementation for any problem in \mathcal{NP} by a reduction to SAT. Sorting is a canonical \mathcal{P} -complete problem [15] and so if we can cast sorting as an MC algorithm, all other problems in \mathcal{P} can be cast to an MC via sorting. The third application, LDPC decoding, is a representative application from the area of communication systems. The fourth application, clustering without prior knowledge of the number of clusters, is a representative application from the domain of unsupervised machine learning.

(a) Boolean satisfiability

Boolean SAT is the problem of determining whether there exists a satisfying assignment to a set of Boolean variables given a set of constraints. If the set of Boolean variables is x_1, x_2, \dots, x_n , one example clause could be

$$x_1 \vee \neg x_3 \vee x_4.$$

For a problem to be satisfiable, there must exist an assignment of the variables x_1, x_2, \dots, x_n such that all clauses evaluate to true. There are two main classes of SAT solvers: *complete* and *incomplete* [16]. Complete SAT solvers will always find a satisfiable assignment if one exists or will report that the problem is unsatisfiable. Incomplete SAT solvers on the other hand only look for a satisfiable solution assuming that it exists and is not able to report if a problem is unsatisfiable.

In the MC sampling framework, we can think of any assignment of the Boolean variables as a state and the particular assignments that result in all clauses being satisfied as the goal states. Our objective then is to construct an MC that has a steady-state distribution over states with peaks at the goal states. As discussed in §3*a*, this can be done by using a suitable transition probability function.

We construct a transition probability function that is based on the stochastic local search mechanism used in WalkSAT [17]—a well-known and effective incomplete SAT solver that has found use in areas such as automated planning [18]. We place a restriction in terms of what state transitions are allowed—from any particular state, the algorithm can only transition to another state that differs only in the assignment of a single variable. Thus, transition probabilities to other states become equivalent to probabilities of different variables in the present state being flipped. We calculate these probabilities over variables and sample from it using the following procedure in each iteration:

- (i) Pick a clause i by sampling from a distribution over clauses where the probability assigned to a clause x is given by

$$p(i | \mathbf{x}) = \frac{1}{Z} \exp \left\{ -\frac{1}{\beta} g_i(\mathbf{x}) \right\}, \quad (4.1)$$

where $g_i(\mathbf{x})$ is either 0 or 1 based on whether the clause is satisfied or not and Z is a normalizing factor ensuring that the probabilities sum up to 1 over all clauses. In this distribution, unsatisfied clauses are assigned a high probability and satisfied clauses a low probability.

- (ii) We assign a probability of zero to any variable that is not present in clause i . For variable x_j present in position j in clause i , let $N(x_j)$ be a function that returns the break count (the number of satisfied clauses that become unsatisfied if that variable is flipped) for the variable. Let G_i be the set of variable indices associated with clause g_i . Using these quantities, we can define the conditional probability of the random variable $j \in G_i$ given clause i as

$$p(j | i, \mathbf{x}) = \frac{1}{Z} \exp \left\{ -\frac{1}{\alpha} U(x_j) \right\}, \quad (4.2)$$

where Z is a normalizing factor and $U(x_j)$ is an energy function defined by

$$U(x_j) = N(x_j)(1 + k \cdot \delta[N(x_j) > 0]).$$

Here, α is our temperature and $k > 0$. The energy function ensures that variables with low break counts are assigned high probabilities. Pick a variable j by sampling from this distribution.

- (iii) Flip the state of variable j .

Figure 4 presents McSAT, the MC algorithm based on this procedure. Every iteration of this algorithm computes a probability distribution over variables, samples from it and flips that variable. This is equivalent to sampling the next state from a transition probability distribution over states. The procedure used to calculate these probabilities results in the steady-state distribution over states that has significant peaks at the goal states. We present the empirical probability distribution over states (the binary assignment \mathbf{x} converted to base-10) when the algorithm is allowed to run for 1 million iterations without the termination condition for a random 3-CNF problem in figure 3. We observe that the steady-state distribution over states indeed has significant peaks at the states that satisfy this problem.




Figure 3. Steady-state distribution over states for McSAT running 1 million iterations on a random 3-CNF problem with 20 variables and 91 clauses. We verified that the distribution indeed has peaks at the satisfied states.

```

procedure McSAT( $C$ )
  Randomly initialize  $\mathbf{x}^{(0)}$ 
  for  $t = 1, 2, \dots$  do
     $i \sim p(i | \mathbf{x}^{(t-1)})$   $\triangleright$  Draw sample from Eqn. 4.1
     $j \sim p(j | i, \mathbf{x}^{(t-1)})$   $\triangleright$  Draw sample from Eqn. 4.2
     $\mathbf{x}^{(t)} \leftarrow \mathbf{x}^{(t-1)}$ 
     $x_j^{(t)} \leftarrow 1 - x_j^{(t-1)}$   $\triangleright$  Flip variable
    if No unsatisfied clause then
      break  $\triangleright$  Termination Condition
    end if
  end for
end procedure

```

Figure 4. McSAT: an MC algorithm for SAT. (Online version in colour.)

Note that, in the case of SAT, any solution that satisfies all clauses is acceptable. Thus, as long as the MC algorithm produces such a solution, there is no effect on the application's output quality.

(b) Low-density parity-check decoding

LDPC codes are linear block codes characterized by sparse parity check matrices. An (N, M) linear block code has a codeword of length N with M parity check bits. It is specified by its parity check matrix H and a generator matrix G . H has the property that $H\mathbf{c} = 0$, where \mathbf{c} is a valid codeword. G is used to encode a message \mathbf{u} into a codeword $\mathbf{c} = G\mathbf{u}^T$. This process is called source encoding. Once the message has been encoded, the codeword is sent from the transmitter to the receiver over a noisy channel.

Source decoding is performed at the receiver by computing the syndrome \mathbf{s} given by $\mathbf{s} = H\mathbf{z}$, where H is the parity check matrix of the code that was used. If \mathbf{z} differs from \mathbf{c} due to channel noise, the result will be a non-zero vector. Thus, when $\mathbf{s} \neq 0$, an error has been detected and error correction will have to be performed. If, on the other hand, $\mathbf{s} = 0$, the received codeword is believed to be valid, and the message can be extracted from it by looking at the appropriate bits.

In an MC setting, the assignment of the N bits in the received vector can be considered a state. An assignment for which $\mathbf{s} = 0$ is a goal state. Our objective is to construct an MC that has a steady-state distribution over states with peaks at the goal states. To do that, we construct a transition probability function that is based on the weighted bit flip (WBF) algorithm [19]—one of the several algorithms for decoding LDPC codes. We place the restriction that the algorithm can only transition to a state that differs in the assignment of a single bit. Thus, transition probabilities to other states become equivalent to probabilities of different bits in the present state being flipped.

```

procedure McWBF( $\mathbf{H}$ ,  $\mathbf{y}$ )
  for  $t = 1, 2, \dots$  do
     $|y|_{\min, m} = \min_{n \in \mathcal{N}(m)} |y_n|$  ▷ Find minimum magnitude of  $y_n$  for each check in  $\mathbf{h}$ 
     $E_n = \sum_{m \in \mathcal{M}(n)} (2s_m - 1) |y|_{\min, m}$  ▷ For each bit calculate an energy function
     $p(N = n | \mathbf{z}) = \frac{1}{Z} \exp \left\{ \frac{1}{\beta} E_n \right\}$  ▷ Probability of that bit being flipped
     $n \sim p(n | \mathbf{z})$  ▷ Sampling from the probability distribution over bits
     $z_n = 1 - z_n$  ▷ Flip that bit
  if  $\mathbf{H}\mathbf{z} = \mathbf{0}$  then
    break ▷ Condition for a valid codeword
  end if
end for
end procedure

```

Figure 5. McWBF: an MC decoding algorithm for LDPC codes.

Before we introduce our procedure for calculating the probabilities of different bits being flipped and sampling from their distribution, let us define some notation. Vectors \mathbf{c} , \mathbf{x} , \mathbf{y} and \mathbf{z} are indices by n and vector \mathbf{s} by m . The element of matrix \mathbf{H} at index (m, n) is referred to as $h_{m,n}$. Also, it is important to keep in mind that row m of \mathbf{H} represents the m th check and all the non-zero bits in that row correspond to the bits in the encoded message that participate in that check. We denote this set of codeword bits participating in the m th check as the neighbourhood system $\mathcal{N}(m) := \{n : h_{m,n} = 1\}$. Similarly, all parity checks in which bit- n participates are denoted by the neighbourhood system $\mathcal{M}(n) := \{m : h_{m,n} = 1\}$.

Using this notation, an MC algorithm is presented in figure 5. We call this algorithm McWBF. In each iteration, the algorithm flips one bit in the received vector \mathbf{z} . To decide which bit to flip, it first computes probabilities for each bit based on an energy function E_n that depends on how many unsatisfied parity checks that particular bit appears in. It then computes a probability of flipping each bit based on the energy function and then samples from the probability distribution to choose a bit and flips it. The procedure used to calculate these probabilities results in a steady-state distribution over states that has significant peaks at the goal states. This, in turn, results in the recovery of a valid codeword.

(c) Sorting

Let us define S as the set of all permutations of the sequence $(1, \dots, n)$ and let $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{R}^n$ be a set of n elements unordered in magnitude. Sorting is simply the problem of finding a subset $S_0 \subseteq S$ such that $x_{c_i} \leq x_{c_{i+1}}$ for all $\mathbf{c} = (c_1, \dots, c_n) \in S_0$. The naive way of solving this problem is to iterate over the set \mathbf{c} n times, where in-place swaps $c_i \leftrightarrow c_{i+1}$ are performed if $x_{c_i} > x_{c_{i+1}}$. We can construct an MC algorithm much similar, by defining a distribution that assigns high probability to indexes pointing to unordered elements, and low probability to indexes pointing to ordered elements. However, instead of iterating, we draw index samples from the distribution and swap according to the samples. As the probability of unordered elements is high, the swapping thereof is more likely and the procedure should converge to a solution. For the sorting problem, one possibility is to construct a Gibbs distribution as follows:

$$p(i | \mathbf{c}) = \frac{1}{Z} \exp \left\{ -\frac{1}{\beta} \delta[x_{c_i} < x_{c_{i+1}}] \right\}, \quad (4.3)$$

where

$$Z = \sum_{i=1}^{n-1} \exp \left\{ -\frac{1}{\beta} \delta[x_{c_i} < x_{c_{i+1}}] \right\}$$


```

procedure MCSORT(x)
  initialize  $\mathbf{c}^{(0)}$ 
  for  $t = 1, 2, \dots$  do
     $i \sim p(i | \mathbf{c}^{(t)}, \mathbf{x})$  ▷ Draw sample from Eqn. 4.3
     $\mathbf{c}^{(t)} \leftarrow \mathbf{c}^{(t-1)}$ 
     $c_i^{(t)} \leftrightarrow c_{i+1}^{(t)}$  ▷ Swap indexes
  end for
end procedure

```

Figure 6. Sorting cast as MCMC. (Online version in colour.)

is the normalization factor, β is the temperature of the distribution and $\delta[a]$ produces a one if proposition a is true and zero otherwise. Given this distribution, indexes are chosen with high probability if they relate to unordered elements and with low probability if not. Given this distribution, it is straightforward to construct the MC using the procedure in figure 6.

(d) Clustering

In clustering, the objective is to group N observed data points (say x_i 's) into multiple clusters based on some similarity between the data points. One method for performing clustering is by using a mixture model which assumes that the overall data were generated from a mixture of several distributions with each cluster representing the subset of the data that originated from the same distribution. Each such distribution generally has the same form $F(\theta_k)$ (say Gaussian) with different *cluster parameters* θ_k . A particular distribution contributes to the overall distribution based on its weight π_k (also called its *mixing proportion*).

The *Dirichlet process mixture model* (DPMM; figure 7) is one such mixture model that assumes specific priors over the cluster parameters and the mixing proportions [20]. A detailed discussion of the model can be found in [21]. Here, we briefly present the details and characteristics of this model that are relevant to understanding its use in clustering.

DPMM assumes that each data point x_i has a corresponding hidden variable z_i that represents the cluster that generated x_i . Hence, z_i takes a value k (that corresponds to a cluster number) with probability π_k . The cluster parameters θ_k are given a common prior distribution $G(\lambda)$ with hyperparameter λ (equation (4.6)). The distribution $G(\lambda)$ is generally chosen to be the conjugate prior of the distribution $F(\theta_k)$. π (a vector of all π_k) is given a Griffiths–Engen–McClosky (GEM) prior $\pi \sim \text{GEM}(1, \alpha)$ (equation (4.4)) [22]. The conditional distributions in the model are presented below.

$$\pi | \alpha \sim \text{GEM}(1, \alpha), \quad (4.4)$$

$$z_i | \pi \sim \pi, \quad (4.5)$$

$$\theta_k | \lambda \sim G(\lambda) \quad (4.6)$$

and
$$x_i | z_i, \{\theta_k\}_{k=1}^{\infty} \sim F(\theta_{z_i}). \quad (4.7)$$

An interesting characteristic of the DPMM is that it allows the model to have an infinite number of clusters *a priori*. However, any finite observed dataset would only contain a finite, but random, number of clusters. Once the data are observed, the number of clusters is inferred from the data using the Bayesian posterior inference framework. This allows the complexity of the model to grow as new data are observed, allowing future data to map to previously unseen clusters. The expected number of clusters grows logarithmically with the size of the dataset.

For clustering, our objective is to construct an MC algorithm that, in addition to inferring the number of clusters, also infers the values of the hidden variables z_i corresponding to each data point x_i . We can think of the assignments of the z_i 's as a state. The goal state is an assignment that results in an acceptable clustering based on a clustering criterion such as m.s.e.




Figure 7. The DPMM.

A variety of inference methods based on Gibbs sampling (which is an MC sampling algorithm) have been proposed for inference in DPMM [23]. The collapsed Gibbs sampling algorithm (algorithm 3 in [23]) is suitable for our use of DPMM for clustering as we are only interested in knowing the cluster assignments (z_i 's) and not the actual cluster parameters (θ_k 's). It is an iterative algorithm that in each iteration updates the values of z_i for each data point one at a time. It does that by (i) removing x_i from its present cluster, (ii) computing the conditional probability of x_i belonging to each of the clusters present in that iteration and also to a potential new cluster, and (iii) samples from this distribution to obtain a cluster assignment for z_i . Thus, the computation in each iteration fits into our model of calculating a probability distribution and sampling from it (§3a).

As we used a dataset generated from Gaussian distributions in our evaluations, we assumed a Gaussian mixture model with $F(\theta_k)$ being a Gaussian distribution. In such a model, the conditional probabilities of x_i 's belonging to different clusters are easy to compute (details in [21]).

5. Methodology

In this section, we present the details of our methodology for evaluating the robustness of the MC implementations of the four applications discussed in §4. We first describe the exact implementations that were used for our evaluations (our code and data are available at <https://github.com/biplabdeka/markov>). Then we describe our fault model and the error injection methodology.

(a) Applications

For SAT, we compare two different implementations: McSAT as presented above, and an iterative version of DPLL which is a deterministic back-tracking algorithm for SAT [24]. DPLL, unlike McSAT, is a complete solver as it is able to convey to the user whether a given problem is unsatisfiable. Both SAT implementations were evaluated with randomly generated 3-CNF problem with 100 variables and 400 clauses. In addition, we further evaluated the robustness of McSAT using SAT-encoded problems from the following domains: graph colouring, planning and all interval series.

For sorting, we compare McSort as presented above with a baseline implementation of QuickSort. These were evaluated using a fully randomized sequence of integers in the range (1,1000). For LDPC decoding, we evaluated the robustness of the McWBF decoding algorithm. We used randomly generated messages that were encoded and decoded using a Gallager (273,82) LDPC code. We used DPMM for clustering and used a collapsed Gibbs sampler for inference. We used code for the collapsed Gibbs sampler that is available at <https://github.com/jacobeisenstein/DPMM>. The dataset used for clustering consisted of 200 two-dimensional data

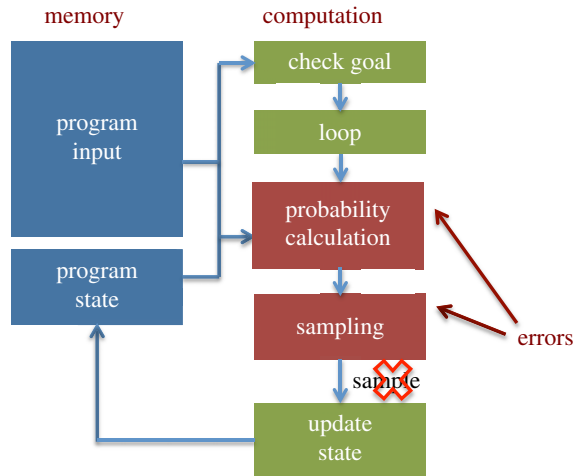


Figure 8. Fault model assumed in our evaluations. Faults that manifest as errors in the transition probability calculation and sampling in each iteration are considered. Memory and other calculations in every iteration are assumed to be fault free. (Online version in colour.)

points generated from a mixture of five Gaussian distributions. Accordingly, the DPMM model assumed the data to be a mixture of Gaussian distributions. The Gibbs sampling algorithm was run for 25 iterations and the m.s.e. of the clustering assignment was used as the quality metric.

(b) Fault model and error injection methodology

The fault model assumed in our evaluations is shown in figure 8. We consider faults that manifest themselves as errors in the transition probability calculation and sampling in each iteration. We assume that memory is fault-free and so are the other parts of the computation (loop iteration, termination checking and state update). This is a reasonable model to work with as the transition probability calculation and sampling constitute the bulk of the computation in each iteration for our applications. As such, other parts of the computation can be made robust using redundancy-based techniques.

We perform error injections at the algorithmic level based on the above model. The worst effect of a fault in the transition probability calculation or in sampling in an iteration is to produce a different sample. Hence, in our injection experiments, we modify the source code to corrupt the sample produced in each iteration with a given error rate. The corrupted sample is assigned a random sample in the state space.

In McSAT, the objective of each iteration is to find a suitable variable to flip. In our error injections, we randomly assign a variable to be flipped (instead of the one selected by the algorithm) at the end of an iteration with a particular error rate. To make as fair a comparison as possible, we inject the same type of error in DPLL. Similar to McSAT, in McWBF, the algorithm selects a bit to flip in the received signal vector in each iteration. We corrupt the result of a particular iteration by selecting a random bit to be flipped. The number of such iterations is again governed by the fault rate. For sorting, we inject errors by randomly choosing elements to swap at the end of an iteration with a particular error rate, instead of the choices made by the algorithms. For clustering, in each iteration, when new cluster assignments are being computed for each of the data points, we assign the data points to one of the existing clusters or to a new cluster randomly.

6. Results

The results of the algorithmic error injections for McSAT are shown in figure 9. McSAT is a randomized algorithm and as such it has a variable runtime even without errors. We show the

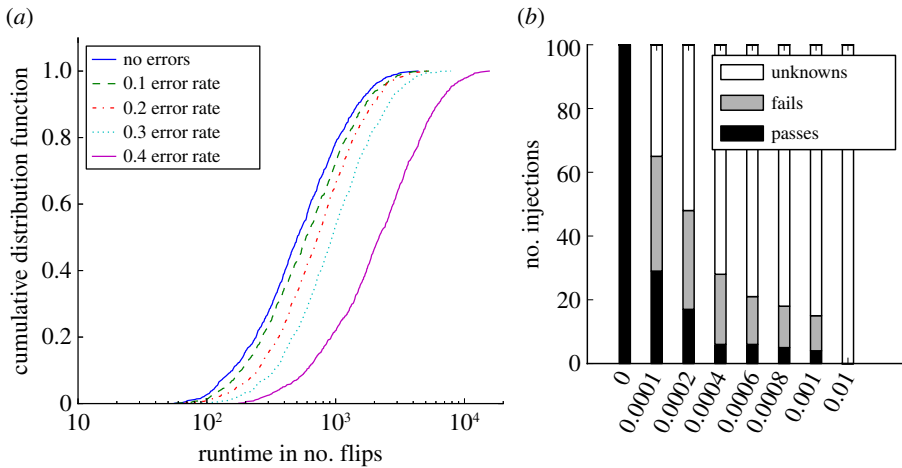


Figure 9. Algorithmic error injection results for SAT. (a) The plot for McSAT shows the runtime distribution for 1000 runs of the algorithm under different error rates as it always produces the correct output at these fault rates. (b) DPLL on the other hand, frequently fails in producing results and as such success rate is shown in the figure for DPLL. (Online version in colour.)




Figure 10. Robustness of McSAT across different inputs from SATLIB. Fault injection results for (a) AIS, (b) blocks world and (c) graph colouring. (Online version in colour.)

distribution of runtimes (in terms of the number of iterations) at different error rates. We observe that the algorithm produces acceptable runtimes even at fault rates as high as 10 or 20% (the median runtime is still within the 75th percentile mark for the error-free case). In addition, the results also exhibit a gradual degradation in runtime as error rates increase.

In comparison, the deterministic algorithm, DPLL, does not show similar robustness to errors. Figure 9 also presents the fraction of runs where DPLL produced correct outputs, incorrect outputs or simply crashed. While DPLL manages to produce correct results in some cases, it either returns an incorrect result (fails) or crashes (unknowns) in the majority of situations.

In order to verify that the observed robustness is not an artefact of a particular input, we repeated our experiments for McSAT with other inputs. We specifically consider SAT-encoded problems from the following domains: graph colouring, planning, all interval series and logistics. All input files were taken from the SATLIB benchmark suite [25]. Results presented in figure 10 show that McSAT displays significant robustness across all inputs.

The results of the algorithmic error injections for McSort are shown in figure 11. McSort also shows acceptable runtimes in the presence of errors with high error rates and shows gradual degradation of results as error rates increase. By contrast, the deterministic algorithm, QuickSort, does not always produce correct outputs in the presence of errors.




Figure 11. Algorithmic error injection results for sorting. (a) The plot for McSort shows the runtime distribution of the algorithm as it always produces the correct output at these error rates. (b) QuickSort on the other hand, frequently fails in producing results and as such success rate is shown in the figure for QuickSort. (Online version in colour.)

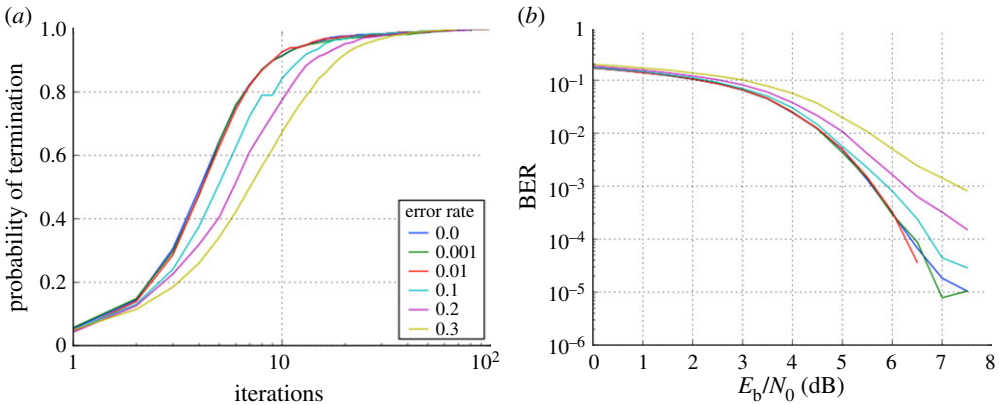


Figure 12. Algorithmic error injection results for LDPC decoding using McWBF algorithm. The runtime distribution is at a particular signal-to-noise ratio (at $E_b/N_0 = 5$ dB). The bit error rate (BER) performance shows the effect of errors on the output quality of the application (lower is better). (a) Runtime distribution and (b) BER performance. (Online version in colour.)

The results of the algorithmic error injections in LDPC decoding using McWBF are shown in figure 12. Similar to McSAT and McSort, LDPC decoding shows acceptable runtimes even in the presence of errors with high error rates and shows gradual degradation in runtime as error rates are increased. Figure 12b shows the bit error rate (BER) performance of the algorithm at different error rates. We observe that at error rates of 0.1 or 1%, the BER performance remains unchanged. At higher error rates, there is a gradual degradation in BER performance.

Figure 13 shows the effect of errors on clustering using a DPMM model and Gibbs sampling. At 0.1% error rate, the clustering is still performed correctly. However, at 1% error rate, we observe a few samples that are clustered incorrectly. Thus, errors have an effect on the output quality in this case. Figure 13d shows the m.s.e. of training data points from the inferred cluster centroids as number of iterations increase for different error rates. We observe that errors do not effect the runtime. In each case, the algorithm produces a cluster with the minimum m.s.e. possible for that particular error rate in about the same number of iterations. However, in terms of the minimum m.s.e. achieved (a measure of output quality), there is a gradual degradation with increasing error rates.

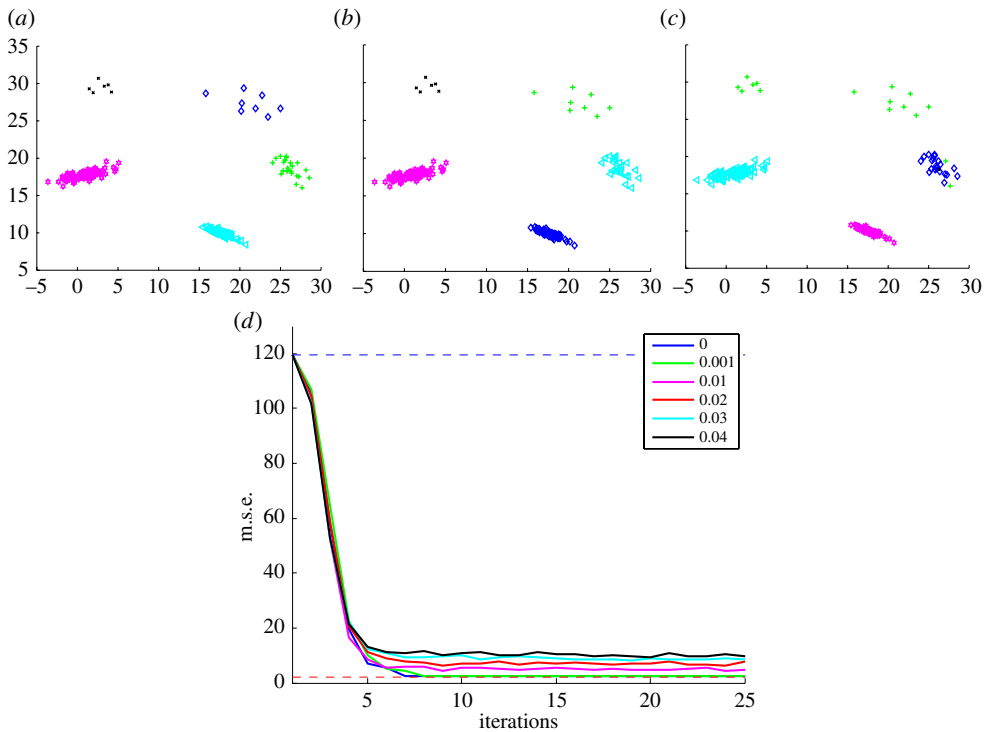


Figure 13. Algorithmic error injection results for clustering using a collapsed Gibbs sampler in a DPMM. (a) No errors, (b) 0.01% error rate, (c) 0.1% error rate and (d) m.s.e. of the training data points from the inferred cluster centroids for different error rates. Horizontal dashed lines represent the m.s.e. if all data points were assigned to a single cluster (top line) or to their correct clusters (bottom line). (Online version in colour.)

7. Discussion and future work

Our results presented in §6 show that MC algorithms can produce acceptable results even in the presence of errors at high error rates. In order to demonstrate the feasibility of using such algorithms as a template for building low-power systems using unreliable post CMOS devices, future work will focus on hardware implementations of MC algorithms for specific applications. These would include applications such as *neural spike sorting* [26] and *stereo matching* [27] for which MC sampling algorithms exist, but have not been implemented in hardware. Using fine-grained fault injection experiments, the robustness of these implementations to errors created by different hardware faults would be evaluated. In addition to robustness, such hardware implementations would also enable us to determine the potential for exploiting different *sampling level parallelism* strategies [4] for increased performance. Recent work has shown that systems built from intentionally stochastic digital hardware can indeed leverage significant parallelism inherent in Bayesian inference applications [28]. We believe such hardware building blocks could be used to implement MC algorithms and could also potentially exploit significant parallelism.

Hardware implementations would also enable us to determine the potential energy savings that can be achieved by different techniques that trade off robustness for energy benefits. This includes techniques such as function under-design and voltage overscaling. In order to appreciate how much energy savings can be achieved from such an approach, let us assume that we operate at a 14% error rate at the hardware level. In the worst case, where none of these faults get masked, this would lead to a 14% error rate at the algorithm level which we showed is acceptable for applications such as SAT and LDPC decoding. From the voltage-error model presented in [5], this translates to operation at 70% of the nominal voltage leading to operation at $0.7^2 = 0.49 \approx 50\%$

of the nominal energy. Energy savings may be possible in non-voltage-overscaling scenarios as well. For example, the high robustness of certain hardware blocks may mean that they could be implemented using low-power, inherently unreliable technologies. Similarly, it may be possible to implement several datapath blocks at reduced precision with minimal effect on the quality of output. Recent work on intentionally stochastic digital hardware also demonstrates that such hardware is robust to errors due to reduced precision of hardware [28]. We believe that MC algorithms implemented using such hardware building blocks could also exhibit additional robustness to errors due to hardware faults.

In the longer term, our research goals are to (i) evaluate the robustness of MC algorithms not just at the algorithmic level but also the robustness of their hardware and software implementations, (ii) cast applications that are generally not implemented as MC algorithms as MC-based implementations to achieve better robustness, and (iii) develop a general execution model and programmable platform that would allow the robust execution of a variety of MC-based implementations on the same low-power system. Another interesting area of research could be to establish a rigorous theoretical framework for understanding and evaluating the robustness of MC algorithms.

8. Conclusion

Based on the insight that MCs can be tolerant to errors in transition probabilities, we investigated building robust applications by implementing them as MC algorithms. Several applications already have MC implementations, whereas others can be converted to MCs. To demonstrate the process, we implemented four applications—SAT, LDPC decoding, sorting and clustering—as MCs and evaluated their robustness using algorithmic fault injections. We demonstrated that these applications, when cast as MCs, are significantly more robust than their deterministic implementation. In fact, these algorithms were evaluated at transition error rates as high as 10–20%. Even at such high error rates, MC algorithms produced acceptable results in terms of runtime and output quality. Also, these algorithms showed gradual degradation in runtime or output quality with increasing error rates. Based on these results and on the generality of MC algorithms, we make a case for using MC algorithms as a template for implementing applications on future robust, low-power hardware.

Data accessibility. A version of this work appeared at the Asilomar Conference on Signals, Systems and Computers, 2013.

Funding statement. This work was supported in part by Systems on Nanoscale Information Fabrics, one of the six SRC STARnet Centres, sponsored by MARCO and DARPA.

References

1. Das S, Tokunaga C, Pant S, Ma W-H, Kalaiselvan S, Lai K, Bull DM, Blaauw DT. 2009 Razor II: in situ error detection and correction for PVT and SER tolerance. *IEEE J. Solid-State Circuits* **44**, 32–48. (doi:10.1109/JSSC.2008.2007145)
2. Shanbhag N, Mitra S, De Veciana G, Orshansky M, Marculescu R, Roychowdhury J, Jones D, Rabaey J. 2008 The search for alternative computational paradigms. *IEEE Des. Test Comput.* **25**, 334–343. (doi:10.1109/MDT.2008.113)
3. Andrieu C, De Freitas N, Doucet A, Jordan MI. 2003 An introduction to MCMC for machine learning. *Mach. Learn.* **50**, 5–43. (doi:10.1023/A:1020281327116)
4. Williamson S, Dubey A, Xing EP. 2013 Parallel Markov chain Monte Carlo for nonparametric mixture models. In *Proc. 30th Int. Conf. on Machine Learning (ICML-13), Atlanta, GA, USA, 16–21 June 2013*, pp. 98–106.
5. Sloan J, Kesler D, Kumar R, Rahimi A. 2010 A numerical optimization-based methodology for application robustification: transforming applications for error tolerance. In *Proc. IEEE/IFIP Int. Conf. on Dependable Systems and Networks (DSN)*, pp. 161–170. (doi:10.1109/DSN.2010.5544923)

6. Huang K-H, Abraham J. 1984 Algorithm-based fault tolerance for matrix operations. *IEEE Trans. Comput.* **C-33**, 518–528. (doi:10.1109/TC.1984.1676475)
7. Sloan J, Bronevetsky G, Kumar R. 2013 An algorithmic approach to error localization and partial recomputation for low-overhead fault tolerance on parallel systems. In *Proc. 43rd Annual IEEE/IFIP Int. Conf. on Dependable Systems and Networks (DSN)*. pp. 1–12. (doi:10.1109/DSN.2013.6575309)
8. Mansinghka V. 2009 Natively probabilistic computation. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, USA.
9. Stroock DW. 2005 *An introduction to Markov processes*. Graduate Texts in Mathematics, vol. 230. Berlin, Germany: Springer. (doi:10.1007/b138428)
10. Besag J. 2000 *Markov chain Monte Carlo for statistical inference*. Technical report. Seattle, WA: Center for Statistics and the Social Sciences, University of Washington.
11. Kirkpatrick S, Gelatt CD, Vecchi MP. 1983 Optimization by simulated annealing. *Science* **220**, 671–680. (doi:10.1126/science.220.4598.671)
12. Morris B, Sinclair A. 1999 Random walks on truncated cubes and sampling 0-1 knapsack solutions. In *Proc. 40th IEEE Symp. on Foundations of Computer Science*, pp. 230–240. (doi:10.1109/SFFCS.1999.814595)
13. Jerrum M, Sinclair A, Vigoda E. 2004 A polynomial-time approximation algorithm for the permanent of a matrix with non-negative entries. *J. ACM* **51**, 671–697. (doi:10.1145/1008731.1008738)
14. Hoos HH. 1998 Stochastic local search: methods, models, applications. PhD thesis, Darmstadt Technische Universität, Darmstadt, Germany.
15. Garey MR, Johnson DS. 1979 *Computers and intractability: a guide to the theory of NP-completeness*. New York, NY: W. H. Freeman and Co.
16. Gu J, Purdom PW, Franco J, Wah BW. 1996 Algorithms for the satisfiability (sat) problem: a survey. In *DIMACS series in discrete mathematics and theoretical computer science*, pp. 19–152. Providence, RI: American Mathematical Society.
17. Selman B, Kautz H, Cohen B. 1995 Local search strategies for satisfiability testing. In *DIMACS series in discrete mathematics and theoretical computer science*, pp. 521–532. Providence, RI: American Mathematical Society.
18. Kautz H, Selman B. 1996 Pushing the envelope: planning, propositional logic, and stochastic search. In *Proc. Natl Conf. on Artificial Intelligence*, pp. 1194–1201. Palo Alto, CA: AAAI Press.
19. Zhang J, Fossorier M. 2004 A modified weighted bit-flipping decoding of low-density parity-check codes. *IEEE Commun. Lett.* **8**, 165–167. (doi:10.1109/LCOMM.2004.825737)
20. Rasmussen CE. 2000 The infinite Gaussian mixture model. In *Advances in neural information processing systems 12*, pp. 554–560. Cambridge, MA: MIT Press.
21. Sudderth EB. 2006 Graphical models for visual object recognition and tracking. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, USA.
22. Pitman J. 2002 *Combinatorial stochastic processes*. Technical report 621. Berkeley, CA: Department of Statistics, University of California.
23. Neal RM. 2000 Markov chain sampling methods for Dirichlet process mixture models. *J. Comput. Graph. Stat.* **9**, 249–265. (doi:10.1080/10618600.2000.10474879)
24. Davis M, Logemann G, Loveland D. 1962 A machine program for theorem-proving. *Commun. ACM* **5**, 394–397. (doi:10.1145/368273.368557)
25. Hoos HH, Stutzle T. 2000 Satlib: an online resource for research on sat. See <http://www.satlib.org/>.
26. Wood F, Black MJ. 2008 A nonparametric Bayesian alternative to spike sorting. *J. Neurosci. Methods* **173**, 1–12. (doi:10.1016/j.jneumeth.2008.04.030)
27. Barbu A, Zhu S-C. 2005 Generalizing Swendsen–Wang to sampling arbitrary posterior probabilities. *IEEE Trans. Pattern Anal. Mach. Intell.* **27**, 1239–1253. (doi:10.1109/TPAMI.2005.161)
28. Mansinghka V, Jonas E. 2014 Building fast Bayesian computing machines out of intentionally stochastic, digital parts. (<http://arxiv.org/abs/1402.4914>)