Several problems in applied mathematics and statistics require integrating a function f over a high-dimensional domain. For example, estimating the partition function of a graphical model for a fixed set of parameters requires integrating (summing) its unnormalized probability function f over all possible configurations (value assignments). The most common methods for performing such integration stochastically involve Markov Chains (MCMC).

We present an alternative to MCMC based on ideas from the theory of modern error-correcting codes. Specifically, we will see that stochastically integrating a non-negative function f over its entire domain can be achieved at the cost of maximizing f over suitable random subsets of its domain. The key lies in choosing these random subsets so that besides conferring good statistical properties, they also do not dramatically increase the difficulty of maximization. Using real-life satisfiability formulas as benchmarks, we see that selecting as subsets the codewords of Low Density Parity Check codes yields dramatic speedup and levels of accuracy that were previously unattainable.