

Counting-based Reliability Estimation for Power-Transmission Grids*

Leonardo Duenas-Osorio

Department of Civil and Environmental Engineering
Rice University

Roger Paredes

Department of Civil and Environmental Engineering
Rice University

Kuldeep S. Meel

Department of Computer Science
Rice University

Moshe Y. Vardi

Department of Computer Science
Rice University

Abstract

Modern society is increasingly reliant on the functionality of infrastructure facilities and utility services. Consequently, there has been surge of interest in the problem of quantification of system reliability, which is known to be #P-complete. Reliability also contributes to the resilience of systems, so as to effectively make them bounce back after contingencies. Despite diverse progress, most techniques to estimate system reliability and resilience remain computationally expensive. In this paper, we investigate how recent advances in hashing-based approaches to counting can be exploited to improve computational techniques for system reliability. The primary contribution of this paper is a novel framework, RelNet, that reduces the problem of computing reliability for a given network to counting the number of satisfying assignments of a Σ_1^1 formula, which is amenable to recent hashing-based techniques developed for counting satisfying assignments of SAT formula. We then apply RelNet to ten real world power-transmission grids across different cities in the U.S. and are able to obtain, to the best of our knowledge, the first theoretically sound a priori estimates of reliability between several pairs of nodes of interest. Such estimates will help managing uncertainty and support rational decision making for community resilience.

1 Introduction

Modern society is increasingly reliant on the availability of critical facilities and utility services, such as power, telecommunications, water, gas, and transportation among others (The White House, Office of the Press Secretary, 2016). To ensure adequate service, it is imperative to quantify system reliability, or the probability of the system to remain functional, as well as system resilience, or the ability of the system to quickly return to normalcy when failure is unavoidable (Bruneau et al. 2003). While resilience assessment requires human decision making principles, it also heavily depends on intrinsic system reliability. Hence, the recent focus on community resilience and sustainability has spurred significant activity in engineering reliability (Zio 2009).

*The author list has been sorted alphabetically by last name; this should not be used to determine the extent of authors' contributions.

Copyright © 2017, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

One of the key challenging problems in the area of engineering reliability is network reliability, wherein the input to the problem consists of a network, represented as a graph, arising out of distribution of water, power, transportation routes and the like. The problem of the network reliability seeks to measure the likelihood of two points of interest being reachable under conditions such as natural disasters. Early theoretical investigations showed that the problem of network reliability is #P complete (Valiant 1979). Although graph contraction strategies combined with DNF counting provide a Fully Polynomial Randomized Approximation Scheme (FPRAS) with error guarantees (Karger 2001), implementation on practical systems does not scale well due to the requirement of a large number of Monte Carlo steps. Consequently, recent investigations have focused on advancing algorithmic strategies that build upon advanced Monte Carlo simulation (Zuev, Wu, and Beck 2015a) and analytical approaches (Lim and Song 2012; Dueñas-Osorio and Rojo 2011). In addition, inventive sampling methods, such as line sampling and variance reduction schemes (Fishman 1986), along with graphical models, especially Bayesian networks, provide versatile strategies to quantify the reliability of complex engineered systems and their dynamics (Bensi, Kiureghian, and Straub 2013).

Despite significant progress, most techniques remain computationally expensive. As an alternative, when invoking approximations, most methods are unable to guarantee the quality of the reliability estimation a priori, barring small instances where exact methods do not time out. Therefore, design of techniques that offer strong theoretical guarantees on the quality of estimates and can scale to large real world instances remains an unattained goal across multiple disciplines.

A promising alternative approach to answer #P queries is to reduce a #P problem to a #SAT problem, where #SAT denotes the problem of computing the number of solutions for a given SAT formula. While computing exact answers to #SAT is known to be computationally hard, recent advances in universal hashing-based approaches to #SAT have received wide interest (Gomes, Sabharwal, and Selman 2007; Chakraborty, Meel, and Vardi 2013; Ivrii et al. 2015). In particular, these techniques can provide probably approximately correct (PAC) guarantees, i.e., they compute an estimate that is within a prescribed error tolerance of ϵ with

confidence at least $1 - \delta$ for a given δ . Furthermore, the expected value of the estimate is equal to the exact count of the number of solutions. The success of recent hashing-based techniques largely stems from the usage of SAT solvers as NP oracles and efficient encoding of universal hash functions as XOR constraints. The recent breakthrough in the reduction of the number of NP oracle calls from linear to logarithmic (in the number of variables of SAT formula) further highlights the promise of the approach (Chakraborty, Meel, and Vardi 2016). This motivates us to ask: *Can we design a counting-based framework that can take advantage of progress in hashing-based techniques to provide theoretically sound estimates for the network reliability problem?*

The primary contribution of this paper is a positive answer to the above question. We present a counting-based framework, called RelNet, that reduces the problem of computing reliability for a given network to counting the number of satisfying assignments of a Σ_1^1 formula, which is amenable to recent hashing-based techniques developed for counting satisfying assignments of SAT formula. RelNet significantly outperforms state of the art techniques and in particular, allowed us to obtain the first theoretically sound estimates of reliability for ten networks representing different cities in the U.S.

The remainder of the paper is organized as follows. We introduce notations and preliminaries in Section 2. We then discuss related work in Section 3 and discuss datasets of interest in Section 4. In Section 5, we present the primary contribution of this paper: RelNet. We then discuss the experimental methodology along with details of real world data sets, to then present experimental results in Section 6. Finally, we conclude in Section 7.

2 Preliminaries

We write $\mathcal{P}[Y : \Omega]$ to denote the probability of outcome Y when sampling from a probability space Ω . For brevity, we omit Ω when it is clear from the context. The expected value of Y is denoted $\mathbb{E}[Y]$. For a set A , \bar{A} denotes the complement of the set A .

Let $G = (V, E)$ be a graph, where V is set of the vertices, also referred as nodes, and E is set of edges. For every edge $e \in E$ from u to v , we define $\text{start}(e) = u$ and $\text{end}(e) = v$. Note that we allow multiple edges between pairs of nodes.

We say that $\pi = (u, w_1, \dots, w_{k-1}, v)$ is a path of length k that connects u and v if $\forall i < k - 1, w_i \in V$ and $\exists e (u = \text{start}(e) \wedge w_1 = \text{end}(e)) \wedge \exists e (w_{k-1} = \text{start}(e) \wedge v = \text{end}(e)) \wedge \forall i < k - 2, \exists e (w_i = \text{start}(e) \wedge w_{i+1} = \text{end}(e))$. We use T_π to denote set of all edges in π . For every subset $\sigma \subseteq E$, we say u and v are connected under σ , denoted by $(u, v) \models \sigma$, if $\exists \pi, k$ such that π is a path of length k that connects u and v and $T_\pi \subseteq \sigma$. For a given graph G , we use $\Gamma_{G,u,v}$ to denote the set of all subsets σ of E that make u and v connected, i.e $\Gamma_{G,u,v} = \{\sigma \subseteq E \mid (u, v) \models \sigma\}$.

For a given graph $G = (V, E)$ and nodes u and v , we use $e(u, v) \cup G$ to denote the augmented graph G' obtained by putting an edge e such that $u = \text{start}(e)$ and $v = \text{end}(e)$. Note that if G has i edges from u to v , then G' has $i+1$ edges from u to v . In this paper, we focus on probabilistic variant of graphs, where probability function is associated

to edges in E . For every edge $e \in E$, we use e^1 to denote the event that edge e does not fail and e^0 to denote the event that edge e fails. We have $\mathcal{P}[e^0] + \mathcal{P}[e^1] = 1$. As discussed in Section 4, the failure of edge corresponds to event in real life when an existing edge is broken due to events such as natural disasters. We assume all e_i^1 to be independent. Without loss of generality, the least significant bit in the representation of $\mathcal{P}[e_i^1]$ is always taken to be 1. We call a graph as unweighted if for all edges $e \in E$, we have $\mathcal{P}[e^0] = 1/2$, otherwise the graph is called weighted. Therefore for $\sigma \subseteq E$, $\mathcal{P}[\sigma] = \prod_{e_i \in \sigma} \mathcal{P}(e_i^1) \times \prod_{e_j \notin \sigma} \mathcal{P}(e_j^0)$. Furthermore, we have $\mathcal{P}[\Gamma_{G,u,v}] = \sum_{\sigma \in \Gamma_{G,u,v}} \mathcal{P}[\sigma]$. For a given graph G , source node u and sink node v , the reliability of $u \rightarrow v$ is defined as $\mathcal{P}[\Gamma_{G,u,v}]$. In this paper, we consider the problem of estimating $r(u, v) = 1 - \mathcal{P}[\Gamma_{G,u,v}]$

We say $F(X)$ is a Σ_1^1 formula if $F(X)$ can be expressed as $(\exists S)\phi(S, X)$, where ϕ is defined over variables in $S \cup X$ and is represented in conjunctive normal form (CNF). Let $\text{Vars}(\phi)$ (resp. $\text{Vars}(F)$) be the set of variables appearing in ϕ (resp. F). An assignment τ of truth values to the variables in $X \cup S$ is called a *satisfying assignment* or *witness* of ϕ if τ makes ϕ evaluate to true. Similarly, an assignment σ of truth values to variables in X is satisfying assignment of F if \exists assignment ρ to variables in S such that $\rho \cup \sigma$ is a satisfying assignment of ϕ . We denote the set of all witnesses of F (resp. ϕ) by R_F (resp. R_ϕ). Given a set of variables $T \subseteq \text{Vars}(\phi)$, we use $R_{\phi \downarrow T}$ to denote the projection of R_ϕ on T . Note that for F and ϕ as defined above, we have $R_F = R_{\phi \downarrow X}$.

The *constrained counting problem* for Σ_1^1 is to compute $|R_F|$ for a given Σ_1^1 formula F . A *probably approximately correct* (or PAC) counter is a probabilistic algorithm $\text{ApproxCount}(\cdot, \cdot, \cdot)$ that takes as inputs a formula F , a sampling set S , a tolerance $\varepsilon > 0$, and a confidence $1 - \delta \in (0, 1]$, and returns a count c such that $\mathcal{P}\left[|R_F|/(1 + \varepsilon) \leq c \leq (1 + \varepsilon)|R_F|\right] \geq 1 - \delta$. The probabilistic guarantee provided by a PAC counter is also called an (ε, δ) guarantee.

Our work uses a special class of graphs, called chain graphs, which are inspired from the work on chain formulas (Chakraborty et al. 2015b). Similar to every edge, every chain graph has start and end node defined as follows. Every edge e is a chain graph, say G , such that $u = \text{start}(G)$ if $u = \text{start}(e)$ and $v = \text{end}(G)$ if $v = \text{end}(e)$, and we represented G as $G := (u \vee v)$. In addition, if $G = (V, E)$ is a chain graph and e is an edge such that (i) $u = \text{start}(e) = \text{start}(G) \in V$ and $v = \text{end}(e) = \text{end}(G) \in V$, we say that $e \cup G$ is a chain formula, represented by $(u \vee G)$ or (ii) $u = \text{start}(e) \notin V$ and $v = \text{end}(e) = \text{start}(G) \in V$, then $e \cup G$ is a chain formula, represented by $(u \wedge G)$. Every chain graph G over nodes a_1, a_2, \dots, a_m and n edges can be represented as $(b_1 C_1 (b_2 C_2 (\dots (b_n C_n b_{n+1}) \dots)))$, where $C_i = \vee$ or \wedge and performing a many to one mapping from $\{b_1, \dots, b_{n+1}\}$ to $\{a_1, a_2, \dots, a_m\}$ such that (i) $b_1 \mapsto a_1 \wedge b_{n+1} \mapsto a_m$, and (ii) $\forall i < m - 1, b_i \mapsto a_j \wedge b_{i+1} \mapsto a_l \rightarrow j < l$ if $C_i = \wedge$ and $j = l$, otherwise.

3 Related Work

Prior Work The problem of computing $r(u, v)$ for a given graph G was shown to be #P-complete by Valiant (1979). Consequently, there has been focus on development of approximate techniques for $r(u, v)$. In his seminal paper, Karger (2001) provided the first Fully Polynomial Randomized Approximation Scheme (FPRAS) such that returned estimate satisfies (ϵ, δ) guarantees while the runtime of algorithm (referred as Karger’s algorithm in rest of the paper) is polynomial in the $|G|, \log(1/\delta), 1/\epsilon$. Our experiments demonstrate that the high requirement of Monte Carlo samples in the above algorithm is a major bottleneck to efficiency and for our benchmarks, Karger’s algorithm times out.

The recent investigations into network reliability have focused on advancing algorithmic strategies that build upon advanced Monte Carlo simulation (Zuev, Wu, and Beck 2015a) and analytical approaches (Lim and Song 2012; Dueñas-Osorio and Rojo 2011). In particular, statistical learning techniques when combined with numerical simulation afford the reliability assessment of complex engineered systems, while unraveling component importance and sensitivities (Hurtado 2013). Also, successful strategies in data science, such as hierarchical clustering, provide novel tools for reliability and risk assessment (Yin and Kareem 2016; Gómez, Sánchez-Silva, and Dueñas-Osorio 2011). Also, state space partition strategies and optimization allow for analytical modeling of system reliability, which also offers, as a by-product, insights on the geometry of the failure space (Alexopoulos 1997; Dotson and Gobien 1979). Classical universal generating functions but combined with optimization also offer fresh alternatives to quantify system reliability approximately (Chang and Mori 2013). Besides, inventive sampling methods, such as line sampling and variance reduction schemes (Fishman 1986), along with graphical models, especially Bayesian networks, provide versatile strategies to quantify the reliability of complex engineered systems and their dynamics (Bensi, Kiureghian, and Straub 2013).

With the advent of resilience engineering, analytical methods are highly regarded in engineering reliability as they provide accurate estimates or, in more challenging instances, they yield lower and upper bounded estimates with 100% confidence. Furthermore, we can classify analytical network reliability methods in two groups based on their algorithmic approach. The first uses prior enumeration of cut sets (or path sets) or boolean algebra to account for non-disjoint events (Aggarwal, Misra, and Gupta 1975; Abraham 1979), whereas the latter uses recursive or iterative decompositions of disjoint events (Dotson and Gobien 1979; Rai and Kumar 1987; Page and Perry 1988). The latter group has proven more practical due to its online decomposition capabilities while not relying on the prior cut (or path) set enumeration and applications of the inclusion-exclusion principle, both NP-hard problems. In particular, research that builds upon the work by Dotson et al. has found wide technical application for medium-size networks (Li and He 2002; Lim and Song 2012) and in this paper we use the Selective path based Recursive Decomposition Algorithm (S-

RDA) as a representative approach of state-of-the-art analytical reliability methods for civil infrastructure systems. Herein, we refer to the gap between upper and lower bound estimates of reliability as the gap error. S-RDA aims at shrinking the gap error as much as possible by finding disjoint path sets that contain the shortest path of maximum likelihood at every decomposition step while prioritizing partitioning subsets of larger likelihood as well allowing it to provide anytime approximation guarantee.

Approximate Counting Complexity theoretic studies of propositional model counting were initiated by Valiant, who showed that the problem is #P-complete (Valiant 1979). Despite advances in exact model counting over the years (Thurley 2006), the inherent complexity of the problem poses significant hurdles to scaling exact counting to large problem instances. The study of approximate model counting has therefore been an important topic of research for several decades.

In (Chakraborty, Meel, and Vardi 2013), a new hashing-based probably approximately correct counting algorithm, called ApproxMC, was shown to scale to formulas with hundreds of thousands of variables, while providing rigorous PAC-style (ϵ, δ) guarantees. The core idea of ApproxMC is to use 2-universal hash functions to randomly partition the solution space of the original formula into “small” enough cells. The sizes of sufficiently many randomly chosen cells are then determined using calls to a specialized SAT solver (CryptoMiniSAT (Soos, Nohl, and Castelluccia 2009)), and a scaled median of these sizes is used to estimate the desired model count. Overall, ApproxMC makes a total of $\mathcal{O}(\frac{n \log(1/\delta)}{\epsilon^2})$ calls to CryptoMiniSAT. The works of (Ermon et al. 2013; Chakraborty et al. 2014; 2015a; 2016) have subsequently extended the ApproxMC approach to finite domain discrete integration. Recent breakthrough by Chakraborty et al (Chakraborty, Meel, and Vardi 2016) require only $\mathcal{O}(\frac{\log n \log(1/\delta)}{\epsilon^2})$ calls to SAT solver. Furthermore, hashing-based techniques, in particular ApproxMC2, have been shown to handle counting over Σ_1^1 formulas as well (Chakraborty, Meel, and Vardi 2016). In this context, we are motivated by success of hashing-based techniques and we demonstrate how we can employ hashing-based counters to design scalable reliability estimation techniques.

4 Datasets

In this paper we use as benchmark 10 power transmission networks powering small to medium size cities in the states of Texas (TX), Florida (FL), California (CA), Tennessee (TN), Georgia (GA), and South Carolina(SC). Such states are susceptible to extreme natural disasters such as flooding, hurricanes, or earthquakes. These cities have populations in the order of tens to hundreds of thousands and the grids connect generators and substations with 110-765 kV transmission-level power lines. Also, as shown in Table 4, networks’ size go from 47 to 112 nodes and the number of edges are of the same order. The raw network data was obtained in GIS format from the “Platts” repository for maps

Index	City Name	$ V $	$ E $
G1	Amarillo, TX	47	62
G2	Lakeland, FL	50	69
G3	El Paso, TX	52	65
G4	San Luis Obispo, CA	57	69
G5	Eureka, CA	61	70
G6	Bulls Gap, TN	62	91
G8	Memphis, TN	66	83
G12	Lubbock, TX	85	106
G22	Athens, GA	103	116
G27	Sumter, SC	112	139

Table 1: Test power networks.

and geospatial data ¹.

Transmission line outages due to random failures are not uncommon in power transmission systems during regular operation. The annualized probability of such failures depends on technical characteristics such as length of lines, supply/demand, temperature, etc. Typical values for ten-hour line outages, based on their annual occurrence rate, range from 60% to 98% for lines of length 50 and 200 kilometers respectively (Billinton and Li 1994). Although these values may appear high, such contingencies can be managed relatively easily. In contrast, extensive and complete damage due to natural disasters have smaller occurrence probabilities but are much more difficult to manage due to increased time of repairs. Even though the likelihood of such extreme natural events is small, conditioned on their occurrence, the probability of failures with significant damage for power transmission lines and facilities can be much larger as is typically depicted in fragility curves that encode probabilities of failure conditioned on some hazard intensity level (Fig. 1). For our experiments, we consider failure probability of 0.125 – a value that is attainable in practice by wide range of extreme natural events.

5 From Network Reliability to Σ_1^1 Counting

In this section, we discuss how the problem of computing reliability can be reduced to counting over Σ_1^1 formulas. In this section, we first discuss how weighted graphs can be reduced to unweighted graphs. We then discuss how the problem of computing reliability for an unweighted graph can be reduced to counting the number of satisfying assignments of a Σ_1^1 formula. We then discuss our proposed framework, RelNet, that combines the two reductions and employs hashing-based techniques to compute reliability for arbitrary graphs.

5.1 From Weighted to Unweighted Graph

The central idea of our reduction is usage of chain graphs to represent weights, which is closely related to usage of chain formulas for weighted counting (Chakraborty et al. 2015b). Let $m > 0$ be a natural number, and $k < 2^m$ be a positive odd number. Let $c_1 c_2 \dots c_m$ be the m -bit binary representation of k , where c_m is the least significant bit. Let z be the

¹<http://www.platts.com/products/gis-data>.

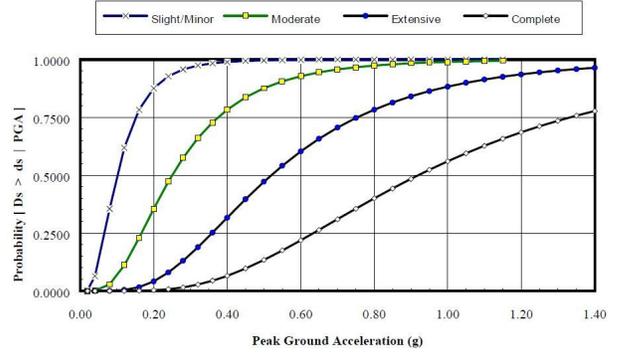


Figure 1: Probability of exceeding a given damage state (DS) for Medium/Large Generation Facilities with Anchored Components as a function of the peak ground acceleration intensity after an earthquake. Different levels of DS are indicated on top. Taken from HAZUS’ technical manual (HAZUS 2003).

number of zeros in the representation of k . Define

$$\psi_{k,m}(b_1, \dots, b_{m+1}) = (b_1 C_1 (b_2 C_2 \dots (b_m C_m b_{m+1}) \dots))$$

where $C_i = \vee$ if $c_i = 1$ and \wedge otherwise. We now construct chain graph $\phi_{k,m}(a_1, \dots, a_{z+2})$ by performing a many to one mapping between $\{b_1, \dots, b_{m+1}\}$ and $\{a_1, a_2, \dots, a_{z+2}\}$ such that (i) $b_1 \mapsto a_1 \wedge b_{m+1} \mapsto a_{z+2}$, and (ii) $\forall i < m - 1, (b_i \mapsto a_j \wedge b_{i+1} \mapsto a_l) \rightarrow j < l$ if $C_i = \wedge$ and $j = l$, otherwise. Note that there is one to one correspondence between $\psi_{k,m}(b_1, \dots, b_{m+1})$ and $\phi_{k,m}(a_1, \dots, a_{z+2})$. For example, consider $k = 3$ and $m = 3$. The binary representation of 3 using 3 bits is 011 and $z = 1$. Therefore, we have $\psi_{3,3}(b_1, b_2, b_3, b_4) = (b_1 \wedge (b_2 \vee (b_3 \vee b_4)))$, which gives us $\varphi_{3,3}(a_1, a_2, a_3) = (a_1 \wedge (a_2 \vee (a_2 \vee a_3)))$. We now first show that $|\varphi_{k,m}|$ is of linear size and then discuss the relationship between k, m and $\Gamma_{G, a_1, a_{z+2}}$.

Lemma 1 *Let $m > 0$ be a natural number, $k < 2^m$, z and $\varphi_{k,m}$ as defined above. Then $|\varphi_{k,m}|$ is linear in m . Furthermore $|\Gamma_{\varphi_{k,m}, a_1, a_{z+2}}| = k$*

Proof By construction, $\varphi_{k,m}(a_1, \dots, a_{z+2})$ is of size linear in m . To prove that $|\Gamma_{\varphi_{k,m}, a_1, a_{z+2}}|$ is of exactly size k , we use induction on m . We apply induction on $\psi_{k,m}$ since $\psi_{k,m}$ and $\varphi_{k,m}$ have 1-1 correspondence. The base case ($m = 1$) is trivial. For $m \geq 1$, let $c_2 \dots c_m$ represent the number k' in binary, and assume that $\psi_{k',m-1} = (b_2 \dots C_m b_{m+1}) \dots$ has corresponding chain graph $\varphi_{k',m-1}$ such that $|\Gamma_{\varphi_{k',m-1}, u, v}| = k'$, where $u = \text{start}(\varphi_{k',m-1})$ and $v = \text{end}(\varphi_{k',m-1})$. If c_1 is 0, then on one hand $k = k'$, and on the other hand we have, $\varphi_{k,m} \equiv e \cup \varphi_{k',m-1}$, where $a_1 = \text{start}(e)$, $\text{end}(e) = \text{start}(\varphi_{k',m-1})$ which has $|\Gamma_{\varphi_{k,m}, a_1, a_{z+2}}| = k$. Otherwise, if c_1 is 1, then on one hand $k = 2^{m-1} + k'$, and on the other hand C_1 is the connector “ \vee ”. Therefore, $\varphi_{k,m} \equiv e \cup \varphi_{k',m-1}$ where $a_1 = \text{start}(e)$, $\text{end}(e) = \text{end}(\varphi_{k',m-1})$, which has $|\Gamma_{\varphi_{k,m}, a_1, a_{z+2}}| = 2^{m-1} + k' = k$. This completes the induction.

5.2 From Graphs to Σ_1^1 Formulas

In this section, we discuss how for a given graph $G = (V, E)$ and nodes u and v , and associated probability function such that $\mathcal{P}[e^1 | e \in E] = 1/2$, we can reduce the problem of computing $r(u, v)$ to the problem of computing $|R_F|$ wherein F is a Σ_1^1 formula.

The central idea of our reduction is based on usage of transitive closure for connectivity. Our reduction has close connection to previously proposed formulations for s-t connectivity (See (Clote and Setzer 1998) for related survey). Let $R(u, v)$ denote the event that \exists path π such that π connects u and v . If $R(u, v)$ occurs and there exists an edge $e \in E$, such that $v = \text{start}(e) \wedge w = \text{end}(e)$, then $R(u, w)$ must occur. For a given graph $G = (V, E)$ and pair of nodes u and v , the goal is to create a Σ_1^1 formula F such that every satisfying assignment to F has one to one correspondence with $\sigma \subseteq E$ such that u and v are not connected under σ . To this end, we define a propositional variables p_u and q_e for every node $u \in V$ and every edge $e \in E$ respectively. Define,

$$\begin{aligned} C_e &= (p_u \wedge q_e \rightarrow p_v) \\ S &= \{p_u | u \in V\} \\ F_{u,v} &= \exists S(p_u \wedge \neg p_v \wedge \bigwedge_{e \in E} C_e) \end{aligned}$$

Lemma 2 *For a given graph $G = (V, E)$ and nodes u and v , let $F_{u,v}$ be as defined above. Then, $|R_{F_{u,v}}| = |\overline{\Gamma_{G,u,v}}|$. Furthermore if $\forall e \in E$, we have $\mathcal{P}[e^1] = \frac{1}{2}$, then $r(u, v) = \frac{|R_{F_{u,v}}|}{2^{|E|}}$*

Proof We defer the proof to technical report (Dueñas-Osorio et al. 2017) for lack of space.

5.3 RelNet

We now describe how the above reductions can be employed to design a counting-based framework, called RelNet, for the problem of network reliability. For a given graph $G = (V, E)$, source node u and sink node v and a probability space Ω over the edges, RelNet consists of the following three steps:

- Step 1: We obtain a transformed graph G' by replacing every $e_i \in E$ with $\phi_{k,m}$ if $\mathcal{P}[e_i^1] = \frac{k_i}{2^{m_i}}$. Let $M = \sum_{e_i \in E} m_i$ where $\mathcal{P}[e_i^1] = \frac{k_i}{2^{m_i}}$.
- Step 2: Construct $F_{u,v}$ as described above for the transformed graph G' , source node u and sink node v
- Step 3: Invoke a hashing-based counting technique to estimate $|R_{F_{u,v}}|$

The following theorem proves the correctness of our framework

Theorem 3 *For a given Graph G , source node u and sink node v , and probability space Ω over the edges, $r(u, v) = \frac{|R_{F_{u,v}}|}{2^M}$*

Proof The proof follows directly from Lemmas 1 and 2.

6 Evaluation

Since the primary objective of this project was to compute connectivity reliability of power transmission grid networks across different cities in U.S., we compared the effectiveness of RelNet vis-a-vis state of the art techniques. Specifically, we sought to answer the following questions:

1. How does the runtime performance of RelNet compare to that of the state-of-the-art techniques on real world power transmission networks?
2. How do estimates computed by RelNet compare to the exact estimates of reliability for networks that could be handled by exact techniques?

6.1 Experimental Methodology

We sought to compute reliability between every pair of nodes for all the ten cities discussed in Section 4. We implemented a Python prototype of RelNet, which invokes ApproxMC2 to perform counting over Σ_1^1 formulas as required by Step 3 of the RelNet. For all our experiments, we used $\varepsilon = 0.8$ and $\delta = 0.2$ as parameters for ApproxMC2, which is consistent with previously reported studies of using hashing-based counting techniques.

For comparison purposes, we considered: (i) Karger's FPRAS algorithm (Karger 2001), (ii) a recently proposed MCMC-based technique (Zuev, Wu, and Beck 2015b) and (iii) selective path based RDA (S-RDA), one of the current state of the art techniques employed by the reliability engineering community. For all our benchmarks, S-RDA outperformed Karger's FPRAS algorithm and the above stated MCMC technique, therefore we omit further discussion of these two techniques in the rest of the section.

We used a high-performance cluster to conduct experiments in parallel. Each node of the cluster had a 12-core 2.83 GHz Intel Xeon processor, with 4GB of main memory, and each experiment was run on a single core. Each experiment consisted of running a given tool on a given graph for a pair of nodes termed as source and sink. The timeout for each experiment was set to 1,000 seconds.

6.2 Results

For lack of space, we present results only on a subset of experiments. We refer the reader to technical report (Dueñas-Osorio et al. 2017) for detailed experimental results.

The analysis of runtime performance of S-RDA and RelNet shows that RelNet dramatically outperforms S-RDA. First of all, RelNet can compute $r(u, v)$ for each pair of source (u) and sink (v) for all the ten cities while S-RDA could handle only G5 and G27 and timed out for almost every pair for rest of the cities. It is worth reiterating before RelNet, no theoretically sound estimates were, to the best of our knowledge, a priori available for rest of the eight cities. Figure 2 presents heat-maps for both S-RDA and RelNet for cities G1, G2, and G3. For every city G_i , the corresponding heatmap is labeled by either G_i (S-RDA) if it presents runtime results for S-RDA or G_i (RelNet), otherwise. For every heat-map, the y-axis represent source node while the x-axis represents sink node. For every pair of source and sink, the runtime for the corresponding tool is represented

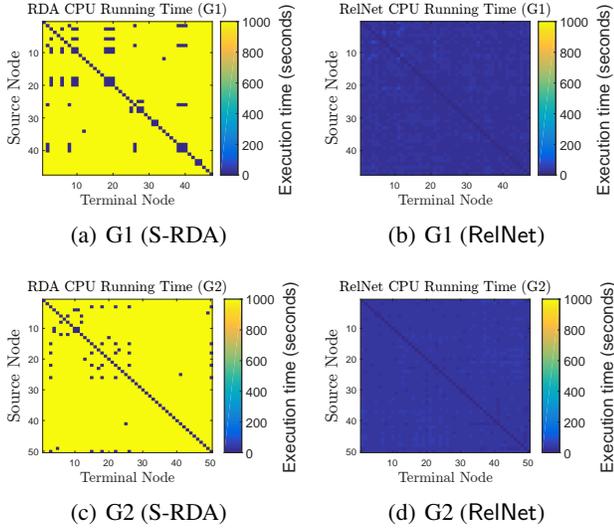


Figure 2: CPU time in seconds using RDA and RelNet for every source and sink pair

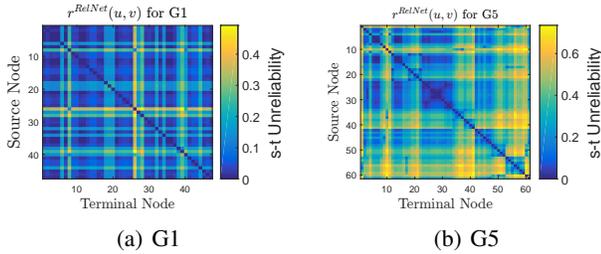


Figure 3: s-t Unreliability estimates for G1 and G5 using RelNet for every pair.

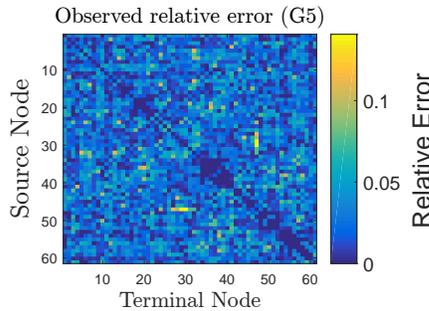


Figure 4: Observed tolerance (ϵ_{obs}) for all pairs of city G5

by the color as specified by the scale next to each heat-map. Overall, the closer the color of the point is to blue, better the method is.

The heat-maps clearly show that while RelNet can compute estimates within few tens of seconds for each pair, S-RDA fails for almost every pair. In this context, it is worth mentioning that runtime of RelNet is very consistent across different pairs of source and sink nodes.

As an illustration, Figure 3 shows heat-maps of reliability estimates between all pairs of nodes for cities G1 and G5 as computed by RelNet. Similar to performance comparison heatmaps, the y-axis of every plot refers to source node while the x-axis refers to sink node. The reliability for (u, v) is represented by the color as per the mapping presented on the right. Looking at these plots, one might wonder about the accuracy of reported results. While RelNet provides theoretical guarantees of accuracy, we sought to measure the quality of our estimates in practice. Given that S-RDA is an exact technique, we use the estimates from S-RDA on G5 to measure the quality of estimates of RelNet. For each pair, the observed tolerance ϵ_{obs} was calculated as $\max(\frac{C}{f} - 1, \frac{f}{C} - 1)$ where C is the estimate from RelNet and f is the exact estimate computed by S-RDA. Figure 4 shows the heat-map of observed tolerance ϵ_{obs} for each pair of G5. First of all, for every pair the observed tolerance is less than 0.14 – far better than the theoretical guarantee of 0.8. Furthermore, the geometric mean of observed tolerance is just 0.023; almost an order of magnitude better than the theoretical guarantee. This highlights conservative nature of theoretical guarantees and the need to strengthen the analysis as part of future work. As this work is part of larger project, where estimates of reliability are required to support decision making for community resilience, the above observations are quite significant as they show how emerging computational algorithms could support analysis and management of infrastructure under uncertainty.

7 Conclusion

Estimation of network reliability is crucial for decision making to ensure availability and resilience of critical facilities. Despite significant interest and long history of prior work, the current state of the art techniques fail to either provide sound theoretical estimates or scale to large networks. In this work, we propose a counting-based framework, RelNet, that provides strong theoretical guarantees of the PAC. Furthermore, unlike the current state of the art techniques, RelNet can scale to real world networks arising from cities across U.S., especially when exact reliability computations are not affordable.

Acknowledgements

We thank Bowen Fu, Dror Fried, Andres Gonzalez, and Jimmy Newmann for helpful discussions during this project. This work was supported in part by NSF grants IIS-1527668, CCF-1139011, CMMI-1436845 and Data Analysis and Visualization Cyberinfrastructure funded by NSF under grant OCI-0959097. Kuldeep S. Meel is supported by IBM Ph.D. Fellowship.

References

- Abraham, J. 1979. An Improved Algorithm for Network Reliability. *IEEE Transactions on Reliability* R-28(1):58–61.
- Aggarwal, K.; Misra, K.; and Gupta, J. 1975. A Fast Algorithm for Reliability Evaluation. *IEEE Transactions on Reliability* R-24(1):83–85.
- Alexopoulos, C. 1997. State space partitioning methods for stochastic shortest path problems. *Networks* 30(1):9–21.
- Bensi, M.; Kiureghian, A. D.; and Straub, D. 2013. Efficient Bayesian network modeling of systems. *Reliability Engineering & System Safety* 112:200–213.
- Billinton, R., and Li, W. 1994. *Reliability Assessment of Electric Power Systems Using Monte Carlo Methods*.
- Bruneau, M.; Chang, S. E.; Eguchi, R. T.; Lee, G. C.; O'Rourke, T. D.; Reinhorn, A. M.; Shinozuka, M.; Tierney, K.; Wallace, W. A.; and von Winterfeldt, D. 2003. A Framework to Quantitatively Assess and Enhance the Seismic Resilience of Communities. *Earthquake Spectra* 19(4):733–752.
- Chakraborty, S.; Fremont, D. J.; Meel, K. S.; Seshia, S. A.; and Vardi, M. Y. 2014. Distribution-aware sampling and weighted model counting for SAT. In *Proc. of AAAI*, 1722–1730.
- Chakraborty, S.; Fremont, D. J.; Meel, K. S.; Seshia, S. A.; and Vardi, M. Y. 2015a. On parallel scalable uniform sat witness generation. In *Proc. of TACAS*, 304–319.
- Chakraborty, S.; Fried, D.; Meel, K. S.; and Vardi, M. Y. 2015b. From weighted to unweighted model counting. In *Proceedings of IJCAI*, 689–695.
- Chakraborty, S.; Meel, K. S.; Mistry, R.; and Vardi, M. Y. 2016. Approximate probabilistic inference via word-level counting. In *Proc. of AAAI*.
- Chakraborty, S.; Meel, K. S.; and Vardi, M. Y. 2013. A scalable approximate model counter. In *Proc. of CP*, 200–216.
- Chakraborty, S.; Meel, K. S.; and Vardi, M. Y. 2016. Algorithmic improvements in approximate counting for probabilistic inference: From linear to logarithmic SAT calls. In *Proc. of IJCAI*.
- Chang, Y., and Mori, Y. 2013. A study on the relaxed linear programming bounds method for system reliability. *Structural Safety* 41:64–72.
- Clote, P., and Setzer, A. 1998. On php, st-connectivity and odd charged graphs. *Proof Complexity and Feasible Arithmetics* 39:93–117.
- Dotson, W., and Gobien, J. 1979. A new analysis technique for probabilistic graphs. *IEEE Transactions on Circuits and Systems* 26(10):855–865.
- Dueñas-Osorio, L., and Rojo, J. 2011. Reliability Assessment of Lifeline Systems with Radial Topology. *Computer-Aided Civil and Infrastructure Engineering* 26(2):111–128.
- Dueñas-Osorio, L.; Meel, K. S.; Paredes, R.; and Vardi, M. Y. 2017. Sat-based connectivity reliability estimation for power transmission grids. Technical report, Rice University.
- Ermon, S.; Gomes, C. P.; Sabharwal, A.; and Selman, B. 2013. Taming the curse of dimensionality: Discrete integration by hashing and optimization. In *Proc. of ICML*, 334–342.
- Fishman, G. S. 1986. A Monte Carlo Sampling Plan for Estimating Network Reliability. *Operations Research* 34(4):581–594.
- Gomes, C.; Sabharwal, A.; and Selman, B. 2007. Near-uniform sampling of combinatorial spaces using XOR constraints. In *Proc. of NIPS*, 670–676.
- Gómez, C.; Sánchez-Silva, M.; and Dueñas-Osorio, L. 2011. Clustering methods for risk assessment of infrastructure network systems. In *International Conference on Applications of Statistics and Probability in Civil Engineering*, 1389–1397.
- HAZUS. 2003. Multi-hazard Loss Estimation Methodology, Earthquake Model, HAZUS-MH MR4 Technical Manual.
- Hurtado, J. E. 2013. *Structural reliability: statistical learning perspectives*, volume 17. Springer Science & Business Media.
- Ivrii, A.; Malik, S.; Meel, K. S.; and Vardi, M. Y. 2015. On computing minimal independent support and its applications to sampling and counting. *Constraints* 1–18.
- Karger, D. R. 2001. A Randomized Fully Polynomial Time Approximation Scheme for the All-Terminal Network Reliability Problem. *SIAM Review* 43(3):499–522.
- Li, J., and He, J. 2002. A recursive decomposition algorithm for network seismic reliability evaluation. *Earthquake Engineering & Structural Dynamics* 31(8):1525–1539.
- Lim, H.-w., and Song, J. 2012. Efficient risk assessment of lifeline networks under spatially correlated ground motions using selective recursive decomposition algorithm. *Earthquake Engineering & Structural Dynamics* 41(13):1861–1882.
- Page, L., and Perry, J. 1988. A practical implementation of the factoring theorem for network reliability. *IEEE Transactions on Reliability* 37(3):259–267.
- Rai, S., and Kumar, A. 1987. Recursive Technique For Computing System Reliability. *IEEE Transactions on Reliability* R-36(1):38–44.
- Soos, M.; Nohl, K.; and Castelluccia, C. 2009. Extending SAT Solvers to Cryptographic Problems. In *Proc. of SAT*.
- The White House, Office of the Press Secretary,. 2016. FACT SHEET: Obama Administration Announces Public and Private Sector Efforts to Increase Community Resilience through Building Codes and Standards.
- Thurley, M. 2006. SharpSAT: counting models with advanced component caching and implicit BCP. In *Proc. of SAT*, 424–429.
- Valiant, L. 1979. The complexity of enumeration and reliability problems. *SIAM Journal on Computing* 8(3):410–421.
- Yin, C., and Kareem, A. 2016. Computation of failure probability via hierarchical clustering. *Structural Safety* 61:67–77.
- Zio, E. 2009. Reliability engineering: Old problems and new challenges. *Reliability Engineering & System Safety* 94(2):125–141.
- Zuev, K. M.; Wu, S.; and Beck, J. L. 2015a. General network reliability problem and its efficient solution by Subset Simulation. *Probabilistic Engineering Mechanics* 40:25–35.
- Zuev, K. M.; Wu, S.; and Beck, J. L. 2015b. General network reliability problem and its efficient solution by subset simulation. *Probabilistic Engineering Mechanics* 40:25–35.