Proceedings of the Twenty-Ninth International
Florida Artificial Intelligence Research Society Conference

# Bayesian Network-Based Extension for PGP — Estimating Petition Support

**Marius Silaghi[†], Song Qin[†], Toshihiro Matsui[‡], Makoto Yokoo[*], Katsutoshi Hirayama[§]**
[†]Florida Institute of Technology, [‡]Nagoya Institute of Technology, [*]Kyushu University, [§]Kobe University

## Abstract

Consider the problem of estimating the expected number of distinct eligible voters among the authors of a set of electronic signatures gathered for a petition (or citizen initiative) that has to pass legally required thresholds.

We formalize this problem and propose an extension to the Pretty Good Privacy Web Of Trust, a mechanism for reciprocally certifying identities between peers. The extension (a) enables agents to certify additional relevant statements about others, and (b) gives agents opportunities for negative authentication statements (e.g., on ineligibility of an identity).

A Bayesian Network model enables inferences on the data provided by the proposed PGP extension. Simulations and an agent-based platform are used to validate the concepts.

## Introduction

Activists that strive to prove popular support for a petition (or citizen initiative) face the challenge of ensuring that the gathered signatures come from *eligible* persons, i.e. people whose right to vote on the raised issue is recognized by the *authority* targeted by the petition. Such authorities have the possibility to verify signatures (by having access to databases and technology that are not public), and they commonly verify the signatures after they are gathered. That is too late for the committees investing the effort of gathering the signatures, who do no longer have time to gather remaining support needed to pass legally required thresholds.

Furthermore, early assesment of the real support can help petition committees better estimate their chances of success and better manage their resources. It can save them money if they find when additional campaigning is not needed due to the reach of sufficient support, or that the petition drive has to be completely abandoned due to a steady opposition.

In this article we describe how this problem can be formalized and addressed in a principled way by using Bayesian Networks as an extension of Pretty Good Privacy (PGP). The introduced framework defines the *expected valid petition support problem*. Simulations, as well as a real platform (Figure 1), are used for its validation with agents that can submit and disseminate petitions and signatures using a peer-to-peer architecture (Silaghi et al. 2013). Each agent




Figure 1: Eligibility Inferences during Petition Drives

evaluates the eligibility of the authors for known signatures based on the available information that signers submit about each-other. The large size of the local databases is addressed by employing approximate inference techniques.

While the described inference is Bayesian, an agent's numeric computation can be sped up with ad-hoc techniques. Here, agents that trust each-other collaborate by exchanging and combining partial computations on their samples. This part is not backed by a belief network model, but the dissemination of false contributions by attackers can be countered using influence inhibition heuristics, such as amortization at each hop, and bound on the weight of external contributions.

## Background

Countries ranging from India to EU provide citizens with opportunities to convey opinions to their representatives. Often the proposed approaches are ad-hoc, ranging from counting SMS messages (in India) to various forums without authentication. The relatively more tested traditional method of gathering signatures on paper has few equivalent

electronic solutions. This is due to the fact that verifiable electronic signature identities are not yet widely available to the population or have prohibitive costs. The most reachable technique, Pretty Good Privacy (PGP), has not seen a sufficiently wide acceptance, likely due to limitations in the usability of available tools and of a limited expressive power for the transitivity of the trust (Ferguson and Schneier 2003).

PGP (Zimmermann 1995) is a mechanism whereby people can certify each other's electronic identities based on visually inspecting each other's government issued identity documents. Key Signing Party events are organized to provide opportunities to potential users to verify each other's documents and issue PGP certifications. A remote agent is trusted in PGP if one finds a chain of identities leading to it, where each node of the chain certifies the identity of the next node and is manually flagged as *trusted* by the user (or if two nodes flagged as *partially trusted* certify the next node). PGP has no mechanism to provide negative certification, denying that a given peer is eligible (aka *legitimate*).

In (Douceur 2002), methods are classified into direct validation and indirect validation approaches. The former suggests that an entity only accepts identities that it has directly validated by some means (Garfinkel 2003). The latter suggests that an entity accepts identities that are vouched for by already accepted identities, as in PGP and X509 certification. Our approach for validation of identities is related to these, but we elicit and exploit both positive validations and negative certificates for each identity, as well as for other kinds of statements.

Bayesian Networks for evaluation of peer properties have been studied in different domains, such as web services (Nguyen, Zhao, and Yang 2010) and anonymous interaction experience in specific contexts (Quercia, Hailes, and Capra 2006; Orman 2013). Research concerning Bayesian network study of groups in file exchanges (Wang and Vassileva 2003) can be seen as an analogy to our witnessing on neighborhoods. The connectivity is suggested by the small world paradigm (Milgram 1967).

## Formalized Problem

To enable the estimation of the expected number of valid supporting signatures among the digital signatures gathered for a petiton, we extend PGP with a mechanism:
(a) enabling agents to certify additional relevant statements about others, besides eligibility (aka legitimacy),
(b) giving agents opportunities for negative authentication statements (e.g., claiming falsity of an identity), and
(c) supporting the evaluation of the probability of someone's eligibility by Bayesian inference.

Authorities require petitions to be supported by a predefined number of valid signatures in order to qualify for further consideration (Obama 2014; Phillips 2014). Signatures are collected together with data identifying their authors, in a *signature-identity* pair. The *identity* information has to uniquely identify the person. It can contain the information typically found in a phone book, and an email. To illustrate what can be an identity, an example is:

**Example 1** *John Doe; jdoe@ddp2p.net; main residence: 1024 $6^{th}$ St., Cambridge MA 02139*

Email addresses alone may be considered insufficient since people can easily have multiple of them, making it difficult to detect repeated signatures. The email can be provided to enable remote peer certification of known people without having to attend a PGP key signing party (Garfinkel 2003). While some kind of address is expected for uniquely identifying an identity, it does not have to be the postal address of the main residence. The identity could also be structured according to different classification hierarchy, such as place and date of birth, or unique identification numbers in countries where these are not considered private. The identity can also contain a public key for digital signatures.

The *signature* can consist of any digital signature of a support statement (Kattamuri et al. 2005) for the petition with a key certified to belong to the associated identity. Here we do not cover text processing (Rozenknop and Silaghi 2001).

**Definition 1 (Validity)** *A signature-identity pair for a given petition is* valid *(aka legitimate) if:*

1. *the author's identity corresponds to a person whose right to vote on the petition (aka eligibility) is recognized by the authority targeted by the petition*
2. *the author's identity signs this petition only once*
3. *the declared author's identity corresponds to the actual author of the signature*

**Definition 2 (EVPSP)** *The Expected Valid Petition Support Problem for a petition, is defined by a tuple $\langle \mathcal{I}, \Sigma, \mathcal{W}, \Omega, \Gamma, B \rangle$, where $\mathcal{I} = \{C_1, ..., C_n\}$ is a set containing $n$ identities, $\Sigma$ is the subset of $m$ identities signing the petition, $\mathcal{W}$ is a set of $w$ certificates, $\Omega$ is a set of properties, and $\Gamma$ is an identity that is fully trusted (the identity of the observer) or $\perp$ (for none). $B$ is a belief network defining the relation between each probability $\Psi(C_k), C_k \in \Sigma$, that $C_k$'s signature is valid, as a function of $\mathcal{W}$ and $\Gamma$. Each certificate is of the type $(a, b, P)$ specifying that the identity $C_a$ certifies the set of statements $P$ concerning the identity $C_b$. $P$ is an assignment of the properties in $\Omega$ with values from their domain, or $\perp$ if unknown.*

*The problem is to find the expected number $E$ of valid gathered digital signatures, defined as the sum over all signatures in $\Sigma$ for the probabilities that the signature is* valid*.*

$$E[[valid(\Sigma)]] = \sum_i (i \times P(|valid(\Sigma)| = i)) = \sum_{(s \in \Sigma)} \Psi(s))$$

Here $valid(\Sigma)$ gives the subset of $\Sigma$ submitting valid signatures. The relation results from Jensen's inequality. The trust of observer $\Gamma$ in its own statements inherits the *ultimate trust* concept of PGP. A belief network $B$ is shown later.

## Agent Framework

In this section we introduce in detail the definitions of the items exchanged by the decentralized petitions platform.

Items used in this work are referenced using global identifiers (GIDs), built in a way that avoids intended and unintended collisions between different items (typically by generating them either as a public key, or as the secure digest of

the item data, alternative to KeyID in PGP). The secure digest function is denoted with $HASH(d)$ where $d$ is the data whose digest is computed. Given a public key $P$, we refer to its secret key as $SK(P)$. The digital signature for data $d$ using secret key $S$ is computed by $SIGN(S, d)$. Identities are validated with regard to a verifying authority, representing an organization, and identified by the a global identifier. Organizations (Orman 2011) where the support is estimated using the proposed decentralized technique are also referred as grassroot organizations (e.g., a city, a county, a state).

**Constituents** The people with right to cast votes that have a predefined weight in an organization form its constituency. A constituent is defined by a tuple $\langle \mathcal{C}, \mathcal{O}, i, d, r, s \rangle$ where $\mathcal{C}$ is its GID, $\mathcal{O}$ is the GID of the relevant organization, $i$ is the set of identity details, and $d$ is the date and time when $i$ was declared. $\mathcal{C}$ is specified as a public key, $r$ is the revocation status of $\mathcal{C}$ used as per the PGP framework, and the constituent data is signed with $s = SIGN(SK(\mathcal{C}), \langle \mathcal{O}, i, r \rangle)$.

**Neighborhood** For ease of accounting, constituents can be organized in tree structures with nodes (called neighborhoods) corresponding to localities, cities, counties, states and countries. They can also be a hierarchy of subdivisions in an organization: university, college, department, center, lab. The leaf of the tree of neighborhoods is the smallest cell of the identity management, and can be configured to correspond to a block, a street or an area/unit small enough (relatively to the population density) such that members can learn and easily verify residency of their neighbors.

Formally a neighborhood is a tuple $\langle \mathcal{N}, n, t, P, c \rangle$, where $\mathcal{N} = HASH(n, t, P, c)$ is the GID of the neighborhood, $n$ is the name of the neighborhood, $t$ is its type/level (e.g., city, block, unit), $P$ is the GID of the parent neighborhood ($\perp$ for a top neighborhood), and $c$ is the list of expected types of descendant levels under this neighborhood.

**Example 2** *A sample neighborhood item is:*

$\mathcal{N}$ = 0x4e1fea12c4... // the hash
n = Lilliput
t = country
P = $\perp$
c = "[county/city/street/block]"

The purpose of the $c$ term is to enable automatic hints for new constituents concerning the type of address fields that they are expected to provide. Note that a child neighborhood can overwrite the expected descendants $c$ suggested by its parent neighborhood, thereby enabling the heterogeneity of addressing scheme existing in the real world (e.g., where some countries have states and other do not). A neighborhood item is of interest only if it is supported (witnessed) by some constituent. Next we give the definition of witness items and their extension to witnessing on neighborhoods.

**Witness** Constituents in a grassroot organization can support or oppose the other constituent items' eligibility for signing petitions. We say that they perform favorable or unfavorable witness stances for those identities. A witness stance can be associated with a set of semantic statements (as epistemological commitments associated to ontological commitments from a set $\Omega$), such as:

- existence versus nonexistence of constituent name-address pair,
- constituent public key (GID) belongs or not to the constituent with declared name-address pair (e.g., checked using WebFinger for a well known email address),
- favored versus disfavored version of a multiply occurring constituent (e.g., at the current residence versus an old residence, or with a correct versus a misspelled name),
- eligibility versus ineligibility of constituent,
- correctness versus inaccuracy of details in identity,
- reliability versus sloppiness of witness.

For example, when a constituent $A$ declares constituent $B$ to be a *sloppy witness*, then $A$ believes that $B$ does not carefully verify all the constituents that it witnesses, unlike a *reliable witness*.

Such a witnessing stance is defined by a tuple $\langle W, \mathcal{O}, S, T, m, e, d, \sigma \rangle$ where $\mathcal{O}$ is an organization identifier, $S$ is the constituent identifier of the witnessing constituent, $T$ is the constituent identifier of the target constituent item and $e$ is an human readable explanation. The set of semantic statements of the witnessing, where each of them can be either *favorable* or *unfavorable*, is captured in $m$. The parameter $d$ represents the creation time of this witnessing stance. The signature is generated as:

$$\sigma = SIGN(SK(S), \langle \mathcal{O}, T, m, e, d \rangle).$$

The GID of the witness stance is generated as:

$$W = HASH(\mathcal{O}, S, T, m).$$

Constituents can also witness about the legitimacy of a whole neighborhood. For example, they can state that no locality called *Mildendo* exists in their county, or that no street called $21^{st}$ *Street* exists in their city. Such a witness stance is represented by a tuple $\langle W, \mathcal{O}, S, \mathcal{N}, m, e, d, \sigma \rangle$ where the only difference with witness stances for constituent items is that a GID of a neighborhood $\mathcal{N}$ is specified instead of the GID of a target constituent $T$. Semantic statements for such witness stances can be of type:

- favored versus disfavored version of a multiply occurring neighborhood (e.g., *New York* vs *New-York city*)
- existing vs nonexistent neighborhood

**Example 3** *A sample witness stance for neighborhoods is:*

$W$ = 0x3e5fa // the hash
$\mathcal{O}$ = 0xe4273f5a6b... // Lilliput's GID
$S$ = 0xfe13ca6... // IDC6_PK
$\mathcal{N}$ = 0x4e1fea12c4... // IDC7
$m$ = "existance=favorable; favored=favorable"
$e$ = "Visited recently"
$d$ = 20130406112015.012Z
$\sigma$ = 0a1e6f7a8bc98e7b9a9e...

## Decentralized Petitioning Concepts

Now let us introduce concepts involved in the decentralized petition drive processes.

**Citizen Interactions** A citizen-driven petition requires participation of individual citizens for actions such as *residence declaration* and *witnessing*. As residence declarations, each individual voluntarily provides identity data not only about herself but also about her neighbors. The neighborhood where a citizen resides is part of its identity details.

**Witness Graph** A graph defined by the witness relations between constituents can be generated in the following way:

- A node is generated for each constituent.
- A directed edge from node $A$ to node $B$ is generated for each semantic statement that $A$ witnesses for $B$.
- Each edge has a color (from a set $\Omega$), given by the type of statement that generated it (ontological commitment).
- An edge has weight 1 if generated for a favorable stance and weight 0 if generated for an unfavorable stance (epistemological commitment).

This graph can be used to reason about the eligibility of the declared identities and implicitly about the petition support.

## Techniques

Here we present techniques used to address the challenge of inferring the expectation of the count of legitimate signers among the gathered digital signatures for a petition, given a witness graph for a grassroot organization. Let us further refine some of the used concepts.

**Eligibility** Although anyone can participate in the debates of a petition drive of a grassroot organization, not everyone is eligible to submit valid support signatures. In a grassroot organization, which is the context of this study, the definition of eligibility can be a function of the constituent (either due to the fact that the decision of the authority is not know, or because it is contested). When the eligibility for a constituent is based on a subjective view, the petition support evaluation result is relevant only to the user (or users) sharing this view. Hence, we define the eligibility as a probabilistic function of several parameters:

- Someone's interpretation of the witness graph, $\mathcal{M}_{\mathcal{S}}$
- Someone's own definition of the eligibility, $\mathcal{O}(\Gamma)$

**Definition 3** *The reference user is the user* $\Gamma$ *who currently computes the support estimation.*

**Definition 4 (Eligible and $\Psi$)** *A constituent item $C$ is* `eligible` *for an organization if it is eligible and new (never counted elsewhere). The* $\Gamma$*'s confidence value in whether $C$ is* `eligible` *is denoted* $\Psi(C)$.

**Definition 5 (Witness Reliability and $\Phi$)** *A constituent item $C$ is a* reliable witness *if* $\Gamma$ *trusts all the witness stances that $C$ issues as she trusts her own.* $\Gamma$ *may not fully trust the stances of another constituent $C$, but only with a confidence value* $\Phi(C)$.

Based on the EVPSP parameters, one can infer a value $\Psi$ for the confidence that observer $\Gamma$ can have on whether a given constituent item $C$ corresponds to an `eligible` user, and a value $\Phi$ for its confidence on whether $C$ is *witness reliable*.

Random variables are used to represent the `eligible` property of each constituent, the *reliable witness* property, and the witnessing stances between each pair of constituents for each quality. All these random variables are Boolean. For each pair of constituents $A$ and $B$ we get the random variables and Bayesian Network in Figure 2. Note that each pair of constituents requires the introduction of $2|\Omega|$ random variables for $|\Omega|$ qualities. With the two considered qualities in Figure 2, $\Phi$ and $\Psi$, a constituent (e.g., $A$) is associated with two hidden random variables: `eligible` ($CS^A$) and *reliable_witness* ($RW^A$). Each pair of constituents items (e.g., A and B) is associated with four evidence (grayed) random variables: A witnesses for B being a reliable witness ($W^{AB\Phi}$), A witnesses for B being `eligible` ($W^{AB\Psi}$), B witnesses for A being a reliable witness ($W^{BA\Phi}$), B witnesses for A being `eligible` ($W^{BA\Psi}$).

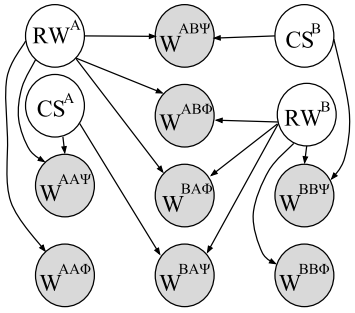


Figure 2: A Bayesian Network that models the EVPSP with two constituent items witnessing each other and themselves. The Bayesian Network of larger communities is similarly built for any configuration of witnessing stances by replicating the corresponding nodes and arcs.

**Property 1** *The number of random variables (all Boolean) in a Bayesian Network modeling a EVPSP is linear in the size of the input.*

**Proof.** Given $n$ constituent items $C_1, ..., C_n$, we get at most $n^2|\Omega|$ random variables:

- $n|\Omega|$ hidden variables modeling the real qualities of each constituent item, and
- $n(n-1)|\Omega|$ modeling the evidence variables about all $|\Omega|$ possible witness stances between each of the $n(n-1)$ possible directed pairs of constituents.

Note that we do not need to model with random variable the nonexistent witness stances. Therefore the actual network size is linear in the size of the input, being proportional to the number $w$ of input semantic statements in witness stances $(w+n)|\Omega|$.

For average sized networks one can perform queries of values for the random variables $CS^{C_i}$, modeling $\Psi(C_i)$ of

$i^{th}$ constituent item, using techniques such as Markov Chain Monte Carlo (MCMC).

**Scalability** The approximate inference technique, does not require to simultaneously load in memory the entire Bayesian Network defined by all signatures and identities available to an agent. It is sufficient to load a node (to be re-estimated) and its Markov blanket, at a time. In fact, we load as many such *node-Markov blanket* pairs as made possible by the available resources of the agent.

The *community-focused heuristic* we use for loading such pairs is to have each of the nodes loaded for re-evaluation, selected to be part of the Markov blanket of many other loaded nodes (identities that witness on each-other). Reference counters are used with each node to only release the memory it takes when all Markov blankets sharing the node are discarded. Multiple re-evaluation rounds are performed on the loaded nodes before loading different communities.

**Efficiency** Peers that trust each other collaborate by exchanging counts obtained for the nodes that they evaluate. This helps new agents to faster converge to good approximations of the probabilities. Data exchanged this way can be used by attackers to influence the authentication results, and therefore received counts are bound to not exceed the weight of the locally obtained counts, as soon as local counts exceed a user-defined threshold. Also, counts received are passed further only after reduction with a further amortization factor, reducing fast the potential influence of the attackers. Past a second threshold, incoming counts are discarded in favor of the result from local computations.

## Experiments

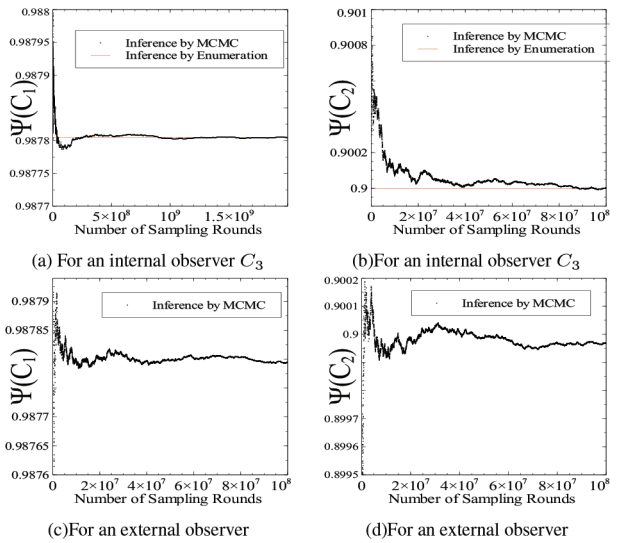


Figure 3: Inference of $\Psi(C_1)$ and $\Psi(C_2)$

With the Bayesian network built following the pattern in Figure 2 and rough manually prepared CPTs we use MCMC to perform queries of values for the random variable $CS^{C_i}$ which models the $\Psi(C_i)$ of the $i^{th}$ constituents. Convergence for a few constituents is shown in Figures 3 (a), (b), (c), (d). The reference (red line) is computed with an exact inference by enumeration. For an external observer the exact inference by enumeration is expected to take 33 days on a computer and is not shown.

We have performed extensive experiments with simulated data. Most of them will not be described here for lack of space. The result of a set of 5 experiments estimating the impact of the percentage $k$ of honest active constituents (HACs) in the global population on the `eligible` properties of the constituents is shown in Figure 4. The true positive rate (TPR) gives the percentage of correctly counted constituents out of the total number of eligible constituents. The false positive rate (FPR) gives the percentage of wrongly counted constituents with respect to the total number of eligible constituents. A robust verification process has a high TPR and a low FPR. The semantic statements for witness stances are only about the eligibility quality. In this experiment, a constituent item $(n, a)$ is eligible if there is someone whose name is $n$ and lives at address $a$. We simulate attackers that declare a number of ineligible constituent items and perform favorable witness stances for a percentage of $h$ ineligible constituent items in their leaf neighborhoods. We assume that witness stances represent all semantic statements. The plotted points are for $h$ of 100%, 93.75%, 87.5%, 75%, 50%, 25%, 12.5%, 6.25%, 0% respectively. The values of $k$ are (0.9, 0.8, 0.7, 0.5 and 0.3).

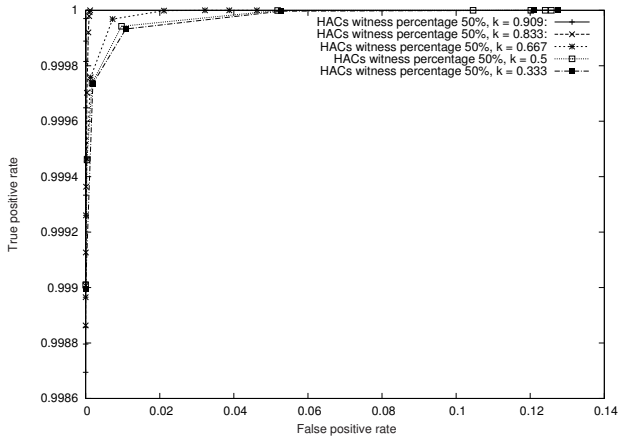


Figure 4: The effects of the percentage of honest active constituent ($k$ value) in global population, for varying $h$.

There are several common parameters for the plotted curves. The total number of eligible constituents is 9300000. The *HACs witness percentage* is 50% out of their leaf neighborhoods. *HACs witness percentage* is the percentage of constituent items that an HAC witnesses honestly (A favorable witness stance is performed if an item is eligible and an unfavorable witness stance is performed if an item is ineligible). The number of *attackers* is 300000. The number of ineligible constituent items declared by each *attacker* is 4.

Since we see in Figure 4 that the curve with parameter

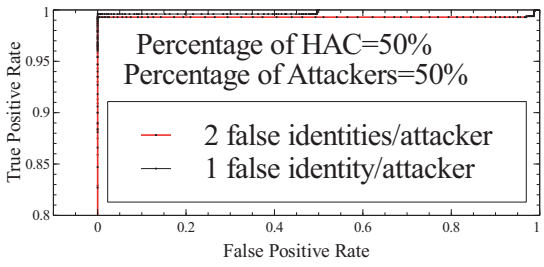


Figure 5: ROC: 1000 MCMC rounds and 1000 constituents

$k$=0.9 is higher than the curve with parameter $k$= 0.8, the curve with parameter $k$=0.8 is higher than the curve with parameter $k$=0.7, the curve with parameter $k$=0.7 is higher than the curve with parameter $k$=0.5 and the curve with parameter $k$=0.5 is higher than the curve with parameter $k$=0.3, we conclude that the $k$ value will affect the robustness of the system positively. That is, the bigger the $k$ value is, the more accurate the system will be.

Figure 5 illustrates the ROC curve for an unstructured experiment (without neighborhoods) based on 1000 real constituents and 1000 MCMC rounds. Each attacker creates $n$ fake identities into the global population ($n \in \{1, 2\}$) and has a favorable witness stance on each of the fake identities. Percentages of HACs and attackers out of all constituents are both 50%. Each HAC has a correct witness on 1% of all identities (real and false). Constituent $C$ is counted here when $\Psi(C)$ is greater than a threshold $t$ (varying between 0 and 1). The ROC curves has a low false positives rate and high true positives rate which indicates a robust system for this configuration of parameters. A second observation is that the curve for more attackers is only slightly below the other one, which is consistent with graceful degradation. We hope to extend this study to many additional configurations.

## Conclusions

To enable an estimation of expected gathered support for decentralized petitions, we extend Pretty Good Privacy (PGP) with additional flexibility to certify properties about peers besides legitimacy, as well as to enable negative certification. Constituents can witness (vote) on each other's qualities, such as: eligibility and witnessing reliability.

The concept of neighborhood is introduced and formalized in order to simplify peer verification. Neighborhoods group constituent addresses in a tree structure. For large organizations, the constituency is organized in a tree of neighborhoods to help evaluating the gathered support for petitions. Eligibility can be verified separately for each neighborhood and it is reasonable to expect most users to be able to verify the existence of the immediate child neighborhoods of all the nodes on the path from their own address to the root. This enables a distributed verification of the existence of declared neighborhoods (and thereby of addresses).

Items for these concepts are identified by global identifiers guaranteed to be unique and that are disseminated among peers based on multi-agent protocols (current experiments being based on the currently implemented platform). A peer is an user acting under one name and public key via multiple agents (e.g., one agent per device that she uses).

We have also proposed and analyzed theoretically and empirically a probabilistic model based on Bayesian Networks that can be used to address the *expected valid petition support problem* in a principled way. Markov Chain Monte Carlo inferences are found to converge within few seconds for reasonable sized neighborhoods, and scale linearly with the input size. The outputs are probabilities that can be used to compute expected support for petitions/initiatives. Experiments with simulated data show that robustness to attackers is possible with a kernel of honest active constituents.

## References

Douceur, J. 2002. The sybil attack. *Peer-to-peer Systems*.

Ferguson, N., and Schneier, B. 2003. *Practical Cryptography*. John Wiley & Sons.

Garfinkel, S. L. 2003. Email-based identification and authentication: An alternative to PKI? *IEEE Security and Privacy* 1(6):20–26.

Kattamuri, K.; Silaghi, M.; Kaner, C.; Stansifer, R.; and Zanker, M. 2005. Supporting debates over citizen initiatives. In *Digital Government Research*, 279–280.

Milgram, S. 1967. The small world problem. *Psychology today* 2(1):60–67.

Nguyen, H. T.; Zhao, W.; and Yang, J. 2010. A trust and reputation model based on bayesian network for web services. In *Web Services*, 251–258.

Obama, B. 2014. We the people. https://petitions.whitehouse.gov.

Orman, L. 2011. Potential advantages of virtual institutions. *IEEE Technology Society Management* 1(30):56–64.

Orman, L. 2013. Bayesian inference in trust networks. *ACM Transactions on Management Information Systems* 4(2):7:1–7:21.

Phillips, M. 2014. European citizens' initiative: Basic facts. http://ec.europa.eu/citizens-initiative/public/basic-facts.

Quercia, D.; Hailes, S.; and Capra, L. 2006. B-trust: Bayesian trust framework for pervasive computing. In *Trust Management, LNCS*, volume 3986. Springer. 298–312.

Rozenknop, A., and Silaghi, M.-C. 2001. Algorithme de décodage de treillis selon le critère du coût moyen pour la reconnaissance de la parole. In *TALN*, 391–396.

Silaghi, M. C.; Alhamed, K.; Dhannoon, O.; Qin, S.; Vishen, R.; Knowles, R.; Hussien, I.; Yang, Y.; Matsui, T.; Yokoo, M.; and Hirayama, K. 2013. DirectDemocracyP2P — decentralized deliberative petition drives —. In *P2P*.

Wang, Y., and Vassileva, J. 2003. Bayesian network-based trust model. In *Web Intelligence*, 372–.

Zimmermann, P. R. 1995. *The Official PGP User's Guide*. Cambridge, MA, USA: MIT Press.