Open Census for Addressing False Identity Attacks in Agent-based Decentralized Social Networks

(Extended Abstract)

Song Qin
Florida Institute of Technology

Marius C. Silaghi
Florida Institute of Technology

Ihsan Hussien
Florida Institute of Technology

Makoto Yokoo
Kyushu University

Toshihiro Matsui
Nagoya Institute of Technology

Katsutoshi Hirayama
Kobe University

Decentralized social networks defined by user actions, e.g., decentralized discussion forums, are expected to be ideal environments for Sybil and false identity attacks (just as in the case of the similar web based systems: YouTube, etc.). In particular, these attacks form a significant impediment for meaningful electronic *deliberative petition drives*, aka. *citizen initiatives*, where the eligibility of voters is required. While many electronic social networks strive for guaranteeing the privacy of their users, existing systems for petition drives ask users to disclose their real identities and are meaningless when users do not follow this request. We describe a framework and investigate techniques for running *decentralized census* processes (DCP) that enable observers to independently verify the identity of participants in social networks for petition signing, `DirectDemocracyP2P.net` [1]. In a synergy, the verification of identities refines the census, while the census information helps to detect false addresses and Sybils, for identity verification.

Identity verification is initialized by PGP-based witnessing of people for their friends, neighbors, and acquaintances. Further, the identity verification is propagated using Bayesian inference.

To enable synergy between census and the authentication process, the census is hierarchical based on the structure of addresses and supports witnessing for one's relevant inactive acquaintances.

Decentralized Census. In `DirectDemocracyP2P.net`, and separately for each jurisdiction, agents continuously disseminate *witness stances* about the *"voting eligibility"* and *"reliability in witnessing"* of agents that they know from the real life. Each agent individually employs the information it has available to evaluate the *probability of the voting eligibility* of each other agent. These probabilities are integrated in a separate census per neighborhood. They are also used as weights in assessing support for current petitions signed by these agents (i.e., whether petitions have chances to pass legal thresholds). We do not see a problem if different agents reach different conclusions based on what they know.

Here we model an agent's census decision making based on Bayesian Networks. Random variables are used to represent the `censable` (i.e., eligible for declared voting weight) property of each constituent, her `reliable` witness property, and the witnessing stances between each pair of constituents for each quality. For each pair of constituents A and B we get the random variables and Bayesian Network in Figure 1. Nodes corresponding to unavailable witness stances are dropped, together with corresponding arcs.

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

Ψ^A	Φ^B	$P(W^{BA\Psi})$
t	t	0.9
t	f	0.5
f	t	0.1
f	f	0.5

Φ^A	Φ^B	$P(W^{BA\Phi})$
t	t	0.9
t	f	0.5
f	t	0.3
f	f	0.5

Ψ^A	Φ^A	$P(W^{AA\Psi})$
t	t	0.99
t	f	0.5
f	t	0.1
f	f	0.5

Φ^A	$P(W^{AA\Phi})$
t	0.99
f	0.5

$P(\Phi)$	$P(\Psi)$
0.5	0.5

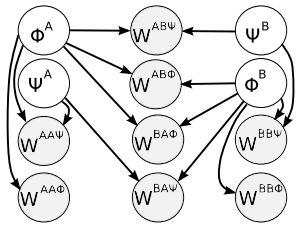

Figure 1: DCP Bayesian Network for a pair of agents

In general, a witness stance of a constituent associates his epistemological committments to ontological commitments from a set of *qualities* Ω. In the above discussion, $\Omega = \{\texttt{censable}(\Psi), \texttt{reliable}(\Phi)\}$. Note that each pair of constituents requires the introduction of $2|\Omega|$ random variables for $|\Omega|$ qualities. With the two considered qualities in Figure 1, Φ and Ψ, a constituent (e.g., A) is associated with two hidden random variables: `censable` (Ψ^A) and `reliable` *witness* (Φ^A). Each pair of constituents items (e.g., **A** and **B**) is associated with at most $2|\Omega| = 4$ evidence (grayed) random variables: **A** witnesses for **B** being a `reliable` witness ($W^{AB\Phi}$), **A** witnesses for **B** being `censable` ($W^{AB\Psi}$), **B** witnesses for **A** being a `reliable` witness ($W^{BA\Phi}$), **B** witnesses for **A** being `censable` ($W^{BA\Psi}$).

Conditional probability tables (CPTs) can be trained from real data once a large amount of data is available. Sample conditional probability tables built manually for the random variables of type Ψ^A, Φ, $W^{BA\Phi}$ and $W^{BA\Psi}$ are shown in Figure 1.

PROPERTY 1. *The number of random variables in a Bayesian Network modeling a DCP is linear.*

For average sized networks one can perform queries of values for the random variables Ψ^{C_i}, the censability of the i^{th} constituent item, using techniques such as Markov Chain Monte Carlo

(MCMC). The MCMC state can be initialized and enhanced with *values exchanged with trusted peers*.

Neighborhood. To simplify peer verification, neighborhoods are used to group constituent addresses in a tree structure. For large organizations, the constituency is organized in a tree of neighborhoods to help with census organization. Census results 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).

Experiments. To evaluate the power of the studied DCP models to represent users reasoning about census, as well as to resist various attacks, we perform two sets of preliminary experiments. One of them is based on a set of volunteers and the other is based on a larger simulated data.

In the experiments based on volunteers we asked 10 people living within an area of a few square kilometers to register themselves as active constituents and to also register others 10 friends as inactive constituents of a regional organization. Each of these volunteers had the opportunity to witness for the other constituent items that they knew. We also introduced 2 obviously wrong constituents at an address that most participants knew to not exist.

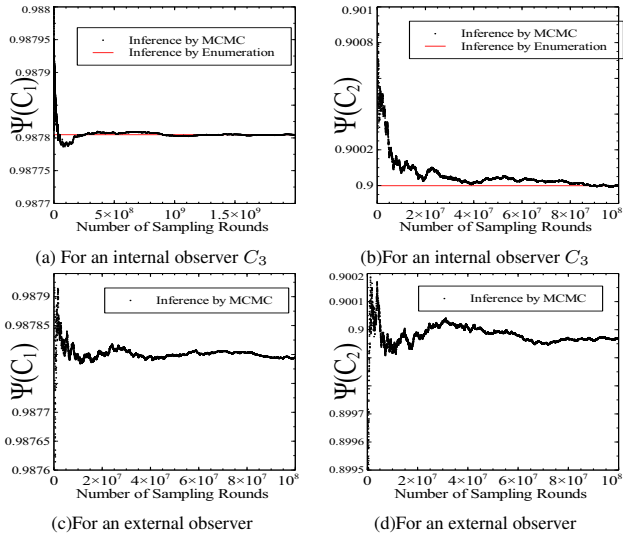

(a) For an internal observer C_3 (b)For an internal observer C_3

(c)For an external observer (d)For an external observer

Figure 2: Inference of $\Psi(C_1)$ and $\Psi(C_2)$

With the Bayesian network with CPTs like in Figure 1, we use MCMC to perform queries for the random variable Ψ^{C_i} of the i^{th} constituent. Convergence for a few constituents is shown in Figures 2 (a), (b), (c), (d). For verification of the convergence, 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.

Unlike for the case of experiments based on volunteers where our preliminary samples are small, 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 `censable` properties of the constituents is shown in Figure 3. 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 census process has a high TPR and a low FPR. The witness stances are only about the `censable` 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. The plotted points are for h of 100%, 93.75%, 87.5%, 75%, 50%, 25%, 12.5%, 6.25%, 0% respectively. The studied values of k are (0.9, 0.8, 0.7, 0.5 and 0.3) respectively.

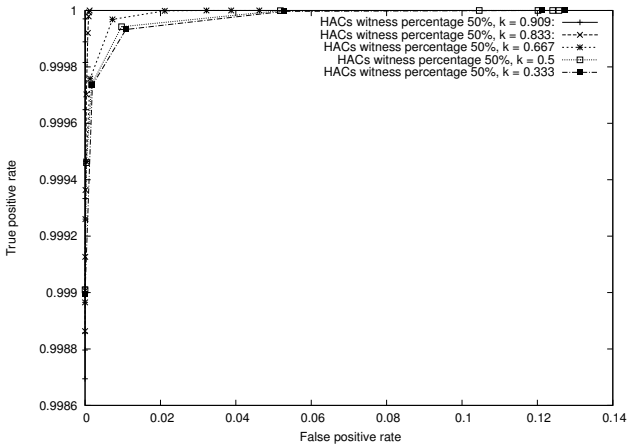

Figure 3: The effects of the percentage of honest active constituent (k value) in global population for various h.

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 *attackers* is 4.

Since we see in Figure 3 that the curve with parameter 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.

Graceful degradation is observed in an experiment with 1000 real constituents and 1000 MCMC rounds. Each attacker creates n fake identities ($n \in \{1, 2\}$). Percentages of HACs and attackers are both 50%. Each HAC correctly witnesses 1% of identities. Constituent C is counted if $\Psi(C)$ is greater than threshold $t \in [0, 1]$.

It can be shown that our proposal enables an extension of Pretty Good Privacy (PGP) authentication mechanism with new flexibility addressing its main recognized weaknesses.

1. REFERENCES

[1] M. C. Silaghi, K. Alhamed, O. Dhannoon, S. Qin, R. Vishen, R. Knowles, I. Hussien, Y. Yang, T. Matsui, M. Yokoo, and K. Hirayama. DirectDemocracyP2P — decentralized deliberative petition drives —. In *Proceedings of IEEE P2P*, Trento, September 2013.