

Brief Announcement: Secure Self-Stabilizing Computation

Shlomi Dolev
Department of Computer Science,
Ben-Gurion University of the Negev
Beersheva, Israel 84105
dolev@cs.bgu.ac.il

Karim Eldefrawy
SRI International
Menlo Park
California, USA
karim@csl.sri.com

Juan Garay
Yahoo Research
Sunnyvale
California, USA
garay@yahoo-inc.com

Muni Venkateswarlu
Kumaramangalam
Department of Computer Science,
Ben-Gurion University of the Negev
Beersheva, Israel 84105
muni@cs.bgu.ac.il

Rafail Ostrovsky
Department of Computer Science and
Department of Mathematics,
University of California
Los Angeles, USA
rafail@cs.ucla.edu

Moti Yung
Snap & Columbia University
New York, USA
moti@cs.columbia.edu

ABSTRACT

Self-stabilization refers to the ability of systems to recover after temporal violations of conditions required for their correct operation. Such violations may lead the system to an arbitrary state from which it should automatically recover. Today, beyond recovering functionality, there is a need to recover security and confidentiality guarantees as well. To the best of our knowledge, there are currently no self-stabilizing protocols that also ensure recovering confidentiality, authenticity, and integrity properties. Specifically, self-stabilizing systems are designed to regain functionality which is, roughly speaking, desired input output relation, ignoring the security and confidentiality of computation and its state. Distributed (cryptographic) protocols for generic secure and privacy-preserving computation, e.g., secure Multi-Party Computation (MPC), usually ensure secrecy of inputs and outputs, and correctness of computation when the adversary is limited to compromise only a fraction of the components in the system, e.g., the computation is secure only in the presence of an honest majority of involved parties. While there are MPC protocols that are secure against a dishonest majority, in reality, the adversary may compromise *all components* of the system for a while; some of the corrupted components may then recover, e.g., due to security patches and software updates, or periodical code refresh and local state consistency check and enforcement based on self-stabilizing hardware and software techniques. It is

currently unclear if a system and its state can be designed to *always* fully recover following such individual asynchronous recoveries. This paper introduces *Secure Self-stabilizing Computation* which answers this question in the affirmative. Secure self-stabilizing computation design ensures that secrecy of inputs and outputs, and correctness of the computation are automatically regained, even if at some point the *entire system is compromised*. We consider the distributed computation task as the implementation of virtual global finite satiate machine (FSM) to present commonly realized computation. The FSM is designed to regain consistency and security in the presence of a minority of Byzantine participants, e.g., one third of the parties, and following a temporary corruption of the entire system. We use this task and settings to demonstrate the definition of *secure self-stabilizing computation*. We show how our algorithms and system autonomously restore security and confidentiality of the computation of the FSM once the required corruption thresholds are again respected.

CCS CONCEPTS

• **Security and privacy** → **Privacy-preserving protocols**; *Public key encryption*; *Security requirements*;

KEYWORDS

Self-stabilization; Secure multi-party computation; Security and privacy;

ACM Reference format:

Shlomi Dolev, Karim Eldefrawy, Juan Garay, Muni Venkateswarlu Kumaramangalam, Rafail Ostrovsky, and Moti Yung. 2017. Brief Announcement: Secure Self-Stabilizing Computation. In *Proceedings of PODC '17, Washington, DC, USA, July 25-27, 2017*, 3 pages.
<https://doi.org/http://dx.doi.org/10.1145/3087801.3087864>

1 PRELIMINARIES

System and Network Model. We consider a system of n parties p_1, \dots, p_n with a completely connected semi synchronous network, that interactively and continuously computes a function $f(\cdot)$ over a finite field \mathbb{F} (below we use a Mealy FSM as an example of $f(\cdot)$). Each party has a True Random Number Generator (TRNG). To simplify

The first author is partially Supported by the Rita Altura Trust Chair in Computer Sciences, by Lynne and William Frankel Center for Computer Sciences, by a grant of the Ministry of Science, Technology and Space, Israel and the National Science Council (NSC) of Taiwan, and by a grant of the Ministry of Science, Technology and Space, Israel, and the Ministry of Foreign Affairs, Italy.

The fifth author's research is supported in part by NSF grant 1619348, US-Israel BSF grant 2012366, by DARPA Safeware program, OKAWA Foundation Research Award, IBM Faculty Research Award, Xerox Faculty Research Award, B. John Garrick Foundation Award, Teradata Research Award, and Lockheed-Martin Corporation Research Award. The views expressed are those of the authors and do not reflect position of the Department of Defense or the U.S. Government.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

PODC '17, July 25-27, 2017, Washington, DC, USA

© 2017 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-4992-5/17/07.

<https://doi.org/http://dx.doi.org/10.1145/3087801.3087864>

the presentation we first assume that the system is set up by a Configuration Authority (CA) that initializes each party p_i with a pair of public- and private-key, PK_{p_i} and SK_{p_i} , respectively; each party is also initialized with public-keys of other parties in the system. Each party's private-key, and the CA's public-key, PK_{CA} are stored in a secure hardware module that only allows the party to decrypt or to sign messages with its private-key, or to verify signatures by the CA using its public-key. Each pair of parties shares private secure authenticated communication channels that can be established based on the initial setup. In fact, the use of TRNG suffices to establish the required keys. Parties can repeatedly establish fresh public and private keys (unknown to the non-present adversary), and agree on the keys (by using Byzantine agreement over each declared public key), and then establish genuine symmetric keys for each pair of the participants and secure broadcast channels from scratch. Alternatively, parties may receive the symmetric keys as a part of the inputs.

The parties periodically receive some global input in secret shared format from an external client (or the environment in general), and use that input in the computation of $f(\cdot)$. The parties are not aware of their state history i.e., they do not know whether they are corrupted or not. The parties exchange messages in sequence of rounds to communicate and compute the function $f(\cdot)$. In each *round*, every party in the system receives messages, performs local computation and sends messages. We consider both synchronous and semi-synchronous network settings during the system implementation. In a synchronous system, all the n parties are synchronized via a global clock and perform a given event on a common clock pulse that occurs periodically. Whereas in a semi-synchronous system, time-bounded events happen in real-time. An event needs to be completed before its time elapses. A self-stabilizing Byzantine clock synchronization [7] can facilitate a (logical) global clock pulse when the system has only semi-synchronous promise. Moreover, it can be used to announce a common round number, that in-turn can be used to start establishing the keys and repeatedly perform the given global FSM transition.

Mixed Adversarial Model. There is a central *mixed adversary* with bounded computational power that is actively trying to prevent the success of the computation. The adversary may corrupt up to n parties corrupting some parties *passively*, to learn their internal information, and/or even *actively*, causing them to deviate arbitrarily from prescribed protocol steps. Because parties are periodically recovered/reset/rebooted to a consistent/clean state (using a hardware watchdog mechanism that is designed to automatically recover following any (time-count) corruption), the adversary is not able to maintain corruption of all the n parties indefinitely. An underlying assumption is that, once the adversary loses control over a party it is unable to immediately compromise the party again.

2 STANDARD SELF-STABILIZING VS SECURE SELF-STABILIZING COMPUTATION

Standard Self-Stabilization. A self-stabilizing system can be started in any possible global state (possibly due to the occurrence of transient unpredictable faults) and is guaranteed to converge to a consistent state and correct operation. If a system is self-stabilizing, then the system must satisfy the following properties [4]:

- **Convergence:** For any arbitrary inconsistent configuration c_r , there exists a safe (legal, arbitrary but consistent) configuration c_f after a finite number of rounds, i.e., starting from any arbitrary configuration, the system is guaranteed to eventually reach a stable legal configuration from which the execution is legal, fulfilling the task requirement.

- **Closure:** Once a system reaches a stable state, the system is guaranteed to stay in the stable state provided that no further unexpected fault happens.

Secure Computation. Secure Multi-Party Computation (MPC) enables n parties to jointly and securely compute a function $f(\cdot)$ of their private inputs, even in the presence of adversarial behavior (typically modeled as a fraction of t parties that can be passively or actively corrupted). MPC ensures that while computing $f(\cdot)$ and when corruption is less than t , parties inputs' remain private, except what is revealed by the output of the computation (*privacy*). MPC also guarantees the correctness of the output of the computation (*correctness*) [10].

Secure Self-Stabilizing Computation. Existing secure MPC protocols [5, 10, 11] guarantee privacy (computation state, inputs and outputs) and correctness of computation only when there are a limited number of corrupted parties. However, in reality, a powerful adversary may even compromise all the parties of a system. The aim of introducing *Secure Self-stabilizing Computation* is to specify the requirements for automatically recovering secrecy of computation state, inputs and outputs, and correctness of a computation even if an entire system is compromised for a finite time. Consider a self-stabilizing system of n parties that perform a computation in a distributed fashion while maintaining the input secrecy and correctness of the output. We assume that the system does not necessarily have an initial state to start with, i.e., it may begin its execution from any arbitrary configuration c_r and does not necessarily stop. An external client or environment, e.g., sensors or a physical process, etc., supplies inputs to the system and reconstructs output(s) from the individual outputs of the parties. Over a long period of time, the considered mixed adversary may corrupt a large number of the parties and may even compromise all the n parties for a finite time. When the system is fully exposed, after a finite number of rounds r_x , n_x parties may recover independently due to some proactive measures, such as security patches and software updates, anti-virus updates, periodical code refreshing based on self-stabilizing hardware and software techniques. From the above *closure property*, ensuring no further corruption, the asynchronous independent recovery establishes local consistency of the recovered parties (for example, see [2, 6] for a self-stabilizing infrastructure that guarantees a safe recovery of a party even in the presence of Byzantine faults). Due to the independent recovery and the *closure property*, the system gradually recovers *itself* autonomously to exhibit the needed behavior. From the *convergence* property, the system may converge to a safe arbitrary consistent configuration c_f , after r_f ($r_f \geq r_x$) rounds, recovering enough number of parties n_f , $n_f \geq n_x$, from Byzantine behavior. The independent recovery continues until the required secure computation requirements (thresholds) hold again. Once the threshold requirements are respected again, the system automatically regains confidentiality of state and inputs/outputs, and security of the computation.

3 SECURE SELF-STABILIZING COMPUTATION OF AN INTERACTIVE FSM

In a nutshell, we use a global clock (obtained by e.g., [7]) to establish a consistent state of the FSM among all the non-Byzantine participants, and then compute transitions using the (secret shared) inputs to obtain (secret shared) outputs. To achieve this, we first harvest randomness from the true random source that each participant owns. Next, we use a secure MPC protocol to check whether the state shares held by all the participants represent a legitimate state (using, say, the Berlekamp-Welch algorithm [1] to eliminate corrupted state shares contributed by the Byzantine participants). If the state is valid, then another MPC instance uses the (secret shared) state and the inputs (represented by secret shares, with the same redundancy against Byzantine participants as the state redundancy) to compute the (secret shares of the) next state and output. In-case the state is found invalid by the first MPC, is followed by assignment of state secret shares of a default state and then continues as in the previous case. The process is described in more details below.

Note that we have two kinds of rounds, one is a clock pulse, and the other is a round of the FSM transition, say starting whenever the clock pulse number is zero, the FSM transition round consists of so many clock pulses to allow (in short): (a) establishing new keys (a1. using true random source the parties obtain fresh public and private keys, a2. sending the public key of each participant to all other participants, a3. agree on each public key using Byzantine agreement [3, 9], a4. establishing a fresh symmetric key between any two participants), (b) using verifiable secret sharing scheme [8] on the secret shares of the current FSM state, (c) executing a secure multi-party computation for enforcing a correct FSM state, both in the polynomial degrees for the representation of each portion of the FSM state and the well formation of the state, (d) applying the transition table, that is shared (hardwired) among the participants, compute a circuit in secure MPC that will compute the next FSM state and the output.

Let F be a publicly known interactive Mealy finite state machine (FSM) $F = \{ST, S_0, \Gamma, \Lambda, T, O\}$ with a finite set of states ST , and initial state S_0 , a finite input alphabet Γ , a finite output alphabet Λ , a transition function $T : ST \times \Gamma \rightarrow ST$ mapping pairs of a state and an input symbol to the corresponding next state, and an output function $O : ST \times \Gamma \rightarrow \Lambda$ mapping pairs of a state and an input symbol to the corresponding output symbol. Computation of the FSM is delegated to the n parties. Every party has two arithmetic circuits (consisting of addition and multiplication gates) that represent the functions T and O , *Transition Circuit (TC)* and *Output Circuit (OC)* respectively. The input is distributed to the parties in a secret shared form [12], denoting $[x]_{p_i}$ as the share of party p_i of secret x . The parties are in agreement (and have necessary software) to execute an MPC protocol Π which can perform any computation represented as addition and multiplication gates using secret shared inputs. Parties are in possession of a circuit representation of a (fault tolerant) error detection algorithm (ξ), e.g., the Berlekamp-Welch algorithm [1], that if given evaluations of an univariate polynomial $P(\cdot)$ of degree t , can interpolate it as long as there are enough correct points given (we use ξ to securely verify that the secret encoded by some shares is one of the legitimate

states of F). Let S_0 also denote the default state of the FSM, while ζ denotes the current state of the FSM. Algorithm 1 is executed when each of the parties, p_i , receives a share of a new input (τ).

Algorithm 1: Secure Self-stabilizing Computation of FSM
 $F = \{ST, S_0, \Gamma, \Lambda, T, O\}$

- 1: Each p_i receives a share of the secret shared input $[\tau]_{p_i}$
 - 2: Each p_i obtains some random seed r_{p_i} from its TRNG
 - 3: Each p_i participates in Π to compute ξ with $[\tau]_{p_i}$ as input and r_{p_i} as its randomness
 - 4: **if** computation of ξ succeeds **then**
 - 5: Each p_i computes via Π the *TC* with $[\zeta]_{p_i}$ and $[\tau]_{p_i}$ and outputs $[T(\zeta, \tau)]_{p_i}$
 - 6: Each p_i computes via Π the *OC* with $[\zeta]_{p_i}$ and $[\tau]_{p_i}$ and outputs $[O(\zeta, \tau)]_{p_i}$
 - 7: **else**
 - 8: Each p_i computes via Π the *TC* with $[S_0]_{p_i}$ and $[\tau]_{p_i}$ and outputs $[T(S_0, \tau)]_{p_i}$
 - 9: Each p_i computes via Π the *OC* with $[S_0]_{p_i}$ and $[\tau]_{p_i}$ and outputs $[O(S_0, \tau)]_{p_i}$
 - 10: **end if**
-

Algorithm 1 ensures that whenever the conditions for stabilization hold, e.g., less than one third of the participants are Byzantine, the adversary's knowledge concerning the state of the FSM is lost, as the adversary cannot obtain the arriving input, and therefore cannot use the revealed state for computing the current state. In particular, when the graph of the transition function of the FSM is a complete graph, then in the eyes of the adversary any state is possible after the first transition following the system convergence. Similarly, when the transition function graph forms an expander, the knowledge concerning the FSM state is totally lost after a number of steps that is logarithmic in number of the FSM states.

REFERENCES

- [1] Elwyn R. Berlekamp. 1984. *Algebraic Coding Theory*. Aegean Park Press.
- [2] Alexander Binun, Mark Bloch, Shlomi Dolev, Ramzi Martin Kahil, Boaz Menuhin, Reuven Yagel, Thierry Coupaye, Marc Lacomte, and Aurélien Wailly. 2014. Self-Stabilizing Virtual Machine Hypervisor Architecture for Resilient Cloud. In *2014 IEEE World Congress on Services, SERVICES 2014, Anchorage, AK, USA, June 27 - July 2, 2014*. 200–207.
- [3] Ran Cohen, Sandro Coretti, Juan A. Garay, and Vassilis Zikas. 2016. Probabilistic Termination and Composability of Cryptographic Protocols. In *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part III*. 240–269.
- [4] Shlomi Dolev. 2000. *Self-Stabilization*. MIT Press.
- [5] Shlomi Dolev, Karim El Defrawy, Joshua Lampkins, Rafail Ostrovsky, and Moti Yung. 2016. Proactive Secret Sharing with a Dishonest Majority. In *Security and Cryptography for Networks - 10th International Conference, SCN 2016, Amalfi, Italy, August 31 - September 2, 2016, Proceedings*. 529–548.
- [6] Shlomi Dolev and Yinnon A. Haviv. 2012. Stabilization Enabling Technology. *IEEE Trans. Dependable Sec. Comput.* 9, 2 (2012), 275–288.
- [7] Shlomi Dolev and Jennifer L. Welch. 2004. Self-stabilizing clock synchronization in the presence of Byzantine faults. *J. ACM* 51, 5 (2004), 780–799.
- [8] Paul Feldman. 1987. A Practical Scheme for Non-interactive Verifiable Secret Sharing. In *Proceedings of the 28th Annual Symposium on Foundations of Computer Science (SFC '87)*. IEEE Computer Society, Washington, DC, USA, 427–438.
- [9] Juan A. Garay and Yoram Moses. 1998. Fully Polynomial Byzantine Agreement for $n > 3t$ Processors in $t + 1$ Rounds. *SIAM J. Comput.* 27, 1 (1998), 247–290.
- [10] O. Goldreich, S. Micali, and A. Wigderson. 1987. How to Play ANY Mental Game. In *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing (STOC '87)*. ACM, New York, NY, USA, 218–229.
- [11] Ueli Maurer. 2006. Secure multi-party computation made simple. *Discrete Applied Mathematics* 154, 2 (2006), 370–381. Coding and Cryptography.
- [12] Adi Shamir. 1979. How to Share a Secret. *Commun. ACM* 22, 11 (Nov. 1979), 612–613.