SAND2020-6401C

Sandia National Laboratories

The Novel Approaches to Anomaly Detection and Surety for Safeguards Data project investigates three core data analysis and management methods and their applicability for international safeguards: Distributed Ledger Technology (DLT) for data authentication, anomaly detection based on Grammar Compression (GC), and how operator data could assist in drawing safeguards conclusions in a Multi-Party Computation (MPC) environment.

DEVELOPMENT OF NOVEL APPROACHES TO ANOMALY DETECTION and SURETY FOR SAFEGUARDS DATA | Year One Results

Natacha Peter-Stein, David Farley, Constantin Brif, Nicholas Pattengale, Chase Zimmerman, Meghan Galiardi | ***Sandia National Laboratories***
Yifeng Gao, Jessica Lin | ***George Mason University***
Mitchell Negus, Rachel Slaybaugh | ***University of California at Berkeley***

Anomaly Detection Using Grammar Compression (GC)

The first approach is that of anomaly detection in data based on Grammar Compression — to provide a practical tool for effective and efficient detection of anomalies in multivariate time-series data obtained from fielded safeguards equipment.

Our efforts have been focused on development, implementation, and testing of four new capabilities to enhance the existing GC method: (1) robust, parameter-free anomaly detection by integrating GC with ensemble learning, (2) anomaly detection on extra-long scale (time series with millions of data points), based on efficient, variable-length motif discovery, (3) anomaly detection in video data, and (4) detection of correlated anomalies in multivariate data.

The standard GC requires preselecting values for at least two parameters at the discretization step. How to choose these parameter values properly is still an open problem. Instead of using a particular combination of parameter values for GC-based anomaly detection, our new method generates the final result based on a set of results obtained using an ensemble of GC algorithm executions with different parameter values. Numerical experiments performed on datasets with known ground truth showed that ensemble GC can outperform existing GC-based approaches with different criteria for selection of parameter values.

Dataset	Ensemble GC	GC-Random	GC-Fix	GC-Select	Discord
TwoLeadECG	0.3951	0.2873	0.0629	0.1663	**0.4931**
ECGFiveDay	0.3903	0.2988	0.2671	0.1050	**0.4794**
GunPoint	**0.4728**	0.3715	0.2411	0.0560	0.4000
Wafer	**0.3179**	0.2126	0.1382	0.2480	0.3090
Trace	**0.5718**	0.2022	0.3601	0.3408	0.2816
StarLightCurve	**0.9369**	0.6930	0.5301	0.8759	0.9161

Table 1. Performance evaluation results: Score averaged over 25 time series with randomly planted anomalies, for Ensemble GC and four baseline methods.

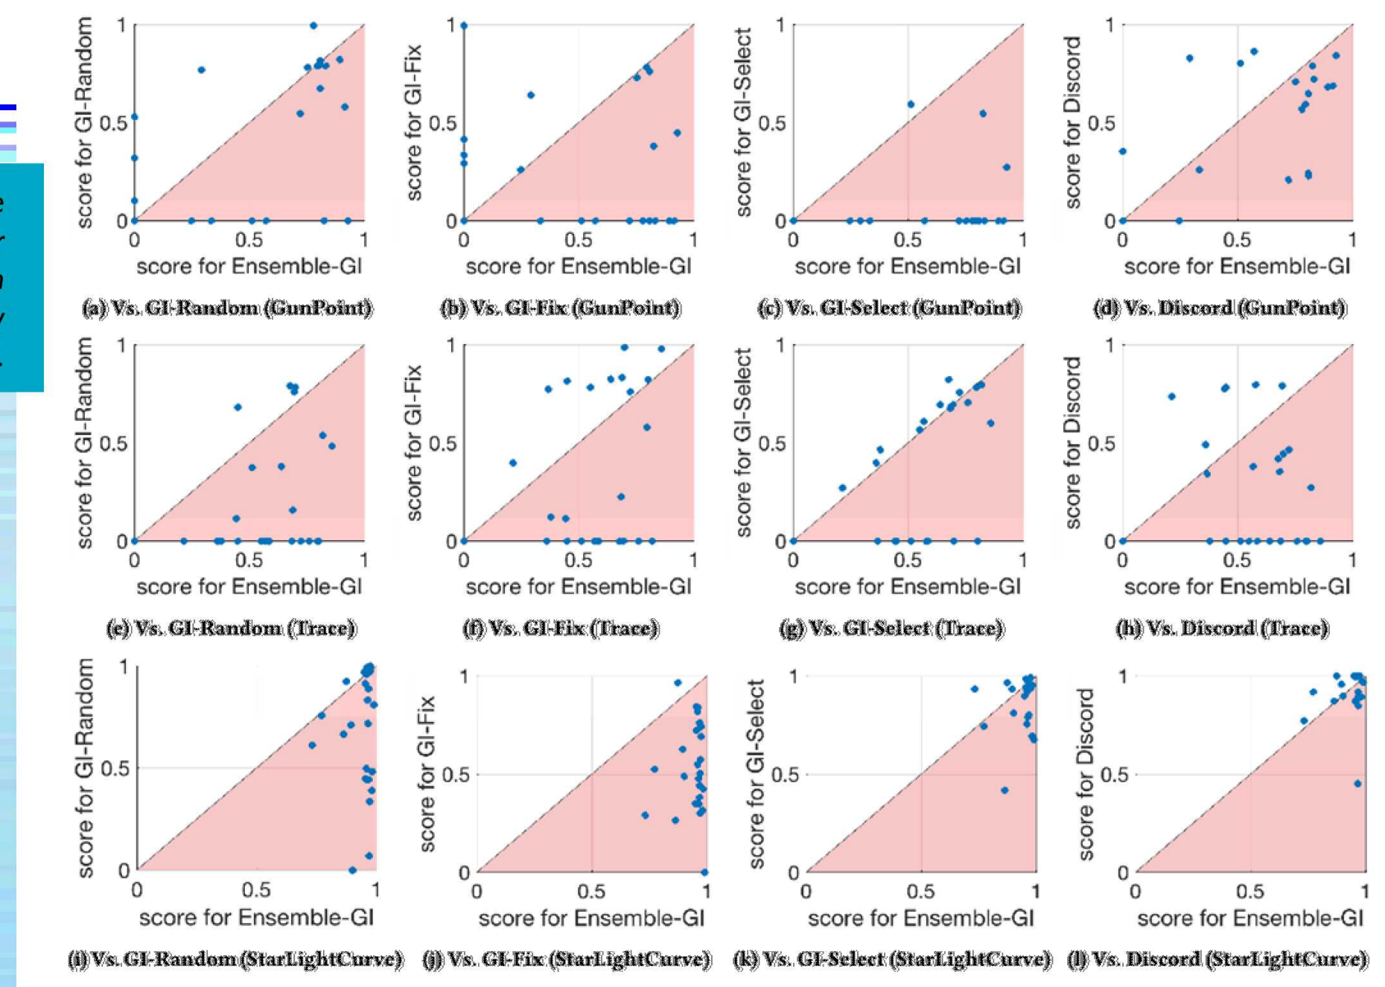

(a) Vs. GI-Random (GunPoint) (b) Vs. GI-Fix (GunPoint) (c) Vs. GI-Select (GunPoint) (d) Vs. Discord (GunPoint)
(e) Vs. GI-Random (Trace) (f) Vs. GI-Fix (Trace) (g) Vs. GI-Select (Trace) (h) Vs. Discord (Trace)
(i) Vs. GI-Random (StarLightCurve) (j) Vs. GI-Fix (StarLightCurve) (k) Vs. GI-Select (StarLightCurve) (l) Vs. Discord (StarLightCurve)

Figure 1. Performance evaluation results: Score values (blue points) for 25 time series with randomly planted anomalies for three datasets (GunPoint, Trace, and StarLightCurve). A point in the lower triangle corresponds to a superior performance by ensemble GC compared to the baseline method.

RESULTS:

We have prioritized anomaly detection methods that would be most valuable in extending and improving the capabilities of the existing GS method, and commenced the work on developing, implementing, and testing these new methods. One of our key results has been the development of a new method that combines GC with ensemble learning to perform robust and efficient anomaly detection in time series data.

Distributed Ledger Technology (DLT) For Data Provenance

The second approach is the creation of a pilot safeguards data tracking system based on recent advances in tracking using Distributed Ledger Technology to improve data authentication.

There are many practical barriers impeding the adoption of DLT: some technical, some policy, and some based on perception. We introduce and describe a framework by which adoption tradeoffs are being objectively evaluated.

To illustrate the DLT approach at a high level, consider the two alternate data topologies in Figure 2. In traditional practice (top of Figure 2), facilities submit NMC&A data to their State authority, typically within 15 days of the end of a month. The State rolls up and investigates facility reports, and submits to the IAEA, typically within 30 days of the end of a month. In the DLT, or shared data, approach (right of Figure 2), the same data (and business processes, and timelines) may be used, however there is only ever a single copy of the data – data exporting and importing is obviated.

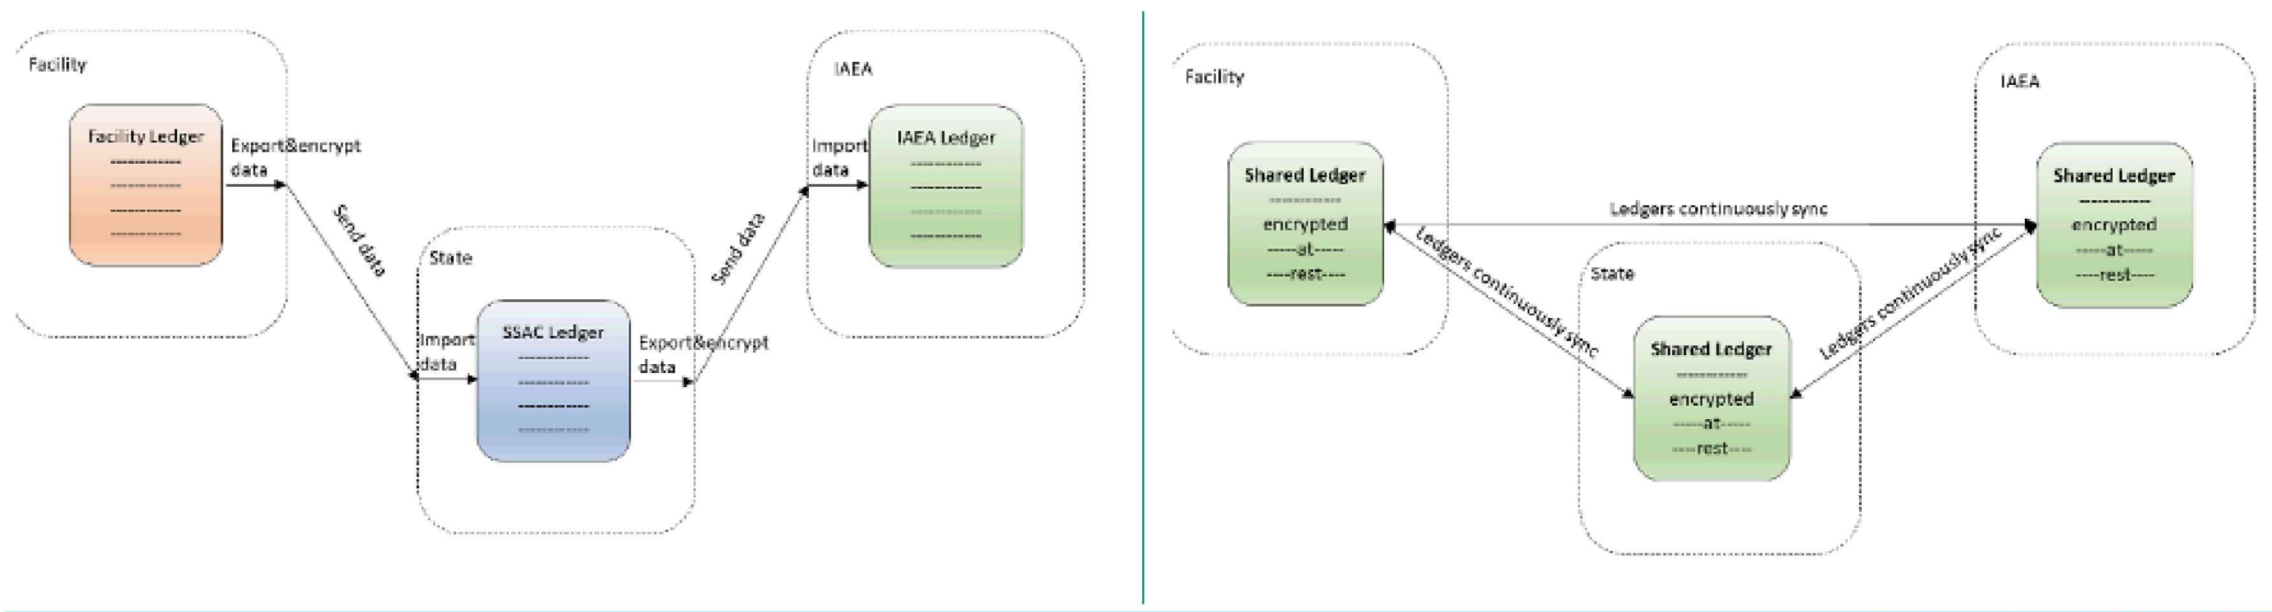

Figure 2. Traditional safeguards practice where organizations export and send data (left) versus a shared ledger concept where all organizations possess a replicated data store (right). Note that in the shared ledger concept, data is encrypted at rest, and decryption keys are only distributed to parties as needed. This enables, e.g. the IAEA, to possess shared data (in an inscrutable form) until deemed releasable by the State.

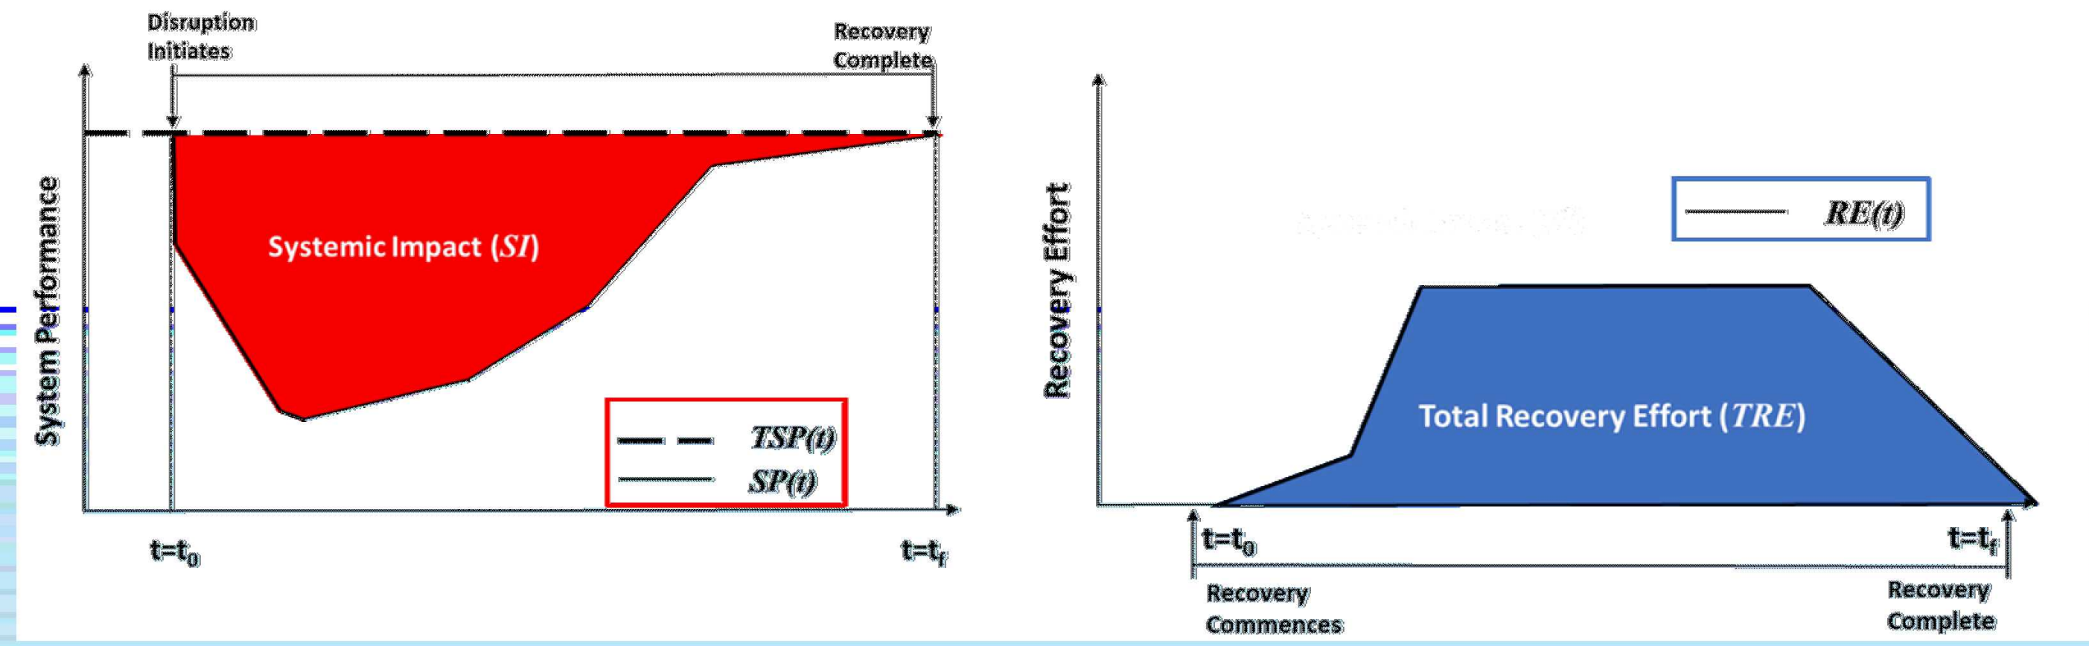

We evaluate the performance of a DLT versus traditional practice. Previous studies have examined metrics of effectiveness of traditional safeguards practice. However, in this work will evaluate using the metrics of resilience. There are many definitions of resilience, but we loosely define resilience as the ability of a system to continue to perform mission essential tasks despite the presence of a disruption or attack. Although related to metrics of effectiveness, the resilience metrics presented here have several advantages: they are quantitative, scenario-based, and temporal-based. This framework is currently being applied in an isolated experimentation network to evaluate resilience of DLT versus current safeguards practice across a variety of disruption scenarios, both cyber and physical.

RESULTS:

We have introduced and described a framework by which adoption tradeoffs of DLT for improved Continuity of Knowledge (CoK) are being objectively evaluated. The described resilience framework should be of independent interest, as it can be adapted to inform adoption tradeoffs of other technical proposals, as well as can quantify the resilience of a system (e.g., other aspects of current safeguards practice) in isolation.

Multi-Party Computation (MPC)

The third approach is the potential use of a Multi-Party Computation environment to examine how operator data, such as safety and physical protection systems' data, could augment traditional safeguards verification data. This could give confidence that anomalous activity within a facility can be attributed to safety or security activities rather than diversion or misuse at the facility. MPC is a method that would allow nuclear facility operators to contribute their data into a combined result with other parties' data, yet never exposing the underlying raw data.

For this effort we are only considering two parties: the IAEA and the State facility. For such a two-party application, the MPC approach of "garbled circuits" makes sense. With garbled circuits, the two participants each take on a role in generating or evaluating the garbled table, and then sharing the combined result with each other (although sharing the result is not required). Some of the advantages of the two-party garbled circuit approach are:

- Straight-forward implementation in Python or other programming language
- Two-party input and roles
- On-line, (near) real-time results
- Fast calculations possible, compared with other MPC implementations

We have incorporated the garbled circuit protocol, including Oblivious Transfer, into a Python library we call CypherCircuit. Two primary focus areas going forward are (1) to increase the speed (efficiency) of our CypherCircuit code, possibly including adjusting the protocol currently being used or perhaps re-programming into C/C++ rather than Python; and (2) investigate enhanced security beyond our current passive security paradigm.

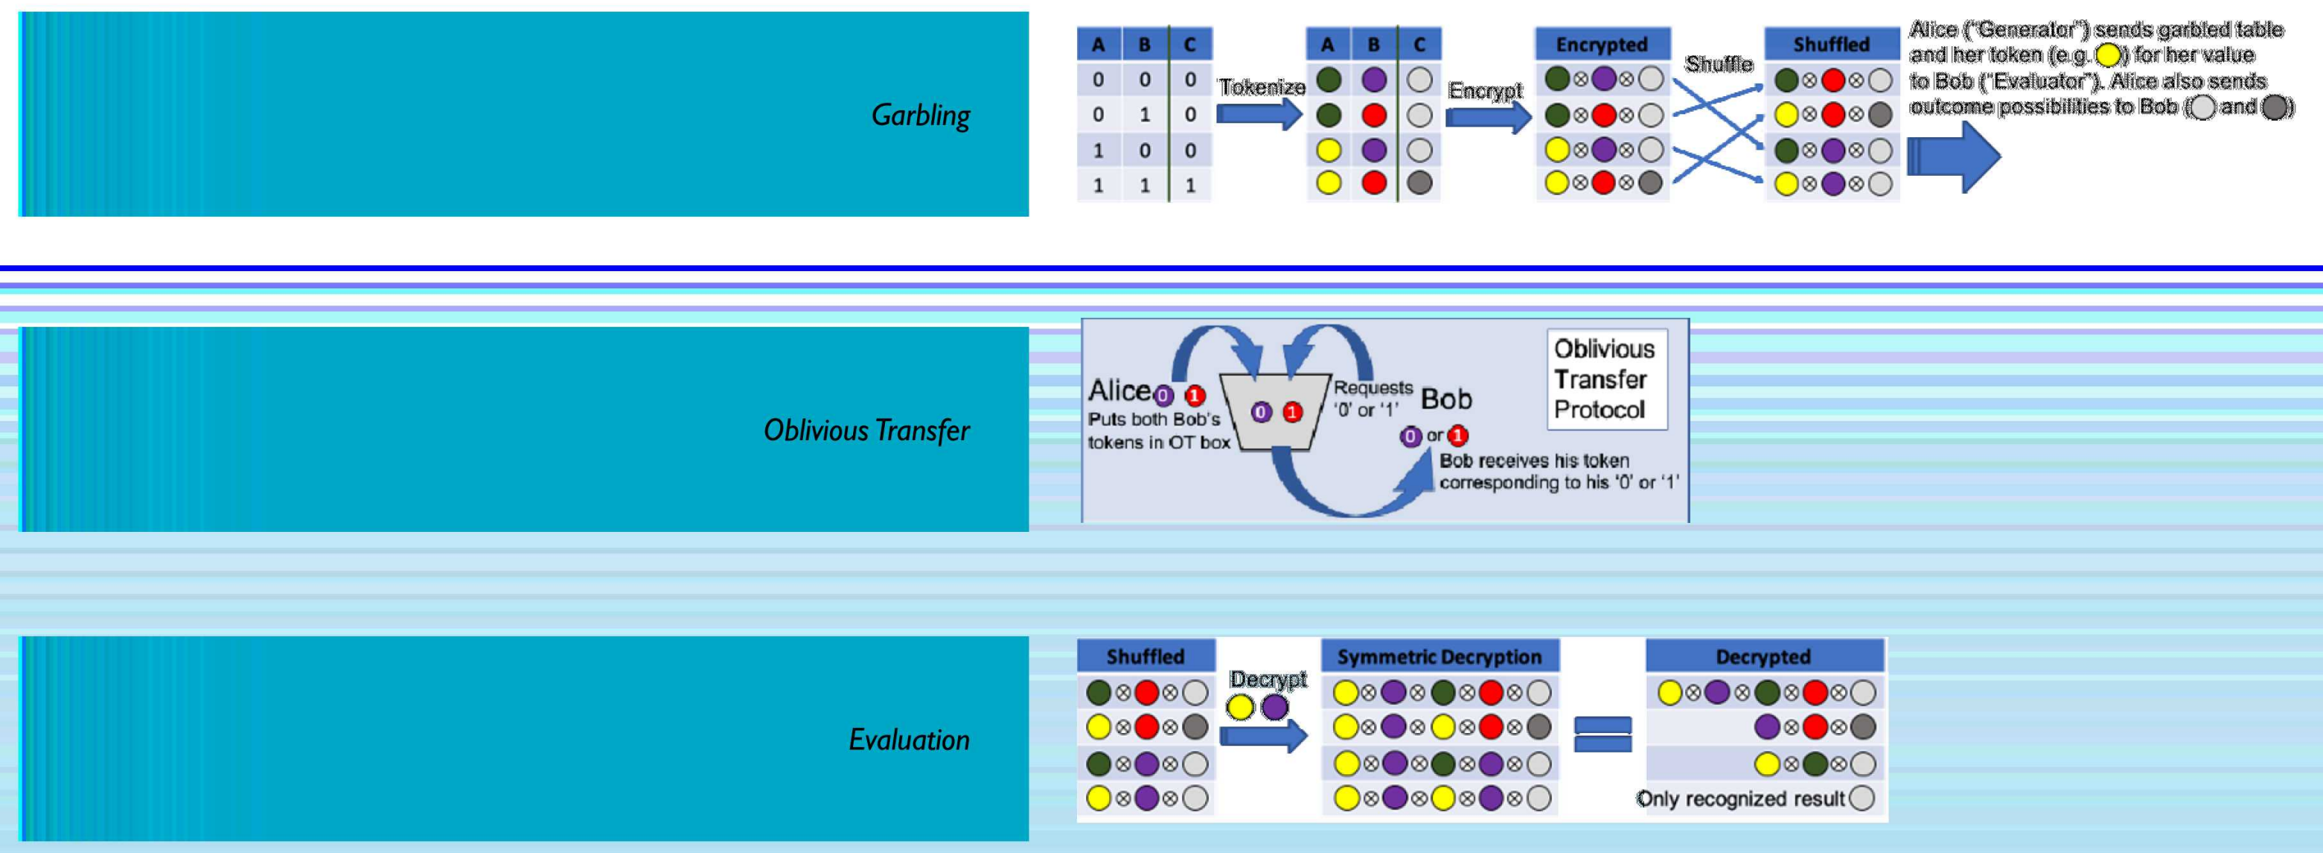
Garbling
A
B
C
0 0 0
0 1 0
1 0 0
1 1 1
Tokenize
Encrypt
Encrypted
Shuffle
Shuffled
Alice ("Generator") sends garbled table and her token (e.g. ○) for her value to Bob ("Evaluator"). Alice also sends outcome possibilities to Bob (○ and ●)
Oblivious Transfer
Alice
Puts both Bob's tokens in OT box
0 1
Requests '0' or '1'
Bob
0 or 1
Bob receives his token corresponding to his '0' or '1'
Oblivious Transfer Protocol
Evaluation
Shuffled
Decrypt
Symmetric Decryption
Decrypted
Only recognized result

Figure 3. Illustration of garbled circuit protocol. Generator ("Alice") creates random tokens for all input and output possibilities; Oblivious transfer is used to provide the Evaluator ("Bob") with his token; and Evaluator decrypts the table to identify the correct output token.

RESULTS:

The two-party formulation of MPC called "garbled circuits" can be used to enable sharing of nuclear data, potentially of safeguards relevance, towards calculating a function result without ever exposing the raw data to the other party. We have developed a Python library called CypherCircuit that automatically creates garbled circuits following the protocol of Kolesnikov, and are using the Oblivious Transfer protocol to provide the circuit evaluator with his decryption token corresponding to his input value.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology &Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.
SAND No. ____________

Sandia National Laboratories