Request for Information (RFI) IARPA-RFI-14-03

Security and Privacy Assurance Research—Multiparty Computation (SPAR-MPC)

Synopsis

The Intelligence Advanced Research Projects Activity (IARPA) is seeking information relevant to developing tools for automatically generating provably secure multiparty computation protocols given specifics of the intended applications. This request for information (RFI) is issued solely for information gathering and planning purposes; this RFI does not constitute a formal solicitation for proposals. The following sections of this RFI define the overall scope of the technical domain of interest, several technical areas in which specific information is sought, and instructions for the preparation and submission of responses.

Problem Background and Scope

Secure Multiparty Computation (MPC) is a solution to the very general problem in which parties A, B, … want to compute a joint function f (a, b, …) on their respective inputs, but no party wants to reveal its input to any other party. The function to be computed can be as simple as the value of a statistic on the data or as complex as the result of an interactive program. For example, MPC protocols can

- Allow a person or corporation to store and manipulate data in an untrusted cloud computing infrastructure as confidentially as if that data were stored in their own network,
- Allow a Web server to customize the presentation of data based on a user's identity, preferences, interests, or other private information, without learning that information, or
- Allow a group of users to determine a course of action that is optimal for the group as a whole, without requiring any user to reveal which course of action they personally prefer.

MPC protocols have been studied for a few decades now, with research continually improving the communication and computational efficiency for specific functions or security definitions [Y86, GMW86, GMW87].

The design of an MPC protocol for a specific application typically considers three main factors: a model of the available computation, an appropriate security threat model, and the specific computation to be performed. First, a computational model includes the number of participants, the computing architecture available to each participant (e.g., Central Processing Unit (CPU) speed, Field Programmable Gate Array (FPGA) chip size, amount of parallelism, and available storage), and the networking architecture (e.g., topology, latency, bandwidth, and the ability to broadcast or synchronize transmissions). Second, a security threat model includes the capabilities and potential actions of adversaries. An adversary may be an eavesdropper on the network, a dishonest participant in the computation, or an otherwise honest participant whose system has been compromised. Finally, the specific computation includes a description of the

function(s) to be evaluated. This could be, e.g., a mathematical expression, a program in C code, or a Boolean circuit.

The wide variety of choices for these factors makes it a challenge to construct an MPC protocol whose performance is optimal for a particular application. Integrating multiple MPC protocols or cryptographic primitives to create a novel solution to a larger secure computation problem is possible [KoSS09, KeSS13], but increases the complexity of the optimization problem even more. Protocol design involves cryptographic experts carefully defining message sequences, data structures, cryptographic primitives, and a host of other parameters that are tuned for a specific application. The design process is very sensitive to changes in the application requirements. It is by no means clear that a protocol designed to be highly efficient for one application will be even close to optimal for a slightly different application.

The most basic requirements for a secure computation are that each input is kept confidential to the party that provided it, and the result is learned only by the intended recipient. Any other information that can be learned by any party during the computation is likely to be a fatal security flaw for some application. Therefore, existing general-purpose MPC protocols go to extraordinary lengths to hide all other information. SPAR-MPC is based on the observation that if the "leakage" of some other information is acceptable or if there is an acceptable bound for the aggregate amount of this other information that can be learned over time, the protocol options may look very different and orders of magnitude gains in efficiency may be achievable. The expected efficiency benefits warrant exploring ways to specify allowed leakage for otherwise secure computations, an important deviation from standard cryptographic threat models for MPC.

End users would be well served to have a language in which to articulate allowed leakage to influence the protocol design process. This means moving protocol design and specification as close to the end application as possible with easy-to-use tools. Tools can help the end user make security priorities explicit, i.e., to decide whether an amount of leakage unacceptable as an average case is nevertheless acceptable as a low likelihood worst case.

The vision is to assist an organization with multiple MPC needs in selecting and deploying efficient protocols for each need, based on the characteristics of each application. MPC applications for the same end user organization might range from big data analysis in the cloud to set intersection on short lists of items. A user knowledgeable about an intended MPC scenario, but not an expert in cryptographic protocols or MPC, could input application data into a future tool which will select or generate the definition of a cryptographic protocol and a proof of the protocol's security or data that can be used to prove the security of the protocol. This does not preclude a generic proof that the generation tool will only produce provably secure protocols on well-formed parameters. The definition of the cryptographic protocol would be in a programming language or in another executable specification [LM03, MNPS04, BNP08, SKM11], able to be compiled to machine code with dependence only on available commercial or open source cryptographic libraries.

Note that the vision is not to have a tool that helps select from among predefined protocols, but to have a tool that can automatically generate tailored protocols that are very efficient for novel applications. This does not preclude the selection of an instance from within a parameterized

family of protocols, but it is desirable that a future tool would support a wide enough range of parameters that no single family of protocols could provide a practically efficient solution for all cases. The envisioned tool might present the user with different protocols that satisfy the application constraints but have different efficiency/security tradeoffs or are efficient under different computation models such as the addition of additional parties.

Note that, for circuits, a general purpose garbled circuit compiler [MNPS04, BNP08, HEKM11] could provide an upper bound of performance from which to measure the sophistication of tools that can take concrete application parameters into consideration. At the lower bound are the naïve protocols that offer no security or privacy. The research goal is to automatically generate tailored protocols that lie in between, and then help the end user make further choices by visualizing the tradeoffs and quantifying the overall cost of alternative protocols. This RFI seeks ideas that illuminate the research challenges in making this vision a reality.

Request to Respondents

In this RFI we seek information about approaches to develop tools that automatically generate provably secure MPC protocols from information about specific applications. To this end, we seek ideas on formalizing security and privacy constraints for MPC participants, formalizing the application characteristics such as tolerance for pre-computation, high bandwidth, or heavy computation loads, and measuring the tradeoffs between protocol efficiency and leakage.

It is very difficult to enumerate all the possible ways information can be leaked in a cryptographic protocol and the consideration of side channel leakage (execution timing, memory usage, CPU usage, etc.) makes this more difficult. It is also difficult to interpret the effects of such leakage on the confidentiality of the application data or the integrity of the overall computation. Formalization and quantification of information available in access patterns, side channels, and other kinds of indirect disclosures are necessary to help non-specialists draw boundaries between intended, acceptable, and strictly prohibited learning by each participant in a specific end-user application. This RFI seeks ideas for exploring, quantifying, interpreting, and articulating the tradeoffs between the efficiency gains and the negative side-effects from allowing a specific leakage in a way that empowers end users to make informed risk acceptance decisions for each particular MPC application.

Additionally, the following ideas are in scope of this RFI: support for verifiable computation [GGP10], function privacy [Y86], malicious adversaries (including a malicious networking infrastructure [C00]), reusable cryptography [GKPVZ13], computation of arbitrary Random Access Machine (RAM) [LO13] or Random Access Stored Program (RASP) machine programs, Fully Homomorphic Encryption (FHE) and Somewhat Homomorphic Encryption (SWHE) techniques [G09]. Restrictions to the two-party case are in scope, as this is a common scenario that may benefit from specialized tools.

Responses should address one or more of the following four topic areas:

1. Formalization of Security and Privacy Constraints for MPC Protocols

   This topic area especially includes identifying, formalizing, and quantifying leakage so that all information flows are explicit.

2. Quantitative Metrics for Security, Privacy, and Efficiency of MPC Protocols

   This topic includes identifying sound metrics for measuring progress and ultimate success of automatic protocol generation tools.

3. Parameterized Families of MPC Protocols

   This topic identifies and discusses current research in MPC protocols that are already parameterized or can be parameterized with respect to the security definition (beyond a simple security parameter), number of participants, available computational, communication, or storage resources, etc.

4. Research Roadmap

   This topic area includes discussing the challenges in developing tools for automatic generation of provably secure MPC protocols, tailored to concrete application parameters, and proposing a research roadmap for getting to this end state. A roadmap is a logical sequence of well-defined intermediate challenges that make cumulative progress toward the envisioned end state.

Responses may additionally discuss any other ideas noted as of interest in this RFI as far as the page limit allows.

Based in part on the results of this RFI, IARPA expects to hold a MPC workshop. Participants will be selected primarily from respondents to this RFI. Such a workshop would likely explore the research challenges of implementing MPC protocol generation tools as described above, particularly as they relate to the structure of a potential future IARPA research program in this area.

Preparation Instructions to Respondents

The Intelligence Advanced Research Projects Activity often selects its research efforts through the Broad Agency Announcement (BAA) process. This request for information is intended to obtain information relevant to a possible future IARPA program, so that feedback from potential participants can be considered prior to the issuance of a BAA. Reponses to this RFI may be incorporated within a future IARPA program BAA, and therefore any information provided must be available for unrestricted public distribution. Although the scope and purpose of any future workshop or program remain to be determined, secure MPC is a topic of strong interest. IARPA requests that submittals briefly and clearly address the topic of this RFI, identify critical technical risks and challenges, describe approaches to address those technical risks and challenges, discuss the range of practical applications that would be affected by program success,

and comment on the durability of the projected payoff. Respondents may also provide a non-proprietary rough order of magnitude (ROM) estimate of costs in terms of time, funding, and other resources expected to be needed to execute the proffered approach. This RFI contains all of the information needed to submit a response. No additional forms, kits, or materials are needed.

IARPA seeks responses from all capable and qualified sources from within and outside the United States. Responses from teams with complementary areas of expertise are encouraged. Responses have the following formatting requirements:

1. A one page cover sheet that identifies the number and title of the RFI responded to, the title of the response, the respondent's organization(s), technical and administrative points of contact (including names, addresses, phone numbers, fax numbers, and email addresses of all co-authors);
2. A one page substantive, focused executive summary;
3. Technical discussions of one or more of the four workshop topic areas described above, up to a maximum of five pages, including any figures and charts, with each topic area discussion clearly titled with the name of the relevant topic area;
4. Optionally, a set of briefing charts, one for each topic area for which a technical discussion was provided, each clearly labeled with the name of the relevant topic area and graphically depicting an overview of the key ideas discussed for a topic area;
5. A list of works cited (any significant claims or reports of success must be accompanied by citations, and all unpublished referenced works must be attached);

All parts of the submittal must be on letter size paper with a minimum 1 inch margin, with a minimum 12 point font size.

Submission Instructions to Respondents

Responses to this RFI are due no later than 4:00pm Eastern Daylight Time on 31 March 2014. All submittals must be electronically submitted to dni-iarpa-rfi-14-03@iarpa.gov in Portable Document Format (PDF) file format. Inquiries to this RFI must be submitted to dni-iarpa-rfi-14-03@iarpa.gov. Do not send questions with proprietary content. No telephone inquiries will be accepted.

DISCLAIMERS AND IMPORTANT NOTES

This RFI is issued solely for information and new program planning purposes and does not constitute a solicitation. IARPA is under no obligation to acknowledge receipt of the information received, or provide feedback to respondents with respect to any information submitted.

Responses to this RFI are not offers and cannot be accepted by the Government to form a binding contract. Respondents are solely responsible for all expenses associated with responding to this RFI. Each respondent is responsible for ensuring that any submitted material or results of research funded by another party has been approved for unrestricted public distribution by the funding organization.

The Government does not intend to award a contract on the basis of this RFI or to otherwise pay for the information requested, nor is the Government obligated to issue a solicitation based on responses received. **Neither proprietary nor classified information should be included in the submittal.** Input on technical aspects of a response may be obtained by IARPA from non-Government consultants who are bound by appropriate non-disclosure requirements.

Contracting Office Address

Office of the Director of National Intelligence
Intelligence Advanced Research Projects Activity
Washington, District of Columbia 20511
United States of America

Primary Point of Contact

W. Konrad Vesey
Program Manager
dni-iarpa-rfi-14-03@iarpa.gov

REFERENCES

[BNP08] Ben-David, Nisan, and Pinkas, "FairplayMP—A System for Secure Multi-Party Computation," Proceedings of the 2008 ACM Conference on Computer and Communications Security (CCS 2008)

[C00] Ran Canetti, "Security and Composition of Multiparty Cryptographic Protocols," J. Cryptology, v. 13 no. 1, 2000, pp. 143-202.

[G09] Gentry, "Fully Homomorphic Encryption Using Ideal Lattices," Proceedings of the 41st Annual ACM Symposium on Theory of Computing Conference (STOC 2009)

[GGP10] Gennaro, Gentry, and Parno, "Non-Interactive Verifiable Computing: Outsourcing Computation to Untrusted Workers," Advanced in Cryptology—CRYPTO 2010, 30th Annual Cryptology Conference

[GKPVZ13] Goldwasser, Kalai, Popa, Vaikuntanathan, and Zeldovich, "Reusable Garbled Circuits and Succinct Functional Encryption," Proceedings of the 45th Annual ACM Symposium on Theory of Computing Conference (STOC 2013)

[GMW86] Goldreich, Micali, and Wigderson, "Proofs that Reveal Nothing but Their Validity and a Methodology of Cryptographic Protocol Design," 27th Annual Symposium on Foundations of Computer Science (FOCS 1986)

[GMW87] Goldreich, Micali, and Wigderson, "How to Play any Mental Game or A Completeness Theorem for Protocols with Honest Majority," Proceedings of the 19th Annual ACM Symposium on Theory of Computing (STOC 1987)

[HEKM11] Huang, Evans, Katz, and Malka, "Faster Secure Two-Party Computation Using Garbled Circuits," Proceedings of the 20th USENIX Security Symposium, 2011

[KeSS13] Kerschbaum, Schneider, and Schröpfer, "Automatic Protocol Selection in Secure Two-Party Computation," 20th Annual Network and Distributed System Security Symposium (NDSS 2013)

[KoSS09] Kolesnikov, Sadeghi, and Schneider, "How to combine Homomorphic Encryption and Garbled Circuits: Improved Circuits and Computing the Minimum Distance Effectively," International Workshop on Signal Processing in the Encrypted Domain (SPEED 2009)

[LM03] Lewis and Martin, "CRYPTOL: High Assurance, Retargetable Crypto Development and Validation," White paper accessible at http://corp.galois.com/storage/files/downloads/Cryptol_Whitepaper.pdf

[LO13] Lu and Ostrovsky, "How to Garble RAM Programs," Advances in Cryptology—EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques

[MNPS04] Malkhi, Nisan, Pinkas, and Sella, "Fairplay—A Secure Two-Party Computation System," Proceedings of the 13th USENIX Security Symposium, 2004

[SKM11] Schröpfer, Kerschbaum, and Müller, "L1—An Intermediate Language for Mixed-Protocol Secure Computation," Proceedings of the 35th Annual IEEE International Computer Software and Applications Conference (COMPSAC 2011)

[Y86] Yao, "How to Generate and Exchange Secrets," 27th Annual Symposium on Foundations of Computer Science (FOCS 1986)