

Accessible Privacy-Preserving Web-Based Data Analysis for Assessing and Addressing Economic Inequalities*

Andrei Lapets
Boston University
lapets@bu.edu

Frederick Jansen
Boston University
fjansen@bu.edu

Kinan Dak Albab
Boston University
babman@bu.edu

Rawane Issa
Boston University
ra1issa@bu.edu

Lucy Qin
Boston University
lucyq@bu.edu

Mayank Varia
Boston University
varia@bu.edu

Azer Bestavros
Boston University
best@bu.edu

ABSTRACT

An essential component of initiatives that aim to address pervasive inequalities of any kind is the ability to collect empirical evidence of both the status quo baseline and of any improvement that can be attributed to prescribed and deployed interventions. Unfortunately, two substantial barriers can arise preventing the collection and analysis of such empirical evidence: (1) the sensitive nature of the data itself and (2) a lack of technical sophistication and infrastructure available to both an initiative's beneficiaries and to those spearheading it. In the last few years, it has been shown that a cryptographic primitive called secure multi-party computation (MPC) can provide a natural technological resolution to this conundrum. MPC allows an otherwise disinterested third party to contribute its technical expertise and resources, to avoid incurring any additional liabilities itself, and (counterintuitively) to *reduce* the level of data exposure that existing parties must accept to achieve their data analysis goals. However, achieving these benefits requires the deliberate design of MPC tools and frameworks whose level of *accessibility* to non-technical users with limited infrastructure and expertise is state-of-the-art. We describe our own experiences designing, implementing, and deploying such usable web applications for secure data analysis within the context of two real-world initiatives that focus on promoting economic equality.

CCS CONCEPTS

• **Security and privacy** → *Privacy-preserving protocols; Human and societal aspects of security and privacy*; • **Information systems** → Information systems applications;

KEYWORDS

secure multi-party computation, usability, web applications

*This material is based upon work partially supported by the NSF (under Grants #1430145, #1414119, #1718135, and #1739000) and the Honda Research Institutes.

Permission to make digital or hard copies of all or part 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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

COMPASS '18, June 20–22, 2018, Menlo Park and San Jose, CA, USA

© 2018 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-1-4503-5816-3/18/06...\$15.00

<https://doi.org/10.1145/3209811.3212701>

ACM Reference Format:

Andrei Lapets, Frederick Jansen, Kinan Dak Albab, Rawane Issa, Lucy Qin, Mayank Varia, and Azer Bestavros. 2018. Accessible Privacy-Preserving Web-Based Data Analysis for Assessing and Addressing Economic Inequalities. In *COMPASS '18: ACM SIGCAS Conference on Computing and Sustainable Societies (COMPASS)*, June 20–22, 2018, Menlo Park and San Jose, CA, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3209811.3212701>

1 INTRODUCTION

Initiatives that aim to address pervasive inequalities within a particular context or of a particular kind (such as wage inequality across race and gender, or inequality of opportunity for women- and minority-owned small businesses) often face an uphill battle due to a variety of factors. One way such an initiative can introduce a level of rigor and legitimacy is by attaining the ability to collect and analyze empirical evidence of both the status quo baseline for relevant metrics, as well as any improvements in those metrics over time that can be attributed to interventions that are being investigated, prescribed, or administered as part of the initiative. Paradoxically, the very same root cause being addressed by the initiative can introduce substantial or even insurmountable ethical, legal, and institutional barriers to such an evidence-based approach.

This quandary can be understood and examined by taking inventory of the parties involved, their concerns, their incentives, and the capabilities and resources available to them. In a typical scenario, an initiative involves at least two kinds of entities: the population of initiative beneficiaries, and the organizations that are spearheading the initiative with the purpose of collectively finding solutions. Both kinds of entities may be disincentivized to collect and analyze data. Beneficiaries may be unwilling to reveal sensitive data to the initiative; this may be especially true for beneficiaries that are members of a vulnerable population and do not want their data collected by third-party organizations. Initiative leaders may be concerned that collection and analysis of data that measures such inequalities might expose participants (or themselves) to negative publicity or risk of litigation. All parties may fear that inadvertent data exposure may be harmful (to beneficiaries, to participant organizations, and to the viability of the initiative as a whole) and that the liability associated with housing it is too burdensome.

Additionally, initiative leaders may not possess the technical skills and resources needed to effect a secure data collection and analysis workflow (e.g., they may be individuals from non-technical fields, human resources departments, social scientists, non-profits, government organizations, and so on). The initiative may have

to resort to using a third party that is tasked with building the appropriate system for securely collecting and storing the data, as well as with performing the desired analytics. However, this requires that all entities strongly trust this third party, presents yet another security risk, and may be costly if the third party wishes to be compensated for incurring liability on its own end. Additionally, this may further disincentivize beneficiaries from participating, since their sensitive data will be revealed to this third party.

Secure multi-party computation (MPC) is a cryptographic primitive for building privacy-preserving analytics over private data. MPC federates data collection and analysis among different parties so that results may be calculated and delivered while preserving the confidentiality of each contributor’s data. Theoretical constructs for MPC have been known for more than 35 years [4, 18, 32, 37]; performance improvements and software implementations in the past decade confirm their feasibility [7, 24, 30, 33, 36]. Available frameworks vary in the maturity of their software implementations, their security guarantees [29], and their communication models (*i.e.*, *peer-to-peer* or *client-server*). Additionally, MPC has been deployed to promote socially desirable outcomes such as fair auctions [8], tax fraud detection [6], and disease surveillance [16].

In the last few years, MPC solutions have been utilized successfully to provide a resolution to the conundrum faced by such initiatives. MPC allows otherwise disinterested third parties to contribute their technical expertise and resources without being privy to the input data or the results of the analysis. Thus, beneficiaries no longer need to trust the third party and the third party can avoid incurring any additional liabilities. However, achieving MPC’s benefits within initiatives that aim to address inequalities requires the deliberate design of highly *accessible* MPC tools and frameworks.

Our contributions. In this paper, we report our experiences designing, implementing, and successfully deploying web-based secure computation frameworks that support two real-world initiatives relying on secure data analysis and promoting economic equality: pay equity and business spending on local, minority-owned businesses. Our frameworks were informed by the interconnected usability, security, and legal requirements of the pay equity application and subsequently adapted to a second use case, demonstrating the ease of generalizing the frameworks to a variety of scenarios.

Our applications and frameworks are distinguished from those used in prior deployments of MPC in that they must drive participation via their *usability* and *accessibility* from the perspective of non-technical users with limited computational resources (*i.e.*, web browsers): users do not need to download an application package or configure a web server. The frameworks also explicitly support introducing resource asymmetries into secure data analysis protocols (in terms of the computations each party must perform). Thus, analysis applications can let resource-rich parties contribute their computational resources to do the heavy lifting, can let resource-constrained parties perform as little computation as possible, and can maintain desired security guarantees throughout the analysis.

2 PRELIMINARIES AND RELATED WORK

Secure MPC allows several parties to learn the output of a shared function while guaranteeing that all inputs remain private even against a colluding coalition of malicious parties. MPC relies on

a “secret-sharing” step in which the private input data of a participant is split into multiple shares (*i.e.*, pieces that on their own are indistinguishable from random noise), and each share is then sent to one of the participating parties. MPC provides generic protocols that transform any typical data analysis program into an equivalent distributed version that operates on the secret shares; each stage of this program (when executed in a distributed manner across all parties) produces a new set of secret shares that represent the intermediate result. At the end, the secret shares of the final output(s) are sent out to the designated parties to allow them to reconstruct the final result. MPC provides mathematically rigorous security guarantees that ensure that (1) the initial input remains private throughout the computation because any subset of parties only sees shares of it and (2) participants can be certain that the distributed MPC protocol was executed correctly by all parties.

The past several years have seen several successful deployments of MPC [6, 8, 13] and the creation of various software frameworks for MPC development and deployment. These frameworks range from proprietary implementations [7, 23] to open-source, proof-of-concept work [3, 9, 14, 15, 24, 26, 30, 36]. They can be divided into two settings: *peer-to-peer* and *client-server*. The *peer-to-peer* setting assumes symmetry in the capabilities and roles of all participating parties: all parties must deploy the software framework on continuously available servers, must remain online throughout protocol execution, and must perform the same amount of computation and communication. Most frameworks in this category are open-source research prototypes. In many cases, the framework’s authors explicitly discourage use in production [14, 24, 26, 30, 36].

In the *client-server* model, one entity acts as a service provider and the other parties can be resource-constrained clients. Unfortunately, this setting has received less attention in the community, with the few available frameworks limited in their flexibility. Work by Schröpfer et al. [31] allows two clients to perform a secure two-party computation using browsers while leveraging a web server for communication and code delivery. The framework is closed-source and restricted to two parties. Canon-MPC [19] offers a web-based system that supports MPC with symmetric binary functions. However, Canon-MPC’s code is delivered as a compiled binary and runs in Google’s proprietary NativeClient, which reduces accessibility and auditability. Canon-MPC also forces all parties to actively participate in the computation and may require an additional round to finish the session if some parties did not submit data.

Existing MPC frameworks fall short of supporting our setting along one or more of three critical dimensions: low-resource computing, accessibility, and usability. Initiative participants often do not possess adequate computational resources or technical expertise, particularly in cases where beneficiaries are members of underserved communities. This is especially problematic because MPC protocols are computationally more expensive than non-secure counterparts and require technical expertise to design and deploy. Successful MPC solutions must be accessible to beneficiaries and must not require any setup (*e.g.*, they must run within a browser as typical web applications) or specialized software or hardware dependencies. These requirements influenced the design of our MPC frameworks, which are, to the best of our knowledge, the only MPC frameworks that are *deployed* in a low-resource user community. We discuss these requirements in greater detail in Section 3.

3 ACCESSIBLE MPC FRAMEWORKS

Successful MPC-enabled data analysis applications must satisfy certain requirements in order to drive participation from data contributors [5]. These requirements are particularly critical in the context of societal initiatives serving communities that have limited technical resources and expertise. *Computational asymmetry*: the stakeholders in an initiative may require assistance from third parties because they must perform as little computation as possible on their end. Solutions should allow expensive parts of the computation to be delegated to the third parties without compromising the security of the computation. *Asynchronicity*: resource constraints may make it impossible for certain parties to be online throughout the computation. Such parties should be able to join and leave a computation at any time while relying on third parties to keep the service available throughout. In our applications, only a single entity (a service provider) must remain online; all other parties can contribute inputs once and remain offline during the rest of the analysis computation. *Idempotence*: contributors must be able to update their data if they discover the data they submitted was corrupted due to human error or software failures. Additionally, parties with more technical resources and expertise should be able to leverage these resources and expertise to help recover from failures and data corruption events without being able to observe or access any of the original input data. *Accessibility*: software solutions must be easily deployable, and should require no setup or specialized software or hardware dependencies. In our applications, users only require a web browser. *Comprehensibility*: MPC protocols must be simple enough to explain to potential users who do not possess technical knowledge, so that initiative leaders and beneficiaries can be confident in the security guarantees of the analysis.

3.1 Web-MPC Data Aggregation Framework

Web-MPC [35] leverages a variant [22] of standard additive secret sharing [10] to enable data analyses that do not require data contributors to reveal their data. This protocol allows a large number of parties with limited computational resources to aggregate their data with the help of at least one party that can operate a web server. Web-MPC supports aggregate analysis of tabular numeric data (supplied by users in the familiar format of a spreadsheet as shown in Figure 1) and responses to multiple-choice questions. Correlations between responses to multiple-choice questions can also be aggregated using the same additive secret sharing approach.

The protocol utilized by Web-MPC groups participating entities into three main roles:

- *contributors* who submit their private data for computation;
- the *service provider* that performs computations necessary for the analysis without being able to observe any of the individual inputs (in our deployment scenarios, this role was assumed by our institution);
- the *analyzer* who receives the result of the analysis computation (in our scenarios, these are the organizations spearheading their respective initiatives).

The described architecture was informed by the resource limitations of both the BWWC and the initiative participants [5]. In addition to its reliance on additive secret sharing, the security of the protocol leverages standard public-key encryption schemes that provide

IND-CPA security; concretely, RSA is used in the implementation [25]. The protocol is secure in the honest-but-curious model [17]: the service provider’s view of the random masks is protected using the analyzer’s public key and the analyzer never sees the individual masked data values unless it violates its promise not to collude with the service provider. The honest-but-curious model is appropriate when the service provider and analyzer have no incentives to collude (*e.g.*, if both are vested in an initiative’s success).

3.2 JavaScript Implementation of Federated Functionalities (JIFF)

JIFF [34] is a general-purpose MPC framework that generalizes Web-MPC along two dimensions: (1) it supports arbitrary computations (as opposed to only aggregation) and (2) it allows designers to customize an application to match the specific resource profiles of participants by defining the roles they play in the computation [21]. JIFF utilizes a server to store and route encrypted messages sent between the various participating parties. The server may also participate in the computation itself to improve performance and usability. The library supports asynchronous and asymmetric computations: parties are not required to remain online throughout the computation, can leave and join a computation dynamically, and can execute customized portions of the computation according to their resources and roles. JIFF assumes the honest-but-curious security model, and uses threshold Shamir’s Secret Sharing [32]. It is compatible with browsers and can run as a Node.js application.

4 APPLICATIONS

4.1 100% Talent Compact

In 2013, Boston Mayor Thomas Menino established a taskforce to advance the interests of women in the workforce, leading to the creation of the Boston Women’s Workforce Council (BWWC) by Mayor Marty Walsh in 2014. The BWWC initiated the 100% Talent Compact to partner with local government and private businesses in addressing issues impacting women in the workplace. Participating companies were required to help analyze wage gaps [12] by anonymously reporting salary data broken down by gender, ethnicity, and level of employment (similarly to the standard EEO-1 form [11]). Initially, no third party was willing to undertake the risk of receiving the raw salary data in order to enable the analysis.

In theory, the City of Boston or BWWC could have implemented an MPC solution such as Web-MPC on their own. However, they lacked the software engineering resources to build and deploy a platform, and did not possess the subject-area expertise needed to explain or execute the cryptographic protocol. The comprehensibility of the additive secret sharing protocol used by Web-MPC played a critical role in our and BWWC’s ability to explain to non-technical participants that each company could contribute cumulative employee earnings to the city-wide analysis, that our institution could provide and manage the computational resources and MPC data analysis application, and that only *aggregate* data would be revealed by the output. Prior work discusses our experiences communicating with participating organizations about the protocol’s capabilities [5, 22]. Ultimately, the BWWC was able to analyze aggregate data across companies while keeping company-specific data private during three successful deployments in 2015, 2016, and 2017.

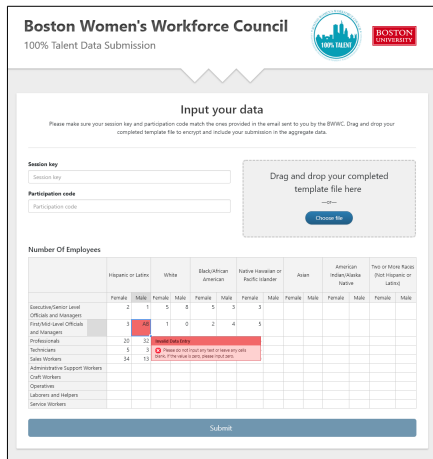


Figure 1: Web-MPC user interface at <https://100talent.org>.

In 2017, the BWWC was able to analyze aggregate data representing 166,705 employees across 114 companies, comprising roughly 16% of the Greater Boston area workforce. The data showed that the gender gap in the Boston area is even larger than previously estimated by the U.S. Bureau of Labor Statistics. Their analysis concluded that white women were earning 75 cents relative to every dollar that a white man earns. Asian women, black women, and Hispanic women made 71, 52, and 49 cents, respectively relative to every dollar made by a white male. The full report [12] documents in more detail the nuances of the study. By having actual employer-reported wage data, these results are a more accurate reflection of the magnitude of the salary inequities in the Greater Boston area. Since these analyses will be repeated at regular intervals, broad comparisons can be made from year to year. The analysis also enables granular insights into how employees of different races and job categories are affected by wage inequities. These outcomes are then used by the BWWC to drive discussions and workshops aimed at helping advance women of all backgrounds in the workplace.

4.2 Pacesetters Initiative

The Greater Boston Chamber of Commerce (GBCC) launched the Pacesetters Initiative [28] in January 2018 [27]. This initiative aims to leverage the purchasing power of large and mid-sized companies to create and promote economic opportunities for local minority-owned businesses. By tracking spending patterns and awarded contracts of participating members over time, the GBCC can measure the initiative’s impact on supplier diversity practices, including ways to increase spending with minority-owned businesses.

During a one-week period in March 2018, nine Pacesetters participated in the data analysis effort. We enabled the aggregation of four data points measuring spend and interactions with local minority-owned businesses; these were later shared by the GBCC with initiative participants. The Pacesetters Initiative will continue to promote equitable spending for minority-owned businesses amongst its members, and these initial results will be used as a baseline to which future analyses can be compared to assess progress.

For this initiative, we deployed a modified version of the Web-MPC platform that was used in the 100% Talent Compact data

analysis. After the many training and feedback sessions we held with participants, we customized the user interface to better accommodate the definitions and number of data points to be analyzed. A large amount of time and effort was spent to ensure every participant felt comfortable interacting with the system. Due to the anonymity guarantees of MPC, mistakes made by any individual affect the overall analysis results in an irreversible way and recovering from incorrect data submissions is non-trivial (and often impossible). Nevertheless, during the collection it became clear that at least one of the data contributors ran into a limitation of our platform. The upper bound for input values was $2^{32} - 1$, which caused an error for one of the participants. We were able to recover from this problem without violating the privacy of all the other submissions (and by learning no additional information about the affected participant other than this lower bound) by changing the finite additive group used for additive secret sharing to $\mathbb{Z}/2^{40}\mathbb{Z}$ in the middle of the analysis session. This was accomplished by rebuilding all secret shares *under MPC* using the general-purpose features of JIFF. While this unexpected issue leaked the lower bound of a data point for one participant, we note that this lower bound can be inferred from publicly available earnings reports.

Future analyses are planned at six month intervals, growing the group of Pacesetters as the initiative continues. One future goal is to perform a more granular analysis to better understand how to improve diversity spending, going beyond simple aggregation to more advanced statistical methods running under MPC.

5 CONCLUSIONS AND FUTURE WORK

By designing accessible MPC-enabled data analysis libraries and applications, we have allowed a third party (our institution) to contribute its resources to empirical studies of inequality without compromising the privacy of data contributors. Our work has received public praise from the BWWC. Speaking on NPR’s OnPoint, BWWC co-chair Evelyn Murphy stated that our framework constitutes a first step in “show[ing] how to use sophisticated computer science research for public programs” [1]. Christina Knowles, executive director of the BWWC [2, 20]: “[We] spent more than a year brainstorming with global experts in fruitless pursuit of a data-gathering method that would ensure employers’ confidentiality. It proved impossible to find a solution — until we were introduced to [the authors] who [were] absolutely vital to our work. The project is the first of its kind in the country and we owe our progress on this innovative and groundbreaking project [to them].”

A framework for privacy-preserving web-based data analysis that properly balances security and usability has far-reaching potential for public initiative research studies. Government agencies, non-profit organizations, and social scientists can identify analytics with social value and address legal and economic barriers to participation, while software engineers and security experts can design a technically sound data analysis application combining cloud computing and thin clients (even if vulnerability of participants is inversely proportional to computing power). Many potentially interesting avenues of research lie ahead in working to understand and model how the availability and accessibility of contemporary cryptographic techniques such as MPC changes the incentives, concerns, and limitations within the research and policy landscape.

REFERENCES

- [1] Tom Ashbrook. 2016. Will Data Help Close the Gender Pay Gap? NPR OnPoint radio show, WBUR 90.9. (March 2016). Retrieved May 8, 2018 from <http://www.wbur.org/onpoint/2016/03/30/gender-gap-pay-gap-boston-amazon>.
- [2] Laura Bassett. 2016. Apple, Facebook, Other Major Companies Commit to Paying Women the Same as Men. HuffPost News. (August 2016). Retrieved May 8, 2018 from http://www.huffingtonpost.com/entry/apple-facebook-obama-pay-gap_us_57bf44ede4b02673444f228f.
- [3] Assaf Ben-David, Noam Nisan, and Benny Pinkas. 2008. FairplayMP: A System for Secure Multi-party Computation. In *Proceedings of the 15th ACM Conference on Computer and Communications Security (CCS '08)*. ACM, New York, NY, USA, 257–266. <https://doi.org/10.1145/1455770.1455804>
- [4] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. 1988. Completeness Theorems for Non-cryptographic Fault-tolerant Distributed Computation. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing (STOC '88)*. ACM, New York, NY, USA, 1–10. <https://doi.org/10.1145/62212.62213>
- [5] Azer Bestavros, Andrei Lapets, and Mayank Varia. 2017. User-Centric Distributed Solutions for Privacy-Preserving Analytics. *Commun. ACM* 60, 2 (January 2017), 37–39. <https://doi.org/10.1145/3029603>
- [6] Dan Bogdanov, Marko Jöemets, Sander Siim, and Meril Vaht. 2015. *How the Estonian Tax and Customs Board Evaluated a Tax Fraud Detection System Based on Secure Multi-party Computation*. Springer Berlin Heidelberg, Berlin, Heidelberg, 227–234. https://doi.org/10.1007/978-3-662-47854-7_14
- [7] Dan Bogdanov, Sven Laur, and Jan Willemsen. 2008. Sharemind: A Framework for Fast Privacy-Preserving Computations. In *Computer Security - ESORICS 2008 (Lecture Notes in Computer Science)*, Sushil Jajodia and Javier Lopez (Eds.), Vol. 5283. Springer Berlin Heidelberg, Berlin, Heidelberg, 192–206. https://doi.org/10.1007/978-3-540-88313-5_13
- [8] Peter Bogetoft, Dan Lund Christensen, Ivan Damgård, Martin Geisler, Thomas P. Jakobsen, Mikkel Kroigaard, Janus Dam Nielsen, Jesper Bus Nielsen, Kurt Nielsen, Jakob Pagter, Michael I. Schwartzbach, and Tomas Toft. 2009. Secure Multiparty Computation Goes Live. In *Financial Cryptography and Data Security (Lecture Notes in Computer Science)*, Roger Dingledine and Philippe Golle (Eds.), Vol. 5628. Springer Berlin Heidelberg, Berlin, Heidelberg, 325–343. https://doi.org/10.1007/978-3-642-03549-4_20
- [9] Martin Burkhart, Mario Strasser, Dilip Many, and Xenofontas Dimitropoulos. 2010. SEPIA: Privacy-preserving Aggregation of Multi-domain Network Events and Statistics. In *Proceedings of the 19th USENIX Conference on Security (USENIX Security '10)*. USENIX Association, Berkeley, CA, USA, 15–15.
- [10] Chris Clifton, Murat Kantarcioglu, Jaideep Vaidya, Xiaodong Lin, and Michael Y. Zhu. 2002. Tools for Privacy Preserving Distributed Data Mining. *SIGKDD Explor. Newsl.* 4, 2 (December 2002), 28–34. <https://doi.org/10.1145/772862.772867>
- [11] U.S. Equal Employment Opportunity Commission. 2018. EEO-1 Survey. (2018). Retrieved May 8, 2018 from <https://www.eeoc.gov/employers/eo1survey/index.cfm>.
- [12] Boston Women's Workforce Council. 2018. Boston Women's Workforce Council Report 2017. (January 2018). Retrieved May 8, 2018 from https://www.boston.gov/sites/default/files/document-file-01-2018/bwvc_2017_report.pdf.
- [13] Ivan Damgård, Kasper Damgård, Kurt Nielsen, Peter Sebastian Nordholt, and Tomas Toft. 2017. Confidential Benchmarking Based on Multiparty Computation. In *Financial Cryptography and Data Security*. Springer Berlin Heidelberg, Berlin, Heidelberg, 169–187. https://doi.org/10.1007/978-3-662-54970-4_10
- [14] Daniel Demmler, Thomas Schneider, and Michael Zohner. 2015. ABY - A Framework for Efficient Mixed-Protocol Secure Two-Party Computation. In *22nd Annual Network and Distributed System Security Symposium (NDSS 2015)*. The Internet Society. <https://doi.org/10.14722/ndss.2015.23113>
- [15] Yael Eijenberg, Moriya Farbstein, Meital Levy, and Yehuda Lindell. 2012. SCAPI: The Secure Computation Application Programming Interface. Cryptology ePrint Archive, Report 2012/629. (November 2012). Retrieved from <https://eprint.iacr.org/2012/629>.
- [16] Khaled El Emam, Jun Hu, Jay Mercer, Liam Peyton, Murat Kantarcioglu, Bradley Malin, David Buckridge, Saeed Samet, and Craig Earle. 2011. A secure protocol for protecting the identity of providers when disclosing data for disease surveillance. *Journal of the American Medical Informatics Association* 18, 3 (May 2011), 212–217. <https://doi.org/10.1136/amiajnl-2011-000100>
- [17] Oded Goldreich. 2004. *The Foundations of Cryptography - Volume 2, Basic Applications*. Cambridge University Press, New York, NY, USA. <https://doi.org/10.1017/CBO9780511721656>
- [18] Oded Goldreich, Silvio Micali, and Avi Wigderson. 1987. How to Play ANY Mental Game or A Completeness Theorem for Protocols with Honest Majority. In *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing (STOC '87)*. ACM, New York, NY, USA, 218–229. <https://doi.org/10.1145/28395.28420>
- [19] Ayman Jarrous and Benny Pinkas. 2013. Canon-MPC, a System for Casual Non-interactive Secure Multi-party Computation Using Native Client. In *Proceedings of the 12th ACM Workshop on Workshop on Privacy in the Electronic Society (WPEC '13)*. ACM, New York, NY, USA, 155–166. <https://doi.org/10.1145/2517840.2517845>
- [20] Katie Johnston. 2015. Mayor Walsh pushes to gather data on gender wage gap. The Boston Globe. (April 2015). Retrieved May 8, 2018 from <https://www.bostonglobe.com/business/2015/04/07/walsh-women-venture-capital-summit-says-female-staff-members-underpaid/nKlivDh1VtOCV8hwVJ5XNM/story.html>.
- [21] Andrei Lapets, Mayank Varia, Azer Bestavros, and Frederick Jansen. 2017. Role-Based Ecosystem Model for Design, Development, and Deployment of Secure Multi-Party Data Analytics Applications. Cryptology ePrint Archive, Report 2017/803. (August 2017). Retrieved from <https://eprint.iacr.org/2017/803>.
- [22] Andrei Lapets, Nikolaj Volgushev, Azer Bestavros, Frederick Jansen, and Mayank Varia. 2016. Secure MPC for Analytics as a Web Application. In *2016 IEEE Cybersecurity Development (SecDev)*, 73–74. <https://doi.org/10.1109/SecDev.2016.027>
- [23] John Launchbury, Iavor S. Diatchki, Thomas DuBuisson, and Andy Adams-Moran. 2012. Efficient Lookup-table Protocol in Secure Multiparty Computation. In *Proceedings of the 17th ACM SIGPLAN International Conference on Functional Programming (ICFP '12)*. ACM, New York, NY, USA, 189–200. <https://doi.org/10.1145/2364527.2364556>
- [24] Chang Liu, Xiao Shaun Wang, Kartik Nayak, Yan Huang, and Elaine Shi. 2015. OblivM: A Programming Framework for Secure Computation. In *2015 IEEE Symposium on Security and Privacy*, 359–376. <https://doi.org/10.1109/SP.2015.29>
- [25] Alfred J. Menezes, Scott A. Vanstone, and Paul C. Van Oorschot. 1996. *Handbook of Applied Cryptography* (1st ed.). CRC Press, Inc., Boca Raton, FL, USA.
- [26] Kartik Nayak, Xiao Shaun Wang, Stratis Ioannidis, Udi Weinsberg, Nina Taft, and Elaine Shi. 2015. GraphSC: Parallel Secure Computation Made Easy. In *2015 IEEE Symposium on Security and Privacy*, 377–394. <https://doi.org/10.1109/SP.2015.30>
- [27] Greater Boston Chamber of Commerce. 2018. GBCC Launches Pacesetters Initiative Aimed at Uniting the Business Community's Response to Economic Inclusion. (January 2018). Retrieved May 8, 2018 from <http://bostonchamber.com/about-us/media-center/gbcc-launches-pacesetters-initiative>.
- [28] Greater Boston Chamber of Commerce. 2018. Pacesetters Initiative. (January 2018). Retrieved May 8, 2018 from <http://bostonchamber.com/programs-events/pacesetters>.
- [29] Jason Perry, Debayan Gupta, Joan Feigenbaum, and Rebecca N. Wright. 2014. Systematizing Secure Computation for Research and Decision Support. In *Security and Cryptography for Networks*, Michel Abdalla and Roberto De Prisco (Eds.). Springer International Publishing, Cham, 380–397. https://doi.org/10.1007/978-3-319-10879-7_22
- [30] Aseem Rastogi, Matthew A. Hammer, and Michael Hicks. 2014. Wysteria: A Programming Language for Generic, Mixed-Mode Multiparty Computations. In *2014 IEEE Symposium on Security and Privacy*. IEEE Computer Society, Washington, DC, USA, 655–670. <https://doi.org/10.1109/SP.2014.48>
- [31] Axel Schroepfer and Florian Kerschbaum. 2011. Demo: Secure Computation in JavaScript. In *Proceedings of the 18th ACM Conference on Computer and Communications Security (CCS '11)*. ACM, New York, NY, USA, 849–852. <https://doi.org/10.1145/2046707.2093509>
- [32] Adi Shamir. 1979. How to Share a Secret. *Commun. ACM* 22, 11 (November 1979), 612–613. <https://doi.org/10.1145/359168.359176>
- [33] Ebrahim M. Songhori, Siam U. Hussain, Ahmad-Reza Sadeghi, Thomas Schneider, and Farinaz Koushanfar. 2015. TinyGarble: Highly Compressed and Scalable Sequential Garbled Circuits. In *2015 IEEE Symposium on Security and Privacy*. IEEE Computer Society, Washington, DC, USA, 411–428. <https://doi.org/10.1109/SP.2015.32>
- [34] Multiparty.org Development Team. 2018. JavaScript Implementation of Federated Functionalities. (May 2018). Retrieved May 8, 2018 from <https://github.com/multiparty/jiff>.
- [35] Multiparty.org Development Team. 2018. Web-MPC. (March 2018). Retrieved May 8, 2018 from <https://github.com/multiparty/web-mpc>.
- [36] VIFF Development Team. 2009. VIFF, the Virtual Ideal Functionality Framework. (December 2009). Retrieved May 8, 2018 from <http://viff.dk/>.
- [37] Andrew C. Yao. 1982. Protocols for Secure Computations. In *Proceedings of the 23rd Annual Symposium on Foundations of Computer Science (SFCS '82)*. IEEE Computer Society, Washington, DC, USA, 160–164. <https://doi.org/10.1109/SFCS.1982.88>