Report from Dagstuhl Seminar 16051

Modern Cryptography and Security: An Inter-Community Dialogue

Edited by

Kristin Lauter[1], Radu Sion[2], and Nigel P. Smart[3]

1 **Microsoft Research – Redmond, US,** klauter@microsoft.com
2 **National Security Institute – Stony Brook, US,** radu@digitalpiglet.org
3 **University of Bristol, GB,** nigel@cs.bris.ac.uk

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 16051 "Modern Cryptography and Security: An Inter-Community Dialogue".

Seminar January 31 to February 5, 2016 – http://www.dagstuhl.de/16051
1998 ACM Subject Classification E.3 Data Encryption, K.6.5 Security and Protection
Keywords and phrases anti-surveillance/anti-censorship systems, homomorphic encryption, post-quantum cryptography, secure hardware design, secure multi-party computation, secure outsourcing, side/covert channels, virtualization security
Digital Object Identifier 10.4230/DagRep.6.1.207

1 Executive Summary

Nigel P. Smart

License Creative Commons BY 3.0 Unported license
© Nigel P. Smart

The seminar aimed to bring together communities with different backgrounds and form a bridge between them.

The outcomes ranged from a series of bridging exercises where participants summarized the current thoughts in existing areas; these included areas such as

- Hardware Attacks: Where we summarized the known attacks in this space.
- Computing on Encrypted Data: Various aspects of this were discussed, including Secure Guard Extensions (SGX), Searchable Symmetric Encryption (SSE), Multi Party Computation (MPC), and Fully Homomorphic Encryption (FHE).

We then went on to discuss more technical aspects, rather than just summarizing work,

- Cyberphysical Systems and IoT: Where the research challenges of performing work in this new area were discussed. A reliance on practical experimental was noted in the current research landscape.
- Mass Surveillance, Trapdoors, Secure Randomness: The recent "backdooring" of the DUAL–EC random number generator formed the background of this discussion. The seminar examined different aspects of this area, both in preventing, creating and detecting backdoors.
- Anonymous Payment Systems: This was a rather broad discussion which examined a number of issues around payments in general, and how cryptography could solve address these issues.

Except where otherwise noted, content of this report is licensed
under a Creative Commons BY 3.0 Unported license

Modern Cryptography and Security: An Inter-Community Dialogue, *Dagstuhl Reports*, Vol. 6, Issue 1, pp. 207–223
Editors: Kristin Lauter, Radu Sion, and Nigel P. Smart
DAGSTUHL REPORTS Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

We also discussed aspects related to the process of research in this field. In particular focusing on the problem of the lack of expository writing. Here we identified a number of disincentives in the research culture which prevents the creation of more discursive writing and expository articles. A number of solutions both existing, and proposed, were discussed to solve this issue. In another small breakout we discussed the lack of incentives to work on the underlying hard problems upon which our security infrastructure rests.

In summary the seminar found more problems with our current research trends, than solutions.

2 Table of Contents

3 Format

Unlike many Dagstuhl seminars this seminar was run on a very different format, with most time devoted to small group discussion, and one-on-one meetings. We thus reduced the total amount of "plenary" time to the minimum. The goal was to foster a dialogue between people working in two distinct but related fields, each with their own methodologies of working and presenting results.

The programme started with a series of one-on-one discussions, followed by a reporting back phase. This phase was to break the ice between participants and get participants to understand the area of research of someone from a very different background. The second phase consisted of the group selecting some common themes from the initial phase and then engaging in breakout discussions, followed by a series of plenary reports back. It is these reports back which we summarize in the abstracts contained in this document. As such the abstracts represent the combined brain storming of all the seminar participants.

This more interactive format was found to be highly successful by the participants, although very tiring, as it required concentration at all points during the week with little down time to "zone out" during someone else's talk. All the plenary sessions were highly interactive with questions and answers coming from the floor, with the main speaker purely leading the discussion.

As can be seen from the abstracts herein, we eventually discussed a wide variety of topics from the mechanics of how our science is performed, through to detailed discussions on specific technical topics. There is no doubt that a number of new collaborations and contacts ensued from the programme, and we hope this more intense style format can be adopted by other seminars at Schloss Dagstuhl in future.

4 Working groups

4.1 Hardware Attacks: Threat Models for Secure Hardware

Ferdinand Brasser (TU Darmstadt, DE), Raad Bahmani (TU Darmstadt, DE), Dieter Gollmann (TU Hamburg-Harburg, DE), Florian Kerschbaum (SAP SE – Karlsruhe, DE), Yongdae Kim (KAIST – Daejeon, KR), Kristin Lauter (Microsoft Research – Redmond, US), and Radu Sion (National Security Institute – Stony Brook, US)

License 🅭 Creative Commons BY 3.0 Unported license
© Ferdinand Brasser, Raad Bahmani, Dieter Gollmann, Florian Kerschbaum, Yongdae Kim, Kristin Lauter, and Radu Sion

The recent developments in the area of secure hardware, in particular the introduction of Intel's Software Guard Extensions (SGX), has yield the question under which condition secure hardware can be useful. To answer this question a threat model for secure hardware is required. This document provides an (incomplete) discussion of different classes and implementation of secure hardware with regard to a set of attack vectors. Attack vectors for secure hardware can be divided into two main groups, software attacks and hardware attacks.

4.1.1 Software Attacks

Software attacks can be carried out without physical proximity to the target system. Fault injection attacks aim to bring a secure hardware system into an invalid state to extract secret information, certain Smartcards are known to be vulnerable; secrets from the Smartcard can be extracted through sequences of interactions with the Smartcard's interface.

Side-channels exist due to the use of shared resources. In the case of SGX, the caches of the CPU are used by the isolated environment (called enclave) and untrusted software on the same system, hence, SGX is vulnerable to cache side-channel attack. However, cache side-channels are dependent on the software executed in the enclave and can be countered by using side-channel resilient algorithm.

Memory access pattern might also leak information about the internal state of a SGX enclave. A malicious OS could observe all memory access of an enclave at page granularity. This attack can be countered by side-channel resilient algorithm, too.

TrustZone can be implemented in a way that cache side-channels are not possible. Cache flushes on transition between normal world and secure world render those attacks ineffective. Page fault side-channel are non-existing in TrustZone due to the fact that the secure world is in charge of handling page faults itself.

Dedicated secure hardware systems, like HSMs and Smartcards, do not share resources with untrusted software and are therefore not vulnerable to software-exploitable side-channels.

4.1.2 Hardware Attacks

Hardware attacks can be further divided into invasive and non-invasive attacks. Physical side-channel, like power consumption, heat, radio emission, etc., are non-invasive and can leak information about secret information processed inside secure hardware.

Protection methods against those attacks exist, however, they are specific to individual attacks. Hence, to achieve comprehensive protection secure hardware needs to implement mechanism against each possible side-channel. Although some HSMs and Smartcards are know to provide certain protections mechanism it is not possible to make general statements about entire classes of devices.

SGX and TrustZone do not provide explicit protection methods against hardware side-channels and therefore must be assumed to be vulnerable.

Destructive physical attacks, like etching of layers of hardware to extract keys strode in hardware, can again be countered by explicit methods. Processors with SGX or TrustZone are produces with state-of-the-art production methods loading to very dense designs impeding those attacks or making them extremely expensive.

Hardware trojan are another threat to secure hardware against which protection methods exist. However, the implementation of hardware trojans requires significant resources on the attackers side (e.g., to manipulate to production process of the hardware). Given that easier attack vectors exist for most secure hardware systems its more likely that an attacker would exploit those.

4.1.3 Conclusion

A general threat model for secure hardware cannot be constructed due to the diversity of secure hardware solution. Even within a class of hardware systems (e.g., Smartcards) the in-homogeneity forbids general statements.

When using secure hardware available solutions must be evaluated against individual requirements.

4.2 Cyberphysical Systems and IoT Security

Dieter Gollmann (TU Hamburg-Harburg, DE), Alex Biryukov (University of Luxembourg, LU), Marc C. Dacier (QCRI – Doha, QA), George Danezis (University College London, GB), Yevgeniy Dodis (New York University, US), Christian Grothoff (INRIA – Rennes, FR), Stefan Katzenbeisser (TU Darmstadt, DE), Yongdae Kim (KAIST – Daejeon, KR), Moni Naor (Weizmann Institute – Rehovot, IL), Claudio Orlandi (Aarhus University, DK), Andreas Peter (University of Twente, NL), and Martina Angela Sasse (University College London, GB)

License 🅭 Creative Commons BY 3.0 Unported license
© Dieter Gollmann, Alex Biryukov, Marc C. Dacier, George Danezis, Yevgeniy Dodis, Christian Grothoff, Stefan Katzenbeisser, Yongdae Kim, Moni Naor, Claudio Orlandi, Andreas Peter, and Martina Angela Sasse

A CPS is a "physical and engineered systems whose operations are monitored, coordinated, controlled and integrated by a computing and communication core" [NSF]. The focus is on control and physical impact. IoT is about the networking of "things" (is a person a thing?), not necessarily via the internet. The focus is on networking.

CPS and IoT security more than a new playground for old techniques? The field is industry and business driven; performance drives demand for networking; previously air-gapped systems are opened to the Internet, but systems are there for 10+ years, old hardware, no patches for 10+ years. Incidents include Stuxnet, using legal commands to centrifuges to gradually decrease their performance, switching off a heart pacemaker by sending the heartbeat signal of a healthy person, and destroying a fail-safe pump by turning it on and off at a frequency too high to be picked up by the safety mechanism. Conclusion: CPS security is more than adapting and deploying familiar security mechanisms.

Distinguishing features? The interplay of safety and security: "fail-safe" systems depend on physical assumptions an attack may break; most work on safety builds on a world model that uses probabilities; this is not warranted in security. Inputs from and attacks on the physical layer need to be integrated in the security design of a system. Input validation is not sufficient, as shown by the pacemaker example. The attacker needs to understand how to manipulate inputs so that the system changes state in a way desired by the attacker. Challenge: One needs to understand the physical processes. Finding failure conditions is hard, process models of engineers cover how systems behave in normal operation but not necessarily in extreme situations.

Which methodology to apply to make progress in CPS security? Current research is hugely experimental, indicative of an early phase of a new field? Simulators, but where to get models and data from? Real data on real systems may be company secrets and not be made available to researchers. Does this matter? Testbeds, allow to observe real physical impact, but are limited to subsystems. Fake products widen the gap between model and reality. How to distinguish fakes from genuine items? How to guarantee that hardware or software was not modified/backdoored? Maybe one can build on the physical structure of objects; there is also the issue of supply-chain management. Final note: Is CPS security a matter of research or a matter of education? Security people need to learn about chemical plants, power grids, and other critical infrastructures. Operators need to learn thinking "security".

4.3 Mass Surveillance, Trapdoors, Secure Randomness

Nadia Heninger (University of Pennsylvania – Philadelphia, US), Alexandra Boldyreva (Georgia Institute of Technology – Atlanta, US), Nikita Borisov (University of Illinois – Urbana Champaign, US), Marc C. Dacier (QCRI – Doha, QA), Yevgeniy Dodis (New York University, US), Stefan Katzenbeisser (TU Darmstadt, DE), Kenneth G. Paterson (Royal Holloway University of London, GB), and Andreas Peter (University of Twente, NL)

License 🅭 Creative Commons BY 3.0 Unported license
© Nadia Heninger, Alexandra Boldyreva, Nikita Borisov, Marc C. Dacier, Yevgeniy Dodis, Stefan Katzenbeisser, Kenneth G. Paterson, and Andreas Peter

4.3.1 Problem

We began by discussing what the scope of mass surveillance is. Surveillance is directed at a target population, and "mass" means the surveillance is directed at the vast majority of the targeted population. This can be done by both governments and companies.

We know that mass surveillance is happening on medium and large scales. More "local" levels of mass surveillance include surveillance of cell phone towers, IMSI catchers, local ISPs, or smaller countries performing surveillance of connections between their countries and the rest of the world.

We discussed the model by which countries would man-in-the-middle connections transiting borders. For HTTPS, there are several examples of countries obtaining fraudulent HTTPS certificates for companies, and using false DNS records or similar to redirect vulnerable traffic to middleboxes who can then impersonate the end site to the users within the country.

Global-scale issues include backdoors being built into communications and cryptographic infrastructure. We discussed the specific case of the Dual-EC DRBG, where the construction of the random number generator allows an entity who generates the input parameters adversarially to recover the state and future outputs from a single output. The standards are known to be influenced by the NSA and GCHQ, and contain a recommended set of parameters that was generated by the NSA, instead of specifying that users generate their own. Additionally, there are widespread rumors from participants on standards committees of agency interference to weaken cryptographic standards. We learned in December 2015 that Juniper used this random number generator in NetScreen products, with parameters they generated themselves and feeding the output into another PRNG that should hide the direct output. However, an unknown party replaced the parameters with different parameters, and the implementation contained a subtle bug that caused the direct output of the RNG cascade to be raw Dual-EC output.

4.3.2 Solutions and ideas for man-in-the-middling

For local levels of surveillance, there are existing technical solutions that can detect and sometimes prevent some kinds of surveillance.

For example, for the problem of countries man-in-the-middling connections that transit fixed exit points, HTTPS with certificate pinning, certificate transparency, and two-factor authentication works in practice. Solutions like certificate pinning don't scale to all types of hosts. If there were a worldwide public-key infrastructure for clients, then TLS could do a two-way authenticated key exchange, but this is not in wide use for clients. Additional client authentication tends to take place after a one-way authenticated TLS session is established, and the client authenticates inside of the channel via a password or temporary code provided via a two-factor authentication device.

We talked about ways to make these schemes more cryptographically integrated to prevent the man-in-the-middle from successfully authenticating as a client even if they are successful in impersonating the site to the client, for example performing a challenge-response protocol combining a host's public key with a secret obtained out-of-band via two-factor authentication. This does not seem to be possible to implement securely in a browser user interface, because a naive implementation would allow the MITM to display a UI element to the host to enter a secret, and there is no way to validate the page/Javascript source implementing this. HTTPS certificate information is also not accessible via the page DOM. It may be possible to implement such a scheme using a browser extension and a TLS cipher suite with PSK authentication.

We also discussed the approach of using plausible deniability to increase the workload of an entity performing mass surveillance. Users could produce a large amount of spurious or cover traffic to try to confuse automated tests. However, it's unclear whether this would actually deter those implementing these schemes, or merely cause more false positives.

4.3.3 Ideas for crypto backdoors

We discussed ways of detecting cryptographic backdoors. We know that theoretical "kleptography" systems can be perfect: one can design a cryptographically undetectable backdoor. But implementors are imperfect. In practice, OpenSSL failed to implement the Dual EC DRBG correctly. (Yet it still passed FIPS certification, raising some questions.) Mistakes in implementations might reveal the presence of backdoors.

Additionally, the Dual EC backdoor is not cryptographically hidden: we can see that the parameters can be backdoored, even if we can't prove that a particular set of parameters was maliciously generated. For this case, there is a known bias that might allow discovery of this traffic in a black-box way given a large quantity of traffic, but for other implementations it's unclear what can be done.

Companies might want to prove to clients that implementations do not contain backdoors. They can publish open-source code, but customers don't have assurance that the code corresponds to the binaries they download. The Tor Project has been doing work on deterministic builds to allow external parties to verify this.

There has been some theoretical work on immunizing schemes against backdoors, or designing cryptographic schemes that cannot be backdoored. This is an area of current and future research.

4.4 Incentivizing Expository Writing

Aaron Michael Johnson (NRL – Washington, US), Allison Bishop (Columbia University – New York, US), Alexandra Boldyreva (Georgia Institute of Technology – Atlanta, US), Nikita Borisov (University of Illinois – Urbana Champaign, US), George Danezis (University College London, GB), Krista Grothoff (GNUNet e. V. – Rennes, FR), Nadia Heninger (University of Pennsylvania – Philadelphia, US), Sarah Meiklejohn (University College London, GB), and Radu Sion (National Security Institute – Stony Brook, US)

License Ⓒ Creative Commons BY 3.0 Unported license
© Aaron Michael Johnson, Allison Bishop, Alexandra Boldyreva, Nikita Borisov, George Danezis, Krista Grothoff, Nadia Heninger, Sarah Meiklejohn, and Radu Sion

4.4.1 Problem

Expository writing helps consolidate research knowledge and communicate it to groups outside of the community of active researchers. However, there is a general lack of incentives for high-quality research exposition of security and cryptography research. The problem is particularly acute for cryptography because it has a higher technical barrier to entry. As a result, research fields are unnecessarily difficult for new researchers to enter, and their results challenging for outside communities to make use of.

Several kinds of expository writing is being undervalued by the security and cryptography research communities. These include (i) systematizations of recent results for other researchers, (ii) writing for practitioners (e.g. people who implement systems, companies looking to commercialize technology), and (iii) popularization of research results for a general audience. Research surveys and lecture notes are at least generally recognized by researchers as having some value, but writing for practitioners or a general audience is infrequently rewarded, and these latter two types of writing have substantial value. For example, implementing cryptographic protocols based on research papers is very difficult for non-researchers. Also, the general audience is vastly larger, and currently it is primarily informed by journalists.

The lack of incentives for expository writing arises primarily within hiring and tenure committees. Tenure committees heavily weight top-tier research publications, which makes expository writing not a good use of time for pre-tenure faculty. Hiring committees also count citations weighted by venue reputation. This issue is similar to the risk of doing interdisciplinary work in that hiring and tenure candidates are frequently evaluated mostly based on their relative publication success within a given research community.

We also note that research communities themselves frequently benefit from the obfuscation of their own results. Appearing simple or easy to understand can lower the perceived value of a paper, particularly among theoreticians. Also, producing quality writing is difficult, but it is not a main consideration for accepting conference submissions, and thus producing high-quality writing is not always a maximally rewarding use of time for researchers. This seems to be a worse problem in some communities than others (e.g. security seems to value simplicity and clarity while PL and theory seems to value complexity more)

We do observe that funding bodies attempt to incentivize impacts beyond citations within a narrow field of research. For example, the NSF in the US explicitly requests "broader impact" statements, DARPA in the US often runs programs with the goal of transitioning technology to industry or government, and REF in the UK values fewer papers with larger impact.

4.4.2 Some existing solutions

There are some types of exposition that are currently working moderately well within the security and cryptography research communities, the successes and failures of which we can learn from. For example,

1. There are journals for survey papers (e.g. ACM Computing Surveys), and the surveys can receive significant numbers of citations. However, such journals are not valued highly by hiring and tenure committees.
2. Professors often produce lecture notes or books for their courses.
3. Certain security and cryptography conferences solicit "systematization of knowledge" (SoK) submissions, including IEEE Security & Privacy and the Proceedings of Privacy Enhancing Technologies.
4. There do exist widely-recognized publications that popularize security research, including USENIX ;login:, Communications of the ACM, and IEEE S&P magazine.
5. Some individual or group research blogs reach wide audiences. Blogs aren't highly valued by research committees, and may be labors of love, but they may also have important second-order benefits such as attracting collaborators, students, and funding.
6. Some researchers serve as contacts for the news media. This can eventually lead to increased funding, but this is highly variable, and talking to reporters can be time consuming.
7. Massive Online Open Courses (MOOCs) have been produced on relatively new topics, and they can contain new texts in addition to videos.
8. Books (e.g. research monographs) are quite valuable when produced, although they can be too slow for fast-moving research fields, and their limited value to the researcher's career means they frequently end up being written by tenured faculty or by people outside the core research community.

4.4.3 Future solutions

We propose several potential solutions to incentivize expository writing:

1. The most direct way to promote expository writing would be for top security and cryptography conferences to request systematization of knowledge (SoK) papers in their calls for papers. As mentioned, some such conferences already do this, and it has resulted in many valuable expository papers on important current topics, including Bitcoin, secure messaging, and website fingerprinting. The short format often required in conferences isn't ideal for exposition, however, and so journals should adopt this strategy as well. However, extra length should not be taken as an invitation for simple laundry lists of previous research, as topic surveys can easily become.
2. Journals or conferences can invite specific researchers to contribute high-quality exposition on a given topic. Conferences could combine this invitation with a keynote or tutorial invitation. It would likely be recognized as a valuable contribution because of the reputation of the journal or conference.
3. Specific "exposition retreats" or "SoK workshops" can be organized with a primary goal of producing a written exposition of a given topic. Contributors would produce different sections or chapters of a cohesive paper or book. The contributors could be invited or could propose beforehand and be selected competitively. This can be combined with summer/winter schools that are already common in cryptography by asking presenters to contribute written versions of their lectures to be combined into a set of lecture notes.

4. Graduate students might be encouraged or expected perform this function as part of their degree. Some universities and professors essentially already require this (e.g. as a qualifying exam or as part of a master's thesis). However, graduate students may lack the perspective of more experienced researchers to produce especially broad or deep exposition, and this may not suffice in fields with a high barrier to entry, such as heavily theoretical areas.

4.5 Computing on Encrypted Data, Secure Databases, Encrypted Cloud

Florian Kerschbaum (SAP SE – Karlsruhe, DE), Melissa Chase (Microsoft Corporation – Redmond, US), Jung Hee Cheon (Seoul National University, KR), Maria Dubovitskaya (IBM Research Zürich, CH), Kristin Lauter (Microsoft Research – Redmond, US), Giuseppe Persiano (University of Salerno, IT), and Benny Pinkas (Bar-Ilan University – Ramat Gan, IL)

License 🅭 Creative Commons BY 3.0 Unported license
© Florian Kerschbaum, Melissa Chase, Jung Hee Cheon, Maria Dubovitskaya, Kristin Lauter, Giuseppe Persiano, and Benny Pinkas

4.5.1 Objective and Methodology

The objective of the session was to discuss different technologies for encrypted computation. The discussion should result in a comparison of advantages and disadvantages, best-fitting use cases, and future direction of research. A single use case of outsourced, private, potentially verifiable computation of arbitrary functions in the cloud was chosen, i.e. no restriction to type of application, e.g. DRM or search. For this use case we compared the technologies of SGX, FHE, MPC and SSE/OPE. Each technology was first discussed in general terms attempting to reach a common understanding of its (security) functionality. Then we compared properties, assumptions and attacks – mostly from a security perspective, but also from economic or functionality aspects. The summary and conclusion of this discussion is listed below.

4.5.2 SGX

SGX provides a unique private, public key pair per processor. This key can be used to send encrypted data, sign the loaded code and messages. The public key must be managed in a PKI by Intel. Messages can be sent to the enclave openable only under the condition the code has been attested.

SGX provides integrity of the loaded code by mechanism comparable to remote attestation. The private key is protected by hardware. The memory of the enclave is protected by encryption, i.e. there is limited interference with other process on the cloud. The public key can also be used to tie a program to a specific processor. Management of keys for server farms can still be challenged. For communication with the client a session key needs to be established. Secure channels need to be implemented. Data stored outside of the enclave – disc or memory – needs to be encrypted.

Intel is trusted to securely generate the private key and not maintain a copy. Intel is trusted to securely manage the root CA key. All code inside the enclave is trusted. Intel and the cloud provider are trusted not to collude.

There seem to be side-channels that can leak either the private key of the processor or data, including the session key. These side channels could be timing, energy consumption or memory access patterns. Access patterns could be read via cache-timing attacks or physically from the bus. These could be combined with known attacks, on e.g. AES, for the session key.

The code inside the enclave could be vulnerable or contain backdoors. This includes code for establishing a secure channel (SSL, etc.). The code must be secure against replay attacks – potentially also from permanent storage, like disk.

There could be simulation or man-in-the-middle attacks, if the PKI fails. Generally, it is not clear how to transfer the key of the processor (or server farm) to the client. Group signatures could help.

There could be hardware attacks, e.g. sniffing the bus. While attacks by sysadmins are harder, they do have such access.

SGX seems only agreeably secure for computationally intensive tasks with little data on the client's self-written code. From an economic point SGX seems well suited. The cloud provider makes an initial investment, but can charge a higher fee.

4.5.3 FHE

There are different type of homomorphic encryption schemes. Partially (1 operation) homomorphic schemes, like Paillier, are known for a long time. Efficient fully, somewhat (2 operation) homomorphic schemes are based on the RLWE or LWE property. They support low-depth circuits and are reasonably fast. For deep circuits the error amplification requires either large parameters or bootstrapping which is complicated. Their challenge are the large key and ciphertext size.

Verification of the function is not included per se. The privacy of the computation is based only on a security assumption (RLWE/LWE). Schemes provide at least IND-CPA style security. Some information about the circuit leaks.

The only assumptions are cryptographic, such as RLWE or LWE.

Attacks could arise from side-information, such as the result or other consequences, used as a partial decryption oracle. This may break the IND-CPA model.

4.5.4 MPC

MPC allows any function to be computed on encrypted data by a set of servers. The servers may be split across organizational or legislative borders or be within a single domain (but different sysadmins). Sysadmins are a frequent target of attacks, hence MPC may help. A split may also be necessitated by legal obligations. The economic motivation for splitting the computation is challenging. A secure computation service may help under certain circumstances. MPC has fault-tolerance (availability) built-in.

MPC provides fully encrypted computation, the data is never in the clear. However, the function is usually known to all servers. A universal circuit can help avoid this leakage. MPC is general, i.e. for any circuit.

The servers share interest in carrying out the joint computation. However, the servers are assumed not to collude. There is no cryptographic assumption for secret shares. Broadcast or secure channels may be required.

There are many security models, such as semi-honest, covert, malicious. The semi-honest model only guarantees confidentiality, if integrity is preserved. This is similar to SGX where the code must be attested or confidentiality of the data cannot be guaranteed.

Obviously there are collusion attacks, but also secure channels are established by cryptographic means and can be attacked.

4.5.5 SSE/OPE

Searchable and order-preserving encryption are limited to search only. They are symmetric key crypto systems and the key is only held by the client. However, they are highly optimized and very, very, fast.

They are extremely fast, but targeted for predefined, limited functionality. SSE requires a specific implementation (search procedure) whereas OPE can be retrofitted into existing applications.

Custom security models, such IND-CKA1, IND-CKA2, IND-OCPA, IND-FAOCPA, are devised for new schemes.

Security against malicious attackers or IND-CCA2 security is likely not achieveable, although desirable. These attacks assume the worst-case and hence make minimal assumptions providing longer lasting security and the strongest security guarantee. Three different type of attacks can distinguished: Attacks based on static leakage have already been demonstrated. Attacks from dynamic information, such as queries and access patterns, are likely. Attacks based on updates are not even yet fully included in the models and algorithms. All attacks can be based on different assumptions about the adversary's knowledge or choice of plaintexts and ciphertexts.

4.6 MPC: killer applications and threat models for applications

Giuseppe Persiano (University of Salerno, IT), Christian Grothoff (INRIA – Rennes, FR), Aaron Michael Johnson (NRL – Washington, US), Yehuda Lindell (Bar-Ilan University – Ramat Gan, IL), Claudio Orlandi (Aarhus University, DK), and Benny Pinkas (Bar-Ilan University – Ramat Gan, IL)

License 🅭 Creative Commons BY 3.0 Unported license
© Giuseppe Persiano, Christian Grothoff, Aaron Michael Johnson, Yehuda Lindell, Claudio Orlandi, and Benny Pinkas

4.6.1 Killer Applications

We have identified the following as applications domain where the need for MPC is needed.

- Statistics over distributed systems.
  Several organizations have a large user base and would like to collect statistics of the users. The type of statistics that one is interested in has a big potential impact on the efficiency of the protocol. Arithmetic statistics (like average, standard deviation) are easier to compute than other more robust statistics (like median).
- Sharing sensitive data.
  One way to protect keys is to share them in a secure way and to distribute the shares to parties (possibly running different OSes). In this way if one of the parties is compromised then the key is still safe. Whenever the key is needed to access encrypted data (or to perform entity/data authentication,...) the parties holding the shares will perform the action required (decryption, authentication,...) by means of MPC.
- Privacy preserving.
  MPC can also be used to protect privacy of users that are really concerned about their private data and would resort to MPC whenever their personal data was needed.
- Auction.

The domain of electronic marketplaces seems to be another area that might benefit from the use of MPC. The first example is action that could be very efficiently implemented in a secure way and would provide an added level of privacy to the users that desire so.

4.6.2 Model threats for applications

Research effort in MPC has focused primarily in obtaining results that would guarantee the best possible security along with several other desirable properties like correctness, fairness, termination,... This approach has been very successful and has led to very general results. We observed though that in several applications not all properties are needed and one could then obtain more efficient/practical construction. The primary example that came up is about termination. If one is running MPC within an organization for the purpose of protecting keys, it would be actually desirable to learn that one of the parties is compromised by observing some execution being aborted.

Dishonest majority arises naturally in the context of secure two-party computation where the problem with honest majority is trivial.

An interesting model that seems to have been adopted by most (if not all) current industrial implementations of MPC relies on a restricted number (as little as two) servers that perform MPC over the shares of data provided by the users. Having a restricted number of players in an MPC has the obvious advantage of increasing efficiency and, in addition, it is much easier to assess the trustworthiness of few parties. This model seems particularly fit for the problem of computing statistics over large distributed system that have already identified a restricted number of nodes for other administrative tasks.

4.7 Anonymous Payment Systems

Nigel P. Smart (University of Bristol, GB), Alex Biryukov (University of Luxembourg, LU), Allison Bishop (Columbia University – New York, US), Bogdan Carbunar (Florida International University – Miami, US), Melissa Chase (Microsoft Corporation – Redmond, US), George Danezis (University College London, GB), Maria Dubovitskaya (IBM Research Zürich, CH), Christian Grothoff (INRIA – Rennes, FR), and Martina Angela Sasse (University College London, GB)

License 🅭 Creative Commons BY 3.0 Unported license
© Nigel P. Smart, Alex Biryukov, Allison Bishop, Bogdan Carbunar, Melissa Chase, George Danezis, Maria Dubovitskaya, Christian Grothoff, and Martina Angela Sasse

The group discussed a number of topics related to payments in general. These ranged from payment systems which are bitcoin like, through to systems based on store loyalty points, credit card systems, smart metering and bartering. An issue with all systems was to define the nature of anonymity, and to whom anonymity is maintained. For example early systems proposed for online banking transactions like SET tried to maintain a cryptographic separation of data based on a "need to know principle". This never took off, and has since been replaced by "best practice" requirements as in the PCI standards. This means that merchants are exposed to identifying information about customers, even when they do not need to be (for example in the purchase of digital goods).

We discussed issues of exchange between various reserves of value; for example altcoin exchanges are already in existence. Members of the group discussed the benefit of usage of exchange between non-currency based reserves of value. For example the trading of frequent flyer miles with, say, store card points.

An example discussed in detail was the anonymity requirements in smart metering, which is essentially a payment system for electricity. The different stakeholders in the system were discussed and how their requirements for visibility of the transactions conflict with each other parties utility. An interesting aspect of the smart metering case study is that anonymity is not required for the identities, but is required for the amounts. This is the opposite of the case in bitcoin. Many of the aspects of the smart metering example also apply to digital goods such as Spotify.

4.8 Cryptographic Hardness Assumptions

Nigel P. Smart (University of Bristol, GB), Yehuda Lindell (Bar-Ilan University – Ramat Gan, IL), and Kenneth G. Paterson (Royal Holloway University of London, GB)

License 🅭 Creative Commons BY 3.0 Unported license
© Nigel P. Smart, Yehuda Lindell, and Kenneth G. Paterson

There are two problems we face at the moment

- People are not working on breaking of hard problems.
- People are not trying to build cryptosystems under minimal assumptions.

As an example of the first problem, few people are seriously working on factoring, discrete logarithms (bar characteristic two fields), or bilinear assumptions. This is a problem of incentives; for example a mathematician who might be interested in working on discrete logarithms would take a long time to come up with any result (if any result is possible), and would end up publishing outside their field. Thus their publication record would be damaged by engaging in working on hard problems. It seems a difficult to counteract this problem of incentives, as it goes to the heart of what is needed to become a successful academic.

The second problem is typified by what we see in the area of iO currently. Researchers seem to be incentivized in coming up with applications of iO and are less incentivized in understanding the underlying problems. This is interesting when compared to the similar "bandwagon" created by FHE: with that bandwagon, researchers worked both on simplifying the underlying constructions and improving the hardness assumptions (e.g. the creation of the LWE and NTRU based schemes compared to the original Gentry scheme), as well as looking at potential applications in areas such as verifiable computation.

There are plenty of problems in the space which are relevant in the real world but are not being addressed. For example, there are plenty of constructions (efficient ZKPoKs) for which we have no analogue in the post-quantum world. Are we likely to end up with a plethora of assumptions in the world of PQC as soon as there is an urgent need for PQC methods beyond encryption and signatures? This need is likely to become a pressing need with the next few years as companies start to look at deploying post-quantum systems.

There seems to be an incentive in research into creating new applications and functionalities within cryptography, as opposed to looking at old problems and finding solutions under better assumptions, or looking at making old applications better (where the metric is "better" is either security, implementation techniques etc). If we look at the top conferences, papers which create new applications (e.g. iO, FHE, etc) seem to get more traction than papers which implement things better. This is despite the CFP for CRYPTO stating implementation or industrially relevant research being welcome for a number of years.

One solution suggested by Angela, would be to have a conference in which half worked out exploratory ideas could be batted around and discussed. Each paper is presented and then discussed by the audience. With post-proceedings which contain the feedback from the audience. This model has apparently worked well in the security community. The question is when did this start becoming a problem? When did "irrational exuberance"[1] take over our field? Perhaps the plethora of pairing based assumptions in the early 2000's led to our current state of affairs. The authors feel that researchers need to get back to cryptographic basics which have high impact:

- Encourage work looking at schemes based on more standard assumptions.
- Encourage work which tries to improve the efficiency of schemes and their practicality.

[1] A phrase borrowed from Phil Rogaway.

Participants

- Raad Bahmani
TU Darmstadt, DE
- Daniel J. Bernstein
University of Illinois – Chicago, US
- Konstantin Beznosov
University of British Columbia – Vancouver, CA
- Alex Biryukov
University of Luxembourg, LU
- Allison Bishop
Columbia University – New York, US
- Alexandra Boldyreva
Georgia Institute of Technology – Atlanta, US
- Nikita Borisov
University of Illinois – Urbana Champaign, US
- Ferdinand Brasser
TU Darmstadt, DE
- Christian Cachin
IBM Research Zürich, CH
- Bogdan Carbunar
Florida International University – Miami, US
- Melissa Chase
Microsoft Corporation – Redmond, US
- Jung Hee Cheon
Seoul National University, KR
- Marc C. Dacier
QCRI – Doha, QA
- George Danezis
University College London, GB
- Yevgeniy Dodis
New York University, US
- Maria Dubovitskaya
IBM Research Zürich, CH
- Dieter Gollmann
TU Hamburg-Harburg, DE
- Christian Grothoff
INRIA – Rennes, FR
- Krista Grothoff
GNUNet e. V. – Rennes, FR
- Nadia Heninger
University of Pennsylvania – Philadelphia, US
- Aaron Michael Johnson
NRL – Washington, US
- Stefan Katzenbeisser
TU Darmstadt, DE
- Florian Kerschbaum
SAP SE – Karlsruhe, DE
- Yongdae Kim
KAIST – Daejeon, KR
- Tanja Lange
TU Eindhoven, NL
- Kristin Lauter
Microsoft Research – Redmond, US
- Yehuda Lindell
Bar-Ilan University – Ramat Gan, IL
- Sarah Meiklejohn
University College London, GB
- Refik Molva
EURECOM – Sophia Antipolis, FR
- Moni Naor
Weizmann Institute – Rehovot, IL
- Claudio Orlandi
Aarhus University, DK
- Kenneth G. Paterson
Royal Holloway University of London, GB
- Adrian Perrig
ETH Zürich, CH
- Giuseppe Persiano
University of Salerno, IT
- Andreas Peter
University of Twente, NL
- Benny Pinkas
Bar-Ilan University – Ramat Gan, IL
- Martina Angela Sasse
University College London, GB
- Vitaly Shmatikov
Cornell Tech NYC, US
- Radu Sion
National Security Institute – Stony Brook, US
- Nigel P. Smart
University of Bristol, GB
- Gene Tsudik
University of California – Irvine, US
- Avishai Wool
Tel Aviv University, IL