Research Statement

Benjamin Fuller
benjamin.fuller@uconn.edu

January 9, 2017

I develop provably secure cryptographic protocols that address real-world problems using a deep understanding of practical considerations. These protocols will improve the security of today's systems by leveraging strong cryptography. My approach is informed by my experience transitioning cryptography to use and solving users' problems. This experience helps me create efficient and usable protocols. I work with researchers in cryptography and system security to design usable protocols that drawing on strong cryptography to improve the security of today's systems.

1 Background

I am currently an Assistant Professor of Computer Science and Engineering at the University of Connecticut. From 2009 to 2014, I designed cryptographic protocols while obtaining my Ph.D. at Boston University.

From 2007 to 2016, I researched computer security at MIT Lincoln Laboratory, a federally funded research and development center. From 2015-2016, I served as a principal investigator at MIT Lincoln Laboratory leading research and software development teams, managing between 5-10 staff and 3 research companies. My responsibilities included project development and management, developing new cryptographic techniques, and leading software development.

MIT Lincoln Laboratory focuses on applying research to government problems. While MIT Lincoln Laboratory publishes in the academic literature, this is a secondary objective. While at UConn and Boston University, research was a primary objective, my publications during those periods are in stronger venues.

2 Prior Research Overview

Cryptographic constructions span a range of security and efficiency guarantees. On one hand lie constructions providing strong security against a variety of threats whose inefficiencies prevent wide scale deployment. On the other hand systems used in practice are efficient but often provide heuristic security claims. I believe this sharp tradeoff is unnecessary.

My research has focused on three areas: 1) how to securely communicate with a group of people 2) how to search over encrypted data 3) and how to provide crytographic authentication from high entropy sources which often suffer from noise between repeated readings. Below I provide a summary of my work in these areas [1–17].

2.1 Secure Group Communication

Public-key encryption schemes allow parties to securely communicate in one-to-one fashion. Group communication can be achieved by individually encrypting to each desired recipient. However, such an approach is prohibitively expensive, especially when the group changes over time. Fiat and Naor [18] formally defined the problem, Naor, Naor, and Lotspiech [19] use a tree based approach to efficiently handle membership changes.

Our team developed and implemented a variant of this scheme [5]. When membership changes, our protocol first changes the key to an intermediate group that is likely to be useful in the future. This creates additional groups with provisioned keys, keeping the average cost of membership change stable.

We demonstrated our implementation using a secure chat application. Our chat system served as the primary communication tool between multiple airplanes flying in and out of range [1–3]. I learned that debugging cryptography on an airplane is higher pressure than a simulation in an office. It made me understand users' focus on executing their task and that cryptography must be unobtrusive and make choices without involving the user.

2.2 Searching over Encrypted Data

The previous section discussed securing data in transit. Received data must then be securely stored. For most users, disk encryption can adequately protect their local data. However, users' data no longer resides on their local machine, it has migrated to remote services. Traditional encryption destroys the search functionality integral to many services. It is possible to balance security and functionality by using encryption that allows for limited functionality.

Deterministic public key encryption [20] allows for a service provided to check equality which can be used to implement keyword search. We constructed a deterministic public-key encryption scheme from any trapdoor function, unifying previous constructions [6, 7].

In 2015, our team deployed searchable symmetric encryption in a government application. Searchable symmetric encryption provides enables searching while protecting the contents at rest. I developed and managed the project, led research and software teams, oversaw subcontracts to three technology providers, and demonstrated and evaluated multiple technologies with users. This role developed new skills: project development and management, financial planning, and supervising researchers. Our team wrote a systemization of knowledge paper on the state of searchable encryption which standardizes terminology, compares existing solutions, and maps a research agenda for the community [17].

My research role included tailoring cryptographic protocols. Some system challenges arise when systems are instantiated. Careful and regular communication can unite research and development. As an example, auditing of secure search was a surprising requirement. Our research team designed a secure audit log system that maintains the security of the underlying cryptographic protocols [15].

2.3 Authentication from Noisy Sources

Communicating parties need to authenticate one another to establish a secure communication channel. Two commonly used methods are passwords and the public key infrastructure (PKI). For example, when a user logs into his or her bank account over TLS, the user's computer authenticates the bank using the PKI, while the bank authenticates the user with a password.

The drawbacks of both methods are numerous. Passwords are known to be insecure: users cannot memorize ones that are strong enough [21, 22] and users often disclose their passwords. PKI, on the other hand, requires centralized trusted authorities. Such authorities are hard to establish and maintain outside of hierarchical organizations. It is attractive to establish a secure communication channel for settings where such infrastructure is not available.

There are many proposals to base authentication on an information source that is known only to communicating parties. Often, such sources of information have higher entropy than passwords and are easier to store. Unfortunately, many of them are noisy and provide similar, but not identical secret values at each reading. Examples of such sources include biometrics [23], keystroke dynamics [24], and hardware devices [25, 26].

Dodis, Ostrovsky, Reyzin, and Smith [27] designed fuzzy extractors to derive keys from noisy sources. Previous constructions of fuzzy extractors have no proofs of security for many sources of practical importance. As an example, the human iris is thought to be strongest biometric [23]. However, in our work, we argue current fuzzy extractors should not be used on irises in high security applications [13].

To improve fuzzy extractors for practical sources, we started by asking which limitations were necessary. Our work described necessary and sufficient conditions for building a fuzzy extractor [14]. We then formulated an alternative definition of fuzzy extractors [8] providing security against adversaries with bounded resources (running in polynomial time). Computational security is often used in cryptography and fuzzy extractors have no compelling need for security against unbounded attackers.

We then constructed novel fuzzy extractors using this definition. Our schemes provide two main new properties: 1) longer keys [8] and 2) reusability across multiple providers [10, 11]. The only known reusable fuzzy extractor assumed that different providers received readings that were correlated in unrealistic ways [28]. Our construction only needs each provider to receive a valid reading. In addition, we proved our construction is secure for an wide class of noisy

sources. I believe our new construction can be used to secure sources of practical importance including the iris.

3 Current and Future Work

My research has dealt with securing message transmission, searching over encrypted data sets, and authenticating users. I am continuing the latter two research areas: searching over encrypted data and designing authentication paradigms. In addition, I am excited to pursue new research topics that will improve today's vulnerable computing environment.

Symmetric Searchable Encryption Symmetric searchable encryption is nearing practical levels of performance. However, the security definitions are not well understood. Most definitions allow the adversary to learn some statistics about the underlying data. Researchers have used learning techniques to recover the underlying data from these statistics [29, 30]. I have worked on information leakage [4, 12] and would like to collaborate with experts in learning theory to understand the impact of leakage on searchable encryption. Our systemization paper includes a review of this topic [17].

Additionally, I will extend search functionality. In many online applications a search returns not only the exact query but "nearby" results. For example, a search for "John" would also return "Jon" as the edit distance between the two words is one. Previous work has added this functionality to searchable encryption schemes [31]. However, this work considered general notions of closeness, resulting in inefficient schemes. I would like to consider common closeness metrics and make the constructions more efficient. This work will combine ideas from the previous two subsections.

Improving Authentication I would like to increase the supported error rate in our work [10] and demonstrate security for real sources including the human iris. I have begun examining iris key derivation with researchers from Boston University and MIT.

Hardware tokens present another promising authentication alternative. Physical unclonable functions (PUFs) are believed to provide strong authentication [25]. I have worked to integrate cryptography and processing of PUFS [9] and implemented a new fuzzy extractor constructor [16]. Our team is implementing other recent fuzzy extractor constructions.

Security as a Property Cryptography is a valuable tool. However, most cryptography does not address the problems in today's systems. Novel constructions in theoretical cryptography are disconnected from the systems used in practice. Several factors contribute to this disconnect: inefficient schemes, misunderstanding of threats, and ignorance of human factors including incentives and usability. Cryptography's maximum impact will come when built with true understanding of the whole system.

I want to build a world where security is a basic design principle, not a feature to be added. I see an internet where users execute their goals with minimal risk to their property, information and lives. This goal will not be achieved in a few years. However, researchers must put forth the effort and work with system designers to understand their problems and needs. I will develop provably secure systems building on my experience with systems in practice. My breadth of experience enables me to design systems that solve today's critical problems.

References

[1] Roger Khazan, Joseph Cooley, Galen Pickard, and Benjamin Fuller. GROK secure multi-user chat at Red Flag 2007-03. In *Military Communications Conference*, pages 1–7, 2008.

[2] Joseph Cooley, Roger Khazan, Benjamin Fuller, and Galen Pickard. GROK: A practical system for securing group communications. In *Network Computing and Applications (NCA)*, pages 100–107. IEEE, 2010.

[3] Benjamin Fuller, Roger Khazan, Joseph Cooley, Galen E Pickard, and Daniil Utin. ASE: Authenticated Statement Exchange. In *Network Computing and Applications (NCA)*, pages 155–161. IEEE, 2010.

[4] Benjamin Fuller and Leonid Reyzin. Computational entropy and information leakage. Technical report, Boston University, 2011.

[5] Galen E Pickard, Roger I Khazan, Benjamin Fuller, and Joseph A Cooley. DSKE: Dynamic Set Key Encryption. In *Local Computer Networks Workshops (LCN Workshops)*, pages 1006–1013. IEEE, 2012.

[6] Benjamin Fuller, Adam ONeill, and Leonid Reyzin. A unified approach to deterministic encryption: New constructions and a connection to computational entropy. In *Theory of Cryptography*, pages 582–599. Springer, 2012.

[7] Benjamin Fuller, Adam O'Neill, and Leonid Reyzin. A unified approach to deterministic encryption: New constructions and a connection to computational entropy. *Journal of Cryptology*, pages 1–47, 2012.

[8] Benjamin Fuller, Xianrui Meng, and Leonid Reyzin. Computational fuzzy extractors. In *Advances in Cryptology-ASIACRYPT 2013*, pages 174–193. Springer, 2013.

[9] Merrielle Spain, Benjamin Fuller, Kyle Ingols, and Robert Cunningham. Robust keys from physical unclonable functions. In *IEEE International Symposium on Hardware-Oriented Security and Trust (HOST)*, pages 88–92. IEEE, 2014.

[10] Ran Canetti, Benjamin Fuller, Omer Paneth, Leonid Reyzin, and Adam Smith. Reusable fuzzy extractors for low-entropy distributions. In *Advances in Cryptology–EUROCRYPT 2016*, pages 117–146. Springer, 2016.

[11] Benjamin Fuller. *Strong Key Derivation from Noisy Sources*. PhD thesis, Boston University, 2015.

[12] Benjamin Fuller and Ariel Hamlin. Unifying leakage classes: Simulatable leakage and pseudoentropy. In *Information Theoretic Security*, pages 69–86. Springer, 2015.

[13] Gene Itkis, Venkat Chandar, Benjamin Fuller, Joseph Campbell, and Robert Cunningham. Iris biometric security challenges and possible solutions: For your eyes only? using the iris as a key. *Signal Processing Magazine, IEEE*, 32(5):42–53, 2015.

[14] Benjamin Fuller, Leonid Reyzin, and Adam Smith. When are fuzzy extractors possible? In *Advances in Cryptology–ASIACRYPT 2016*, pages 277–306. Springer, 2016.

[15] Robert Cunningham, Benjamin Fuller, and Sophia Yakoubov. Catching MPC cheaters: Identification and openability. In *Submission*, 2016.

[16] Chenglu Jin, Phuong Ha Nguyen, Ling Ren, Charles Herder, Benjamin Fuller, Marten van Dijk, and Srini Devadas. Practical cryptographically-secure PUFs based on learning parity with noise. In *Submission*, 2016.

[17] Benjamin Fuller, Mayank Varia, Arkady Yerukhimovich, Emily Shen, Ariel Hamlin, Vijay Gadepally, Richard Shay, John Darby Mitchell, and Robert Cunningham. SoK: Cryptographically protected database search. In *Submission*, 2016.

[18] Amos Fiat and Moni Naor. Broadcast encryption. In *Advances in Cryptology-CRYPTO93*, pages 480–491. Springer, 1994.

[19] Dalit Naor, Moni Naor, and Jeff Lotspiech. Revocation and tracing schemes for stateless receivers. In *Advances in Cryptology-CRYPTO 2001*, pages 41–62. Springer, 2001.

[20] Mihir Bellare, Alexandra Boldyreva, and Adam O'Neill. Deterministic and efficiently searchable encryption. In *Advances in Cryptology–CRYPTO 2007*, pages 535–552, 2007.

[21] Jeff Jianxin Yan, Alan F Blackwell, Ross J Anderson, and Alasdair Grant. Password memorability and security: Empirical results. *IEEE Symposium on Security and Privacy*, 2(5):25–31, 2004.

[22] Matt Weir, Sudhir Aggarwal, Michael Collins, and Henry Stern. Testing metrics for password creation policies by attacking large sets of revealed passwords. In *ACM Conference on Computer and Communications Security*, pages 162–175. ACM, 2010.

[23] John Daugman. How iris recognition works. *Circuits and Systems for Video Technology, IEEE Transactions on*, 14(1):21 – 30, January 2004.

[24] Fabian Monrose, Michael K Reiter, and Susanne Wetzel. Password hardening based on keystroke dynamics. *International Journal of Information Security*, 1(2):69–83, 2002.

[25] Ravikanth Pappu, Ben Recht, Jason Taylor, and Neil Gershenfeld. Physical one-way functions. *Science*, 297(5589):2026–2030, 2002.

[26] Pim Tuyls, Geert-Jan Schrijen, Boris Skoric, Jan Geloven, Nynke Verhaegh, and Rob Wolters. Read-proof hardware from protective coatings. In *Cryptographic Hardware and Embedded Systems - CHES 2006*, pages 369–383. 2006.

[27] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM Journal on Computing*, 38(1):97–139, 2008.

[28] Xavier Boyen. Reusable cryptographic fuzzy extractors. In *Proceedings of the 11th ACM conference on Computer and Communications Security*, CCS '04, pages 82–91, New York, NY, USA, 2004. ACM.

[29] Muhammad Naveed, Seny Kamara, and Charles V Wright. Inference attacks on property-preserving encrypted databases. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 644–655. ACM, 2015.

[30] David Cash, Paul Grubbs, Jason Perry, and Thomas Ristenpart. Leakage-abuse attacks against searchable encryption. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 668–679. ACM, 2015.

[31] Alexandra Boldyreva and Nathan Chenette. Efficient fuzzy search on encrypted data. In *Fast Software Encryption*, pages 613–633. Springer, 2014.