

Cryptography encompasses many intriguing topics beyond the better-known, currently deployed encryption and digital signatures. It also includes subjects such as electronic voting, on-line auctions, private database queries, and secure multi-party computation. The theoretical cryptography research I undertake is based on precise definition of the desired security properties. Furthermore, coherent, logical arguments in the form of proofs are employed to guarantee that solutions satisfy these properties. In contrast, information security focuses on the design and implementation of secure systems in practice.

Security, the inability of an attacker to cause harm or act in an unprescribed manner, is an inherently negative property. Hence, whenever possible, a security proof is preferable to loose claims or unsupported heuristics. A project I worked on with a researcher at a large financial services firm illustrates this. He had developed a digital signature protocol based on one-time password tokens and wanted our group's help to prove it secure. Despite his many years of practical experience and information security knowledge, our theoretical analysis identified several necessary modifications.

My first publication was in the area of steganography: the study of undetectable communication over a publicly observable channel. It resulted from a survey of papers in this area undertaken with my advisor's help. One way to get undergraduates involved in research on intriguing topics like this is to help them read and present existing results. Our research considered the case where communicating parties know very little about the channel, specifically only a bound on the channel entropy, and have the ability to sample from it as desired, while adversarial observers trying to detect the dialog do so passively. Surprisingly, steganography is possible over almost any channel: the higher the entropy the higher the possible rate. We presented two efficient, provably-secure encoding schemes: one stateful and one stateless. Furthermore, my coauthors showed in the same work that the efficiency of our schemes is, in fact, optimal for this setting. This experience helped teach me the importance of persistence in research, the usefulness of literature searches, and how to critically read research results.

My second publication explored enhancements to the traditional secure channel primitive used for confidential communication. The security of cryptographic primitives usually hinges upon the secrecy of keys or other information. Our intrusion-resilient secure channel primitive minimizes the consequences of compromised secret keys and, in some cases, even facilitates the complete restoration of confidentiality. We precisely defined this primitive and described a provably-secure generic construction that uses only standard public-key encryption schemes. We also proved results about the benefits of composing intrusion-channels with arbitrary two-party protocols.

Another yet to be published work that is a significant part of my dissertation describes protocols for binary search and simple range queries with mutual privacy: the user learns only the query result and nothing else about the database, while the honest-but-curious database server learns nothing about the query (or its result). We developed query protocols for honest and malicious users. The work emerged from discussions with database colleagues about private nearest-neighbor queries for spatial datasets. Solving this problem, improving and extending our range query result, and investigating other private queries will be the focus of my ongoing research in the immediate future.

I believe that the most effective security solutions require a combination of sound technology together with practical policies and legislation grounded in realistic expectations. I hope to work toward such solutions with the help of other computer scientists and like-minded researchers from complementary disciplines such as government and business. I also strive to educate students, my colleagues, and the public about what security can and cannot be achieved by technology alone.

Involving undergraduates in cryptography and information security research will help them better appreciate the importance and limits of these areas. Students are often enthusiastic, early adopters of new technology who push its limits by using it in unexpected ways. Because of this, I think they are well-suited to this kind of research and will naturally be drawn to it. Increased media coverage of digital security and privacy issues such as identity theft and lost personal information will also help attract them to this research. Encouraging them to explore for themselves the technological and societal impacts of security and privacy issues will spur them to learn even more about these areas, computer science in general, and public policy. I will also encourage them to follow and take an active part in the public policy debates related to privacy and security. I welcome their participation in my private database query research and eagerly anticipate being involved in numerous student-initiated projects in the near future.