

Research Statement

Chris Peikert

December 22, 2008

1 New Foundations for Cryptography

Most cryptographic tasks must inherently rely on *assumptions* about the difficulty of some computational problem. Over the past three decades, *number theory* has served as the primary source of seemingly hard problems for cryptography; for instance, a prototypical conjecture is that it is infeasible to factor the product of two large, random prime numbers. Many such number-theoretic problems have a common underlying structure, so the resulting cryptographic schemes frequently have similar characteristics and limitations. Moreover, this shared structure means that unforeseen developments could render many schemes less secure or useful than had been believed.

A major theme of my research is to develop *new mathematical foundations for cryptography*, with a special focus on objects called *lattices*. (A lattice is essentially a periodic “grid” of points in \mathbb{R}^n .) Compared to conventional number theory, lattices offer a host of intriguing properties and potential advantages:

- Lattice-based schemes **can be quite efficient, especially when exploiting parallelism**. Their core operations usually involve just *adding small integers*, whereas number-theoretic schemes typically require *exponentiating very large integers*.
- Lattice problems have so far **resisted attacks by subexponential-time and quantum algorithms**. In contrast, most number-theoretic problems used in cryptography can be solved in much better than exponential time [LL93], and *all* of them can be solved efficiently by quantum algorithms [Sho97].
- It is possible to design schemes that are **provably secure under “worst-case” assumptions** about the difficulty of well-studied lattice problems. The schemes are therefore “*as hard as possible*” to break, in a precise sense. This kind of strong guarantee is entirely unique in cryptography, where “average-case” assumptions about the difficulty of *random* instances are the norm.

Lattices have a very different mathematical structure than other objects that are more commonly used in cryptography, and this has been a double-edged sword: while they contribute some much-needed diversity, they also require entirely different perspectives and techniques to be of use. As a result, for many years the literature on lattice-based schemes was technically complex and limited to just a few basic primitives.

My research over the past few years has *significantly advanced the state of the art* in lattice-based cryptography. My contributions have included new conceptual perspectives and simplifying abstractions, powerful algorithmic and analytic tools, and novel cryptographic schemes having rich features and strong security properties. As a bonus, some of these perspectives and techniques have also led to significant progress in conventional *number-theoretic* cryptography. In the following, I will describe a few highlights of these research efforts and how I plan to continue these lines of inquiry.

Encryption from the Shortest Vector Problem. The (approximate) *shortest vector problem* GapSVP is perhaps the most fundamental problem related to lattices; essentially, it asks whether the shortest nonzero vector in a given lattice is “long” or “short,” where the precise quantities are separated by some approximation factor. Since the seminal work of Ajtai [Ajt04] in 1996, a handful of basic cryptographic primitives, such as one-way functions, have been based upon the conjectured worst-case hardness of GapSVP (for a gap that is a small polynomial in the dimension of the lattice). However, richer notions like *public-key encryption* have so far required stronger assumptions about either a *special case* of the shortest vector problem [AD97, Reg04] or the difficulty of lattice problems for *quantum algorithms* [Reg05]. Resolving this disparity has been one of the central open problems in the area.

In a recent work [Pei08b], I solve this problem by constructing public-key encryption schemes whose security is based on the *standard* GapSVP (with small polynomial gap), thereby placing encryption on a foundation comparable to that of other primitives. My solution involves a new design approach and proof technique, which comes with two additional advantages:

1. It yields *very simple* cryptosystems that are secure not only under the basic concept of a *passive* attack, but even under the much stronger “gold standard” notion of an active *chosen-ciphertext attack*.
2. It provides a stronger foundation for several other recent lattice-based schemes (e.g., [Reg05, GPV08, PW08, PVW08]), in the sense that they are now known to be secure under *classical* assumptions, whereas conjectures about the limits of *quantum* algorithms were needed previously.

How to Use a Short Basis. From the very early years of lattice-based cryptography, there were proposals for cryptographic schemes (e.g., GGH encryption and signatures [GGH97]) in which the secret key is a “short” or “high-quality” basis for a lattice, and the public key is a very “low-quality” basis for the same lattice. However, none of these schemes came with security proofs, and recently it was discovered that some of them are *completely insecure*, because they implicitly expose the secret key over time [NR06].

My recent work from STOC '08 (with C. Gentry and V. Vaikuntanathan) [GPV08] provides a new approach for *securely using a short basis* in cryptography and other contexts. The core innovation is a general technique for “obliviously sampling” lattice points in a way that *provably leaks no information* about the secret basis. Upon this foundation we design several desirable cryptographic schemes, including simple “hash-and-sign” *digital signatures* and even *identity-based encryption* (which has been notoriously difficult to obtain under *any* reasonable assumption). My subsequent work has built further upon these ideas to construct other important cryptographic protocols such as *noninteractive zero-knowledge proofs* [PV08] and *oblivious transfer* [PVW08], and has shown how to construct optimally short bases for hard lattices [AP09].

Lossy Trapdoor Functions. *Trapdoor functions* (TDFs) were one of the first abstract notions described in the modern cryptographic literature [DH76], and have been a central linchpin in many applications since. Yet despite over three decades of research, every proposed TDF from broadly accepted assumptions (e.g., the RSA function [RSA78]) has relied on the conjectured difficulty of *factoring* for its security.

In work from STOC '08 (with B. Waters) [PW08], we introduce a powerful new abstract concept called *lossy trapdoor functions*. These are special kinds of TDFs that come with a complementary “information-losing” mode that greatly aids the design and security analysis of higher-level applications. We show that lossy TDFs imply a host of fundamental cryptographic notions, including *oblivious transfer*, *collision-resistant hash functions*, and public-key cryptosystems that are secure under *chosen-ciphertext attack*. Moreover, we show how to construct lossy TDFs under any of several common cryptographic assumptions. Two notable outcomes of this research are the first chosen ciphertext-secure cryptosystem based on a worst-case assumption, and the first known TDFs that are unrelated to the factoring problem.

Future Directions. Lattices (and more generally, objects in continuous spaces) are a plentiful source of intriguing problems in mathematics and computer science. A few general questions I plan to investigate in the near future include: How can additional algebraic structure be exploited to design practical schemes, and how does it affect security? (See [PR06, PR07, ADL⁺08] for some initial forays.) How are the many different lattice problems related across different dimensions and norms? (See [RR06, Pei08a] for recent results.) Which other important cryptographic notions (or useful relaxations thereof) have natural lattice-based realizations? Might certain lattice problems be at least as hard as conventional number-theoretic problems like factoring? More broadly, what essential theoretical notions in discrete settings (e.g., hardness amplification, randomness extraction, learnability) have natural and useful analogs in continuous contexts?

2 Applying Cryptography Across Disciplines

Broadly construed, cryptography is about designing systems that resist malicious behavior. Another major theme of my research is *to apply the cryptographic methodology to other contexts where malicious or otherwise faulty behavior can arise*. Two such areas where I believe great contributions are possible include *coding theory* and *game theory*.

Error Correction. A central goal of *coding theory* is to reliably transmit data over a noisy channel. The main objects of study for this purpose are *error-correcting codes*, which are typically measured according to certain combinatorial properties (e.g., distance, rate). The noisy channel is also typically combinatorial, limited only in the number of errors it may introduce. While this model has yielded many elegant constructions, it also seems to be quite constraining. For example, it is *impossible* to unambiguously correct errors in even one-quarter of the transmitted bits, if the errors are chosen maliciously.

In work with S. Micali, M. Sudan, and D. A. Wilson [MPSW05], we advocate viewing the channel as an arbitrary (possibly malicious) entity that is limited to *feasible computation*, and design a very effective error-correcting scheme in that model. Our solution blends, in an elegant and modular way, the notion of *efficient list decoding* from coding theory with the concept of *authentication* from cryptography. This yields a dramatic improvement in the ability to correct errors — for example, our binary codes can withstand error rates up to *one-half*, which vastly exceeds the classical limit of one-quarter (and is in fact optimal for *any* reasonable model of noise). Moreover, our technique results in *efficient and explicit* codes whose information rates match the Shannon limit for the channel. Prior to our work, such codes were known only for the much easier case of random (non-adversarial) noise, and even then the codes were non-constructive.

Game Theory. Game theory is also concerned with players who have antagonistic objectives. Players in games are assumed to act *rationally* so as to increase their utilities; while this behavior may not rise to the level of outright malice, it may involve a wide class of devious actions that cannot be ruled out *a priori*. In addition, players are usually assumed to have *unbounded* rationality, which amounts to unlimited computational resources.

In work with M. Lepinski, S. Micali, and A. Shelat [LMPS04], we reconsider game theory in a model where players are limited to *efficient* computation. In this setting, we view games as cryptographic protocols involving potentially adversarial players, and define a new equilibrium concept that captures the notion that no *feasible* deviation from equilibrium should result in higher utility, even when *colluding* with other deviating players. Using the cryptographic notions of *secure function evaluation* and *fairness*, we then demonstrate how to design player strategies that satisfy our new notion of equilibrium. These strategies yield large and robust utilities that are *impossible* to achieve (in general) via traditional solution concepts like Nash equilibrium.

References

- [AD97] Miklós Ajtai and Cynthia Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *STOC*, pages 284–293, 1997.
- [ADL⁺08] Yuriy Arbitman, Gil Dogon, Vadim Lyubashevsky, Daniele Micciancio, Chris Peikert, and Alon Rosen. SWIFFTX: A proposal for the SHA-3 standard. Submitted to NIST SHA-3 competition, 2008.
- [Ajt04] Miklós Ajtai. Generating hard instances of lattice problems. *Quaderni di Matematica*, 13:1–32, 2004. Preliminary version in *STOC* 1996.
- [AP09] Joël Alwen and Chris Peikert. Generating shorter bases for hard random lattices. In *STACS*, 2009. Accepted.
- [DH76] Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, IT-22(6):644–654, 1976.
- [GGH97] Oded Goldreich, Shafi Goldwasser, and Shai Halevi. Public-key cryptosystems from lattice reduction problems. In *CRYPTO*, pages 112–131, 1997.
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008.
- [LL93] A. K. Lenstra and H. W. Lenstra, editors. *The development of the number field sieve*. Springer-Verlag, August 1993.
- [LMPS04] Matt Lepinski, Silvio Micali, Chris Peikert, and Abhi Shelat. Completely fair SFE and coalition-safe cheap talk. In *PODC*, pages 1–10, 2004.
- [MPSW05] Silvio Micali, Chris Peikert, Madhu Sudan, and David A. Wilson. Optimal error correction against computationally bounded noise. In *TCC*, pages 1–16, 2005.
- [NR06] Phong Q. Nguyen and Oded Regev. Learning a parallelepiped: Cryptanalysis of GGH and NTRU signatures. In *EUROCRYPT*, pages 271–288, 2006.
- [Pei08a] Chris Peikert. Limits on the hardness of lattice problems in ℓ_p norms. *Computational Complexity*, 17(2):300–351, May 2008. By invitation to special issue on CCC '07 (Conference on Computational Complexity).
- [Pei08b] Chris Peikert. Public key cryptosystems from the worst-case shortest vector problem. Submitted, 2008.
- [PR06] Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, pages 145–166, 2006.
- [PR07] Chris Peikert and Alon Rosen. Lattices that admit logarithmic worst-case to average-case connection factors. In *STOC*, pages 478–487, 2007.
- [PV08] Chris Peikert and Vinod Vaikuntanathan. Noninteractive statistical zero-knowledge proofs for lattice problems. In *CRYPTO*, pages 536–553, 2008.

- [PVW08] Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. In *CRYPTO*, pages 554–571, 2008.
- [PW08] Chris Peikert and Brent Waters. Lossy trapdoor functions and their applications. In *Proceedings of STOC '08 (Symposium on Theory of Computing)*, pages 187–196, 2008. Invited to SIAM Journal on Computing special issue on STOC '08.
- [Reg04] Oded Regev. New lattice-based cryptographic constructions. *J. ACM*, 51(6):899–942, 2004.
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, pages 84–93, 2005.
- [RR06] Oded Regev and Ricky Rosen. Lattice problems and norm embeddings. In *STOC*, pages 447–456, 2006.
- [RSA78] Ronald L. Rivest, Adi Shamir, and Leonard M. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM*, 21(2):120–126, 1978.
- [Sho97] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, 1997.