Chris Peikert – Research Statement

My research is dedicated to developing *new, stronger mathematical foundations for cryptography*, with a particular focus on geometric objects called *lattices*. Informally, a lattice is a periodic 'grid' of points in n-dimensional real space $\mathbb{R}^n$. Lattices have been studied since the early 1800s, and their apparent simplicity belies many deep connections and applications across mathematics, physics, and computer science.

In cryptography, the security of most systems necessarily relies upon computational problems that are conjectured to be *intractable*, i.e., infeasible to solve with any realistic amount of computational resources. Over the past three decades, the most useful candidate hard problems have come from an area of mathematics called *number theory*. For instance, a commonly made conjecture is that it is infeasible to compute the prime factors of huge random integers. However, the relatively high computational cost, and largely sequential nature, of operating on such enormous numbers inherently limits the efficiency and applicability of number-theoretic cryptography. Even more worrisome is that *quantum* algorithms, which work in a model of computation that exploits quantum mechanics to dramatically speed up certain kinds of computations, can efficiently solve *all* the number-theoretic problems commonly used in cryptography! Therefore, the future development of a practical, large-scale quantum computer would be devastating to the security of today's cryptographic systems. Alternative foundations are therefore sorely needed.

The seminal works of Ajtai [Ajt96] and Ajtai-Dwork [AD97] in the mid-1990s used conjectured hard problems on lattices as a basis for cryptography. Since then, it has been broadly recognized that lattices have the potential to yield cryptographic schemes with unique and attractive security guarantees—including "worst-case" hardness (explained in the next section) and resistance to quantum attacks—and high levels of *asymptotic* efficiency and parallelism. However, until 2007 only a few very basic lattice-based objects (having limited applicability) were known, and in practice they were very inefficient and so mainly of theoretical interest. Over the past several years, my research has contributed ground-breaking progress toward realizing the full potential of lattices in cryptography, by (1) strengthening the theoretical foundations of the area, (2) designing new cryptographic constructions that enjoy rich functionality and strong security properties, and (3) making lattice cryptography efficient and practical via new design paradigms, fast algorithms, and optimized implementations.

Impact. My work in lattice cryptography has greatly contributed to a recent explosion of activity and interest in the area, and to other exciting developments in cryptography more broadly. As one indication of impact, cryptographic research on lattices in the two decades prior to 2006 was relatively rare. Most of these works were focused on using lattices to *break* various cryptographic proposals, e.g., "knapsack" cryptosystems and variants of RSA. A large fraction of the remaining works were substantial (and important) refinements of the seminal papers, but did not provide any *new* cryptographic constructions or concepts. As explained in detail below, my work has shown that lattices are unexpectedly versatile and efficient for a wide variety of important cryptographic goals—in fact, for several tasks the *only* known solutions are based on lattices! As a result, top-tier conferences like Crypto and Eurocrypt now frequently have at least one session dedicated entirely to lattice cryptography; many conferences (both theoretical and applied) such as Crypto, Theory of Cryptography, and Security & Cryptography for Networks have featured invited tutorials on the area by myself and others; and several workshops have been held around the world to disseminate new developments.

More generally, my work on lattices has introduced new perspectives and tools to cryptography, which in some cases have been applied to solve long-standing problems having nothing at all to do with lattices! Therefore, I find this beautiful subject to be not merely a niche subarea of cryptography, but rather one with the potential to continue shedding light on some of the central questions of the field.

Theoretical Foundations

Cryptography inherently requires a certain kind of intractability, called *average-case* hardness: schemes must be hard to break not only for some possibly rare or hard-to-generate choices of their "keys" (i.e., instances of an underlying problem), but rather for *almost all* choices of *random* keys. This notion of hardness is quite different from, and much less well understood in general than, the *worst-case* notion most commonly studied in the context of algorithms, NP-completeness, etc. For the problems used in cryptography, there is the possibility that they could be hard in the worst case, but easy (or substantially easier) for any usable probability distribution over their instances [Imp95].

The groundbreaking work of Ajtai [Ajt96] proved that several problems related to finding "relatively short" vectors in lattices have a striking *worst-case/average-case equivalence*: solving them on even a tiny fraction of randomly generated lattices (under a particular probability distribution) is provably as hard as solving them on *any* lattice of a related dimension. Moreover, Ajtai's work and follow-ups (e.g., [GGH96, AD97, Reg03, Reg05]) showed that random lattices can be used for basic cryptographic tasks like secure hashing and public-key encryption. Because finding short vectors in high-dimensional lattices has been a notoriously hard algorithmic question for hundreds of years—even when one allows for the power of quantum algorithms (see, e.g., [LLL82, Sch87, SE94, Reg02])—we have solid and unique evidence that lattice-based cryptoschemes are secure. However, many aspects of the seminal results leave room for improvement in various ways, and my work has further cemented the foundations of the area.

Minimizing assumptions. A core endeavor in theoretical cryptography is to determine the weakest intractability assumptions that suffice for achieving various security and functionality goals. As a first example, all the seminal worst-case/average-case reductions for lattices were designed to work for the Euclidean (ℓ_2) norm, i.e., they measured "shortness" in ℓ_2. Of course, ℓ_2 and other norms like ℓ_p are related up to some polynomial factor in the dimension, but this alone leads to worse quantitative bounds for norms other than ℓ_2. Because ℓ_2 is, in a precise sense, the "easiest" norm for lattice problems [RR06], the security of lattice cryptography therefore relied on the qualitatively *strongest* assumptions, at least with respect to the choice of norm. In a work [Pei07] that was selected as one of the top papers at the Conference on Computational Complexity '07, I showed that existing worst-case/average-case reductions for lattice problems actually work for *all* ℓ_p norms simultaneously, for $p \geq 2$, with essentially *no loss* in quality versus the ℓ_2 norm. In other words, in order to break existing lattice-based cryptosystems, one must be able to solve not only all instances of certain lattice problems, but all instances in *all* ℓ_p norms, simultaneously. At a technical level, this work exhibited new properties of "*discrete Gaussian*" probability distributions on lattices, which are at the heart of many of my later works (e.g., [GPV08, Pei10, OPW11, MP12, MP13]), as discussed below.

As a second example, Ajtai's seminal work [Ajt96] constructed a *one-way function*—the 'minimal' cryptographic object—assuming the worst-case hardness of (among others) the approximate Shortest Vector Problem (GapSVP), which is usually considered the most fundamental computational problem on lattices. By contrast, public-key *encryption* schemes like [AD97, Reg03, Reg05] relied either on more specialized (and so possibly easier to break) lattices having some additional geometric structure, or on the explicit assumption that lattice problems are indeed hard for quantum algorithms.[1] For over a decade, a central open question was whether such stronger assumptions were necessary for secure encryption. I resolved this problem, in a work [Pei09] awarded **Best Paper** at the top-tier theory conference STOC'09 (and cited at least 280 times), by giving an encryption scheme that is provably as hard as GapSVP for *classical* algorithms, on *arbitrary* lattices having no extra structure. Building on this work, several coauthors and I recently showed [BLP$^+$13] (in a work at STOC'13) that a wide class of previously incomparable problems used in lattice-based encryption are in fact *equivalent*, thus unifying many disparate assumptions into one.

[1] Even if lattice problems turn out to be easy for quantum algorithms, lattice cryptography might still be secure against *classical* (non-quantum) attacks. That is, quantum-hardness is not a *necessary* assumption for security, but rather a "bonus" property.

Algebraic lattices. I have been a primary investigator of so-called *ideal lattices* and their use in efficient and highly functional cryptographic schemes. Ideal lattices arise naturally in the classical study of number fields, and have algebraic symmetries that enable compact representations and efficient operations (i.e., quasi-linear $\tilde{O}(n)$ space and time), via algorithms like the Fast Fourier Transform. Building on [Mic02], my initial work [PR06] in lattice cryptography (with Rosen) provided the first cryptographically secure hash function based on ideal lattices. Our follow-up work [PR07] from STOC'07 showed the surprising fact that certain families of ideal lattices admit worst-case/average-case reductions with approximation factors as small as $O(\sqrt{\log n})$, versus $\Omega(n)$ for general lattices.

A capstone of my research on ideal lattices is a work [LPR10] with Lyubashevsky and Regev from Eurocrypt'10 and the *Journal of the ACM*, in which we proved the worst-case hardness of a natural average-case problem, called "learning with errors over rings" (ring-LWE). As a first application of ring-LWE, we gave a public-key cryptosystem that can encrypt and decrypt in just polylogarithmic $\tilde{O}(1)$ time per message bit. Subsequently, ring-LWE has served as the foundation for numerous efficient cryptographic constructions, such as pseudorandom functions and fully homomorphic encryption (described in the upcoming sections).

Ongoing and future work. In a recent work [CDLP13] with Chung, Dadush, and Liu, we study the complexity of approximating the "smoothing parameter" of a lattice, which is a fundamental quantity at the heart of the best known worst-case/average-case reductions [MR04, GPV08, MP13] and other important analytical results on lattices (e.g., [Ban93, AR04]). As one consequence, our work gives security proofs of existing cryptographic schemes under the mildest lattice assumptions to date. Another outcome of our work, which I am eager to pursue, is the tantalizing possibility that certain lattice problems may be *complete* for the class SZK of languages having statistical zero-knowledge proofs (see [SV97]). The class SZK contains all problems that have ever been used as a foundation for cryptography, and so SZK-completeness of lattice problems would be exceptionally strong evidence for the security of lattice cryptography, and unprecedented among problems used in cryptography.

Cryptographic Constructions

Broadly speaking, cryptography is concerned with defining and achieving security for complex interactions among potentially malicious entities. As such, the field deals with a wide spectrum of objects—ranging from "one-way" functions (which are necessary for almost any nontrivial cryptography) to full-featured encryption schemes and interactive protocols—and various notions of security for them. A major research endeavor is to design and analyze schemes that combine desirable security, functionality, and efficiency properties. The following describes a few of my main research threads along these lines.

Secure encryption against active attacks. For encryption, two types of security are commonly considered: (1) a basic notion that guarantees the secrecy of messages from any *passive* eavesdropper (who does not alter the ciphertext or interact with the communicating parties), and (2) a "gold standard" notion which preserves secrecy from an *active* adversary that may alter ciphertexts and decrypt arbitrary ones of its choice. Most real-world applications of encryption actually require the latter notion, but rigorously obtaining it (under any natural assumption) is quite challenging. In particular, no lattice-based cryptosystem had active security until my work in STOC'08 with Waters [PW08] (which was selected as one of the top papers of the conference, and has been cited at least 290 times). Our construction worked by way of a novel abstraction called a "*lossy*" trapdoor function. The concept of lossiness has since found many other applications by myself and others, including my work [PVW08] on very simple, efficient, and securely composable instantiations of "oblivious transfer" (a central component of secure protocols), and solutions to the long-standing open problems of deterministic encryption [BFO08] and security under so-called "selective-opening" attacks [BHY09]. In addition, my further work [Pei09, MP12] on active security for lattice-based encryption has refined the somewhat complex and inefficient scheme from [PW08] to be nearly as simple and efficient as passively secure encryption.

Lattice trapdoors and applications. A second major thrust has been to construct lattice-based cryptographic schemes that support flexible and expressive functionality, going far beyond basic authentication and/or confidentiality. A seminal example of this is my work from STOC'08 with Gentry and Vaikuntanathan [GPV08], which introduced techniques for securely using "trapdoors" for lattices in cryptographic schemes. This work, which has been cited more than 540 times, opened up a major new line of research that has led to a vast array of powerful applications by myself and many others, e.g., [PVW08, PV08, CHKP10, ABB10a, GHV10, Boy10, Pei10, ABB10b, GKV10, BF11b, BF11a, OPW11, AFV11, MP12, AP12, GVW13, GSW13, BGG$^+$14, Wic14, GV14], to name just a few.

A main technical result from [GPV08] is an efficient algorithm that randomly samples from a *discrete Gaussian* distribution over a desired lattice coset, given a suitable "trapdoor" for the lattice (roughly analogous to the factorization of a composite integer). A main conceptual contribution is that such discrete Gaussian samples are, in a precise sense, "zero knowledge:" they reveal essentially no information about the trapdoor that was used to generate them. This property can be exploited in many ways to design a variety of cryptographic applications (and it has applications outside cryptography as well).

One main cryptographic application from [GPV08] is a simple "hash-and-sign" digital signature (i.e., message authentication) scheme, which was the first direct construction of lattice-based signatures with a security proof. Indeed, the zero-knowledge property of Gaussian sampling is the key that circumvents certain attacks [NR06, DN12] that had rendered previous heuristic signature proposals [GGH97, HPS01, HHGP$^+$03] *completely insecure*! A second main application is a realization of powerful idea called "identity-based" encryption (IBE), originally conceived of by Shamir [Sha84], which allows a user's name or other identifying string (such as an email address) to serve as her public encryption key. Despite first being envisioned in 1984, no candidate scheme appeared until 17 years later [BF01], using mathematical structures that were completely new to cryptography at the time. In [GPV08] we used discrete Gaussian sampling to give the first IBE from standard lattice problems (in particular, from worst-case assumptions), which remains the only known IBE that has withstood quantum attacks.

A large portion of my work in recent years has been dedicated to further *expanding the range of applications* made possible by the ideas introduced in [GPV08], and giving *more efficient and practical algorithms* to instantiate them. A highlight is my work [CHKP10] with Cash, Hofheinz, and Kiltz, which was awarded **Best Paper** at the top-tier Eurocrypt'10 conference and has been cited at least 270 times.[2] In this work we resolved the main open problem from [GPV08], which was to eliminate the use of an idealized object called a "random oracle" from the digital signature and IBE schemes. In addition, we demonstrated a secure *delegation* mechanism for our IBE, which allows users to pass restricted capabilities to subordinates, in a hierarchical fashion. These techniques are at the heart of countless works that extend the flexibility of our schemes even further.

A few more examples of my works that build upon the foundation of [GPV08] include:

- A work with visiting student Alwen [AP09], which was selected as one of the top papers of STACS'09, that gives a new construction (improving on a rather complex and quantitatively loose one of Ajtai [Ajt99]) for *generating* a worst-case-hard random lattice together with a trapdoor. Notably, our construction is simultaneously optimal in all relevant parameters.
- A work [Pei10] from Crypto'10, which addresses the practical inefficiency and inherent *sequentiality* of the discrete Gaussian sampling algorithm from [GPV08]. It introduces a new analytical technique for discrete Gaussians, which yields an algorithm that is very simple and fast, and optimally parallelizable.
- A work with O'Neill and Waters [OPW11] from Crypto'11 that uses trapdoor techniques to solve a fourteen-year-old question about the existence of "deniable" encryption [CDNO97].

[2] The results from this work were discovered independently by myself and the other three authors, and our papers were merged at the request of the Eurocrypt program committee.

- A work with Micciancio from Eurocrypt'12 [MP12], which was selected as one of the top papers of the conference. It gives a *new trapdoor notion* which is more compact and easy to work with than the one used in [GPV08], along with specialized efficient and parallel algorithms for generating and using such trapdoors. (Our notion is also implicitly at the heart of a recent wave of functional encryption and signature schemes [GSW13, BGG$^+$14, AP14, Wic14, GV14].)
- A work with my PhD student Krehbiel and visiting student Bendlin [BKP13] that gave the first "threshold" versions of trapdoor-using applications, in which trust is split among several users (some of whom may be malicious) who collectively run a distributed protocol to perform privileged operations.

Fully homomorphic encryption. A third area of interest is the exciting concept of "fully homomorphic encryption" (FHE), which allows an untrusted worker to perform computations on encrypted data. Due to its many applications, and the lack of any plausible candidate construction for decades, FHE was often called a "holy grail" of cryptography. In 2009, Gentry [Gen09b, Gen09a] proposed the first candidate FHE construction (based on lattices), which quickly set off a flurry of interest and led to several improvements (e.g., [BV11b, BV11a, BGV12, Bra12, GSW13]). Today, the most efficient FHE schemes are based upon the ring-LWE problem from my foundational work on encryption from algebraic lattices [LPR10, LPR13], discussed in the previous section. However, these FHE schemes are still exceedingly inefficient, due mainly to the high cost of an auxiliary "bootstrapping" operation.

My PhD student Alperin-Sheriff and I have devised new methods for bootstrapping that yield major efficiency improvements. In work from Crypto'13 [AP13], which was selected as one of the top papers of the conference, we gave the fastest bootstrapping algorithm to date, which runs in only quasi-linear $\tilde{O}(n)$ time in the bit length of the ciphertext. In addition, the hidden factors are relatively small, and the method seems suitable for practice. Our method builds upon a "ring-switching" technique that I and coauthors devised for entirely different purposes [GHPS12], and homomorphically evaluates a Fast Fourier-like transform by switching through a carefully designed sequence of "hybrid" rings.

More recently, a new kind of FHE scheme and bootstrapping approach [GSW13, BV14] have emerged, based on quantitatively much milder lattice assumptions, but at the cost of vastly more expensive bootstrapping operations (i.e., very large polynomial runtimes). In work from Crypto'14 [AP14], we gave a substantially different bootstrapping method that is based on very mild assumptions and is quasi-optimal in the number of homomorphic operations required. It has subsequently been implemented [DM14], and is many hundred times faster than the previous state of the art.

Pseudorandom functions and future work. A final research direction that I am especially excited about concerns *pseudorandom functions* (PRFs), which are the central objects in symmetric-key cryptography. While PRFs are theoretically much "easier" to obtain (via general-purpose but inefficient and sequential transformations) than rich objects like IBE described above, until recently they were one of the last fundamental objects that lacked a direct lattice-based realization. In a work [BPR12] from Eurocrypt'12, my PhD student Banerjee, Rosen and I devised new approaches and proof techniques to construct the first PRFs directly from lattice problems, yielding efficient and highly parallel constructions. One interesting aspect of our work is that for our lowest-depth constructions, our security proof requires quite large parameters and strong hardness assumptions. However, the constructions themselves appear to be secure (they resist all known attacks) for very small parameters! This mismatch may be an artifact of our proof technique, and obtaining a security proof that yields the "right" parameters is a fascinating theoretical question that I continue to pursue.

A second research thread relates to PRFs having additional properties that can be useful in applications. A nice work of Boneh *et al.* [BLMR13] showed that a variant of our PRF from [BPR12] is the first known one to have "key homomorphism," a very useful property with several applications. However, their construction was highly inefficient, with large polynomial key sizes and runtimes. In follow-up work [BP14], Banerjee and I gave much more efficient, compact, and general constructions of key-homomorphic PRFs. As an unexpected side benefit, these are also the first known PRFs having both good parallelism and quasi-optimal key sizes. Looking forward, a sizable body of recent research has demonstrated the power and importance of so-called "functional" or "constrained" PRFs in a variety of contexts. Building on our work from [BPR12, BP14], we have obtained some initial results [BFP$^+$14] in constructing such objects, and there is a vast horizon waiting to be explored.

Efficiency and Towards Practice

In terms of performance, lattice cryptography has long been recognized for its promise of high *asymptotic* efficiency. In practice, however, most theoretically sound schemes suffer from impractically large keys and runtimes. The main reason is that in general, it requires at least quadratic $\Omega(n^2)$ space and time to specify and operate on an n-dimensional lattice, and n needs to be at least in the few hundreds for sufficient security. As already mentioned above in the section on algebraic lattices, my research has started to bridge this gap between theory and practice by constructing compact, fast cryptographic objects with sound proofs of security (under reasonable complexity assumptions). Going further, a major goal of mine is to see lattice cryptography brought successfully to practice in the real world. My work in this direction falls into a few broad, mutually reinforcing threads.

Algorithms and analysis. First, lattice cryptography in practice will require highly efficient specialized *algorithms* and sharp *analysis* to minimize runtimes, parameters, key sizes, etc. For example, my work with Micciancio [MP12] completely reworked and simplified all the central operations and analytical tools used by all "trapdoor" applications discussed above (like digital signatures and IBE), yielding major practical improvements over previous methods. As another example, in a companion work [LPR13] to the ring-LWE paper [LPR10], we designed a toolkit of very fast, customized algorithms and analytical techniques for optimal use of ring-LWE and ideal lattices in cryptography, across all applications. These tools have paid unexpected dividends, as they have become central to the unexpected functionality and substantial efficiency improvements we exhibited [GHPS12, AP13] for fully homomorphic encryption schemes (described in the previous section).

Constructions and security estimates. Second, moving to practice will require not just general-purpose constructions and asymptotic bounds, but *optimized schemes* for particular applications, with *concrete parameters and estimated security levels*. To that end, my work [LP11] with PhD student Lindner showed that concrete instantiations of ring-LWE encryption (from [LPR10]) can have key sizes and security levels comparable to those of the widely used RSA cryptosystem, while running many times faster. In addition, a recent work of mine [Pei14] proposed "drop-in" lattice-based components for core Internet protocols like key exchange, along with various optimizations that improve ciphertext lengths by about a factor of two over the prior best.

Implementations. Third, practical lattice cryptography obviously will require *implementations* and benchmarks. For example, with colleagues I expanded my early work [PR06] on ideal lattice-based hash functions into a fast implementation [LMPR08] and a proposal [ADL$^+$08] to NIST's SHA-3 hash function competition. More recently, with colleagues I implemented and benchmarked [BBL$^+$14] the PRF design from [BPR12], and showed that with only basic optimizations, it runs only 2-3 times slower than highly optimized software implementations of AES, the national standard. It also turns out that the algebraic nature of our PRF is highly amenable to *homomorphic* evaluation, which has many attractive applications in theory and practice (see, e.g., [GGI$^+$14]). Using techniques from my work on bootstrapping [AP13], we have demonstrated [ACPR14] that our PRF can be evaluated homomorphically with more than a *thousand-fold* speedup and memory reduction versus AES (cf. [GHS12]). Finally, my students and I have made major progress on the design and implementation of a general-purpose, modular, and highly parallel library for lattice cryptography, which we intend to serve as a platform to bring the many exciting theoretical applications of lattices to practice.

Additional Research Interests

I also have significant interests in *algorithms* for hard lattice problems like the ones used in cryptography, and in the use of lattices for *error correction* and reliable communication.

Because many lattice problems are conjectured to be hard, and some are even NP-complete, we do not expect to find efficient algorithms for them. However, due to their utility in areas like optimization (e.g., integer programming), their usefulness in low dimensions as subroutines to other algorithms, and their importance in cryptanalysis, it is very important to investigate the most efficient algorithms for hard lattice problems. For instance, several problems are not yet even known to be solvable in simply exponential $2^{O(n)}$ time! A major development by Micciancio and Voulgaris [MV10] gave the first deterministic, exponential-time algorithm for the shortest and closest lattice vector problems, but it was highly specialized to the Euclidean norm. In a work [DPV11] with former GT PhD student Dadush and colleague Vempala, which appeared in the top-tier theory conference FOCS'11, we gave a new "enumerative" algorithm that lists all the points of a given lattice inside an arbitrary convex body. We applied this powerful tool to give the first deterministic, $2^{O(n)}$-time algorithm for the exact shortest lattice vector problem, in *any* norm.

In the theory of error-correcting codes, one of the most exciting modern developments was the discovery that it is possible to efficiently recover from significantly more error if one allows the decoder to output a short *list* of all possible original codewords within the error bound, rather than insisting on a unique answer. Beyond its natural applications to reliable communication and storage, this discovery of efficient "*list decoding*" algorithms has had profound implications for cryptography, learning theory, and computational complexity. In work [GP12] at the Conference on Computational Complexity '12, my former postdoc Grigorescu and I initiated the study of list decoding for lattices, using the Euclidean (ℓ_2) norm to measure the amount of error. Our main result is an efficient list decoder for the well-studied Barnes-Wall family of lattices: our decoder runs in time polynomial in the list size for any desired error tolerance, is perfectly parallelizable, and is even quite efficient in practice. We continue to investigate list decoding algorithms for other families of lattices, and new families of lattices that admit good tradeoffs between list size and error tolerance (along with efficient decoding algorithms).

References

[ABB10a] S. Agrawal, D. Boneh, and X. Boyen. Efficient lattice (H)IBE in the standard model. In *EUROCRYPT*, pages 553–572. 2010.

[ABB10b] S. Agrawal, D. Boneh, and X. Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In *CRYPTO*, pages 98–115. 2010.

[ACPR14] G. Alberini, E. Crockett, C. Peikert, and A. Rosen. Fast homomorphic evaluation of symmetric-key primitives, 2014. Submitted.

[AD97] M. Ajtai and C. Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *STOC*, pages 284–293. 1997.

[ADL$^+$08] Y. Arbitman, G. Dogon, V. Lyubashevsky, D. Micciancio, C. Peikert, and A. Rosen. SWIFFTX: A proposal for the SHA-3 standard, 2008. Submitted to NIST SHA-3 competition.

[AFV11] S. Agrawal, D. M. Freeman, and V. Vaikuntanathan. Functional encryption for inner product predicates from learning with errors. In *ASIACRYPT*. 2011.

[Ajt96] M. Ajtai. Generating hard instances of lattice problems. *Quaderni di Matematica*, 13:1–32, 2004. Preliminary version in STOC 1996.

[Ajt99] M. Ajtai. Generating hard instances of the short basis problem. In *ICALP*, pages 1–9. 1999.

[AP09] J. Alwen and C. Peikert. Generating shorter bases for hard random lattices. *Theory of Computing Systems*, 48(3):535–553, April 2011. By invitation as a top paper from STACS '09.

[AP12] J. Alperin-Sheriff and C. Peikert. Circular and KDM security for identity-based encryption. In *Public Key Cryptography*, pages 334–352. 2012.

[AP13] J. Alperin-Sheriff and C. Peikert. Practical bootstrapping in quasilinear time. In *Proceedings of CRYPTO '13*, pages 1–20. 2013. **One of a few top papers invited to J. Cryptology**.

[AP14] J. Alperin-Sheriff and C. Peikert. Faster bootstrapping with polynomial error. In *CRYPTO*, pages 297–314. 2014.

[AR04] D. Aharonov and O. Regev. Lattice problems in NP $\cap$ coNP. *J. ACM*, 52(5):749–765, 2005. Preliminary version in FOCS 2004.

[Ban93] W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(4):625–635, 1993.

[BBL$^+$14] A. Banerjee, H. Brenner, G. Leurent, C. Peikert, and A. Rosen. SPRING: Fast pseudorandom functions from rounded ring products. In *Proceedings of FSE '14 (Fast Software Encryption)*, pages ??–?? 2014.

[BF01] D. Boneh and M. K. Franklin. Identity-based encryption from the Weil pairing. *SIAM J. Comput.*, 32(3):586–615, 2003. Preliminary version in CRYPTO 2001.

[BF11a] D. Boneh and D. M. Freeman. Homomorphic signatures for polynomial functions. In *EUROCRYPT*, pages 149–168. 2011.

[BF11b] D. Boneh and D. M. Freeman. Linearly homomorphic signatures over binary fields and new tools for lattice-based signatures. In *Public Key Cryptography*, pages 1–16. 2011.

[BFO08] A. Boldyreva, S. Fehr, and A. O'Neill. On notions of security for deterministic encryption, and efficient constructions without random oracles. In *CRYPTO*, pages 335–359. 2008.

[BFP$^+$14] A. Banerjee, G. Fuchsbauer, C. Peikert, K. Pietrzak, and S. Stevens. Key-homomorphic constrained pseudorandom functions, 2014. Submitted.

[BGG$^+$14] D. Boneh, C. Gentry, S. Gorbunov, S. Halevi, V. Nikolaenko, G. Segev, V. Vaikuntanathan, and D. Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In *EUROCRYPT*, pages 533–556. 2014.

[BGV12] Z. Brakerski, C. Gentry, and V. Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. In *ITCS*, pages 309–325. 2012.

[BHY09] M. Bellare, D. Hofheinz, and S. Yilek. Possibility and impossibility results for encryption and commitment secure under selective opening. In *EUROCRYPT*, pages 1–35. 2009.

[BKP13] R. Bendlin, S. Krehbiel, and C. Peikert. How to share a lattice trapdoor: Threshold protocols for signatures and (H)IBE. In *ACNS*, pages 218–236. 2013.

[BLMR13] D. Boneh, K. Lewi, H. W. Montgomery, and A. Raghunathan. Key homomorphic PRFs and their applications. In *CRYPTO*, pages 410–428. 2013.

[BLP$^+$13] Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D. Stehlé. Classical hardness of learning with errors. In *STOC*, pages 575–584. 2013.

[Boy10] X. Boyen. Lattice mixing and vanishing trapdoors: A framework for fully secure short signatures and more. In *Public Key Cryptography*, pages 499–517. 2010.

[BP14] A. Banerjee and C. Peikert. New and improved key-homomorphic pseudorandom functions. In *CRYPTO*, pages 353–370. 2014.

[BPR12] A. Banerjee, C. Peikert, and A. Rosen. Pseudorandom functions and lattices. In *EUROCRYPT*, pages 719–737. 2012.

[Bra12] Z. Brakerski. Fully homomorphic encryption without modulus switching from classical GapSVP. In *CRYPTO*, pages 868–886. 2012.

[BV11a] Z. Brakerski and V. Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In *FOCS*, pages 97–106. 2011.

[BV11b] Z. Brakerski and V. Vaikuntanathan. Fully homomorphic encryption from ring-LWE and security for key dependent messages. In *CRYPTO*, pages 505–524. 2011.

[BV14] Z. Brakerski and V. Vaikuntanathan. Lattice-based FHE as secure as PKE. In *ITCS*, pages 1–?? 2014.

[CDLP13] K.-M. Chung, D. Dadush, F.-H. Liu, and C. Peikert. On the lattice smoothing parameter problem. In *Proceedings of CCC '13 (Conference on Computational Complexity)*. 2013.

[CDNO97] R. Canetti, C. Dwork, M. Naor, and R. Ostrovsky. Deniable encryption. In *CRYPTO*, pages 90–104. 1997.

[CHKP10] D. Cash, D. Hofheinz, E. Kiltz, and C. Peikert. Bonsai trees, or how to delegate a lattice basis. In *Proceedings of EUROCRYPT '10*, pages 523–552. 2010. **Awarded Best Paper.**

[DM14] L. Ducas and D. Micciancio. FHE bootstrapping in less than a second. Cryptology ePrint Archive, Report 2014/816, 2014. `http://eprint.iacr.org/`.

[DN12] L. Ducas and P. Q. Nguyen. Learning a zonotope and more: Cryptanalysis of NTRUSign countermeasures. In *ASIACRYPT*, pages 433–450. 2012.

[DPV11] D. Dadush, C. Peikert, and S. Vempala. Enumerative lattice algorithms in any norm via M-ellipsoid coverings. In *FOCS*, pages 580–589. 2011.

[Gen09a] C. Gentry. *A fully homomorphic encryption scheme*. Ph.D. thesis, Stanford University, 2009. `http://crypto.stanford.edu/craig`.

[Gen09b] C. Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178. 2009.

[GGH96] O. Goldreich, S. Goldwasser, and S. Halevi. Collision-free hashing from lattice problems. *Electronic Colloquium on Computational Complexity (ECCC)*, 3(42), 1996.

[GGH97] O. Goldreich, S. Goldwasser, and S. Halevi. Public-key cryptosystems from lattice reduction problems. In *CRYPTO*, pages 112–131. 1997.

[GGI^{+}14] C. Gentry, J. Groth, Y. Ishai, C. Peikert, A. Sahai, and A. Smith. Using fully homomorphic hybrid encryption to minimize non-interactive zero-knowledge proofs. *Journal of Cryptology*, 2014. Accepted, to appear.

[GHPS12] C. Gentry, S. Halevi, C. Peikert, and N. Smart. Ring switching in BGV-style homomorphic encryption. In *Proceedings of SCN '12 (Security and Cryptography for Networks)*, pages 19–37. 2012. **One of a few top papers invited to Journal of Computer Security**.

[GHS12] C. Gentry, S. Halevi, and N. P. Smart. Homomorphic evaluation of the AES circuit. In *CRYPTO*, pages 850–867. 2012.

[GHV10] C. Gentry, S. Halevi, and V. Vaikuntanathan. A simple BGN-type cryptosystem from LWE. In *EUROCRYPT*, pages 506–522. 2010.

[GKV10] S. D. Gordon, J. Katz, and V. Vaikuntanathan. A group signature scheme from lattice assumptions. In *ASIACRYPT*, pages 395–412. 2010.

[GP12] E. Grigorescu and C. Peikert. List decoding Barnes-Wall lattices. In *Proceedings of CCC '12 (Conference on Computational Complexity)*. 2012.

[GPV08] C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206. 2008.

[GSW13] C. Gentry, A. Sahai, and B. Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In *CRYPTO*, pages 75–92. 2013.

[GV14] S. Gorbunov and V. Vaikuntanathan. (Leveled) fully homomorphic signatures from lattices. Cryptology ePrint Archive, Report 2014/463, 2014. `http://eprint.iacr.org/`.

[GVW13] S. Gorbunov, V. Vaikuntanathan, and H. Wee. Attribute-based encryption for circuits. In *STOC*, pages 545–554. 2013.

[HHGP^{+}03] J. Hoffstein, N. Howgrave-Graham, J. Pipher, J. H. Silverman, and W. Whyte. NTRUSIGN: Digital signatures using the NTRU lattice. In *CT-RSA*, pages 122–140. 2003.

[HPS01] J. Hoffstein, J. Pipher, and J. H. Silverman. NSS: an NTRU lattice-based signature scheme. In *EUROCRYPT*, pages 211–228. 2001.

[Imp95] R. Impagliazzo. A personal view of average-case complexity. In *Structure in Complexity Theory Conference*, pages 134–147. 1995.

[LLL82] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, December 1982.

[LMPR08] V. Lyubashevsky, D. Micciancio, C. Peikert, and A. Rosen. SWIFFT: A modest proposal for FFT hashing. In *FSE*, pages 54–72. 2008.

[LP11] R. Lindner and C. Peikert. Better key sizes (and attacks) for LWE-based encryption. In *CT-RSA*, pages 319–339. 2011.

[LPR10] V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. *Journal of the ACM*, 60(6):43:1–43:35, November 2013. Preliminary version in EUROCRYPT '10.

[LPR13] V. Lyubashevsky, C. Peikert, and O. Regev. A toolkit for ring-LWE cryptography. In *EUROCRYPT*, pages 35–54. 2013.

[Mic02] D. Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Computational Complexity*, 16(4):365–411, 2007. Preliminary version in FOCS 2002.

[MP12] D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *Proceedings of EUROCRYPT '12*, pages 700–718. 2012. **One of three top papers invited to J. Cryptology**.

[MP13] D. Micciancio and C. Peikert. Hardness of SIS and LWE with small parameters. In *CRYPTO*, pages 21–39. 2013.

[MR04] D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007. Preliminary version in FOCS 2004.

[MV10] D. Micciancio and P. Voulgaris. A deterministic single exponential time algorithm for most lattice problems based on Voronoi cell computations. In *STOC*, pages 351–358. 2010.

[NR06] P. Q. Nguyen and O. Regev. Learning a parallelepiped: Cryptanalysis of GGH and NTRU signatures. *J. Cryptology*, 22(2):139–160, 2009. Preliminary version in Eurocrypt 2006.

[OPW11] A. O'Neill, C. Peikert, and B. Waters. Bi-deniable public-key encryption. In *CRYPTO*, pages 525–542. 2011.

[Pei07] C. Peikert. Limits on the hardness of lattice problems in ℓ_p norms. *Computational Complexity*, 17(2):300–351, May 2008. By invitation as a top paper from CCC '07.

[Pei09] C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *Proceedings of STOC '09 (Symposium on Theory of Computing)*, pages 333–342. 2009. **Awarded Best Paper.**

[Pei10] C. Peikert. An efficient and parallel Gaussian sampler for lattices. In *CRYPTO*, pages 80–97. 2010.

[Pei14] C. Peikert. Lattice cryptography for the internet. In *PQCrypto*, page ?? 2014.

[PR06] C. Peikert and A. Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, pages 145–166. 2006.

[PR07] C. Peikert and A. Rosen. Lattices that admit logarithmic worst-case to average-case connection factors. In *STOC*, pages 478–487. 2007.

[PV08] C. Peikert and V. Vaikuntanathan. Noninteractive statistical zero-knowledge proofs for lattice problems. In *CRYPTO*, pages 536–553. 2008.

[PVW08] C. Peikert, V. Vaikuntanathan, and B. Waters. A framework for efficient and composable oblivious transfer. In *CRYPTO*, pages 554–571. 2008.

[PW08] C. Peikert and B. Waters. Lossy trapdoor functions and their applications. *SIAM J. Comput.*, 40(6):1803–1844, December 2011. By invitation to special issue on STOC '08.

[Reg02] O. Regev. Quantum computation and lattice problems. *SIAM J. Comput.*, 33(3):738–760, 2004. Preliminary version in FOCS 2002.

[Reg03] O. Regev. New lattice-based cryptographic constructions. *J. ACM*, 51(6):899–942, 2004. Preliminary version in STOC 2003.

[Reg05] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):1–40, 2009. Preliminary version in STOC 2005.

[RR06] O. Regev and R. Rosen. Lattice problems and norm embeddings. In *STOC*, pages 447–456. 2006.

[Sch87] C.-P. Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theor. Comput. Sci.*, 53:201–224, 1987.

[SE94] C.-P. Schnorr and M. Euchner. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Mathmatical Programming*, 66:181–199, 1994.

[Sha84] A. Shamir. Identity-based cryptosystems and signature schemes. In *CRYPTO*, pages 47–53. 1984.

[SV97] A. Sahai and S. P. Vadhan. A complete problem for statistical zero knowledge. *J. ACM*, 50(2):196–249, 2003. Preliminary version in FOCS 1997.

[Wic14] D. Wichs. Leveled fully homomorphic signatures from standard lattices. Cryptology ePrint Archive, Report 2014/451, 2014. `http://eprint.iacr.org/`.