Research Statement

Salil Vadhan

September 2006

1 Summary

My research centers around the interface between *computational complexity theory* and *cryptography*. Complexity theory studies the power and limitations of efficient computation, and cryptography aims to design protocols that withstand adversarial behavior. Within these general areas, I have focused on the topics of *pseudorandomness*, the theory of efficiently generating objects that 'look random' despite being constructed with little or no randomness, and *zero-knowledge proofs*, which are interactive proofs that are convincing yet reveal nothing other than the validity of the assertion being proven. My most significant research contribution in the area of pseudorandomness has been the discovery of the *zig-zag graph product* for constructing *expander graphs* [RVW], which has been used in or inspired the solution of several long-standing open questions in theoretical computer science [RVW, ALW, CRVW, Rei, Din]. In the area of zero-knowledge proofs, my work began with my Ph.D. thesis [Vad1], which provided a comprehensive complexity-theoretic understanding of a subclass of zero-knowledge proofs known as *statistical zero-knowledge proofs*. In the last few years, my students and I have returned to this topic, and have shown that the techniques and results of my thesis have significance far beyond the scope of statistical zero knowledge. Specifically, in a series of papers [MV, Vad2, NV, NOV], we have managed to reduce or eliminate the complexity assumptions used in many fundamental results about zero-knowledge proofs, resolving at least one long-standing open problem.

2 General Research Areas

Computational Complexity. The goal of computational complexity theory is, broadly speaking, to understand the power of efficient computation.

That is, it asks:

What problems can and cannot be solved with limited computational resources?

The 'problems' studied include not only ones from computer science but ones from mathematics (factoring integers), economics (finding Nash equilibria), physics (sampling random configurations of a physical system), communications (decoding error-correcting codes), operations research (network flow); and the 'resources' studied include time, space/memory, randomness, interaction, and quantum mechanics. The above question is partly addressed by exhibiting algorithms that solve particular problems efficiently (e.g. an algorithm that factors n-digit numbers in $2^{\sqrt{n}}$ steps), but a major objective of complexity theory is to also understand what *cannot* be computed efficiently (e.g. prove that there does not exist an algorithm that, no matter how cleverly designed, factors n-digit numbers in n^3 steps).

The questions of computational complexity are some of the most basic in computer science, yet many remain unresolved despite several decades of effort, particularly those requiring *lower bounds* on the resources needed to solve problems. However, computational complexity has been and continues to be extremely successful in establishing *relationships* between seemingly unrelated questions. One of the first examples was the beautiful theory of **NP**-completeness [Coo, Lev, Kar], which showed that thousands of natural problems are computationally equivalent, in that all of them have efficient algorithms or all of them do not. And surprising new relationships continue to be discovered; some examples are discussed below in the context of my work on the theory of pseudorandomness.

In addition to its role in understanding the power and limitations of computer technology, complexity theory provides an illuminating perspective on questions arising in other fields, explaining why certain problems resist analysis, and on basic philosophical questions, such as whether finding solutions to problems is ever harder than verifying their correctness. The latter is the famous **P** vs. **NP** question and amounts to asking to what extent can the work of mathematicians be automated (posed by Gödel in a letter to von Neumann in 1956). Thus it is also viewed as one of the most important open problems in mathematics, and is one of the seven "millennium prize problems" of the Clay Math Institute. In terms of practical computing technology, however, the most significant impact of complexity theory has probably come through its role in the foundations of cryptography, my other main research area.

Cryptography. The aim of cryptography is to

Design protocols that withstand adversarial behavior.

That is, we wish to construct algorithms and protocols that guarantee privacy, authenticity, and integrity of data when parties are communicating or computing in an insecure environment. This subject provides the technology that allows us to make on-line purchases without hackers being able to learn our credit card numbers. Historically, when cryptography was used mainly for military purposes, it was more of an art than a science, with a seemingly never-ending cycle of schemes being broken and repaired. Indeed, Shannon [Sha] explained the lack of mathematically rigorous guarantees in cryptography, showing that (essentially) no encryption scheme can be "perfectly secure." It was the development of computational complexity theory in the 1970's that opened the door to bypassing this barrier [DH]. Specifically, we could hope to have cryptosystems that provably require an infeasible amount of computational resources to break. Unfortunately, as mentioned above, proving strong lower bounds on the resources needed to solve computational problems (in this case, breaking a cryptosystem) still seems beyond the reach of current techniques in complexity theory. Indeed, proving the security of a cryptosystem requires, at a minimum, resolving the **P** vs. **NP** question.

Nevertheless, the language and framework provided by complexity theory has enabled cryptography to flourish over the past three decades, and develop into a mature science. It has formulated very convincing and precise definitions of security for a plethora of cryptographic tasks (ranging from basic primitives such as encryption to complex protocols such as electronic voting) in a variety of adversarial environments (including coordinated attacks on many concurrently executing protocols). And it has provided cryptosystems and protocols whose security can be provably reduced to the intractability of some basic computational problem (e.g. factoring large integers). The part of cryptography that remains an art is the choice of these underlying computational problems. Whether we use well-studied and mathematically 'clean' problems such as integer factorization [RSA, Rab] or man-made constructs such as the Advanced Encryption Standard [AES] their intractability, and thus the security of our cryptosystems, remains ultimately based on conjecture. The reality of this threat is evidenced by the recent breaks of the widely used cryptographic hash functions MD5 and SHA-0 [WLF$^+$, WY, BCJ$^+$]. Thus a major goal in the foundations of cryptography is to base cryptography on complexity assumptions that are as weak and general as possible, so that if one problem turns out to be

easy (e.g. a fast algorithm for integer factorization is discovered), we can easily replace it with a variety of alternatives. Much of my recent work on zero-knowledge proofs, described in Section 4, fits within this project of minimizing the complexity assumptions in cryptography.

The Interface. As described above, complexity theory provides the conceptual framework on which modern cryptography is built, and resolving the major open questions of complexity theory is necessary to have cryptosystems whose security does not rely on unproven conjectures.[1] However, the interaction between complexity theory and cryptography has not been unidirectional. Rather, questions arising in cryptography have led to the development of entirely new areas and have sparked some of the most exciting developments in complexity theory. This interface has remained extremely rich and fertile, with continuing benefits for both cryptography and complexity theory. Two of the most active areas of interaction are the topics of *pseudorandomness* and *zero-knowledge proofs*, which are the focuses of my own research.

3 Pseudorandomness and Expander Graphs

Pseudorandomness is the theory of

Efficiently generating objects that "look random"
despite being constructed using little or no randomness.

The modern form of this theory originated from research in cryptography, where one often needs to generate a large number of unpredictable bits (e.g. for encryption or authentication) from a short random key [BM, Yao]. However, it has also become a major topic within computational complexity theory, because of the insight it gives us into the power of randomness for efficient computation. In the 70's and 80's, it was realized that randomization is extremely useful in the design of algorithms and protocols, as researchers found randomized solutions to a wide variety of problems for which no deterministic solutions were known. However, in many cases, it was not known (and is still not known) to what extent the randomness is really necessary. Thus an intriguing question, of both theoretical and practical

[1] There are certain models in which unconditionally secure cryptosystems can be constructed, such as the "bounded-storage model" [Mau, ADR], but if we have a standard communication channel and measure the adversary's resources by computation time, then almost any cryptographic task implies the existence of one-way functions and $\mathbf{P} \neq \mathbf{NP}$.

importance, is whether we can reduce the amount or quality of randomness required for solving these problems — ideally to full *derandomization*, where we eliminate the randomness entirely. Pseudorandomness provides a general approach for doing this.

Through two decades of research, pseudorandomness has proved to be a fundamental concept in theoretical computer science, finding applications in areas such as cryptography, computational complexity, learning theory, and algorithm design. In recent years, however, our understanding of this area has increased dramatically. Specifically, through the work of myself and others (notably Trevisan [Tre]), a number of previously distinct research directions were unified. We discovered that four fundamental and heavily studied objects — pseudorandom generators, randomness extractors, expander graphs, and error-correcting codes — are all essentially *equivalent*. It was the exciting potential of this unified viewpoint that prompted me to make pseudorandomness my main research area during my first few years at Harvard.

My most significant contribution in this topic was the discovery, with Reingold and Wigderson [RVW], of the *zig-zag graph product* for constructing *expander graphs*. Expander graphs are networks that are sparse yet very highly connected. Expanders have a wide variety of applications in computer science, ranging from network design to coding theory to data storage, and thus there is a long and celebrated body of work on constructions of expander graphs. We introduced the zig-zag product as a new tool for constructing expanders, based on the connection between expander graphs and randomness extractors mentioned above. Since then, it has found a variety of applications, in some cases to resolving long-standing open problems. Examples include:

- In our original paper [RVW], we showed how to use the zig-zag product to obtain a simple, combinatorial construction of expanders, achieving a goal that had eluded researchers for decades. (Previous constructions were algebraic and provided little control over or intuition for the expansion property.)

- The zig-zag product was a crucial component of Reingold's breakthrough logarithmic-space algorithm for connectivity in undirected graphs [Rei], which resolved the space complexity of one of the most basic problems in computer science.

- The uses of the zig-zag product in [RVW, Rei] inspired Dinur's beautiful new combinatorial proof of the celebrated PCP Theorem [Din].

(The PCP Theorem shows that mathematical proofs can be encoded in such a way that one needs to read only a constant number of randomly chosen bits of the proof to verify correctness with high confidence; it has many applications to understanding the complexity of approximation problems.)

- With Capalbo, Reingold, and Wigderson [CRVW], I used a variant of the zig-zag product to give the first explicit construction of constant-degree expanders with near-optimal expansion, bypassing a barrier of previous methods ("eigenvalue methods") and thereby enabling a number of the applications of expanders to be efficiently realized.

- Alon, Lubotzky, and Wigderson [ALW] used the zig-zag product to disprove a group-theoretic conjecture of Lubotzky and Weiss [LW].

I have done a substantial amount of additional work on other aspects of pseudorandomness and its applications; I mention a few highlights here:

- With Lu, Reingold, and Wigderson [LRVW], I gave the first construction of randomness extractors that are "optimal up to constant factors," reaching a milestone in a decade of work on the subject. (Randomness extractors are algorithms for extracting almost-uniform bits from sources of biased and correlated bits, and have many applications beyond their original purpose of simulating randomized algorithms with weak random sources.)

- With Sudan and Trevisan [STV, TV2], I established essentially optimal relationships between the worst-case complexity and average-case complexity of exponential time (**EXP**), first for circuit complexity and then for uniform complexity. Such relationships are the starting point for pseudorandom generator constructions. These works also established a link between pseudorandom generators and list-decodable error-correcting codes, which has played an important role in subsequent developments in the area.

- With graduate students Healy and Viola [HVV], I have obtained the strongest known results on amplifying the average-case complexity of **NP** (rather than **EXP**), going beyond barriers that were proven to hold for previous approaches.

- With Trevisan, Zuckerman, Kamp, and Rao [TV1, TVZ, KRVZ], I developed a theory of randomness extraction (and data compression)

from *samplable sources*, which are random sources generated by an (unknown) efficient algorithm. The advantage of the restriction to samplable sources is that randomness extractors for such sources do not require a 'seed' and thus can be used to purify randomness when one cannot afford to choose the seed uniformly at random or enumerate over all choices, e.g. when purifying a physical source of randomness for use in cryptography.

- With Barak and then-undergraduate Ong [BOV], I gave the first applications of "Nisan–Wigderson-type" pseudorandom generators in cryptography. In particular, we give the first construction of one-message witness-indistinguishable proofs for NP. (Witness-indistinguishable proofs are a natural weakening of zero-knowledge proofs that suffice for many applications.) With Barak and Lindell [BLV], I used these same techniques to give some of the first lower bounds on the round complexity of general zero-knowledge proofs (as opposed to "black-box zero knowledge," a restriction needed by previous lower bounds.)

4 Zero-Knowledge Proofs

Zero-knowledge proofs are interactive protocols whereby a "prover" can convince a "verifier" that some assertion is true, with the remarkable property that

The verifier learns nothing other than the fact that the assertion being proven is true.

Since their introduction two decades ago [GMR], zero-knowledge proofs have taken on a central role in the design and study of cryptographic protocols. They provide a powerful building block for secure protocols, because they can be used in any situation where one participant needs to convince another of some fact (e.g. that it has not deviated from the specified protocol) without revealing its secret information (e.g. encryption keys) [GMW]. In addition, they are the most common testbed for studying new issues in cryptographic protocols, such as composability and concurrency.

In addition to their role in cryptography, however, zero-knowledge proofs have also provided one of the main avenues of interaction between complexity theory and cryptography. This stems in part from the central role that the "efficiently verifiable proofs" play in complexity theory (e.g. they are the subject of the **P** vs. **NP** question); zero knowledge enriches this study with such fascinating ingredients as interaction, randomness, knowledge,

and secrecy. Indeed, the study of zero knowledge has led to some of the most exciting developments in complexity theory, such as the construction of probabilistically checkable proofs (proofs that can be verified by reading just a constant number of random locations), which in turn revolutionized our understanding of the complexity of finding approximate solutions to **NP**-complete optimization problems.

My work in this area has aimed to use complexity-theoretic methods to understand the power and limitations of zero-knowledge proofs. In particular, I have sought to characterize the classes of assertions that can be proven with various types of zero-knowledge proofs, to prove general theorems about zero knowledge, and to minimize or eliminate complexity-theoretic assumptions used in the study of zero knowledge. (Recall from Section 2 that one of the major goals in the foundations of cryptography is to base cryptographic tasks on assumptions that are as weak and general as possible.) In my Ph.D. thesis [Vad1], I carried out such a study for *statistical zero-knowledge proofs*, which are zero-knowledge proofs where all of the security conditions are information-theoretic. That is, even a computationally unbounded prover cannot convince the verifier of a false assertion and even a computationally unbounded verifier cannot learn anything other than the fact that the assertion being proven is true. In a series of papers by myself and others, we gained a very thorough understanding of the class **SZK** of problems possessing statistical zero-knowledge proofs. We characterized the class **SZK** via natural complete problems [SV, GV], proved that **SZK** is closed under various operations (e.g. complementation [Oka]) and proved numerous general theorems about this class (e.g. showing how to effectively eliminate the verifier's ability to gain knowledge by deviating from the protocol [GSV]). Moreover, all of these results were unconditional, in that they did not rely on any unproven complexity assumptions such as the existence of one-way functions. However, while **SZK** turned out to be very interesting and useful from both a complexity-theoretic and cryptographic point of view, it still represents a small subclass of the notions of zero knowledge considered in the literature. In particular, in many cryptographic applications, one can afford to use zero-knowledge proofs whose security conditions are only *computational*, that is only hold with respect to computationally feasible (i.e. polynomial time) adversarial strategies. In my Ph.D. thesis, I suggested that the work on statistical zero knowledge could be a useful stepping stone to understanding zero knowledge in general, but carrying out this idea seemed beyond reach at the time.

After spending my first few years at Harvard focused almost entirely on pseudorandomness, I started to return to zero knowledge, because I saw

an opportunity to extend the study begun in my thesis work far beyond the confines of **SZK**. Indeed, in a series of papers done mostly in collaboration with my Ph.D. students [MV, Vad2, NV, NOV, OV], we have demonstrated that the theory of statistical zero knowledge could be leveraged to understand zero-knowledge proofs in general, even those incorporating computational security conditions. Specifically, we have obtained a characterization of the class **ZK** of problems possessing general, computational zero-knowledge proofs (and even "arguments") in terms of the class **SZK** [Vad2, OV]. This has allowed us to translate most of the general theorems known about **SZK** to all of **ZK**. The resulting theorems are unconditional, in contrast to most previous works on **ZK**, which rely on the assumption that one-way functions exist. As part of this effort, we have obtained a result that is new for **SZK**, too — showing that, for problems in **NP**, we can transform any zero-knowledge proof in which the honest prover strategy is computationally unbounded into one where the prover runs in polynomial time (given an **NP** witness) [NV]. This closes a significant gap between the complexity-theoretic study of zero knowledge (which often allows computationally unbounded prover strategies) and the cryptographic applications of zero knowledge (which require polynomial-time prover strategies). Finally, the methods we developed in [NV] enabled us in [NOV] to resolve a long-standing open problem posed in [NOVY] — that every language in **NP** has a statistical zero-knowledge "argument system" assuming only the existence of one-way functions. (Here we do not expect an unconditional result; indeed, work of Ostrovsky and Wigderson [Ost, OW] shows that one-way functions are essentially the minimal assumption possible.)

5 Other Research Highlights

In addition to my work on pseudorandomness and zero knowledge, two other highlights of my research are the following:

- With Barak et al. [BGI$^+$], I initiated a theoretical study of *program obfuscation*, which aims to make programs unintelligible while preserving their functionality. There are many heuristics for obfuscation and it is widely used in practice, but prior to our work, it had never received a formal cryptographic treatment. Our main result is a negative one — for a natural but weak formulation of the security goal, it is *impossible* to have a general-purpose obfuscator that works for all programs. The impact of this work has been to guide further work on obfuscation (both theory and practice) in more fruitful directions, e.g. by restrict-

ing the class of programs considered or by seeking alternative notions of security.

- With undergraduate Saurabh Sangvhi [SV], I obtained the first non-constant and tight lower bound on the number of rounds needed for two mutually distrusting parties to generate a random n-bit string such that even if one party deviates arbitrarily from the protocol, the outcome will still not be too "biased". Specifically, we show that $\Theta(\log^* n)$ rounds are both necessary and sufficient. Such random selection protocols are a basic building block for solving other tasks in cryptography and distributed computation. Recently, we have obtained similar results for the multiparty random selection protocols where a majority of participants are dishonest [GVZ], and are currently exploring the possibility of obtaining lower bounds for the case of an honest majority, which is a longstanding open problem in the literature on "collective coin-flipping" and "leader election."

6 Future Directions

Ultimately, my goal as a researcher is can be described simply.

I aim to understand the nature of efficient computation.

I am motivated in this effort both by the illuminating lens it provides on the universe in which we live and by its potential impact on technology and society. For me, computational complexity and cryptography provide rich domains in which to pursue this goal, and I expect to continue working in these areas in the forseeable future, while remaining open to new opportunities. Below are just a few examples to illustrate the kinds of problems I expect to pursue in the coming years.

Randomness vs. Space. Despite the apparent usefulness of randomization in algorithm design, the theory of pseudorandomness has provided strong evidence that actually every randomized algorithm can be derandomized with only a small loss in efficiency, in the sense that other widely believed conjectures in complexity theory are known to imply derandomization. However, for the case where efficiency is measured by *space* (i.e. memory), then there is hope for an unconditional proof that randomness saves at most a constant factor ($\mathbf{RL} = \mathbf{L}$). Indeed, there was substantial progress on this problem in the early 90's [Nis, SZ], but then progress stalled

for a decade. In a breakthrough last year, Reingold [Rei] used our earlier work on expander graphs and the zig-zag product [RVW] to fully derandomize the classic example of a space-efficient randomized algorithm, namely the random-walk algorithm for connectivity in undirected graphs. Given this development as well as our overall improved understanding of pseudorandomness, there is hope that we now might be able to derandomize *all* space-bounded algorithms. Indeed, with Reingold and Trevisan [RTV], we have made progress, reducing the gap between Reingold's specific result and the general **RL** vs. **L** problem to just a single "technical condition" in one theorem. While this technical condition may prove to be a major obstacle, I am optimistic that some significant progress can be made on the problem.

Complexity between P and NP. The magnificent theory of **NP**-completeness [Coo, Lev, Kar] has provided a powerful tool for showing that a computational problem is likely to be intractable. Namely, if we can prove a problem to be **NP**-complete, then we know it is computationally equivalent to the thousands of other **NP**-complete problems that have resisted attempts at finding efficient algorithms, and thus is unlikely to have an efficient algorithm itself. However, not all intractable computational problems are **NP**-complete (or even **NP**-hard). Indeed, it is known that there are computational problems in **NP** that are neither **NP**-complete nor solvable efficiently (assuming $\mathbf{P} \neq \mathbf{NP}$) [Lad]. Unfortunately, complexity theory has not made much progress in classifying problems in this region, even though numerous important problems seem to lie here, such as Factoring and Graph Isomorphism. Ideally, we would like a theory of completeness for some class whose complexity seems to be strictly between **P** and **NP**; this would allow us to provide evidence of intractability for problems that are unlikely to be **NP**-complete. Some beautiful progress has recently been made for problems involving equilibria and fixed points, such as Nash Equilibrium, as these have been shown to be complete for a class called **PPAD** [GP, DGP, CD]. However, this class does not seem to be the right one to capture problems such as Graph Isomorphism, Factoring, as well as numerous other problems of cryptographic significance. The class **SZK** of problems having statistical zero-knowledge proofs seems to have potential for playing this role instead. In my Ph.D. work, I gave a start by showing that two natural problems about estimating statistical properties of probability distributions are complete for **SZK** [SV, GV]. Just recently, I have started to develop an approach to showing **SZK**-completeness for other types of problems. Regardless of whether my approach works, under-

standing complexity between **P** and **NP** is an important goal that is likely to remain an interest of mine in the future.

The Assumptions for Cryptography. As discussed earlier, the aspect of cryptography that remains more of an art than a science is the identification of underlying hard problems on which to build cryptosystems. Of course, one way to resolve this is to develop techniques for directly proving lower bounds on the complexity of such problems. While lower bounds are extremely difficult and having ones strong enough for cryptography requires resolving the famous **P** vs. **NP** question, they are at the very core of understanding efficient computation, and thus it is important that researchers continue to work on them (as I might do one day, if I see an opportunity to make a significant contribution). In the meantime, however, we might at least look for ways to relate the assumptions needed for cryptography on better-understood questions in complexity theory. There has been substantial effort to base cryptography on **NP**-completeness, and unfortunately most of the evidence so far has been negative. An alternative, however, is to base cryptography on completeness for some class between **P** and **NP**, as discussed in the previous paragraph. In particular, if an approximate version of the Shortest Vector problem in high-dimensional integer lattices could be shown to be **SZK**-complete, then by work of Ajtai [Ajt], we could base cryptography on the (worst-case) intractability of **SZK**.

Branching Out. While I continue to focus on my core interests in complexity theory and cryptography, I am always looking out for good opportunities to branch out, whether motivated by scientific, philosophical, societal, or technological reasons. For example, through my involvement in the DEAS Center for Research on Computation and Society (CRCS), I am developing an interest in privacy — finding mathematical formulations for what it means for an individual's privacy to be compromised (through use of data such as in medical databases) and designing algorithms that ensure that such compromises cannot occur. I have also am starting to explore the interface of computer science with economics and game theory. This area has significance for complexity theory because it gives rise to computational problems, such as Nash Equilibrium, that lie in the area between **P** and **NP** described above. And it has significance for cryptography because it may be possible to design better protocols if we combine the rational behavior model typically used in economics with the honest/malicious behavior model typically used in cryptography. I have been exploring this possibility

in collaboration with Prof. David Parkes and CRCS postdoc Alon Rosen.

References

[AES] The Advanced Encryption Standard. *Federal Information Processing Standards*, 179, November 2001.

[Ajt] M. Ajtai. Generating hard instances of lattice problems (extended abstract). In *Proceedings of the Twenty-eighth Annual ACM Symposium on the Theory of Computing (Philadelphia, PA, 1996)*, pages 99–108, New York, 1996. ACM.

[ALW] N. Alon, A. Lubotzky, and A. Wigderson. Semi-direct product in groups and zig-zag product in graphs: connections and applications (extended abstract). In *42nd IEEE Symposium on Foundations of Computer Science (Las Vegas, NV, 2001)*, pages 630–637. IEEE Computer Soc., Los Alamitos, CA, 2001.

[ADR] Y. Aumann, Y. Z. Ding, and M. O. Rabin. Everlasting Security in the Bounded Storage Model. *IEEE Transactions on Information Theory*, 48(6):1668–1680, June 2002.

[BGI+] B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S. Vadhan, and K. Yang. On the (Im)possibility of Obfuscating Programs. In J. Kilian, editor, *Advances in Cryptology—CRYPTO '01*, volume 2139 of *Lecture Notes in Computer Science*, pages 1–18. Springer-Verlag, 19–23 August 2001.

[BLV] B. Barak, Y. Lindell, and S. Vadhan. Lower bounds for non-black-box zero knowledge. *Journal of Computer and System Sciences*, 72(2):321–391, March 2006. Special Issue on FOCS '03.

[BOV] B. Barak, S. J. Ong, and S. Vadhan. Derandomization in Cryptography. In D. Boneh, editor, *Advances in Cryptology—CRYPTO '03*, volume 2729 of *Lecture Notes in Computer Science*, pages 299–315. Springer-Verlag, 17–21 August 2003. Full version accepted to *SIAM J. Computing*, pending minor revisions.

[BCJ+] E. Biham, R. Chen, A. Joux, P. Carribault, C. Lemuet, and W. Jalby. Collisions of SHA-0 and Reduced SHA-1. In *EUROCRYPT*, volume 3494 of *Lecture Notes in Computer Science*, pages 36–57. Springer, 2005.

[BM] M. Blum and S. Micali. How to Generate Cryptographically Strong Sequences of Pseudo-Random Bits. *SIAM Journal of Computing*, 13(4):850–864, 1984.

[CRVW] M. Capalbo, O. Reingold, S. Vadhan, and A. Wigderson. Randomness Conductors and Constant-Degree Lossless Expanders. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC '02)*, pages 659–668, Montréal, CA, May 2002. ACM. In joint session with *CCC '02*.

[CD] X. Chen and X. Deng. Settling the Complexity of 2-Player Nash-Equilibrium. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS '06)*, Berkeley, CA, 22–24 October 2006. To appear. Full version posted as *ECCC* TR05-140.

[Coo] S. A. Cook. The Complexity of Theorem-Proving Procedures. In *Conference Record of Third Annual ACM Symposium on Theory of Computing*, pages 151–158, Shaker Heights, Ohio, 3–5 1971 1971.

[DGP] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a Nash equilibrium. In *STOC '06: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 71–78, New York, NY, USA, 2006. ACM Press.

[DH] W. Diffie and M. E. Hellman. New Directions in Cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, 1976.

[Din] I. Dinur. The PCP Theorem via Gap Amplification. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, 2006.

[GP] P. W. Goldberg and C. H. Papadimitriou. Reducibility among equilibrium problems. In *STOC '06: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 61–70, New York, NY, USA, 2006. ACM Press.

[GMW] O. Goldreich, S. Micali, and A. Wigderson. Proofs that yield nothing but their validity, or All languages in NP have zero-knowledge proof systems. *Journal of the Association for Computing Machinery*, 38(3):691–729, 1991.

[GSV] O. Goldreich, A. Sahai, and S. Vadhan. Honest-Verifier Statistical Zero-Knowledge Equals General Statistical Zero-Knowledge. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing (STOC '98)*, pages 399–408, Dallas, TX, May 1998. ACM.

[GV] O. Goldreich and S. Vadhan. Comparing Entropies in Statistical Zero-Knowledge with Applications to the Structure of SZK. In *Proceedings of the Fourteenth Annual IEEE Conference on Computational Complexity (CCC '99)*, pages 54–73, Atlanta, GA, May 1999. IEEE Computer Society Press.

[GMR] S. Goldwasser, S. Micali, and C. Rackoff. The Knowledge Complexity of Interactive Proof Systems. *SIAM Journal on Computing*, 18(1):186–208, February 1989.

[GVZ] R. Gradwohl, S. Vadhan, and D. Zuckerman. Random Selection with an Adversarial Majority. In C. Dwork, editor, *Advances in Cryptology—CRYPTO '06*, Lecture Notes in Computer Science. Springer-Verlag, 20–24 August 2006. To appear.

[HVV] A. Healy, S. Vadhan, and E. Viola. Using nondeterminism to amplify hardness. In *Proceedings of the 36th ACM Symposium on Theory of Computing (STOC '04)*, pages 192–201, Chicago, IL, 13–15 June 2004. ACM. Full version accepted to *SIAM J. Computing* Special Issue on STOC '04.

[KRVZ] J. Kamp, A. Rao, S. Vadhan, and D. Zuckerman. Deterministic Extractors for Small-Space Sources. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC '06)*, pages 691–700, 21–23 May 2006.

[Kar] R. M. Karp. Reducibility among Combinatorial Problems. In J. W. Thatcher and R. E. Miller, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, Inc., 1972.

[Lad] R. E. Ladner. On the structure of polynomial time reducibility. *Journal of the Association for Computing Machinery*, 22:155–171, 1975.

[Lev] L. A. Levin. Universal'nyĭe perebornyĭe zadachi (Universal search problems : in Russian). *Problemy Peredachi Informatsii*, 9(3):265–266, 1973.

[LRVW] C.-J. Lu, O. Reingold, S. Vadhan, and A. Wigderson. Extractors: optimal up to constant factors. In *Proceedings of the 35th ACM Symposium on Theory of Computing (STOC '03)*, pages 602–611. ACM, 2003.

[LW] A. Lubotzky and B. Weiss. Groups and expanders. In *Expanding graphs (Princeton, NJ, 1992)*, pages 95–109. Amer. Math. Soc., Providence, RI, 1993.

[Mau] U. Maurer. Conditionally-Perfect Secrecy and a Provably-Secure Randomized Cipher. *Journal of Cryptology*, 5(1):53–66, 1992.

[MV] D. Micciancio and S. Vadhan. Statistical zero-knowledge proofs with efficient provers: lattice problems and more. In D. Boneh, editor, *Advances in Cryptology—CRYPTO '03*, volume 2729 of *Lecture Notes in Computer Science*, pages 282–298. Springer-Verlag, 17–21 August 2003.

[NOVY] M. Naor, R. Ostrovsky, R. Venkatesan, and M. Yung. Perfect Zero-Knowledge Arguments for NP Using Any One-Way Permutation. *J. Cryptology*, 11(2):87–108, 1998. Preliminary version in *CRYPTO '92*.

[NV] M. Nguyen and S. Vadhan. Zero Knowledge with Efficient Provers. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC '06)*, pages 287–295, 21–23 May 2006.

[NOV] M.-H. Nguyen, S. J. Ong, and S. Vadhan. Statistical Zero-Knowledge Arguments for NP from Any One-Way Function. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS '06)*, Berkeley, CA, 22–24 October 2006. To appear. Full version posted as *ECCC* TR06-075.

[Nis] N. Nisan. Pseudorandom Generators for Space-bounded Computation. *Combinatorica*, 12(4):449–461, 1992.

[Oka] T. Okamoto. On Relationships Between Statistical Zero-Knowledge Proofs. *Journal of Computer and System Sciences*, 60(1):47–108, February 2000.

[OV] S. J. Ong and S. Vadhan. The Complexity of Zero-Knowledge Arguments. In preparation, September 2006.

[Ost] R. Ostrovsky. One-Way Functions, Hard on Average Problems, and Statistical Zero-Knowledge Proofs. In *Proceedings of the Sixth Annual Structure in Complexity Theory Conference*, pages 133–138, Chicago, Illinois, 30 June–3 July 1991.

[OW] R. Ostrovsky and A. Wigderson. One-Way Functions are Essential for Non-Trivial Zero-Knowledge. In *Proceedings of the Second Israel Symposium on Theory of Computing and Systems*, pages 3–17, 1993.

[Rab] M. O. Rabin. Digital signatures and public-key functions as intractable as factorization. Technical Report MIT/LCS/TR-212, MIT Laboratory for Computer Science, January 1979.

[Rei] O. Reingold. Undirected ST-Connectivity in Log-Space. In *Proceedings of the 37th ACM Symposium on Theory of Computing*, pages 376–385, 2005.

[RTV] O. Reingold, L. Trevisan, and S. Vadhan. Pseudorandom Walks in Regular Digraphs and the RL vs. L Problem. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC '06)*, pages 457–466, 21–23 May 2006. Preliminary version on *ECCC*, February 2005.

[RVW] O. Reingold, S. Vadhan, and A. Wigderson. Entropy Waves, the Zig-Zag Graph Product, and New Constant-Degree Expanders. *Annals of Mathematics*, 155(1), January 2001. Extended abstract in *FOCS '00*.

[RSA] R. L. Rivest, A. Shamir, and L. M. Adleman. A Method for Obtaining Digital Signatures and Public-Key Cryptosystems. *Commun. ACM*, 21(2):120–126, 1978.

[SV] A. Sahai and S. Vadhan. A complete problem for statistical zero knowledge. *Journal of the ACM*, 50(2):196–249, March 2003. Extended abstract in *FOCS '97*.

[SZ] M. Saks and S. Zhou. $\mathrm{BP_HSPACE}(S) \subseteq \mathrm{DSPACE}(S^{3/2})$. *Journal of Computer and System Sciences*, 58(2):376–403, 1999. 36th IEEE Symposium on the Foundations of Computer Science (Milwaukee, WI, 1995).

[SV] S. Sanghvi and S. Vadhan. The Round Complexity of Two-Party Random Selection. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC '05)*, pages 338–347, 22–24 May 2005. Invited to *SIAM J. Computing* Special Issue on STOC '05.

[Sha] C. E. Shannon. Communication theory of secrecy systems. *Bell System Technical Journal*, 28:656–715, 1949.

[STV] M. Sudan, L. Trevisan, and S. Vadhan. Pseudorandom Generators without the XOR Lemma. *Journal of Computer and System Sciences*, 62:236–266, 2001. Special issue on CCC '99. Extended abstract in *STOC–CCC '99* joint session.

[Tre] L. Trevisan. Extractors and pseudorandom generators. *Journal of the ACM*, 48(4):860–879 (electronic), 2001.

[TV1] L. Trevisan and S. Vadhan. Extracting Randomness from Samplable Distributions. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science (FOCS '00)*, pages 32–42, Redondo Beach, CA, 17–19 Oct. 2000. IEEE.

[TV2] L. Trevisan and S. Vadhan. Pseudorandomness and Average-Case Complexity via Uniform Reductions. In *Proceedings of the 17th Annual IEEE Conference on Computational Complexity (CCC '02)*, pages 129–138, Montréal, CA, May 2002. IEEE. Full version accepted to *Computational Complexity*, pending minor revisions.

[TVZ] L. Trevisan, S. Vadhan, and D. Zuckerman. Compression of Samplable Sources. *Computational Complexity*, 14(3):186–227, December 2005. Special Issue on CCC '04.

[Vad1] S. P. Vadhan. *A Study of Statistical Zero-Knowledge Proofs*. PhD thesis, Massachusetts Institute of Technology, August 1999. Updated version to be published by Springer-Verlag for winning the *ACM Doctoral Dissertation Award 2000*.

[Vad2] S. P. Vadhan. An Unconditional Study of Computational Zero Knowledge. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS '04)*, pages 176–185, Rome, Italy, 17–19 October 2004. Full version accepted to *SIAM J. Computing* Special Issue on Randomness & Complexity.

[WLF+] X. Wang, X. Lai, D. Feng, H. Chen, and X. Yu. Cryptanalysis of the Hash Functions MD4 and RIPEMD. In *EUROCRYPT*, volume 3494 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2005.

[WY] X. Wang and H. Yu. How to Break MD5 and Other Hash Functions. In *EUROCRYPT*, volume 3494 of *Lecture Notes in Computer Science*, pages 19–35. Springer, 2005.

[Yao] A. C. Yao. Theory and Applications of Trapdoor Functions (Extended Abstract). In *23rd Annual Symposium on Foundations of Computer Science*, pages 80–91, Chicago, Illinois, 3–5 Nov. 1982. IEEE.