Research Statement

Vinod Vaikuntanathan

The main focus of my research is the theoretical foundations of *cryptography* and *distributed protocols*. Thanks to the impressive developments in cryptography over the past three decades, we now have a rich framework to precisely define the security of cryptographic tasks, as well as mechanisms to prove the security of candidate constructions. Barring a handful of exceptions, the security of cryptographic schemes relies on assumptions: either *computational assumptions* such as the hardness of factoring large numbers, or *physical assumptions* such as the existence of private communication channels.

This state of affairs places cryptography under perpetual threat – sometimes, the assumptions we make are rendered false, the most recent and striking example being the breaking of the widely used cryptographic hash functions SHA-0 and SHA-1 [WYY05b, WYY05a]. A major goal in theoretical cryptography is thus to base cryptographic schemes on assumptions that are as weak as possible (or none at all).

> My primary research goal is to build cryptography on a strong foundation
> by eliminating, or minimizing the dependence on unproven assumptions.

In addition, I have an abiding interest in studying and exploiting the interaction between cryptography and the closely allied fields of computational complexity and applied security.

From Average-Case to Worst-Case Assumptions

Traditional complexity theory deals with the worst-case hardness of computational problems (NP-hardness, for instance, provides the guarantee that *some* instance of a problem is hard). Cryptography, on the other hand, requires problems to be hard on the average (for instance, security of an encryption scheme requires hardness of breaking it for a *random* key). Bridging this discrepancy – constructing cryptographic schemes whose security is based on worst-case hardness – is considered to be one of the main open questions in cryptography.

Thanks to the seminal work of Ajtai [Ajt96], cryptography based on *lattices* enjoys the unique and rather enticing feature of being based on a *worst-case* hardness assumption. Unfortunately, despite many advances in the study of lattices, it is not known how to construct sophisticated and extremely useful primitives such as trapdoor functions and identity-based encryption schemes based on the worst-case hardness of lattice problems.

In joint work with Craig Gentry and Chris Peikert [GPV07], I construct algorithms and techniques to exploit natural and innate "trapdoors" for lattices. As a result, we obtain constructions of a variety of trapdoor functions, efficient signature schemes and identity-based encryption schemes based on the worst-case hardness of lattice problems. In subsequent work with Chris Peikert and Brent Waters [PVW07], we utilize these techniques to construct efficient, and maximally secure (universally composable, in technical terms) protocols for oblivious transfer and secure two-party

and multi-party computation. Together, these two works significantly extend the range of cryptography that can be based on worst-case hardness assumptions.

FUTURE OBJECTIVES A natural and intriguing question is whether it is possible to base cryptography on the worst-case hardness of *any* problem in NP (as opposed to a particular problem). Unfortunately, results in this direction have so far been negative. Nevertheless, it seems possible (in fact, quite likely) that cryptography can be based on the (worst-case) hardness of a large class of problems that are intermediate between P and NP-complete (in fact, the problems that underlie lattice-based cryptography are precisely of this kind). I plan to devote a considerable amount of energy to understanding and investigating this possibility. In addition, I plan to continue working on lattice-based cryptography, constructing new primitives and making them efficient and practical.

Removing Physical Assumptions

Physical assumptions, such as the existence of private communication channels, often considerably simplify cryptographic protocol design. Such assumptions, while convenient, are sometimes too strong to be realistic: thus, it is a fundamental question to understand the extent to which they are necessary. My work studies this question in the context of Byzantine Agreement, a central problem in the theory of fault-tolerant distributed systems.

Informally, Byzantine Agreement is the problem of agreeing on a common view of the world starting from differing local views, in the presence of faults that try to thwart this agreement. Aside from the theoretical significance, it is also a valuable tool used within multiparty protocols to implement a reliable broadcast channel when only pairwise channels are available.

It has been known for a long time [FL82] that deterministic algorithms for Byzantine agreement are necessarily inefficient: any such algorithm that tolerates n faulty participants has to run for $\Omega(n)$ rounds of communication. Randomized algorithms for Byzantine agreement, on the other hand, can be dramatically efficient: in fact, Feldman and Micali [FM88] showed how to reach agreement in an expected constant number of rounds! The algorithm of [FM88] assumes *perfect randomness* and perfectly *private channels* between each pair of players.

A private channel is a strong physical assumption: it is impossible to observe or inject messages into the channel, or even detect that a message is being transmitted! The best algorithm that does not rely on this assumption, due to Chor and Coan [CC85], runs in $O(\frac{n}{\log n})$ rounds, not significantly more efficient than the deterministic algorithms. This raises the intriguing question, originally posed in [FM88] and unanswered for two decades, whether private channels are necessary for efficient Byzantine agreement. In a series of works with Michael Ben-Or, Shafi Goldwasser and Elan Pavlov [BPV06, GPV06], I answer this question in the negative by constructing a variety of efficient Byzantine Agreement algorithms that run in $O(\log n)$ rounds without assuming private channels.

Most randomized distributed algorithms assume access to a source of perfectly random – unbiased and independent – bits. In joint work with Madhu Sudan and Shafi Goldwasser [GSV05], I show that this assumption is not necessary: I propose models of imperfect randomness in distributed systems, and show how to modify any distributed protocol (in particular, any Byzantine agreement protocol) into one that works even with "defective randomness".

FUTURE OBJECTIVES A large body of work on distributed algorithms assumes a network that is static and fully-connected. Today's networks, on the other hand, are huge, dynamic and sparsely

connected (think of the internet). I plan to spend a considerable amount of time studying these more demanding and realistic models and designing new distributed algorithms for them.

A key question in the theory of distributed algorithms is the connection between *local* and *global* knowledge (indeed, Byzantine Agreement is an important manifestation of such a connection). Knowledge is a key notion in cryptography too – privacy essentially means not leaking any knowledge about secret data! I plan to investigate the synergistic connections between cryptography and the theory of distributed algorithms, and more specifically, the concepts of knowledge, privacy and reliability.

From Weak Security to Strong Security

Over time, researchers propose new and stronger notions of security for cryptographic schemes to reflect the power of an adversary in the real-world. Yet another direction in minimizing assumptions is to construct cryptographic schemes that satisfy strong notions of security from a scheme that satisfies only a basic definition of security.

Public-key encryption is a case in point, where the definition of security has evolved from the basic to the "super-strong". Classically, encryption is viewed as a mechanism that protects the privacy of data transmitted over an insecure channel, which is captured by the formal notion of semantic security [GM84]. In practice, when encryption is used as part of a complex system, stronger notions of security become indispensable. Consider the setting of an electronic "sealed-bid" auction where the bidders send the encrypted bids to the auctioneer. In this setting, the encryption scheme would certainly be considered insecure if the adversary could consistently bid exactly one more dollar than the previous bidders (without actually knowing what their bids are)! *Non-malleability*, a notion defined by Dolev, Dwork and Naor [DDN00] protects against precisely this class of attacks: informally, an encryption scheme is non-malleable if it is infeasible for an adversary to *maul* the encryption of a message to the encryption of a "related" message. In addition, within a complex system, the adversary might be able to observe the output of the decryption machinery on inputs of her choice – this is called a chosen ciphertext attack (CCA) [RS93, NY90, DDN00]. A significant, long-standing open problem in cryptography is: *Can we achieve these strong notions of security from the weakest one?*

In a series of works with Rafael Pass, Abhi Shelat and others, I make significant progress towards answering this question in the affirmative. Our first result [PSV06] is a construction of a non-malleable encryption scheme from any semantically secure encryption scheme, without making any additional assumptions. In subsequent work [CHH$^+$07], we relax the notion of CCA-security to (what we call) *bounded* CCA-security. Whereas full CCA-security allows the adversary to observe the decryption machinery on an unbounded number of inputs, bounded CCA-security restricts the adversary to an a-priori bounded number of invocations of the decryption device. We then show how to construct a bounded CCA-secure encryption scheme from any semantically secure encryption scheme.

FUTURE OBJECTIVES Indeed, it is known how to achieve full (as opposed to bounded) CCA-security from semantic security (thus, answering the open problem) in an *idealized model* called the *random oracle model*. The random oracle model hypothesizes the existence of an *efficiently computable* and *random* function (a "random oracle") that all parties can access. A cryptographic scheme designed in this idealized model does not, in general, translate to a scheme in the real world, where there is no random oracle. I plan to investigate *sufficient* conditions for a protocol in the "random oracle

world" to be secure also in the "real world". This research direction has already led to a number of significant results, in joint work with Rafael Pass [PV07].

Additional Work

Program Obfuscation. A recent line of research in cryptography aims to understand whether it is possible to *obfuscate* programs, that is, modify the program so that it becomes unintelligible, while its functionality remains unchanged. In joint work with Susan Hohenberger, Guy Rothblum and Abhi Shelat [HRSV07], I construct obfuscators for a complex, practical functionality called re-encryption. I am also interested in the practical implementation of obfuscators: an ongoing work deals with constructing and implementing obfuscators for access-control functionalities such as passwords and biometrics.

The theoretical results on obfuscation have been depressingly negative: general-purpose obfuscation is impossible, and only simple functionalities such as point functions (a.k.a passwords) are known to be obfuscatable. In contrast, there are a number of "obfuscating heuristics" designed by practitioners that "seem to work well". This state of affairs raises a couple of possibilities to obtain positive results for obfuscation, which I plan to investigate in the near future.

1. RELAX THE SECURITY REQUIREMENTS: Obfuscation is hard, perhaps because we demand too much of it in terms of security. One of the problems I plan to investigate is coming up with new definitions of security for obfuscation as well as designing practical schemes that satisfy these (relaxed) definitions.

2. RELAX THE MODEL OF COMPUTATION: I plan to study the possibility of obfuscation in specialized, yet realistic models. One of the possibilities is to design obfuscators using trusted hardware or mechanisms (such as virtual machines) that provide built-in separation of software execution environments. Yet another possibility is to build a quantum obfuscation mechanism.

UNCONDITIONALLY SECURE COMPUTATION. Secure Computation, roughly stated, enables a set of parties with private inputs to jointly and securely compute some function of their inputs. Ben-Or, Goldwasser and Wigderson [BGW88] (BGW), in a seminal work, showed that secure computation can be achieved *without any assumptions* if a majority of parties are honest. A key component in their protocol is a scheme for sharing a secret among n players such that only a threshold number of players can recover the secret. BGW – and indeed, most known multiparty protocols – use secret-sharing schemes that satisfy additional properties (such as verifiability, and multiplicativity), that are constructed based on Reed-Solomon codes.

I show, in joint work with Ronald Cramer, Hao Chen, Shafi Goldwasser and Robbert de Haan [CCG+07], that unconditional secure computation can be based on *any* "good" error-correcting code, thus allowing the use of versatile machinery from coding theory for the design of multiparty protocols. An immediate consequence of our result is the construction of very efficient multiparty protocols based on algebraic-geometry codes.

To summarize, my research goals (at least for the foreseeable future) are to better understand and solve fundamental issues in cryptography, drawing from the arsenal of theoretical computer science and mathematics.

References

[Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *STOC*, pages 99–108, 1996.

[BGW88] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *STOC*, pages 1–10, 1988.

[BPV06] Michael Ben-Or, Elan Pavlov, and Vinod Vaikuntanathan. Byzantine agreement in the full-information model in $O(\log n)$ rounds. In *STOC*, pages 179–186, 2006.

[CC85] Benny Chor and Brian A. Coan. A simple and efficient randomized byzantine agreement algorithm. *IEEE Trans. Software Eng.*, 11(6):531–539, 1985.

[CCG+07] Hao Chen, Ronald Cramer, Shafi Goldwasser, Robbert de Haan, and Vinod Vaikuntanathan. Secure computation from random error correcting codes. In *EUROCRYPT*, pages 291–310, 2007.

[CHH+07] Ronald Cramer, Goichiro Hanaoka, Dennis Hofheinz, Hideki Imai, Eike Kiltz, Rafael Pass, Abhi Shelat, and Vinod Vaikuntanathan. Bounded CCA2-secure encryption. In *ASIACRYPT*, pages 502–518, 2007.

[DDN00] Danny Dolev, Cynthia Dwork, and Moni Naor. Nonmalleable cryptography. *SIAM Journal of Computing*, 30(2):391–437, 2000.

[FL82] Michael J. Fischer and Nancy A. Lynch. A lower bound for the time to assure interactive consistency. *Inf. Process. Lett.*, 14(4):183–186, 1982.

[FM88] Paul Feldman and Silvio Micali. Optimal algorithms for byzantine agreement. In *STOC*, pages 148–161, 1988.

[GM84] Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *J. Comput. Syst. Sci.*, 28(2):270–299, 1984.

[GPV06] Shafi Goldwasser, Elan Pavlov, and Vinod Vaikuntanathan. Fault-tolerant distributed computing in full-information networks. In *FOCS*, pages 15–26, 2006.

[GPV07] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. ECCC TR07-133, `http://eccc.hpi-web.de/eccc-reports/2007/TR07-133/index.html`, December 2007.

[GSV05] Shafi Goldwasser, Madhu Sudan, and Vinod Vaikuntanathan. Distributed computing with imperfect randomness. In *DISC*, pages 288–302, 2005.

[HRSV07] Susan Hohenberger, Guy N. Rothblum, Abhi Shelat, and Vinod Vaikuntanathan. Securely obfuscating re-encryption. In *TCC*, pages 233–252, 2007.

[NY90] M. Naor and M. Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *STOC '90: Proceedings of the twenty-second annual ACM symposium on Theory of computing*, pages 427–437, New York, NY, USA, 1990. ACM Press.

[PSV06] Rafael Pass, Abhi Shelat, and Vinod Vaikuntanathan. Construction of a non-malleable encryption scheme from a any semantically secure one. In *CRYPTO*, pages 271–289, 2006.

[PV07] Rafael Pass and Vinod Vaikuntanathan. New-age cryptography. Manuscript, 2007.

[PVW07] Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. Cryptology ePrint Archive, Report 2007/348, 2007. `http://eprint.iacr.org/2007/348`.

[RS93] Charles Rackoff and Daniel R. Simon. Cryptographic defense against traffic analysis. In *STOC '93: Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, pages 672–681, New York, NY, USA, 1993. ACM Press.

[WYY05a] Xiaoyun Wang, Yiqun Lisa Yin, and Hongbo Yu. Finding collisions in the full SHA-1. In *CRYPTO*, pages 17–36, 2005.

[WYY05b] Xiaoyun Wang, Hongbo Yu, and Yiqun Lisa Yin. Efficient collision search attacks on SHA-0. In *CRYPTO*, pages 1–16, 2005.