RESEARCH STATEMENT
ANIKET KATE

My research interests lie at the intersection of cryptographic research, and systems security and privacy research.

With the rise of personal computers and the Internet in the last three decades, cryptography has received a tremendous amount of attention, which has led to its rapid and extensive development. It is now considered a full-fledged academic subject rather than an applied field in algebra and complexity theory. However, only a small fraction of this extensive cryptographic research is being used in practice. Practitioners and systems researchers prefer to build their systems using the basic encryption and signature schemes, and generally reliable but theoretically unsound security assumptions as most existing elaborate cryptographic protocols are not designed with careful consideration of real-world systems issues and threats. With few exceptions, these systems issues have remained largely unaddressed in the cryptography research community.

My work aims at bridging this gap between cryptographic research, and system security (and privacy) research: Along with producing theoretically elegant cryptographic results, I endeavor to make them useful in real-world scenarios. In the long run, I wish to resolve real-world security, privacy, and robustness issues with ever-growing Internet-based systems by developing advanced-yet-practical cryptographic tools.

My current projects focus on developing cryptographic systems for privacy and decentralized trust. During my PhD at Waterloo and my postdoc at MPI-SWS, I concentrated mainly on bridging the gap between theoretical and practical research on privacy enhancing technologies (PETs) and secure distributed computing. At Saarland University, I have also started to work on minimal trusted hardware assumptions for practical secure systems, and recently on decentralized payment systems. In the following sections, I present a brief overview of these projects.

I. Design and Analysis of Anonymous Communication Networks

Anonymous communication networks (ACNs) address the important privacy threat that arises from linking individuals to their communication over the Internet. In particular, they provide anonymity by decoupling communication and sender/receiver identities, and unlinkability by dissociating two correlated communication patterns. Several ACNs have been proposed in the literature. Among these the onion routing (OR) network Tor has turned out to be a definite success: It protects privacy of millions of users using more than six thousand dedicated proxies all over the world. We have made the following contributions to the design and analysis of Tor as well as ACNs in general.

Efficient OR Constructions. Despite its success, Tor still lacks efficiency and may require a few seconds to answer a simple HTTP request. Tor's interactive circuit construction protocol involving computationally expensive public key cryptographic operations has been one of the factors behind this latency. We observed that identity-based cryptography (IBC), with its scalable key rotation feature, can mitigate the efficiency challenge with the existing Tor protocol. In a series of results, we designed three pairing-based onion routing (PB-OR) circuit construction protocols in an IBC setting [1–3]. In parallel, we found a variant of our PB-OR protocols to be applicable to achieve anonymity in delay tolerant networks (DTNs) [4]. Recently, using multi-exponentiation techniques for the discrete log setting, we proposed a significantly more efficient OR protocol (Ace) [5], which is under consideration to be incorporated in Tor.

Analyzing Tor and Quantifying Anonymity. A comprehensive *security analysis* of the OR protocol has been lacking. This has resulted in a significant gap between research work on OR protocols and existing OR anonymity analyses. In ongoing work, we address these issues with onion routing: As a first step, we define a provably secure OR protocol in the universal composability (UC) framework [6]. While our UC-secure OR protocol is practical for deploying in the next generation Tor network, our OR definition (i.e., ideal functionality) greatly simplifies the process of analyzing OR anonymity metrics. We then define a framework (AnoA) that generalizes the notion of adjacency from differential privacy to anonymity properties [7]; in particular, we formalize sender anonymity, unlinkability and relationship anonymity properties, and using our OR definition, demonstrate that those properties carry over to secure cryptographic instantiations of the Tor protocol. Finally, we extend our formal AnoA framework to develop a client-side real-time light-weight monitor [8] for rigorously assessing the degree of various anonymity properties in practice, and use the monitor to study the temporal evolution of anonymity provided by Tor.

Anonymity with Accountability. Tor lacks an ability to detect or prevent anonymous criminal activities. This has not only hindered many organizations from actively participating in the Tor network, but also led to proxy nodes risking sanctions by law enforcement. A widespread use of any PET necessitates establishing liability under the protection of privacy, and for Tor, this calls for combining the seemingly contradictory goals of *anonymity and accountability*. In this direction, our reactive accountability mechanism, BackRef [9], provides practical repudiation for the OR proxy nodes by tracing back the selected outbound traffic to the predecessor node through a novel verifiable chain of signatures. It also provides a full backward traceability option when all intermediate nodes are cooperating.

The area of ACNs is still full of challenging systems issues. Existing low-latency ACNs are far from being robust against the known system-level attacks, and far from being sufficiently scalable for future requirements. For example, Tor has been found to be highly vulnerable to traffic-analysis attacks, and its anonymity guarantees are not satisfactory for many privacy scenarios. Moreover, the current Tor network design will not remain bandwidth-efficient once its user-base reaches several millions. In the future, I look forward to tackling these systems issues using advanced cryptographic concepts such as *distributed* differential privacy, and private information retrieval over an ACN.

II. Establishing Trust using Secure Distributed Computing

A trusted authority, in some form, is essential for every cryptographic system. However, this requirement always leads to the issue of single point of failure and sometimes to the more undesirable issue of key escrow. Resolving these two issues is of paramount importance when designing secure systems for use over the Internet where denial-of-service attacks, large-scale surveillance, and malicious entities are widespread. Although the concept of secure distributed computing (SDC) has emerged as a natural choice to mitigate these problems, the cryptography literature has failed to provide protocols suitable for use over the Internet. Namely, the aspects related to the practicality of these protocols have been largely ignored, and real-world implementations for most SDC primitives are not yet available. This need for practical SDC protocols motivated most of my PhD thesis work.

Distributed Key Generation—DKG. DKG allows a set of nodes to collectively generate a secret with its shares spread over all the nodes in a malicious fault-tolerant manner and without a trusted entity. It was a prominent example of a SDC primitive that lacked a practical design: all of the known DKG protocols were based on (non-existent) reliable broadcast channels. We observed the *need for Byzantine agreement* for DKG in an asynchronous communication setting, and designed a DKG protocol in a system model that arguably represents the Internet [10]. We implemented and extensively tested our DKG protocol on the PlanetLab platform, and found it to be efficient and reliable for use over the Internet [11]. Furthermore, we demonstrated its utility by developing the following two illustrative applications:

(a) Application to Identity-Based Encryption. The key escrow and single-point-of-failure properties of identity-based encryption (IBE) necessitate the use of a distributed private-key generator (distributed PKG) to make IBE suitable for systems outside closed organizational settings. As the first application, we formalized distributed PKG setup and private key extraction primitives for IBE, and designed provably secure distributed PKG constructions for three prominent IBE schemes [12].

(b) Application to Distributed Hash Tables. For quorum-based distributed hash tables (DHTs), there exist several analytical results that can tolerate malicious faults. Unfortunately, these results avoided the use of SDC, and incur significant communication cost in order to achieve message routing. As our second application, we obtained two robust communication protocols for DHTs by adding threshold signatures on top of our DKG system [13, 14]. Both of these protocols asymptotically reduce the communication costs of previous solutions against a computationally bounded Byzantine adversary, and importantly, without using a trusted third party. Our microbenchmarks have shown that the protocols are practical for deployment under significant levels of churn and adversarial behavior.

Verifiable Secret Sharing—VSS. VSS is an important primitive in SDC systems such as DKG and multi-party computation (MPC). It allows an untrusted dealer to share a secret among a set of nodes in the presence of a malicious adversary controlling some of those nodes. We have made a few theoretical and practical contributions to VSS in the computational complexity setting [15–17].

Cryptographic commitments to polynomial coefficients/evaluations are used to achieve verifiability for VSS in the computational complexity setting. We defined the concept of *polynomial commitments*, and presented a constant-size polynomial commitment scheme [16] that significantly reduces communication complexity for computational VSS [17] and several cloud computing applications. The three-round Pedersen synchronous VSS scheme has been the norm for distributed cryptography for the last twenty years. In another effort, we showed that *two rounds* are actually necessary and sufficient for synchronous VSS, and defined a new two-round VSS scheme [15]. Our contributions to VSS not only complement our DKG system but also are of interest to MPC. In particular, we combined our two-round VSS with floating-point MPC primitives to define PrivaDA [18], a distributed noise generation framework for differential privacy (DP) mechanisms. This allows PrivaDA to achieve the best possible tradeoff between privacy and utility for DP mechanisms in the truly distributed data aggregation setting.

The field of secure distributed computing has much larger application potential. In upcoming projects, I will apply the SDC primitives to the emerging areas of privacy-preserving biometric authentication and genome privacy.

III. Trusted Hardware-Based Security and Privacy

Trusted hardware modules are becoming prevalent in computing devices of all kinds. A broad trusted hardware assumption purports to solve almost all security problems in a trivial and uninteresting manner. However, relying entirely on hardware assumptions to achieve security goals of a system (a) can be impractical given the limited memory, bandwidth and CPU capabilities of available hardware modules, and (b) makes the designed system vulnerable to even a tiny overlooked or undiscovered flaw/side-channel in the employed module. Thus, the key challenge to me while designing a trusted hardware-based system is to determine a *minimal hardware assumption* required to achieve the system's goals, and justify the assumption for an available hardware module.

Privacy Preserving Online Behavioral Advertising. Online behavioral advertising (OBA) involves the tracking of users' activities in order to deliver tailored advertisements. It is typically conducted by third-party data analytics firms (brokers). They track user behaviors across web sessions and maintain user profiles on their ends, which raises significant privacy concerns among users and privacy advocates alike.

Although the standard private information retrieval (PIR) for keywords primitive appears to solve this problem in a straightforward way, existing PIR protocols are far from being efficient in practice. We observe that (server-side) hardware module-aided oblivious RAM (ORAM) can instead solve the problem in a realistic manner: server-side trusted hardware runs user queries on server-side encrypted data in a data-insensitive fashion. Efficiency comes from the fact that ORAM algorithms incur only a poly-logarithmic (in data size) computation overhead rather than a (nearly) linear overhead required by PIR protocols. Based on this observation, we proposed ObliviAd [19], a provably secure architecture for privacy preserving OBA, which provides brokers an economical alternative that preserves the privacy of users without hampering the precision of ad selection using a trusted hardware-based oblivious RAM.

In the near future, I plan to study other data-oblivious algorithms and use them to solve novel privacy problems. For example, we are now designing an oblivious algorithm for landmark routing in a credit network graph, and executing it on a server-side hardware module for performing credit network transactions in a privacy preserving manner [20].

Non-equivocation in SDC. A higher replication factor has been an important reason behind practitioners' lack of enthusiasm towards Byzantine fault tolerant distributed computing (DC) systems. In recent years, there have been a few proposals (e.g., Trusted Incrementer—TrInc) to add a small amount of trusted hardware at each node in a Byzantine fault tolerant DC system to cut back the replication factor. These trusted components eliminate the ability of a node to perform *equivocation*, which intuitively means sending conflicting messages to different nodes. To study the power of non-equivocation in the context of SDC, we first defined *non-equivocation* formally. This led us to find that in contrast to previous perceptions non-equivocation alone does not allow for reducing the replication factor in DC primitives [21]; however, it is possible to use non-equivocation along with signatures to transform any crash tolerant protocol into a protocol that tolerates Byzantine faults, without requiring an increase in the replication factor.

We then observed that both our transformation [21] and previous work consider only the basic DC properties of safety and liveness, and they do not capture the secrecy/privacy property required in SDC primitives VSS or MPC. We filled this gap by composing non-equivocation with public-key encryptions and zero-knowledge proofs, which has resulted in novel VSS and MPC protocols with a reduced replication factor [22].

IV. Emerging Research Goal: Decentralized Payment Systems

Over the last five years we have been observing a rather unexpected growth of crypto-currencies such as Bitcoin and payment networks such as Ripple. Their decentralized and pseudonymous nature, ability to perform transactions across the globe in a matter of seconds, and potential to monetize everything regardless of jurisdiction have been pivotal to their success so far. Despite some major hiccups, their market capitalizations are increasing steadily, and many now believe that the concepts of decentralized payment systems and the internet of value are here to stay.

Nevertheless, the current decentralized payment systems are far from perfect. Along with several apparent regulatory and economic issues, they are facing major challenges in the form of transaction privacy, double spending, and mining vulnerabilities. In the long term, along with continuing my research on ACNs, SDC and trusted hardware, I am looking forward to tackling the security, privacy, and economic challenges with the decentralized payment systems. As a first step in this direction, we have recently introduced CoinShuffle [23] a decentralized Bitcoin mixing protocol that allows users to utilize Bitcoin in a truly anonymous manner by mixing their coins with those of others. Unlike its predecessors (i.e., Zerocoin and Zerocash), CoinShuffle is perfectly compatible with the current Bitcoin system, and a few Bitcoin wallet developers are working towards incorporating CoinShuffle in their client software.

Decentralized payment systems are also evolving as the internet of value: by disincentivizing adversarial threats economically, they are found to be useful towards resolving fairness and robustness issues with some distributed systems. For example, Bitcoin has been used towards defining fair MPC protocols, or payment networks (in the form of credit networks) have been employed towards designing sybil-tolerant systems. The concept of the internet of value can improve a significantly larger variety of distributed protocols, and I plan to explore this potential in the future.

As these payment systems grow further, tackling the regulatory and economic issues will become critical: Use of the Bitcoin network for money laundering and illegal trade is already becoming a global issue. This calls for interdisciplinary research between security and legal studies, and working on this interdisciplinary topic will be an excellent chance for me to evolve as a practical security researcher. All in all, this is an exciting time to work on the distinctive challenges provided by the decentralized payment systems, and I will like to take on this unique opportunity.

References

[1] **A. Kate**, G. M. Zaverucha, and I. Goldberg. Pairing-Based Onion Routing. In *PETS '07*, pages 95–112, 2007.

[2] **A. Kate**, G. M. Zaverucha, and I. Goldberg. Pairing-Based Onion Routing with Improved Forward Secrecy. *ACM Transactions on Information and System Security (TISSec)*, 13(4), 2010.

[3] **A. Kate** and I. Goldberg. Using Sphinx to Improve Onion Routing Circuit Construction. In *Financial Cryptography (FC) '10)*, pages 359–366, 2010.

[4] **A. Kate**, G. M. Zaverucha, and U. Hengartner. Anonymity and Security in Delay Tolerant Networks. In *SecureComm '07*, pages 504–513, 2007.

[5] M Backes, **A. Kate**, and E. Mohammadi. Ace: an efficient key-exchange protocol for onion routing. In *ACM WPES '12*, pages 55–64, 2012.

[6] M. Backes, I. Goldberg, **A. Kate**, and E. Mohammadi. Provably Secure and Practical Onion Routing. In *IEEE CSF '12*, pages 369–385, 2012.

[7] M. Backes, **A. Kate**, P. Manoharan, S. Meiser, and E. Mohammadi. AnoA: A Framework for Analyzing Anonymous Communication Protocols. In *IEEE CSF '13*, pages 163–178, 2013.

[8] M. Backes, **A. Kate**, S. Meiser, and E. Mohammadi. (Nothing else) MATor(s): Monitoring the Anonymity of Tor's Path Selection. In *ACM CCS '14*, pages 513–524, 2014.

[9] M. Backes, J. Clark, **A. Kate**, M. Simeonovski, and P. Druschel. BackRef: Accountability in Anonymous Communication Networks. In *ACNS '14*, pages 380–400, 2014.

[10] **A. Kate** and I. Goldberg. Distributed Key Generation for the Internet. In *IEEE ICDCS '09*, pages 119–128, 2009.

[11] **A. Kate**, Y. Huang, and I. Goldberg. Distributed Key Generation in the Wild. *IACR Cryptology ePrint Archive*, 2012:377, 2012. Project Webpage: `http://crysp.uwaterloo.ca/software/DKG`.

[12] **A. Kate** and I. Goldberg. Distributed Private-Key Generators for Identity-Based Cryptography. In *SCN '10*, pages 436–453, 2010.

[13] M. Young, **A. Kate**, I. Goldberg, and M. Karsten. Practical Robust Communication in DHTs Tolerating a Byzantine Adversary. In *IEEE ICDCS '10*, pages 263–272, 2010.

[14] M. Young, **A. Kate**, I. Goldberg, and M. Karsten. Towards Practical Communication in Byzantine-Resistant DHTs. *IEEE/ACM Trans. Netw.*, 21(1):190–203, 2013.

[15] M. Backes, **A. Kate**, and A. Patra. Computational Verifiable Secret Sharing Revisited. In *Advances in Cryptology - ASIACRYPT '11*, pages 590–609, 2011.

[16] **A. Kate**, G. M. Zaverucha, and I. Goldberg. Constant-Size Commitments to Polynomials and Their Applications. In *Advances in Cryptology - ASIACRYPT '10*, pages 177–194, 2010.

[17] M. Backes, A. Datta, and **A. Kate**. Asynchronous Computational VSS with Reduced Communication Complexity. In *CT-RSA '13*, pages 259–276, 2013.

[18] F. Eigner, **A. Kate**, M. Maffei, F. Pampaloni, and I. Pryvalov. Privacy-preserving Data Aggregation with Optimal Utility. In *ACSAC '14*, pages 316–325, 2014.

[19] M. Backes, **A. Kate**, M. Maffei, and K. Pecina. ObliviAd: Provably Secure and Practical Online Behavioral Advertising. In *IEEE Symposium on Security and Privacy (Oakland '12)*, pages 257–271, 2012.

[20] P. Moreno-Sanchez, **A. Kate**, M. Maffei, and K. Pecina. Privacy Preserving Payments in Credit Networks. In *NDSS*, 2015. To appear.

[21] A. Clement, F. Junqueira, **A. Kate**, and R. Rodrigues. On the (limited) power of non-equivocation. In *ACM PODC '12*, pages 301–308, 2012.

[22] M. Backes, F. Bendun, A. Choudhury, and **A. Kate**. Asynchronous MPC with a strict honest majority using non-equivocation. In *ACM PODC '14*, pages 10–19, 2014.

[23] T. Ruffing, P. Moreno-Sanchez, and **A. Kate**. CoinShuffle: Practical Decentralized Coin Mixing for Bitcoin. In *ESORICS '14*, pages 345–364, 2014.