Research statement

Prashant Agrawal
Department of Computer Science and Engineering, IIT Delhi
Centre for Digitalisation, AI and Society, Ashoka University
https://sites.google.com/view/agrawalprash
prashant@cse.iitd.ac.in

My research interests lie at the intersection of computers and society. During my PhD research, I landed on a rich problem space at this intersection — secure electronic voting — which is replete with many adversarial entities, conflicting design requirements and creative approaches to navigate these conflicts. Going forward, I wish to continue exploring similar problem areas that enhance social good by securing our increasingly digital social processes.

1 PhD research highlights

My PhD research focused on enhancing public trust in large public elections like the Indian elections. Towards this goal, I pursued the following main research directions:

Bringing transparency and recoverability to verifiable voting systems[1]. Existing voting protocols provide public verifiability via two main approaches — cryptographic voting and risk-limiting audits — but both fail to provide the necessary transparency and robustness properties required for large public elections. Cryptographic voting, though mathematically sound, is complex for voters and administrators and vulnerable to social engineering. Risk-limiting audits depend on a reliable paper audit trail, which is problematic given the history of ballot-box tampering and booth capture in Indian elections. A hybrid dual voting system, combining the benefits of cryptographic and paper-based voting, offers a promising alternative, but it poses the tricky dilemma of how to resolve discrepancies between paper and electronic tallies? In large elections like the Indian elections, such discrepancies seem inevitable given the wide range of potential adversaries at various polling booths. Our approach solves this dilemma by introducing "recoverability" in dual voting systems: identifying individual vote-level mismatches causing the tally-level disputes, resolving them if possible, and identifying the polling booths causing them otherwise. Thus, our approach enables localised recovery by potentially rerunning the election only in the compromised booths. The recovery is done provably and in zero-knowledge, hiding all individual voters' votes and booth-level voting patterns. The overall approach considerably improves the transparency and robustness of large elections without compromising vote secrecy.

Traceable mixnets[2]. A key technical component for our recovery protocol is a cryptographic primitive I introduced called "traceable mixnets." A traceable mixnet allows controlled reveal of partial information about the input ciphertexts and output plaintexts of a mixnet, without leaking any additional information. In essence, a traceable mixnet is a distributed zero-knowledge

[1] Published in E-VOTE-ID 2023. Link: https://bit.ly/4bwvMNL
[2] Published in PETS 2024. Link: https://bit.ly/3ROxpLv

proof where the mix-servers act as the provers and jointly answer set-membership queries — membership of a ciphertext's decryption in a set of plaintexts and of a plaintext's encryption in a set of ciphertexts — to a querier, without learning the answer themselves. The primitive is useful not only in recoverable voting, but also as a general tool in applications such as privacy-preserving public health analytics. I also developed a comprehensive implementation[3] of our traceable mixnet construction, which showed that our approach is orders of magnitude faster than existing generic techniques for batched execution of such distributed membership proofs.

Publicly auditable yet privacy-preserving electoral rolls[4]. For the final facet of my thesis, I turned to a relatively unexplored but critical area of election security: the manipulation and privacy threats with electoral rolls (voter lists) and polling-booth eligibility verification processes. Common threats include ineligible voters' registration, duplicate registration, eligible voters' denial, ballot stuffing, and voters' privacy loss from electoral rolls. Existing deployments commonly make the entire electoral roll public, which leads to various voter profiling and microtargeting attacks by political parties. Alternatively, eligibility credentials are issued to voters, but these methods lack public verifiability: there is no way to confirm that only eligible voters received credentials and no eligible voters were excluded; also, these solutions suffer from various disenfranchisement, coercion and impersonation issues. I proposed a protocol that statistically audits that no large-scale spurious vote-casting or voter denial happened, while protecting against the profiling attacks and without requiring voters to carry any special eligibility credentials, smart cards, secret keys, etc. The solution acts as the first to achieve the dual goals of public verifiability and privacy of electoral rolls in large public elections.

2 Other projects

Below I outline some additional projects I explored outside the scope of my thesis problem.

An operational architecture for privacy[5]. I have worked on developing a general operational architecture for balancing privacy and utility in large public-service applications. The architecture comprises a language-based framework to express an ideal functionality for a proposed application, enabling abstract, design-level analysis of its privacy risk impact. It also offers a rigorous definition of *purpose limitation* in terms of protocols emulating this ideal functionality.

Security and privacy analysis of India's Unified Payments Interface (UPI). Although UPI has been immensely successful in India, anecdotal instances of "payments getting stuck" abound, leading to the payer and the payee in a deadlock situation and confused whether the payment would eventually succeed or not. There are privacy threats too, mainly because of a permanent identifier used in all the payments. We are modelling these threats into a formal security definition and analysing the current UPI deployment with respect to the proposed definitions. The exercise hopes to inform the design of next-generation secure UPI deployments.

A lightweight anonymous proof of account ownership[6]. An anonymous proof of account ownership (PAO) is a variant of anonymous credentials that allows a prover to anonymously prove to a verifier that they hold an account with an account provider. This is done in a way that utilises existing network infrastructure and avoids any change from the account provider.

[3] Link: https://bit.ly/3UYr18s
[4] Published in CSF 2024 (to appear). Personal link: https://bit.ly/3WRDzBe
[5] Links to early design documents: https://bit.ly/3Xek7yX and https://bit.ly/3VenhAZ
[6] Link to an initial preprint version: https://bit.ly/3yySIxf

In fact, the provider remains entirely unaware of the proof being conducted. This characteristic makes these proofs easily adoptable, especially in applications where the account provider's business interests do not align with users' privacy (e.g., whistleblowing). We proposed a novel lightweight anonymous PAO that is considerably simpler and more general than existing solutions based on multiparty computation.

Traceable mixnets for anonymous identity systems. I am also working on a novel anonymous identity system based on the traceable mixnet paradigm. Current anonymous credential systems struggle with deployment as anonymous national identity systems because of conflicts with traceability, duplicate identities, and recovery of lost credentials. A traceable mixnet based approach, where users' identities in two different contexts are modelled as input and output lists of a mixnet and anonymous credentials are modelled as zero-knowledge set membership proofs about these identities, can address many of these issues.

3 Future directions

Looking ahead, I aim to broaden my research scope and tackle other impactful social problems in our digital age. Below, I highlight a few specific problems that have piqued my interest. I am keen to learn any new techniques required to solve these or other similarly challenging interdisciplinary problems.

Secure digital identity systems. Secure digital identity systems must address multiple challenges: bootstrapping of identities, maintaining interoperability across different domains while preventing their cross-linking, deduplication, and prevention of exclusion and recovery from lost or stolen identities. While existing cryptographic solutions address some of these issues, no holistic solution exists and deployments often run into ethical and legal challenges. Can we develop comprehensive and inclusive solutions for next-generation national digital identity systems?

Countering deepfakes and disinformation. In the age of generative AI, distinguishing real from artificial content has become increasingly difficult. Can we leverage cryptographic techniques to distinguish photographs or videos captured by real cameras, perhaps using signature schemes that are malleable enough to authenticate not only original images but also common image transformations and filters? I am also interested in combating fake news on social media.

Integrating economic and cryptographic arguments. Cryptography is fundamentally about designing protocols to achieve desirable social outcomes and thus shares its goals with mechanism design in economic theory. Can we integrate economic and cryptographic arguments to arrive at clean solutions to oft-corruptible social processes like political financing?

Zero-knowledge protocols under new constraints. I am eager to apply my experience in designing and implementing zero-knowledge proof systems in new and innovative settings. For example, I can imagine interesting applications of human-verifiable or human-provable zero-knowledge proofs in voting. Also, just how anonymous PAOs enable practical deployment of anonymous credentials by avoiding changes to the existing network infrastructure, can we create network-variants of exotic zero-knowledge proofs and signature schemes known in cryptography?