Research Statement

Hou Ruomu

Nov 2024

My research focuses on designing **robust distributed systems** in **adversarial settings**. Such systems play a critical role in enabling secure and reliable computing in an increasingly interconnected and decentralized world. For example, cryptocurrencies have been securely processing billions of dollars of transactions daily.

The development of adversarial-resilient distributed systems exhibits a dynamic interplay with real-world applications—a continuous push-and-pull relationship. On one hand, advancements in resilience mechanisms drive new application scenarios. For instance, progress in Sybil defence has made **permissionless systems** like blockchains feasible, coordinating hundreds of thousands of mutually distrusting nodes globally to deliver reliable cryptocurrency services for high-stakes transactions. On the other hand, the increasing **scale, complexity, and mission-critical nature** of emerging systems—such as federated learning platforms, smart grids, and drone networks—demands ever more robust and efficient protocols.

My research leverages this dynamic relationship between theoretical discoveries and practical needs, treating it as a source of inspiration for solving fundamental challenges. As systems scale, they inevitably become distributed. However, the heterogeneity of **trust models, network assumptions, and application priorities** makes a one-size-fits-all solution impractical.

Tapping into this observation, my work has the following focuses:

1. **Enabling key capabilities**: I design distributed systems with novel features that make entirely new applications feasible. For example, we proposed the first practical blockchain system that could tolerate more than 50% failure, making secure federated computing with dominant players possible.
2. **Systems for the future**: In our research, we build future-proof systems. Many new systems are fantastic in enabling the applications that suits today's needs, but have obstacles in sustainability, scale or performance that limits the long-term suitability in application. We propose principle-driven, and application-driven system designs that overcome such obstacles.

By balancing foundational innovations with practical considerations, my research bridges the gap between theory and deployment, paving the way for resilient distributed systems that meet the demands of evolving real-world applications.

1. Create Systems with new Capability

Blockchain systems are designed to tolerate Byzantine failures, where some nodes may behave maliciously, collude, or disseminate false information to disrupt protocol integrity. This failure model is particularly relevant to real-world deployments of decentralized systems, where trust among participants cannot be assumed.

However, most blockchain systems tolerate a minority of faulty nodes, largely due to fundamental limits associated with Byzantine agreement—a commonly used building block in blockchains. The limit on tolerance makes certain important application scenario impossible. One example is federated computing where a dominant player who controls a majority of resources work with a bunch of smaller players. Such scenario is often seen among cooperates where a large company form a logistics/settlement system with many supplier/clients. Systems with minority tolerance, if deployed in such scenario, offers guarantees that is no more than a centralized system controlled by the dominant party.

My research addresses this challenge by proposing the first **practical blockchain systems with majority fault tolerance**. In my works [1], I leverage Byzantine Broadcast (BB), a distributed computing primitive that is not subject to the same theoretical limits. However, existing BB protocols suffer from low throughput, making them impractical for blockchain applications. I designed a novel BB protocol incorporating **message pipelining**, **multi-hop propagation**, and **look-forward security**, significantly improving throughput. Using this protocol, I developed blockchain systems capable of tolerating faults in majority of nodes, providing stronger guarantees for high-stake applications.

We further designed a general optimistic execution framework [2] for majority tolerance. The capability of being able to execute optimistically is powerful in distributed systems with minority tolerance, for it greatly improves performance when the actual number of faults is relatively small. It turns out that it is far from trivial to extend such capability to majority tolerance. Through usage of carefully designed certificates and quorums, we coupled traditional byzantine agreement protocol with byzantine broadcast, and created the first system that supports optimistic execution while tolerating majority failure. It improved latency performance by about 30x over the state-of-the-arts.

2. Create Systems for the future

One important capability that enables the large-scale deployment of blockchain systems is the permissionless participation. Since Nakamoto consensus, many permissionless systems adopt the paradigm of proof-of-work (PoW) mining, which limit the power of the adversary to its computational power.

Such capability does not come for free. One of the major criticisms for all proof-of-work systems is the monstrous consumption of electricity. For example, Bitcoin's energy consumption rivals that of a medium-sized country, and is only growing. Such trend is not sustainable. Hence, we need better ways of defending permissionless systems, which should be as secure as proof-of-work, but at the same time less energy hungry.

Our research [3] investigated such sustainable sybil defences for building future-proof distributed systems. We observed that existing high-performance blockchain systems necessarily consume certain resources, for example, bandwidth. We successfully used this existing resource consumption to do sybil defence through proof-of-bandwidth. By reshaping protocol communication patterns to enable secure and efficient bandwidth accounting, we developed a proof-of-bandwidth blockchain Selfied [3]. Different from existing systems, Selfied achieves practical and provable sybil defence guarantees and **does not consume any external resources**. Our contribution makes permissionless systems more sustainable and removed a major obstacle to the long-term applicability of such systems.

3. Other Works

Combine Resources for Sybil Defence

Although alternative proof-of-X systems have gained traction to replace proof-of-work, many existing systems still rely on proof-of-work for its proved safety and maturity. To help with the transition, we also investigate combining multiple resource types for Sybil defense. In [4],we proposed a novel hybrid protocol that integrates PoW with proof-of-stake to harness the strengths of both mechanisms, achieving better security and reduced waste. The main challenge in designing the protocol is that PoW solutions can be selectively revealed, while stake distribution is a public knowledge. By combining two resources properly, my research does not only achieve a better security guarantee than any single-resource protocol, but also open up a way to transitions between different resource models, ensuring long-term sustainability.

Simple Blockchain Protocol Scaling

Scaling up the throughput of distributed system is always good. However, many designs use complicated protocols for such task, which not only hinders understanding but also make correct and efficient implementation harder. In [5], we proposed an extremely simple parallel-chain construction, where multiple instances of Nakamoto chains operate concurrently, forming a directed acyclic graph (DAG) of blocks. By designing a rigorous yet simple ordering mechanism to coordinate blocks across chains, the system is very simple while achieving over 100x throughput improvement over the plain Nakamoto consensus. We believe a simple scaling

construction helps not only in the immediate application, but also serve as a pedagogical inspiration to future protocols.

Broader Research Interest in Distributed Algorithms

I also enjoy working on problems in distributed algorithms. In [4], we gave a generalized solution to the feasibility of combing several resources in sybil defence, and gave the first matching upper bound to previous lower bound result. We also proposed new randomized algorithms that gives exponentially better asymptotic improvement than previous works in dynamic networks [6] and distributed view reconciliation problems [7].

4. Future Research

Vision 1 - Expanding Sybil Defence Mechanisms

We have explored Sybil defence using computational power [5], stake [1] [2], and bandwidth [3]. Each resource type has unique security properties and deployment advantages. For example, proof-of-stake relies on publicly verifiable stake distribution, while proof-of-bandwidth depends on localized observations. I plan to further investigate **resource piggybacking**, particularly in areas like **storage**, which is already a significant overhead for blockchain systems. This would open up new capabilities that enables storage-focused applications. A key challenge in using storage for Sybil defence lies in mitigating compression and deduplication, which could allow malicious nodes to bypass storage requirements. My work will focus on principled approaches to address this challenge and design systems with minimal overhead.

Vision 2 - Blockchain Resilience to Network Partition:

A key motivation for decentralized systems is their ability to reduce reliance on trusted entities, such as governments or centralized institutions. However, today's distributed systems are vulnerable to network partitions, which can be instigated by state-level adversaries. A partition-resistant system is thus useful for applications such as central-bank digital currency, which should be reliable even when global interconnectivity is disturbed. Existing research shows that a general, strong sense of security is impossible under catastrophic network partition. Nevertheless, not all network partitions are unpredictable and happen in the worst case, and sometimes a weaker sense of security is sufficient for certain application scenarios. Hence, my work seeks to provide new capabilities in this context:

- **Partition-Resilient Topologies**: Using blockchain state information (e.g., stake distribution), we demonstrated in [8] that it is possible to construct efficient connection topologies that is harder to partition. In the next, I want to

investigate whether other information (e.g. AS ownership) to create topologies that offer better guarantees.

- **Committee-Based Resilience**: I am investigating systems where multiple committees handle different decisions. If one committee remains connected during a partition, it can continue to process decisions securely.
- **Partition Recovery Mechanisms**: Developing predictable recovery processes for short- or long-term partitions is essential for building user confidence in blockchain adoption.

References

[1] **R. Hou**, H. Yu and P. Saxena, "Using Throughput-Centric Byzantine Broadcast to Tolerate Malicious Majority in Blockchains," in *Proceedings of the IEEE Symposium on Security and Privacy (SP)*, 2022.

[2] **R. Hou** and H. Yu, "Optimistic Fast Confirmation While Tolerating Malicious Majority in Blockchains," in *Proceedings of the IEEE Symposium on Security and Privacy (SP)*, 2023.

[3] **R. Hou**, H. Yu and Y. Sun, "Selfied: Sybil defense in permissionless blockchains via in-protocol bandwidth consumption," *Computer Networks*, vol. 256, 2024.

[4] Y. Sun, **R. Hou** and H. Yu, "Using Multi-dimensional Quorums for Optimal Resilience in Multi-resource Blockchains," in *Proceedings of the 28th IEEE Pacific Rim International Symposium on Dependable Computing (PRDC)*, 2023. **Winner of 2023 IEEE Distinguished Paper Award on Dependable Computing.**

[5] H. Yu, I. Nikolic, **R. Hou** and P. Saxena, "OHIE: Blockchain Scaling Made Simple," in *Proceedings of the IEEE Symposium on Security and Privacy (SP)*, 2020.

[6] **R. Hou**, I. Jahja, Y. Sun, J. Wu and H. Yu, "Achieving Sublinear Complexity under Constant T in T-interval Dynamic Networks," in *Proceedings of the 34th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 2022.

[7] **R. Hou**, I. Jahja, L. Luu, P. Saxena and H. Yu, "Randomized View Reconciliation in Permissionless Distributed Systems," *IEEE/ACM Transactions on Networking*, vol. 28, 2020.

[8] Y. Sun, **R. Hou** and H. Yu, "Robust and Low-degree Overlay for Secure Flooding Against Resource-bounded Adversaries," in *IEEE Pacific Rim International*

Symposium on Dependable Computing (PRDC), 2024. **Winner of Best Paper Award**.