

# SONGZE LI

The Hong Kong University of Science and Technology

Tower 13, Flat 3B, Senior Staff Quarters

Web: <https://songzli.github.io/>

Google Scholar: <https://scholar.google.com/citations?hl=en&user=zX8LAXkAAAAJ>

Github: <https://github.com/songzLi>

e-mail: [songzeli8824@gmail.com](mailto:songzeli8824@gmail.com)

cell phone: +852 97991376

---

## RESEARCH INTERESTS

- Algorithms and systems for distributed/federated data analytics
  - Techniques to improve resiliency, security and privacy, and computation efficiency of distributed machine/deep learning tasks
  - Attacks in federated learning systems
  - Techniques for gradient and model compression/quantization in distributed deep learning
- Scalable blockchain systems
  - Protocol designs for high-throughput and low-latency blockchain systems for permissioned and permissionless settings
  - Vertical and horizontal scaling of blockchain performance with compute and network resources
  - Efficient protocol designs for data availability proof

## EDUCATION

**University of Southern California (USC)**, Los Angeles, California, USA

- PhD, Electrical Engineering, Aug. 2011 – Aug. 2018
- Advisor: Prof. Salman Avestimehr
- GPA: 3.9/4

**University of Southern California (USC)**, Los Angeles, California, USA

- MSc, Electrical Engineering, Aug. 2011 – May 2013
- Advisor: Prof. Urbashi Mitra
- GPA: 3.9/4

**Polytechnic Institute of New York University (NYU-Poly)**, Brooklyn, New York

- BSc, Electrical Engineering, Jan. 2008 – May 2011
- Advisor: Prof. Elza Erkip
- GPA: 3.91/4

## APPOINTMENTS

- **Affiliate Assistant Professor, CSE Department, School of Engineering, HKUST** (Aug. 2021 – present).
- **Visiting Assistant Professor, ECE Department, School of Engineering, HKUST** (Dec. 2020 – Nov. 2021).
- **Assistant Professor, IoT Thrust, Information Hub, HKUST** (Dec. 2020 – present).
  - Secure and scalable federated learning.

- Blockchain-based secure data market.
- **Research scientist at Stanford University** (Jan. 2020 – Aug. 2020), Hosting professor: Prof. David Tse.
  - Manage a multi-university effort to develop a smart contract public blockchain testnet based on the **Prism** consensus protocol.
  - Design and implement database and state recovery mechanism when new nodes join the network.
  - Design and implement peer discovery mechanism.
- **Research scientist and blockchain developer at Trifecta Blockchain Inc.** (July 2019 – Dec. 2019).
  - Designed and developed communication-efficient mechanism to verify data availability for **Trifecta** sharding protocol.
  - Invented **Coded Merkle Tree**, which is a novel cryptographic accumulator that encodes a block using LDPC codes. The commitment of a **Coded Merkle Tree** allows a light client to verify data availability of a block with *constant* sampling cost, and to detect incorrectly coded blocks from a *constant-size* fraud proof. Filed two patents.
- **Research scientist at Applied Protocol Research** (Jan. 2019 – June 2019).
  - Invented **PolyShard**, which is a novel information-theoretic approach to simultaneously scale the security and throughput performance of a blockchain sharding system with the number of nodes in the network.
  - Designed and developed *Proof-of-Stake* version of the **Prism** consensus protocol.
- **Postdoctoral Researcher, Communication Sciences Institute, USC** (Sept. 2018 – Dec. 2018), Advisor: Prof. Salman Avestimehr.
  - *Scaling blockchain systems*: proposed and developed **Coded State Machine** to simultaneously scale the computation, storage, and security of transaction processing of blockchain systems.
  - *Coding for distributed ML systems*: invented **Lagrange Coded Computing** for distributed computing systems, which simultaneously achieve optimal resiliency to stragglers, security to adversaries, and data privacy to colluding eavesdroppers. Filed one patent.
- **Research Assistant, Communication Sciences Institute, USC** (Aug. 2014 – Aug. 2018), Advisor: Prof. Salman Avestimehr.
  1. *Coding for distributed computing systems*
    - Introduced a general distributed computing framework motivated by MapReduce. Characterized the optimal tradeoff between computations of the Map functions and the communication to exchange intermediate results.
    - Invented **coded distributed computing (CDC)**, to substantially reduce the communication load to move intermediate data of distributed computing applications, which takes advantage of redundant task executions.
    - Developed and implemented a coded distributed sorting algorithm **CodedTeraSort**, which improves the run-time performance of the sorting benchmark **TeraSort** by a factor of  $1.97 \sim 3.39\times$  on Amazon EC2 clusters.
    - Characterized a tradeoff between computation latency and communication

- load in distributed computing environments with straggling servers.
- Extended the principles of CDC into wireless distributed computing platforms, achieving a scalable design that can accommodate an unlimited number of mobile users.
  - Developed a communication-aware mobile edge computing framework, in which coded computations are designed at the edge nodes to simultaneously minimize the computation load and maximize the spectral efficiency to communicate to mobile users.
  - Developed a coded distributed computing scheme for linear computation tasks, named “compressed coded distributed computing” (compressed CDC). Compressed CDC utilizes both the combining technique to compress intermediate values for the same computation, and the coded multicast technique to combine intermediate values from different computations.
  - Designed the Batched Coupon’s Collector (BCC) scheme to optimally place data batches on multiple learners, minimizing the overall computation latency for running distributed gradient descent algorithms.
  - Minimized the computation latency for a class of distributed least-squares regression problems, by designing coded storage at each worker node.
2. *Speeding up decentralized deep learning systems*
    - Designed and implemented PCA-based gradient compressor **GradiVeQ** to reduce the communication time in a decentralized deep learning system using ring all-reduce communication.
    - Designed a pipelined training algorithm **Pipe-SGD** on a decentralized system using ring all-reduce communication, where the computation and communication of consecutive iterations are parallelized to minimize the wall-clock training time.
  3. *Scalable blockchain systems*
    - Proposed **PolyShard** for blockchain systems, which designs coded storage/computation for blockchain sharding to simultaneously achieve efficiency and security scalability.
  4. *Topological Interference Management*
    - Designed and optimized the communication schemes from the rovers on the surface of Mars to Mars-orbiting satellites (orbiters). The particular scheme exploits the time-varying characteristic of the topology of the rover-to-orbiter communication network by optimally *coding across topologies*.
    - Designed pre-coding matrices for partially connected interference networks with multiple receive antennas when no channel state information at transmitters (CSIT) is available.
- **Interim Engineering Intern at Qualcomm Technologies, Inc.** (May 2014 – Aug. 2014).  
Project: *Latency Improvements on UMTS/HSPA*.
  - **Research Assistant, Communication Sciences Institute, USC** (Aug. 2011 – May 2014), Advisor: Prof. Urbashi Mitra.
    - Proposed transmitter cooperation schemes and derived the achievable rate region (capacity region) for Gaussian multiple access channel, and optimized power allocation.

- Designed new spectrum sharing scheme for hierarchical networks enabling mutual cooperation.
- **Research Assistant, Integrated Information Systems Laboratory, NYU-Poly, Brooklyn, NY** (May 2009 – May 2011), Advisor, Prof. Elza Erkip
  - Analyzed interference mitigation methods for the Z interference channel.

## Teaching

- **IoT Thrust, HKUST.**
  - *IOTA 5101: Fog/Edge/Cloud Computing for IoT (Spring 2021, Fall 2021).*
    - Graduate level course
    - Comprehensive introduction of Fog/Edge/Cloud computing paradigms, technical challenges in these systems, and state-of-the-art research directions and solutions.
  - *IOTA 6910A: Security in IoT Systems: Theory and Applications (Spring 2022).*
    - Graduate level course
    - Fundamentals of modern cryptography and cryptographic primitives. Security vulnerabilities of practical IoT systems, and recent advancements in technologies to defend IoT security.
- **Teaching Assistant, Department of Electrical Engineering, USC.**
  - *EE 441: Applied Linear Algebra for Engineering (Fall 2013).*
    - Graduate level course
    - Responsibilities: prepared materials, led discussion sessions, provided assistance on coursework.
  - *EE 562a: Random Process in Engineering (Spring 2014).*
    - Graduate level course
    - Responsibilities: prepared materials, led discussion sessions, provided assistance on coursework.

## DISTINCTIONS & AWARDS

- Best Paper Award at NeurIPS-20 Workshop on Scalability, Privacy, and Security in Federated Learning.
- Graduation Day Talk at ITA, 2018.
- Qualcomm Innovation Fellowship Finalists, 2017.
- Graduation Day Poster at ITA, 2017.
- USC Viterbi School of Engineering Doctoral Fellowship, 2011.
- Poster presentation at the 2010 NYU-Poly convocation.
- Polytechnic Institute of NYU Promise Scholarship, 2008.

## PUBLICATIONS

### Monographs

- M 1. S. Li and S. Avestimehr, "Coded Computing", Foundations and Trends in Communications and Information Theory: Vol. 17: No. 1, pp 1-148. Aug. 2020.

## Journal Articles and Preprints

- J 7. A. R. Elkordy, S. Li, M. A. Maddah-Ali, and A. S. Avestimehr, “*Compressed Coded Distributed Computing*,” IEEE Transactions on Communications, vol. 69, no. 5, pp. 2773-2783, May 2021.
- J 6. S. Li, M. Yu, C. Yang, A. S. Avestimehr, S. Kanna, and P. Viswanath, “*PolyShard: Coded Sharding Achieves Linearly Scaling Efficiency and Security Simultaneously*,” IEEE Transactions on Information Forensics & Security, vol. 16, pp. 249-261, July, 2020.
- J 5. S. Li, M. Mousavi Kalan, Q. Yu, M. Soltanolkotabi, and A. S. Avestimehr, “*Polynomially Coded Regression: Optimal Straggler Mitigation via Data Encoding*,” e-print arXiv:1805.09934.
- J 4. S. Li, M. A. Maddah-Ali, and A. S. Avestimehr, “*Coding for distributed Fog computing*,” IEEE Communications Magazine, vol. 55, no. 4, pp. 34-40, Apr. 2017.
- J 3. S. Li, Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, “*A Scalable Framework for Wireless Distributed Computing*,” IEEE/ACM Transactions on Networking, vol. 25, no. 5, pp. 2643-2654, Oct. 2017.
- J 2. S. Li, Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, “*A Fundamental Trade-off between Computation and Communication in Distributed Computing*,” IEEE Transactions on Information Theory, vol. 64, no. 1, pp. 109-128, Jan. 2018.
- J 1. S. Li, D. Kao, and A. S. Avestimehr, “*Rover-to-Orbiter Communication in Mars: Taking Advantage of the Varying Topology*,” IEEE Transactions on Communications, Vol. 64, No. 2, Feb. 2016.

## Conference Articles

- C 27. C. Yang, J. So, C. He, S. Li, Q. Yu, and S. Avestimehr, “*LightSecAgg: Rethinking Secure Aggregation in Federated Learning*,” IEEE Information Theory Workshop (ITW 2021), Oct. 2021.
- C 26. J. Liang, W. Jiang, and S. Li, “*OmniLytics: A Blockchain-based Secure Data Market for Decentralized Machine Learning*,” ICML International Workshop on Federated Learning for User Privacy and Data Confidentiality (FL-ICML), July 2021.
- C 25. C. He, S. Li, J. So, M. Zhang, X. Zeng, H. Wang, X. Wang, P. Vepakomma, A. Singh, H. Qiu, L. Shen, P. Zhao, Y. Kang, Y. Liu, R. Raskar, Q. Yang, M. Annavaram and A. S. Avestimehr, “*FedML: A Research Library and Benchmark for Federated Machine Learning*,” NeurIPS Workshop on Scalability, Privacy, and Security in Federated Learning (SpicyFL), Dec. 2020. (Best Paper Award)
- C 24. S. Li, M. Yu, C. Yang, A. S. Avestimehr, S. Kanna, and P. Viswanath, “*PolyShard: Coded Sharding Achieves Linearly Scaling Efficiency and Security Simultaneously*,” IEEE International Symposium on Information Theory (ISIT 2020), June 2020.

- C 23. M. Yu, S. Sahraei, S. Li, A. S. Avestimehr, S. Kanna, and P. Viswanath, “*Coded Merkle Tree: Solving Data Availability Attacks in Blockchains*,” Financial Cryptography and Data Security, Feb. 2020.
- C 22. S. Li, S. Sahraei, M. Yu, A. S. Avestimehr, S. Kanna, and P. Viswanath, “*Coded State Machine - Scaling State Machine Execution under Byzantine Faults*,” ACM Symposium on Principles of Distributed Computing (PODC 2019), July 2019.
- C 21. Q. Yu, S. Li, N. Raviv, M. Mousavi Kalan, M. Soltanolkotabi, and A. S. Avestimehr, “*Lagrange Coded Computing: Optimal Design for Resiliency, Security, and Privacy*,” International Conference on Artificial Intelligence and Statistics (AISTATS 2019), Apr. 2019.
- C 20. Q. Yu, N. Raviv, S. Li, M. Mousavi Kalan, M. Soltanolkotabi, and A. S. Avestimehr, “*Lagrange Coded Computing: Optimal Design for Resiliency, Security, and Privacy*,” NeurIPS MLSystems workshop, Dec. 2018.
- C 19. Y. Li, M. Yu, S. Li, A. S. Avestimehr, NS Kim, and A. Schwing, “*Pipe-SGD: A Decentralized Pipelined SGD Framework for Distributed Deep Net Training*,” Neural Information Processing Systems (NeurIPS 2018), Dec. 2018.
- C 18. M. Yu, Z. Lin, H. V. Narra, S. Li, Y. Li, NS Kim, A. Schwing, M. Annavaram, and A. S. Avestimehr, “*GradiVeQ: Vector Quantization for Bandwidth-Efficient Gradient Aggregation in Distributed CNN Training*,” Neural Information Processing Systems (NeurIPS 2018), Dec. 2018.
- C 17. S. Li, M. A. Maddah-Ali, and A. S. Avestimehr, “*Compressed Coded Distributed Computing*,” IEEE International Symposium on Information Theory (ISIT 2018), June 2018.
- C 16. S. Li, M. Mousavi Kalan, A. S. Avestimehr, and M. Soltanolkotabi, “*Near-Optimal Straggler Mitigation for Distributed Gradient Methods*,” the 7th International Workshop on Parallel and Distributed Computing for Large Scale Machine Learning and Big Data Analytics, May 2018.
- C 15. S. Li, M. A. Maddah-Ali, and A. S. Avestimehr, “*Architectures for Coded Mobile Edge Computing*,” Fog World Congress, Oct. 2017.
- C 14. S. Li, M. A. Maddah-Ali, and A. S. Avestimehr, “*Communication-Aware Computing for Edge Processing*,” IEEE International Symposium on Information Theory (ISIT 2017), June 2017.
- C 13. S. Li, S. Supittayapornpong, M. A. Maddah-Ali, and A. S. Avestimehr, “*Coded Terasort*,” the 6th International Workshop on Parallel and Distributed Computing for Large Scale Machine Learning and Big Data Analytics, May 2017.
- C 12. Q. Yu, S. Li, M. A. Maddah-Ali, and A. S. Avestimehr, “*How to optimally allocate resources for coded distributed computing?*,” IEEE ICC, May 2017.
- C 11. S. Li, M. A. Maddah-Ali, and A. S. Avestimehr, “*A Unified Coding Framework for Distributed Computing with Straggling Servers*,” IEEE NetCod, Dec. 2016.
- C 10. S. Li, Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, “*Edge-Facilitated Wireless Distributed Computing*,” IEEE GLOBECOM, Dec. 2016.

- C 9. S. Li, Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, “*Coded Distributed Computing: Fundamental Limits and Practical Challenges,*” IEEE Asilomar Conference on Signals, Systems, and Computers, Nov. 2016.
- C 8. S. Li, Q. Yu, M. A. Maddah-Ali and A. S. Avestimehr, “*A Scalable Coded Computing Framework for Edge-Facilitated Wireless Distributed Computing,*” The First IEEE/ACM Symposium on Edge Computing, Oct. 2016.
- C 7. S. Li, M. A. Maddah-Ali, and A. S. Avestimehr, “*Coded Distributed Computing: Straggling Servers and Multistage Dataflows,*” 54rd Annual Allerton Conference on Communication, Control, and Computing, Sept. 2016.
- C 6. S. Li, M. A. Maddah-Ali, and A. S. Avestimehr, “*Fundamental Tradeoff between Computation and Communication in Distributed Computing,*” IEEE International Symposium on Information Theory (ISIT 2016), July 2016.
- C 5. S. Li, M. A. Maddah-Ali, and A. S. Avestimehr, “*Coded MapReduce,*” 53rd Annual Allerton Conference on Communication, Control, and Computing, Sept. 2015.
- C 4. S. Li, D. Kao, and A. S. Avestimehr, “*Rover-to-Orbiter Communication in Mars: Taking Advantage of the Varying Topology,*” IEEE International Symposium on Information Theory (ISIT 2015), June 2015.
- C 3. S. Li, E. Akyol and U. Mitra, “*Power Allocation for Gaussian Multiple Access Channel with Noisy Cooperative Links,*” IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2014), May 2014.
- C 2. S. Li, U. Mitra and A. Pandharipande, “*Cooperative Spectrum Sharing with Joint Receiver Decoding,*” IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2013), May 2013.
- C 1. S. Li, U. Mitra, V. Ratnam and A. Pandharipande, “*Jointly Cooperative Decode-and-Forward Relaying for Secondary Spectrum Access,*” Conference on Information Sciences and Systems (CISS 2012), Mar. 2012.

## News Articles

- N 1. S. Li, and A. S. Avestimehr, “*Coding for Distributed Computing on the Edge: Enabling Robust and Resilient Edge Computing in a Service Oriented Network,*” IEEE ComSoc Technology News (CTN) August Issue, 2018. Online at <https://www.comsoc.org/ctn/coding-distributed-computation-edge-enabling-robust-and-resilient-edge-computing-service>.

## Posters & Abstracts

- P 6. S. Li, Q. Yu, “*Accelerating Cloud Computing via Coding,*” Qualcomm Innovation Fellowship Finalists Poster, Apr. 2017.
- P 5. S. Li, Q. Yu, M. A. Maddah-Ali and A. S. Avestimehr, “*Coded Distributed Computing: Fundamental Limits and Practical Impacts,*” Information Theory and Applications Workshop (ITA) Graduation Day Poster, Feb. 2017.

- P 4. S. Li, M. A. Maddah-Ali and A. S. Avestimehr, “*Coded MapReduce: Trading Computation for Bandwidth via Coding*,” 2015 EE Research Festival, University of Southern California, Nov. 2015.
- P 3. S. Li and U. Mitra, “*A Jointly Cooperative Scheme for Secondary Spectrum Access*,” 2013 EE Research Festival, University of Southern California, Feb. 2013.
- P 2. S. Li and U. Mitra, “*Cooperative Spectrum Sharing with joint receiver decoding*,” CSI’s 30th Anniversary Conference and Celebration, USC Davidson Conference Center, Nov. 2012.
- P 1. S. Li and U. Mitra, “*Jointly cooperative decode-and-forward relaying for secondary spectrum access*,” 2012 North American School of Information Theory, Cornell University, June 2012.

PROFESSIONAL  
SERVICE

- **Technical Committee Member**

- AAAI International Workshop on Trustable, Verifiable and Auditable Federated Learning (FL-AAAI-22), 2022.
- AAAI Conference on Artificial Intelligence, 2022.
- ICML International Workshop on Federated Learning for User Privacy and Data Confidentiality (FL-ICML’21), 2021.
- MobiCom Technologies for the Wireless Edge Workshop, 2018.

- **Reviewer for**

- Journal: IEEE Transactions on Information Theory, IEEE Transactions on Signal Processing, IEEE Transactions on Communications, IEEE Journal on Selected Areas in Communications, IEEE Transactions on Wireless Communications, IEEE Transactions on Information Forensics and Security, IEEE Transactions on Knowledge and Data Engineering, IEEE Wireless Communications Magazine, IEEE Communication Letters, EURASIP
- Conference: ISIT, ITW, ICC, NeurIPS, ICML, AAAI, ICASSP, SPAWC, WUWNET.

SELECTED  
COURSE WORK

- **Graduate level courses at the University of Southern California**

- Random Processes in Engineering
- Random Processes in Engineering II
- Information Theory
- Network Information Theory
- Error Correcting Codes
- Mobile Communications
- Estimation Theory
- Real Analysis
- Convex Optimization
- Network Flows and Combinatorial Optimization



- Analysis of Algorithms
- Data Mining
- Advanced Distributed Systems

- **Advanced undergraduate level courses at NYU-Poly**

Digital Communication, Wireless Communication Laboratory, Advanced Topics in Telecommunications.

TECHNICAL  
SKILLS

- **Programming Languages:** C++, Rust, Python, GO, Matlab, LaTeX, LabVIEW.
- **Libraries and frameworks:** Apache Hadoop, Apache Spark, Apache MXNet, TensorFlow, Ray
- **Platforms:** Amazon EC2, Digital Ocean.
- **VLSI Design:** Spice.

LANGUAGES

- **Chinese:** Native
- **English**

PROFESSIONAL  
MEMBERSHIPS

- IEEE Information Theory Society
- IEEE Communication Society
- IEEE Computer Society
- Tau Beta Pi Member NY R Chapter