

John Sonchack

Associate Research Scholar
jsonch@princeton.edu

Princeton University
Department of Computer Science
35 Olden Street
Princeton, NJ 08540

EDUCATION

Ph.D. Computer Science <i>University of Pennsylvania, 2020</i> Advisor: Jonathan M. Smith	M.S.E. Computer Science <i>University of Pennsylvania, 2011</i>	B.S. Mathematics <i>Villanova University, 2009</i>
---	--	---

RESEARCH INTERESTS

Networking, distributed systems, programming languages, verification, ML, and security.

PROJECTS

Lucid [5, 2, 25] is a programming language that makes complicated high-throughput (12.8 terabits per second) network switches easy for any software engineer to program reliably. It combines correct-by-construction primitives, static analysis, and compiler optimizations to enable high-level programming that avoids common bugs while maintaining the efficiency of hand-optimized low-level code. Lucid is used at 6+ universities and research labs, in over a dozen projects from independent studies to top conference papers and NSF/DARPA projects.

Mantis [8, 3, 15] pushes the limit of low-latency network control. It lets commodity devices react to network events in $100\mu s$ with consistent sequential updates, by applying driver-level optimizations and domain-specific synchronization algorithms. Applications include load balancing for dynamic workloads (e.g., LLM training), security, and fair rate limiting.

StarFlow [10, 11, 26] demonstrates how to support fine-grained monitoring efficiently at scale. It combines a compute-friendly intermediate data representation with a hybrid architecture to give analytics pipelines and ML-based decision processes visibility into every event without adding major infrastructure overhead. It has been used on Princeton's network for 3+ years.

Speedlight [12, 7, 1] revives the classical distributed systems primitive of a *global snapshot* for modern data center environments. It gives users the ability to sample the state of an entire distributed system at a consistent point in time. An ML-driven profiler stitches snapshots into traffic patterns to reveal otherwise hidden phenomena such as hotspots and stragglers.

Flightplan [6, 7, 19] automatically partitions and distributes a network program across heterogeneous hardware accelerators to optimize for user objectives such as power, throughput, and latency. A suite of hardware-accelerated primitives give Flightplan programs powerful new capabilities such as forward error correction and real-time payload inspection.

ROCK [28, 16] applies machine learning to enable proactive security through collaboration. It recommends security rules before local deployment, based on predicted effectiveness at other networks with similar latent profiles. The underlying algorithms are domain-specific adaptations of collaborative filtering and stochastic simulation.

PUBLICATION HIGHLIGHTS AND AWARDS

Citations: 990

h-index: 15

i10-index: 18

Reproducibility Badges and Awards at OSDI (Beaver 2024); SIGCOMM (Cebinae 2022, Lucid 2021, Mantis 2020); NSDI (OrbWeaver 2022, Flightplan 2021); and CONEXT (DeepMatch 2020)

Best Student Paper Award at EuroSys 2018 for TurboFlow

PUBLICATIONS

Refereed conference papers:

1. Liangcheng Yu, Xiao Zhang, Haoran Zhang, John Sonchack, Dan Ports, and Vincent Liu (2024). Beaver: Practical Partial Snapshots for Distributed Cloud Services. In: 18th USENIX Symposium on Operating Systems Design and Implementation (OSDI 24). Santa Clara, CA: USENIX Association, pp.233–249.
2. Vaibhav Mehta, Devon Loehr, John Sonchack, and David Walker (2023). SwitchLog: A Logic Programming Language for Network Switches. In: Practical Aspects of Declarative Languages. Springer Nature Switzerland, pp.180–196.
3. Liangcheng Yu, John Sonchack, and Vincent Liu (2022). Cebinae: scalable in-network fairness augmentation. In: Proceedings of the ACM SIGCOMM 2022 Conference, pp.219–232.
4. Liangcheng Yu, John Sonchack, and Vincent Liu (2022). OrbWeaver: Using IDLE Cycles in Programmable Networks for Opportunistic Coordination. In: 19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22), pp.1195–1212.
5. John Sonchack, Devon Loehr, Jennifer Rexford, and David Walker (2021). Lucid: A Language For Control in the Data Plane. In: ACM Conference on Special Interest Group on Data Communication (SIGCOMM), pp.296–309.
6. Nik Sultana, John Sonchack, Hans Giesen, Isaac Pedisich, Zhaoyang Han, Nishanth Shyamkumar, Shivani Burad, André DeHon, and Boon Thau Loo (2021). Flightplan: Dataplane Disaggregation and Placement for P4 Programs. In: 18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21). USENIX Association, pp.571–592.
7. Joel Hypolite, John Sonchack, Shlomo HersHKop, Nathan Dautenhahn, André DeHon, and Jonathan M. Smith (2020). "DeepMatch: Practical deep packet inspection in the data plane using network processors." In Proceedings of the 16th International Conference on emerging Networking EXperiments and Technologies, pp. 336-350.
8. Nofel Yaseen, John Sonchack, and Vincent Liu (2020). tpprof: A Network Traffic Pattern Profiler. In: USENIX Symposium on Networked Systems Design and Implementation (NSDI), pp.1015–1030.
9. Liangcheng Yu, John Sonchack, and Vincent Liu (2020). Mantis: Reactive Programmable Switches. In: ACM Conference on Special Interest Group on Data Communication (SIGCOMM), pp.296–309.
10. Nikos Vasilakis, Ben Karel, Yash Palkhiwala, John Sonchack, André DeHon, and Jonathan M. Smith (2019). Ignis: scaling distribution-oblivious systems with light-touch distribution. In: ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI).
11. John Sonchack, Adam J. Aviv, Eric Keller, and Jonathan M. Smith (2018). Turboflow: Information rich flow record generation on commodity switches. In: European Conference on Computer Systems (EuroSys). **[best student paper]**
12. John Sonchack, Oliver Michel, Adam J. Aviv, Eric Keller, and Jonathan M. Smith (2018). Scaling hardware accelerated network monitoring to concurrent and dynamic queries with *flow. In: USENIX Annual Technical Conference (ATC).
13. Nofel Yaseen, John Sonchack, and Vincent Liu (2018). Synchronized network snapshots. In: ACM Conference on Special Interest Group on Data Communication (SIGCOMM).
14. John Sonchack, Anurag Dubey, Adam J. Aviv, Jonathan M. Smith, and Eric Keller (2016). Timing-based Reconnaissance and Defense in Software-defined Networks. In: Annual Computer Security Applications Conference (ACSAC).

15. John Sonchack, Adam J. Aviv, Eric Keller, and Jonathan M. Smith (2016). Enabling Practical Software-defined Networking Security Applications with OFX. In: Network and Distributed System Security Symposium (NDSS).
16. John Sonchack and Adam J. Aviv (2014). LESS Is More: Host-Agent Based Simulator for Large-Scale Evaluation of Security Systems. In: Annual European Symposium on Research in Computer Security (ESORICS).

Refereed workshop papers:

17. Henri Maxime Demoulin, Nikos Vasilakis, John Sonchack, Isaac Pedisich, Vincent Liu, Boon Thau Loo, Linh Thi Xuan Phan, Jonathan M. Smith, and Irene Zhang (2019). TMC: Pay-as-you-Go Distributed Communication. In: Asia-Pacific Workshop on Networking (AP- NET).
18. Oliver Michel, John Sonchack, Eric Keller, and Jonathan M. Smith (2019). PIQ: Persistent Interactive Queries for Network Security Analytics. In: ACM International Workshop on Security in Software Defined Networks & Network Function Virtualization (SDN-NFV Sec.)
19. Hans Giesen, Lei Shi, John Sonchack, Anirudh Chelluri, Nishanth Prabhu, Nik Sultana, Latha Kant, Anthony J. McAuley, Alexander Poylisher, Andre DeHon, et al. (2018). In-network computing to the rescue of faulty links. In: ACM SIGCOMM 2018 Morning Workshop on In-Network Computing (NetCompute).
20. Oliver Michel, John Sonchack, Eric Keller, and Jonathan M. Smith (2018). Packet-level analytics in software without compromises. In: USENIX Workshop on Hot Topics in Cloud Computing (HotCloud).
21. John Sonchack and Jonathan M. Smith (2018). PathMiner Powered Predictable Packet Processing. NDSS Workshop on Binary Analysis Research (BAR).
22. John Sonchack, Adam J. Aviv, and Eric Keller (2016). Timing SDN Control Planes to Infer Network Configurations. In: ACM International Workshop on Security in Software Defined Networks and Network Function Virtualization (SDN-NFV Sec.)
23. John Sonchack and Adam J. Aviv (2013). Bridging the Data Gap: Data Related Challenges in Evaluating Large Scale Collaborative Security Systems. In: USENIX Workshop on Cyber Security Evaluation and Testing (CSET).
24. John Sonchack and Jonathan M. Smith (2011). Signature Correlations in Multiple Honey-pot Defense System. In: Proceedings of the Future Internet Workshop.

Journals and magazine articles:

25. Mary Hogan, Devon Loehr, John Sonchack, Shir Landau Feibish, Jennifer Rexford, David Walker. (2025). Automated Optimization of Parameterized Data-Plane Programs with Parasol. IEEE/ACM Transactions on Networking (ToN).
26. Oliver Michel, John Sonchack, Greg Cusack, Maziyar Nazari, Eric Keller, and Jonathan M. Smith (2021). Software packet-level network analytics at cloud scale. IEEE transactions on network and service management 18, no. 1: 597-610.
27. John Sonchack and Adam J. Aviv (2016). Exploring Large Scale Security System Reproducibility with the LESS Simulator. Journal of Computer Security.
28. John Sonchack, Adam J. Aviv, and Jonathan M. Smith (2015). Cross-Domain Collaboration for Improved IDS Rule Set Selection. Journal of Information Security and Applications.

In submission:

29. Pamela Zave, John Sonchack, and Jennifer Rexford. Verifiable Lucid: Language-based Data Plane Program Verification.
30. Andrew Johnson, John Sonchack, Srinivas Narayana, and David Walker. Mario: A Language for Developer-guided Packet Pipeline Parallelization.

FUNDING AND GRANT MANAGEMENT

Tools for Programming Distributed Data plane Measurements

NSF IMR, 2022-2025

Role: Senior Personnel

Contribution: Co-authored proposal based on Lucid [5], mentored Ph.D. dissertations, supervised student research, developed tutorials, and reported key results to NSF.

Collaborators: Princeton University, University of Virginia

ProD3

DARPA OPS-5G, 2020-2024

Role: Researcher (2020-2022), Principle Investigator (2022-2024)

Contribution: Led research projects, developed open-sourced tools, supervised and mentored students, and communicated key results to DARPA at site visits and PI meetings.

Collaborators: Princeton University, Peraton Labs

Active Security

NSF SaTC, 2014-2019

Role: Ph.D. Student

Contribution: Authored key section of proposal, designed and published multiple systems for low-latency control and efficient fine-grained telemetry that led to project success.

Collaborators: University of Pennsylvania, University of Colorado Boulder

WORK EXPERIENCE

Associate Research Scholar

Princeton University, 2022-current

- Evolved Lucid into a practical language for research prototypes and teaching by continuously integrating feedback from students and researchers at multiple universities and labs.
- Co-wrote an NSF grant proposal based on Lucid that was awarded \$600,000.
- Served as a Principle Investigator on the PRoD3 project, part of DARPA's OPS-5G program.
- Mentored 4 Ph.D. dissertation projects, 2 MSE independent studies, and 7 undergrad projects.
- Established and maintained collaborations with industry and academia, including: Peraton Labs, Charles River Analytics, CAIDA, The University of Pennsylvania, The University of Virginia, Hebrew University, Microsoft Research, and Nvidia.

Postdoctoral Researcher

Princeton University, 2020-2022

- Designed, implemented, and published research on languages and compilers for high-performance programmable networks.
- Mentored students on topics including: experimental design, engineering principles, identifying novel topics, testbed infrastructure, and improving communication.
- Deployed a privacy-preserving telemetry system to the campus network based on StarFlow [12]. It operated for over 4 years and enabled multiple measurement studies.
- Developed core parts of a 5G DDoS defense demo for DARPA, which led to the successful funding of Pronto, a large (\$30M) project across multiple universities and companies.

Network Programming and Cybersecurity Consultant

Independent, 2012-current

Selected projects:

Control Protocol Compression in Constrained Tactical Networks (*Peraton Labs, 2025*). Analyzed OSPF protocol and designed model and aggregation based techniques to reduce OSPF bandwidth utilization by 4 orders of magnitude in a DoD prototype.

Luxon P4 (*Stateless, Inc., 2019*). Designed and implemented a custom 6.4 Tb/s P4 data plane for the Luxon rack-scale product. Features included: VXLAN & IPSec classification, custom load balancing, telemetry, and run-time configurable QoS.

Cybersecurity Insurance Market Research (*Chubb Insurance, 2012*). Identified and distilled major cybersecurity threats relevant to large organizations and best practices for policyholder risk mitigation.

Engineering Intern

Financial Software Systems, Hatboro, PA, USA 2010

- Implemented high-performance analytic pricing models in C for security futures.
- Extended core product to support user-defined pricing models in Python.

TEACHING EXPERIENCE

Dissertation and thesis mentor

2012-current

- Mentored and identified research directions for dissertation and thesis projects.
 1. Mary Hogan. Language Expressiveness Under Extreme Scarcity in Programmable Data Planes. Ph.D. Dissertation. Princeton, 2024.
 2. Devon Loehr. Lucid: A High-Level, Easy-To-Use Dataplane Programming Language. Ph.D. Dissertation. Princeton, 2024.
 3. Liangchen Yu. Towards Zero-waste Terabit Networked Systems. Ph.D. Dissertation. University of Pennsylvania, 2024.
 4. Nofel Yaseen. Network-wide Monitoring and Debugging. Ph.D. Dissertation. University of Pennsylvania, 2022.
 5. Oliver Michel. Packet-Level Network Telemetry and Analytics. Ph.D. Dissertation. University of Colorado Boulder, 2019.
 6. Sumanth Sathyanarayana. Software Defined Network Defense. MSE Thesis. University of Pennsylvania, 2012.

Data-plane engineering advisor

2022-current

- Introduced students implementing data plane systems in Lucid to relevant engineering principles, e.g., performant data structures, algorithmic optimizations, and efficient idioms.
 1. Yufei Zheng, Huacheng Yu, and Jennifer Rexford. Detecting tcp packet reordering in the data plane. arXiv:2301.00058. (*In submission*)
 2. Mary Hogan, Shir Landau Feibish, Jennifer Rexford, David Hay. Multipath Syn Flood defense. (*In submission*)
 3. Yoo, Sophia, Xiaoqi Chen, and Jennifer Rexford. SmartCookie: Blocking Large-Scale SYN Floods with a Split-Proxy Defense on Programmable Data Planes. USENIX Security. 2024.
 4. Henry Birge-Lee, Sophia Yoo, Benjamin Herber, Jennifer Rexford, and Maria Apostolaki. TANGO: Secure Collaborative Route Control across the Public Internet. NSDI. 2024.
 5. Andrew Johnson, Ryan Beckett, Xiaoqi Chen, Ratul Mahajan, and David Walker. Sequence Abstractions for Flexible, Line-Rate Network Monitoring. NSDI. 2024.
 6. Satadal Sengupta, Hyojoon Kim, and Jennifer Rexford. Continuous in-network round-trip time monitoring. SIGCOMM. 2022.
 7. Hyunkeun O. Real-time Pivot Attack Detection. Independent study, University of Virginia. 2024.
 8. Lucas Kohler. A Library for Line-rate Traffic Anonymization. Independent study, University of Virginia. 2024.
 9. Hieu Vu. Real-time QUIC round-trip-time measurement. Independent study, University of Virginia. 2024.
 10. John Bloch. Automated testing for network switch programs. Princeton/Intel REU. 2024.
 11. Iris Shi. Language-supported Data Plane Security. Princeton/Intel REU. 2023.
 12. Ifesi Dimma Onubogu. Verifiable Syn Flood Defenses. Princeton/Intel REU. 2022.

Research supervisor, Princeton/Intel REU Program

2022-2024

- Supervised and mentored annual undergraduate projects related to network security.

Guest lecturer, CIS 553: Networked Systems, University of Pennsylvania

2016-2018

- Introduced Software-defined Networking and programmable data plane fundamentals.

Teaching assistant, CIS 551: Computer & Net. Security, University of Pennsylvania

2011-2012

- Updated and virtualized projects, automated project submission and grading, set up and moderated discussion forums, held weekly office hours, supervised grading assistants.

PROFESSIONAL SERVICE

Journal reviewer: Transactions on Networking, Letters of the Computer Society, IEEE Communications Letters, Transactions on Dependable and Secure Computing, Transactions on Information Forensics & Security

Program committee member: USENIX Symposium on Networked Systems Design and Implementation (2025), The International Conference on Internet Monitoring and Protection, (2018-2019), IEEE Global Communications Conference (2018-2019)

Workshop program committee member: Security in Softwarized Networks: Prospects and Challenges (SecSoN) (2018), ACM International Workshop on Security in Software Defined Networks & Network Function Virtualization (2017), Workshop on Cyber Security Experimentation and Testing (2013-2015)

External reviewer: ACM/IEEE Symposium on Architectures for Networking and Communications Systems (2019), Annual Computer Security Applications Conference (2013-2014)

NSF panelist and visioning: Information and Intelligent Systems Division (2019-current), NSF Visioning Workshop on Programmable Security in a Software Defined World (2018)

OPEN SOURCE SOFTWARE

Lucid, an easy-to-use language for programming multi-terabit switches.

<https://github.com/PrincetonUniversity/lucid>

Parasol, a framework for self-optimizing Lucid libraries.

<https://github.com/mhogan26/Parasol>

Mantis, a platform for sub-millisecond network control.

<https://github.com/eniac/Mantis>

Orbweaver, a framework for zero-overhead data plane control protocols.

<https://github.com/eniac/OrbWeaver>

Speedlight, a primitive for consistent network snapshots in line-rate data planes.

<https://github.com/eniac/Speedlight>

tpprof, a machine-learning based profiler for snapshot traces.

<https://github.com/eniac/tpprof>

Beaver, practical partial snapshots for cloud environments.

<https://github.com/eniac/Beaver>

Cebinae, a fairness-augmenting rate limiting protocol.

<https://github.com/eniac/Cebinae>

TurboFlow, performance-optimized flow telemetry for commodity switches.

<https://github.com/jsonch/TurboFlow>

StarFlow, a hybrid system for efficient packet-level telemetry.

<https://github.com/jsonch/StarFlow>