

David Darais

darais@galois.com

david.darais.com

Appointments

Sep 2020–present: Principal Scientist, Galois, Inc.
Jan 2018–Aug 2020: Assistant Professor (tenure-track), Computer Science, University of Vermont, VT
Jun 2017–Dec 2017: Postdoctoral Researcher, University of Maryland, College Park, MD

Research Interests

Data Privacy, Differential Privacy, Secure Multiparty Computation, Oblivious Computation, Software Verification, Type Systems, Program Analysis, Abstract Interpretation, Mechanized Proofs, Parsing.

Education

Jun 2017 PhD, Computer Science, University of Maryland
PhD Thesis: “Mechanizing Abstract Interpretation”
PhD Advisor: David Van Horn
Jan 2015 MS, Computer Science, Harvard University
PhD Qualifying Exam (passed): “Abstract Control in Program Analysis”
PhD Advisor: Greg Morrisett
May 2011 BS, Computer Science, University of Utah
BS Thesis: “Extracting the Essence of Type Classes”
BS Advisors: Matthew Might & Matthew Flatt

Peer-reviewed Conference Publications

- 2021 [11] **DDUO: General-Purpose Dynamic Analysis for Differential Privacy.**
Chike Abuah, Alex Silence, David Darais, Joseph P. Near.
Computer Security Foundations (CSF). IEEE, 2021.
- 2020 [10] **Types and Abstract Interpretation for Authorization Hook Advice.**
Christian Skalka, David Darais, Trent Jaeger, Frank Capobianco.
Computer Security Foundations (CSF). IEEE, 2020.
- 2020 [9] **Abstracting Faceted Execution.**
Kris Micinski, David Darais, Thomas Gilray.
Computer Security Foundations (CSF). IEEE, 2020.
- 2020 [8] **A Language For Probabilistically Oblivious Computation.**
David Darais, Ian Sweet, Chang Liu, Michael Hicks.
Principles of Programming Languages (POPL). ACM, 2020.
- 2019 [7] **Proof Carrying Network Code.**
Christian Skalka, John Ring, David Darais, Minseok Kwon, Sahil Gupta, Kyle Diller, Steffan Smolka, Nate Foster.
Computer and Communications Security (CCS). ACM, 2019.
- 2019 [6] **Duet: An Expressive Higher-order Language and Linear Type System for Statically Enforcing Differential Privacy.**
Joseph P. Near, David Darais, Chike Abuah, Tim Stevens, Pranav Gaddamadugu, Lun Wang, Neel Somani, Mu Zhang, Nikhil Sharma, Alex Shan, Dawn Song.
Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA). ACM, 2019.
«*ACM SIGPLAN Distinguished Paper Award*»

- 2017 [5] **Abstracting Definitional Interpreters.**
David Darais, Nicholas Labich, Phc C. Nguyn, David Van Horn.
International Conference on Functional Programming (ICFP). ACM, 2017.
- 2016 [4] **Constructive Galois Connections: Taming the Galois Connection Framework for Mechanized Metatheory.**
David Darais, David Van Horn.
International Conference on Functional Programming (ICFP). ACM, 2016.
- 2015 [3] **Galois Transformers and Modular Abstract Interpreters: Reusable Metatheory for Program Analysis.**
David Darais, Matthew Might, David Van Horn.
Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA). ACM, 2015.
- 2013 [2] **Monadic Abstract Interpreters.**
Ilya Sergey, Dominique Devriese, Matthew Might, Jan Midtgaard, David Darais, Dave Clarke, Frank Piessens.
Programming Language Design and Implementation (PLDI). ACM, 2013.
- 2011 [1] **Parsing with Derivatives: A Functional Pearl.**
Matthew Might, David Darais, Daniel Spiewak.
International Conference on Functional Programming (ICFP). ACM, 2011.

Peer-reviewed Journal Publications

- 2019 [2] **Constructive Galois Connections.**
David Darais, David Van Horn.
Journal of Functional Programming (JFP). Cambridge University Press, 2019.
- 2012 [1] **Macros that Work Together: Compile-time Bindings, Partial Expansion, and Definition Contexts.**
Matthew Flatt, Ryan Culpepper, David Darais, Robert Bruce Findler.
Journal of Functional Programming (JFP). Cambridge University Press, 2012.

Peer-reviewed Workshop Publications

- 2020 [3] **DuetSGX: Differential Privacy with Secure Hardware.**
Phillip Nguyen, Alex Silence, David Darais, Joseph P. Near.
Theory and Practice of Differential Privacy (TPDP). 2020.
- 2020 [2] **Short Paper: Probabilistically Almost-Oblivious Computation.**
Ian Sweet, David Darais, Michael Hicks.
Programming Languages and Analysis for Security (PLAS). ACM, 2020.
- 2016 [1] **Compositional and Mechanically Verified Program Analyzers.**
David Darais.
European Conference on Object-Oriented Programming Doctoral Symposium (ECOOP-DS). Dagstuhl Publishing, 2016.

Theses

- 2017 [2] **Mechanizing Abstract Interpretation.**
David Darais.
PhD Thesis. University of Maryland, College Park, 2017.
- 2011 [1] **Abstracting the Essence of Type Classes.**
David Darais.
BS Thesis. University of Utah, 2011.

Invited Talks

- 2021 [6] **Data-oblivious Computation.**
Hot Topics in the Science of Security (HotSoS) Symposium - Hard Problems Special Session.
Virtual. 2021.
- 2019 [5] **Data Privacy by Programming Language Design.**
Tech Talk - Millennium Institute - Foundational Research on Data.
University of Chile, Santiago, Chile. 2019.
- 2017 [4] **A Simple and Extensible Approach to Program Analysis.**
International Federation for Information Processing Working Group 2.4 on Software Implementation Technology (IFIP WG2.4).
Essex, Vermont. 2017.
- 2016 [3] **Constructive Galois Connections.**
New Jearsey Programming Languages and Systems Seminar (NJPLS).
Philadelphia, Pennsylvania. 2016.
- 2016 [2] **Adventures in Abstract Interpretation.**
Department Research Seminar.
University of Utah, Salt Lake City, Utah. 2016.
- 2016 [1] **Constructive Galois Connections and Applications to Gradual Typing.**
Department Research Seminar.
University of Chile, Santiago, Chile. 2016.

Funded Projects

- 2021 [9] **DARPA: SocialCyber: LAGOON: Leveraging AI to Guard Online Open Source Networks.**
Funded by DARPA. Galois, Inc. award \$499,718 Phase 1. Galois, Inc. PI: David Darais; Co-PI Walt Woods. 9 months, starting TBD 2021. Scientific peer reviewed.
The LAGOON project develops AI techniques for securing online open source software from social threats to development and maintenance workflows. The project develops new data modeling and fusion techniques, NLP algorithms, predictive models for social behavior, and threat analysis for code dependencies.
- 2021 [8] **ONR: SBIR: Phase 1: 5STARS: 5G SDN Tools for Automated and Reliable Security.**
Funded by ONR. Galois, Inc. award \$139,905 Phase 1. Galois, Inc. PI: David Darais. 6 months, starting TBD 2021. Scientific peer reviewed.
The 5STARS project develops verification tools for software defined networking applications applied within 5G infrastructure. The project develops new verification techniques for data plane and control plane functionalities, and new case studies for SDN verification within 5G.
- 2020 [7] **CLAMPED: Collaborative Learning Architecture with Mathematical Privacy over Embedded Data.**
Funded by DARPA. UVM Award \$323,717. UVM PI: Joe Near. UVM Co-PI: David Darais. 1.5 years, starting 08/28/2020. Scientific peer reviewed. Sub-award transferred to Galois, Inc.. Galois PI: David Darais.
In this project we develop data privacy techniques suitable for collaborative secure learning. Our techniques enable collaborative machine learning scenarios involving sensitive data, including differentially private training of deep neural networks.

- 2020 [6] **DARPA: SIEVE: Wizkit: Wide-scale Zero-Knowledge Interpreter Toolkit.Analyzers for Critical Software.**
 Funded by DARPA. UVM Award \$405,858. UVM PI: Joe Near. Key Personnel: David Darais. 4 years, starting 05/01/2020. Scientific peer reviewed.
In this project we develop a system, called Wizkit, that efficiently compiles and executes zero-knowledge proofs for a wide range of applications. Wizkit will enable the adoption of zero-knowledge proofs in new domains with high societal benefit, such as secure auctions, auditing and voting.
- 2020 [5] **Amazon: ARA: Provable Fairness for Deep Learning via Automatic Differentiation.**
 Funded by Amazon Research Awards (ARA). UVM Award \$91,749. UVM PI: Joe Near. UVM co-PI: David Darais. 1 year, starting 04/15/2020. Amazon internal review.
In this project we propose a new approach for enforcing fairness in the context of deep neural networks that provides a provable guarantee by analyzing the network’s architecture directly.
- 2020 [4] **DoD: AVATAR: Army Visual and Tactical Arctic Reconnaissance.**
 Funded by DoD. UVM Award \$3,700,000. UVM PI: Jeff Marshall. Investigator: David Darais. 3 years, starting 04/01/2020. DoD internal review.
In this project we develop defense techniques to secure machine learning models from physical corruption, data poisoning and adversarial examples.
- 2019 [3] **NSF: SHF: Medium: Collaborative Research: Synthesizing Verified Analyzers for Critical Software.**
 Funded by UVM. UVM Award \$598,434. UVM PI: David Darais. 4 years, starting 06/03/2019. Scientific peer reviewed.
In this project we develop a system, called PANTHEON, that transforms the description of a computation involving sensitive data into a cryptographically secure, “end-to-end” solution that satisfies specified security goals. PANTHEON will greatly expand the number of programmers capable of developing secure-computation protocols, and the number of users who will benefit from their deployment.
- 2019 [2] **IARPA: HECTOR: PANTHEON: Programming Architecture iNtegrated Toolchain for compiling Homomorphic Encryption and ONline Secure Computation.**
 Funded by IARPA. UVM Award \$527,262. UVM PI: David Darais. 5 years, starting 10/01/2019. Subcontract with Stealth Software, Inc.. Scientific peer reviewed.
In this project we investigate both formal verification and automated synthesis of program analyzers using interactive proof assistants. Our results will both accelerate existing approaches for designing high assurance analyzers, as well as enable developers without expertise in formal verification to prototype them.
- 2019 [1] **UVM: REACH: Data Privacy for Deep Learning via Language Design.**
 Funded by UVM. UVM Award \$29,409. UVM PIs: David Darais and Joe Near. 1.25 years, starting 05/01/2019. UVM internal review.
We propose to discover new techniques for applying differential privacy to deep learning. In contrast to previous efforts, we plan to leverage our experience in programming languages to develop novel techniques that address this challenge.

Teaching

- 2020 Spring UVM CS 225: Programming Languages.
 2019 Fall UVM CS 295A: Software Verification.
 2019 Spring UVM CS 225: Programming Languages.
 2018 Fall UVM CS 295A: Software Verification.
 2018 Spring UVM CS 225: Programming Languages.

Institutional Activities

- 2017–2020: University of Vermont Center for Computer Security and Privacy. Core member.
<https://compsec.w3.uvm.edu>.
- 2013–2015: Harvard Computer Science Graduate Council. Founding Chair.
- 2012–2013: Harvard School of Engineering and Applied Science Graduate Council. Founding Co-President.

Professional Activities

- NSF Panels: 2 SHF; (dates confidential)
- Steering Committee: TyDe 2020, 2021 (chair), 2022
- Organizing Chair: TyDe 2019 (co), OOPSLA SRC 2019 (co), ECOOP DS 2017 (co)
- Program Committee (PC): PLDI 2021, TrustNLP 2021, ICFP 2020, TyDe 2018, IFL 2018
- External/Extended Review Committee (ERC): PLDI 2020, ICFP 2018
- Posters Committee: ICFP SRC 2019 (+ judge); ECOOP 2019 Posters
- Mentoring Workshops: PLMW 2018 (panel)
- Video Chair: OOPSLA 2017; POPL 2017; PLDI 2017, 2016; ECOOP 2017, 2016; ICFP 2013
- Artifact Evaluation Committee (AEC): POPL 2016
- Student Volunteer: POPL 2016, ICFP 2013 (chair)
- Logo Designer: ICFP 2013

Professional Training

- Process Oriented Guided Inquiry Learning (POGIL). Workshop, May 21.

Awards

- Lin Fellowship. Harvard University. 2012.
- GSAS Graduate Fellowship. Harvard University. 2011.
- Magna Cum Laude Graduation Honors (3.97 GPA). University of Utah. 2011
- College of Engineering Scholarship. University of Utah. 2007, 2008, 2009.