

# Michelle Goodstein

✉ [resume@michellegoodstein.com](mailto:resume@michellegoodstein.com)

🌐 [www.michellegoodstein.com](http://www.michellegoodstein.com)

---

## Professional Experience

- May 2018 – present **Senior Software Engineer**, *Google*, Kirkland, WA.  
Improving BigQuery's ability to ingest and query externally partitioned data stored in open-source formats.
- Led end-to-end design and development of BigQuery's ability to [query](#) and [ingest](#) externally partitioned data stored in default hive-partitioned layouts on Google Cloud Storage (GCS).
  - Personally drove backend implementation as well as API design, with significant involvement in UI and DDL design.
  - Generally available as of [March 2020](#). Demoeed at [Cloud NEXT Tokyo 2019](#) as well as a ["A Very Special BigQuery Day \(The Data Show\)"](#).
  - Tech lead for ongoing related improvements in this area.
- November 2014 – March 2018 **Research Scientist**, *Facebook*, Seattle, WA.  
Large-scale distributed systems development in data infrastructure, spanning data warehouse and real-time monitoring.
- Led project to improve real-time alerting systems by surfacing additional information.
  - Led prioritization of reliability improvements across real-time detection and alerting.
- Summer 2006 **Software Engineering Intern**, *Google*, Kirkland, WA.  
Mobile application development.
- Summer 2005 **Software Engineering Intern**, *Google*, Kirkland, WA.  
Web application development, focusing on AJAX.

---

## Education

- August 2014 **PhD, Computer Science**, *Carnegie Mellon University*, Pittsburgh, PA.  
Thesis Title: *Dataflow Analysis-Based Dynamic Parallel Monitoring*  
Advisor: Prof. Todd Mowry
- May 2008 **MS, Computer Science**, *Carnegie Mellon University*, Pittsburgh, PA.
- June 2005 **BS, Computer Science** and **BS, Mathematics**, *Univ. of Washington*, Seattle, WA.  
Magna Cum Laude with College Honors in Computer Science

---

## Programming Languages

Proficient with: C/C++, Java. Familiar with: PHP(Hack), Python.

---

## Professional Service

- Google **Diversity, Equity and Inclusion**
- Women in Data Analytics - WA*: Creator/organizer. March 2019 - present  
*CoreInfra DEI Steering Committee*: Member. March 2020 - present  
*Data Analytics DEI Working Group*: Member. May 2020 - present

## Academic Service

*Program Committee: PPOPP 2018, PPOPP 2020*

*External Review Committee: PPOPP 2017, PACT 2016*

*Student Research Competition Judge: PACT 2015*

---

## Publications

### Dissertation

**Michelle L. Goodstein.** Dataflow Analysis-Based Dynamic Parallel Monitoring. Carnegie Mellon University Technical Report: CMU-CS-14-132, August 2014.

### Peer Reviewed Conferences

- PACT 2015 **Michelle L. Goodstein**, Phillip B. Gibbons, Michael A. Kozuch and Todd C. Mowry. Tracking and Reducing Uncertainty in Dataflow Analysis-Based Dynamic Parallel Monitoring. In *Proceedings of the Twenty-Fourth International Conference on Parallel Architectures and Compilation Techniques*, October 2015.
- PACT 2012 **Michelle L. Goodstein**, Shimin Chen, Phillip B. Gibbons, Michael A. Kozuch, and Todd C. Mowry. Chrysalis Analysis: Incorporating Synchronization Arcs in Dataflow-Analysis-Based Parallel Monitoring. In *Proceedings of the Twenty-First International Conference on Parallel Architectures and Compilation Techniques*, September 2012.
- ICRA 2011 Michael P. Ashley-Rollman, Padmanabhan S. Pillai, and **Michelle L. Goodstein**. Simulating multi-million-robot ensembles. In *IEEE International Conference on Robotics and Automation*, May 2011.
- ASPLOS 2010 **Michelle L. Goodstein**, Evangelos Vlachos, Shimin Chen, Phillip B. Gibbons, Michael A. Kozuch, and Todd C. Mowry. Butterfly Analysis: Adapting Dataflow Analysis to Dynamic Parallel Monitoring. In *Proceedings of the Fifteenth International Conference on Architectural Support for Programming Languages and Operating Systems*, March 2010.
- ASPLOS 2010 Evangelos Vlachos, **Michelle L. Goodstein**, Michael A. Kozuch, Shimin Chen, Babak Falsafi, Phillip B. Gibbons, and Todd C. Mowry. ParaLog: Enabling and Accelerating Online Parallel Monitoring of Multithreaded Applications. In *Proceedings of the Fifteenth International Conference on Architectural Support for Programming Languages and Operating Systems*, March 2010.