

RESEARCH INTERESTS

(Interests: Program Analysis and Verification, Developer Tools, ML Compilers.)

My interests broadly span programming languages, compilers, and developer tools. Specifically, I am interested in program analysis, verification, and testing problems that improve the productivity of programmers. I also have experience working at the intersection of Machine Learning (ML), compilers, and programming languages. I also enjoy building and working with large-scale systems.

- At Google, I have led and worked on a variety of projects spanning my interests:
 - My recent focus is the use of formal methods to ensure privacy of data in ML applications.
 - I worked on the Swift for TensorFlow project that improved the productivity of ML engineers.
 - I have also worked on ML models/embeddings for Code [ICML2020].
 - I was a founding member of the team that developed program analysis tools to improve the security of Android Apps.
- When I was at NEC Labs, my latest focus was on tools for analyzing and testing C/C++ programs. In particular, I worked on an automated test-case generation tool for C/C++ programs, and a source-to-source translation tool that models inheritance and exceptions in such a way that makes it easier to statically analyze C++ programs [ECOOP11, FoVeOOS11, CC12, ICSE13, ISSTA14].
- At NEC Labs, I was also involved in building a framework for testing large-scale distributed systems such as HBase, Zookeeper, and so on [TRIOS 2013].
- I was also involved in building NEC's internal bug finding tool called Varvel that is based on static analysis and model checking technologies [SAS08, SEFM08, EMSOFT09, FMCAD10, ASE11]. Varvel is used internally at NEC Japan on millions of lines of C/C++ code on a daily basis.
- For my Ph.D. dissertation, I developed static-analysis algorithms to address the problem of analyzing executables without debugging information [CC04, CC05, CAV05, SAS06, VMCAI07, TACAS08, TOPLAS10].

RESEARCH AND WORK EXPERIENCE

Aug 2014 - Present	Staff Engineer	Google, Mountain View, CA, USA
Jan 2014 - Jul 2014	Software Engineer	Facebook, Menlo Park, CA, USA
Oct 2007 - Jan 2014	Research Staff Member	NEC Laboratories Inc., Princeton, NJ, USA
Aug 2001 - Aug 2007	Research Assistant	University of Wisconsin, Madison, WI, USA

AWARDS AND DISTINCTIONS

- 2016 Google TI Award
- 2012 NECLA Technology Commercialization Award
- 2008 Outstanding Graduate Research Award (UW Computer Sciences Department)
- 2004 EAPLS Best Paper Award at ETAPS 2004 (with T. Reps)

TECHNOLOGY TRANSFER

- Technology transfer to NEC Corporation, Japan. Led the design, implementation, and transfer of an analysis component that lowers a C++ program into an analysis-friendly C program. (cf. my CC12 and ECOOP11 papers.) It is being currently used within NEC Corporation, Japan.
- Technology transfer to GrammaTech, Inc., Ithaca, NY, USA. Led the design, implementation, and transfer of the core analysis component of CodeSurfer/x86. (cf. my doctoral thesis.)

EDUCATION

- Aug 2007 **Ph.D. (Computer Science)**
University of Wisconsin, Madison, USA
Advisor: Prof. Thomas. W. Reps
Thesis: "WYSINWYX: **W**hat **Y**ou **S**ee **I**s **N**ot **W**hat **Y**ou **eX**ecute"
(Static Analysis of Executables)
- May 2003 **M.S. (Computer Science)**
University of Wisconsin, Madison, USA
- May 2001 **B.E. (Computer Science and Engineering)**
College of Engineering, Guindy
Anna University, Chennai, India

PUBLICATIONS

Machine Learning and Programming Languages

While at Google, I have dabbled with building ML models or embeddings for code.

- Aditya Kanade, Petros Maniatis, [Gogul Balakrishnan](#), Kensen Shi, **Learning and Evaluating Contextual Embedding of Source Code**, In Proc. of International Conference on Machine Learning (ICML), 2020.

Analysis of C++ programs

At NEC labs, I led an effort to build tools for analyzing and testing C/C++ programs. In particular, I worked on an automated test-case generation tool for C/C++ programs, and a source-to-source translation tool that models inheritance and exceptions in such a way that makes it easier to statically analyze C++ programs. It has been incorporated into NEC's internal bug finding tool called Varvel.

- Xiao, X., Balakrishnan, G., Ivancic, F., Maeda, N., Gupta, A., and Chhetri, D. **ARC++: Effective Typestate and Lifetime Dependency Analysis**, In Proc. Intl. Symposium on Software Testing and Analysis (ISSTA), 2014 (to appear).
- Garg, P., Ivancic, F., Balakrishnan, G., Maeda, N., and Gupta, A., **Feedback-directed unit test generation for C/C++ using concolic execution**, In Proc. Intl. Conf. on Software Engg. (ICSE), May 2013.
- Yang, J., Balakrishnan, G., Maeda, N., Ivancic, F., Gupta, A., Sinha, N., Sankaranarayanan, S., Sharma, N., **Object Model Construction for Inheritance in C++ and Its Applications to Program Analysis**, In Proc. Compiler Construction (CC), March 2012.
- Prabhu, P., Maeda, N., Balakrishnan, G., Ivancic, F., Gupta, A., **Interprocedural Exception Analysis for C++**, In Proc. European Conf. on Object Oriented Prog. (ECOOP), 2011.
- Balakrishnan, G., Maeda, N., Sankaranarayanan, S., Ivancic, F., Gupta, A., Pothengil, R., **Modeling and Analyzing the Interaction of C and C++ Strings**, In Formal Verification of Object-Oriented Software (FoVeOOS), 2011.

Static Analysis and Model Checking

At NEC Labs, I was involved in building NEC's internal bug finding tool called Varvel that is based on static analysis and model checking technologies. Varvel is used internally at NEC Japan on millions of lines of C/C++ code on a daily basis.

- Ghorbal, K., Ivancic, F., Balakrishnan, G., Maeda, N., Gupta, A., **Donut Domains: Efficient Non-convex Domains for Abstract Interpretation**, In Proc. Conf. on Verification Model Checking and Abstract Interpretation (VMCAI), Jan 2012.
- Ivancic, F., Balakrishnan, G., Gupta, A., Sankaranarayanan, S., Maeda, N., Tokuoka, H., Imoto, T., Miyazaki, Y., **DC2: A framework for scalable, scope-bounded software verification**, In Automated Software Engg. (ASE), 2011.
- Ganai, M., Arora, N., Wang, C., Gupta, A., Balakrishnan, G., **BEST: A symbolic testing tool for predicting multi-threaded program failures**, (Tool Paper), In Automated Software Engg. (ASE), 2011.
- Balakrishnan, G., Ganai, M., Gupta, A., Ivancic, F., Kahlon, V., Li, W., Maeda, N., Papakonstantinou, N., Sankaranarayanan, S., Sinha, N., Wang, C., **Scalable and precise program analysis at NEC**, In Formal Methods in Computer-Aided Design (FMCAD), 2010.
- Balakrishnan, G., Sankaranarayanan, S., Ivancic, F., and Gupta, A., **Refining the Control Structure of Loops using Static Analysis**, In Proc. Intl. Conf. on Embedded Software (EMSOFT), Oct 2009.
- Balakrishnan, G., Ganai, Malay., **PED: Proof-guided Error Diagnosis by Triangulation of Program Error Causes**, In Proc. Software Engineering and Formal Methods (SEFM), Nov 2008.
- Balakrishnan, G., Sankaranarayanan, S., Ivancic, F., Wei, O., and Gupta, A., **SLR: Path-Sensitive Analysis Through Infeasible-Path Detection and Syntactic-Language Refinement**. In Proc. Static Analysis Symposium (SAS), Jul 2008 .

Analysis of executables

For my Ph.D. dissertation, I developed static-analysis algorithms to address the problem of analyzing executables without debugging information.

- Balakrishnan, G., **WYSINWYX: What You See Is Not What You eXecute**, Ph.D. dissertation (TR-1603), Computer Sciences Department, University of Wisconsin, Madison, WI, August 2007. **[Ph.D. thesis.]**
- Balakrishnan, G., and Reps, T., **WYSINWYX: What You See Is Not What You eXecute**. In *Proc. of ACM Transactions on Programming Languages and Systems (TOPLAS)*, Vol. 32, Nr. 6, Aug 2010. **[Journal Publication.]**
- Balakrishnan, G. and Reps, T., **Analyzing Stripped Device-Driver Executables**, In *Proc. 14th Int Conf. on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, Mar. 2008.
- Balakrishnan, G. and Reps, T., **Recency-Abstraction for Heap-Allocated Storage**. In *Proc. 13th Static Analysis Symposium (SAS)*, pages 221-239, Aug. 2006.
- Balakrishnan, G., Reps, T., Melski, D., and Teitelbaum, T., **WYSINWYX: What You See Is Not What You eXecute**. In *Proc. IFIP Working Conference on Verified Software: Theories, Tools, Experiments (VSTTE)*, Oct. 2005.
- Balakrishnan, G., Reps, T., Kidd, N., Lal, A., Lim, J., Melski, D., Gruian, R., Yong, S., Chen, C.-H., and Teitelbaum, T., **Model checking x86 executables with CodeSurfer/x86 and WPDS++**, (Tool Demonstration Paper). In *Proc. 17th Int. Conf. on Computer Aided Verification (CAV)*, pages 158-163, Jul. 2005.
- Lal, A., Reps, T., and Balakrishnan, G., **Extended Weighted Pushdown Systems**. In *Proc. 17th Int. Conf. on Computer Aided Verification (CAV)*, pages 434-448, Jul. 2005.
- Balakrishnan, G., Gruian, R., Reps, T. and Teitelbaum, T., **CodeSurfer/x86 - A Platform for Analyzing x86 Executables**, (Tool Demonstration Paper). In *Proc. 14th Int. Conf. on Compiler Construction (CC)*, pages 250-254, Apr. 2005.
- Balakrishnan, G. and Reps, T., **Analyzing Memory Accesses in x86 Executables**. In *Proc. 13th Int. Conf. on Compiler Construction (CC)*, pages 5 - 23, Apr. 2004. **(Won the EAPLS best paper award at ETAPS 2004.)**

Testing of large-scale systems

At NEC Labs, I was also involved in building a framework for testing large-scale distributed systems such as HBase, Zookeeper, and so on.

- Joshi, P., Ganai, M., Balakrishnan, G., Gupta, A., Papakonstantinou, N., **SETSUDO: Perturbation-based Testing Framework for Scalable Distributed Systems**, In *Conf. on Timely Results in Operating Systems (TRIOS)*, November 2013.

Invited Papers

- Reps, T., Lim, J., Thakur, A., Balakrishnan, G., and Lal, A., **There's plenty of room at the bottom: Analyzing and verifying machine code (Invited Tutorial)**, In *Proc. Computer Aided Verification (CAV)*, July 2010.
- Reps, T. and Balakrishnan, G., **Improved Memory-Access Analysis for x86 Executables**, In *Proc. Conf. on Compiler Construction (CC)*, Budapest, Hungary, Mar 29-Apr 6, 2008. **(Paper accompanying a unifying invited talk at ETAPS08.)**
- Balakrishnan, G. and Reps, T., **DIVINE: Discovering Variables IN eXecutables**. In *Proc. Conf. on Verification Model Checking and Abstract Interpretation (VMCAI)*, Nice, France, Jan

14-17, 2007.

- Reps, T., Balakrishnan, G., and Lim, J., **Intermediate-representation recovery from low-level code**. In *Proc. Workshop on Partial Evaluation and Program Manipulation (PEPM)*, Charleston, SC, Jan. 9-10, 2006.
- Reps, T., Balakrishnan, G., Lim, J., and Teitelbaum, T., **A Next-Generation Platform for Analyzing Executables**. In *Proc. 3rd Asian Symposium on Programming Languages and Systems*, Tsukuba, Japan, Nov. 3-5, 2005.

Book Chapters

- **Malware Detection**, volume 27 of *Advances in Information Security*. Springer-Verlag, Oct. 2006. [Editors: M. Christodorescu, S. Jha, D. Maughan, D. Song, and C. Wang]
- **Analysis of COTS for Security Vulnerability Remediation**, G. Balakrishnan, M. Christodorescu, V. Ganapathy, J. T. Giffin, S. Rubin, H. Wang, S. Jha, B. P. Miller, and T. Reps. In *Information Security Research: New Methods for Protecting against Cyber Threats*, C. Wang, S. King, R. Wachter, R. Herklotz, C. Arney, G. Toth, D. Hislop, S. Heise, and T. Combs, editors, Wiley Publishing Inc., July 2007.

Reprinted in Collections

- Reps, T., Balakrishnan, G., Lim, J., and Teitelbaum, T., **A Next-Generation Platform for Analyzing Executables**. In *Proc. of the ARO-DHS Malware Detection Workshop*, (Arlington, VA, Aug 10-11, 2005), *Advances in Information Security* series, Springer-Verlag. (Reprinted from *Proc. 3rd Asian Symposium on Programming Languages and Systems*, Tsukuba, Japan, Nov. 3-5, 2005.)

Other Publications and Reports

- Balakrishnan, G., Reps, T., Kidd, N., Lal, A., Lim, J., Melski, D., Gruian, R., Yong, S., Chen, C.-H., and Teitelbaum, T., **Model Checking x86 Executables with CodeSurfer/x86 and WPDS++**. In *Proc. Workshop on the Evaluation of Software Defect Detection Tools*, June 2005. [Co-located with PLDI05.]

TECHNICAL PRESENTATIONS

Conference and External Seminar Presentations

- “Scalable Security Analysis of Android Apps,” University of Wisconsin-Madison, September 2016.
- “Scalable Security Analysis of Android Apps,” Workshop on Formal Methods and Security (co-located with PLDI), June 2016.
- “Scalable Program Analysis at NEC,” Dagstuhl Seminar on Pointer Analysis (13162), April 2013.
- “Object Model Construction for Inheritance in C++ and Its Applications to Program Analysis,” Intl. Conf. on Compiler Construction (CC), Tallinn, Estonia, March 2012.
- “Tale of two tools: BEST & GIRA,” Dagstuhl Seminar on Executable Analysis (12051), Jan 2012.
- “DC2: A framework for scalable, scope-bounded software verification,” Intl. Conf. on Automated Software Engg. (ASE), Lawrence, Kansas, November 2011.

- “Refining the control structure of loops using static analysis,” Intl. Conf. on Embedded Software (EMSOFT), Grenoble, France, Oct 2009.
- “SLR: Path-Sensitive Analysis Through Infeasible-Path Detection and Syntactic-Language Refinement,” Static Analysis Symposium (SAS), Valencia, Spain, June 2008.
- “Analyzing Stripped Device Driver Executables,” Intl. Conf. on Tools and Algorithms for the Construction and Analysis of Systems (TACAS), Budapest, Hungary, March 2008.
- “WYSINWYX: What You See Is Not What You eXecute,”
 - ROSAEC Center, Seoul National University, June 02, 2009.
 - IMDEA, Madrid, Spain, June 2008.
 - North East Verification (NEVER) Seminar, U Penn, Philadelphia, PA, Nov 2007.
 - Ph.D Defense at UW Madison, Aug 2, 2007.
 - IBM Research, Hawthorne, NY, May 2007.
 - GrammaTech, Inc., Ithaca, NY, April 2007.
 - NEC labs, Princeton, NJ, March 2007.
 - Microsoft Research, Bangalore, India, August 2006.
 - IBM India Research Lab, New Delhi, India, August 2006.
- “Recency-Abstraction for Heap-Allocated Storage,” Static Analysis Symposium (SAS), Seoul, South Korea, August 2006.
- “Model Checking x86 Executables with CodeSurfer/x86 and WPDS++,” Conference on Computer Aided Verification (CAV), Scotland, UK, July 2005.
- “CodeSurfer/x86 - A Platform for Analyzing x86 Executables” Conference on Compiler Construction (CC), Scotland, UK, April 2005.
- “Analyzing Memory Accesses in x86 Executables,”
 - CSA Dept., Indian Institute of Science (IISc), Bangalore, India, August 2004.
 - CSE Dept., Anna University, Chennai, India, August 2004.
 - CS Dept., Indian Institute of Technology (IIT), Chennai, India, August 2004.
 - Conference on Compiler Construction (CC), Barcelona, Spain, April 2004.
 - IBM Watson Research Center, Hawthorne, NY, USA, December 2003.
 - Dagstuhl Seminar on Language-based Security, Germany, October 2003.

Presentations to Funding Agencies

- “Identifying Variables in x86 Executables,” ONR MURI Workshop, Arlington, Virginia, February 14, 2005.
- “Analyzing Memory Accesses in x86 Executables,”
 - ONR MURI Workshop, Pittsburgh, Pennsylvania, July 23, 2003.
 - ONR MURI Workshop, Williamsburg, Virginia, January 28, 2003.
 - ONR MURI Workshop, Harpers Ferry, West Virginia, July 12, 2002.
 - ONR MURI Workshop, Washington, DC, January 15, 2002.

PROFESSIONAL ACTIVITIES

- Program Committee Member, Static Analysis Symposium (SAS) 2023.
- Program Committee Member, Static Analysis Symposium (SAS) 2021.
- Program Committee Member, Intl. Conf. on Computer Aided Verification (CAV), 2020.
- Program Committee Member, 14th International Conference on Information Systems Security (ICISS), 2018.
- Program Committee Member, Intl. Symposium on Automated Technology for Verification and Analysis (ISSTA), 2018.
- Program Committee Member, Intl. Conf. on Computer Aided Verification (CAV), July 2014.
- Program Committee Member, Intl. SPIN Symposium on Model Checking of Software, July 2013.
- Program Committee Member, Intl. Conf. on Compiler Construction (CC) March 2013.
- Program Committee, Workshop on Tools for Automatic Program Analysis (TAPAS) 2012.
- Reviewer (as an external referee) for POPL 10, SAS 09, VMCAI 09, TACAS 09, ATVA 08, Computer Aided Verification (CAV) 2008, Programming Language Design and Implementation (PLDI) 2008, Tools and Algorithms for the Construction and Analysis of Systems (TACAS) 2007, CAV 2006.

PATENTS

- Control structure refinement of loops using static analysis (US Patent 8,522,226)
- Program verification through symbolic enumeration of control path programs (US Patent 8,402,440)
- Path-sensitive analysis through infeasible-path detection and syntactic language refinement (US Patent 8,365,152)

REFERENCES

Available on request.