Vaggelis Atlidakis

(+30) 698 992 9982 | v.atlidakis@gmail.com | https://vatlidak.github.io/ | vatlidak

Areas

Operating Systems, Software Testing, Program Analysis, Machine Learning

Education

Ph.D. in Computer Science — *New York, USA*
COLUMBIA UNIVERSITY — *2013 - 2020*
- Dissertation: "Structure and Feedback in Cloud Service API Fuzzing"
- Supervisors: Roxana Geambasu & Jason Nieh

M.Phil. in Computer Science — *New York, USA*
COLUMBIA UNIVERSITY — *2018*
- Candidacy: "Automated Software Testing: Approaches and Tradeoffs"
- Supervisors: Roxana Geambasu & Jason Nieh

M.Sc. in Computer Science (GPA 3.85/4) — *New York, USA*
COLUMBIA UNIVERSITY — *2013 - 2016*

B.Sc. in Informatics & Telecommunications (GPA 9/10) — *Athens, Greece*
UNIVERSITY OF ATHENS (NKUA) — *2005 - 2011*
- Thesis: "Efficient Data Pipelining in the Bittorrent Protocol"
- Supervisors: Alex Delis & Mema Roussopoulos

Professional Appointments

Lacework — *London, UK*
FULL-TIME SOFTWARE ENGINEER — *February 2023 - Present*
- Code Security Team

Lacework — *San Jose, CA*
VISITING SOFTWARE ENGINEER — *August 2022 - November 2022*
- Confidential

Brown University — *Providence, RI*
POSTDOCTORAL RESEARCH ASSOCIATE — *January 2021 - January 2023*
- Brown Secure Systems Lab
- Mentor: Vasileios P. Kemerlis

Microsoft Research — *Redmond, WA*
RESEARCH INTERN — *June-August 2017 and 2018*
- Microsoft Security Risk Detection: "Testing of Cloud Services"
- Mentors: Patrice Godefroid & Marina Polishchuk

Microsoft Research — *New York, NY*
CONTRACTOR — *October 2017 - December 2017*
- Microsoft Security Risk Detection: "Testing of Cloud Services"
- Mentors: Patrice Godefroid & Marina Polishchuk

European Organization for Nuclear Research (CERN) — *Geneva, Switzerland*
RESEARCH ASSOCIATE — *September 2012 - September 2013*
- IT Department: "High Availability DNS Load Balancing for the OpenStack, Private Cloud of CERN"
- Supervisor: Ignacio Reguero

Research Experience

Automated Vulnerability Detection in Deep Learning Frameworks. Deep Learning (DL) frameworks involve code that spans different low- and high-level languages. In this context, it is no surprise that, due to missing sanity checks and mismatched security assumptions, untrusted inputs may transfer through the stack of APIs and reach memory-unsafe code. Given that, how can we automatically detect such code defects and report them to framework developers in an actionable manner? To address this question, we introduced a two-fold bottom-up approach, implemented in our IvySyn framework [C4]. IvySyn leverages the statically-typed nature of native APIs in order to automatically perform type-aware mutation-based fuzzing on low-level kernels. Next, given a set of offending inputs that trigger memory safety (and fatal runtime) errors in low-level, native DL (C/C++) code, it automatically synthesizes code snippets in managed languages (e.g., Python), which propagate offending input through high(er)-level APIs. Such code snippets essentially act as "Proof of Vulnerability" (PoV), as they demonstrate the existence of bugs in native, C/C++ code that attackers can target (and potentially abuse) via high-level APIs. IvySyn has already helped TensorFlow and PyTorch DL framework developers identify and fix numerous previously-unknown security vulnerabilities, implicitly and explicitly affecting millions, if not billions, of users worldwide.

Testing of Cloud Services. Today, most cloud services—like those running on Amazon Web Services (AWS) and Microsoft Azure—are programmatically accessed through REST APIs. Yet, tools for automatically testing cloud services through their REST APIs and checking whether those services are reliable and secure are still in their infancy. Along with my collaborators Patrice Godefroid and Marina Polishchuk, from Microsoft Research, we did work in the area of cloud service API testing. We introduced the idea of stateful REST API fuzzing and built RESTler: the first Stateful REST API Fuzzer [C7]. RESTler has found hundreds of bugs in several production-scale, open-source and proprietary cloud services through their APIs. Next, we showed how RESTler can be extended with active checkers that automatically test and detect violations of security rules that capture desirable properties of REST APIs in a modular and efficient way [C6]. Using security checkers, RESTler has found new bugs in several deployed, production Azure and Office365 cloud services. In my most recent work, which was a collaborative effort among my Microsoft Research mentors and my Ph.D. advisors, we built Pythia, a new fuzzer that augments stateful REST API fuzzing with coverage-guided feedback and learning-based mutations. Our experimental evaluation showed that Pythia can report previously-unknown errors on production-scale cloud services that were beyond the reach of baseline stateful REST API fuzzing [A1].

POSIX Abstractions in Modern OSes. We conducted a broad measurement study that shed light into a number of questions regarding the use of POSIX abstractions by modern application running on Android, OSX, and Ubuntu [C10, J1]. We found out that modern applications rely on abstractions not supported by the POSIX API and therefore custom, user-space libraries, providing all necessary abstractions, are being implemented on top of POSIX. This layering causes mismatches, inefficiencies, and even security risks. For example, we observed that new abstractions heavily use POSIX extension APIs (such as ioctl) to implement their functionality, suggesting that POSIX lacks appropriate abstractions for modern workloads. Extension APIs are problematic because their invocations cannot be mediated by the OS, putting pressure on user-space libraries and kernel device drivers to implement correct and coherent protections of these invocations. Our findings have broad implications related to the future of POSIX-compliant OS portability, which the systems research community and the relevant standard bodies will likely need to address in the near future.

Security of Machine Learning Systems in Adversarial Settings. Deep Neural Networks (DNNs) perform exceptionally well on many machine learning tasks, including safety- and security-sensitive applications, such as self-driving cars, malware classification, face recognition, and critical infrastructure. Robustness against malicious behavior is important in many of these applications. Yet, in recent years it has become clear that DNNs are vulnerable to a broad range of attacks, including adversarial examples, where the adversary finds small perturbations to correctly classified inputs that cause a DNN to produce an erroneous prediction. Adversarial examples pose a serious threat to security-critical applications. Based on a novel connection between robustness to adversarial examples and differential privacy we proposed PixelDP: the first certifiably robust defense against adversarial examples that scales to large, real-world DNNs and datasets (e.g., Google's Inception network for ImageNet) and applies broadly to arbitrary model types [C8]. PixelDP also enables a firewall-like security architecture, where a small model is prepended to an existing, already trained one to make it more robust. Such an architecture is common in traditional software systems but unique in ML workloads.

Testing Tools for Data-Driven Applications A key aspect characterizing modern applications is their increased reliance on data and data-driven decision-making. While often beneficial, this practice can have subtle detrimental consequences, such as discriminatory or racially offensive effects. We argued that such effects are bugs that should be tested for and debugged in a manner similar to functionality, reliability, and performance bugs. To this end, we developed FairTest: a testing toolkit for data-driven applications that identifies unwarranted association between application outputs and user subpopulations, including sensitive groups (e.g., minorities defined by race or gender) [C9].

Publications

Preprints

A1 **Vaggelis Atlidakis**, Roxana Geambasu, Patrice Godefroid, Marina Polishchuk, Baishakhi Ray. *"Pythia: Grammar-Based Fuzzing of REST APIs with Coverage-guided Feedback and Learning-based Mutations."*

Conference Publications

C1 Alexander Gaidis, **Vaggelis Atlidakis**, and Vasileios P. Kemerlis. *"SysXCHG: Refining Privilege with Adaptive System Call Filters."* In Proceedings of the 30th ACM Conference on Computer and Communications Security (CCS '23).

C2 Alexander Gaidis, Joao Moreira, Ke Sun, Alyssa Milburn, **Vaggelis Atlidakis**, and Vasileios P. Kemerlis. *"FineIBT: Fine-grain Control-flow Enforcement with Indirect Branch Tracking."* In Proceedings of the 26th International Symposium on Research in Attacks, Intrusions and Defenses (RAID '23).

C3 Di Jin, **Vaggelis Atlidakis**, and Vasileios P. Kemerlis. *"EPF: Evil Packet Filter."* In Proceedings of the 2023 USENIX Annual Technical Conference (USENIX ATC '23).

C4 Neophytos Christou, Di Jin, **Vaggelis Atlidakis**, Baishakhi Ray, and Vasileios P. Kemerlis. *"IvySyn: Automated Vulnerability Discovery in Deep Learning Frameworks."* In Proceedings of the 32nd USENIX Security Symposium (USENIX Sec '23).

C5 Thodoris Sotiropoulos, Stefanos Chaliasos, **Vaggelis Atlidakis**, Dimitris Mitropoulos and Diomidis Spinellis. *"Data-Oriented Differential Testing of Object-Relational Mapping Systems."* In Proceedings of the 43rd International Conference on Software Engineering (ICSE '21).

C6 **Vaggelis Atlidakis**, Patrice Godefroid, and Marina Polishchuk. *"Checking Security Properties of Cloud Service REST APIs."* In Proceedings of the 13th IEEE International Conference on Software Testing, Verification and Validation (ICST'20).

C7 **Vaggelis Atlidakis**, Patrice Godefroid, and Marina Polishchuk. *"RESTler: Stateful REST API Fuzzing."* In Proceedings of the 41st International Conference on Software Engineering (ICSE'19).

C8 Mathias Lecuyer, **Vaggelis Atlidakis**, Roxana Geambasu, Daniel Hsu, and Suman Jana. *"Certified Robustness to Adversarial Examples with Differential Privacy."* In Proceedings of the 40th IEEE Symposium on Security and Privacy (S&P'19).

C9 Florian Tramer, **Vaggelis Atlidakis**, Roxana Geambasu, Daniel Hsu, Jean-Pierre Hubaux, Mathias Humbert, Ari Juels, and Huang Lin. *"Discovering Unwarranted Associations in Data-Driven Applications with the FairTest Testing Toolkit."* In Proceedings of the 2nd European Symposium on Security and Privacy (Euro S&P '17).

C10 **Vaggelis Atlidakis**, Jeremy Andrus, Roxana Geambasu, Dimitris Mitropoulos, and Jason Nieh. *"POSIX Abstractions in Modern Operating Systems: The Old, the New, and the Missing."* In the 11th European Conference on Computer Systems (EuroSys '16).

C11 **Vaggelis Atlidakis**, Mema Roussopoulos, and Alex Delis. *"Changing the Unchoking Policy for an Enhanced Bittorrent."* In Proceedings of International European Conference on Parallel and Distributed Computing (EuroPar '12).

Journal & Magazine Publications

J1 **Vaggelis Atlidakis**, Jeremy Andrus, Roxana Geambasu, Dimitris Mitropoulos, and Jason Nieh. *"POSIX has become outdated."* USENIX ;login: Magazine, 41(3), Fall 2016.

J2 **Vaggelis Atlidakis**, Mema Roussopoulos, and Alex Delis. *"EnhancedBit: Unleashing the Potential of the Unchoking Policy in the BitTorrent Protocol."* Journal of Parallel and Distributed Computing (JPDC), Vol. 74, Issue 1, pp. 1959-1970, January 2014.

Honors & Awards

2021 - 2023	**Fellowship,** Computing Innovation Fellowship (CIFellow 2020), Brown University	*Providence, USA*
2013 - 2020	**Fellowship,** Graduate Research Assistantship (GRA), Columbia University	*New York, USA*
2015	**Scholarship (for Ph.D. studies),** Gerondelis Foundation	*New York, USA*
2013	**Scholarship (for Ph.D. studies),** Computer Science Chair's Distinguished Award	*New York, USA*
2005	**Scholarship (for B.S. studies),** National Scholarship Foundation of Greece	*Athens, Greece*

Patents

P1 **Vaggelis Atlidakis**, Patrice Godefroid, and Marina Polishchuk *"Automatic intelligent cloud service testing tool."* US Patent 20190370152A1, 2019.