Theofilos Petsios

(+1) 6462587246 | theofilos.pe@gmail.com | www.cs.columbia.edu/~theofilos | nettrino | @theofilospe

Research Interests

I am interested in all aspects of systems and software security, with a focus on application security, binary analysis, and privacy.

Education

Ph.D. in Computer Science — *New York, USA*
Columbia University — *2012 - Present*
- Advisor: Prof. Angelos D. Keromytis

M.Phil. in Computer Science — *New York, USA*
Columbia University — *2015*

M.Sc. in Computer Science (GPA 3.9 / 4) — *New York, USA*
Columbia University — *2012 - 2014*

B.S. in Electrical Engineering & Computer Science (GPA 8.2/10) — *Athens, Greece*
National Technical University of Athens (NTUA) — *2005 - 2011*
- Thesis: "Term suggestion mechanisms for Scientific Database Systems". Advisor: Prof. Timos Sellis

Research Experience

Software Testing

Automatically discovering bugs in security-sensitive code is hard. In my recent work, I designed and developed Nezha [C2], an evolutionary-based, input-format-agnostic fuzzer which focuses on semantic bugs. This fuzzer utilizes novel black-box and gray-box input generation mechanisms and has found multiple bugs in security critical applications like SSL libraries (e.g., LibreSSL, GnuTLS, WolfSSL) and anti-virus software, reporting semantic bugs missed by state-of-the-art fuzzers like AFL and libFuzzer.

When auditing an application for bugs and vulnerabilities, it is vital to determine what constitutes a true bug, rather than a developer-intended violation of a specification. In previous work, I co-developed IntFlow [C6], a compiler extension which uses taint analysis to report integer errors, taking into account common developer practices to reduce false positives. IntFlow achieves 89% reduction in false positives compared to the Clang-based Integer Overflow Checker (IOC).

My latest research in this field attempts to enhance static analysis tools to further improve fuzzing, as well as to detect new types of bugs such as resource exhaustion bugs [C1].

Software Hardening

Applications and Operating Systems can be separated into discreet parts, each facing different types of threats. In the past, I have worked on software hardening defenses, aiming at efficiently addressing such security issues. One such example of my work is DynaGuard [C4], which secures canary-based protections against brute-force attacks. DynaGuard comes in two flavors, a compiler-based, inducing a 1.2% overhead, and a version implemented using Dynamic Binary Instrumentation (DBI), which can applied as-is to black-box binaries.

Another defense in this space which I have co-developed is, kR^X a practical and comprehensive kernel hardening scheme that protects the Linux kernel from code-reuse attacks [C3]. The above scheme combines an execute-only memory principle with code diversification defenses, and can benefit from hardware support (e.g., MPX on modern Intel CPUs) to optimize performance.

Privacy

Modern users face a multitude of privacy threats. In my research, I seek to evaluate the privacy guarantees that modern services provide, as well as to develop tools that enable users to use these services in a transparent, privacy-preserving manner. In past work, I co-performed a formal analysis of the defenses that major Location Based Services (LBSes)[C5,J1] deploy against location disclosure attacks. I also co-developed novel attacks to bypass these defenses, across all major services, like Facebook, Foursquare and others, and was the primary author of an open-source Python framework for auditing Location Based Services with respect to their location privacy guarantees.

With respect to Web transparency, in previous work I co-developed XRay [C7], a framework that utilizes differential correlation to determine which particular actions of users (emails, YouTube views and clicks *etc.*) are linked with their targeting by online services.

My latest work in this field involves developing new methods to fingerprint users, using behavioral patterns derived form their online presence, and implementing viable, real-world defenses for the users' protection against such attacks.

Web Security

In 2014 I developed a compiler extension offering SQL injection protection to C/C++ applications. Currently, I am co-developing a symbolic execution engine for PHP, which operates at the PHP interpreter level and fully supports database operations and loose types. Finally, some of my current research in the field involves examining the discrepancies in the way the various entities of the Web ecosystem handle content, developing novel attacks due to content handling confusion and suggesting real-world defenses for the presented attacks.

Professional Appointments

Microsoft Research — *Cambridge, UK*

Research Intern — *October 2017 - December 2017*

- Systems and Networking Group

Trail of Bits — *New York, NY, USA*

Research Intern — *May 2017 - August 2017*

- Worked on the Manticore symbolic execution engine, adding support for symbolic execution of Binary Ninja IL, and participated in a number of security audits.

Symantec Corporation — *Herndon, VA, USA*

Research Intern — *June 2015 - August 2015*

- Developed an ELF binary rewriting library for ARM. Designed and implemented interfaces supporting binary rewriting, injection of anti-debugging features, injection of new modules into an existing binary, as well as detection of packers and backdoors.

Greek Army (Obligatory Service) — *Athens, Greece*

IT Support & Website Administrator — *August 2011 - June 2012*

- Responsible for the backup & maintenance of the servers hosting the website of the Greek Ministry of National Defence.

Cybex S.A. — *Athens, Greece*

System Administrator & Web Developer — *January 2011 - July 2011*

- Responsible for the administration of the server infrastructure, as well as the hosting and domain name registration services of the company.

Talks & Media Coverage

Invited Talks

August 2017 — Empire Hacking NYC: "Extending Manticore with Binary Ninja" — *MongoDB, New York, USA*

October 2016 — Empire Hacking NYC: "Differential Fuzzing with LLVM's LibFuzzer" — *TwoSigma, New York, USA*

Media Coverage

August 2014 — NYT Bits: "XRay: A New Tool for Tracking the Use of Personal Data on the Web" — *The New York Times*

Conference and Workshop Talks

May 2017 — "Nezha : Efficient Domain - Independent Differential Testing" — *S&P, San Jose, USA*

December 2015 — "DynaGuard : Armoring Canary-Based Protections against Brute-force Attacks" — *ACSAC, Los Angeles, USA*

Skills

Programming C/C++, Python, LaTeX

Technologies LLVM & GCC internals, Intel PIN

Languages Greek (*Native*) English (*Proficient*) French, Spanish (*Elementary proficiency - ILR scale 1+*)

Honors & Awards

2012-2017 **Fellowship**, Graduate Research Assistantship (GRA), Columbia University — *New York, USA*

2014 **Bug Bounty Grant**, Facebook — *New York, USA*

2014 **Scholarship (for Ph.D. studies)**, Gerondelis Foundation — *New York, USA*

2005 **Scholarship (for B.S. studies)**, Eurobank EFG — *Athens, Greece*

Teaching and Mentorship

Instructor, Introduction to Programming in C *New York, USA*

Columbia University *Summer 2014*

- Designed and taught a three week summer intensive course for the School of Continuing Education at Columbia University (Students: 16)

Teaching Assistant *New York, USA*

Columbia University *Summer 2013 - Fall 2016*

- Fall 2016: Head Teaching Assistant (TA) for Network Security (Graduate level. Instr. Debbie Cook. Students: 34)
- Spring 2015: TA for Network Security (Graduate level. Instr. Debbie Cook. Students 33)
- Fall 2013: Head TA for Data Structures & Algorithms (Undergraduate level. Instr.: Shlomo Hershkop. Students: 70)
- Spring 2013: Head TA for Advanced Programming (Undergraduate level. Instr.: Shlomo Hershkop. Students: 14)

Student Mentor *New York, USA*

Columbia University *2016*

- Fall 2016: Jason Zhao (undergraduate student). Project: Guided fuzzing for resource exhaustion bugs.
- Spring 2016: Benjamin Low (undergraduate student). Project: Towards a taxonomy of the security properties of major OSes.

Service

External Reviewer

USENIX USENIX Security Symposium: 2017
CCS ACM Conference on Computer and Communications Security: 2013, 2014
IWSEC International Workshop on Security: 2014
MTD ACM Workshop on Moving Target Defense: 2015
IET IET Information Security: 2014

Publications

Conference Publications

C1 **Theofilos Petsios**, Jason Zhao, Angelos D. Keromytis, and Suman Jana. *"SlowFuzz : Automated Domain-Independent Detection of Algorithmic Complexity Vulnerabilities. "*. In Proceedings of the 24th ACM Conference on Computer and Communications Security (CCS), Dallas, TX, November 2017.

C2 **Theofilos Petsios**, Adrian Tang, Salvatore Stolfo, Angelos D. Keromytis, and Suman Jana. *"Nezha : Efficient Domain - Independent Differential Testing"*. In 38th IEEE Symposium on Security and Privacy (S&P), San Jose, CA, May 2017.

C3 Marios Pomonis,**Theofilos Petsios**, Angelos D. Keromytis, Michalis Polychronakis, Vasileios P. Kemerlis. *"kR^X : Comprehensive Kernel Protection against Just-In-Time Code Reuse"*. In Proceedings of the 12th European Conference on Computer Systems (EuroSys), April 2017.

C4 **Theofilos Petsios**, Vasileios P. Kemerlis, Michalis Polychronakis, Angelos D. Keromytis. *"DynaGuard: Armoring Canary-based Protections against Brute-force Attacks"*. In Proceedings of the 31th Annual Computer Security Applications Conference (ACSAC), December 2015.

C5 Iasonas Polakis, George Argyros, **Theofilos Petsios**, Suphannee Sivakorn, Angelos D. Keromytis *"Where's Wally? Precise User Discovery Attacks in Location Proximity Services"* . In Proceedings of the 22nd ACM Conference on Computer and Communications Security (CCS), October 2015.

C6 Marios Pomonis, **Theofilos Petsios**, Kangkook Jee, Michalis Polychronakis, and Angelos D. Keromytis. *"IntFlow: Improving the Accuracy of Arithmetic Error Detection Using Information Flow Tracking"*. In Proceedings of the 30th Annual Computer Security Applications Conference (ACSAC), December 2014.

C7 M. Lecuyer, G. Ducoffe, F. Lan, A. Papancea, **T. Petsios**, R. Spahn, A. Chaintreau, and R. Geambasu, *"XRay: Enhancing the Web's Transparency with Differential Correlation."*, in Proceedings of the USENIX Security Symposium, August 2014.

Journal Publications

J1 George Argyros, **Theofilos Petsios**, Suphannee Sivakorn, Angelos D. Keromytis, and Jason Polakis. *"Evaluating the Privacy Guarantees of Location Proximity Services"*. In ACM Transactions on Privacy and Security (TOPS) (formerly known as TISSEC).