Research Statement

Maria Apostolaki, ETH Zurich

As a security and network researcher, my goal is to design and build networked systems that are secure, reliable, and performant. This goal is fascinatingly challenging because real-world systems consist of a complex series of *interdependent* components and layers. Thus, these systems rely not only on their application-layer components but also on the underlying network layer, the used hardware, and the incentives of the involved parties. Still, researchers often focus on a single layer or component *in isolation*, making *assumptions* about the rest. To avoid the dire consequences of broken assumptions on a system's properties, I proactively challenge them and build systems that have *provable* behavior under all *realistic* settings.

To that end, I go beyond traditional disciplinary barriers and balance between theory and practice. My research draws from networking, security, and blockchain. I look at networked systems from a *cross-layer* perspective, observing the interactions across layers and network components. Doing so enables me to uncover *interdependencies* that leave systems susceptible to exploits and inefficiencies. Having a clear view of potential interdependencies, I design and build systems with *provable* properties that are also *deployable* considering economic incentives, Internet policies, and available hardware.

I have used this approach to secure blockchain applications from networking adversaries. To do so, I first uncovered critical vulnerabilities in Bitcoin, rooted in its *implicit-yet-dangerous* assumptions about the underlying network. I then combined *domain-specific* knowledge from both Bitcoin and networking to design a *practical* system to protect Bitcoin-like protocols working atop the Internet (§1). Further, I used this approach to improve Internet-based communications in terms of reliability and performance by modifying routing decisions while respecting both the available networking hardware and economic incentives of Internet providers (§2). Finally, I demonstrated that the inefficiencies of the state-of-the-art buffer-management algorithm impose harmful dependencies on seemingly independent traffic. I then designed a novel algorithm that offers *provable* isolation and performance guarantees while being *deployable* in today's hardware (§3).

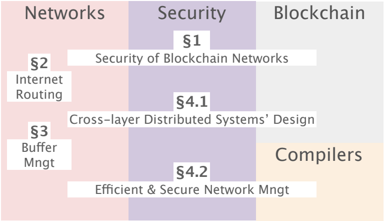
Networks
Security
Blockchain
§1
Security of Blockchain Networks
§2
Internet Routing
§4.1
Cross-layer Distributed Systems' Design
§3
Buffer Mngt
Compilers
§4.2
Efficient & Secure Network Mngt

Figure 1: My prior (§1, §2, §3) and future (§4) research draw mostly from Networks, Security, and Blockchain. I have also recently started leveraging Compilation techniques.

Impact: Owing to my inclination to *interdisciplinary* research and my intrigue for *real-world* systems, my work had a significant impact on the blockchain community. Particularly, my research on routing attacks on Bitcoin challenged long-standing assumptions about the network, triggered modifications in the Bitcoin implementation, received widespread media coverage, and was awarded an Applied Networking Research Prize by IETF.

1. *Security of Blockchain Networks*

Blockchain technology has received tremendous attention from the research community and the corporate world. Among all blockchain applications, cryptocurrencies (*e.g.,* Bitcoin) are the most widely-used, offering an *open* platform for *secure* and *anonymous* transactions. With millions of users per day, Bitcoin is a target of choice for attackers.

Bitcoin relies on a distributed network of nodes that communicate over the Internet infrastructure. Thus, Bitcoin traffic is forwarded via multiple networks called Autonomous Systems (ASes), which compose the Internet, as shown in Fig. 2. While Bitcoin's security properties have been extensively studied, two questions were never asked:

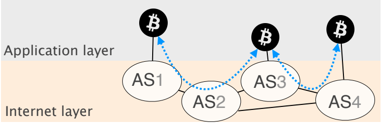

Figure 2: Bitcoin nodes communicate over the Internet which is composed of smaller networks, namely ASes.

Q1: Can a malicious or compromised AS attack Bitcoin?
Q2: How can we protect such systems from AS-level adversaries?

In my dissertation work, I answered both questions. First, I proved that a single AS is able to compromise Bitcoin's security, pseudonymity, and availability. Second, I showed that we can protect Bitcoin-like systems from AS-level adversaries by leveraging Internet policies.

1.1 *Routing Attacks on Bitcoin*

To gauge the power of AS-level adversaries, I did the first analysis of the Bitcoin network from the routing perspective.[1] My findings contradict a core assumption about the Bitcoin network, namely its decentralization, unavoidably challenging many of its security properties. Particularly, my analysis revealed two unintuitive insights:

I1: A few ASes intercept a large portion of Bitcoin connections.
I2: A few ASes host a large portion of the nodes and the mining power.

Driven by these insights, I introduced a new attack vector, namely *routing attacks*. To prove its menace, I uncovered and ethically performed in the wild three novel routing attacks against Bitcoin: the Partitioning, the Perimeter, and the Delay attack. Notably, the first two attacks *generalize to any blockchain* system working atop the Internet.

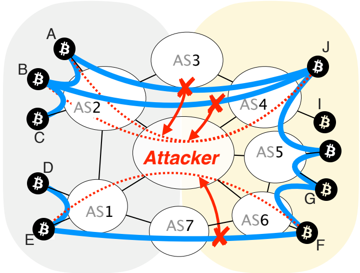

Figure 3: In the Partitioning attack, the AS-level adversary splits the Bitcoin network into the grey and yellow components, effectively preventing consensus.

Partitioning Attack.[1] I showed that an AS-level adversary can compromise the Bitcoin system as a whole by splitting it into two disjoint components, as shown in Fig. 3. In effect, nodes can no longer reach consensus, *i.e.,* agree on the latest distribution of funds. The Partitioning attack is the most effective and general routing attack to date as: *(i)* it is an unavoidable denial of service attack; and *(ii)* it can cause revenue loss of thousands of dollars to miners.

Perimeter Attack.[2,3] I showed that an AS-level adversary can compromise Bitcoin's anonymity properties by mapping a user's Internet address to one of their transactions. The Perimeter attack is dangerous for users and relevant today. First, finding a user's transaction is often equivalent to learning their entire transaction history. Second, transaction surveillance companies already sell such data.

Delay attack.[1] I showed that an AS-level adversary can prevent individual clients from securely using Bitcoin by delaying the delivery of new information to them, leaving them vulnerable to fraud.

[1] Maria Apostolaki, Aviv Zohar, and Laurent Vanbever. Hijacking Bitcoin: Routing Attacks on Cryptocurrencies. In *Security and Privacy (IEEE S&P'17)*, 2017. IETF/IRTF ANRP Award

[2] Maria Apostolaki, Cedric Maire, and Laurent Vanbever. PERIMETER: A Network-layer Attack on the Anonymity of Cryptocurrencies. *Under Submission*

[3] iPoster for the Perimeter attack

1.2 *Protecting Bitcoin Against Routing Attacks*

To secure Bitcoin from the Partitioning attack, I have built SABRE,[4] a relay network that is the first and only defense against this attack. By leveraging Internet policies in locating relay nodes, SABRE's design allows nodes of *any blockchain system* to exchange information, even in the presence of an AS-level adversary. SABRE's relay nodes are able to sustain Tbps of benign or malicious load owing to SABRE's hardware/software co-design, illustrated in Fig. 4.

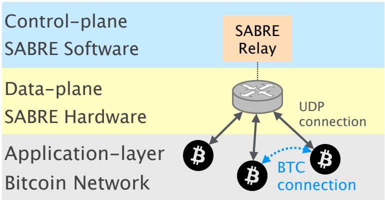

Figure 4: SABRE is a full-stack relay network working alongside the Bitcoin network to protect it against the Partitioning attack. Each relay node has a software and a hardware component.

1.3 *Broader Impact*

My work on the security of Bitcoin had a significant impact on the blockchain community. First, it raised awareness of the need for *application-layer* defenses against routing attacks. As a result, it triggered a modification in the Bitcoin Core[5] *i.e.*, the reference implementation of the Bitcoin client. Second, my work is one of the critical motivations for asynchronous blockchain systems. Finally, it has received widespread media coverage.[6]

Beyond *application-layer* defenses, my work highlights the need for secure Internet routing. As a recognition of this contribution, my work was awarded an Applied Networking Research Prize by IETF, the premier Internet standards body. To further highlight the need for secure Internet routing, I joined forces with researchers focusing on other services vulnerable to routing attacks, *e.g.*, anonymity systems, certificate authorities. Our work[7] targets a broader scientific community and provides vital momentum for the deployment of network-layer defenses.

2. *Reliability & Performance in Internet Routing*

The default Internet routing protocol, namely BGP, statically selects a single path per destination along which to direct traffic, despite the rich path diversity of today's Internet. As BGP is oblivious to performance, the selected paths are often suboptimal. Indeed, my measurements showed that BGP selects a path of higher delay than the best available 90% of the times. Worse yet, in case of a link failure, BGP takes minutes to change its selection, resulting in minutes of downtime for users.

To address BGP's shortcomings, I used a *cross-layer approach*: I leveraged information exposed by the application-layer protocol (*i.e.*, TCP) to revise the decisions of the networking layer (*i.e.*, routing). This approach is practical due to the recent development of programmable switches, which allow fine-grained monitoring and flexible forwarding at line-rate. Using this approach, I designed RouteScout and co-designed Blink.

[4] Maria Apostolaki, Gian Marti, Jan Müller, and Laurent Vanbever. SABRE: Protecting Bitcoin against Routing Attacks. In *Proceedings of the 26th Network and Distributed System Security Symposium (NDSS'19)*, 2019

[5] Bitcoin Core PR Review Club

[6] Bitcoin.com, The Morning Paper, The Register, Coin Telegraph, Naked Security, Coindesk

[7] Yixin Sun, Maria Apostolaki, Henry Birge-Lee, Laurent Vanbever, Jennifer Rexford, Mung Chiang, and Prateek Mittal. Securing Internet Applications from Routing Attacks. *Communications of the ACM (CACM)*, 2020

RouteScout[8] solves BGP's performance suboptimality problem by testing equal-cost paths and rerouting traffic accordingly. RouteScout compares the performance across paths online by observing TCP signals (*e.g.,* packet re-transmissions) and forwards traffic to the most performant one cautiously to avoid oscillations.

Blink[9] is the first data-driven fast reroute framework, effectively solving BGP's reliability problem. In particular, Blink reroutes traffic to an alternative path upon detection of a link failure. Blink detects failures directly in the data plane by monitoring TCP-induced signals. In effect, Blink restores connectivity in less than a second, instead of waiting minutes for BGP to converge.

3. *Performance Guarantees in Buffer Management*

To reduce costs and maximize utilization, network devices often rely on a shared buffer whose allocation across ports is dynamically adjusted by a buffer management algorithm, as illustrated in Fig. 5.

During my time at Microsoft, I observed that while sharing avoids wasting buffer space during times of low utilization, it also creates harmful interferences across seemingly independent traffic (*e.g.,* traffic going to different ports). By analyzing the state-of-the-art buffer management algorithm, namely Dynamic Thresholds (DT), I found that the root cause of the harmful interferences was DT's slow convergence on the buffer allocation under abrupt changes in demand. In effect, DT cannot offer any performance guarantees. To address DT's shortcomings, I designed FAB and Plasticine.

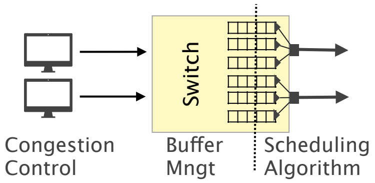

Figure 5: Buffer management influences both host-level optimizations (*e.g.,* congestion control) and port-level optimizations (*e.g.,* scheduling algorithm).

FAB[10] is a flow-aware buffer-sharing scheme built on top of DT. FAB makes the buffer allocation more predictable by preventing long-lived flows (often associated with less critical tasks) from monopolizing the buffer. This is beneficial for shorter flows whose performance becomes less dependent on the presence of other flows.

Plasticine[11] is a novel and practical buffer-sharing scheme that offers *provable* isolation guarantees without keeping the buffer idle. Plasticine achieves this by taking into consideration the buffer's temporal behavior and content. Notably, unlike the common belief that buffer management can only be devised by the chip manufacturer, I implemented Plasticine on a programmable device (Barefoot Tofino).

[8] Maria Apostolaki, Ankit Singla, and Laurent Vanbever. Performance-Driven Internet Path Selection. *CoRR abs/2001.07817*, 2020

[9] Thomas Holterbach, Edgar Costa Molero, Maria Apostolaki, Alberto Dainotti, Stefano Vissicchio, and Laurent Vanbever. Blink: Fast Connectivity Recovery Entirely in the Data Plane. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI'19)*, 2019

[10] Maria Apostolaki, Laurent Vanbever, and Manya Ghobadi. FAB: Toward Flow-aware Buffer Sharing on Programmable Switches. In *ACM Workshop on Buffer Sizing*, 2019

[11] Maria Apostolaki, Vamsi Addanki, Manya Ghobadi, and Laurent Vanbever. Plasticine: A Flexible Buffer Management Scheme for Data Center Switches. *Under Submission*

4. *Future Directions*

I intend to leverage my knowledge on networking and blockchain towards two main research directions: facilitating secure cross-layer distributed systems' design (§4.1), and systematizing network management following the general-purpose computing paradigm (§4.2).

4.1 *Cross-layer Distributed Systems' Design*

Distributed systems are increasingly used today. Still, they are often designed without explicitly stating their requirements from the underlying network. This gap leaves critical systems at risk and prevents them from leveraging networking advancements. In the future, I intend to bridge this gap in three ways: *(i)* by demonstrating the effectiveness of a general network adversary to various distributed applications; *(ii)* by leveraging networking advancements to improve the performance and security of such systems; and *(iii)* by building an Internet-wide testbed to facilitate interdisciplinary research.

General network adversary. In my dissertation, I uncovered critical network vulnerabilities of blockchain systems. In the future, I intend to extend this research direction along two dimensions: *(i)* unexplored network attacks against blockchain, *e.g.*, malicious increase of the propagation delay or loss; and *(ii)* network attacks against other distributed systems, *e.g.*, VoIP,[12] smart environments, and distributed file sharing.

Cross-layer design. Distributed systems working atop the Internet naturally inherit its vulnerabilities and inefficiencies. To prevent this, we need to challenge all traditional *abstractions* and adopt a *cross-layer* design approach that *(i)* limits the effectiveness of known Internet attacks (*e.g.*, BGP hijacking) on the newly-designed systems; *(ii)* leverages relevant networking advancements (*e.g.*, emerging networking hardware) to increase the systems' performance; and *(iii)* considers rational (not necessarily benign) behavior from network-infrastructure stakeholders.

A PlanetLab for Blockchain. Modeling the Internet is challenging, especially for emerging technologies, such as blockchain applications, whose interaction with the underlying network is not well understood. Still, these interactions must be extensively tested before applications are widely deployed. I intend to build an Internet-wide testbed composed of *both computing and networking devices* to allow researchers to measure the security and performance of their applications in practice. To promote collaborations, I intend to organize worldwide hackathons in which participants will be given access to a slice of the infrastructure in a mission to break other participants' systems.

[12] Ege Cem Kirci, Maria Apostolaki, Roland Meier, Ankit Singla, and Laurent Vanbever. Mass Surveillance of VoIP Calls Does Not Require a Nation-state Budget. *Work in Progress*

4.2 *Efficient & Secure Network Management*

Networking devices were recently made programmable. Yet, their management remains rather ad-hoc, thus error-prone and inefficient. For instance, network operators often program each device in isolation and leave resource management, privacy, and security issues as an afterthought. Instead, I believe we should first design an abstraction that will allow us to manage a network of switches as we manage a general-purpose computing system. Such an abstraction will enable the use of a broader spectrum of well-studied techniques for improving our networks' efficiency and security.

Efficient resource management. Programmability increases the networks' design space, making efficient resource management increasingly challenging. Today, a network operator is overloaded with tasks, including programming each switch considering its hardware constraints, splitting functionality across multiple devices, and tuning the network topology and routing accordingly. To alleviate the operator's burden, I intend to draw inspiration from general-purpose computing. So far, I have leveraged profile-guided optimizations (a widely-used technique in general-purpose compilers) to automatically minimize the hardware requirements of a given network (P4) program.[13] Yet, this barely scratches the surface of beneficial tech-transfer from general-purpose computing to networking. For instance, one could also draw a parallel from a network of switches to a multi-core system. Doing so will allow us to directly apply well-studied techniques used in parallelization and scheduling problems to manage our networks more scalably and efficiently.

Secure network management. We often forward packets from networking hardware to general-purpose software components, where they are stored until they are further processed. While storing packets eases the processing and allows for complex operations, it also exposes sensitive user data to all the software and hardware vulnerabilities of general-purpose systems. Instead, I believe we should start seeing networking hardware as our minimum trusted computing base and perform needed operations on sensitive user traffic directly and exclusively there. While programs running on networking hardware might unintentionally leak information and the networking hardware itself might be unreliable, they are both more restrictive and thus easier to verify than general-purpose systems. With this in mind, I intend to work towards verifying the security and privacy properties of P4 programs and networking hardware, *e.g.*, ensure that they do not store sensitive information.

[13] Patrick Wintermeyer, Maria Apostolaki, Alexander Dietmüller, and Laurent Vanbever. P2GO: P4 Profile-Guided Optimizations. In *Proceedings of the 19th ACM Workshop on Hot Topics in Networks (HotNets'20)*, 2020

References

Click on the title to access the corresponding PDF. All my publications, talks, and CV are also available online at www.apostolaki.net.

Maria Apostolaki, Aviv Zohar, and Laurent Vanbever. Hijacking Bitcoin: Routing Attacks on Cryptocurrencies. In *Security and Privacy (IEEE S&P'17)*, 2017. IETF/IRTF ANRP Award .

Maria Apostolaki, Cedric Maire, and Laurent Vanbever. PERIMETER: A Network-layer Attack on the Anonymity of Cryptocurrencies. *Under Submission*.

Maria Apostolaki, Gian Marti, Jan Müller, and Laurent Vanbever. SABRE: Protecting Bitcoin against Routing Attacks. In *Proceedings of the 26th Network and Distributed System Security Symposium (NDSS'19)*, 2019.

Yixin Sun, Maria Apostolaki, Henry Birge-Lee, Laurent Vanbever, Jennifer Rexford, Mung Chiang, and Prateek Mittal. Securing Internet Applications from Routing Attacks. *Communications of the ACM (CACM)*, 2020.

Maria Apostolaki, Ankit Singla, and Laurent Vanbever. Performance-Driven Internet Path Selection. *CoRR abs/2001.07817*, 2020.

Thomas Holterbach, Edgar Costa Molero, Maria Apostolaki, Alberto Dainotti, Stefano Vissicchio, and Laurent Vanbever. Blink: Fast Connectivity Recovery Entirely in the Data Plane. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI'19)*, 2019.

Maria Apostolaki, Laurent Vanbever, and Manya Ghobadi. FAB: Toward Flow-aware Buffer Sharing on Programmable Switches. In *ACM Workshop on Buffer Sizing*, 2019.

Maria Apostolaki, Vamsi Addanki, Manya Ghobadi, and Laurent Vanbever. Plasticine: A Flexible Buffer Management Scheme for Data Center Switches. *Under Submission*.

Ege Cem Kirci, Maria Apostolaki, Roland Meier, Ankit Singla, and Laurent Vanbever. Mass Surveillance of VoIP Calls Does Not Require a Nation-state Budget. *Work in Progress*.

Patrick Wintermeyer, Maria Apostolaki, Alexander Dietmüller, and Laurent Vanbever. P2GO: P4 Profile-Guided Optimizations. In *Proceedings of the 19th ACM Workshop on Hot Topics in Networks (HotNets'20)*, 2020.