

Research Statement

John V. Monaco

Email: contact@vmonaco.com

Website: www.vmonaco.com

July 19, 2019

Modern computing devices are equipped with an unprecedented range of sensing capabilities, enabling novel insights to be gained about the user, the device, and the environment. While these insights permit applications to be more responsive to user actions, operate with greater efficiency, and react to complex stimuli, heterogeneous sensing capabilities have also led to a new class of side channel attacks aimed to undermine user privacy and compartmentalized security. Interestingly, the patterns that emerge from sensed input, such as through keyboard, microphone, and numerous other modalities, can be used to both increase security, such as by continuously verifying the user's identity, and decrease security, with cross-sensor eavesdropping being a prime example. It is not clear which force, if any, will prevail, and more than likely there will remain a balance between security and vulnerability; however, the sensors are here to stay.

My research program aims to better understand the security and privacy implications that underlie human-computer interaction and emerging computer architectures, especially in the presence of heterogeneous sensing. Prominently, this includes the user's temporal behavior, which gives rise to multiple human-based timing side channel attacks targeting user privacy and application content. Such attacks are largely enabled by two related phenomena: between-user variability, whereby differences in behavior can reveal the user or device identity, and between-user consistency, whereby sameness in behavior across a population can reveal device usage patterns, such as specific keyboard keys that were pressed. My work has the goal to better understand the interplay between these opposing factors and develop novel applications and proofs of concept that leverage each.

More generally, I examine the dynamics of both user and device behavior from a security perspective in order to create a safe and dependable human-computer ecosystem. My current focus on temporal behavior is driven by the ubiquity of timestamped events that may be observed both on the device and over a network. Operating under the assumption of strong encryption, the time of transmission represents, perhaps, the minimal amount of information a single encrypted message must reveal to an eavesdropping adversary. By itself, a single timestamp seems relatively benign. However, in aggregate form a distinct temporal pattern can emerge. As these patterns propagate to the network level, an adversary is granted the unique opportunity to remotely observe user behavior

and device operation. This dilemma calls for a better understanding and management of information leakage from user and device behavioral patterns.

My current efforts follow three parallel lines of inquiry. These include: 1) modeling user and device behavior with a focus on security; 2) the development of techniques to counter attacks that leverage exposed behavior, such as to preserve anonymity and mitigate information leakage; and 3) from a computational perspective, investigating the impact of emerging non-von Neumann architectures on security and privacy.

1 Modeling user and device behavior

While encryption aims to protect a message from eavesdropping third parties, the time of transmission may itself reveal some unintended information about the sender or the message contents. This situation can arise during interactive real-time client/server applications, such as ssh in interactive mode or Google search, whereby network traffic flows from the client to the server upon each input event, and is due in part to two phenomena in human-computer interaction (HCI):

1. *Between-user variability*: Individual users exhibit unique and identifiable behavior interacting with a computer interface, resulting from idiosyncratic factors such as motor skill and physiology. These differences can be observed during interaction with keyboard, touchscreen, and other HCI modalities, and as the events propagate to the network level, user identification can be performed passively and remotely.
2. *Between-user consistency*: An entire population may operate a computer interface within certain constraints, enabling an adversary to make general inferences about a particular user's actions. HCI largely follows general rules, such as Fitts' Law in navigating the mouse pointer on a computer screen and the inverse scaling of keystroke latencies with key distance: shorter time intervals usually correspond to keys that are far apart compared to longer time intervals between neighboring keys.

These two phenomena are generally seen as opposing one another, and quantifying their tradeoff remains an active area of my current research. To date, I have focused on three areas that either leverage or compensate for each phenomenon, including: 1) accurate modeling of user and device behavior; 2) behavior-based identification and verification; and 3) increasing model robustness to variability in environment conditions.

Modeling: In order to better understand user behavior in an HCI context, I developed the partially observable hidden Markov model (POHMM), an extension of the hidden Markov model in which the hidden state is conditioned on an independent Markov chain [8]. The POHMM is motivated by processes comprised of point events in time that are accompanied by discrete metadata, such as an event type. Applied to keystroke dynamics, the POHMM leverages the second above phenomenon, namely that different time intervals are typically observed between different keys. The independent Markov chain is formed by the keyboard keys, which conditions the hidden state, enabling the model to capture keystroke dynamics independent of linguistic structure. This work, combined with nonlinear dynamical systems techniques, has achieved first place in several

international biometric competitions^{1,2}, received recognition from local media^{3,4}, and led to one patent [1].

Identification and verification: Moving beyond motor behavior, and towards the cognitive and social bands in Newell’s time scale, I have examined how time intervals alone can be leveraged for user identification and verification in the absence of other reliable information [2]. Using a nonlinear dynamical systems approach, I examined the extent to which lower-frequency actions are able to reliably distinguish between users, such as the identification of users by Bitcoin transaction behavior [3]. The consideration of other sources of temporal behavior, including those which may be partially censored, i.e., an open-sided interval is observed, remains ongoing work [5].

Real-world conditions: Non-stationarity can arise in any scenario where the environment conditions present during training are different than the conditions observed during evaluation, such as a user who is impaired from typing normally [6]. This problem can be approached within a transfer learning framework, and the development of methods to increase the robustness of existing models to real-world conditions is a focus of my work. My recent approach involved forming a bipartite graph between samples in different domains and learning a representative transfer function across domains [10]. This function can then be applied during evaluation in order to adapt a black-box model to previously-unseen conditions.

2 Privacy-enabling technologies

The interactive real-time client/server model, where the role of the client is to transmit the user’s actions to the server in an event-driven fashion, represents a rather broad class of web-based applications. As HCI patterns manifest at the network level, an adversary can remotely observe user and device behavior without the victim’s cooperation or knowledge. Leveraging each phenomenon described in the previous section, the ability to remotely identify a user or determine a user’s actions presents a privacy concern. From this perspective, my work in privacy-enabling technologies has focused on two areas: 1) online and adaptive behavior obfuscation; and 2) behavior spoofing.

Behavior obfuscation: The goal of behavior obfuscation is to prevent identifying behavioral patterns from being publicly exposed. There are actually several different objectives to meet this criteria, such as concealing the identity of the user, masking the user’s actions, and limiting the ability to reproduce the user’s behavior (i.e., prevent spoofing). As each objective has different quantitative requirements, the obfuscation strategy that should be employed depends on the particular user’s privacy objectives. My previous work has focused on quantifying each objective and developing a set of techniques that can be applied online and adaptive to the user’s behavior with minimal impact to the application experience [7]. For real-time applications, I have shown that there is a direct

¹<http://web.archive.org/web/20171221175245/http://www.kasprowski.pl/emvic/results.php>

²<http://web.archive.org/web/20171221175327/https://sites.google.com/site/btas16kboc/>

³<http://web.archive.org/web/20171221180722/http://www.westchestermagazine.com/914-INC/Q2-2013/Wunderkinds-2013-John-Vincent-Monaco-24/>

⁴<http://web.archive.org/save/https://www.orau.org/arlfellowship/profiles/monaco.htm>

tradeoff between obfuscation ability and responsiveness of the application. This work resulted in the development of kloak (keystroke-level online anonymization kernel), an input device-level tool for obfuscating keystroke time intervals with strict guarantees on application responsiveness. kloak has received interest from the privacy community and is currently being integrated into privacy-centric Linux distributions Whonix⁵ and Tails.

Behavior spoofing: Closely related to obfuscation is the problem of spoofing, or modifying user and device behavior to appear as though it is something it is not. Spoofing can provide anonymity by masking the true behavior but also brings to light the possibility of abuse whereby an adversary spoofs a user’s behavior with malicious intent. This also remains an active area of my research, especially spoofing behavior in real time. Previously, I developed a model that enabled an adversary to spoof a user’s keystroke dynamics with only remote web-based observations of typing behavior [9]. This work demonstrated that with minimal user-specific knowledge, and combined with a model trained on a general population, the security of a behavioral biometric system is substantially weakened.

3 Challenging computational hardness

Computing architectures are on the cusp of a fundamental shift towards concurrency as sequential models struggle to process large amounts of data in real time. Alongside this shift, a range of non-von Neumann architectures, some designed to operate within size, weight, and power constrained environments, have begun to emerge. Many of these architectures are inspired by the mammalian brain, built on principles such as massive parallelism and event-driven computation. Though much interest has been invested in leveraging these devices for intelligent and autonomous systems, their impact to security and privacy remains an open question.

A recent focus of my research has been evaluating the impact of emerging non-von Neumann architectures to fundamental security and optimization problems. I am specifically interested in neuromorphic and neural-inspired architectures and how these devices can be leveraged to challenge some of the computational hardness assumptions underlying encryption protocols. My work on brain-inspired integer factorization has led to several theoretical insights, such as the lossy encoding of polynomial sequences on a neural architecture and the efficient detection of smooth numbers by a neural network. Practically, I was able to demonstrate that a neuromorphic processor can provide an asymptotic speedup to the sieving portion of the quadratic sieve and number field sieve integer factorization algorithms [12]. Further efforts established techniques for mitigating the adverse effect of some common neuromorphic architectural constraints [11]. This line of inquiry remains an active area of research with future directions described below.

4 Future work

In the near term, I plan to continue to examine the interplay between behavioral attacks and defenses. This will involve the development of new modeling techniques, attack

⁵<http://web.archive.org/web/20171221185732/https://phabricator.whonix.org/T596>

vectors, and methods of defense. I see the integration of concepts from adversarial machine learning with behavior obfuscation as a promising area to pursue for online defenses. At the same time, quantifying exactly what information might be leaked from exposed behavior, and the impact of such leakage, remains a primary objective.

In the long term, I would like to enable a shift from reactive to proactive defense models. Currently, it is more common to patch systems and build tools that respond to a particular threat, especially those that arise from user behavior; in the long term, I would like to shift this paradigm to a proactive design, with security and privacy concepts baked into the user experience. I believe that part of this shift will require a better understanding of what types of behavior patterns pose a threat and at what scales.

Quantifying information leakage: Recognizing user activities based on behavioral patterns remains an area of my active research. I am especially interested in human-based timing side channels, such as the ability to discern keyboard keys based on time intervals. This is analogous to the way a timing side channel enables encryption keys to be efficiently recovered, relying on the notion that certain math operations take shorter or longer to complete based on the encryption key bit pattern. Initial results and a systemization of knowledge of decades of research suggest that temporal keylogging side channels potentially remain an underdeveloped threat [4]. An immediate next step is to develop more robust action recognition models, e.g., temporal keylogging in the presence of noise.

Dynamically adapting behavior: I will further develop methods of dynamically adapting user and device behavior to meet various privacy criteria, such as to conceal actions or identity. Doing this effectively in near-real-time remains a challenge, and there is a tradeoff between the usability of the method and the ability to obfuscate behavior. I plan to construct a framework for dynamically adapting behavior that supports multiple simultaneous privacy objectives, such as to preserve anonymity or to mimic the behavior of a different device. Such a framework would have a broad impact for privacy-enabling technologies, wherein behavior leakage remains a valid concern. At the same time, the techniques should be usable, such as introducing minimal delay to real-time applications.

Emerging computer architectures: Emerging neural-inspired architectures have sparked a wave of interest for machine learning and artificial intelligence applications. However, with vastly different computational capabilities, their potential impact to security and privacy is less understood. My future work will utilize emerging architectures to efficiently solve optimization problems, challenge protocols that rely on computational hardness, and design new protocols that leverage specific architectural properties. Analog architectures built with novel materials, such as excitable graphene lasers and magnetic tunnel junctions, are especially of interest, as these devices operate on extremely small time scales and are naturally stochastic. I plan to investigate the extent to which these architectures can be used to increase and/or decrease security.

References

- [1] Jordan A. Berger and John V. Monaco. Universal keyboard. US Patent No. 9,864,516. Filed on 27 July 2015. Published on 9 Jan 2018.
- [2] John V. Monaco. Classification and authentication of one-dimensional behavioral biometrics. In *Proc. International Joint Conference on Biometrics (IJCB)*. IEEE, IAPR, 2014.
- [3] John V. Monaco. Identifying bitcoin users by transaction behavior. In *Proc. Defense, Security, and Sensing: Biometric and Surveillance Technology for Human and Activity Identification XII*. SPIE, 2015.
- [4] John V. Monaco. Sok: Keylogging side channels. In *Proc. 39th IEEE Symposium on Security and Privacy (S&P)*. IEEE, 2018.
- [5] John V. Monaco, Malka Gorfine, and Li Hsu. General semiparametric shared frailty model estimation and simulation with frailtysurv. *Journal of Statistical Software*, 2018.
- [6] John V. Monaco, Gonzalo Perez, Charles C. Tappert, Patrick Bours, Soumik Mondal, Sudalai Rajkumar, Aythami Morales, Julian Fierrez, and Javier Ortega-Garcia. One-handed keystroke biometric identification competition. In *Proc. 8th IAPR International Conference on Biometrics (ICB)*. IAPR, 2015.
- [7] John V. Monaco and Charles C. Tappert. Obfuscating keystroke time intervals to avoid identification and impersonation. In *Proc. 9th IAPR International Conference on Biometrics (ICB)*. IAPR, 2016.
- [8] John V. Monaco and Charles C. Tappert. The partially observable hidden markov model and its application to keystroke dynamics. *Pattern Recognition*, 2017.
- [9] John V. Monaco, Charles C. Tappert, and Md Liakat Ali. Spoofing key-press latencies with a generative keystroke dynamics model. In *Proc. 7th IEEE International Conference on Biometrics Theory, Applications and Systems (BTAS)*. IEEE, 2015.
- [10] John V. Monaco and Manuel M. Vindiola. Crossing domains with the inductive transfer encoder: Case study in keystroke biometrics. In *Proc. 8th IEEE International Conference on Biometrics Theory, Applications and Systems (BTAS)*. IEEE, 2016.
- [11] John V. Monaco and Manuel M. Vindiola. Factoring integers with a brain-inspired computer. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 2017.
- [12] John V. Monaco and Manuel M. Vindiola. Integer factorization with a neuromorphic sieve. In *Proc. 50th IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, 2017.