Editors: Davide Balzarotti, davide.balzarotti@eurecom.fr | William Enck, whenk@ncsu.edu | Samuel King, kingst@ucdavis.edu | Angelos Stavrou, astavrou@gmu.edu

Credit Card Fraud Is a Computer Security Problem

Samuel T. King | UC Davis and Bouncer Technologies
Patrick Traynor and Christian Peeters | University of Florida and Skim Reaper
Nolen Scaife | University of Colorado
Zainul Abi Din and Hari Venugopalan | UC Davis

Most apps and merchants do not want to deal with financial fraud, but, if they accept payments, they will eventually have to. Our position is that credit card fraud prevention is a technical problem that needs technical solutions.

Industry research for credit card fraud estimates that merchants will lose US$130 billion to fraud between 2018 and 2023.[9] Four trends drive this rise in these losses. First, fintech innovations empower more apps and online merchants to accept payments, creating opportunities for fraudsters to use services directly and fraudulently pay for them, e.g., drug dealers use Uber to run drugs and pay for these rides with stolen credit card numbers.[1]

Second, attackers can establish an agent service where they sell discounted goods and services to people while paying for them stolen credit card numbers.[3] For example, fraudsters could order food for people and collect money from them directly, paying the food delivery app using these stolen credit cards. Consumers still get the goods, but they pay the fraudster instead of the merchant; when the rightful owner of the card disputes the transaction, the merchant is stuck with the bill.

Third, two-sided marketplaces, such Airbnb with hosts and guests and Uber with drivers and passengers, create opportunities for a 21st-century version of money laundering.[5] In these schemes, attackers occupy both sides of the marketplace and use the company to move money between the two. For example, fraudsters could create both guest and host accounts on a vacation rental app and pay for a fake stay using stolen credit cards, causing a guest transaction with a stolen card to result in a deposit to the host.

©SHUTTERSTOCK/JARIRIYAWAT

Fourth, fintech innovations have also provided attackers with more options to collect money themselves. In the previous decade, antispam researchers could track spammers down to a small number of merchant

Digital Object Identifier 10.1109/MSEC.2021.3050247
Date of current version: 15 March 2021

accounts that they used to collect payments.[7] In contrast, modern-day fraudsters can move money using peer-to-peer payments apps, like Venmo or WeChat Pay, or cryptocurrencies, like Bitcoin or ZCash, so there is not a merchant account to go after. As payments continue to become more distributed, this basic trend of more difficult financial accountability will continue. Accountability and fraud prevention are, therefore, critical features and differentiators for credit cards in the long term.

As companies find more ways to integrate payments into their apps, fraudsters will continue to find more ways to steal money. Although apps can make progress by creating large antifraud teams that focus solely on stopping scams, coming up with principles and general techniques is the only way that we will stop this threat completely.

A Holistic View of Antifraud Systems

In our vision, we will eliminate credit card fraud. However, to do this, one must be able to navigate consumers, merchants, point-of-sale devices, card-not-present transactions, payment processors, payment networks, and issuing banks—plus all of the various levers that each principal uses to limit fraud. Furthermore, antifraud systems make heavy use of device, payment method, and user tracking to identify repeat scammers as well as machine learning (ML) to predict fraud, so end-user privacy and ethics must be a fundamental part of any long-term solution.

Given the complexity of credit card fraud, our position is that a defense-in-depth approach is the most practical way to solve fraud holistically (Figure 1). At the base level, we advocate cutting off the supply of card details before they become stolen. Since fraudsters will still steal some card details, we advocate designing and implementing defensive systems that merchants can use to stop these fraudsters from being able to use stolen cards. Fraudsters will still find services to use stolen cards, so we advocate empowering law enforcement to catch them. By viewing this problem at multiple key points in the fraud lifecycle and applying defense-in-depth principles, we believe that we can eliminate stolen-card financial fraud.

Cutting Off the Supply of Stolen Card Numbers

Attackers use stolen account data to make fraudulent transactions. One cause of this is static account data; newer technologies, such as phone-based near-field communication payments and Europay, Mastercard, and Visa (EMV) cards, attempt to avoid fully static transaction data. However, dynamic data require active circuitry on the payment mechanism, which increases the cost of these systems and slows their deployment. It remains an open challenge to accept dynamic transaction data in card-not-present transactions (e.g., online) where the consumer cannot physically present the card.

Aside from standard data breaches, attackers acquire account data through one of two mechanisms, each leading to an open research direction:

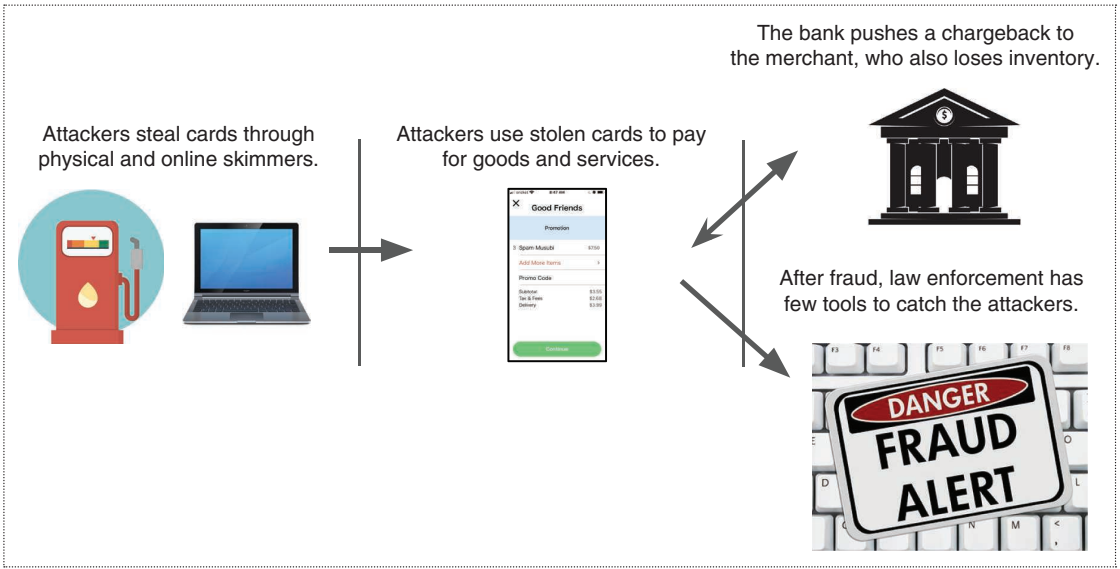

Figure 1. The overall lifecycle for credit card fraud. It starts with attackers stealing card numbers and using these to purchase goods and services, resulting in chargebacks where law enforcement has little or no recourse. (Sources: gas pump: https://commons.wikimedia.org/wiki/File:Creative-Tail-Objects-gas-station.svg (creativetail.com); laptop: http://www.pngall.com/laptop-png/download/5426; fraud alert: https://www.new-york-city-travel-tips.com/renting-apartments-new-york-frauds/.)

- *Physical skimming*: In this scenario, the attacker compromises the point-of-sale hardware. This device is additive and either performs an independent read of the payment card or captures the data electronically as it passes through the terminal. For ATMs and devices that accept personal identification numbers (PINs), these attacks can include PIN pad overlays or cameras to capture the victim's entry. Worse, EMV cards permit capture of the account number and other sensitive data via the chip, leading to EMV "shimming" attacks. There is a critical need for systems that detect tampering and verify the state of payment terminals. Our previous work has made substantial progress in this direction, but new techniques for modifying terminals continue to emerge.
- *Online skimming*: When buying goods or services online, consumers often enter payment card details manually, including the account number, cardholder name, and expiration date. Through compromised web sites, phishing, or social engineering, attackers capture the data. These attacks can go undetected for long periods of time due to advanced evasion techniques and victim selection (making the attack occur on only a subset of users). New approaches are needed for protecting data entry, monitoring and detecting compromised payment flows, and providing visible indicators to users about the safety of a payment process.

Empower Merchants to Reduce Fraud

Empowering merchants has two main advantages. First, merchants have full semantic information for a transaction and are in the best position to identify anomalies or potentially fraudulent transactions. For example, a merchant would know which products or transaction amounts tend to lead to fraud and the full transaction history of an individual user. These signals enable merchants to make the most accurate fraud predictions. Second, merchants can stop a transaction before a fraud happens, which is the last point in time when they can prevent damage.

In recent years, merchants have turned increasingly to user-facing challenges to verify users and payment methods for stopping fraud. These are a modern take on step-up authentication where an app shows only suspicious users a flow to collect additional information. With this, apps can separate good users whom the app flagged accidentally from fraudsters trying to attack the app. Examples of this new style of verification include Apple's FaceID, Uber's and Boxer's credit card scanning systems,[4] and Coinbase's ID verification flow.

Unfortunately, user-facing challenges open a new set of design constraints. They must trade off privacy by choosing to run on the server or client, cope with vastly different performance characteristics for ML running on client devices, and withstand a broad range of network speeds. Security challenges must deal with these subtleties of practical deployments or else they will block users unethically.

To provide an environment for apps to run user-facing challenges that are trustworthy and provide equal access for all users, we advocate the following two open research directions:

- *Ethical client-side ML*: The first research direction is designing methodologies and systems for security-related ML models that run on the client side. The key challenge is supporting advanced ML while still running on devices with a wide range of performance capabilities, like we see in practice today. For example, prediction times for the same ML model with identical inputs and system architectures varies by one to four orders of magnitude, based on data from our production deployment of Boxer.[4] The core issue is that security-related ML must be able to execute efficiently, or else it will block people using low-end phones due to their inability to run the ML models. Although there is recent work on running image classification models on mobile phones efficiently (e.g., SqueezeNet[6]), ML for security often requires architectures outside of the traditional image-classification paradigm, even when solving a computer vision problem.[4]
- *Secure Hardware*: A second research direction is designing hardware-based abstractions to enable trustworthy ML and other security logic to run on a device. Current abstractions are either too low level, like trusted execution environments that let apps execute raw instructions,[8] or too high level, like Apple's DeviceCheck, which exposes two persistent bits. The ideal abstraction would provide enough flexibility to run meaningful security checks, have a narrow enough interface to support asserting security invariants, and provide runtime support to access data that persist across device resets while still respecting end-user privacy. To be practical, these abstractions must be able to withstand both malicious humans using a legitimate device and malicious system-level software.

Empower Law Enforcement to Reduce Fraud

Developing new tools and algorithms is critical to assist businesses in their deterrence and detection of fraud. However, what happens next (especially when merchants and consumers detect fraud after the fact) is often unclear. While common advice often includes alerting law enforcement of an incident,

which agencies to contact and how to enable effective investigations are less well understood.

The challenges in this space have many similar constraints as those faced by retailers. For instance, outside of the largest law enforcement agencies, few such organizations have the expertise or person power to perform tasks such as skimmer reverse engineering. Therefore, we recommend that researchers look to make contributions to the following open problems.

- *Improved detection*: Law enforcement simply cannot be in every place at all times. Tools that are embedded or anchored to a location may help retailers know when to call for assistance from law enforcement, but they do little to alert patrol units that something is out of place. Tools such as the Skim Reaper[11] work well against deep-insert and overlay devices, but they do not detect other classes of skimmers. Wireless detection techniques have shown some success[2] but have, thus far, been easily evadable via configuration changes.[10] Finding strong indicators of skimming in environments where spectrum use is high and diverse (e.g., wireless cameras, customer phones, Bluetooth headphones, and so on) increases the difficulty. Tools providing fast, easy-to-use detection for a range of increasingly deceptive skimmers are critical.
- *Identification of campaigns*: Laws governing the penalties for skimming vary significantly from state to state. For instance, whereas Nevada law allows for a maximum fine of US$250,000 for each count of use/possession of skimming devices, the maximum penalty in Florida is only US$5,000. As such, being able to demonstrate that multiple devices participate in a single campaign is critical in many states to justify dedicating resources to an investigation. Tools including software and hardware similarity measures would help make such attribution possible.
- *Identification of skimming paraphernalia*: Attempts to use cloned cards have become more sophisticated. Whereas those attempting to use a cloned card traditionally used blank cards (i.e., plain white cards with no logos), many attacks now create realistic-looking, embossed clones. Even in the absence of seeing individual cards, behaviors such as attempting to perform transactions with multiple cards served as a strong indicator of fraud. However, there are many reasons that customers would legitimately behave in such a fashion, including customers using the remaining balance on each of multiple cards or attempting to split purchases across multiple cards to manage a combination of credit limits, overdraft fees, and interest rates. Finding customers with multiple cards alone is therefore a noisy indicator of fraud. Advances in this space could include tools to rapidly measure hard-to-fake manufacturing characteristics, including the consistency of magnetic encoding[12] or image quality on the physical card.

To solve credit card fraud, we need collaboration from industry, government, and academia to solve this problem holistically. This collaboration ranges from joint projects and NSF workshops dedicated specifically to financial fraud to funding opportunities from industry and governments. By working together, we can solve credit card fraud. ■

References

1. B. Anderson. "How uber is changing drug dealing." Vice. https://www.vice.com/en_us/article/qkjjwb/how-uber-is-changing-drug-dealing (accessed Jan. 27, 2021).
2. N. Bhaskar, M. Bland, K. Levchenko, and A. Schulman, "Please pay inside: Evaluating Bluetooth-based detection of gas pump skimmers," in *Proc. 28th USENIX Security Symp. (USENIX Security 19)*, Santa Clara, CA: USENIX Association, Aug. 2019, pp. 373–388.
3. T. Chen. "Advanced technologies for detecting and preventing fraud at uber." Uber Engineering. https://eng.uber.com/advanced-technologies-detecting-preventing-fraud-uber/ (accessed Jan. 27, 2021).
4. Z. A. Din et al., "Bouncer: Preventing fraud by scanning credit cards," in *Proc. USENIX Security Symp. (USENIX Security 2020)*, 2020, pp. 1571–1588.
5. K. Fazzini. "How criminals use Uber and Airbnb to launder money stolen from your credit card." CNBC. https://www.cnbc.com/2019/02/07/how-criminals-use-airbnb-uber-launder-stolen-credit-card-money.html (accessed Jan. 27, 2021).
6. F. N. Iandola, M. W. Moskewicz, K. Ashraf, S. Han, W. J. Dally, and K. Keutzer, "Squeezenet: Alexnet-level accuracy with 50x fewer parameters and <1mb model size," *CoRR*, Feb. 2016. [Online]. Available: http://arXiv:1602.07360
7. K. Levchenko et al., "Click trajectories: End-to-end analysis of the spam value chain," in *Proc. 2011 IEEE Symp. Security Privacy (SP '11)*, IEEE Computer Society, pp. 431–446. doi: 10.1109/SP.2011.24.
8. J. M. McCune, B. Parno, A. Perrig, M. K. Reiter, and A. Seshadri, "Minimal TCB code execution," in *Proc. 2007 IEEE Symp. Security Privacy, (SP'07)*, IEEE Computer Society, pp. 267–272. doi: 10.1109/SP.2007.27.
9. "Retailers to lose $130bn globally in card-not-present fraud over the next 5 years." Juniper Research. https://www.juniperresearch.com/press/press-releases/retailers-to-lose-130-bn-globally-in-card-fraud (accessed Jan. 27, 2021).
10. N. Scaife et al., "Kiss from a rogue: Evaluating detectability of pay-at-the-pump card skimmers," in *Proc.*

2019 IEEE Symp. Security Privacy (SP), pp. 1000–1014. doi: 10.1109/SP.2019.00077.

11. N. Scaife, C. Peeters, and P. Traynor, "Fear the reaper: Characterization and fast detection of card skimmers," in *Proc. 27th USENIX Security Symp. (USENIX Security 18)*, USENIX Association, Baltimore, MD, 2018, pp. 1–14.
12. N. Scaife, C. Peeters, C. Velez, H. Zhao, P. Traynor, and D. Arnold, "The cards aren't alright: Detecting counterfeit gift cards using encoding jitter," in *Proc. IEEE Symp. Security Privacy (S&P)*, 2018, pp. 1063–1076. doi: 10.1109/SP.2018.00042.

Samuel T. King is an associate professor in the Computer Science Department at the University of California, Davis, California, USA, and founder of Bouncer Technologies, an antifraud company. His research interests include fighting financial fraud from a first-principles perspective and security for digital identities in the 21st century. Contact him at kingst@ucdavis.edu.

Patrick Traynor is a professor in the Department of Computer and Information Science and Engineering at the University of Florida, Gainesville, Florida, 32611, USA, where he is also the John and Mary Lou Dasburg Preeminent Chair in Engineering. Contact him at traynor@ufl.edu.

Christian Peeters is a Ph.D. student in the Department of Computer and Information Science and Engineering at the University of Florida, Gainesville, Florida, 32611, USA. His research interests include cellular security and fraud detection. Contact him at cpeeters@ufl.edu.

Nolen Scaife is an assistant professor at the University of Colorado Boulder, Boulder, Colorado, USA. His research interests include systems and network security with a focus on real-world impact. Contact him at scaife@colorado.edu.

Zainul Abi Din is pursuing a Ph.D. in computer science at the College of Engineering, University of California, Davis, California, USA. His research interests include user-facing machine learning for security applications. Contact him at zdin@ucdavis.edu.

Hari Venugopalan is a Ph.D. student in computer science at the University of California, Davis, California, USA. His research interests include the practical applications of machine learning for security. Contact him at hvenugopalan@ucdavis.edu.

Digital Object Identifier 10.1109/MSEC.2021.3060559