

FinTechSec: Addressing the Security Challenges of Digital Financial Services

Patrick Traynor, Kevin Butler, and Jasmine Bowers | University of Florida
Bradley Reaves | North Carolina State University

You probably don't think about traditional banking very often. Many of us would be hard-pressed to say why we picked one financial institution over another—it could be one that partners with an alma mater, one used by friends and neighbors, or simply the one closest to home. Most important, most of us also have options and could easily take our business (and our money) down the street to another institution should we fail to receive the terms, service, and security we feel necessary to grow and protect our assets.

Now that we have you thinking about traditional banking, it's easy to begin enumerating the many things it enables: our employers electronically deposit our paychecks (and they become immediately available), we make payments using credit and debit cards (reducing the need to physically carry and protect cash), and we even have an array of protections against fraud. This infrastructure extends far beyond national borders and is now so pervasive in the developed world that travelers think nothing of withdrawing money at foreign ATMs. In short, traditional banking makes payments (and most of the challenges around it) largely frictionless to the consumer.

It would be easy to assume that everyone has access to traditional banking given its seeming ubiquity. Unfortunately, that assumption is simply wrong.

Billions around the world lack access to even the most basic



banking services for many reasons. Many simply lack physical access. Even more lack the ability to maintain the relatively high minimum balances required by traditional financial institutions. The practical impacts are significant. In the US alone, only 68 percent of homes were “fully banked” in 2015, meaning that the remaining 32 percent required the use of so-called “alternative financial services” including check cashing and payday loans.¹ Worldwide, some two billion people remain unbanked.

The lack of basic banking services makes tasks most of us take for granted, such as saving, electronic payment, and short-term loans,

essentially out of reach for huge portions of the population. Technology might provide a real path to so-called financial inclusion; however, as our research shows, security and privacy remain significant impediments to future progress in this space.

Our goal in this article is to discuss our experience in securing *mobile money*, a digital financial system that uses mobile phones to transfer currency without the need for a bank. Our efforts began in 2011 and have resulted in extensive collaboration with organizations including the US Department of State, the International Telecommunications Union (ITU), the GSM

Association (GSMA), the World Bank, and many individual providers and vendors. These transformative systems have already demonstrated the power to raise populations out of poverty, and we believe that they will soon be deeply intertwined with the traditional global financial infrastructure. This means that we have a chance to get security and privacy correct now instead of looking back with regret when the above systems are made manifest.

What Is Mobile Money?

In the mid-2000s, the Kenyan cellular provider Safaricom noticed an interesting trend. For some time, customers in its network could send minutes to their friends and families, and often did so to ensure that those with access to funds could talk to those without. However, a few enterprising customers began sending minutes in exchange for goods and services. This was no small innovation—at the time, the vast majority of Kenyans didn't have a bank account, and electronic payment was beyond most citizens' reach. In contrast, nearly eight out of 10 citizens had mobile phones. Seeing this tremendously unmet need for electronic payment being approximated with “top up” minutes, Safaricom launched M-Pesa in 2007 and allowed subscribers to send actual money to one another via SMS.

M-Pesa was an overnight success. Urban residents who would often travel long distances to physically transport money to their rural family members (often at the literal risk of highway robbery) could simply transfer those funds at the press of a few buttons. Moreover, M-Pesa overcame the problem of physical access by making virtually every vendor the equivalent of an ATM—capable of both depositing funds to and withdrawing funds from the network. Finally, M-Pesa charged

extremely low transaction rates, further enticing those unable to use traditional banking services to join.

M-Pesa now claims more than two-thirds of the Kenyan population as its customers. Moreover, this model has been copied and attempted widely across the globe (especially in the developing world). In 2016, there were more than half a billion mobile money accounts around the world, and the industry processed an estimated US\$22 billion.² These numbers continue to increase by staggering amounts each year.

What we've described here might sound somewhat familiar. After all, the past few years have seen the rise of peer-to-peer payment systems such as Apple Pay, Google Wallet, Samsung Pay, Venmo, and a handful of others. However, none of these are mobile money because they're all backed by the traditional banking infrastructure. That means that unless you acquire a credit or debit card, you really can't use these systems. Think of mobile money instead in the following way: rather than Bank of America or HSBC, AT&T or Orange now becomes your “bank,” and you deposit money or checks at your local gas station, corner market, or grocery store.

Mobile money is also not the same thing as cryptocurrency (for example, Bitcoin and Ethereum). Speaking very broadly, these two systems solve decidedly different problems. Whereas cryptocurrencies strive to create alternative money outside of centralized control, mobile money systems operate using traditional nation-state-backed fiat currency. While some researchers and start-ups have attempted to deploy cryptocurrencies in the context of mobile money, they haven't met much success. Moreover, mobile money is being used by a far greater number of people: M-Pesa alone reported 6 billion transactions in 2016,³ compared to Bitcoin's 184 million over

its entire lifetime (blockchain.info/charts/n-transactions-total). Given this number of transactions, we believe that those who care about cryptocurrencies should also understand mobile money.

What Is the State of Security?

Most first-generation mobile money systems were built on widely deployed 2G GSM cellular networks. These services relied on either SMS or Unstructured Supplementary Service Data (USSD) channels for communication. These channels are ideal from the perspective of rapid deployment in that they're nearly omnipresent. However, they're problematic from the perspective of security. First, 2G networks generally rely on cipher suites that are known to be weak. Specifically, the A5/1 and A5/2 algorithms protecting the wireless portion of GSM networks can both be cracked with relatively little effort by an adversary. Although A5/1, the stronger of the two ciphers, was believed to provide significantly improved protection, software-defined radio systems capable of cracking this cipher in real time are now available in backpack-sized setups. That means these first-generation services are vulnerable. To make matters worse, many providers instead rely on the A5/0 (that is, no encryption) standard, removing the already low barrier to attack.

Second, even if providers were to universally upgrade their over-the-air cipher suite to A5/3 (a stronger cipher also known as KASUMI, with known theoretical weaknesses but no practical attacks at this time), encryption protecting data in the SMS and USSD channels ends at the base station. That means that in the core network (and potentially over wireless backhauled used to connect remote towers to that core network), an attacker can easily observe and modify transactions

without detection. Moreover, because authentication in GSM networks is unidirectional (that is, device to network, but not network to device), an adversary could easily deploy a so-called “rogue base station” in a busy area and force all connections to pass unencrypted through it.

The most discussed solution in this space has been the SIM Application Toolkit (also known as SIM Toolkit). SIM Toolkit lets providers develop applications directly on SIM cards, thereby overcoming the need to build applications for a massive set of feature phone platforms. Many have proposed adding application-layer encryption to mobile money via SIM Toolkit, but these efforts have largely failed in practice. Providers privately express frustration in ensuring the correct operation of such a solution. Moreover, there’s great difficulty in replacing the massively deployed number of SIM cards, and over-the-air updates haven’t proven to be a successful path for upgrade.

Network and device upgrades represent a second, more viable path to security. The use of 3G and 4G cellular standards (with better encryption options) and smartphones offer the potential for strong protections from both core network and end-to-end perspectives. The first suggestion, while slowly happening, is unlikely to be universal in the near future. The return on investment for ripping out the massively deployed infrastructure and replacing it with an expensive new network is low. That’s not to say that more 3G and 4G infrastructures aren’t being deployed; rather, the pace at which they’re being rolled out is slow in the developing world. More critically, these networks don’t provide end-to-end cryptographic protection of user data flows, meaning that a total replacement of all 2G networks alone wouldn’t solve the security problems discussed earlier.

Much of our research has focused on mobile money applications for smartphones because they represent the most practical and rapid path to security. Smartphones come equipped with libraries containing an array of strong encryption algorithms, making it possible for developers to quickly and correctly provide end-to-end security for their applications. In 2015, we undertook a major effort to measure how well such mechanisms were being used.⁴ What we found was disheartening. Using a combination of automated and manual analysis, we discovered widespread misuse of insecure protocols, failure to properly authenticate users and mobile money entities, and poor SSL/TLS configuration on back-end servers (among many other issues). Our comprehensive teardown of seven applications revealed that we could steal money from six of them with ease. Moreover, the terms of service in all these applications made customers responsible for all fraud, even though we demonstrated that funds could be stolen without any negligence (for instance, giving out their PIN) on the consumers’ part.

These weaknesses were covered in news outlets including the *Wall Street Journal*, and we worked diligently behind the scenes to provide each of the at-risk companies with detailed vulnerability reports. We also worked with the GSMA and the ITU to spread word of the problems as well as how they could be addressed at low cost (for instance, correct configuration or code updates). However, when we remeasured the applications a year later, we saw not only that the majority of vulnerabilities hadn’t been fixed (in spite of promises to the contrary) but also that development of new features and interfaces had proceeded significantly.⁵

Much remains to be done by the research community. We need

to make it harder to design applications that use insecure communications. Although Android took significant steps forward in this space, the amount of insecure code and security bypasses discovered in the recovered code means that we aren’t there yet. Mechanisms that prevent the submission of applications that fail to properly use TLS would be great, but creating tools to do this will require extremely careful design. Moreover, because of the lack of an obvious push to replace feature phones and 2G networks, easy-to-deploy protocols and solutions are critical. Too many academics view GSM networks and feature phones as “solved” problems, but the reality is that like any massively deployed infrastructure (think COBOL in banking or the magnetic stripe on credit cards), they will never fully be removed from service, especially in the developing world.

What Is the State of Privacy?

Mobile money creates new privacy challenges. Whereas traditional banks are limited to seeing exchanges between their customers and vendors, the peer-to-peer nature of mobile money systems means that providers can observe additional social interactions. For instance, a group of people eating a meal together might send money to one another. Whereas traditional payment systems would have allowed a bank to see that all such attendees were at a restaurant at the same time, mobile money transaction data could be used to definitively link these attendees. Smartphone platforms also offer mobile money applications access to a wealth of additional information, including GPS location.

We don’t believe that collecting such data is inherently problematic. In fact, it’s being used as a means of bootstrapping emerging

credit offerings. In settings in which traditional metrics for determining credit-worthiness aren't available (for example, citizens might not file tax returns or have an official address, a mortgage, or an official history of payments), such data is beginning to act as a substitute. M-Shwari, which offers interest-bearing savings and loans to M-Pesa customers in Kenya, uses M-Pesa usage history to develop credit scores. Such loans have proven critical to merchants, who can eliminate the cash flow issues that traditionally made fully stocking their shelves a challenge.

We believe that consumers should be made aware of how their data is being collected and used and, therefore, be able to make informed decisions when selecting a mobile money or digital credit service. As such, our most recent research efforts have focused on a comprehensive study of privacy policies for mobile money applications.⁶ We collected privacy policies for all 54 mobile Android-based money applications listed by the GSMA and compared these policies to those of the top 50 US financial institutions as listed by the Federal Deposit Insurance Corporation (FDIC; an independent government body in the US responsible for providing regulation for the nation's banks, insurance for deposits, and consumer protection). Although many in the privacy community have opined about what financial privacies should look like ideally, in our evaluation, we relied instead on GSMA and FDIC recommendations. This was important because it let us measure compliance with their communities' published standards.

The results of this were similarly discouraging. Of the 54 studied mobile money applications, only 30 (54 percent) had privacy policies at all. A full third of those that had policies weren't written in either of the two most common languages spoken

in the country, meaning that many in the targeted customer demographics would simply be unable to read such terms. Finally, in the cases in which privacy policies were available, many were too short to contain meaningful content (for instance, EcoCash and TigoPesa's policies were 68 and 268 words long, respectively), or they lacked any mention of critical issues (for instance, fewer than half of those with policies had definitions of terms, mentions of accountability and enforcement, or data retention policies). Finally, mobile money privacy policies also tended to be more difficult to read according to several grade-level readability tests (for example, the Gunning-Fogg index). Given lower literacy rates in many of the populations served by mobile money applications, these results were troubling.

These results were in stark comparison to the traditional financial institutions, which were directly regulated by the FDIC. Mobile money systems, however, generally don't fall under the same regulatory bodies as financial institutions. Adding regulations isn't a simple solution. Many mobile money applications offer low transaction costs because their compliance costs are low. Moreover, these systems exist across a wide array of countries, each with cultures that hold different values to individual data privacy. Accordingly, creating a single set of strong privacy standards that meet universal approval is unlikely to be successful. We instead recommend that the industry push for stronger enforcement of the ideals put forth by the GSMA. Methods and tools for ensuring such compliance, however, remain a research challenge.

The rate at which mobile money systems are bringing traditionally unbanked populations into the global financial infrastructure is unprecedented and absolutely requires new ways of reasoning about and enforcing consumer protection.

We're firm believers in the transformative power of mobile money systems. We also believe that they will connect the finances of the developed and developing worlds in the most meaningful way yet in human history. Accordingly, the price for getting security and privacy wrong is extremely high.

Meaningfully addressing these challenges will require the efforts of our large community. We're trying to expand our engagement through an upcoming NSF-sponsored workshop entitled "Addressing the Technical Security Challenges of Emerging Digital Financial Services." Here, we hope to engage some of the top academic and industrial minds in the details of the challenges we've listed here. Other issues are also critical to address, including how to establish programming interfaces that let developers securely perform critical financial functions in mobile applications; how to ensure the security of legacy 2G infrastructure; and how to address the usability gap when populations with limited literacy and exposure to finance, who represent some of the populations most vulnerable to fraud, are using mobile money. Successfully addressing these problems will require a unique and sustained effort among academia, industry, and nongovernmental organizations. ■

Acknowledgments

This work was supported in part by the NSF under grants CNS-1526718 and CNS-1540217. Any opinions, findings, and conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

References

1. "2015 FDIC National Survey of Unbanked and Underbanked Households," Fed. Deposit Insurance Corporation, Oct. 2016; www.fdic.gov/householdsurvey/2015/2015execsumm.pdf.

2. "GSMA State of the Industry Report on Mobile Money, Decade Edition: 2006–2016," GSM Assoc., 2017; www.gsma.com/mobilefordevelopment/wp-content/uploads/2017/03/GSMA_State-of-the-Industry-Report-on-Mobile-Money_2016.pdf.
3. B. Ngugi, "M-Pesa Global Transactions Hit Six Billion in 2016," Business Daily, 26 Feb. 2017; www.businessdailyafrica.com/markets/MPesa-global-transactions-hit-six-billion-2016/539552-3828662-7h9g3xz/index.html.
4. B. Reaves et al., "Mo(bile) Money, Mo(bile) Problems: Analysis of Branchless Banking Applications in the Developing World," *Proc. USENIX Security Symp. (SEC 15)*, 2015, pp. 17–32; www.cise.ufl.edu/~traynor/papers/reaves-usenix15a.pdf.
5. B. Reaves et al., "Mo(bile) Money, Mo(bile) Problems: Analysis of

Branchless Banking Applications in the Developing World," *ACM Trans. Privacy and Security*, vol. 20, no. 3, 2017 article 10.

6. J. Bowers et al., "Regulators, Mount Up! Analysis of Privacy Policies for Mobile Money Services," *Proc. USENIX Symp. Usable Privacy and Security (SOUPS 17)*, 2017; www.cise.ufl.edu/~traynor/papers/bowers-soups17.pdf.

Patrick Traynor is the John and Mary Lou Dasburg Preeminent Chair in Engineering and associate professor at the University of Florida. He is also a Fellow of the Center for Financial Inclusion at Accion. Contact him at traynor@ufl.edu.

Kevin Butler is an associate professor at the University of Florida. He also serves as the vice chairman and leader of Security Workstream

for the International Telecommunication Union's Focus Group on Digital Financial Services. Contact him at butler@ufl.edu.

Jasmine Bowers is a PhD student at the University of Florida. Contact her at jdbowers@ufl.edu.

Bradley Reaves is an assistant professor at North Carolina State University. Contact him at bgreaves@ncsu.edu.

myCS


Read your subscriptions through the myCS publications portal at <http://mycs.computer.org>



Call for Software Engineering Award Nominations

Established in memory of Harlan D. Mills to recognize researchers and practitioners who have demonstrated long-standing, sustained, and impactful contributions to software engineering practice and research through the development and application of sound theory. The award consists of a \$3,000 honorarium, plaque, and a possible invited talk during the week of the annual International Conference on Software Engineering (ICSE), co-sponsored by the IEEE Computer Society Technical Council on Software Engineering.

Deadline for 2018 Nominations:
1 October 2017

Nomination site:
awards.computer.org
IEEE  computer society

*The award nomination requires at least 3 endorsements.
Self-nominations are not accepted.
Nominees/nominators do not need
to be IEEE or IEEE Computer Society members.*