Failures of Tamper-Proofing in PIN Entry Devices

Bank customers are forced to rely on PIN entry devices in stores and bank branches to protect account details. The authors examined two market-leading devices and found them easy to compromise owing to both their design and the processes used to certify them as secure.

Saar Drimer, Steven J. Murdoch, and Ross Anderson
University of Cambridge

Smart cards are replacing magnetic strip cards for point-of-sale and ATM payments in many countries. The leading system, EMV (originally from Europay, MasterCard, and Visa), has been deployed throughout most of Europe and is currently being rolled out in Canada. With EMV, customers authorize a transaction by inserting a bank smart card and entering a PIN into a PIN entry device (PED); the smart card verifies this PIN, and then a public-key certificate authenticates it to the PED. The card issuer might further authenticate transactions online.

The move from magnetic strip to chip has reduced the use of counterfeit cards domestically, but rising fraud abroad has more than compensated. The deployed system's inadequacies have thus affected many people. According to the Association for Payment Clearing Services (APACS), the UK banks' trade association, 2008 saw £169.8 million of fraud due to counterfeit cards, up 18 percent from the 2007 figure.[1]

To explore the causes of some of this fraud, we examined market-leading PEDs to assess their abilities to resist tampering and protect cardholders. This work is part of a larger research program to examine the EMV system's strengths and weaknesses, and how fraud patterns have changed in response to its introduction. We present a shortened version of our original report[2] in this article.

Real-World Failures in Tamper-Proofing

For backward compatibility, cards in the UK have both a chip and magnetic strip; the strip is used in ATMs without chip readers or when the chip is unreadable. Thus, a criminal who learns both the magnetic strip's contents and a cardholder's PIN can make a magnetic-strip copy and withdraw cash by causing an ATM to fall back to the older system, or by using the copy in a country that hasn't adopted EMV, such as the US. The chip stores a copy of the magnetic strip in its public-key certificate, which is sent to terminals with every transaction. So, PEDs' anti-tampering mechanisms must protect not only PINs that cardholders enter but also card details. (PEDs can also contain symmetric keys that protect communication between the PED and the bank, but these are outside the EMV protocol.)

Merchants and corrupt employees have free access to PEDs, and customers sometimes have access long enough to tamper with them. We must assume, therefore, that the PED operates in an uncontrolled environment and must thus protect card details and PINs, subject to assumptions about attacker capabilities defined in certification criteria. Because European bank customers don't, in general, enjoy the consumer protection that US law affords to American bank customers, they're routinely accused of negligence when they complain of fraud: the bank will often say "your card and your PIN were used, so you must have let someone get hold of your card and learn your PIN, contrary to our terms and conditions." This creates a moral hazard: the PED must protect the cardholder, yet the merchant purchases it from a list of bank-approved devices.

We examined the most widely deployed PEDs in the UK—the Ingenico i3300 (www.ingenico.com/

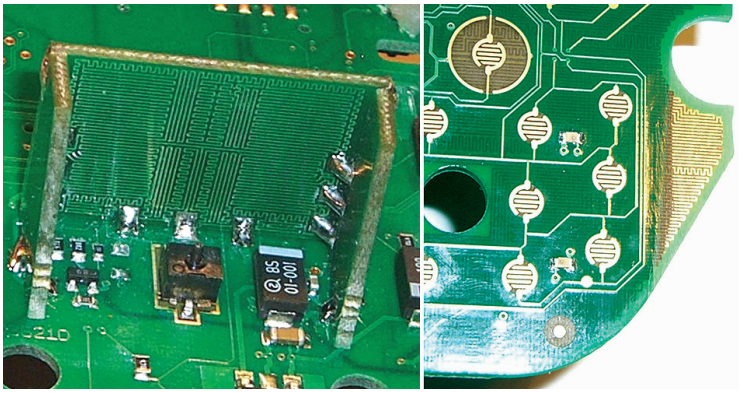

(a) (b)

Figure 1. Tamper-response mechanisms in the Ingenico PIN entry device. A sensor mesh (a) extends to a wall that protects a lid switch and (b) spans an entire circuit board layer.

i3300-i3300_28.html?lg=UK&productId=14#0) and the Dione Xtreme (now branded as VeriFone; www.verifone.com/products/devices/countertop/xtreme.html), each obtained online for less than US$20. We found that they both appear to protect bank and merchant secrets well, yet leave customer card details and PINs inadequately protected.

Both terminals have passed the Visa PED evaluation, which requires that the terminal meet one of four alternative requirements (that defeating tamper-detection would cost more than $25,000 per PED; that the PED would detect the insertion of a PIN-stealing bug, or that such an insertion would take more than 10 hours or cost more than $25,000).[3] Neither terminal actually meets any of these requirements. The Ingenico device also passed the APACS PED Common Criteria evaluation, which requires that "the [security function] shall resist physical attacks based on addition of any PIN-tapping device to the PIN Entry Device and Card Reader by (selection: providing the capability to detect such attacks with a high probability, automatically responding such that the [security policy] is not violated)."[4] Again, the Ingenico device clearly fails this evaluation criterion. We'll next examine how.

Anti-Tampering Mechanisms

The Ingenico PED's enclosure is made from two plastic shells attached to each other by four Torx 6 starhead screws, possibly intended to discourage casual opening. Opening the shell releases a tamper-response switch and breaks a supervisory circuit (Figure 1a). One entire internal circuit board layer is a dense sensor mesh intended to detect drilling from the PED's rear. This mesh extends to a three-sided wall that protects the switch from drilling through a user-accessible compartment (Figure 1a). Additionally, the top shell presses on four contacts (one of which is shown in Figure 1b) to detect the keypad panel's removal. The contacts are surrounded by a conductive ring connected to the battery supply, presumably to prevent attackers from defeating the mechanism by injecting a conductive liquid. The processing module is gift-wrapped with a coarse sensor mesh, then potted.

The Dione PED is ultrasonically sealed at seven interlocking plastic joints and has a simple pad that shorts a contact to detect if it's opened. Unlike the Ingenico PED, it has no mechanisms to detect drilling from the rear (the designers even provide easily accessible circuit board pads to bypass the tamper-detection mechanism). However, the main processing unit and the keypad are potted together, which makes it harder to capture PIN keystrokes between the keypad and the processor.

In both designs, the secure storage for cryptographic keys appears fairly well protected. However, in each case, you can tap the data line of the PED-smart-card interface. The data exchanged on this line isn't encrypted;[5] it yields both the information we need to create a fake magnetic-strip card and the PIN to use with it.

Signal Eavesdropping Attack

We defeated the Ingenico PED with a simple tapping attack thanks to a succession of design flaws. Its rear has a user-accessible compartment, shown in Figure 2a, that was intended to accommodate optional SIM-sized cards to expand its functionality. This space isn't intended to be tamper-proof, and, when it's covered, the cardholder can't inspect it even if she handles the PED. This compartment gives access to signals routed on its bottom layer, although the sensor-mesh layer mentioned earlier prevents drilling through the circuit board to access the smart card's data line. Curiously, however, there is no need to drill. The PED's designers provide holes one millimeter in diameter and other vias through the circuit board. The holes are used for positioning optional surface-mount sockets, none of which was populated in the PEDs we examined. Through one of these holes, a simple metal hook can tap the serial data line between the microprocessor and the card interface chip. We preferred, however, to tap the signal before the interface chip, and found that we could easily access a 1-mm via carrying the data signal using a bent paperclip. We can insert this through a hole in the plastic surrounding the internal compartment without leaving external marks.

Having tested this attack in the laboratory, we repeated it in the field for the *BBC Newsnight* program; we tapped a terminal in a London shop and, during a transaction, extracted the card and PIN details for a journalist's card without triggering the tamper-detection system.

(a) (b)

Figure 2. Tapping attacks on working Ingenico and Dione PIN entry devices (PEDs). (a) A paperclip inserted through a hole in the Ingenico's concealed compartment wall can intercept the smart card's data. The inset shows the front of the PED. (b) We inserted a needle through the rear of the Dione PED for data interception, attaching it to a ribbon cable connector shown on the bottom left inset; the top-right inset shows a mounted PED.

The Dione PED doesn't provide a concealed compartment to hide the wiretap but is still vulnerable. By drilling a 0.8-mm hole from the rear, we can insert a 4-cm needle into a flat ribbon connector socket (shown in Figure 2b).

What should have required $25,000 needed just a bent paperclip, a needle, a short length of wire, and some creative thinking; attaching a probe to the data line takes minutes with some practice. A small field-programmable gate array (FPGA) or microcontroller board with some nonvolatile memory can easily fit inside the Ingenico PED's compartment and record thousands of transaction details without the cardholder's knowledge. A wire routed from the back of a mounted Dione PED to a recorder under the counter won't be detected unless the cardholder conducts a very close inspection—and knows what to look for.

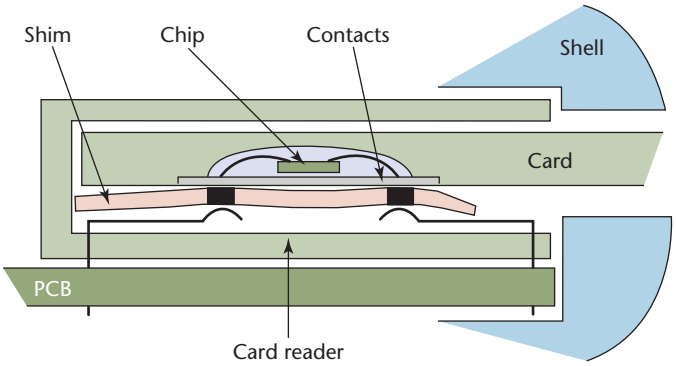

Figure 3. Shim-in-the-middle attack. A flexible circuit board placed between the card and card-reader contacts transmits transaction details to a nearby receiver. A fraudster could use low-profile components in the reader to create a simple transmitter.

Shim-in-the-Middle Attack

We postulate, but have yet to implement, an attack in which we insert a thin, flexible circuit board into the card slot so that it lodges between the reader and the card's contacts. This attack completely bypasses all tamper protections and doesn't even require the participation of anyone in the store. Figure 3 illustrates this "shim in the middle"; a very basic circuit that can transmit the signal on the data line to a nearby receiver wouldn't be easily detected, being within the PED itself. The fraudster can create an "inserter card" with the shim attached to it so that, when inserted into a particular device, the shim locks into place as the crook removes the carrier card. He would then place a receiver nearby to record card details and PINs; this receiver could easily include a mobile phone to SMS the data back to its master.

Defenses and Attack Extensions

We believe that the interface between smart card and PED simply can't be adequately protected and that EMV is flawed in that it permits unencrypted PIN transfer over this exposed communication line. Essentially, the vulnerabilities we exploited aren't a consequence only of hardware design but also of many banks opting to disable PIN encryption when they implemented EMV. Thus, some upgrade options and

mitigation techniques are possible—although not all are as effective as they might first appear. Let's look at some of these options.

Encrypted PIN

Our attack reads data as it passes between the PED and the card. If both card and PED support it, EMV lets the PED encrypt the PIN under the card's public key. Cards currently issued in the UK don't support this because banks chose low-cost cards that can't do asymmetric cryptography. Upgraded cards, with encrypted PIN capability, will prevent a passive eavesdropper from observing the PIN (though card details still pass unencrypted in the other direction).

However, due to a quirk in the EMV implementation, attackers can sometimes bypass PIN encryption. A card advertises that it supports encrypted PIN verification by placing an appropriate entry in the cardholder verification method (CVM) list, which it sends at the start of the transaction. In eight out of 15 cards we examined, the CVM list isn't signed and so can be modified—causing the PED to send the PIN unencrypted.

Fraudsters can conduct this attack using an active tap that selectively alters the communication, forcing a `HIGH` bit to `LOW` so that the PED thinks the card can't process encrypted PINs and sends the cardholder's PIN in the clear.

If we can implement a full middleman scenario, a more sophisticated attack might defeat even signed CVM lists. Here, the attack device impersonates an entirely different card to the PED at the transaction's start and presents a CVM list that allows unencrypted PIN entry. Once the customer has entered his or her PIN and it's been intercepted, the attack device causes the PED to restart the normal transaction. At worst, this looks like an intermittent error; in some PED implementations, it might be possible to avoid alerting the customer at all.

This attack shows that evaluators should consider active attacks, too. All the specifications we've examined appear to consider only passive taps. But there might be some mileage in anti-tampering measures that prevent the communication path from being broken, and where the card or PED checks if the data sent has been corrupted—a more feasible task than detecting passive taps. Protocol defenses are also possible: displaying the cardholder name from the card's certificate on the PIN entry prompt would let alert customers detect some middleman attacks.

CVV for Integrated Circuit Cards

The backward-compatibility feature whereby the card certificate contains a copy of the magnetic-strip data is a serious vulnerability. Visa has therefore proposed replacing the card verification value (CVV)—a cryptographic checksum stored on the magnetic strip—with a different one in the certificate: the CVV for Integrated Circuit Cards (iCVV).

With iCVV implemented, a fraudster will have to swipe the card in addition to reading the chip, thus reducing the risk from some of the vulnerabilities we discuss here. We strongly support iCVV deployment, but despite Visa's making its recommendation in 2002, and APACS stating that iCVV was mandatory beginning in January 2008, banks were still issuing cards in 2008 that store an exact copy of the magnetic strip on the chip.

The Certification Process

Until recently, market forces could exercise some discipline on vendors. For example, in 2006, Shell withdrew EMV terminals from its UK fuel stations following a fraud that involved PED tampering and fell back for some months on magnetic-strip processing.[6] Its PED vendor, Trintech, sold its terminal business to VeriFone and left the market.[7] Since then, however, rapid consolidation has occurred, with Ingenico and VeriFone now apparently controlling most of the market. In addition, all but the largest merchants tend to get their terminals from their bank, many of which offer only one make of terminal.

So, customers and now merchants depend critically on the certification of terminals, PEDs, smart cards, and other system components that the EMV system uses. Some certification schemes merely ensure compatibility, such as EMV level 1 (see www.emvco.com), but there are also extensive security evaluations. As mentioned, both PEDs we examined are certified under the Visa PED approval scheme,[8] and the Ingenico PED passed the APACS PED Common Criteria evaluation,[9] despite the vulnerabilities we identified. What does that tell us about the evaluation and certification process?

Why Evaluations Fail

A security failure in an evaluated product can have numerous causes. The Common Criteria (or other framework) might be defective; the protection profile might not specify adequate protection; the evaluator might miss attacks or estimate their cost and complexity as too high. One known problem with the Common Criteria is the proliferation of protection profiles. Anyone can propose a protection profile for any purpose and get a lab to evaluate it. The result is a large number of profiles giving little assurance of anything—for example, the profile for ATMs is written in management-speak (complete with clip art), states that it "has elected not to include any security policy," and misses many of the problems that were well known when it was written in 1999. Indeed, it states that it relies on developers to document vulnerabilities

and includes the vague statement that "the evaluator shall determine that the [target of evaluation] is resistant to penetration attacks performed by an attacker possessing a moderate attack potential."[10]

A deeper problem in the security evaluation process is the economics involved. Since the demise of the philosophy behind the Orange Book,[11] device manufacturers now select and pay laboratories to perform evaluations. The vendor will naturally select the lab that will give its product the easiest ride and charge the least money. What's more, the same process applies to the protection profiles against which the product is evaluated.

Market competition might help reduce evaluation costs, but it promotes a race to the bottom between the labs. To mitigate this, vendors must use approved labs, selected by Visa in the case of PED approval, or by a national body such as the US National Institute of Standards and Technology or Britain's Government Communications Headquarters (GCHQ) for the Common Criteria. In principle, this might provide some quality control, but in practice the agencies appear to have never revoked a lab's license for fear of undermining confidence in "the system." Government agencies might also feel reluctant to drive evaluation work abroad. Are the evaluation failures described here systemic or the fault of an individual evaluator?

Government and Industry Response

We wrote to GCHQ, Visa, APACS, Ingenico, and VeriFone (Dione) in November 2007 and asked them to comment on our findings. We asked for copies of the evaluation reports, why these reports weren't public, whether the insecure PEDs would be decertified, whether the labs that negligently certified them as secure would lose their licenses, and whether the evaluation system should be changed. The imminent broadcast, on 26 February 2008, of the BBC program where we demonstrated our research prompted responses from GCHQ, APACS, and VeriFone, while Ingenico and Visa remained silent.

VeriFone's response was evasive, pushing responsibility to APACS, Visa, and GCHQ, but the replies from APACS and GCHQ were more instructive. APACS, the bankers' trade association, claimed that previous evaluations "did not identify any specific vulnerabilities in the devices that required additional mitigation"; it denied that the evaluations were defective and said it wouldn't withdraw the devices from use because it disagreed with our risk assessment: according to APACS, the attack was harder than we described and uneconomical for criminals, and there were simpler fraud methods.

APACS claimed that "the numbers of PED compromise that have taken place in the UK are minimal, however, and the banking industry's standard fraud prevention measures have meant that these frauds and their location were detected quickly." (In one recent court case, the defendant was convicted of £2 million worth of fraud from PED tampering, with the potential for a further £16 million.[12]) APACS refused to name the evaluation labs and insisted that evaluations must be carried out under nondisclosure agreements. It said, "we are not aware of any widely recognized and credible evaluation methodology process, in security or otherwise, which makes evaluation reports publicly available."

GCHQ's response was equally uncompromising but totally different. It informed us that evaluation reports for Common-Criteria-certified devices must be made public as a condition of the mutual-recognition arrangement under which evaluations performed in one of the Common Criteria countries are recognized in others. It transpired that the Ingenico device was merely "evaluated" under Common Criteria, not "certified" and hence wasn't subject to GCHQ oversight or that of any other country's certification body (CB). All certified products are listed on the Common Criteria Portal (www.commoncriteriaportal.org), although it was confusingly titled "List of Evaluated Products." Following the initial publication of our article, this was renamed the "Certified Product List," and as of November 2009, no PEDs are present on it.

In short, APACS performed the certification for the Ingenico PED on the basis of a secret report by an undisclosed laboratory. A CB licensed this laboratory to perform certifications, but APACS refused to identify the country in which the lab was registered and hence which CB was responsible. Had the devices been through certification, the CB would have been responsible for ensuring that the security target was appropriate and that the lab had conducted proper testing. APACS said that the decision about whether to revoke the laboratory's license is the responsibility of the CB that registered it. But because the PED evaluation was done outside the Common Criteria system—and as far as we know without any CB's knowledge—it's unclear how an errant lab could ever be disciplined.

As a CB, GCHQ doesn't object to anyone calling any device "Common Criteria Evaluated" and will merely object if a false claim is made that a device is "Common Criteria Certified." This undermines its brand and enables organizations such as APACS to free-ride by exploiting the Common Criteria name without either evaluating products rigorously or publishing the results. GCHQ admits that as the licensing authority it has an interest: "The CB then has a direct involvement in maintaining the quality of each of the individual evaluations for certification. These mechanisms counter any tendency for such a 'race to the bot-

tom.'" Regrettably, its confidence is inconsistent with our research results.

For both devices, the proximate cause of evaluation failure was that the equipment didn't meet the protection goals set out in either the Visa certification requirements or the APACS Common Criteria protection profile. A deeper cause was that these requirements were unrealistic; given the shim attack, it's just not clear that anyone could construct a compact, low-cost device that meets either set of requirements, so the labs might have faced an impossible task. We'd argue that the protection profile should never have assumed that it was possible to protect the card–PED interface at all.

The banks clearly had an incentive to pretend that it could be protected—by using cheap smart cards rather than the more expensive ones, which are capable of asymmetric cryptography, they saved perhaps $1 per card over 70 million accounts. GCHQ's failure to protect the Common Criteria brand let the banking industry describe insecure terminals as Common Criteria Evaluated without legal penalty.

EMV's failure was multifactorial: too many protocol options, liability dumping, an overly optimistic protection profile, vendor-funded evaluations, and failures of both markets and regulation at several levels. What should be done about it?

Fixing the Evaluation Process

Unfortunately, Common Criteria evaluations seem to be most prevalent where incentives are skewed—that is, where one principal operates a system but others bear the costs of failure. This tempts operators to be careless, a phenomenon known to economists as "moral hazard." It's now well known that moral hazard is a major cause of security failure; organizations exposed to it might seek evaluation as a way to avoid blame for failures by demonstrating due diligence.[13]

We believe that the certification process should be re-engineered to take into account incentives and accountability. In an ideal world, representatives for users would conduct evaluations, but in the real world, cardholders and small merchants aren't in a position to act collectively. Where evaluation by the relying party is impractical, the next best option might be a hostile laboratory. The closest we often get to this ideal is an academic evaluation, such as the results we report here. But these evaluations' quantity and timeliness falls far short of the optimum: manufactures have offered more than 200 types of PEDs for sale in Europe, and our work is the first open evaluation.

The industry's attitude toward independent evaluation is at best unhelpful and at worst actively obstructive. Merchants fear that if they're discovered to have assisted in confirming security vulnerabilities, they might face retribution from their bankers. In many ways, criminals are in a better position because they can easily set up fake merchants and be more anonymous than an independent researcher cooperating with a legitimate merchant.

One possible solution is to have a market for PED vulnerabilities, much like the thriving operating system vulnerability market. Furthermore, given the very strong incentives for vendors to shop around for the easiest evaluation lab, the resulting race to the bottom, and the lack of institutional incentives for CBs to exercise proper discipline, we propose that evaluations of equipment the public is forced to rely on should in the future come with a sufficient reward to motivate independent evaluation.

For an evaluation at level EAL3, for example, we propose a mandatory reward of $10,000 for each vulnerability, whereas for EAL4, the reward should be $100,000. Introducing real money will call forth a more socially optimal level of attack effort, while making the rewards conditional on responsible disclosure could control any increase in exposure. What's more, we propose that the rewards be paid not by the vendors nor even the evaluation labs, but by the CBs that license those labs (which must be regulators rather than trade associations such as Visa). This way, careless evaluators will cost their regulators real money and are more likely to be disciplined. (The CBs might in turn require vendors to post performance bonds.)

Criminals are actively exploiting the failings we've described here. Aside from the 2006 Shell case we previously mentioned, Irish criminals were caught in August 2008 installing Bluetooth bugs inside supermarket PEDs while pretending to be service engineers;[14] they would read out the recorded data simply by passing by the till later on. In October 2008, the police exposed a large-scale operation in which criminals installed tiny GSM modules inside PEDs during or soon after manufacture, before they were shipped to merchants.[15] Every so often, these modules called Pakistan to deliver the recorded card details and PINs by SMS and could even receive instructions back. This type of "supply-chain attack" is very hard to prevent or detect.

In our extended report,[2] we provide a security analysis of EMV's failings in more detail and conclude that its several-thousand-page specification is a prime contributor to the failures we've identified. We recommend that such standards' promoters also publish a concise security document that describes not just the threat model but the protection requirements on each component, sufficient for engineers to understand and manage the final product's security complexities. Systems engineering—and indeed computer science—are increasingly about managing complexity;

this will be a growing concern for security engineers, and EMV provides a good case study of how things go wrong in a major security-critical system deployed in the real world.

The lessons we learned aren't limited to banking. Devices in other fields, such as voting machines, suffer from the same combination of stupid mistakes, sham evaluations, and obstructive authorities. Finally, our findings have policy implications as well. The global financial collapse has showed that bank regulators credulously accepted banks' financial models at face value; our experience with EMV shows that they were just as gullible when it came to the banks' security models. Claims that terminals were Common Criteria Evaluated turned out to be misleading, if not fraudulent. The devices in question weren't Common Criteria Certified, and the certification body wasn't interested in protecting its brand. If the Common Criteria brand is to retain any value, its promoters had better defend it. And if information security agencies aren't interested in regulating bank security, we'd better figure out who's going to do it. Will it be the bank regulator or some other government agency, or laws that repair incentives and facilitate a private-sector solution? □

References

1. "2008 Fraud Figures Announced by APACS," Assoc. for Payment Clearing Services, Mar. 2009; www.ukpayments.org.uk/media_centre/press_releases/-/page/685/.
2. S. Drimer, S.J. Murdoch, and R. Anderson, *Thinking Inside the Box: System-Level Failures of Tamper Proofing*, tech. report UCAM-CL-TR-711, Computer Laboratory, Univ. of Cambridge, Feb. 2008.
3. "PIN Entry Device Security Requirements Manual," Visa Int'l Service Assoc., Mar. 2004; https://partnernetwork.visa.com/vpn/global/retrieve_document.do?documentRetrievalId=35.
4. "PIN Entry Device Protection Profile," Assoc. for Payment Clearing Services, July 2003; www.commoncriteriaportal.org/public/files/ppfiles/PED_PPv1_37.pdf.
5. M. Bond, "Chip & PIN (EMV) Interceptor," Mar. 2006; www.cl.cam.ac.uk/research/security/banking/interceptor/.
6. J. Bale, "Shell Halts Chip-and-PIN after Fraud," *The Times*, May 2006; http://business.timesonline.co.uk/tol/business/law/article714402.ece.
7. "VeriFone to Acquire Trintech's Payment Systems Business," Trintech, Aug. 2006; www.trintech.com/verifone-to-acquire-trintechs-payment-systems-business/.
8. "Approved PIN Entry Devices," Visa Int'l Service Assoc.," Oct. 2007; http://partnernetwork.visa.com/dv/pin/pedapprovallist.jsp.
9. "PIN Entry Device Protection Profile Common Criteria Evaluation," the UK Cards Assoc., Sept. 2007, www.theukcardsassociation.org.uk/about_us/what_we_do/technical_services_and_standards/common_criteria_evaluation/.
10. Bull, Dassault, Diebold, NCR, Siemens Nixdorf, and Wang Global, "Protection Profile: Automatic Cash Dispensers/Teller Machines," 1999, www.commoncriteriaportal.org/files/ppfiles/PP9907.pdf.
11. S.L. Brand, "DoD 5200.28-STD Department of Defense Trusted Computer System Evaluation Criteria (Orange Book)," *Nat'l Computer Security Center*, Dec. 1985.
12. S. Bird, "'Catch Me If You Can,' Said Student Behind Biggest Chip and PIN Fraud," *The Times*, Oct. 2009; www.timesonline.co.uk/tol/news/uk/crime/article5034185.ece.
13. R.J. Anderson, *Security Engineering: A Guide to Building Dependable Distributed Systems*, John Wiley & Sons, 2001.
14. C. Lally, "9,000 Credit Cards Illegally Copied in Scam on Stores," *Irish Times*, Aug. 2008; www.irishtimes.com/newspaper/ireland/2008/0819/1218868120438.html.
15. H. Samuel, "Chip and Pin Scam 'Has Netted Millions from British Shoppers,'" *Telegraph*, Oct. 2008; www.telegraph.co.uk/news/newstopics/politics/lawandorder/3173346/Chip-and-pin-scam-has-netted-millions-from-British-shoppers.html.

Saar Drimer *is a researcher at the University of Cambridge Computer Laboratory. His research interests include secure content distribution for reconfigurable systems, hardware security, and banking security. Drimer has a PhD in computer science from the University of Cambridge. Contact him at saar.drimer@cl.cam.ac.uk.*

Steven J. Murdoch *is a senior research associate at the University of Cambridge Computer Laboratory. His research interests include privacy, anonymous communications, and banking security. Murdoch has a PhD from the University of Cambridge. He is also a fellow of Christ's College, Cambridge, and a member of the Tor Project. Contact him at steven.murdoch@cl.cam.ac.uk.*

Ross Anderson *is the professor of security engineering at the University of Cambridge. His research interests range from security protocols and APIs through hardware tamper-resistance and critical national infrastructure to security economics and security psychology. Anderson is a fellow of the Royal Society, the Royal Academy of Engineering, the Institute of Engineering and Technology, and the Institute of Mathematics and its Applications. He's also the author of the standard textbook* Security Engineering—A Guide to Building Dependable Distributed Systems *(Wiley, 2001). Contact him at ross.anderson@cl.cam.ac.uk.*

A version of this article won the IEEE S&P *Outstanding Papers Award at the 2008 IEEE Security and Privacy Symposium.*