Gmail **Richard Lamb <naticklamb@gmail.com>**

CRYPTO-GRAM, June 15, 2021

2 messages

Bruce Schneier <schneier@schneier.com> Mon, Jun 14, 2021 at 9:14 PM
Reply-To: Bruce Schneier <schneier@schneier.com>
To: lamb@ieee.org

Crypto-Gram

June 15, 2021

by Bruce Schneier
Fellow and Lecturer, Harvard Kennedy School
schneier@schneier.com
https://www.schneier.com

A free monthly newsletter providing summaries, analyses, insights, and commentaries on security: computer and otherwise.

For back issues, or to subscribe, visit Crypto-Gram's web page.

Read this issue on the web

These same essays and news items appear in the Schneier on Security blog, along with a lively and intelligent comment section. An RSS feed is available.

** *** ***** ******* ********** **********

In this issue:

If these links don't work in your email client, try reading this issue of Crypto-Gram on the web.

** *** ***** ******* ********** **********

Is 85% of US Critical Infrastructure in Private Hands?

[2021.05.17] Most US critical infrastructure is run by private corporations. This has major security implications, because it's putting a random power company in -- say -- Ohio -- up against the Russian cybercommand, which isn't a fair fight.

When this problem is discussed, people regularly quote the statistic that 85% of US critical infrastructure is in private hands. It's a handy number, and matches our intuition. Still, I have never been able to find a factual basis, or anyone who knows where the number comes from. Paul Rosenzweig investigates, and reaches the same conclusion.

So we don't know the percentage, but I think we can safely say that it's a lot.

** *** ***** ******* ********** **********

Adding a Russian Keyboard to Protect against Ransomware

[2021.05.18] A lot of Russian malware -- the malware that targeted the Colonial Pipeline, for example -- won't install on computers with a Cyrillic keyboard installed. Brian Krebs wonders if this could be a useful defense:

> In Russia, for example, authorities there generally will not initiate a cybercrime investigation against one of their own unless a company or individual within the country's borders files an official complaint as a victim. Ensuring that no affiliates can produce victims in their own countries is the easiest way for these criminals to stay off the radar of domestic law enforcement agencies.
>
> [...]
>
> DarkSide, like a great many other malware strains, has a hard-coded do-not-install list of countries which are the principal members of the Commonwealth of Independent States (CIS) -- former Soviet satellites that mostly have favorable relations with the Kremlin.
>
> [...]
>
> Simply put, countless malware strains will check for the presence of one of these languages on the system, and if they're detected the malware will exit and fail to install.
>
> [...]
>
> Will installing one of these languages keep your Windows computer safe from all malware? Absolutely not. There is plenty of malware that doesn't care where in the world you are. And there is no substitute for adopting a defense-in-depth posture, and avoiding risky behaviors online.
>
> But is there really a downside to taking this simple, free, prophylactic approach? None that I can see, other than perhaps a sinking feeling of capitulation. The worst that could happen is that you accidentally toggle the language settings and all your menu options are in Russian.

** *** ***** ******* ********** **********

Apple Censorship and Surveillance in China

[2021.05.19] Good investigative reporting on how Apple is participating in and assisting with Chinese censorship and surveillance.

** *** ***** ******* ********** **********

Bizarro Banking Trojan

[2021.05.20] Bizarro is a new banking trojan that is stealing financial information and crypto wallets.

> ...the program can be delivered in a couple of ways -- either via malicious links contained within spam emails, or through a trojanized app. Using these sneaky methods, trojan operators will implant the malware onto a target device, where it will install a sophisticated backdoor that "contains more than 100 commands and allows the attackers to steal online banking account credentials," the researchers write.
>
> The backdoor has numerous commands built in to allow manipulation of a targeted individual, including keystroke loggers that allow for harvesting of personal login information. In some instances, the malware can allow criminals to commandeer a victim's crypto wallet, too.

Research report.

** *** ***** ******* ********** **********

Double-Encrypting Ransomware

[2021.05.21] This seems to be a new tactic:

> Emsisoft has identified two distinct tactics. In the first, hackers encrypt data with ransomware A and then re-encrypt that data with ransomware B. The other path involves what Emsisoft calls a "side-by-side encryption" attack, in which attacks encrypt some of an organization's systems with ransomware A and others with ransomware B. In that case, data is only encrypted once, but a victim would need both decryption keys to unlock everything. The researchers also note that in this side-by-side scenario, attackers take steps to make the two distinct strains of ransomware look as similar as possible, so it's more difficult for incident responders to sort out what's going on.

** *** ***** ******* ********** **********

AIs and Fake Comments

[2021.05.24] This month, the New York state attorney general issued a report on a scheme by "U.S. Companies and Partisans [to] Hack Democracy." This wasn't another attempt by Republicans to make it harder for Black people and urban residents to vote. It was a concerted attack on another core element of US democracy -- the ability of citizens to express their voice to their political representatives. And it was carried out by generating millions of fake comments and fake emails purporting to come from real citizens.

This attack was detected because it was relatively crude. But artificial intelligence technologies are making it possible to generate genuine-seeming comments at scale, drowning out the voices of real citizens in a tidal wave of fake ones.

As political scientists like Paul Pierson have pointed out, what happens between elections is important to democracy. Politicians shape policies and they make laws. And citizens can approve or condemn what politicians are doing, through contacting their representatives or commenting on proposed rules.

That's what should happen. But as the New York report shows, it often doesn't. The big telecommunications companies paid millions of dollars to specialist "AstroTurf" companies to generate public comments. These companies then stole people's names and email addresses from old files and from hacked data dumps and attached them to 8.5 million public comments and half a million letters to members of Congress. All of them said that they supported the corporations' position on something called "net neutrality," the idea that telecommunications companies must treat all Internet content equally and not prioritize any company or service. Three AstroTurf companies -- Fluent, Opt-Intelligence and React2Media -- agreed to pay nearly $4 million in fines.

The fakes were crude. Many of them were identical, while others were patchworks of simple textual variations: substituting "Federal Communications Commission" and "FCC" for each other, for example.

Next time, though, we won't be so lucky. New technologies are about to make it far easier to generate enormous numbers of convincing personalized comments and letters, each with its own word choices, expressive style and pithy examples. The people who create fake grass-roots organizations have always been enthusiastic early adopters of technology, weaponizing letters, faxes, emails and Web comments to manufacture the appearance of public support or public outrage.

Take Generative Pre-trained Transformer 3, or GPT-3, an AI model created by OpenAI, a San Francisco based start-up. With minimal prompting, GPT-3 can generate convincing seeming newspaper articles, résumé cover letters, even Harry Potter fan fiction in the style of Ernest Hemingway. It is trivially easy to use these techniques to compose large numbers of public comments or letters to lawmakers.

OpenAI restricts access to GPT-3, but in a recent experiment, researchers used a different text-generation program to submit 1,000 comments in response to a government request for public input on a Medicaid issue. They all sounded unique, like real people advocating a specific policy position. They fooled the Medicaid.gov administrators, who accepted them as genuine concerns from actual human beings. The researchers subsequently identified the comments and asked for them to be removed, so that no actual policy debate would be unfairly biased. Others won't be so ethical.

When the floodgates open, democratic speech is in danger of drowning beneath a tide of fake letters and comments, tweets and Facebook posts. The danger isn't just that fake support can be generated for unpopular positions, as happened with net neutrality. It is that public commentary will be completely discredited. This would be bad news for specialist AstroTurf companies, which would have no business model if there isn't a public that they can pretend to be representing. But it would empower still further other kinds of lobbyists, who at least can prove that they are who they say they are.

We may have a brief window to shore up the flood walls. The most effective response would be to regulate what UCLA sociologist Edward Walker has described as the "grassroots for hire" industry. Organizations that deliberately fabricate citizen voices shouldn't just be subject to civil fines, but to criminal penalties. Businesses that hire these organizations should be held liable for failures of oversight. It's impossible to prove or disprove whether telecommunications companies knew their subcontractors would create bogus citizen voices, but a liability standard would at least give such companies

an incentive to find out. This is likely to be politically difficult to put in place, though, since so many powerful actors benefit from the status quo.

This essay was written with Henry Farrell, and previously appeared in the *Washington Post*.

EDITED TO ADD: CSET published an excellent report on AI-generated partisan content. Short summary: it's pretty good, and will continue to get better. Renee DeRista has also written about this.

This paper is about a lower-tech version of this threat. Also this.

EDITED TO ADD: Another essay on the same topic.

** *** ***** ******* ********** **********

New Disk Wiping Malware Targets Israel

[2021.05.26] Apostle seems to be a new strain of malware that destroys data.

> In a post published Tuesday, SentinelOne researchers said they assessed with high confidence that based on the code and the servers Apostle reported to, the malware was being used by a newly discovered group with ties to the Iranian government. While a ransomware note the researchers recovered suggested that Apostle had been used against a critical facility in the United Arab Emirates, the primary target was Israel.

** *** ***** ******* ********** **********

The Story of the 2011 RSA Hack

[2021.05.27] Really good long article about the Chinese hacking of RSA, Inc. They were able to get copies of the seed values to the SecurID authentication token, a harbinger of supply-chain attacks to come.

** *** ***** ******* ********** **********

The Misaligned Incentives for Cloud Security

[2021.05.28] Russia's Sunburst cyberespionage campaign, discovered late last year, impacted more than 100 large companies and US federal agencies, including the Treasury, Energy, Justice, and Homeland Security departments. A crucial part of the Russians' success was their ability to move through these organizations by compromising cloud and local network identity systems to then access cloud accounts and pilfer emails and files.

Hackers said by the US government to have been working for the Kremlin targeted a widely used Microsoft cloud service that synchronizes user identities. The hackers stole security certificates to create their own identities, which allowed them to bypass safeguards such as multifactor authentication and gain access to Office 365 accounts, impacting thousands of users at the affected companies and government agencies.

It wasn't the first time cloud services were the focus of a cyberattack, and it certainly won't be the last. Cloud weaknesses were also critical in a 2019 breach at Capital One. There, an Amazon Web Services cloud vulnerability, compounded by Capital One's own struggle to properly configure a complex cloud service, led to the disclosure of tens of millions of customer records, including credit card applications, Social Security numbers, and bank account information.

This trend of attacks on cloud services by criminals, hackers, and nation states is growing as cloud computing takes over worldwide as the default model for information technologies. Leaked data is bad enough, but disruption to the cloud, even an outage at a single provider, could quickly cost the global economy billions of dollars a day.

Cloud computing is an important source of risk both because it has quickly supplanted traditional IT and because it concentrates ownership of design choices at a very small number of companies. First, cloud is increasingly the default mode of computing for organizations, meaning ever more users and critical data from national intelligence and defense agencies ride on these technologies. Second, cloud computing services, especially those supplied by the world's four largest providers -- Amazon, Microsoft, Alibaba, and Google -- concentrate key security and technology design choices inside a small number of organizations. The consequences of bad decisions or poorly made trade-offs can quickly scale to hundreds of millions of users.

The cloud is everywhere. Some cloud companies provide software as a service, support your Netflix habit, or carry your Slack chats. Others provide computing infrastructure like business databases and storage space. The largest cloud companies provide both.

The cloud can be deployed in several different ways, each of which shift the balance of responsibility for the security of this technology. But the cloud provider plays an important role in every case. Choices the provider makes in how these

technologies are designed, built, and deployed influence the user's security -- yet the user has very little influence over them. Then, if Google or Amazon has a vulnerability in their servers -- which you are unlikely to know about and have no control over -- you suffer the consequences.

The problem is one of economics. On the surface, it might seem that competition between cloud companies gives them an incentive to invest in their users' security. But several market failures get in the way of that ideal. First, security is largely an externality for these cloud companies, because the losses due to data breaches are largely borne by their users. As long as a cloud provider isn't losing customers by the droves -- which generally doesn't happen after a security incident -- it is incentivized to underinvest in security. Additionally, data shows that investors don't punish the cloud service companies either: Stock price dips after a public security breach are both small and temporary.

Second, public information about cloud security generally doesn't share the design trade-offs involved in building these cloud services or provide much transparency about the resulting risks. While cloud companies have to publicly disclose copious amounts of security design and operational information, it can be impossible for consumers to understand which threats the cloud services are taking into account, and how. This lack of understanding makes it hard to assess a cloud service's overall security. As a result, customers and users aren't able to differentiate between secure and insecure services, so they don't base their buying and use decisions on it.

Third, cybersecurity is complex -- and even more complex when the cloud is involved. For a customer like a company or government agency, the security dependencies of various cloud and on-premises network systems and services can be subtle and hard to map out. This means that users can't adequately assess the security of cloud services or how they will interact with their own networks. This is a classic "lemons market" in economics, and the result is that cloud providers provide variable levels of security, as documented by Dan Geer, the chief information security officer for In-Q-Tel, and Wade Baker, a professor at Virginia Tech's College of Business, when they looked at the prevalence of severe security findings at the top 10 largest cloud providers. Yet most consumers are none the wiser.

The result is a market failure where cloud service providers don't compete to provide the best security for their customers and users at the lowest cost. Instead, cloud companies take the chance that they won't get hacked, and past experience tells them they can weather the storm if they do. This kind of decision-making and priority-setting takes place at the executive level, of course, and doesn't reflect the dedication and technical skill of product engineers and security specialists. The effect of this underinvestment is pernicious, however, by piling on risk that's largely hidden from users. Widespread adoption of cloud computing carries that risk to an organization's network, to its customers and users, and, in turn, to the wider internet.

This aggregation of cybersecurity risk creates a national security challenge. Policymakers can help address the challenge by setting clear expectations for the security of cloud services -- and for making decisions and design trade-offs about that security transparent. The Biden administration, including newly nominated National Cyber Director Chris Inglis, should lead an interagency effort to work with cloud providers to review their threat models and evaluate the security architecture of their various offerings. This effort to require greater transparency from cloud providers and exert more scrutiny of their security engineering efforts should be accompanied by a push to modernize cybersecurity regulations for the cloud era.

The Federal Risk and Authorization Management Program (FedRAMP), which is the principal US government program for assessing the risk of cloud services and authorizing them for use by government agencies, would be a prime vehicle for these efforts. A recent executive order outlines several steps to make FedRAMP faster and more responsive. But the program is still focused largely on the security of individual services rather than the cloud vendors' deeper architectural choices and threat models. Congressional action should reinforce and extend the executive order by adding new obligations for vendors to provide transparency about design trade-offs, threat models, and resulting risks. These changes could help transform FedRAMP into a more effective tool of security governance even as it becomes faster and more efficient.

Cloud providers have become important national infrastructure. Not since the heights of the mainframe era between the 1960s and early 1980s has the world witnessed computing systems of such complexity used by so many but designed and created by so few. The security of this infrastructure demands greater transparency and public accountability -- if only to match the consequences of its failure.

This essay was written with Trey Herr, and previously appeared in *Foreign Policy*.

** *** ***** ******* ********** **********

Security Vulnerability in Apple's Silicon "M1" Chip

[2021.06.01] The website for the M1racles security vulnerability is an excellent demonstration that not all vulnerabilities are exploitable. Be sure to read the FAQ through to the end.

EDITED TO ADD: Wired article.

** *** ***** ******* ********** **********

The DarkSide Ransomware Gang

[2021.06.02] The *New York Times* has a long story on the DarkSide ransomware gang.

> A glimpse into DarkSide's secret communications in the months leading up to the Colonial Pipeline attack reveals a criminal operation on the rise, pulling in millions of dollars in ransom payments each month.
>
> DarkSide offers what is known as "ransomware as a service," in which a malware developer charges a user fee to so-called affiliates like Woris, who may not have the technical skills to actually create ransomware but are still capable of breaking into a victim's computer systems.
>
> DarkSide's services include providing technical support for hackers, negotiating with targets like the publishing company, processing payments, and devising tailored pressure campaigns through blackmail and other means, such as secondary hacks to crash websites. DarkSide's user fees operated on a sliding scale: 25 percent for any ransoms less than $500,000 down to 10 percent for ransoms over $5 million, according to the computer security firm, FireEye.

** *** ***** ******* ********** **********

Security and Human Behavior (SHB) 2021

[2021.06.04] Today is the second day of the fourteenth Workshop on Security and Human Behavior. The University of Cambridge is the host, but we're all on Zoom.

SHB is a small, annual, invitational workshop of people studying various aspects of the human side of security, organized each year by Alessandro Acquisti, Ross Anderson, and myself. The forty or so attendees include psychologists, economists, computer security researchers, sociologists, political scientists, criminologists, neuroscientists, designers, lawyers, philosophers, anthropologists, business school professors, and a smattering of others. It's not just an interdisciplinary event; most of the people here are individually interdisciplinary.

Our goal is always to maximize discussion and interaction. We do that by putting everyone on panels, and limiting talks to six to eight minutes, with the rest of the time for open discussion. The format translates well to Zoom, and we're using random breakouts for the breaks between sessions.

I always find this workshop to be the most intellectually stimulating two days of my professional year. It influences my thinking in different, and sometimes surprising, ways.

This year's schedule is here. This page lists the participants and includes links to some of their work. As he does every year, Ross Anderson is liveblogging the talks.

Here are my posts on the first, second, third, fourth, fifth, sixth, seventh, eighth, ninth, tenth, eleventh, twelfth, and thirteenth SHB workshops. Follow those links to find summaries, papers, and occasionally audio recordings of the various workshops. Ross also maintains a good webpage of psychology and security resources.

** *** ***** ******* ********** **********

The Supreme Court Narrowed the CFAA

[2021.06.07] In a 6-3 ruling, the Supreme Court just narrowed the scope of the Computer Fraud and Abuse Act:

> In a ruling delivered today, the court sided with Van Buren and overturned his 18-month conviction.
>
> In a 37-page opinion written and delivered by Justice Amy Coney Barrett, the court explained that the "exceeds authorized access" language was, indeed, too broad.
>
> Justice Barrett said the clause was effectively making criminals of most US citizens who ever used a work resource to perform unauthorized actions, such as updating a dating profile, checking sports scores, or paying bills at work.
>
> What today's ruling means is that the CFAA cannot be used to prosecute rogue employees who have legitimate access to work-related resources, which will need to be prosecuted under different charges.
>
> The ruling does not apply to former employees accessing their old work systems because their access has been revoked and they're not "authorized" to access those systems anymore.

More.

It's a good ruling, and one that will benefit security researchers. But the confusing part is footnote 8:

> For present purposes, we need not address whether this inquiry turns only on technological (or "code-based") limitations on access, or instead also looks to limits contained in contracts or policies.

It seems to me that this is exactly what the ruling does address. The court overturned the conviction because the defendant was not limited by technology, but only by policies. So that footnote doesn't make any sense.

I have written about this general issue before, in the context of adversarial machine learning research.

** *** ***** ******* ********** **********

Vulnerabilities in Weapons Systems

[2021.06.08] "If you think any of these systems are going to work as expected in wartime, you're fooling yourself."

That was Bruce's response at a conference hosted by US Transportation Command in 2017, after learning that their computerized logistical systems were mostly unclassified and on the Internet. That may be necessary to keep in touch with civilian companies like FedEx in peacetime or when fighting terrorists or insurgents. But in a new era facing off with China or Russia, it is dangerously complacent.

Any twenty-first century war will include cyber operations. Weapons and support systems will be successfully attacked. Rifles and pistols won't work properly. Drones will be hijacked midair. Boats won't sail, or will be misdirected. Hospitals won't function. Equipment and supplies will arrive late or not at all.

Our military systems are vulnerable. We need to face that reality by halting the purchase of insecure weapons and support systems and by incorporating the realities of offensive cyberattacks into our military planning.

Over the past decade, militaries have established cyber commands and developed cyberwar doctrine. However, much of the current discussion is about offense. Increasing our offensive capabilities without being able to secure them is like having all the best guns in the world, and then storing them in an unlocked, unguarded armory. They just won't be stolen; they'll be subverted.

During that same period, we've seen increasingly brazen cyberattacks by everyone from criminals to governments. Everything is now a computer, and those computers are vulnerable. Cars, medical devices, power plants, and fuel pipelines have all been targets. Military computers, whether they're embedded inside weapons systems or on desktops managing the logistics of those weapons systems, are similarly vulnerable. We could see effects as stodgy as making a tank impossible to start up, or sophisticated as retargeting a missile midair.

Military software is unlikely to be any more secure than commercial software. Although sensitive military systems rely on domestically manufactured chips as part of the Trusted Foundry program, many military systems contain the same foreign chips and code that commercial systems do: just like everyone around the world uses the same mobile phones, networking equipment, and computer operating systems. For example, there has been serious concern over Chinese-made 5G networking equipment that might be used by China to install "backdoors" that would allow the equipment to be controlled. This is just one of many risks to our normal civilian computer supply chains. And since military software is vulnerable to the same cyberattacks as commercial software, military supply chains have many of the same risks.

This is not speculative. A 2018 GAO report expressed concern regarding the lack of secure and patchable US weapons systems. The report observed that "in operational testing, the [Department of Defense] routinely found mission-critical cyber vulnerabilities in systems that were under development, yet program officials GAO met with believed their systems were secure and discounted some test results as unrealistic." It's a similar attitude to corporate executives who believe that they can't be hacked -- and equally naive.

An updated GAO report from earlier this year found some improvements, but the basic problem remained: "DOD is still learning how to contract for cybersecurity in weapon systems, and selected programs we reviewed have struggled to incorporate systems' cybersecurity requirements into contracts." While DOD now appears aware of the issue of lack of cybersecurity requirements, they're still not sure yet how to fix it, and in three of the five cases GAO reviewed, DOD simply chose to not include the requirements at all.

Militaries around the world are now exploiting these vulnerabilities in weapons systems to carry out operations. When Israel in 2007 bombed a Syrian nuclear reactor, the raid was preceded by what is believed to have been a cyber attack on Syrian air defenses that resulted in radar screens showing no threat as bombers zoomed overhead. In 2018, a 29-country NATO exercise, Trident Juncture, that included cyberweapons was disrupted by Russian GPS jamming. NATO does try to test cyberweapons outside such exercises, but has limited scope in doing so. In May, Jens Stoltenberg, the NATO secretary-general, said that "NATO computer systems are facing almost daily cyberattacks."

The war of the future will not only be about explosions, but will also be about disabling the systems that make armies run. It's not (solely) that bases will get blown up; it's that some bases will lose power, data, and communications. It's not that self-driving trucks will suddenly go mad and begin rolling over friendly soldiers; it's that they'll casually roll off roads or into water where they sit, rusting, and in need of repair. It's not that targeting systems on guns will be retargeted to 1600 Pennsylvania Avenue; it's that many of them could simply turn off and not turn back on again.

So, how do we prepare for this next war? First, militaries need to introduce a little anarchy into their planning. Let's have wargames where essential systems malfunction or are subvertednot all of the time, but randomly. To help combat siloed military thinking, include some civilians as well. Allow their ideas into the room when predicting potential enemy action. And militaries need to have well-developed backup plans, for when systems are subverted. In Joe Haldeman's 1975 science-fiction novel *The Forever War,* he postulated a "stasis field" that forced his space marines to rely on nothing more than Roman military technologies, like javelins. We should be thinking in the same direction.

NATO isn't yet allowing civilians not employed by NATO or associated military contractors access to their training cyber ranges where vulnerabilities could be discovered and remediated before battlefield deployment. Last year, one of us (Tarah) was listening to a NATO briefing after the end of the 2020 Cyber Coalition exercises, and asked how she and other information security researchers could volunteer to test cyber ranges used to train its cyber incident response force. She was told that including civilians would be a "welcome thought experiment in the tabletop exercises," but including them in reality wasn't considered. There is a rich opportunity for improvement here, providing transparency into where improvements could be made.

Second, it's time to take cybersecurity seriously in military procurement, from weapons systems to logistics and communications contracts. In the three year span from the original 2018 GAO report to this year's report, cybersecurity audit compliance went from 0% to 40% (those 2 of 5 programs mentioned earlier). We need to get much better. DOD requires that its contractors and suppliers follow the Cybersecurity Maturity Model Certification process; it should abide by the same standards. Making those standards both more rigorous and mandatory would be an obvious second step.

Gone are the days when we can pretend that our technologies will work in the face of a military cyberattack. Securing our systems will make everything we buy more expensive -- maybe a lot more expensive. But the alternative is no longer viable.

The future of war is cyberwar. If your weapons and systems aren't secure, don't even bother bringing them onto the battlefield.

This essay was written with Tarah Wheeler, and previously appeared in Brookings TechStream.

** *** ***** ******* ********** **********

Information Flows and Democracy

[2021.06.09] Henry Farrell and I published a paper on fixing American democracy: "Rechanneling Beliefs: How Information Flows Hinder or Help Democracy."

> It's much easier for democratic stability to break down than most people realize, but this doesn't mean we must despair over the future. It's possible, though very difficult, to back away from our current situation towards one of greater democratic stability. This wouldn't entail a restoration of a previous status quo. Instead, it would recognize that the status quo was less stable than it seemed, and a major source of the tensions that have started to unravel it. What we need is a dynamic stability, one that incorporates new forces into American democracy rather than trying to deny or quash them.
>
> This paper is our attempt to explain what this might mean in practice. We start by analyzing the problem and explaining more precisely why a breakdown in public consensus harms democracy. We then look at how these beliefs are being undermined by three feedback loops, in which anti-democratic actions and anti-democratic beliefs feed on each other. Finally, we explain how these feedback loops might be redirected so as to sustain democracy rather than undermining it.
>
> To be clear: redirecting these and other energies in more constructive ways presents enormous challenges, and any plausible success will at best be untidy and provisional. But, almost by definition, that's true of any successful democratic reforms where people of different beliefs and values need to figure out how to coexist. Even when it's working well, democracy is messy. Solutions to democratic breakdowns are going to be messy as well.

This is part of our series of papers looking at democracy as an information system. The first paper was "Common-Knowledge Attacks on Democracy."

** *** ***** ******* ********** **********

Detecting Deepfake Picture Editing

[2021.06.10] "Markpainting" is a clever technique to watermark photos in such a way that makes it easier to detect ML-based manipulation:

> An image owner can modify their image in subtle ways which are not themselves very visible, but will sabotage any attempt to inpaint it by adding visible information determined in advance by the markpainter.
>
> One application is tamper-resistant marks. For example, a photo agency that makes stock photos available on its website with copyright watermarks can markpaint them in such a way that anyone using common editing software to remove a watermark will fail; the copyright mark will be markpainted right back. So watermarks can be made a lot more robust.

Here's the paper: "Markpainting: Adversarial Machine Learning Meets Inpainting," by David Khachaturov, Ilia Shumailov, Yiren Zhao, Nicolas Papernot, and Ross Anderson.

> **Abstract:** Inpainting is a learned interpolation technique that is based on generative modeling and used to populate masked or missing pieces in an image; it has wide applications in picture editing and retouching. Recently, inpainting started being used for watermark removal, raising concerns. In this paper we study how to manipulate it using our markpainting technique. First, we show how an image owner with access to an inpainting model can augment their image in such a way that any attempt to edit it using that model will add arbitrary visible information. We find that we can target multiple different models simultaneously with our technique. This can be designed to reconstitute a watermark if the editor had been trying to remove it. Second, we show that our markpainting technique is transferable to models that have different architectures or were trained on different datasets, so watermarks created using it are difficult for adversaries to remove. Markpainting is novel and can be used as a manipulation alarm that becomes visible in the event of inpainting.

** *** ***** ******* ********** **********

FBI/AFP-Run Encrypted Phone

[2021.06.11] For three years, the Federal Bureau of Investigation and the Australian Federal Police owned and operated a commercial encrypted phone app, called AN0M, that was used by organized crime around the world. Of course, the police were able to read everything -- I don't even know if this qualifies as a backdoor. This week, the world's police organizations announced 800 arrests based on text messages sent over the app. We've seen law enforcement take over encrypted apps before: for example, EncroChat. This operation, code-named Trojan Shield, is the first time law enforcement managed an app from the beginning.

If there is any moral to this, it's one that all of my blog readers should already know: trust is essential to security. And the number of people you need to trust is larger than you might originally think. For an app to be secure, you need to trust the hardware, the operating system, the software, the update mechanism, the login mechanism, and on and on and on. If one of those is untrustworthy, the whole system is insecure.

It's the same reason blockchain-based currencies are so insecure, even if the cryptography is sound.

** *** ***** ******* ********** **********

TikTok Can Now Collect Biometric Data

[2021.06.14] This is probably worth paying attention to:

> A change to TikTok's U.S. privacy policy on Wednesday introduced a new section that says the social video app *"may collect biometric identifiers and biometric information"* from its users' content. This includes things like *"faceprints and voiceprints,"* the policy explained. Reached for comment, TikTok could not confirm what product developments necessitated the addition of biometric data to its list of disclosures about the information it automatically collects from users, but said it would ask for consent in the case such data collection practices began.

** *** ***** ******* ********** **********

Upcoming Speaking Engagements

[2021.06.14] This is a current list of where and when I am scheduled to speak:

- I'll be part of a European Internet Forum virtual debate on June 17, 2021. The topic is "Decrypting the encryption debate: How to ensure public safety with a privacy-preserving and secure Internet?"

- I'm speaking at the all-online Society for Philosophy and Technology Conference 2021, June 28-30, 2021.
- I'm keynoting the 5th International Symposium on Cyber Security Cryptology and Machine Learning (via Zoom), July 8-9, 2021.
- I'm speaking (via Internet) at SHIFT Business Festival in Finland, August 25-26, 2021.
- I'll be speaking at an Informa event on September 14, 2021. Details to come.

The list is maintained on this page.

** *** ***** ******* ********** **********

Since 1998, CRYPTO-GRAM has been a free monthly newsletter providing summaries, analyses, insights, and commentaries on security technology. To subscribe, or to read back issues, see Crypto-Gram's web page.

You can also read these articles on my blog, Schneier on Security.

Please feel free to forward CRYPTO-GRAM, in whole or in part, to colleagues and friends who will find it valuable. Permission is also granted to reprint CRYPTO-GRAM, as long as it is reprinted in its entirety.

Bruce Schneier is an internationally renowned security technologist, called a security guru by the Economist*. He is the author of over one dozen books -- including his latest,* We Have Root *-- as well as hundreds of articles, essays, and academic papers. His newsletter and blog are read by over 250,000 people. Schneier is a fellow at the Berkman Klein Center for Internet & Society at Harvard University; a Lecturer in Public Policy at the Harvard Kennedy School; a board member of the Electronic Frontier Foundation, AccessNow, and the Tor Project; and an Advisory Board Member of the Electronic Privacy Information Center and VerifiedVoting.org. He is the Chief of Security Architecture at Inrupt, Inc.*

** *** ***** ******* ********** **********

Bruce Schneier <schneier@schneier.com> Mon, Jun 14, 2021 at 9:15 PM
Reply-To: Bruce Schneier <schneier@schneier.com>
To: slamb@xtcn.com

[Quoted text hidden]

[Quoted text hidden]