1 CORE CASES

1.1 Chaff: defeating military radar

During the Second World War, a radar operator tracks an airplane over Hamburg, guiding searchlights and anti-aircraft guns in relation to a phosphor dot whose position is updated with each sweep of the antenna. Abruptly, dots that seem to represent airplanes begin to multiply, quickly swamping the display. The actual plane is in there somewhere, impossible to locate owing to the presence of "false echoes."[1]

The plane has released chaff—strips of black paper backed with aluminum foil and cut to half the target radar's wavelength. Thrown out by the pound and then floating down through the air, they fill the radar screen with signals. The chaff has exactly met the conditions of data the radar is configured to look for, and has given it more "planes," scattered all across the sky, than it can handle.

This may well be the purest, simplest example of the obfuscation approach. Because discovery of an actual airplane was inevitable (there wasn't, at the time, a way to make a plane invisible to radar), chaff taxed the time and bandwidth constraints of the discovery system by creating too many potential targets. That the chaff worked only briefly as it fluttered to the ground and was not a permanent solution wasn't relevant under the circumstances. It only had to work well enough and long enough for the plane to get past the range of the radar.

As we will discuss in part II, many forms of obfuscation work best as time-buying "throw-away" moves. They can get you only a few minutes, but sometimes a few minutes is all the time you need.

The example of chaff also helps us to distinguish, at the most basic level, between approaches to obfuscation. Chaff relies on producing echoes—imitations of the real thing—that exploit the limited scope of the observer. (Fred Cohen terms this the "decoy strategy."[2]) As we will see, some forms of obfuscation *generate genuine but misleading signals*—much as you would protect the contents of one vehicle by sending it out accompanied by several other identical vehicles, or defend a particular plane by filling the sky with other planes—whereas other forms *shuffle genuine signals*, mixing data in an effort to make the extraction of patterns more difficult. Because those who scatter chaff have exact knowledge of their adversary, chaff doesn't have to do either of these things.

If the designers of an obfuscation system have specific and detailed knowledge of the limits of the observer, the system they develop has to work for only one wavelength and for only 45 minutes. If the system their adversary uses for observation is more patient, or if it has a more comprehensive set of capacities for observation, they have to make use of their understanding of the adversary's internal agenda—that is, of what useful information the adversary hopes to extract from data obtained through surveillance—and undermine that agenda by manipulating genuine signals.

Before we turn to the manipulation of genuine signals, let's look at a very different example of flooding a channel with echoes.

1.2 Twitter bots: filling a channel with noise

The two examples we are about to discuss are a study in contrasts. Although producing imitations is their mode of obfuscation, they take us from the Second World War to present-day circumstances, and from radar to social networks. They also introduce an important theme.

In chapter 3, we argue that obfuscation is a tool particularly suited to the "weak"—the situationally disadvantaged, those at the wrong end of asymmetrical power relationships. It is a method, after all, that you have reason to adopt if you can't be invisible—if you can't refuse to be tracked or surveilled, if you can't simply opt out or operate within professionally secured networks. This doesn't mean that it isn't also taken up by the powerful. Oppressive or coercive forces usually have better means than obfuscation at their disposal. Sometimes, though, obfuscation becomes useful to powerful actors—as it did in two elections, one in Russia and one in Mexico. Understanding the choices faced by the groups in contention will clarify how obfuscation of this kind can be employed.

During protests over problems that had arisen in the 2011 Russian parliamentary elections, much of the conversation about ballot-box stuffing and other irregularities initially took place on LiveJournal, a blogging platform that had originated in the United States but attained its greatest popularity in Russia—more than half of its user base is Russian.[3] Though LiveJournal is quite popular, its user base is very small relative to those of Facebook's and Google's various social systems; it has fewer than 2 million active accounts.[4] Thus, LiveJournal is comparatively easy for attackers to shut down by means of distributed denial of service (DDoS) attack—that is, by using computers

scattered around the world to issue requests for the site in such volume that the servers making the site available are overwhelmed and legitimate users can't access it. Such an attack on LiveJournal, in conjunction with the arrests of activist bloggers at a protest in Moscow, was a straightforward approach to censorship.[5] When and why, then, did obfuscation become necessary?

The conversation about the Russian protest migrated to Twitter, and the powers interested in disrupting it then faced a new challenge. Twitter has an enormous user base, with infrastructure and security expertise to match. It could not be taken down as easily as LiveJournal. Based in the United States, Twitter was in a much better position to resist political manipulation than LiveJournal's parent company. (Although LiveJournal service is provided by a company set up in the U.S. for that purpose, the company that owns it, SUP Media, is based in Moscow.[6]) To block Twitter outright would require direct government intervention. The LiveJournal attack was done independently, by nationalist hackers who may or may not have the approval and assistance of the Putin/Medvedev administration.[7] Parties interested in halting the political conversation on Twitter therefore faced a challenge that will become familiar as we explore obfuscation's uses: time was tight, and traditional mechanisms for action weren't available. A direct technical approach—either blocking Twitter within a country or launching a worldwide denial-of-service attack—wasn't possible, and political and legal angles of attack couldn't be used. Rather than stop a Twitter conversation, then, attackers can overload it with noise.

During the Russian protests, the obfuscation took the form of thousands of Twitter accounts suddenly piping up and users posting tweets using the same hashtags used by the protesters.[8] Hashtags are a mechanism for grouping tweets together; for example, if I add #obfuscation to a tweet, the symbol # turns the word into an active link—clicking it will bring up all other tweets tagged with #obfuscation. Hashtags are useful for organizing the flood of tweets into coherent conversations on specific topics, and #триумфальная (referring to Triumfalnaya, the location of a protest) became one of several tags people could use to vent their anger, express their opinions, and organize further actions. (Hashtags also play a role in how Twitter determines "trending" and significant topics on the site, which can then draw further attention to what is being discussed under that tag—the site's Trending Topics list often draws news coverage.[9])

If you were following #триумфальная, you would have seen tweet after tweet from Russian activists spreading links to news and making plans. But those tweets began to be interspersed with tweets about Russian greatness, or tweets that seemed to consist of noise, gibberish, or random words and phrases. Eventually those tweets dominated the stream for #триумфальная, and those for other topics related to the protests, to such a degree that tweets relevant to the topic were, essentially, lost in the noise, unable to get any attention or to start a coherent exchange with other users. That flood of new tweets came from accounts that had been inactive for much of their existence. Although they had posted very little from the time of their creation until the time of the protests, now each of them was posting dozens of times an hour. Some of the accounts' purported users had mellifluous names, such as imelixyvyq, wyqufahij, and hihexiq; others had more conventional-seeming names, all built on a firstname_lastname model—for example, latifah_xander.[10]

Obviously, these Twitter accounts were "Twitter bots"—programs purporting to be people and generating automatic, targeted messages. Many of the accounts had been created around the same time. In numbers and in frequency, such messages can easily dominate a discussion, effectively ruining the platform for a specific audience through overuse—that is, obfuscating through the production of false, meaningless signals.

The use of Twitter bots is becoming a reliable technique for stifling Twitter discussion. The highly contentious 2012 Mexican elections provide another example of this strategy in practice, and further refined.[11] Protesters opposed to the front-runner, Enrique Peña Nieto, and to the Partido Revolucionario Institucional (PRI), used #marchaAntiEPN as an organizing hashtag for the purposes of aggregating conversation, structuring calls for action, and arranging protest events. Groups wishing to interfere with the protesters' organizing efforts faced challenges similar to those in the Russian case. Rather than thousands of bots, however, hundreds would do—indeed, when this case was investigated by the American Spanish-language TV network Univision, only about thirty such bots were active. Their approach was both to interfere with the work being done to advance #marchaAntiEPN and to overuse that hashtag. Many of the tweets consisted entirely of variants of "#marchaAntiEPN #marchaAntiEPN #marchaAntiEPN #marchaAntiEPN #marchaAntiEPN #marchaAntiEPN." Such repetition, particularly by accounts already showing

suspiciously bot-like behavior, triggers systems within Twitter that identify attempts to manipulate the hashtagging system and then remove the hashtags in question from the Trending Topics list. In other words, because the items in Trending Topics become newsworthy and attract attention, spammers and advertisers will try to push hashtags up into that space through repetition, so Twitter has developed mechanisms for spotting and blocking such activity.[12]

The Mexican-election Twitter bots were deliberately engaging in bad behavior in order to trigger an automatic delisting, thereby keeping the impact of #marchaAntiEPN "off the radar" of the larger media. They were making the hashtag unusable and removing its potential media significance. This was obfuscation as a destructive act. Though such efforts use the same basic tactic as radar chaff (that is, producing many imitations configured to hide the real thing), they have very different goals: rather than just buying time (for example, in the run-up to an election and during the period of unrest afterward), they render certain terms unusable—even, from the perspective of a sorting algorithm, toxic—by manipulating the properties of the data through the use of false signals.

1.3 CacheCloak: location services without location tracking

CacheCloak takes an approach to obfuscation that is suited to location-based services (LBSs).[13] It illustrates two twists in the use of false echoes and imitations in obfuscation. The first of these is making sure that relevant data can still be extracted by the user; the second is trying to find an approach that can work indefinitely rather than as a temporary time-buying strategy.

Location-based services take advantage of the locative capabilities of mobile devices to create various services, some of them social (e.g., FourSquare, which turns going places into a competitive game), some lucrative (e.g., location-aware advertising), and some thoroughly useful (e.g., maps and nearest-object searches). The classic rhetoric of balancing privacy against utility, in which utility is often presented as detrimental to privacy, is evident here. If you want the value of an LBS—for example, if you want to be on the network that your friends are on so you can meet with one of them if you and that person are near one another—you will have to sacrifice some privacy, and you will have to get accustomed to having the service provider know where you are. CacheCloak suggests a way to reconfigure the tradeoff.

"Where other methods try to obscure the user's path by hiding parts of it," the creators of CacheCloak write, "we obscure the user's location by surrounding it with other users' paths"[14]—that is, through the propagation of ambiguous data. In the standard model, your phone sends your location to the service and gets the information you requested in return. In the CacheCloak model, your phone predicts your possible paths and then fetches the results for several likely routes. As you move, you receive the benefits of locative awareness—access to what you are looking for, in the form of data cached in advance of potential requests—and an adversary is left with many possible paths, unable to distinguish the beginning from the end of a route and unable to determine where you came from, where you mean to go, or even where you are. From an observer's perspective, the salient data—the data we wish to keep to ourselves—are buried inside a space of other, equally likely data.

1.4 TrackMeNot: blending genuine and artificial search queries

TrackMeNot, developed in 2006 by Daniel Howe, Helen Nissenbaum, and Vincent Toubiana, exemplifies a software strategy for concealing activity with imitative signals.[15] The purpose of TrackMeNot is to foil the profiling of users through their searches. It was designed in response to the U.S. Department of Justice's request for Google's search logs and in response to the surprising discovery by a *New York Times* reporter that some identities and profiles could be inferred even from anonymized search logs published by AOL Inc.[16]

Our search queries end up acting as lists of locations, names, interests, and problems. Whether or not our full IP addresses are included, our identities can be inferred from these lists, and patterns in our interests can be discerned. Responding to calls for accountability, search companies have offered ways to address people's concerns about the collection and storage of search queries, though they continue to collect and analyze logs of such queries.[17] Preventing any stream of queries from being inappropriately revealing of a particular person's interests and activities remains a challenge.[18]

The solution TrackMeNot offers is not to hide users' queries from search engines (an impractical method, in view of the need for query satisfaction), but to obfuscate by automatically generating queries from a "seed list" of terms. Initially culled from RSS feeds, these terms evolve so that different users develop different seed lists. The precision of the imitation is continually refined by repopulating the seed list with new terms generated from returns to search

queries. TrackMeNot submits queries in a manner that tries to mimic real users' search behaviors. For example, a user who has searched for "good wi-fi cafe chelsea" may also have searched for "savannah kennels," "freshly pressed juice miami," "asian property firm," "exercise delays dementia," and "telescoping halogen light." The activities of individuals are masked by those of many ghosts, making the pattern harder to discern so that it becomes much more difficult to say of any query that it was a product of human intention rather than an automatic output of TrackMeNot. In this way, TrackMeNot extends the role of obfuscation, in some situations, to include plausible deniability.

1.5 Uploads to leak sites: burying significant files

WikiLeaks used a variety of systems for securing the identities of both visitors and contributors. However, there was a telltale sign that could undercut the safety of the site: uploads of files. If snoops could monitor the traffic on WikiLeaks, they could identify acts of submitting material to WikiLeaks' secure server. Especially if they could make informed guesses as to the compressed sizes of various collections of subsequently released data, they could retroactively draw inferences as to what was transmitted, when it was transmitted, and (in view of failures in other areas of technical and operations security) by whom it was transmitted. Faced with this very particular kind of challenge, WikiLeaks developed a script to produce false signals. It launched in the browsers of visitors, generating activity that looked like uploads to the secure server.[19] A snoop would therefore see an enormous mob of apparent leakers (the vast majority of whom were, in actuality, merely reading or looking through documents already made available), a few of whom might really be leakers. It didn't seek to provide *particular* data to interfere with data mining or with advertising; it simply sought to imitate and conceal the movements of some of its users.

Even encrypted and compressed data contain pertinent metadata, however, and the proposal for OpenLeaks—an ultimately unsuccessful variant on WikiLeaks, developed by some of the disaffected participants in the original WikiLeaks system—includes a further refinement.[20] After a statistical analysis of the WikiLeaks submissions, OpenLeaks developed a model of fake uploads that would keep to the same ratios of *sizes* of files typically appearing in the upload traffic of a leak site. Most of the files ranged in size from 1.5 to 2

megabytes, though a few outliers exceeded 700 megabytes. If an adversary can monitor upload traffic, form can be as telling as content, and as useful in sorting real signals from fake ones. As this example suggests, obfuscation mechanisms can gain a great deal from figuring out all the parameters that can be manipulated—and from figuring out what the adversary is looking for, so as to give the adversary a manufactured version of it.

1.6 False tells: making patterns to trick a trained observer

Consider how the same basic pattern of obfuscation can be called to service in a context lighter than concealing the work of whistleblowers: poker.

Much of the pleasure and much of the challenge of poker lies in learning to infer from expressions, gestures, and body language whether someone is bluffing (that is, pretending to hold a hand weaker than the one he or she actually holds) in hopes of drawing a call. Central to the work of studying one's opponents is the "tell"—some unconscious habit or tic that an opponent displays in response to a strong or a weak hand, such as sweating, glancing worriedly, or leaning forward. Tells are so important in the informational economy of poker that players sometimes use *false tells*—that is, they create mannerisms that may appear to be parts of a larger pattern.[21] In common poker strategy, the use of a false tell is best reserved for a crucial moment in a tournament, lest the other players figure out that it is inaccurate and use it against you in turn. A patient analysis of multiple games could separate the true tells from the false ones, but in the time-bound context of a high-stakes game the moment of falsehood can be highly effective. Similar techniques are used in many sports that involve visible communication. One example is signaling in baseball—as a coach explained to a newspaper reporter, "Sometimes you're giving a sign, but it doesn't even mean anything."[22]

1.7 Group identity: many people under one name

One of the simplest and most memorable examples of obfuscation, and one that introduces the work of the *group* in obfuscation, is the scene in the film *Spartacus* in which the rebel slaves are asked by Roman soldiers to identify their leader, whom the soldiers intend to crucify.[23] As Spartacus (played by Kirk Douglas) is about to speak, one by one the others around him say "I am Spartacus!" until the entire crowd is claiming that identity.

Many people assuming the same identity for group protection (for example, Captain Swing in the English agricultural uprisings of 1830, the ubiquitous "Jacques" adopted by the radicals in Dickens's *A Tale of Two Cities*, or the Guy Fawkes mask in the graphic novel *V for Vendetta*, now associated with the hacktivist group known as Anonymous) is, at this point, almost a cliché.[24] Marco Deseriis has studied the use of "improper names" and collective identities in the effacement of individual responsibility and the proliferation of action.[25] Some forms of obfuscation can be conducted solo; others rely on groups, teams, communities, and confederates.

1.8 Identical confederates and objects: many people in one outfit

There are many examples of obfuscation by members of a group working in concert to produce genuine but misleading signals within which the genuine, salient signal is concealed. One memorable example from popular culture is the scene in the 1999 remake of the film *The Thomas Crown Affair* in which the protagonist, wearing a distinctive Magritte-inspired outfit, is suddenly in a carefully orchestrated mass of other men, dressed in the same outfit, circulating through the museum and exchanging their identical briefcases.[26] The bank-robbery scheme in the 2006 film *Inside Man* hinges on the robbers' all wearing painters' overalls, gloves, and masks and dressing their hostages the same way.[27] Finally, consider the quick thinking of Roger Thornhill, the protagonist of Alfred Hitchcock's 1959 film *North By Northwest*, who, in order to evade the police when his train arrives in Chicago, bribes a redcap (a baggage handler) to lend him his distinctive uniform, knowing that the crowd of redcaps at the station will give the police too much of something specific to look for.[28]

Identical objects as modes of obfuscation are common enough and sufficiently understood to recur in imagination and in fact. The *ancilia* of ancient Rome exemplify this. A shield (*ancile*) fell from the sky—so the legend goes—during the reign of Numa Pompilius, Rome's second king (753–673 BCE), and was interpreted as a sign of divine favor, a sacred relic whose ownership would guarantee Rome's continued imperium.[29] It was hung in the Temple of Mars along with eleven exact duplicates, so would-be thieves wouldn't know which one to take. The six plaster busts of Napoleon from which the Sherlock Holmes story gets its title offers another example. The villain sticks a black pearl into the wet plaster of an object that not only has five duplicates but also

is one of a larger class of objects (cheap white busts of Napoleon) that are ubiquitous enough to be invisible.[30]

A real-world instance is provided by the so-called Craigslist robber. At 11 a.m. on Tuesday, September 30, 2008, a man dressed as an exterminator (in a blue shirt, goggles, and a dust mask), and carrying a spray pump, approached an armored car parked outside a bank in Monroe, Washington, incapacitated the guard with pepper spray, and made off with the money.[31] When the police arrived, they found thirteen men in the area wearing blue shirts, goggles, and dust masks—a uniform they were wearing on the instructions of a Craigslist ad that promised a good wage for maintenance work, which was to start at 11:15 a.m. at the bank's address. It would have taken only a few minutes to determine that none of the day laborers was the robber, but a few minutes was all the time the robber needed.

Then there is the powerful story, often retold though factually inaccurate, of the king of Denmark and a great number of Danish gentiles wearing the Yellow Star so that the occupying Germans couldn't distinguish and deport Danish Jews. Although the Danes courageously protected their Jewish population in other ways, the Yellow Star wasn't used by the Nazis in occupied Denmark, for fear of arousing more anti-German feeling. However, "there were documented cases of non–Jews wearing yellow stars to protest Nazi anti–Semitism in Belgium, France, the Netherlands, Poland, and even Germany itself."[32] This legend offers a perfect example of cooperative obfuscation: gentiles wearing the Yellow Star as an act of protest, providing a population into which individual Jews could blend.[33]

1.9 Excessive documentation: making analysis inefficient

Continuing our look at obfuscation that operates by adding in genuine but misleading signals, let us now consider the overproduction of documents as a form of obfuscation, as in the case of over-disclosure of material in a lawsuit. This was the strategy of Augustin Lejeune, chief of the General Police Bureau in the Committee of Public Safety, a major instrument in the Terror phase of the French Revolution. Lejeune and his clerks produced the reports that laid the groundwork for arrests, internments, and executions. Later, in an effort to excuse his role in the Terror, Lejeune argued that the exacting, overwhelmingly detailed quality of the reports from his office had been deliberate: he had instructed his clerks to overproduce material, and to report "the most minor

details," in order to slow the production of intelligence for the Committee without the appearance of rebellion. It is doubtful that Lejeune's claims are entirely accurate (the numbers he cites for the production of reports aren't reliable), but, as Ben Kafka points out, he had come up with a bureaucratic strategy for creating slowdowns through oversupply: "He seems to have recognized, if only belatedly, that the proliferation of documents and details presented opportunities for resistance, as well as for compliance."[34] In situations where one can't say No, there are opportunities for a chorus of unhelpful Yeses—for example, don't send a folder in response to a request; send a pallet of boxes of folders containing potentially relevant papers.

1.10 Shuffling SIM cards: rendering mobile targeting uncertain

As recent reporting and some of Edward Snowden's disclosures have revealed, analysts working for the National Security Agency use a combination of signals-intelligence sources—particularly cell-phone metadata and data from geolocation systems—to identify and track targets for elimination.[35] The metadata (showing what numbers were called and when they were called) produce a model of a social network that makes it possible to identify particular phone numbers as belonging to persons of interest; the geolocative properties of mobile phones mean that these numbers can be situated, with varying degrees of accuracy, in particular places, which can then be targeted by drones. In other words, this system can proceed from identification to location to assassination without ever having a face-to-face visual identification of a person. The closest a drone operator may come to setting eyes on someone may be the exterior of a building, or a silhouette getting into a car. In view of the spotty records of the NSA's cell-phone-metadata program and the drone strikes, there are, of course, grave concerns about accuracy. Whether one is concerned about threats to national security remaining safe and active, about the lives of innocent people taken unjustly, or about both, it is easy to see the potential flaws in this approach.

Let us flip the situation, however, and consider it more abstractly from the perspective of the targets. Most of the NSA's targets are obligated to always have, either with or near them, a tracking device (only the very highest-level figures in terrorist organizations are able to be free of signals-generating technology), as are virtually all the people with whom they are in contact. The calls and conversations that sustain their organizations also provide the

means of their identification; the structure that makes their work possible also traps them. Rather than trying to coordinate anti-aircraft guns to find a target somewhere in the sky, the adversary has complete air superiority, able to deliver a missile to a car, a street corner, or a house. However, the adversary also has a closely related set of systemic limitations. This system, remarkable as it is in scope and capabilities, ultimately relies on SIM (subscriber identity module) cards and on physical possession of mobile phones—a kind of narrow bandwidth that can be exploited. A former drone operator for the Joint Special Operations Command has reported that targets therefore take measures to mix and confuse genuine signals. Some individuals have many SIM cards associated with their identity in circulation, and the cards are randomly redistributed. One approach is to hold meetings at which all the attendees put their SIM cards into a bag, then pull cards from the bag at random, so that who is actually connected to each device will not be clear. (This is a time-bound approach: if metadata analysis is sufficiently sophisticated, an analyst should eventually be able to sort the individuals again on the basis of past calling patterns, but irregular re-shuffling renders that more difficult.) Re-shuffling may also happen unintentionally as targets who aren't aware that they are being tracked sell their phones or lend them to friends or relatives. The end result is a system with enormous technical precision and a very uncertain rate of actual success, whether measured in terms of dangerous individuals eliminated or in terms of innocent noncombatants killed by mistake. Even when fairly exact location tracking and social-graph analysis can't be avoided, using obfuscation to mingle and mix genuine signals, rather than generating false signals, can offer a measure of defense and control.

1.11 Tor relays: requests on behalf of others that conceal personal traffic

Tor is a system designed to facilitate anonymous use of the Internet through a combination of encryption and passing the message through many different independent "nodes." In a hybrid strategy of obfuscation, Tor can be used in combination with other, more powerful mechanisms for concealing data. Such a strategy achieves obfuscation partially through the mixing and interleaving of genuine (encrypted) activity. Imagine a message passed surreptitiously through a huge crowd to you. The message is a question without any identifying information; as far as you know, it was written by the last person to hold it,

the person who handed it to you. The reply you write and pass back vanishes into the crowd, following an unpredictable path. Somewhere in that crowd, the writer receives his answer. Neither you nor anyone else knows exactly who the writer was.

If you request a Web page while working through Tor, your request will not come from your IP address; it will come from an "exit node" (analogous to the last person who hands the message to its addressee) on the Tor system, along with the requests of many other Tor users. Data enter the Tor system and pass into a labyrinth of relays—that is, computers on the Tor network (analogous to people in the crowd) that offer some of their bandwidth for the purpose of handling Tor traffic from others, agreeing to pass messages sight unseen. The more relays there are, the faster the system is as a whole. If you are already using Tor to protect your Internet traffic, you can turn your computer into a relay for the collective greater good. Both the Tor network and the obfuscation of individuals on the network improve as more people make use of the network.

Obfuscation, Tor's designers point out, augments its considerable protective power. In return for running a Tor relay, "you do get better anonymity against some attacks. The simplest example is an attacker who owns a small number of Tor relays. He will see a connection from you, but he won't be able to know whether the connection originated at your computer or was relayed from somebody else."[36] If someone has agents in the crowd—that is, if someone is running Tor relays for surveillance purposes—the agents can't read a message they pass, but they can notice who passed it to them. If you are on Tor and not running a relay, they know that you wrote the message you gave to them. But if you are letting your computer operate as a relay, the message may be yours or may be just one among many that you are passing on for other people. Did that message start with you, or not? The information is now ambiguous, and messages you have written are safe in a flock of other messages you pass along. This is, in short, a significantly more sophisticated and efficient way to render particular data transactions ambiguous and to thwart traffic analysis by making use of the volume of the traffic. It doesn't merely mix genuine signals (as shaking up SIM cards in a bag does, with all the consequent problems of coordination); it gets each message to its destination. However, each message can serve to make the sources of other messages uncertain.

1.12 Babble tapes: hiding speech in speech

An old cliché about mobsters under threat from the FBI involved a lot of talking in bathrooms: the splash and hiss of water and the hum of the ventilation fan, so the story went, made conversations hard to hear if the house was bugged or if someone in the room was wearing a wire. There are now refined (and much more effective) techniques for defeating audio surveillance that draw more directly on obfuscation. One of these is the use of so-called babble tapes.[37] Paradoxically, babble tapes have been used less by mobsters than by attorneys concerned that eavesdropping may violate attorney-client privilege.

A babble tape is a digital file meant to be played in the background during conversations. The file is complex. Forty voice tracks run simultaneously (thirty-two in English, eight in other languages), and each track is compressed in frequency and time to produce additional "voices" that fill the entire frequency spectrum. There are also various non-human mechanical noises, and a periodic supersonic burst (inaudible to adult listeners) engineered specifically to interfere with the automatic gain-control system of an eavesdropping device configures itself to best pick up an audio signal. Most pertinent for present purposes, the voices on a babble tape used by an attorney include those of the client and the attorney themselves. The dense mélange of voices increases the difficulty of discerning any single voice.

1.13 Operation Vula: obfuscation in the struggle against Apartheid

We close this chapter with a detailed narrative example of obfuscation employed in a complex context by a group seeking to get Nelson Mandela released from prison in South Africa during the struggle against Apartheid. Called Operation Vula (short for Vul'indlela, meaning Opening the Road), it was devised by leaders of the African National Congress within South Africa who were in contact with Mandela and were coordinating their efforts with those of ANC agents, sympathizers, and generals around the world.

The last project of this scale that the ANC had conducted had resulted in the catastrophe of the early 1960s in which Mandela and virtually all of the ANC's top leaders had been arrested and the Liliesleaf Farm documents had been captured and had been used against them in court. This meant that Operation Vula had to be run with absolutely airtight security and privacy practices. Indeed, when the full scope of the operation was revealed in the 1990s, it came

as a surprise not just to the South African government and to international intelligence services but also to many prominent leadership figures within the ANC. People purportedly receiving kidney transplants or recovering from motorcycle accidents had actually gone deep underground with new identities and then had returned to South Africa, "opening the road" for Mandela's release. Given the surveillance inside and outside South Africa, the possible compromise of pre-existing ANC communications channels, and the interest of spies and law-enforcement groups around the world, Operation Vula had to have secure ways of sharing and coordinating information.

The extraordinary tale of Operation Vula has been told by one of its chief architects, Tim Jenkin, in the pages of the ANC's journal *Mayibuye*.[38] It represents a superb example of operations security, tradecraft, and managing a secure network.

Understanding when and how obfuscation came to be employed in Operation Vula requires understanding some of the challenges its architects faced. Using fixed phone lines within South Africa, each linked to an address and a name, wasn't an option. The slightest compromise might lead to wiretaps and to what we would now call metadata analysis, and thus a picture of the activist network could be put together from domestic and overseas phone logs. The Vula agents had various coding systems, each of them hampered by the difficulty and tedium of doing the coding by hand. There was always the temptation to fall back on "speaking in whispers over phones again," especially when crises happened and things began moving fast. The operation had to be seamlessly coordinated between South Africa (primarily Durban and Johannesburg) and Lusaka, London, Amsterdam, and other locations around the world as agents circulated. Postal service was slow and vulnerable, encrypting was enormously time consuming and often prone to sloppiness, use of home phones was forbidden, and coordinating between multiple time zones around the world seemed impossible.

Jenkin was aware of the possibilities of using personal computers to make encryption faster and more efficient. Based in London after his escape from Pretoria Central Prison, he spent the mid 1980s working on the communications system needed for Operation Vula, which ultimately evolved into a remarkable network. Encryption happened on a personal computer, and the ciphered message was then expressed as a rapid series of tones recorded onto a portable cassette player. An agent would go to a public pay phone and

dial a London number, which would be picked up by an answering machine that Jenkin had modified to record for up to five minutes. The agent would play the cassette into the mouthpiece of the phone. The tones, recorded on the cassette's other side, could be played through an acoustic modem into the computer and then decrypted. (There was also an "outgoing" answering machine. Remote agents could call from a pay phone, record the tones for their messages, and decrypt them anywhere they had access to a computer that could run the ciphering systems Jenkin had devised.)

This was already an enormously impressive network—not least because large parts of its digital side (including a way of implementing error-handling codes to deal with the noise of playing back messages over international phone lines from noisy booths) had to be invented from scratch. However, as Operation Vula continued to grow and the network of operatives to expand, the sheer quantity of traffic threatened to overwhelm the network. Operatives were preparing South Africa for action, and that work didn't leave a lot of time for finding pay phones that accepted credit cards (the sound of coins dropping could interfere with the signal) and standing around with tape players. Jenkin and his collaborators would stay up late, changing tapes in the machines as the messages poured in. The time had come to switch to encrypted email, but the whole system had been developed to avoid the use of known, owned telephone lines within South Africa.

Operation Vula needed to be able to send encrypted messages to and from computers in South Africa, in Lukasa, and in London without arousing suspicion. During the 1980s, while the network we have described was taking shape, the larger milieu of international business was producing exactly the kind of background against which this subterfuge could hide itself. The question was, as Jenkin put it, "Did the enemy have the capacity to determine which of the thousands of messages leaving the country every day was a 'suspicious' one?" The activists needed a typical user of encrypted email—one without clear political affiliation—to find out if their encrypted messages could escape notice in the overall tide of mail. They needed, Jenkin later recalled, to "find someone who would normally use a computer for communicating abroad and get that person to handle the communications."

They had an agent who could try this system out before they switched their communications over to the new approach: a native South African who was about to return to his homeland after working abroad for many years as a

programmer for British telecommunications companies. Their agent would behave just as a typical citizen sending a lot of email messages every day would, using a commercial email provider rather than a custom server and relying on the fact that many businesses used encryption in their communications. "This was a most normal thing for a person in his position to do," Jenkin recalled. The system worked: the agent's messages blended in with the ordinary traffic, providing a platform for openly secret communications that could be expanded rapidly.

Posing as computer consultants, Tim Jenkin and Ronnie Press (another important member of the ANC Technical Committee) were able to keep abreast of new devices and storage technologies, and to arrange for their purchase and delivery where they were needed. Using a combination of commercial email providers and bulletin-board services run off personal and pocket computers, they were able to circulate messages within South Africa and around the world, and also to prepare fully formatted ANC literature for distribution. (The system even carried messages from Mandela, smuggled out by his lawyer in secret compartments in books and typed into the system.) The ordinary activity of ordinary users with bland business addresses became a high-value informational channel, moving huge volumes of encrypted data from London to Lukasa and then into South Africa and between Vula cells in that country. The success of this system was due in part to historical circumstance—personal computers and email (including encrypted email) had become common enough to avoid provoking suspicion, but not so common as to inspire the construction of new, more comprehensive digital surveillance systems such as governments have today.

The Vula network, in its ultimate stage, wasn't naive about the security of digital messages; it kept everything protected by a sophisticated encryption system full of inventive details, and it encouraged its users to change their encryption keys and to practice good operations security. Within that context, however, it offers an excellent example of the role obfuscation can play in building a secure and secret communications system. It illustrates the benefits of finding the right existing situation and blending into it, lost in the hubbub of ordinary commerce, hidden by the crowd.