Safecracking for the computer scientist*

Matt Blaze
Department of Computer and Information Science
University of Pennsylvania
blaze@cis.upenn.edu

DRAFT – 7 December 2004 (Revised 21 December 2004) – DRAFT
The latest version of this document can be found at http://www.crypto.com/papers/safelocks.pdf
This document contains medium resolution photographs and should be printed in color.

Abstract

This paper is a general survey of safe and vault security from a computer science perspective, with emphasis on the metrics used to evaluate these systems and the weaknesses that cause them to fail. We examine security against forced, covert and surreptitious safe opening, focusing on the mechanical combination locks most commonly used on commercial safes in the US. Our analysis contrasts the philosophy and tools of physical security with those of information security, especially where techniques might be profitably applied across these disciplines.

1 Safe and vault security: a computer science perspective

There is an undeniable mystique surrounding safes and vaults. Containers to safeguard valuables and secrets from theft and prying eyes have existed almost as long as the concepts of valuables and secrets themselves, and yet in spite of the "Internet age," details of safes and the methods used to defeat them remain shrouded in obscurity and even a certain amount of mystery. Safe security is a delicate, almost perilous subject, protected by a near reverence that extends, in our imaginations at least, across both sides of the law. *Safecrackers* are perhaps the most romantic and "professional" of thieves, conjuring images of meticulously planned and executed exploits straight out of Hollywood screenplays. And among the law-abiding, safe and vault technicians *(safe men* in the traditional parlance) are perceived as an elite, upper echelon of the locksmithing community whose formidable trade is surely passed on only to the most trustworthy and dedicated.

Reverence for safe work can even be found in the trade's own internal literature, with an almost unavoidable, if subtle, swagger accompanying mastery of safe opening technique. The title of a venerable locksmithing treatise on the subject – *The Art of Manipulation*[LK55] — signals a discipline that demands artistry, not mere craft. Its text begins with a warning to faithfully guard the material in its pages, as well as the suggestion that the book be destroyed completely after its techniques are learned. (Fortunately, some readers have ignored that advice, and a few copies remain available through interlibrary loan). The ambiguity in the term *manipulation* itself seems oddly appropriate here, evoking perhaps a "lock whisperer," with the safe somehow persuaded to open against its better judgment, only to regret it later.

"Security-by-obscurity," if viewed rather dismissively by those in information security, remains a central tenet of the safe and vault trade. It isn't easy to learn how safes work or what makes one better than another, and while the basic techniques and designs are available to those who search persistently enough, few professionals (on either side of the law) openly discuss the details of safe opening with the unindoctrinated. Consequently, it can be difficult for a potential user to judge independently whether a given container is sufficiently secure for its intended application; that role is left primarily to the safe industry itself (although standards bodies and the insurance industry have some influence here as well).

For all the reticence surrounding the subject, however, safes and safe locks (and how they are defeated) are worthy topics of study for students not only of locksmithing but of information security. An unfortunate side effect of the obscurity of safe and vault technology is the obscurity of tools and techniques that deserve to be better known and more widely applied to other disciplines. The attack models against which safes are evaluated, for example, are far more sophisticated than their counterparts in computer science. Many of the attacks, too, will remind us of similar vulnerabilities in computer systems, in spite of having been discovered (and countermeasures developed against them) decades earlier.

The mechanical combination locks used to control access to safes and vaults are among the most interesting and elegant examples of security engineering and design available today. The basic internal structure of (and user interface to) the modern safe lock long predates computers and networks, and yet a careful study of these devices reveals a rich history of threats and countermeasures that mimic the familiar cycles of attacks and patches that irk practitioners of computer and network security.

One of the most striking differences between the physical and information security worlds is the relative sophistication of the threat models against which mechanical security systems are measured. Perhaps owing to its long history and relatively stable technological base, the physical security community – and especially the safe and vault community – generally seeks remarkable precision in defining the expected capabilities of the adversary and the resources required for a successful attack to occur. Far more than in computers or networks, security here is recognized to be a tradeoff, and a quantifiable one at that. The essence of the compromise is time.

1.1 Safe and vault construction

For the purposes of this discussion, a safe or vault is a container designed to resist (or leave evidence of) unauthorized entry by force. (That is, we are discussing *burglary* safes. Many consumer products marketed as "safes" do not actually meet this definition, being intended to resist only very casual pilfering or to protect contents from fire damage; we do not consider such safes here). The difference between a safe and a vault is scale; safes are small containers designed to store objects, while vaults are essentially room-sized safes with features (such as lighting and ventilation) that support human activity.

Many different safe and vault designs are in use, including stand-alone "box like" containers, in-floor safes, in-wall safes, prefabricated vaults and custom made containers; even a superficial survey would be beyond the scope of this document. All share certain common characteristics, however.

Normal access to a safe or vault is via a *door*, which is usually hinged to the container walls. The door is locked shut by one or more *door bolts* (comprising the *boltwork),* which generally are extended or retracted by an external *opening lever,* which can only be operated if a *lock bolt* has been retracted by the locking mechanism (e.g., after dialing the correct combination). Most modern burglary safes accept a standard *lock package* (with an externally-mounted dial), consisting of an internally-mounted lock module with a small retracting lock bolt designed to mate with the door bolts and handle. See Figure 1. (We will discuss these locks in more detail later). Some older safes (as well as certain contemporary low security

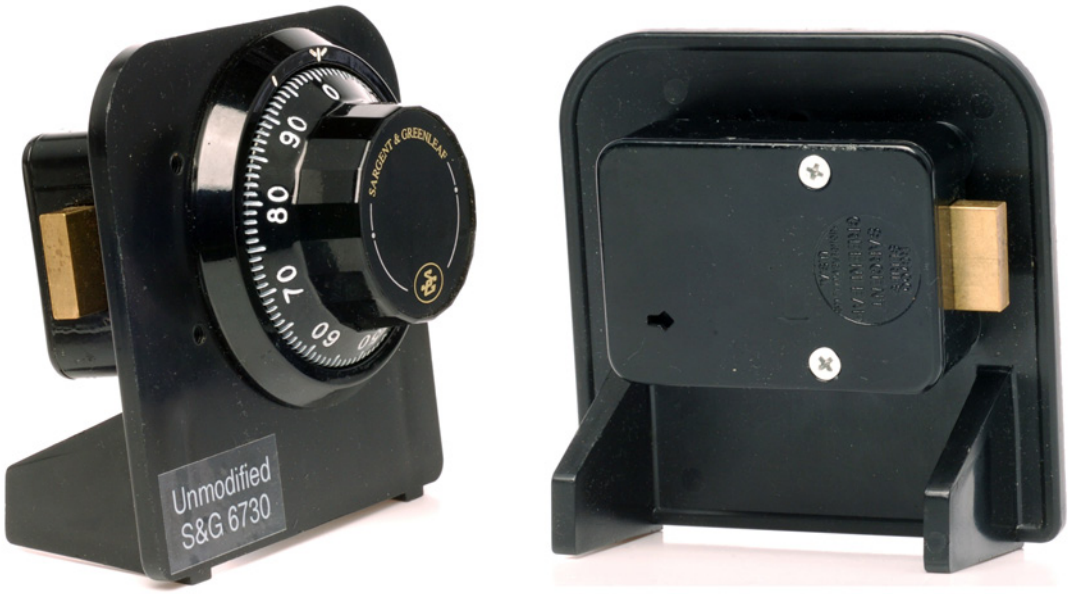

Figure 1: Standard lock package (in this case, a Sargent & Greenleaf model R6730), shown mounted on a display stand. The *dial* (left image) is accessible on the outside of the container. The internal *lock module* (right image), in a standard form factor, contains the lock mechanism and retractable *lock bolt* (the brass tab at the far right). Note the *change key hole* on the back of the lock case, into which the user can insert a tool to change the combination when the container is open. The dial is connected to the lock module via a *spindle* running through a small hole in the container wall.

safes) incorporate a customized lock as an integral component of the boltwork and use the lock bolt directly as the door latch.

The main function of the safe or vault container is to resist opening by force and to protect the lock package from tampering. Container walls and doors usually consist of several layers of material. The outer layer is typically of conventional mild steel, intended to resist blunt force and prying. Resistance to more specialized attacks is provided by *barrier layers*, which are fabricated from materials that resist penetration by various kinds of tools. Barrier materials intended to thwart drilling, called *hardplate*, protect the parts of the safe (such as the lock package) that might be profitably drilled in an opening.

Barrier materials may protect all six sides of a container or, more often, only one (typically the door itself). In-wall and in-floor safes are often protected at the door only, under the assumption that the surrounding environment will prevent access from other directions. To prevent the container itself from being stolen as a whole, stand-alone safes (especially less heavy models) are often designed to be bolted to a floor or wall.

Many safes and vaults (including most burglary safes, but, interestingly, not GSA containers intended for storage of classified materials) include one or more internal *relockers* (also known as *relock devices)* that trigger when certain conditions consistent with an attack are detected. Once triggered, the relockers prevent the door bolts from moving even after the lock bolt is retracted. Several kinds of relockers are in common use. The most common detect *punching* attacks, in which the back of the door is damaged (e.g., dislodging the internal lock package by applying force to the external dial). *Thermal links*, used in some safes, melt and trigger a relock under the high temperatures that might be induced by cutting torches. Some of the highest-end safes include *tempered glass plates* that trigger relock devices when breached by a drill. Lock

packages themselves often have internal relock triggers that prevent retraction of the lock bolt if the lock case is forced open.

Any attack that aims to open the container door must therefore avoid triggering relock devices. The chief value of many relockers seems to be thwarting novice burglars unaware of their existence. Especially on mass-produced safes (the majority of the market), the types and locations of relockers can be predicted and triggering them thereby avoided. On higher-end safes and vaults, however, especially those incorporating tempered glass plates, relockers might be randomly placed as a unique parameter of each instance of the container. Here the relockers force the attacker to employ a more conservative opening technique (e.g., one that involves drilling through more hardplate), making the best-case penetration time slower (and more predictable), even against the expert.

1.2 Container security metrics

Even the best safes and vaults are not absolutely impenetrable, of course; their strength is constrained by both physics and economics. Safes are distinguished from one another not by whether they can be penetrated, but by how long it would be expected to take, the resources required, and the evidence it produces.

The basic security metrics for safes attempt to measure resistance to the kinds of tools that attackers of varying degrees of sophistication might be expected to wield. At the bottom of the attack-tool hierarchy are ordinary hand tools, against which even a low-end safe might be expected to give at least some resistance, then portable motorized power tools, then cutting torches, and finally (presumably for those concerned with international jewel thieves from Hollywood movies), explosives.

We can also measure attacks according to the obviousness of the evidence left behind. Here the terminology is at its most cloak-and-dagger; an attack is said to be *surreptitious* if it leaves behind no evidence at all, *covert* if it leaves behind evidence that would not be noticed in normal use (although it might be noticed in an expert inspection), and *forced* if the evidence is obvious (of course, force might be involved in surreptitious or covert entry as well, so the term is a bit of a misnomer). These distinctions are mainly of interest for safes used to store confidential (or classified) information, where prompt discovery of successful attacks can be almost as important as preventing them in the first place.

Safe and vault rating categories aim to provide a multi-dimensional picture that allows the potential user to evaluate protection according to the perceived threat: a given safe might be rated for a very long time against surreptitious entry aided only by the simple tools of the most casual thief, but for shorter times as the tools used become more sophisticated, heavy, conspicuous, and expensive or as the evidence of attack becomes more pronounced. (Several organizations publish ratings according to various criteria, including, in the U.S., Underwriters Laboratories (UL) for commercial safes and the General Services Administration (GSA) for federal government safes).

Because the materials and mechanical designs from which safes and vaults are manufactured have rather well understood physical properties, relatively simple procedures are used to estimate time bounds on resistance to attack. The usual approach is to make rather generous assumptions about the skill and tools of the attacker and the conditions under which an unauthorized opening might be carried out. For example, a sample safe might be drilled (under laboratory conditions and with the best commercially available equipment and techniques), and the time for penetration considered to be the minimum required for a drilling-based opening by a burglar.

These tests produce safe ratings that may seem disturbingly weak at first blush. The best UL rating categories are for only 15, 30 and 60 minutes, and GSA ratings against forced attack are for either zero(!) or

10 minutes. Yet opening even a zero-minute rated GSA container may require an hour or longer under field conditions (and attract considerable attention in the process).

Observe that safe testing as described here does not produce upper or lower bounds on security in the sense usually used in information security. They are clearly not lower bounds, since better tools or techniques not known when a safe was tested might substantially reduce the required penetration time. The results are not especially meaningful as upper bounds, either, since the conditions are sufficiently generous to the attacker to make it very unlikely that they could be achieved under field conditions. Instead they are less formal "guidelines," intended mainly for comparison, and useful as approximate lower bounds only under the (perhaps tenuous) assumption that improved tools and techniques will not become available in the future.

1.3 Lock security metrics

Time is also the essential metric by which the locks used on safes and vaults are measured. Here, however, we are less concerned with attacks by force, since the sensitive components of the lock are protected by the container itself. Instead, the primary attacks involve exploiting poorly-chosen combinations (birthdays are said to be popular), finding the working combination through exhaustive search, or interpreting incidental feedback given through a lock's user interface to make inferences about its internal state. The latter approach is usually called *manipulation* within the safe and vault trade, although, as we will later see, the techniques involve careful observation more than outright manipulation.

Mechanical combination dial locks are the most common access control devices used on burglary safes and vaults in the United States, and these locks will be the focus of our attention here. Such locks are opened by demonstrating knowledge of the combination by rotating a dial, reversing direction at specific places on the dial; we will discuss the user interface and dialing procedure in detail in Section 2. Electronic combination locks (using a keypad or rotary-encoder dial) are becoming increasingly popular at the low- and high-ends of the safe market, but we will not consider them here; analyzing such locks is essentially a software and embedded system security problem beyond the scope of this paper. Keyed safe locks (usually of a lever-tumbler design) are more common in Europe and elsewhere, but again, they are beyond our scope here.

Nondestructive attacks against the combination itself are usually considered to be in the "surreptitious" category; they leave little or no forensic evidence. (Electronic and electro-mechanical locks may incorporate logs and audit trails, but we are considering strictly mechanical locks here). Many lock attacks, including manipulation, can be performed across several (interrupted) sessions, making them an especially serious threat in some environments.

1.3.1 The combination keyspace

The most obvious lock security factor is the number of distinct combinations; it provides a bound on the time required for exhaustive search. Most safe and vault lock dials are divided into 100 graduations (see Figure 2), with three (or occasionally four) dialed numbers in the combination. This implies 100^3 (1,000,000) possible combinations for a three number lock and 100^4 (100,000,000) possible combinations in a four number lock.

The number of *effectively distinct* combinations is usually considerably lower, however. Most locks have a wider *dialing tolerance* than the dial graduations would suggest, allowing an error of anywhere between $\pm.75$ and ±1.25 in each dialed number, depending on the lock model. So although there may be 100 marked positions on the dial, there may be as few as 40 mechanically distinct positions. A three number

Figure 2: External dial (user interface) of common Group 2 lock (again, a Sargent & Greenleaf model R6730). The actual *dial*, with 100 graduations, rotates; the surrounding *dial ring* is fixed to the container. The main index mark at 12 o'clock (shown here at dial position 2) is used for dialing the combination to open the lock; the smaller index mark at 11 o'clock (at 94) is used only when setting a new combination.

lock would thus have between 40^3 (64,000) and 67^3 (300,763) effective combinations. Other restrictions reduce the combination keyspace a bit further: the selection of the last number is usually constrained to about 80% of the dial, depending on the lock design. With 20% of the last number's space lost, the effective number of distinct combinations on three wheel locks is in practice between 51,200 (with a tolerance of ± 1.25) and 242,406 (with a tolerance of $\pm .75$).[1]

Clearly, even 51,200 combinations would render manual exhaustive search by an unaided attacker infeasible. However, commercially available robotic dialers (a servo motor attached to the dial controlled by a simple microcontroller) can search the effective keyspace of most three-number locks, as well as some four-number locks, overnight or over a weekend (however, this is still generally longer than the expected required penetration time for the container itself; the repeated, high speed dialing also introduces significant wear on the lock).

The size of the combination keyspace is one of the most important metrics used in the certification of safe locks by various standards bodies. In the United States, UL rating standards for Group 2 safe locks (the most common commercial locks) specify that there must be at least 1,000,000 different combinations and that the dialing tolerance be at most ± 1.25. (The standard does not explicitly address the number of *usable* combinations that can actually be set by the user, however, and so three-number locks can be certified even when the last number is constrained). Comparable standards in other countries demand specific minimum sizes for the combination keyspace more directly. CEN, the European standards body, requires at least 80,000 distinct usable combinations for "Class A" locks (roughly equivalent to UL Group 2). VdS, a similar German standard, has the same requirement for its "Class 1" rating (again, roughly equivalent to UL Group 2).

Unfortunately – and ironically – a significant further reduction in combination keyspace comes from overly broad "guidelines" concerning the choice of "good" combinations. Presumably to compensate for

[1] One of the better examples on the market in this regard, the (Group 2) Sargent and Greenleaf R6730 lock, has a dialing tolerance of $\pm .75$ and allows the use of 94% of the dial for the last combination number, yielding a usable combination keyspace of roughly 282,807 distinct combinations.

the notoriously poor ability of users to select sufficiently "random" combinations, many lock manufacturers recommend avoiding selection of combinations that do not "look random." A typical example is Sargent and Greenleaf[Cos01], which recommends for its three-number locks the combination as a whole not consist of a monotonically increasing or decreasing series, that adjacent numbers differ by at least ten graduations[2], and that 25% of the dial be avoided for the final number (although the mechanism itself on S&G locks requires avoiding only 6% of the dial). Acceptable combinations under these recommendations comprise less than 50% of the usable combination keyspace. For example, while the S&G R6730 lock has 282,807 distinct usable combinations according to its mechanical specifications, only 111,139 of them are considered "good" according to the manufacturer's recommendations. For locks with the full ± 1.25 dialing tolerance allowed under UL Group 2, these recommendations seem especially misguided, leaving only 22,330 distinct "good" combinations. Observe that this is less than 2.5% of the apparent keyspace of 1,000,000.

Similar reductions in effective keyspace will be familiar to observers of many computer password authentication systems.

1.3.2 Manipulation resistance

Some combination lock designs, including those used on burglary safes, are subject to imperfections that leak information about their internal state through the external dial user interface. It may be feasible for an attacker to exploit this information to discover a working combination by *manipulation*, the systematic entry of trial combinations and interpretation of state information.

An obvious security metric for a combination lock, therefore, is whether the design (and its fabrication processes) resists manipulation attacks. Because elaborate equipment is not generally required to perform these attacks, the most significant variable in the threat model is whether the attacker is familiar with and practiced in the technique. Ratings for these locks distinguish between "expert" and "non-expert" attacks.

There are two classes of commercial (UL-rated) safe locks in the United States. "Group 1" locks are intended to resist expert manipulation for at least twenty hours; in practice this means the best attack against such locks should be exhaustive search. (A sub-category, "Group 1R," also requires resistance to radiological analysis, perhaps the only lock attack assumed to involve special tools). "Group 2" locks provide only "moderate" resistance to manipulation, but are considered secure against non-experts. (Locks in a recently introduced sub-category, "Group 2M," are said to resist expert manipulation for up to two hours).

The vast majority of commercial safes use Group 2 (and sometimes even unrated) locks. Even apparently formidable containers that might require significant effort to penetrate by force are often equipped with locks that can be manipulated open with no evidence by anyone familiar with the procedure. Group 1 locks are usually found only on high-end safes and vaults intended specifically for the storage of high-value items or classified materials.

The relative rarity of mechanical locks designed to resist expert manipulation seems somewhat surprising, especially given that the containers on which they are used are often quite secure against penetration by experts. Group 1 locks are not substantially more expensive than their Group 2 counterparts, and certainly would not represent a significant increase in the overall cost of a container. The likely explanation is that mechanical Group 1 locks typically have a more complex user interface, usually requiring an additional step before unlocking, and are less forgiving of dialing errors. Just as in computing systems, many users are willing to exchange even a significant degree of security for improved usability and convenience.

[2] This recommendation also may slightly improve resistance to manipulation.

1.4 Using security metrics

Time for successful attack is used as a metric of safe and vault quality not only because it is somewhat measurable, but because it is exactly the property that the security engineer must know in order to design and evaluate the system as a whole. If a safe or vault is trusted to resist a particular kind of attack (e.g., drilling) for a particular amount of time, (e.g., 30 minutes), we can conclude that the system is secure as long as the conditions for the attack cannot occur for longer than the period specified (e.g., guards and alarm sensors that prevent people with drills from having unsupervised access to the safe for more than 30 minutes).

It is notable that analogous security metrics do not generally exist for information systems. In particular, while some measures of required attack resources do exist (e.g., for cryptographic work factors), in practice the resources required to attack most information security mechanisms are either completely unknown or are known only with low confidence. When such metrics are available, the usual design principles of computer and communications security consider a system to be secure only when the work factor is so large to make any possibility of attack completely infeasible (e.g., requiring turning every molecule in the solar system into a supercomputer). In physical security, perhaps because the security metrics are believed to be more realistic, much smaller "safety margins" are generally tolerated.

In other words, the tools of information security generally group systems into one of three categories (completely secure, completely insecure, or, most commonly, unknown security), with few meaningful ways to compare systems within a given category. The measurement tools of physical security, on the other hand, recognize finer shades of security, allowing comparisons to be made in which one system might be considered more secure than another for a given purpose. It is unclear which approach is sounder; attacks occur in both domains, of course. In any case, the principles of physical security design and evaluation richly repay careful study by computer scientists, and the development of similar metrics for information security would represent a significant advance in the field.

2 Group 2 mechanical combination locks

The modern dial combination lock mechanism is relatively simple, and its basic design has remained essentially unchanged for at least a century. There are relatively few variations from the standard design, although some models incorporate extra security features (e.g., to meet Group 1 standards).

The most common are the Group 2 locks, and among them, most current products use a "spring loaded lever-fence, key changeable" design. The "standard" such lock is the Sargent & Greenleaf model R6730. Other current locks employing a virtually identical design as the R6730 include the Kaba-Ilco model 673 and the LaGard model 3330. Because the design is so common, and also because it illustrates the basic principles of operation (and security pitfalls) of mechanical combination locks well, it will be the focus of our attention here. (We will discuss variants on the design later).

The standard external user interface is via a roughly 3 inch diameter rotating dial mounted to the door of the container and graduated into 100 positions. A dial ring, fixed to the door's wall, has a primary index mark (usually at 12 o'clock) for dialing the numbers of combination, plus a second index mark (usually at about 11 o'clock) that is used only when changing the combination. See Figure 2.

The dial is connected to the internal lock module via a *spindle* running through the container wall that serves as the dial's axis and that rotates along with it. The major internal components of (Group 2) lock modules are shown in Figure 3.

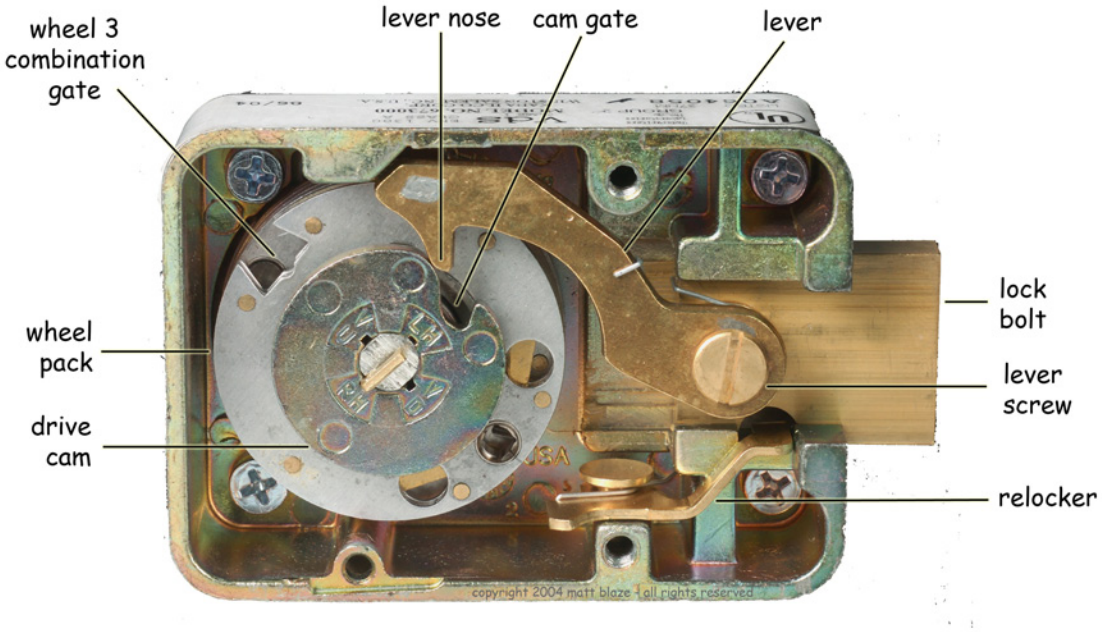

Figure 3: Major components of Group 2 lever-fence lock, as seen from the back (Kaba-Ilco model 673)

Although many of the lock components serve more than one purpose, with complex interactions that depend on the lock state, the design is simpler than it might first seem. Recall that the purpose of the lock is to retract the lock bolt (and thereby release the door bolts) only after a correct combination has been entered. It is easier to understand the design as a whole by studying its two basic functions separately – retracting the lock bolt and enforcing the combination.

2.1 Retracting the lock bolt: the drive cam and lever

The two main internal components involved in retracting the lock bolt are the *drive cam* and the *lever*.

Within the lock module, the spindle terminates at a *drive cam* (also known as the *cam wheel* or simply the *cam*). The cam moves with the external dial, with all rotational movement of the dial transmitted directly to the cam. (On most locks, including that shown in Figure 3, the cam is the rear-most element, but that is not essential to the design.) Observe that the cam is circular with a wedge-shaped notch cut in its circumference; the notch is called the *cam gate*.

The lock bolt slides partly into or out of the lock within a channel in the side of the module's housing (i.e., to the left or the right in the figures here). The bolt is attached within the lock module to the *lever*. The lever is attached to the bolt with a *lever screw*, which acts as a pivot point for the lever, allowing it to move upward and downward across a range of a few degrees. The lever is pressed downward by a *lever spring*, which is usually wound around around the lever screw.

The lever runs within the module from the lock bolt to near the spindle axis. At the far end of the lever, the *lever nose* rests along the edge of the cam, held down by the pressure of the lever spring. Observe that the lever nose is in the same wedge shape as the cam gate. The bolt is moved by allowing the lever nose to mate with the cam gate.

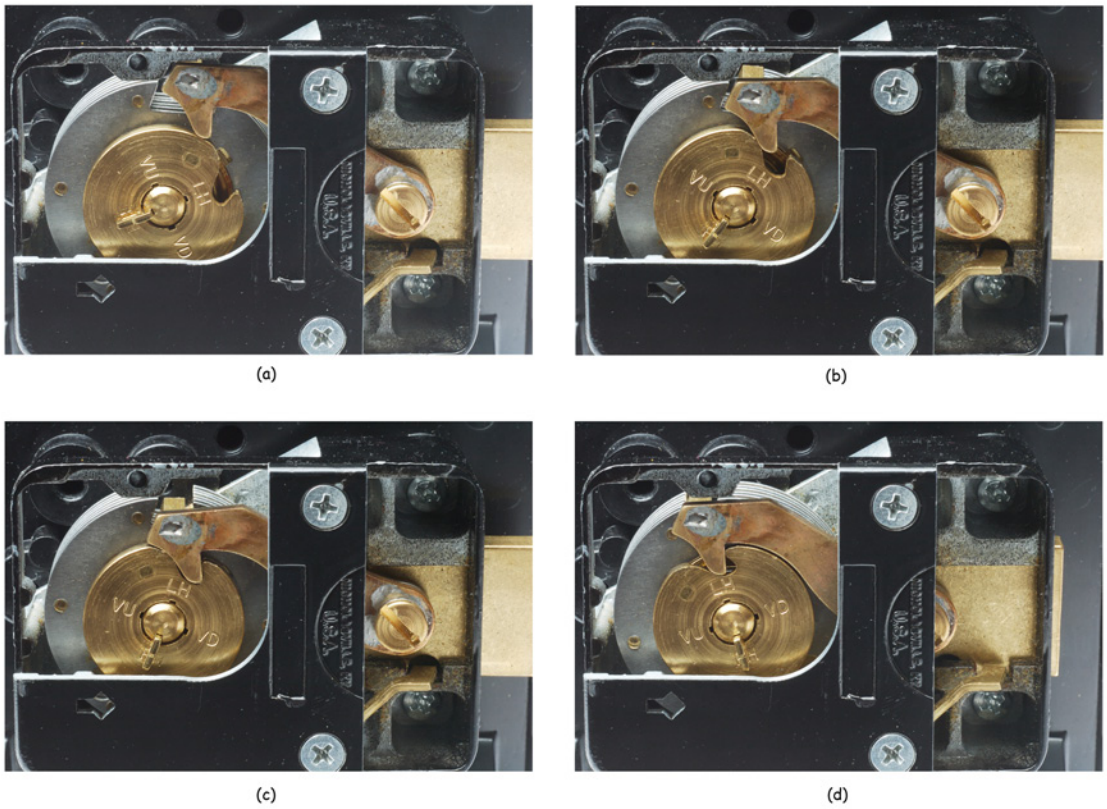

Figure 4: Opening a lock. In this cutaway view from the back of the lock, the dial is rotated clockwise to retract the lock bolt. In *(a)* through *(c)*, the cam rotates toward the lever, allowing the lever to lower as the nose mates with the cam gate. In *(d)*, further rotation of the cam has pulled the lever to the left, which retracted the bolt.

In Figure 4 the dial (and hence the cam) is rotated clockwise[3] and the cam gate approaches the lever nose. As the cam gate moves under the lever nose, the nose is pushed downward in to the gate. Continued clockwise rotation, with the lever nose fully mated with the cam gate, pulls the lever, which in turn retracts the bolt. Once the bolt is fully retracted, the dial cannot be turned further clockwise; counterclockwise dial rotation extends the bolt back to the locked position.

2.2 Enforcing the combination: the fence and wheel pack

As described so far, our lock can retract and extend its bolt but does not have any security; it is opened by simple clockwise dial rotation. Two additional components, the *wheel pack* and the *fence*, interact with the lever and cam to allow the bolt to retract only after a correct combination has been dialed.

The *wheel pack* is the set of "security tumblers" for the lock. It is mounted behind the cam around the spindle, but does not make direct contact with the spindle itself. The wheel pack consists of a collection of disks (called *wheels)*, of larger diameter than the cam and that can rotate independently of the cam and of one another. In the edge of each wheel is a notched *combination gate* (or simply a *gate)*. See Figure 5. (On

[3] In most figures here the view is from the *back* of the lock, and so clockwise and counterclockwise rotation appear reversed.

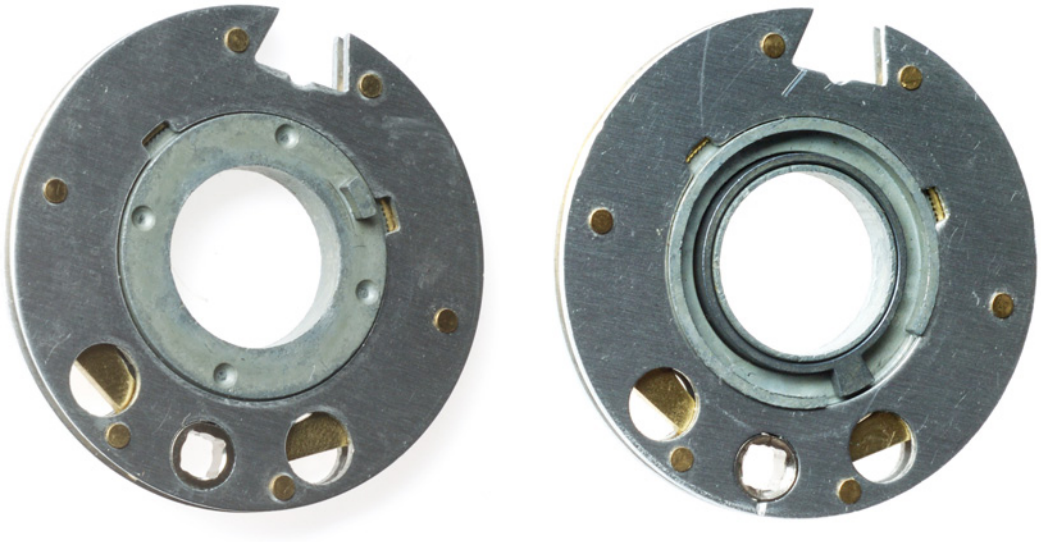

Figure 5: Wheels (front and back). *Left:* Front; note the *drive pin* protruding near the axis. *Right:* Back; note the *fly* near the axis, which mates with the adjacent wheel's drive pin. This wheel (for a Kaba-Ilco 673 lock) has a *movable* fly, designed to rotate within a fixed range before moving the wheel; this allows the same combination number to be dialed either clockwise or counterclockwise.

most locks, the wheels are actually a three-layer "sandwich"; see Figure 6).

The combination gates are tested by the lever. There is more to the lever than Figure 3 suggested; as we can see in Figure 7, above and behind the lever nose is a horizontally protruding *fence* positioned just above and across the top of the wheel pack. When the cam gate is under the nose, the lever lowers, allowing the fence to make contact with the wheel pack. If all of the combination gates in the wheel pack are aligned under the fence, the lever nose continues to lower and can engage the cam gate. See Figure 8. But if even one of the combination gates is not under the fence, the lever is prevented from falling low enough for the nose to engage the cam gate. See Figure 9.

2.3 Dialing the combination

The ability to retract the bolt depends on positioning all of the combination gates of the wheels in the wheel pack under the fence. Each wheel configures one number of the combination; a three number lock has three wheels in its wheel pack. An ingenious coupling arrangement among the cam and the wheels allows the use of the external dial to position each wheel separately.

Attached to and protruding behind the cam is a small tab, called the *cam drive pin*, which rotates with the cam around the spindle and which rides in a slotted ring in the wheel directly behind the cam. Attached to the wheel, and within this slot, is another tab, called the *fly*. As the dial, and hence the cam, rotates, the cam drive pin moves within the slot of the wheel behind it, and, at some point in the revolution, hits the wheel's fly. Thus after one complete rotation of the dial, not only does the cam rotate with the dial, the wheel behind the cam does as well.

Protruding behind and attached to that wheel is another drive pin, which rides within a slot on the wheel behind that one (and which has its own fly). So while after one rotation of the dial, only the first wheel

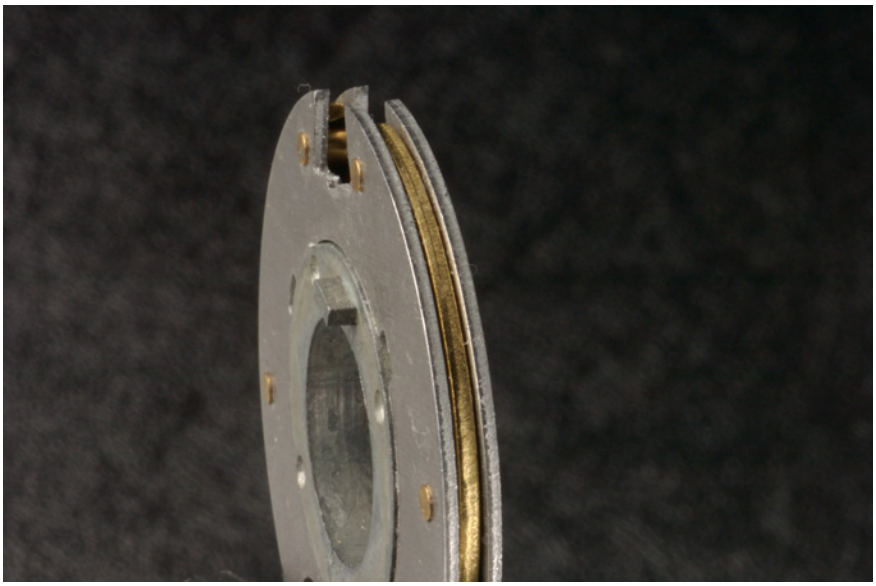

Figure 6: Edge view of a typical wheel. Note the three-layer "sandwich" construction, which facilitates changing the combination.

.

Figure 7: Lever. Note the *fence* protruding from above and behind the nose.

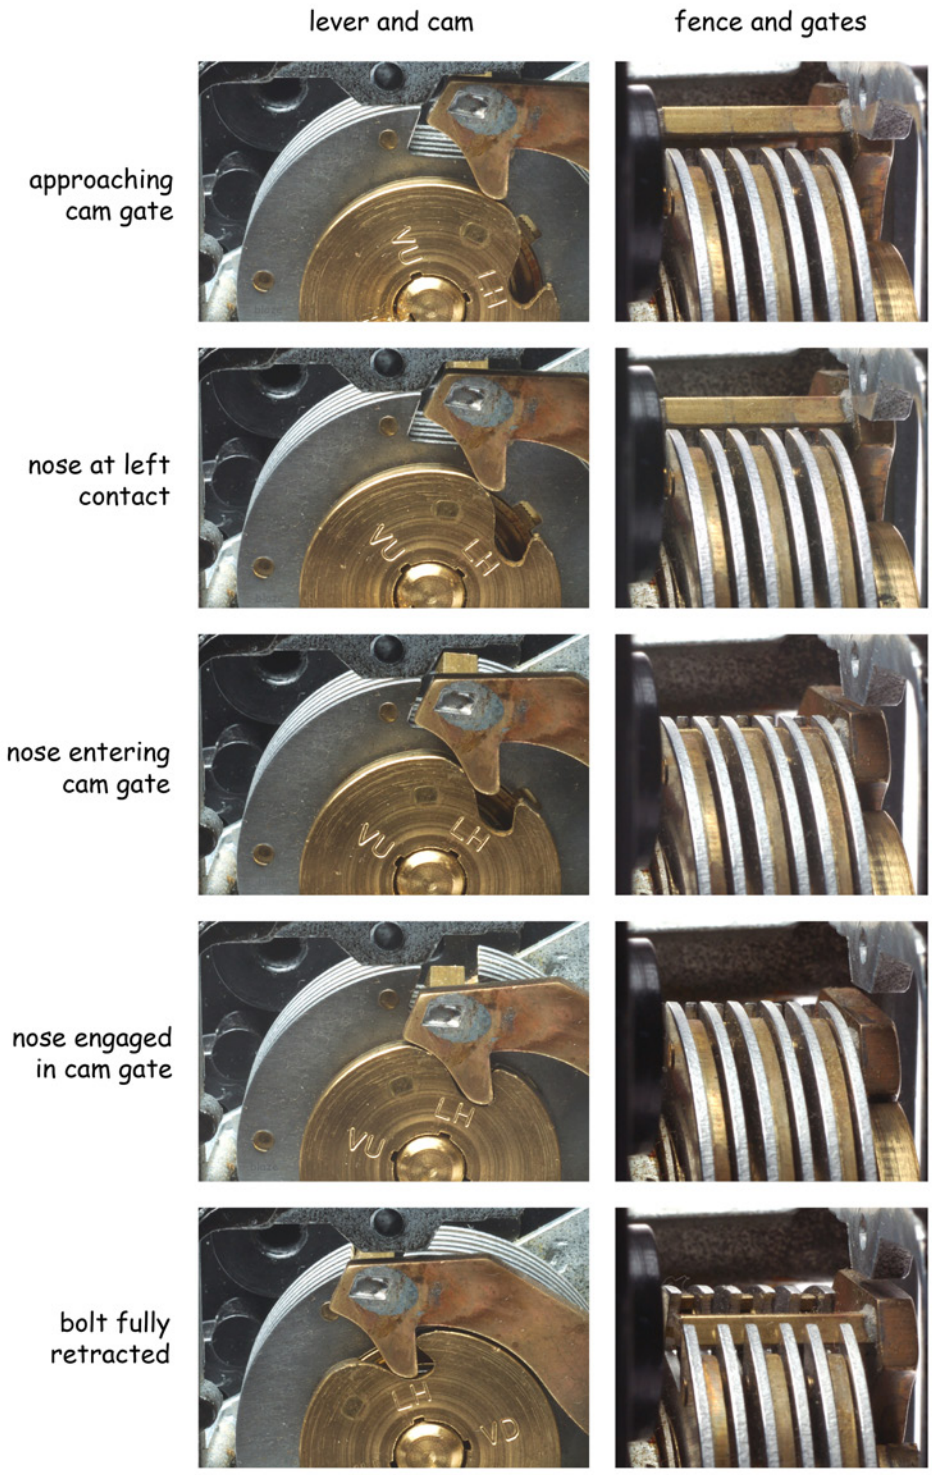

Figure 8: Interaction of lever and fence with cam and gates, correct combination set

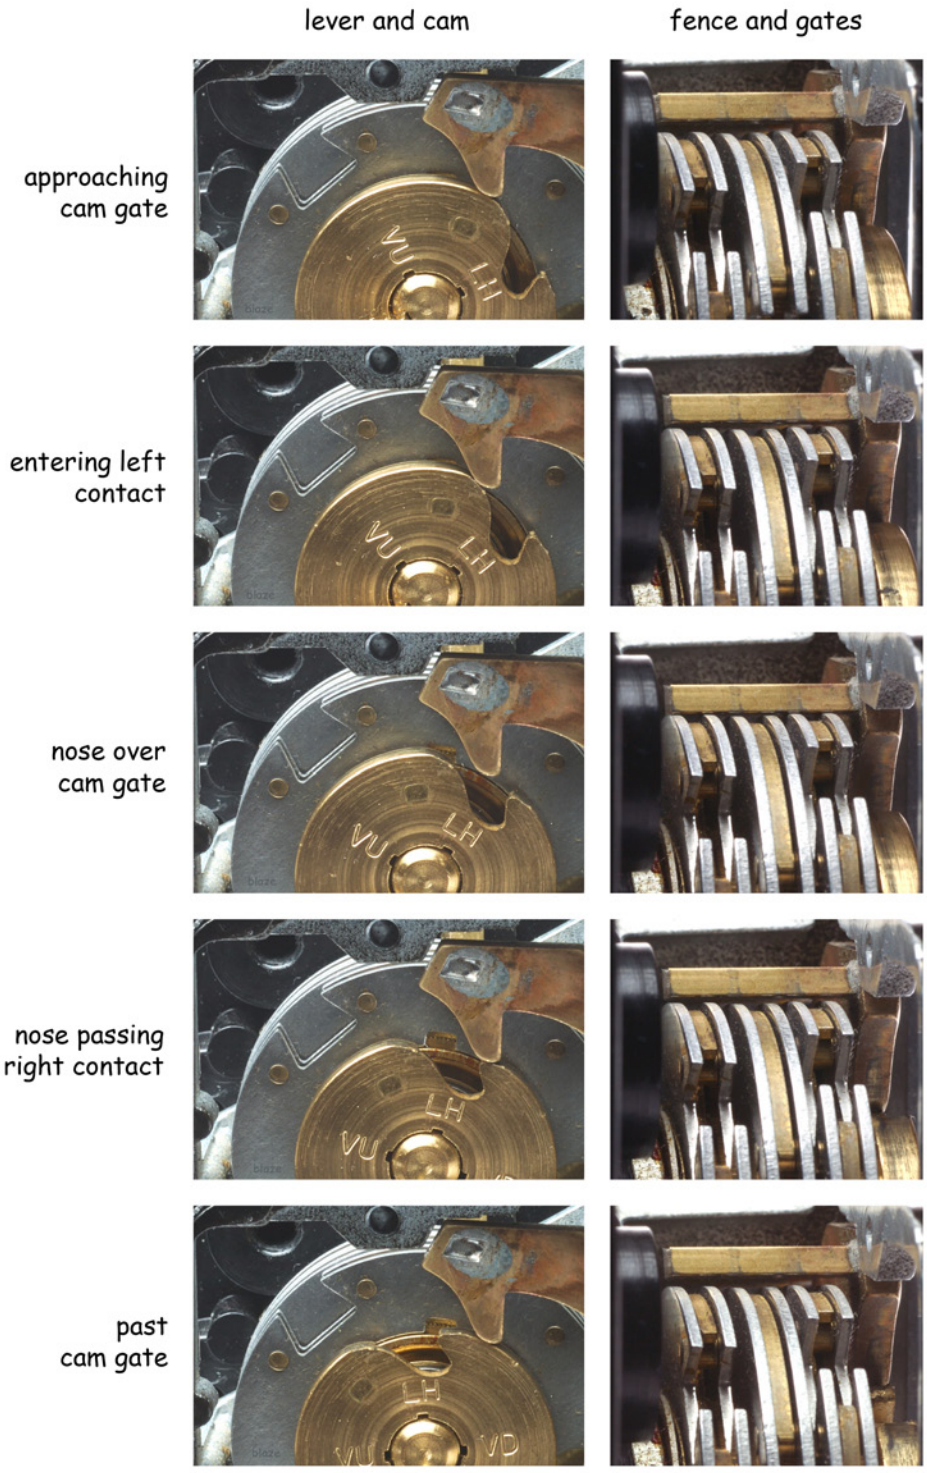

Figure 9: Interaction of lever and fence with cam and gates, combination not set

rotates, after two rotations, the second wheel rotates as well. Arbitrarily many wheels can be "stacked" this way (although most locks have three wheels, with some having four and a few specialized high security locks having more than that). When a wheel's drive pin is in contact with another wheel's fly and is able to rotate it, we say that the new wheel has been "picked up."

Note that no matter how many wheels have been picked up, reversing the direction of rotation (e.g., from clockwise to counterclockwise) breaks the cycle, and additional complete rotations are required in the new direction to begin picking up the wheels again. The wheel farthest from the drive cam must therefore be put into position first, followed by the next, with the wheel adjacent to the drive cam positioned last.

Using the index mark and dial graduations, the drive pin and fly pickup scheme allows rotation of the dial to set the wheels to any specific configuration desired. Combination entry requires the user to keep track of the numbers being entered and manage multiple dial rotations; entry of all three numbers may require turning the dial as many as ten complete rotations. The usual procedure is roughly:

1. Turn dial counterclockwise (to the left), stopping when the first combination number is aligned with the top index mark the *fourth* time.
2. Turn dial clockwise (to the right), stopping when the second combination number is aligned with the top index mark the *third* time.
3. Turn dial counterclockwise (to the left), stopping when the third combination number is aligned with the top index mark the *second* time.
4. Turn dial clockwise (to the right) slowly, until the lock bolt retracts and the dial will turn no further. (If dial does not stop after two complete revolutions, the combination was dialed incorrectly and must be re-entered from the beginning.)

On an N-wheel lock, the wheel directly behind the cam is usually considered wheel N, because although that wheel picks up first, it corresponds to the last dialed number of the combination. So on a three number lock the wheel behind the cam is wheel three, the wheel behind that wheel two, and the last, wheel one.

A wheel's combination is determined by the relative position of its fly and its gate. Changing a lock's combination is a matter of changing this relationship. On some (mostly older) locks, the wheels must be removed and swapped out to change the combination. The wheels used in many locks (including those shown here) allow the user to change the combination without disassembly. A small tool inserted in the back of the lock case "unlocks" the inner parts of the wheels (with the fly and drive pin) from the outer parts (with the combination gate), allowing a new combination to be entered. The procedure is a bit cumbersome (it usually requires dialing to a different index mark), but is considered an end-user operation.

2.4 Other considerations and design variants

The lever-fence design is subject to somewhat anomalous behavior if the combination of the last wheel is set too near the point at which the nose enters the drive cam gate. Usually, the lever nose will become trapped in the cam gate, preventing the bolt from being re-locked. More rarely, the lock will fail to open altogether. This is the reason that the range of numbers allowable for the last combination is restricted, avoiding those that would position the last wheel gate too close to the cam gate. This region of the dial is usually called the *forbidden zone*, and applies only to the last number of the combination.

In Figure 3 (and most of the other figures here), the lock is shown (from the back) with the lock bolt on the right, which is the mounting position used in many safes and vaults. However, depending on the container, the lock may also be mounted with the bolt facing downward, upward, or to the left. Most locks allow the dial to be attached at any of four 90° rotations so that the opening position remains near zero (with the index at 12 o'clock) even when the lock is mounted in one of these ways. (A removable *spline key* attaches the spindle to the cam inside the lock case; it is installed in one of the four slots marked RH, VU, LH or VD to determine the dial orientation with respect to the cam.)

Different vendors of these locks produce some of the parts to slightly different specifications than shown here and may position the components somewhat differently. For example, Sargent and Greenleaf and Kaba-Ilco locks use three-layer "sandwich" wheels with the gate cut in the larger diameter outer layers (i.e., as shown here), while LaGard locks put the gate in a larger diameter inner layer. Sargent and Greenleaf, Kaba-Ilco and LaGard locks have the drive cam as the rear-most element, behind the wheel pack (e.g, away from the dial), while Mosler locks make it the front-most element (between the dial and the wheel pack). A number of different lever designs are in use, and on some locks, the dialing tolerance can be adjusted by swapping the lever for one with a wider fence (for a tighter tolerance) or a narrower one (for a looser tolerance).

An older variant on the design, the *direct-entry fence*, does not use a spring-loaded lever-fence or the lever/cam arrangement as described here. Instead, the door bolt handle itself moves the fence into the wheel pack to engage the gates. Direct-entry locks are usually considered to have lower security than the Group 2 design shown here and are found primarily in older containers as well as in many of the more inexpensive current-production fire safes.

Locks designed for use on thick-walled vault doors often do not use the direct spindle-cam coupling show here, with an indirect geared drive arrangement employed instead.

Locks with "high security" features might have additional components. Most Group 1 locks have some mechanism to further restrict the interaction of the lever nose and the cam gate. Some Group 2M locks add an irregularly-shaped ball or wheel to the lever nose. The security implications of these features will be discussed in Section 3.3.6.

Most Group 1 and Group 2 locks use wheels stamped from brass or other easily-machined metals. Group 1R locks, on the other hand, use wheels made from low-density materials (that is, plastics such as Delrin) to resist X-ray analysis.

GSA containers and vaults used to protect DoD-classified materials have historically used Group 1R locks, but current GSA containers no longer use mechanical locks at all. The only locks now approved for this application are electro-mechanical, with the lock mechanism controlled by an embedded microcontroller. (Only one vendor, Kaba-Mas, has products certified for this application as of this writing).

2.5 Observations

With small variations, the design described here is almost ubiquitous among contemporary safes that use Group 2 locks. The basic components are also similar to those used in other kinds of high- and low- security combination locks; a working knowledge of this mechanism can serve as a foundation for a more general study of lock security.

Any safe or vault lock has two basic operational requirements: it must not open if the correct combination has not been entered, and it must reliably open if it has. We have paid attention primarily to how this design achieves the former, but the latter requirement is just as critical, since a lock that fails to un-

lock in response to its combination is effectively trapped in a sealed container. In practice, locks of this design are remarkably reliable, although because of the somewhat baroque user interface (especially with respect to combination changing), failures caused by user error are not uncommon. Potential failure modes, and procedures for diagnosing and ameliorating them, are beyond the scope of this paper but are treated comprehensively in [Tob00], [Oeh97] and [Cos01].

There is much to admire from an engineering perspective about these locks. Many of the components serve more than one security purpose. The lever has at least two distinct functions: testing the wheel pack for correct gate alignment and retracting the bolt. The cam has at least three: transmitting dial rotation to the wheel pack, controlling contact between the fence and the wheel pack, and engaging the lever to retract the bolt.

Some information about this mechanism's internal state can be learned through the external dial interface, but much of it is of little use to the attacker. Even a relatively inexperienced user can easily feel the wheels pick up one by one, as well as the falling and rising of the lever as the cam gate moves under, and then away from, the lever nose. This feedback reveals the position of the cam gate as well as the position at which each wheel picks up[4].

It is therefore straightforward to use the dial to distinguish among three different states of the drive cam and wheel pack:

1. Nose not over drive cam gate.
2. Nose over drive cam gate; fewer than N wheel gates under fence.
3. Nose over drive cam gate; all N wheel gates under fence (bolt retracts).

At least on the surface, the security of the basic Group 2 lock design appears to be quite good. Although the user interface leaks some state information, the design obscures the most important information that an attacker seeks, which is the position of the gates. The lever nose and drive cam lift the fence above the edge of the wheel pack and do not allow the gates to be felt as they pass below the fence.

However, in spite of these security features, this lock design is not perfect. In particular, small (and largely unavoidable) manufacturing imperfections can introduce significant vulnerabilities by leaking more state information about the wheel pack than our analysis so far might otherwise suggest is possible. We will discuss these imperfections, and how they can be exploited, in Section 3.3.

3 Attacks against containers

3.1 Forced entry: brute force

Forced entry is the most obvious attack against a safe or vault, and countermeasures against it are relatively straightforward. In an attack by force, the goal is to open the container as quickly as possible, without regard to the evidence left behind or the future repairability of the container. The primary defense against brute force is brute strength: heavy and/or hardened materials that resist prying, cutting, and bending.

[4]Knowledge of the positions at which the wheels pick up is not usually useful to the attacker, since this reveals only the positions at which the wheels were last set, not the locations of the gates. However, if the last user to open the lock failed to rotate the dial a sufficient number of times before closing it again, this may reveal some or all numbers of the combination. A version of this attack against locks used at Los Alamos during the Manhattan Project was amusingly reported by Richard Feynman in [FHL85]

Most, although certainly not all, successful forced attacks do not aim to open the door but rather breach an unprotected side of the container. While forced attack is most commonly associated with amateur, opportunistic burglary, a notable professional application of forced entry against an unprotected side is in bank vault penetration. Bank vault doors are very expensive and notoriously difficult to breach. When they fail (which is unusual), it is sometimes less expensive to bore a large hole in the side of the building than to attack the door itself. (This itself is a non-trivial operation, since the vault walls are typically constructed from thick, steel-reinforced concrete).

Analysis of, and countermeasures against, forced attack is primarily a matter of metallurgy, mechanical engineering and materials science, and is beyond our scope here.

3.2 Covert entry: drilling

A more refined use of force, by *drilling*, aims not to breach a large opening into a container, but rather to release its boltwork and open the door as if a correct combination were entered. Carefully-planned drilling attacks that open the door are not only often faster than those against the container, they may be sufficiently nondestructive that the container can be repaired and put back into service. Such attacks are therefore said to be *covert*, since although there may be evidence of the opening, it may escape casual notice once repaired[5].

Because the aim is to open the door normally by releasing the boltwork, drilling usually attacks the lock package. A drilled container is opened by one of three methods, depending on the location of the hole or holes. The lock state might be observed and the combination *decoded*, so that it can then be dialed in the normal way. The lock might be probed to manually reset it to a known combination or to the unlocked state. And finally, parts of the lock or lock bolt might be destroyed outright, to allow the bolt to retract as if the lock had been opened, *bypassing* the lock altogether.

3.2.1 Principles of destructive decoding and bypass

A small opening that allows visual (or small tool) access to the lock module can be very powerful, especially against containers that are operating normally (i.e., those with relockers that are not triggered, etc.).

On containers with properly operating locks, the combination can be decoded easily by observing the edge of the wheel pack as the dial is rotated. A single small inspection hole drilled through the container wall and lock case is usually sufficient for this purpose. Any opening that gives a view of the gates is sufficient for this purpose.

If the inspection hole provides a view of the fence as well as the wheel pack edge, decoding is simply a matter of lining up the gates under the fence by turning the dial (or spindle) in the usual way.

As long as the inspection hole allows viewing the gates as the wheels rotate, the combination can be decoded, even if the fence's position is unknown, although it requires two steps rather than one. First we find the combination that aligns all gates to the same position. Each number in this combination will be "out of phase" from the true opening combination by a fixed offset. The true opening combination can be found by simple exhaustive search of the possible offsets. (If the location of the fence with respect to the inspection hole is known, there is no need for the exhaustive search – the combination can be calculated simply by adding the distance between the inspection hole and the fence location).

[5]Drilling is a destructive activity that is not undertaken lightly, especially against safes that are to be repaired and put back into service. Also, drilling of some older fire and burglary safes can disturb hazardous materials, such as asbestos and tear gas.

Figure 10: A 10 inch by .200 inch 90° fiber optic borescope with portable 30 Watt halogen light source. An angled borescope provides a good view of the gates and/or the fence through a small opening.

The inspection hole is usually drilled into the lock case through the front of the container door, but need not be. Recall that the back of most locks have a "change key" hole for inserting the tool that changes the combination. An opening that provides a view of the change key hole may provide a sufficient view of the wheel pack to allow decoding, especially with the aid of an angled borescope such as that shown in Figure 10.

A small opening that exposes the wheel pack is useful even when the lock is not operating properly and where rotation of the spindle does not properly engage the wheel pack. A sharp probe (such as an ice pick) can often be inserted through the inspection hole to force the wheels into position under the fence and to retract the lever.

Careful and selective destruction of various lock components can obviate the need to decode the combination at all. The simplest such attack is to drill into the lock module through the front of the container to destroy the fence. This allows the lever to engage the cam without first lining up the gates. (Performing this kind of attack requires knowledge of the exact layout of the lock module on the container door and precise measurement of the drilling location). Another possible attack drills out the lock bolt, although this area is often difficult to reach or protected by relockers.

3.2.2 Determining drilling points

Nowhere is the sense of secrecy surrounding safe and vault security greater than it is around the selection and location of effective drilling points on various safes. Where to drill is perceived as one of the darkest and most carefully guarded secrets of the safe trade, something that can be learned only through careful analysis and cataloging of hundreds of different containers.

Determining where to drill is actually quite straightforward. While the dimensions and internal structure of the container are indeed useful to know when drilling (and there are commercially available databases of safe diagrams and drilling points sold in the trade), effective drilling locations for most safes can be determined without any special knowledge of the container. Only very high security safes – those with glass plate relockers – require special consideration in calculating where to drill. For other safes, the optimum drill point is determined entirely by the lock itself.

In fact, the same drilling point is effective against virtually all modern commercial safes that use direct-

drive (non-geared) Group 1 and Group 2 mechanical locks of the kind discussed here (except for the rare high security containers with glass plates). Any security supposedly deriving from the obscurity of workable drilling points for such containers is illusory at best; their location is a fixed and almost universal constant.

Recall that the goal of a drilling-based decoding attack is to gain visual access to the edge of the wheel pack. The wheels in most locks are of a standardized diameter - about $1\ 3/4$ inches. Therefore, because the wheels are mounted around the spindle, the edge of the wheel pack on virtually all modern locks can be found $7/8$ of an inch – in any direction – from the center of the spindle.

Most locks are mounted "RH" (with the bolt on the right as viewed from the back). The lever on RH-mounted S&G R6730 locks, and many similar locks, is at about dial position 97 (with the dial at 0). This means that drilling a hole centered $7/8$ of an inch from the center of the spindle, at a compass point of about $350°$ will provide a good view of the edge of the wheel pack and of the fence. If the lock is not mounted RH or if the fence on a particular lock model happens to be in a different position, the lock can still be decoded with an offset, as discussed in Section 3.2.1. A hole anywhere from $1/4$ to $3/8$ of an inch in diameter is sufficient for this purpose. (Some safe technicians prefer to drill slightly farther from the spindle – say, 1.0 inches – to allow the top of the wheel pack to be viewed with a $90°$ borescope).

Drilling is more complex if the lock is not mechanical or is of an unusual design, if relockers have been triggered (as in an unsuccessful burglary attempt), or if glass plate relockers protect the lock case. In such cases, reference to commercial databases of relocker and bolt locations and diagrams of the container greatly facilitate efficient opening. But for most safes, a hole $7/8$ of an inch from the spindle center is completely sufficient.

More detailed information about the specific safe model is also valuable when drilling is done through the container's side, back, or top, (e.g., to gain access to the change key hole). Side drilling is usually done to avoid hardplate or glass relockers in the door for safes that have barrier layers in the door only. Effective side drilling requires knowing the mounting orientation of the lock and the precise thickness of the door. It also requires the use of a sufficiently long borescope to reach the change hole (or other opening in the lock case).

Note that the standard drilling point of $7/8$ of an inch from the spindle center puts the hole within the dial. The outer part of the dial must therefore be removed (e.g., with a $1\ 5/8$ inch hole saw) for access to this drill point. There are also commercially available "dial pullers" that can remove the dial from the spindle, but they do not work effectively on every kind of safe dial. Once the dial is removed, the wheels can be rotated with a small "emergency" dial or by turning the spindle directly.

Other lock designs require slightly different drill points, but the basic principles are the same. Geared drive locks require knowledge of the location of the wheel pack, which is off-center from the dial spindle. Direct-entry fence locks of the kind used on inexpensive safes often do not have a wheel pack enclosed in a separate lock housing and are therefore very forgiving of small errors in drilling location. They can usually be drilled almost anywhere near the wheel pack that allows a view of the gates.

3.2.3 Hardplate drilling techniques

The most important tools involved in safe drilling are specialized bits intended for cutting through hardplate.

The standard mild steel on the outside of most safes can be easily penetrated with a conventional cobalt drill bit designed for metal; these are widely available on the consumer and contractor market. Medium-speed drilling at about 2000 RPM is sufficient for this material, and special tools are not required.

Figure 11: Drill bits optimized for penetrating safe hardplate. (a) is a low-speed diamond-tipped bit designed for use at 200-500 RPM. Note the hollow core, which dispenses the blue lubricant under the heat of drilling. (b) and (c) have a hardened tungsten-carbide tip with a shape designed for cutting through steel (most carbide drills are designed for masonry). These can be used from about 400-3000 RPM depending on the specific hardplate material and user preference.

However, all but the lowest security safes have some kind of barrier material – hardplate – protecting the lock, especially on the door. A variety of hardplates are used, generally a proprietary alloy or composite formulation specific to a particular safe manufacturer. All such materials are very difficult to penetrate with conventional drill bits, which are not hard enough to cut through hardplate effectively.

Most hardplate can be relatively easily penetrated with special drill bits designed for this purpose. High- and low-speed lubricated diamond-tipped bits are occasionally used, but most hardplate can be efficiently drilled with tungsten-carbide bits shaped for use on metal. (These are not the same as the inexpensive carbide bits used in the general contractor market, which are optimized for cutting masonry, not metal). See Figure 11 for examples of $1/4$ inch bits intended specifically for safe and vault hardplate.

Most hardplate drilling is done under relatively high pressure at low or medium speed. (High speed drilling and drilling with diamond-tipped bits is done at moderate or low pressure.) Maintaining control of pressure and depth is easiest with a portable *drill rig* attached to the container that holds the bit in the correct position. A magnetic drill rig designed specifically for safe and vault opening is shown in Figure 12.

There are many different hardplate materials, and some require special techniques to drill efficiently. Some require high speed and moderate pressure, while others require lower speed and higher pressure. A full survey is beyond our scope here; some of the techniques require experience and practice to use successfully. For example, some hardplate barrier materials, particularly those made from composite carbide alloys, are best penetrated by a combination of drilling and *punching.* Before drilling and at periodic intervals a small hardened punch is struck against the work surface to break and dislodge hardened pieces of material and to smooth the surface to be drilled. A hardened pin (or old drill bit) is then used to further smooth and create a small dimple in the surface before drilling is resumed. Materials requiring punching are identified by their

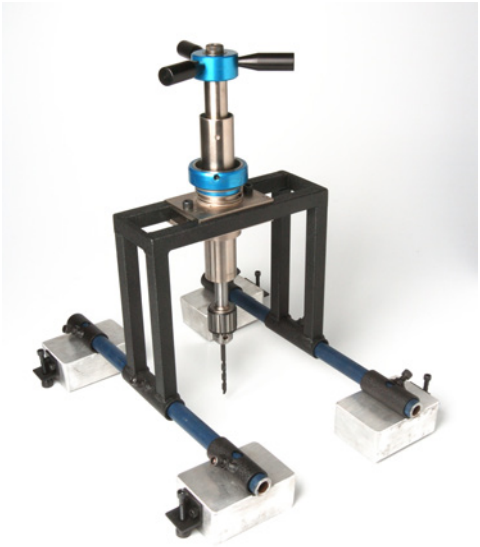

Figure 12: A portable drill rig designed for safes and vaults. As shown, the base is attached to the side of the container by powerful permanent magnets; it can also be attached with machine screws or nylon straps. An ordinary low- or medium- speed handheld electric drill (not shown) supplies power.

rough and irregular surface.

Although hardplate materials resist drilling, they are not impervious to it, and the equipment required is surprisingly ordinary[6]. With the proper bits, ordinary handheld drills (of the kind used by contractors and sold in home improvement stores) provide more than sufficient power for hardplate penetration, especially when used in conjunction with a drill rig. However, while hardplate may not be impenetrable, it succeeds at its intended function: increasing the minimum time required for drilling attacks, even by experts.

3.3 Surreptitious entry: lock manipulation

The precision with which combination locks can be produced is subject to inherent physical (and economic) limitations. It is inevitable that any components fabricated in a conventional industrial process will have small, variable imperfections. Wheels that are supposed to be round in fact will have a slightly irregular shape, pivot points and rotational axes will be slightly off center, angles will be slightly wrong, component sizes will vary, and parts intended to fit snugly will allow some play. Although the standard Group 2 combination lock design tolerates many of these imperfections quite well, others make it possible to construct detailed inferences about the state of the wheel pack at various points of the dial's rotation. By careful analysis of this state information, it is often possible to discover efficiently a lock's combination. Use of these analytical techniques is usually called *manipulation.*

Manipulation attacks are especially powerful because they are inherently surreptitious; they interact with the container through its normal user interface and leave behind essentially no forensic evidence. In fact, unlike attacks by drilling, they require no elaborate tools (beyond a pencil and paper)[7].

[6] A more esoteric tool, not discussed here, is the *thermal lance* (also called the *thermic lance* or *burning bar),* which can penetrate most hardplate with great speed as well as sufficient precision to make it useful for decoding attacks.

[7] Computers and signal processing techniques may be helpful, however. A recent commercial product, the *Soft Drill,* uses a laptop computer to control a servo and a transducer and is said to be able to perform manipulation attacks against most Group 2

Figure 13: Imperfections in wheel pack and fence. Only the "largest" wheel actually determines the depth to which the fence lowers. Sometimes the wheels are of slightly different diameter (as shown here), and sometimes the fence is not exactly parallel with the wheel pack. In this cutaway view of an off-the-shelf S&G 6730, the middle wheel is slightly larger.

3.3.1 Manipulation principles

Two properties of the Group 2 lock design render it vulnerable to manipulation attacks.

The first property is *imperfect wheel/fence alignment.* Recall that the combination is "tested" by lowering the fence along the edge of the wheel pack at a fixed position, allowing the nose to engage the cam only if the fence can enter the gates. If at least one wheel in the wheel pack has its gate elsewhere, the fence can go no lower than the edge of the wheel pack. If the lock were perfectly manufactured, when no gate is under the fence the fence would rest on all three wheels simultaneously. But since the lock cannot be perfectly manufactured, in fact the wheels will be of slightly different diameter and the fence will not be perfectly parallel with the axis on which the wheels ride. This means that, in practice, the fence is blocked from lowering not by all wheels, but only by an *effectively largest* wheel. When that wheel is rotated so that its gate is under the fence, the fence will be able to lower slightly more, but will then be prevented from lowering further by the *next* "largest" wheel. That is, although a complete lowering of the fence requires positioning the gates of all wheels, the exact depth to which the fence can lower at any given time is actually determined by only a single wheel. See Figure 13 for an example of this phenomenon in a typical commercial lock.

The second property is the *amplification of fence depth through the nose and cam gate.* Recall that the lever nose and cam gate are roughly wedge shaped. When the nose is fully engaged in the cam gate, it is a snug fit, with very little lateral play. But when the nose only partially lowers into the cam gate, there is considerable play from side to side. In fact, the total amount of play is inversely proportional to the depth of the nose in the cam gate (and hence the depth at which the fence touches the largest wheel in the wheel pack).

The play of the nose in the cam gate is readily observable through the external dial interface (as

locks in about 20 minutes

discussed in Section 2.5, even casual lock users often notice the change in resistance as the nose enters and leaves the cam gate region). More specifically, if the dial is rotated to the left with the nose over the cam gate, we can feel the *left contact point* as the cam edge begins to raise the lever. If the dial is rotated to the right, we can similarly feel the *right contact point.* (We usually call the part of the dial between the left and right contact points the *contact region*). See Figure 14 for a (somewhat exaggerated) comparison of the play between the left and right contact points with a wheel pack with no gates aligned and with the gate of the largest wheel aligned.

The precise amount of dial travel between the left and right contact points gives us the ability to measure – from outside the lock – the relative height of the largest wheel in the wheel pack with respect to the lever. (Note that this ability does not depend on any manufacturing imprecision; it is a natural consequence of the nose and cam gate design.)

Accurate measurement of changes in the amount of dial travel between the left and right contact points often allows us to determine the location of the gate on the largest wheel. When that wheel's gate is positioned below the fence, the nose rides lower in the cam gate and the amount of dial travel in the contact region is reduced. That is, at "low points" in the wheel pack the range between the left and right contact points narrows. Manipulation is simply a systematic, adaptive probing and analysis of the distance between the left and right contact points that derives part or all of a lock's combination.

In effect, we can use the lock's own user interface to perform a kind of "differential analysis" against the wheel pack, in which we observe how changes in the wheel pack configuration introduce corresponding changes in the distance between the left and right contact points. Careful observation and analysis of these differences allows us to deduce more about the internal state of a lock than just the three states in Section 2.5. On a three-wheel lock we can often distinguish among five states:

1. Nose not over drive cam gate.
2. Nose over drive cam gate; no wheel gates under fence.
3. Nose over drive cam gate; one wheel gate under fence.
4. Nose over drive cam gate; two wheel gates under fence.
5. Nose over drive cam gate; all three wheel gates under fence (bolt retracts).

3.3.2 Measuring relative fence depth

Manipulation is primarily an analytical activity; in spite of its name, success does not depend on unusual manual dexterity or great sensitivity (popular myths about stethoscopes and sandpapered finger notwithstanding). Indeed, many people find it easier to master the basic mechanical skills of manipulation than, e.g., those of pin tumbler lock picking.

However, manipulation does require some non-intuitive mechanical and observational technique that must be learned and practiced.

Manipulation exploits small but observable changes in the size of the contact region, and so the most important manipulation skill is reading the left and right contact points to measure the relative fence depth. There are three elements to this skill. The first (and simplest) is finding the approximate location of the contact region. The second is consistent recognition of exactly where the left and right contacts occur. The third is reading the precise dial position (to about $1/4$ or, preferably, $1/8$ of a dial graduation) of the contact points.

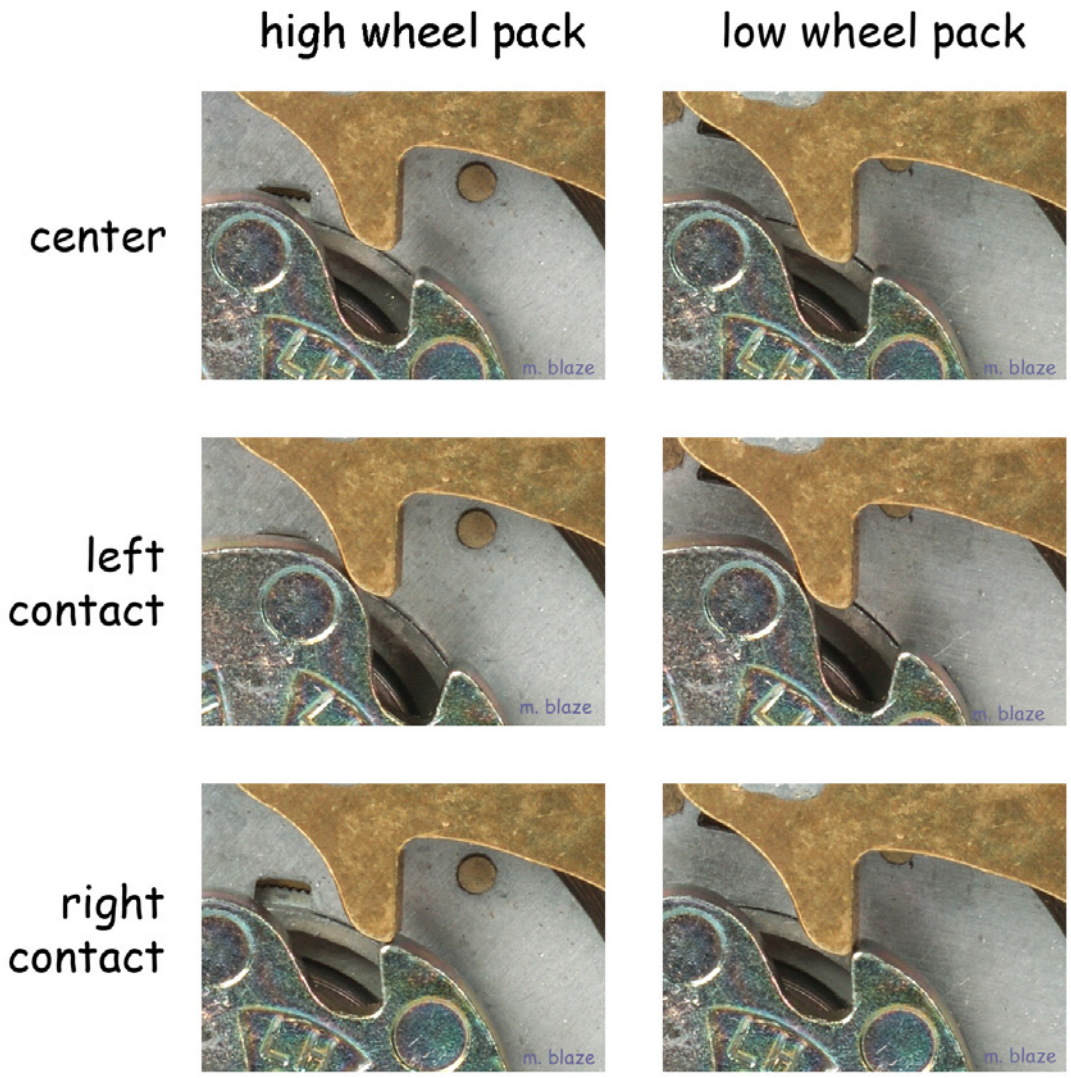

Figure 14: Play of lever nose within cam gate, high wheel pack vs. low wheel pack. The contact region (the range between the left and right contact points) is narrower when the lever can fall lower.

The first, and easiest, required psychomotor skill is locating the contact region. Various locks put the contact region in slightly different places. For example, current production S&G R6730 locks usually have the contact region between about 4 and about 13, while Kaba-Ilco locks usually use between about 99 and 7. Regardless of where the contact region is, it is not difficult to locate. The contact points can be distinguished from the other main source of dial feedback – wheels being picked up – by the fact that it remains at about the same dial location regardless of the state of the wheel pack.

The next psychomotor skill, consistently recognizing exactly when contact has occurred, requires practice. The contact point is determined by dialing from *within* the contact region *outward* (e.g., if the contact region is from 5-10, we would check the left contact point by dialing counterclockwise toward 10 and the right contact point by dialing clockwise toward 5). Note that the exact location of the "contact point" is somewhat ambiguous; different people will consider them to occur in slightly different places. Since manipulation requires only noting *changes* in the contact points, not measurement against an absolute value, this is not a problem. The important thing here is consistency; given the same lock and wheel pack state, one must be able to recognize the contact point as being at the same point every time.

A fluid dialing motion and slow, steady dial movement are critical here. The dial must be free to move easily; the dial ring on some locks will occasionally slip and bind, but this can be corrected by re-centering the ring. Note that there will sometimes be an inward/outward play in the dial, which can cause inconsistent results. One can correct this by pulling gently on the dial before taking each reading.

The final psychomotor skill is reading the precise contact point from the dial. Again, consistency is the most important problem here, not accuracy against an absolute standard. A precision of $1/4$ of a graduation or better is needed to distinguish wheel low points on most locks. It is often helpful to use the index marks themselves as a guide. On many locks the marks are about $1/4$ of a graduation thick.

A small complication in reading the relative fence depth is when the last wheel is set to a number in the contact region. Since moving the dial across the contact region would disturb the last wheel's position, readings must be taken at the one side, and the dial rotated around to take readings at the other side. Alternatively, since having the last wheel set to such a number would place it in the forbidden zone and such combinations would be illegal, the last wheel could simply be moved out of the way to a fixed position before readings are taken.

Most Group 2 locks use a cam gate and lever nose that is more gently sloped on one side than on the other. Consequently, a given change in the relative fence depth will often cause a larger change in the contact point on the more gently sloping side (usually the left) than on the other. Some locks will not show any changes in the right contact point at all during the manipulation process, with meaningful readings appearing only on the left.

In general, mechanical aids are not required for manipulation, although a few simple devices can make it easier to read precise dial locations. The simplest is a sticker that fits around the dial ring showing $1/4$ graduation increments, as shown in Figure 15.

An even better device, unfortunately no longer commercially marketed but relatively easily homemade, is a Vernier scale that fits to the dial and ring, making it easy to read the dial to a precision of $1/8$ of a graduation.

Audio aids are not usually considered to be especially helpful, although amplifiers with magnetically-attached transducers are sold for this purpose by safe and vault supply houses. In general, however, contact points are most reliably identified by feel rather than by sound.

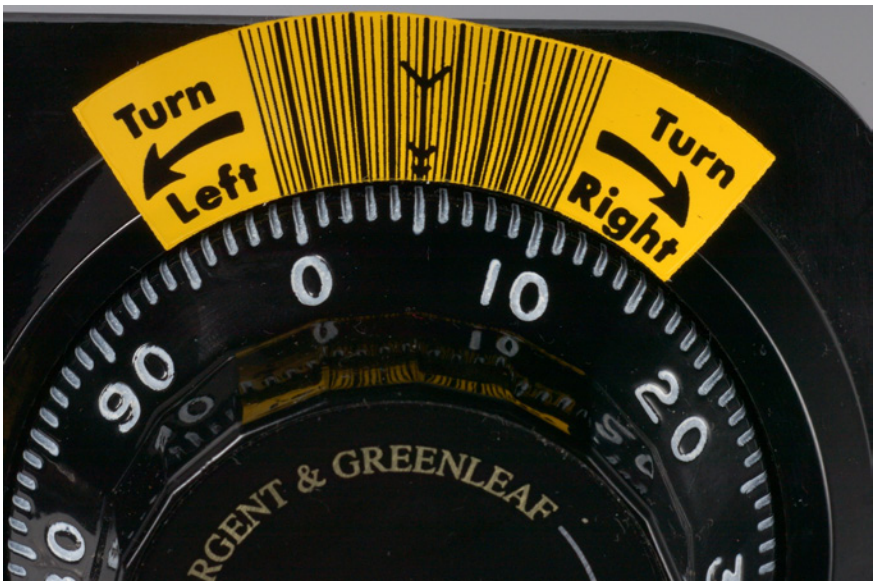

Figure 15: A simple manipulation aid – a sticker marked with a $1/4$ graduation scale.

3.3.3 Wheel pack analysis: the simple case

The goal of the manipulation process is to find successively "lower" wheel pack configurations as the fence lowers into each gate.

The best procedures to decode a given lock's combination through manipulation depend not only on the lock design, but also on the specific imperfections present in it. Different procedures that take into account how different classes of imperfections affect the state of the wheel pack can make the manipulation process more efficient. Although many different imperfections will be present simultaneously, it is simplest to regard the aggregate error behavior of a given lock as if it were due to a single *dominant imperfection.*

We first examine approaches suitable for locks with the most common dominant imperfection. An adapted version of the procedure that is effective against locks with most other dominant imperfections will be discussed in Section 3.3.4.

The most common dominant imperfection is a single wheel that is "largest" with respect to the fence regardless of its rotational position or that of the other wheels in the wheel pack, and where this size difference is sufficient to detect unambiguously the gate location. This might be caused by a fence that is not exactly parallel with the wheel axis or by wheels that are actually of different diameter. Fortunately, in addition to being the most common case, this is also the simplest situation from an analytical perspective. No matter what position the other wheels are in, the relative fence depth is always determined by the same wheel.

This means that we can find the gate of the largest wheel by testing the relative fence depth around the circumference of all wheels simultaneously. This, of course, tells us the location of the low point (the gate) of the largest wheel but will not tell us on which wheel it located. Additional tests will quickly identify the wheel, however. We can then move on to find the gate on the next largest wheel, and so on.

Step 1: First we catalog the relative fence depth at sufficiently many wheel positions to construct a "map" of the largest wheel that identifies its approximate gate location (as given by the size of the contact region). Even on locks with relatively tight dialing tolerances, it is usually sufficient to test the wheel pack at

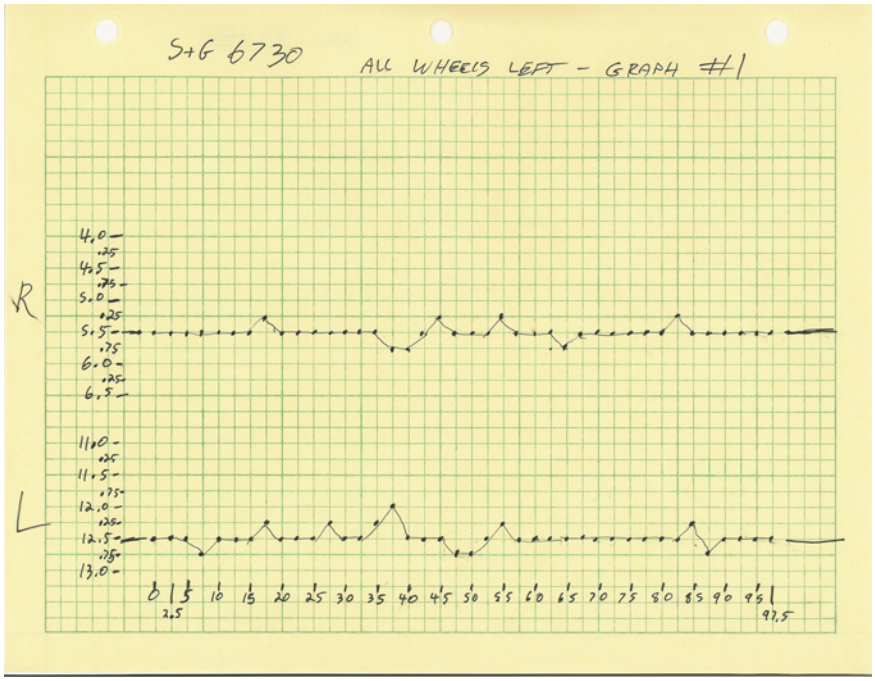

Figure 16: Traditional graph of the left and right contact points (taken every 2.5 dial graduations with contact points recorded to a precision of $1/4$ graduation) for a typical lock whose dominant imperfection is a single largest wheel. Note the prominent narrow region between the graphs – corresponding to a low relative fence depth – at 37.5.

40 positions, at intervals of 2.5 dial graduations (e.g., at 0, 2.5, 5, 7.5, 10, 12.5, etc)[8]. Each test is performed (on a three wheel lock) by rotating all wheels to the left four times, stopping at the number under test, and measuring the relative fence depth by returning to the contact region using the techniques of Section 3.3.2. On locks with (correctly operating) movable flys, wheels are positioned at the same place whether rotated to the left or to the right, so this procedure is sufficient to map wheels with combinations dialed in either direction[9]. (As noted in the previous section, testing wheel positions that are within the contact region requires entering the contact region twice – once from the left and once from the right).

Traditionally, the map is produced on graph paper, with the left and right contact points graphed (to a precision of between $1/4$ and $1/8$ of a graduation) at each tested number. Adjacent graphs of the left and right contact points facilitate identification of the lowest point, as shown in Figure 16.

Step 2: At this stage we know the approximate location of the low point (and hence the gate) of the largest wheel. The next step is to narrow down the exact position of the low point, to find the center of the gate. A simple and efficient way to do this is to simply repeat the mapping procedure of the previous step at the eleven graduated dial positions ± 5 from the identified low point.

Step 3: Finally, we identify the wheel on which the low point was found. For a three wheel lock, we can measure the relative fence depth with three test combinations, one for each wheel. In each test combination, we dial the low point on two of the wheels but some number *other than* the low point on the

[8]Even if the dialing tolerance is tighter than ± 1.25, it is usually adequate to test only at every 2.5 graduations. The fence will usually partially lower into the gate even if it is too far away to drop far enough into it to allow the nose to fully engage the cam gate.

[9]Most modern Group 2 locks have movable flys; many direct-entry fence locks, however, have fixed flys and stop at slightly different positions depending on whether they were dialed from the left or the right.

wheel under test. We should measure a higher relative fence depth with the test combination that moves the largest wheel away from the discovered low point.

Once we know one number of the combination we can discover the low point (and gate location) of the *next* largest wheel. We repeat steps 1–3 as before, except that instead of dialing all wheels to the same number when mapping the wheel pack, we must *park* the known wheel at its low point. For example, if for the largest wheel we had discovered a gate at 38 located on the second wheel, we would repeat the mapping process of Step 1 as before, except this time testing 0–38–0, 2.5–38–2.5, 5–38–5, and so on. We similarly park the previously mapped wheel at its known gate location when we narrow down the next wheel's exact low point and when we identify which wheel it is.

The combination of the final wheel is most efficiently discovered by a straightforward exhaustive search; there is no need to explicitly map its low points, since dialing a trial combination with the correct number at this stage will retract the bolt.

Observe that the complexity of decoding the combination under this dominant imperfection is bounded by $MN + N(N-1) - 1$ trials, where M is the number of distinct wheel positions and N is the number of wheels. On a three wheel lock with 66 distinct positions this requires at most 203 trial combinations, as opposed to M^N (287,496) under naive exhaustive search. Furthermore, these trial combinations can be performed (and contact points recorded) across several interrupted sessions over an arbitrarily long period of time, assuming that the combination remains unchanged.

Finally, note that while in practice mapping the left and right contact points onto a graph gives an easily interpreted visual representation of the high and low points of the wheel pack, producing the actual graphs is not always essential. It may be faster to simply record the maximum right contact points and minimum left contact points (noting when a new maximum or minimum is discovered as the dial is worked around).

3.3.4 Wheel pack analysis: complex cases

If the dominant imperfection allows the wheels to be mapped independently of one another, the procedures of Section 3.3.3 will efficiently discover the combination in one pass for each wheel. Other dominant imperfections do not allow independent wheel mapping, however. If the axis is off center, for example, the wheels will have an eccentric rotation, with different wheels being largest with respect to the fence at different rotational positions. If the wheels are not perfectly round, there may be one or more significant low and high points not associated with any actual gates. (Some lock vendors use wheels with *false gates* that are deep enough to allow a lower fence depth but not deep enough to allow the nose to engage the cam gate; such wheels will map with multiple low points.) As we will see, however, locks with such imperfections can still usually be decoded efficiently. In fact, a slight refinement on the procedure of Section 3.3.3 will decode most locks regardless of the exact nature of their dominant imperfections.

The solution for locks with these kinds of dominant imperfections is to simply iterate the mapping process with wheels parked at their lowest points, the aim being to find successively lower relative fence depths. For example, if the initial mapping of all wheels indicates a low point on the third wheel at position 53, we would park the third wheel at position 53 and find the low point on the next wheel. Assume the next low point is at position 25 on wheel two. Now we could re-map wheel three with wheel two parked at 25, hoping to find a lower point than 53.

Several different combinations of wheels may have to be parked when wheels are re-mapped. However, a systematic mapping of successively lower points will usually yield the locations of all three real gates with surprising efficiency.

Figure 17: Studying mounted “cutaway” locks can increase proficiency with wheel pack operations, but such locks are of only limited use in practicing manipulation *per se.*

3.3.5 Learning manipulation

Manipulation is a practical threat only to the extent that an attacker is able to become proficient in its skills and techniques. The threat would be minimal if the technique required years of practice or if only a few individuals were gifted with the requisite sensitivity. In fact, compared with many other surreptitious attacks against locks (such as lock picking), manipulation is not especially difficult to master.

It is usually easiest to learn manipulation by focusing on its three elements separately: setting the lock state, the analysis process, and consistent reading of contact points.

The most basic, and arguably most important, element of lock manipulation is having enough familiarity with the lock mechanism to be able to visualize the current state of the wheel pack and to confidently enter trial combinations. This is more important than it may seem, since a single incorrectly entered trial combination can easily ruin the entire analysis of a lock’s state. Mounted “cutaway” locks, such as the one shown in Figure 17, are helpful for this purpose.

Next, it is usually easiest to learn and practice the manipulation analytical process in a way that does not depend on a high degree of psychomotor skill in reading the contact points. Mounted (but not cutaway) practice locks, with exaggerated dominant imperfections, can be very useful here. The simplest way to produce such locks is to bend the fence slightly upward.

Finally, reading the relative fence depth should be practiced with a variety of lock samples. Individual locks and those of different manufacturers and production runs will have different characteristics and dominant imperfections. These variations become readily apparent with experience and practice.

Successful decoding of Group 2 locks does not require extraordinary talent, but it does demand a systematic approach and a moderate degree of practiced skill. The procedure usually requires about 30 to 60 minutes (in one or more sessions). Interestingly, a novice manipulator may not require appreciably more time to decode a given lock than an expert, but will be less successful against locks with more subtle imperfections. The number of test combinations that must be dialed depends on the lock, not the manipulator, and

so while the accuracy and consistency of contact readings (and therefore the ability to open more difficult locks) depends heavily on skill, the time required to decode a given lock largely does not.

3.3.6 Design variants and manipulation countermeasures

Not every Group 2 lock can be manipulated, of course. Some individual samples will have been made with sufficient precision that meaningful differences in relative fence depth cannot be detected.

Other lock designs besides that shown here can be manipulated but may require different techniques. For example, direct-entry fence locks do not have a "contact region" from which the relative fence depth can be derived. Instead, the relative fence depth on such locks is proportional to the amount of handle movement. The usual manipulation technique involves attaching a long needle pointer to the handle and a scale to the container and reading the relative fence depth by noting the precise amount of handle movement.

Some wheel packs (especially those used in direct-entry fence locks) include shallow false gates. Once recognized, they present little difficulty against manipulation, since the false gates serve as low points from which the other wheels' gates can be be detected.

Group 1 locks include additional features to make manipulation more difficult even when significant imperfections are present. The design of these locks is beyond our scope here, but the general approach involves preventing the contact points from being read accurately. Many Group 1 locks employ a secondary mechanism to hold the fence well above the wheel pack until the nose is already within the contact region, and prevent the dial from moving again until the fence has been retracted back into position. These mechanisms appear to make manipulation by conventional manual techniques infeasible. See Figure 18 for examples of manipulation-resistant locks.

4 Conclusions

Are safes and vaults secure by computer science standards? To be sure, most containers can be attacked, and indeed, some of the vulnerabilities are subtle and surprising. And yet, when compared against their counterparts in information security, the mechanisms that protect safes are remarkably successful. Few weaknesses in physical security admit the kinds of catastrophic failures common in computers and networks, in which a low-risk, low-cost attack can yield a high-value and easily replicated benefit. Even the most sophisticated attacks against safes, whether involving force or lock manipulation, almost always entail at least some risk of exposure. Relatively accurate estimates of the time and other resources required for various kinds of attacks make it possible to tightly optimize effective physical security systems and to coordinate complementary security mechanisms.

Physical security has been studied for far longer than information security, of course, and the tradeoffs between resistance to attack and the cost of protection are relatively well understood. The situation in computer security is quite different, with new mechanisms, attacks, countermeasures, and threat models being invented and made obsolete in a dizzyingly fast cycle that lacks the luxury of generations of hindsight.

There is much that information security can learn from physical security, and a careful study across the two disciplines should strengthen both of them. One of the most interesting aspects of physical security's methodology is its ability to very closely measure both the capabilities of the attacker and the resistance of various mechanisms to specific threats, as well as to *compose* these metrics in useful ways (e.g., to determine the required response time of an alarm system). Nothing approaching these kinds of metrics exists in information security.

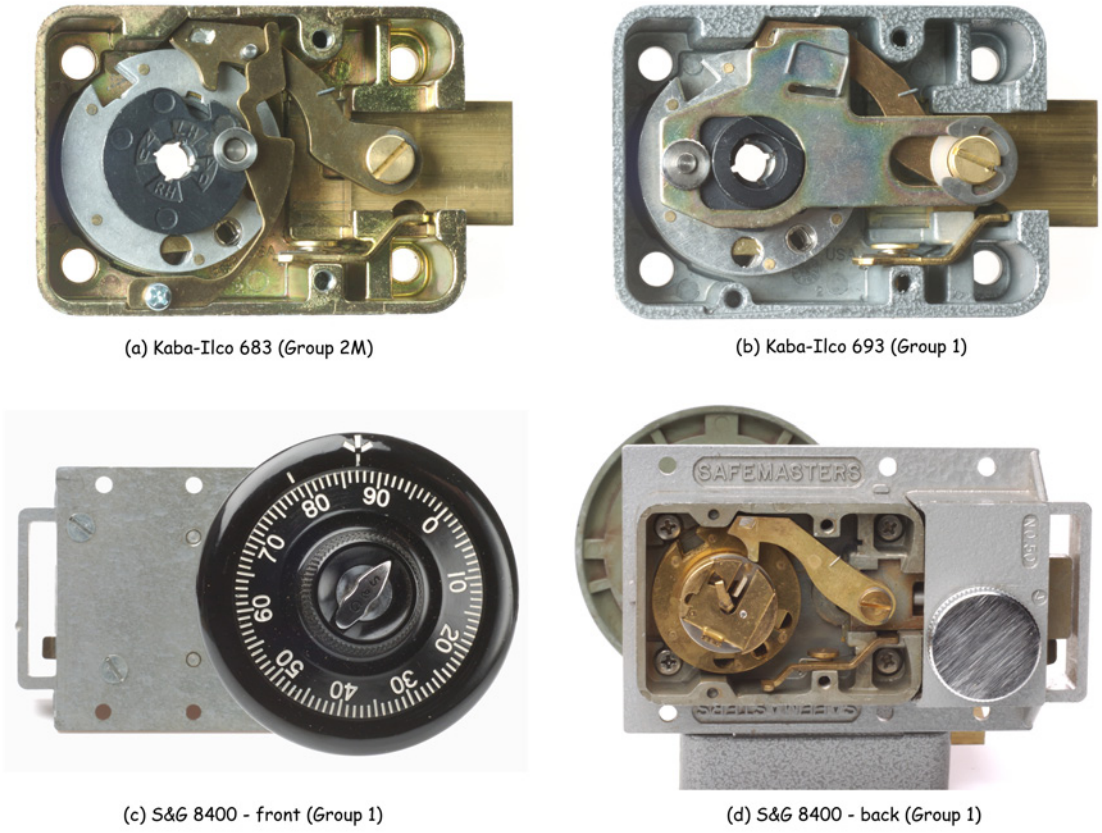

Figure 18: Manipulation-resistant locks. *Top left:* Kaba-Ilco 683 Group 2M lock, with mechanism that adds irregularity to contact points. *Top right:* Kaba-Ilco 693 Group 1 lock, with secondary mechanism that holds the fence off the wheel pack until the nose is already within the contact region. *Bottom left & right:* Sargent & Greenleaf 8400 Group 1 lock, with "butterfly" in dial. A secondary mechanism holds the fence off the wheel pack until the "butterfly lever" in the dial knob is rotated, which also locks the dial into position. This lock is shown in a mount for use on a US DoD "SCIF" vault. The Group 1R version of this lock has plastic Delrin wheels.

The two disciplines are already converging. Many modern safes and vaults, at both the high- and low-security ends of the market, now use electronically controlled locks. This may not represent unmitigated forward progress for security. These locks depend not only on a sound physical design and manufacturing process, but on a sound software architecture and implementation as well. It is not at all clear how to even measure the security of such devices, let alone make them as trustworthy as the mechanical systems they replace.

Mechanical safe lock design is an especially relevant aspect of the study of physical security for the computer scientist. The devices are, after all, analog authentication systems that effectively accept or reject passwords. But there is a deeper reason to study them than this superficial analogy. The mechanisms are remarkable not because their components are especially well made, but rather because their design assumes that they are not. Safe locks are designed not to *eliminate* imperfections, but to *tolerate* them, because it is recognized that the manufacturing processes that produce them cannot be perfect. (They still can fail when, as in the Group 2 locks discussed here, the design turns out not to tolerate the imperfections as well as expected.) Contrast this with contemporary research in software security, which has the Herculanean goal of completely eliminating any bugs that might have security implications. Perhaps we would do better learning instead to design systems that recognize the inevitability of software errors, tolerating them as safe locks tolerate inevitable mechanical imperfections.

Appendix: Acknowledgments and further reading

This survey grew from my own interest in what I have come to refer to as *human-scale security*, and particularly from my attempt to understand how the tools and techniques of the physical security world might strengthen and sharpen the tools and techniques of information security. My thoughts on this are explored in more detail in [Bla03] and [Bla04] and in various documents on my web site, at `www.crypto.com`.

The depth of the relationship between physical security and cryptology was made indelibly clear to me in an extended conversation with Gus Simmons a decade ago. (One of my most valued treasures is a dual-key Abloy cylinder he gave me; it eloquently demonstrates the basic concept of public key encryption). My thinking about the philosophy, principles, and technology of physical security has been further shaped and informed by discussions with many people, but especially David Chaum, Mark Seiden, Marc Tobias and Barry Wels. This paper owes a considerable debt to all of these people.

Many of the opening techniques surveyed here are widely practiced in the safe and vault trade. Readers seeking greater depth than provided here are referred to the literature of that field. An unfortunate consequence of the closed nature (and obsession with secrecy) of physical security is that many of its most powerful tools and techniques are passed along only anecdotally or as folklore. There are notable exceptions, however, and the subject has produced several excellent written resources.

The most comprehensive reference on physical security of which I am aware is Tobias' *Locks, Safes and Security*[Tob00]. Written primarily for the law enforcement and intelligence communities, its emphasis on security metrics and investigation is of particular interest to those seeking to understand the subject from an information security perspective. It should be regarded as required reading for any serious study of human scale security.

There are several worthwhile books aimed at the safe and vault technician community that catalog the practical details of safe penetration, a close reading of which yields insight into how physical security succeeds and fails. The classic text on lock manipulation is Lentz and Kenton's *Art of Manipulation*[LK55]. Although written in 1953 (with a revised edition in 1955), its treatment remains relevant - mechanical safe

locks haven't changed much in half a century. A more recent reference on manipulation is Sieveking's *Guide to Manipulation*[Sie93], which includes special techniques applicable to specific models of locks. Hardplate drilling methods and theories of drill point selection are well covered in Cloud's *Guide to Drilling Safes*[Clo91]. Oehlert's *Safe Technician's Reference Manual*[Oeh97] is a broad and comprehensive study of safe opening techniques, drawing on many specific examples of real safes and vaults.

References

[Bla03] M. Blaze. Cryptology and physical security: Rights amplification in master-keyed mechanical locks. *IEEE Security and Privacy*, 1(2), March/April 2003.

[Bla04] M. Blaze. Toward a broader view of security protocols. In *Proceedings of the 12th Cambridge International Security Protocols Workshop (to appear)*, April 2004.

[Clo91] C. Cloud. *Guide to Drilling Safes*. National Publishing, 1991.

[Cos01] B. Costley. *Sargent and Greenleaf Mechanical Safe Lock Guide*. Sargent and Greenleaf (available from `http:www.sglocks.com/)`, 2001.

[FHL85] R.P. Feynman, E. Hutchings, and R. Leighton. *"Surely You're Joking, Mr. Feynman!"*. W. W. Norton, 1985.

[LK55] C. Lentz and W. Kenton. *The Art of Manipulation*. HPC (reprint edition), 1955.

[Oeh97] M. Oehlert. *Safe Technican's Reference Manual*. AOLA, 1997.

[Sie93] R.G. Sieveking. *Guide to Manipulation*. National Publishing, 1993.

[Tob00] M. W. Tobias. *Locks, Safes and Security*. C. Thomas, Ltd, also `www.security.org`, 2000.