Cryptology and Physical Security: Rights Amplification in Master-Keyed Mechanical Locks

Matt Blaze
AT&T Labs – Research
mab@crypto.com, mab@research.att.com

PREPRINT — 15 Sept 2002 (revised 6 Feb 2003). To appear in *IEEE Security and Privacy*, March/April 2003. This paper can be found at
http://www.crypto.com/papers/mk.pdf

Abstract

This paper examines mechanical lock security from the perspective of computer science and cryptology. We focus on new and practical attacks for amplifying rights in mechanical pin tumbler locks. Given access to a single master-keyed lock and its associated key, a procedure is given that allows discovery and creation of a working master key for the system. No special skill or equipment, beyond a small number of blank keys and a metal file, is required, and the attacker need engage in no suspicious behavior at the lock's location. Countermeasures are also described that may provide limited protection under certain circumstances. We conclude with directions for research in this area and the suggestion that mechanical locks are worthy objects for study and scrutiny.

1 Introduction

In the United States and elsewhere, mechanical locks are the most common mechanisms for access control on doors and security containers. They are found in (and guard the entrances to) the vast majority of residences, commercial businesses, educational institutions, and government facilities, and often serve as the primary protection against intrusion and theft.

As important as locks are in their own right, their design and function has also influenced much of how we think about security generally. Computer security and cryptology borrow much of their language and philosophy from metaphors that invoke mechanical locksmithing. The concept of a "key" as a small secret that allows access or operation, the notion that system security should be designed to depend only on the secrecy of keys, and even the reference to attackers as "intruders," can all be traced back to analogies that long predate computers and modern cryptology.

Conversely, the design of mechanical locks could well be informed by the philosophy and methodology of computer security and cryptology. For example, formal notions of the computational complexity and other resources required to attack a system could be applied to the analysis and design of many aspects of mechanical locks. In general, however, these concepts have not enjoyed widespread adoption by locksmiths or lock designers. Computer security specialists, for their part, are often surprisingly unskeptical in evaluating claims of physical security.

This paper examines the security of the common master-keyed pin-tumbler cylinder lock against an insider threat model more commonly associated with computing systems: unauthorized *rights amplification*. As we shall see, not only is this threat of practical concern in physical security, there are simple attacks that render many real-world lock systems quite vulnerable to it.

2 Background: Mechanical Locks

A complete review of lock technology is well beyond the scope of this paper. For an excellent discussion of physical security design and evaluation, the reader is referred to [3]. For the purposes of consistent terminology, a brief overview follows.

Broadly speaking, mechanical locks fall into two general categories: *combination* locks, which operate upon demonstration of a secret procedure, and *keyed* locks, which operate with use of a secret token. Combination locks are most frequently used to control access to safes and vaults and on some padlocks; most commercial and residential doors and entrances use keyed locks.

There are many different keyed lock designs that have been invented and used throughout the industrial age; among currently manufactured schemes there are *warded* locks, *lever* locks, *wafer* locks, and *dimple* locks. More recently, electronic locks and computer-based access control systems have found application in certain commercial environments. By far the most common medium- and high-security mechanical keyed lock mechanism in the U.S. and many other countries, however, is the mechanical *pin-tumbler* lock cylinder.

2.1 Evaluating Lock Security

Mechanical locks must resist a much wider range of threats than those associated with computing or communications systems.

First, of course, locks function in the physical world and must therefore be sufficiently mechanically strong to withstand forceful attack. Evaluation of this aspect of lock security focuses on such issues as the strength of materials, the accessibility of weak points, resistance to various tools, and so forth. There are industry and government standards that require specific physical characteristics of locks for various applications, which vary depending on the expected resources of the attacker and the likely ease of alternative methods of entry (e.g., through a broken window).

A related issue is the ease with which the locking mechanism itself can be bypassed. It may be possible to open a lock without interacting with the keyed mechanism at all: door latches can often be wedged or pried open, for example. Here, security depends not only on the lock but also the soundness and correctness of its installation.

It is also possible that a lock might be manipulated to operate without a key or that a key can be fabricated without knowledge of its parameters. The most common (or at least famous) manipulation method involves *picking,* which exploits small manufacturing imperfections and mechanical tolerances to set a lock to a keyed state without using a key. A related method, *impressioning,* fabricates a working key directly. Manipulation is generally non-destructive and may leave behind only minimal external evidence. Both picking and impressioning require finesse and skill, however, and are much more difficult to carry out against locks of better quality, especially designs that employ security features intended specifically to thwart manipulation.

Evaluating and protecting against most of the above threats focuses more on the details of a lock's mechanical and physical construction than on abstractly quantifiable security metrics. A computer science and cryptologic security analysis, on the other hand, might take a more abstract, idealized view of locks and their operation. In particular, we might be especially concerned with the security of the key space against various threats.

The most basic design goal of all keyed locks is that a correct key is required for operation; ideally, it should not be possible to operate a lock without possession of the key. (This is rarely achieved in practice due to the factors discussed above, but that is not critical for the purposes of this discussion). Among the most quantifiable security parameters for discussing locks, therefore, is the number of possible unique keys (called the number of *differs* or *changes* in the terminology of the trade), which gives the probability that a randomly cut key will operate a given lock and an upper bound on the resources required to find a

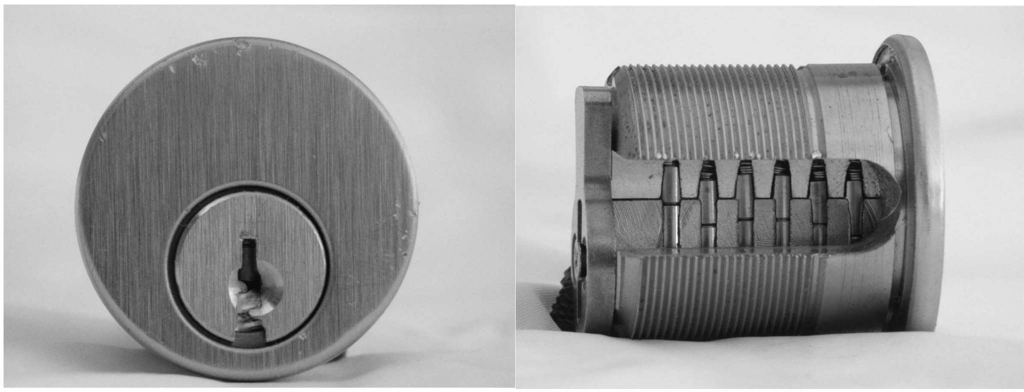

Figure 1: A pin tumbler lock cylinder. *Left:* The cylinder face. Note the *keyway*, which is cut into the *plug*, which in turn sits inside the *shell. Right:* Side view, with part of the shell and plug cut away to expose the six *pin stacks.* Note the border between the plug and shell, which forms the *shear line,* and the *cuts* in each pin stack resting within the plug.

working key by exhaustive search. On typical commercial locks, there are between several thousand and several million possible distinct keys. While these numbers may seem very small by computational security standards, mechanical locks perform on a more human scale. Testing a key against a lock, after all, is an "online" operation requiring seconds, not microseconds, and carries with it at least some risk of discovery if the lock is not one to which the attacker has legitimate access.

If exhaustive search is not feasible, it may still be possible to analyze and exploit a lock's key space in other ways.

2.2 The Pin-Tumbler Lock

The modern pin tumbler lock is quite simple, dating back to ancient Egypt but not commercially mass-produced until the middle of the 19th century. The basic design consists of a rotatable cylinder tube, called a *plug*, that operates the underlying locking mechanism. Around the circumference of the plug is a *shell*, which is fixed to the door or container. Rotation of the plug within the shell operates the locking mechanism. In the locked state the plug is prevented from rotating by a set of movable *pin stacks*, typically under spring pressure, that protrude from holes in the top of the opening in the shell into corresponding holes drilled into the top of the plug. Each pin stack is *cut* in one or more places perpendicular to its length. See Figure 1. (In practice, the cuts are produced by stacking pin segments of particular sizes, not by actually cutting the pins; hence the term "pin stack.")

With no key in the lock, all the pin stack cuts rest within the plug. When a key is inserted into the *keyway slot* at the front of the plug, the pin stacks are raised within the plug and shell. The plug can rotate freely only if the key lifts every pin stack's cut to align at the border between the plug and shell. The plug/shell border is called the *shear line.* See Figure 2. The plug will be blocked from rotating if any pin stack is lifted either not far enough (with the cut still in the plug below the shear line) or too far (with the cut pushed above the shear line and into the shell); to rotate, all pin stacks must have a cut at the shear line. See Figure 3. The height (or *cut depth*) of a key under each pin stack position is called its *bitting*; the bitting of a key is the "secret" needed to open a lock. A key that is bitted to the wrong depth in even one pin position will not allow the lock to operate.

Generally, a lock manufacturer will choose from among only a small number of standard bitting depths at each pin position. This allows keys to be described concisely: typically, the bitting depth number is written starting from the shoulder (handle) of the key to the tip, giving the standard depth number at each

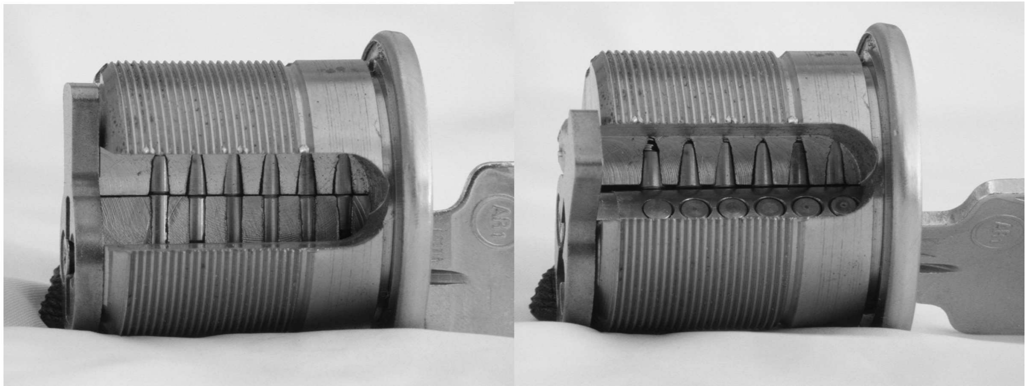

Figure 2: Pin tumbler lock with a correct key inserted. *Left:* The correct key lifts the pin stacks to align the cuts at the shear line. *Right:* With all of the cuts at the shear line, the plug can rotate freely within the shell. Here the plug has been turned slightly toward the camera, so that the tops of the pins in the plug are visible.

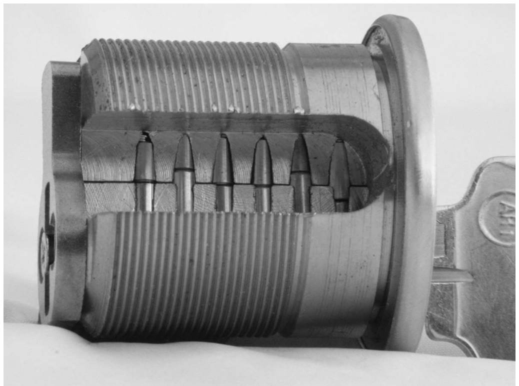

Figure 3: A lock with an incorrect key. Observe that while three of the pin stacks' cuts are at the shear line, two stacks have the cut too high and one stack has the cut too low.

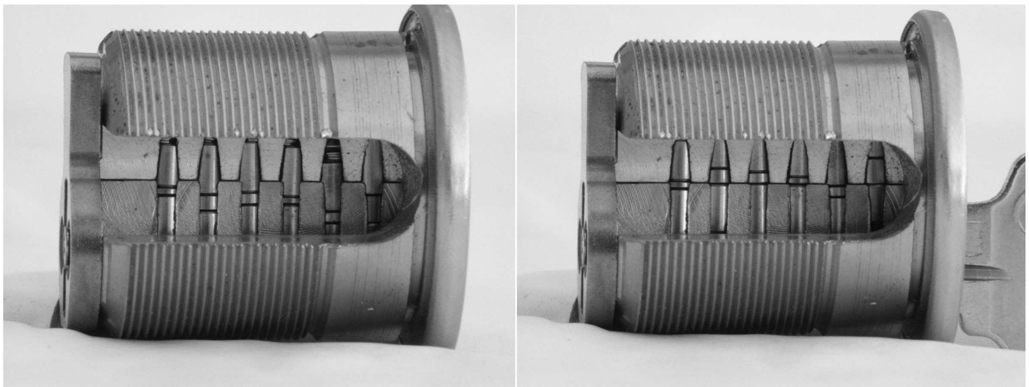

Figure 4: A master keyed pin tumbler lock. *Left:* Each of the six pin stacks has two cuts. *Right:* With the correct *change key* inserted, one of the cuts on each pin stack is aligned at the shear line. Observe that the other cut is sometimes above and sometimes below the shear line.

position. So a key for a five pin lock denoted "12143" would be cut to depth "1" nearest the shoulder, and proceeding toward the tip cut at depths "2," "1," "4" and "3." (The exact specifications of the depths and positions for most commercial locks are widely published in the trade or could be discovered easily by disassembling a sample lock or measuring a small number of cut keys.) Typically, the number of pins is in the range of four to seven, and the number of possible depths ranges from four to ten, depending on the lock model. Better quality locks generally employ more pins and use more distinct cut depths on each.

Pin tumbler locks can often be defeated in various ways, although a discussion of lock picking and other bypass techniques that require specialized skills or tools or that exploit mechanical imperfections is beyond the scope of this paper. In practice, however, even very modest products are often sufficiently secure (or offer the perception of being sufficiently secure) to discourage the more casual would-be intruder from attempting to operate a lock without a key. Probably the most commonly used techniques for unauthorized entry, aside from brute force, involve procuring a working key.

2.3 Master Keying

Complicating the analysis of pin tumbler lock security is the fact that, especially in larger-scale installations, there may be more than one key bitting that operates any given lock. The most common reason for this phenomenon is the practice of *master keying,* in which each lock in a group is intended to be operated not only by its own unique key (the *change key* in trade parlance) but also by "master" keys that can also operate some or all other locks in the system.

Master keying in pin tumbler locks can be accomplished in several ways, with the earliest systems dating back over 100 years. The conceptually simplest master key method entails two cylinders on each lock, one keyed individually and the other keyed to the master bitting; a mechanical linkage operates the lock when either cylinder is turned. Other master keying schemes employ an independently keyed *master ring* around the lock core, and still others depend on only a subset of pin positions being used in any given lock. All of these approaches have well-known advantages and disadvantages and are not considered in this paper. Most importantly, these schemes require the use of special locks designed specifically for master keying.

The most common master keying scheme – the subject of consideration of this paper – can be used with virtually any pin tumbler lock. Recall that in a pin tumbler lock, each pin stack is cut in one place, defining exactly one depth to which the stack must be lifted by the key bitting to align with the shear line. In the conventional *split pin* mastering scheme, some or all pin stacks are cut in more than one place (typically in two places), allowing additional bittings that align such pins. See Figure 4.

Consider for example, a lock *A,* which has five pin stacks with four possible cut positions in each. Suppose pin stacks 1 through 5 are each cut in two places, corresponding to bittings “1” and “4”. Observe that this lock can be opened by at least two keys, one with bitting 11111 and another with bitting 44444. We could create a second lock *B*, this time with pin stacks 1 through 5 each cut at depth “2” and depth “4”. This lock can be operated by keys cut 22222 and 44444. If these are the only two locks in the system, keys 11111 and 22222 can be said to be the change keys for locks *A* and *B*, respectively, while key 44444 is a master key that operates both.

There are a number of different schemes for master keying; the subject is surprisingly subtle and complex, and the trade has developed standardized practices in recent years. For in-depth treatments, the reader is referred to [1] and [2].

For the purposes of our discussion, it is sufficient to note that modern master systems fall into two broad categories: *Total Position Progression (TPP)* and *Rotating Constant (RC)*. In TPP schemes, every pin stack has a single separate master cut, which is never used in that position on any change keys. In RC schemes, change keys do share the master bitting for a fixed number of pin stack positions, although the positions will vary (rotate) from lock to lock. Both these schemes can implement a directed graph with several levels of master keys: “sub-master” keys that open a subset of locks in the system and “grand master” keys that open more[1]. The highest-level master key, which opens all locks in a multi-level system, is sometimes called the *Top Master Key (TMK)*.

The astute reader will note that master keying reduces security in several important ways. First, of course, the master key represents a very valuable target; compromise of the master key compromises the entire system. Even if the master keys are well protected, security is still somewhat degraded. Because each mastered pin stack aligns with the shear line in several positions, mastered systems are more susceptible to unintentional *cross keying,* in which keys from the same or other systems will operate more locks than intended. For the same reason, mastered locks tend to be more vulnerable to manipulation by picking and impressioning. These weaknesses can be mitigated to some extent through careful planning, improved mechanical construction, and the use of additional pin stacks and possible cut depths.

In this paper, however, we consider methods for discovering the master key bitting in conventional pin-tumbler systems given access to a single change key and its associated lock. No special skills or tools are required on the part of the attacker, nor is it necessary to disassemble any lock or engage in any inherently conspicuous or suspicious activity. We also suggest countermeasures and alternative lock designs that can frustrate these attacks to at least some extent under certain circumstances.

3 Rights Amplification: Reverse-Engineering Master Keys

Clearly, the most valuable, sensitive secret in any lock system is the bitting of the top-level master key (TMK). Insiders, who possess legitimate change keys and have physical access to locks, represent perhaps the most serious potential threat against master keyed systems. The primary purpose of assigning locks unique change key bittings, after all, is to allow operating privileges to be granted to only specific locks; if a change key can be converted into a master key, a major security objective of the system is compromised. In the terminology of computer security, master key systems should resist unauthorized *rights amplification* (also called *privilege escalation*). Unfortunately, most deployed master key systems are quite vulnerable in this regard.

[1]There are also *Selective Key* systems, in which any lock can be keyed to operate with an arbitrary subset of keys, using techniques similar to master keying, and *Maison Key* schemes, in which certain locks are keyed to all keys in a group. We do not consider such systems here.

3.1 Background

Several time-honored methods convert change keys into master keys, with different techniques applicable depending on the particular system and resources available to the attacker.

The simplest approach to master key discovery involves direct decoding of an original master key, e.g., from visual inspection, photographs, photocopies, or measurement. A trained observer may be able to recall the cut depths with surprising accuracy after being allowed to look only briefly at a key.

Another direct technique involves disassembly of a master keyed lock and measurement of the pins in each pin stack to determine the bittings that will operate each pin position. Without access to the lock's change key, this does not yield complete information about the master bitting; there will be exponentially many potential master key bittings, only one of which will correspond to the true master key. If every pin is mastered according with a standard TPP scheme, disassembly of a single lock will reveal 2^P potential master keys, where P is the number of pin stacks. (This exponent is still small enough to make exhaustive search of these keys feasible in many cases). Disassembly of additional locks from the same system can narrow this search space significantly. If the change key to a disassembled lock is available, the cuts corresponding to its bitting can be eliminated from each pin stack, making the correct bitting of the true master unambiguously clear from a single sample. (More secure lock designs make it difficult to non-destructively remove a lock without the key, e.g., by placing set screws in locations that are inaccessible when a door is closed and locked). Padlocks are especially vulnerable to these sorts of attacks, since they can be stolen easily when they are left unlocked.

A sufficiently large group of change key holders in TPP-based systems may be able to reverse engineer a master key without disassembling any locks. Recall that in these systems change keys never have the same bitting at a given pin position as the master. By measuring their change keys, a conspiracy of key holders may discover a single depth not used at each pin position on the change keys; this will correspond to the master bitting. Several correspondents have noted that this technique is occasionally employed by enterprising university students, especially at better engineering schools.

None of these approaches is completely satisfactory from the point of view of the attacker, however. Direct decoding from the true master key entails limited access to such a key and is not possible if no master key is available for measurement. Lock disassembly for pin measurement may expose the attacker to suspicion and could be difficult to perform in secret (and carries the risk that the lock may be damaged in reassembly). Comparing a large number of different keys requires, in the first case, a large number of different keys, which may not be available, and is ineffective against RC-based systems.

A more powerful attack requires only one change key and is effective against all standard TPP- and RC-based systems.

3.2 An Adaptive Oracle-Based Rights Amplification Attack

It is useful now to consider a lock in more abstract terms. From a cryptologic point of view, we might observe that a lock is really an online "oracle" that accepts or rejects keys presented to it. In this sense, the oracle gives a single bit answer for each key presented to it; the lock either turns or it does not.

A natural question to ask about any online oracle is whether it is feasible to issue a small number of queries that force the oracle to leak its secrets. In particular, can we exploit the oracle to test efficiently single "bits" of a possible key or must we exhaustively search the entire key space?

Recall that a pin tumbler lock will operate when each of its pin stacks is raised (by a key) to a position where one of its cuts is aligned at the shear line. There is no "communication" among pins; the lock will operate not only with all pin stacks aligned at the change key depth or all pin stacks at the master key depth, but also by keys that align some stacks at the change depth and others at the master depth. That is, consider our five pin lock *A* from the previous section, with key bitting 11111 representing *A*'s change key and 44444 representing the system's master key. This lock can be operated not only by the obvious keys cut 11111 and 44444, but by a total of 2^5 different keys, including, e.g., 11114, 11141, etc.

It is straightforward to exploit this phenomenon to discover the master key bitting given access to a single change key and its associated lock, plus a small number of blank keys milled for the system keyway.

In our new[2] attack, we use the operation or non-operation of a lock as an "oracle" to determine, pin by pin, the complete bitting of the TMK.

3.2.1 Notation

Let P denote the number of pin stacks in a lock, with stack 1 representing the first stack (e.g., the one closest to the shoulder of the key) and stack P representing the last (e.g., the stack at the tip of the key).

Let D denote the number of distinct key bitting depths in a pin stack, where 1 is the highest bitting (in which the pin stack is raised the most) and D is the lowest (in which the pin stack is raised the least).

Assuming that the physical properties of the system place no restrictions on the bitting depth of adjacent pin positions, observe that the number of distinct keys is D^P.

3.2.2 The Attack

For each pin position, p from 1 to P, prepare $D-1$ test keys cut with the change key bitting at every position except position p. At position p, cut each of the $D-1$ keys with each possible bitting depth *excluding* the bitting of the change key at that position. Attempt to operate the lock ("query the oracle") with each of these test keys, and record which keys operate the lock.

In a TPP-based system with every pin mastered, exactly one of the $D-1$ test keys for each pin position will operate the lock; the depth of the test key at that position represents the master bitting at that position. If none of the test keys for a particular position operates the lock, then either that pin is not mastered or it is an RC-based system. In either of these cases, the master key bitting at that position is the same as that of the original change key.

Once the master bitting has been determined at each of the P positions, a complete top-level master key can be cut easily.

Observe that our attack consumes $P(D-1)$ key blanks and requires $P(D-1)$ probes of the lock, in the worst case. If it is possible for the attacker to cut keys between probes of the lock, however, a simple optimization reduces the number of blanks consumed to P in the worst case. Rather than cutting $D-1$ separate blanks per position, the attacker need use a single key, initially cutting the position under test to the highest depth and re-cutting the same blank successively lower after probing the lock. This reduces the total cost of carrying out the attack to less than about two US dollars in the worst case. This optimized attack still requires $P(D-1)$ probes of the lock in the worse case, of course.

3.2.3 Practical Considerations

In some lock designs, not all of the D^P possible keys are "legal". In particular, with some lock models it is not possible on a standard key to have a very high cut immediately adjacent to a very low cut if the angle at which the bittings are cut reaches across to the next pin position. A lock's *Maximum Adjacent Cut Specification* (*MACS*) might require, for example, in a system with 7 different cut depths that adjacent cuts

[2] It is always difficult to be sure that something is novel in the sense of not having previously been discovered independently; the lack of a coherent and open body of literature on locks makes it especially so here. Our attack surely is not new in this sense. Several correspondents have suggested that similar approaches to master key reverse engineering have been discovered and used illicitly in the past and the method occasionally circulated informally, e.g., on Internet message boards. (We subsequently found a message originally sent to a private mailing list in 1987 from Doug Gwyn that describes a similar method.) However, there do not appear to be references to this particular attack in the published literature of either the locksmith or underground communities.

be no more than 4 steps apart, disallowing a depth "1" cut next to a depth "7" cut. Even if both the change key and the master key do not violate the MACS rule for a particular lock, this attack employs test keys that mix change key cuts with potential master cuts. If the original change key has very high or very low cuts, it may therefore be necessary for the attacker to create some test keys that do violate MACS. In practice, on the locks we examined with MACS restrictions, it is generally still possible to cut working test keys by using a steeper than usual angle and with cuts occupying slightly narrower than usual space on the key. Although insertion and removal of such keys is more difficult, they are sufficient for this limited (single-use) purpose. Alternatively, previously discovered master depths could be used in adjacent positions on subsequent test keys.

Also complicating our attack is the possibility that the master cuts lie somewhere between the "standard" depths ordinarily used by the lock manufacturer. This is more likely in older systems or those keyed by private locksmiths who may not not follow manufacturer-standardized practices. When this is suspected to be the case, the attacker must probe the lock at more test cut depths, removing only a small amount of key material (.005 inches or so) from the position under test between probes. (This is similar to the procedure used when creating a key by the "impressioning" technique and could be performed with a fine metal file.)

Some systems, especially in older installations, use master cuts that are consistently higher or lower than the change key cuts. This practice makes it especially easy to discover the master key with this attack.

Multi-level master systems may or may not present a special challenge. In standard TPP and RC systems, every pin stack has at most two cuts; "submasters" are implemented by using fixed change key bitting on certain pins for locks within each submaster group. In such cases, the attack proceeds as described and yields the TMK. It is also possible, however, to implement hierarchical submastering by using more than two cuts on each pin stack. In such cases the TMK bitting of a given pin may be ambiguous. An attacker can distinguish the true TMK cuts in such systems by conducting the attack on locks from different submaster groups. This may not always be necessary, however. It is common for such systems to employ the convention that all of the TMK cuts are either above or below the submaster cuts.

Some larger installations put different groups of locks on distinct keyways, such that a change key for a lock in one group does not fit into the keyway of locks from others. The TMK is cut on a special "master" blank that fits all the keyways in the system. This practice, called *Sectional Mastering*, expands the number of effective differs in the system and reduces cross keying between different lock groups. Sectionally mastered systems are especially attractive targets for attack, since the TMK works for a very large number of locks across groups that would otherwise have to be keyed on different master systems. The attacker simply cuts the TMK bitting (derived from a lock in any section) onto a blank milled for the master section.

It is worth noting that even "high security" pin tumbler lock designs, including those that use sidebar cuts and rotating pins, are usually in principle vulnerable to this attack; the only question is whether the attacker can obtain or fabricate the required blanks. Furthermore, our attack can be generalized to many other lock schemes, including, for example, certain high security lever lock and disk wafer designs (such as Abloy).

3.3 Experimental Results

It is easy to see that this attack is effective against the standard master keying schemes we described. It is natural to ask, then, whether master key systems deployed in practice follow these schemes and are therefore vulnerable. Unlike computing systems that can be tested relatively easily and safely in isolated testbed environments running standard software, such a question can only be answered by attempting the attack against real installations. The reader is cautioned that reproduction of these experiments should be carried out only with the cooperation of the owner of the lock systems on which the attack is attempted.

We tested our attack against a variety of medium- and large- scale institutional master keyed installations, including both educational and commercial environments. Systems tested were both relatively new and relatively old, had been both factory-keyed as well as privately rekeyed, and included locks manufactured by Arrow (SFIC), Best (SFIC), Corbin-Russwin, Schlage, and Yale. For the Best SFIC, Arrow SFIC and Schlage systems, we used portable key punches and a supply of blank keys brought to the facilities tested. For the Corbin-Russwin and Yale systems, we pre-cut six test keys on a general purpose code machine (based on measurements previously taken from a change key) and used a metal file at the test site to progressively cut the test keys and finally to cut the full master bitting onto a fresh blank key.

All required key blanks were procured from standard commercial sources (which can be found easily on the Internet with a search engine). Cost per blank ranged from US$0.14 to US$0.35 depending on the particular lock type, plus shipping. We used, for convenience in some of the attacks, key cutting machines, also available widely from commercial sources for a few hundred dollars. In other cases, we used a fine metal file and a dial caliper or micrometer to cut the keys to the correct bitting depth. None of the equipment or supplies we used are restricted in any way. (Such restrictions, even if they existed, would not be especially effective at preventing potential attackers from obtaining blank keys, given the vast number of small businesses that have legitimate need for them (hardware stores, etc.)).

In every case, the attack yielded the top master key bitting, as expected. In general, it required only a few minutes to carry out, even when using a file to cut the keys.

All six Arrow SFIC and Best SFIC systems we tested had all (six or seven) pin stacks mastered with a TPP format. The two Corbin-Russwin (system 70) systems each had three pin stacks (out of six) mastered, again with a TPP format. The Schlage system used an RC-based scheme, with every pin mastered and two master cuts used on each change key. The Yale system was also RC-based, with one master cut used on each change key. Several of the systems had multi-level mastering hierarchies; the attack yielded the TMK in all cases.

Notably, although some of the complications discussed in the previous section (such as more than one master cut per pin stack, selective keying, or non-standard master depths) are possible in principle, we did not encounter them. Every system we tested was keyed according to standard (TPP or RC) industry practice, had at most one master cut per pin and employed standard depths, making the attacker's job especially straightforward. Although our experiments hardly constitute an exhaustive survey, they were conducted across a wide variety of facilities that seem reasonably representative of a large segment of US institutional lock installations. A check of several other lock vendors' standard master keying practices further supports this conclusion.

4 Countermeasures

Our adaptive oracle attack is only effective against locks that have a single shear line used by both master and change keys. Although this is the case with the majority of mastered locks, there are commercially available designs that do not have this property. Locks with a separate master ring, for example, require that all pin stacks be aligned to the same one of two distinct master or change shear lines, and therefore do not provide feedback about the master bitting of a pin given the change bittings of the other pins[3] . (Master ring locks, however, are actually *more* vulnerable to reverse engineering from lock disassembly by an attacker without access to the change key).

This attack assumes that the attacker has access to a modest supply of blank keys for the system. Whether this is a practical assumption depends on the particular system, of course, and some "restricted keyway" lock products may make it more difficult for the attacker to obtain blanks from commercial sources. However,

[3]A master ring lock has two concentric plugs, with the keyway cut into the inner plug. Two distinct shear lines are formed. The pin stacks are correspondingly taller, with one cut on each stack designed to be able to reach one shear line and another cut designed to reach the other. A few master ring locks are still commercially manufactured, but the design has largely fallen out of favor for most applications.

blanks for many so-called restricted systems are in fact readily available from aftermarket vendors. Even when an exact blank is not commercially available, often a different key can be milled down to fit. Unusual or patent-protected key designs, such as those employing a sidebar cut, may be more difficult to procure directly or modify from commercial sources, but blanks can still usually be fabricated in small quantities relatively easily by casting (especially since the attacker already possesses a working change key cut on the correct blank).

In medium-scale master systems, it may be possible to limit the information contained in any given lock, at the expense of somewhat increased vulnerability to cross keying and picking. In standard master schemes, each pin stack is cut only at the master and change depths. The attacker exploits the fact that any working depths not corresponding to the change key must be on the master. A natural way to frustrate the attack, therefore, is to add "false" cuts to some pin stacks that do not correspond to the master and that do not appear in the majority of other locks in the system. If one "extra" cut is added to each pin stack, the attacker will learn 2^P different possible master keys from one lock, only one of which will correspond to the "true" TMK bitting. These extra cuts must be selected very carefully, however, since each such cut reduces the number of unique differs available in the system. Effectively, the extra cuts create new subclasses of sub-master keys among locks that share the same false cuts, which the attacker must eliminate before learning the true high-level master key.

5 Conclusions and Lessons Learned

In this paper, we have shown a very simple rights amplification attack that is effective against virtually all conventional master-keyed pin tumbler locks, including many so-called "high-security" products. This attack is an especially serious threat to the security of such systems because it is easy to carry out, leaves no forensic evidence, requires no special skills and uses only very limited resources (a few blank keys and a file, in the case of the most frugal attacker). Compounding the threat are the facts that the attacker need engage only in apparently ordinary behavior – operating the lock to which he or she already has legitimate access – and that the attack can be carried out over a period of time in several (interrupted) sessions.

Any successful compromise of a master keyed installation can be very difficult and costly to remedy (assuming it is even discovered). Every mastered lock must be rekeyed and, depending how the keying is done, new keys distributed to the key holders. Not only is this very expensive, but system-wide re-keying can also require a considerable period of time to complete, during which all the old locks remain exposed. In light of the inherent security vulnerabilities introduced by master keying, owners of lock systems should consider carefully whether the security risks of mastering outweigh its convenience benefits. (Unfortunately, the computing world is not alone in often putting a premium on convenience over security.)

If master keying must be used, simple countermeasures, especially the use of false cuts in mastered pin stacks, can frustrate the adaptive oracle attack. If it is not possible to employ lock designs, such as master rings, that resist such attacks, these countermeasures should be considered seriously.

It is worth noting that these attacks become rather obvious when the basic analysis techniques of cryptology and computer security are employed. (In fact, as noted previously, these attacks appear to have been discovered and rediscoverd independently several times, occasionally passed on as underground engineering and locksmithing folklore but never documented in the literature). One of the first questions asked about any proposed cryptosystem, for example, is whether it is possible to test the value of one key bit independently from the others. If it is, the system would be considered hopelessly insecure, since an attack would take time only linear in the number of key bits, instead of exponential. The same question readily translates into the mechanical lock domain by substituting "pin stack" for "key bit." (In fact, our master key discovery scheme bears a striking resemblance to a famous character-by-character attack against the Tenex password mechanism.) Similarly, the notion of an online service as an authentication oracle is familiar in the analysis of cryptographic systems. Mechanical locks can likewise be modeled as online oracles that accept or reject keys, and security analysis conducted accordingly. Finally, the attack against TPP systems that compares many different change keys is reminiscent of “related key” attacks against cryptosystems, with a threat model much like “traitor tracing” in broadcast encryption. Perhaps other aspects of the analysis of mechanical and physical security would benefit from similar analogies to computing systems and cryptology.

On the other side of the coin, the vulnerability to rights amplification in master keying of mechanical locks recalls similar weaknesses in cryptographic systems that attempt analogous capabilities. Consider, for example, the vulnerabilities inherent in “key escrow” systems that attempt to facilitate emergency decryption by a central third party of data encrypted with many different users’ keys. Even more direct analogies can be found in digital rights management schemes and smartcard-based digital cash systems that contain but aim to hide, as master keyed locks do, global secrets from their users.

6 Acknowledgments

The author is grateful to David Chaum, Niels Ferguson, A.J. Hoffman, Dave Korman, Avi Rubin, Mark Seiden, Lloyd Seliber, Adi Shamir, Jonathan Smith, Marc W. Tobias and Barry Wels for comments on this paper and interesting conversations about locks generally. John Ioannidis made the cutaway lock shown in the figures. We are also indebted to managers at several master-keyed installations who allowed us to conduct our experiments, but who, in the interests of protecting the security of their facilities, cannot be thanked publicly.

References

[1] J. Andrews. *Fundamentals of Master Keying.* Associated Locksmiths of America. 1990.

[2] B.B. Edwards Jr. *Master Keying by the Numbers (2/e)*. Security Resources. Pensacola, FL, USA. 1997.

[3] M. W. Tobias. *Locks, Safes and Security (2/e).* Charles Thomas Publisher, Ltd. Springfield, IL, USA. 2000.