Authentication

From Passwords to Public Keys



"On the Internet, no one knows you're a dog."

Richard E. Smith



Authentication

From Passwords to Public Keys

Richard E. Smith



Addison-Wesley
Boston • San Francisco • New York • Toronto • Montreal
London • Munich • Paris • Madrid
Capetown • Sydney • Tokyo • Singapore • Mexico City

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and Addison-Wesley, Inc. was aware of a trademark claim, the designations have been printed with initial capital letters or in all capitals.

The author and publisher have taken care in the preparation of this book, but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information or programs contained herein.

The publisher offers discounts on this book when ordered in quantity for special sales. For more information, please contact:

Pearson Education Corporate Sales Division
One Lake Street
Upper Saddle River, NJ 07458
(800) 382-3419
corpsales@pearsontechgroup.com

Visit AW on the Web: www.informit.com/aw

Library of Congress Cataloging-in-Publication Data

Smith, Richard E., 1952–
Authentication: from passwords to public keys / Richard E. Smith.
p. cm.
Includes bibliographical references and index.
ISBN 0-201-61599-1 (pbk.).
1. Computer security. 2. Authentication. I. Title.
QA76.9.A25 S65 2001
005.8–dc21

2001046001

Copyright © 2002 by Addison-Wesley

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form, or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior consent of the publisher. Printed in the United States of America. Published simultaneously in Canada.

ISBN 0-201-61599-1
This product is printed digitally on demand.
1 2 3 4 5 6 7 8 9 10—CRS—0504030201
First printing, October 2001

CHAPTER 1

THE AUTHENTICATION LANDSCAPE

Open, Sesame!

— Scheherazade (attr.), *Ali Baba and the Forty Thieves*

IN THIS CHAPTER

This chapter provides an overview of *authentication*, the problem of verifying identities, and the major issues in making it work.

- Elements of authentication systems
- Early developments in password authentication
- Attacks via cleverness, theft, and trickery
- Authentication factors: passwords, tokens, biometrics
- Judging attack prevalence
- Summary of the chapter's attacks and defenses

1.1 A VERY OLD STORY

For centuries, people have relied on guards, spoken passwords, and hard-to-forge seals to prove their identity to other people and to verify important messages. Unattended authentication by mechanical devices is also quite old: key-operated locks date back to the ancient Egyptians. Practical mechanisms did their job with as little human interaction as possible.

Unattended authentication is essential with today's computer-based systems. It may be a cliché to call the Internet the "information superhighway," but here it captures truth: we can't afford to post a policeman at every cloverleaf or on the countless

interconnected streets and driveways. We must depend on mechanized protection.

The notion of an unattended, password-controlled lock appeared centuries ago in *Ali Baba and the Forty Thieves*, the Mideastern folk tale. In the well-known story, the narrator Scheherazade told of a great treasure hidden in a cave behind a stone. The password "Open, Sesame" caused the stone to move out of the way. Guards in cities of that era also used passwords to allow citizens through the city gates. But the thieves' cave didn't need a human guard to recognize faces, voices, or styles of dress. Instead, there was an unexplained and probably magical device that mechanically responded to the spoken words. Most importantly, the mechanism didn't discriminate between different people speaking the words. It responded to the words themselves, just like modern combination locks or password-protected workstations, which admit anyone knowing the secret. *see Note 1.*

The point of Scheherazade's tale is that magic (or mechanism) always follows its own logic, independent of people's wishes or intentions. The same thing plagues us today with computer-based authentication systems. We have a wealth of technical alternatives, each following its own logic and falling to its own distinctive weaknesses. But if we understand the logic and the weaknesses of a given method, we stand a better chance of bending technology to fulfill our real needs.

The passwords and other authentication mechanisms used with computers today cover a broad range of techniques and technologies. Web site designers, e-commerce planners, and other system developers must choose from numerous products and make numerous configuration decisions within each product. Systems like Windows NT and Windows 2000 by themselves incorporate several password alternatives to provide interoperability with other products. Some organizations need the extra security of smart cards or authentication tokens like Safeword or SecurID. A major motivation behind the "public key infrastructure" (or PKI) is to someday revolutionize, strengthen, and simplify individual authentication. But, as with any evolving technology, it's hard to predict how much of its promise it will ultimately achieve.

TABLE 1.1: *Examples of the Five Elements in an Authentication System*

Authentication Element	Cave of the 40 Thieves	Password Login	Teller Machine	Web Server to Client
Person, principal, entity	Anyone who knew the password	Authorized user	Owner of a bank account	Web site owner
Distinguishing characteristic, token, authenticator	The password "Open, Sesame"	Secret password	ATM card and PIN	Public key within a certificate
Proprietor, system owner, administrator	The forty thieves	Enterprise owning the system	Bank	Certificate authority
Authentication mechanism	Magical device that responds to the words	Password validation software	Card validation software	Certificate validation software
Access control mechanism	Mechanism to roll the stone from in front of the cave	Login process, access controls	Allows banking transactions	Browser marks the page "secure"

1.2 ELEMENTS OF AN AUTHENTICATION SYSTEM

Regardless of whether an authentication system is computer based or not, there are several elements usually present, and certain things usually take place. First of all, we have a particular *person* or group of people to be authenticated. Next, we need a *distinguishing characteristic* that differentiates that particular person or group from others. Third, there is a *proprietor* who is responsible for the system being used and relies on mechanized authentication to distinguish authorized users from other people. Fourth, we need an *authentication mechanism* to verify the presence of the distinguishing characteristic. Fifth, we grant some privilege when the authentication succeeds by using an *access control mechanism*, and the same mechanism denies the privilege if authentication fails. Table 1.1 gives examples of these elements.

For example, the person of interest to the thieves' cave might be Ali Baba, his brother, or whichever thief wanted the door to open. The distinguishing characteristic was knowledge of the password, "Open, Sesame." The cave's proprietors were obviously the gang of the Forty Thieves. There was some unexplained authentication

CHAPTER 7

BIOMETRICS

By the pricking of my thumbs
Something wicked this way comes
Open locks,
Whoever knocks.

— William Shakespeare, *Macbeth, Act IV, Scene I*

IN THIS CHAPTER

Biometrics use unique personal traits to authenticate people. There are a variety of biometric techniques, each measuring a different personal trait. Biometrics can provide incredibly convenient authentication, but can also suffer from a variety of shortcomings.

- Biometric uses and basic limitations
- Summary of common biometric techniques
- How biometric authentication works
- Measuring biometric accuracy
- Improving biometric security with cryptographic techniques

7.1 BIOMETRICS: SOMETHING YOU ARE

Biometrics use personal features instead of facts to authenticate a person. When we step away from computer technology, we see that biometric authentication has been used for centuries in one form or another. Passports almost always carry a physical description of the bearer, and have carried photographs as well for the past several decades. Photo ID cards have been used at least since the 1930s; during World War II they were used in various high-security military projects. Although there was no automation involved, these passports and cards served as a portable credential. People could inspect the credential and satisfy themselves regarding the identity of its owner.

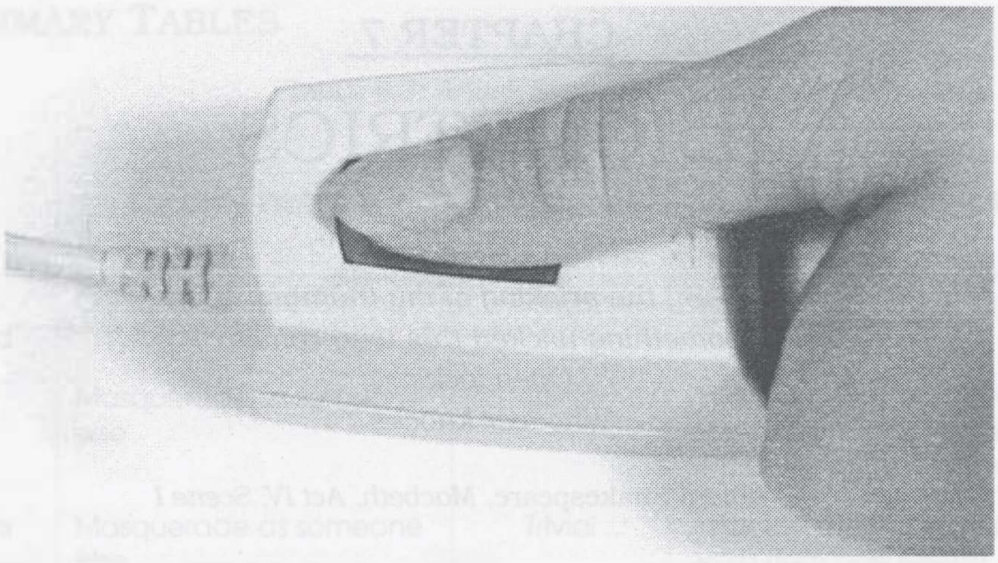


FIGURE 7.1: *Biometric authentication based on fingerprints.* The user has placed her thumb on a device that reads the thumbprint. To authenticate her, the system compares this reading with a previously collected thumbprint. (Scanner courtesy of Informer Systems, Ltd.)

Of course, it's not practical for computer-based authentication to rely on people to manually examine and verify biometric credentials. Instead, computers use special devices to perform the biometric reading (Figure 7.1) and use a computational procedure to verify the user's identity. The procedure compares a biometric reading against a corresponding reading stored in that person's user record. If the two readings match closely, then the system authenticates the person. If the two do not match very well, then the system refuses to recognize the person.

Unlike passwords and other base secrets, the biometric data rarely matches perfectly. Instead, the authentication mechanism measures how close the latest reading matches the reading in the user record. Readings from the right person are supposed to match closely while readings from other people should not match closely.

PROMISE AND REALITY

Many people and organizations are enthusiastic about the promise of biometrics for authentication. This attitude is summed up by a popular slogan among biometrics proponents: "The body is the password." With biometrics, people won't be locked out by forgotten or automatically changed passwords. Biometrics also inspire enthusiasm because they seem to represent the ultimate in personal authentication. Since biometrics are based on a person's unique traits, some observers believe the techniques offer the highest possible confidence in a person's identity. *see Note 1.*

But appearances can deceive. First, there is the problem of matching biometric readings. The process generally takes a picture of the person's trait and then tries to match it. The match may fail to authenticate a legitimate user if the biometric has changed or been injured, if it is presented in an unusual fashion, or if dirt or noise works itself into the system. In other words, the user can suffer a *false rejection* by the authentication system. If we adjust the system to be more tolerant of noise and poor matches, then we increase the risk of authenticating a user based on a biometric trait from the wrong person. In other words, we risk *false acceptance* of a masquerade, the same problem we have when an attacker guesses the right password.

Another problem is the threat of replay attacks. An attacker can collect a victim's biometric reading, either through an external recorder or by copying the reading's binary data representation. Too often, an attacker can fool the biometric system by presenting the copy of the biometric reading. Some biometric systems are sophisticated enough to distinguish between a live measurement and a fake (that is, the difference between a fingerprint on a real finger and a reproduction on paper, rubber, or plastic). Sophisticated sensors are always more expensive. Moreover, attackers may generate a plausible binary version of a biometric reading and use such data for replay attacks on networks. There is no way to combat such attacks on networks except with additional authentication measures that typically require a base secret (see Section 7.7).

Biometrics are safest and easiest to use with the local deployment pattern described in Chapters 4 and 5. The local pattern places the entire authentication mechanism inside a physically secure perime-

ter. This eliminates the risk of an attacker's capturing and replaying the digital representation of a legitimate user's biometric reading. This does not prevent forgeries or recordings of the original biometric reading; the biometric reading device must be designed to detect replayed readings. Biometrics can also be safely used in conjunction with random base secrets as described in Section 7.7. However, the costs of a biometric system can outweigh the benefits if the system still requires us to manage base secrets.

USES OF BIOMETRICS

Since biometrics are based on unique personal characteristics, we can use them for three different but related applications:

- **Authentication**

  This is the problem addressed by most other mechanisms in this book: given a user name, can we confirm that it belongs to the person who presents it? The biometric reading serves the role of the base secret or verifier, and it must closely match the entry in the corresponding user record.

- **Identification**

  This is the classic law-enforcement use of fingerprints: given a biometric sample, can we associate it with a unique human being or, at worst, a small number of "suspects"? This application requires a comprehensive database of biometric samples that is likely to contain the sample being searched for. The most effective identification systems are those with the largest database of biometric samples. The U.S. Federal Bureau of Investigation (FBI) has systematically collected fingerprints since the early 20th century, and this collection is now the centerpiece of their *automated fingerprint identification system* (AFIS).

- **Uniqueness**

  This variant of the identification application has become practical with the advent of low-cost computer-based biometrics: given a biometric sample, can we determine if its owner is already in the database? The technique is used by government-sponsored benefits organizations to ensure that recipients don't register more than once for welfare checks or other valuable services.

The distinction between identity and uniqueness is subtle but important. Both applications use the same technology but they strive for different goals. Identity applications strive for completeness: the ideal database contains every living person, or at least everyone within a particular population. Uniqueness applications involve people associated with a particular institution or enterprise and use biometrics to ensure that each person is registered no more than once. Identification systems have a compulsion for completeness that uniqueness systems usually lack. An important technical difference between the two applications is in their scale: biometric pattern matching gets much harder as the size of the pattern collection grows. Identification systems may grow arbitrarily large. *see Note 2.*

In 1990, Los Angeles County deployed a pilot project that authenticated welfare recipients with fingerprints and verified uniqueness. The three districts using the new system saw an immediate decrease in welfare applicants while surrounding districts continued to see an increase. This was considered a promising result, so the project was rapidly expanded to other districts, eventually covering 32 district offices and over 700,000 clients. By the end of 1998, changes in welfare enrollment patterns led the county to claim that the $42.4 million system had yielded a $87.8 million savings. Other governments, including the state of Connecticut in 1996, have deployed similar systems with similar results. *see Note 3.*

These applications underscore another issue with biometrics: privacy. Since each biometric measurement is supposed to be unique to each person, it becomes possible to track individuals by following their biometric readings. For example, anyone who happened to have visited the scene of a crime might be called in for questioning simply because his or her fingerprint or other biometric is similar to one collected at the scene. If e-commerce companies routinely collected biometrics, one could track an individual by searching on-line transaction records for that person's biometric pattern. Some people worry that fingerprints collected for biometric applications might find their way into an identification database like the FBI's AFIS. *see Note 4.*

At present, few promoters of biometric systems are intent on violating people's privacy. In fact, public concern about privacy has led many biometrics promoters to emphasize security and privacy features of their systems in order to ensure their adoption. This does

not always satisfy the critics. There is always a risk of "function creep" in that biometric information could find itself serving new and different functions as the system evolves over time. While it is possible to omit features from a system (like the ability to look up a fingerprint collected from a crime scene), there is no way to guarantee that the feature never appears later. Once an institution possesses the biometric readings, there is no foolproof technique to prevent those readings from being used for other purposes. The systems know *how* to match a biometric reading to the patterns they have collected, but they can't possibly understand *why* a particular reading has been presented.

7.2 BIOMETRIC TECHNIQUES

Just about every device ever built to measure a unique human characteristic has been pressed into service for biometric authentication. Some devices cost tens or hundreds of thousands of dollars, while others try to use equipment already available on a personal computer, like the keyboard or perhaps a microphone. A recent book on biometric technology identified 14 different types of biometrics that have been used, experimentally or in practice, for authentication. These different techniques fall into two categories: those that measure behavior and those that measure physical traits. *see Note 5.*

MEASURING PHYSICAL TRAITS

These techniques measure physical features of human bodies that should be unique among most or all of the population. Ideally, the measured feature is not affected by normal human behavior or cosmetics, and does not significantly change in a matter of days or weeks. Described below are the five major techniques.

- **Fingerprints**

  These systems typically use a relatively small image scanner to read the user's fingerprint (Figures 7.1 and 7.2). The technology evolved from earlier systems used to match printed fingerprints in law enforcement applications. Today's fingerprint biometric systems yield a very practical combination of cost and reliability.

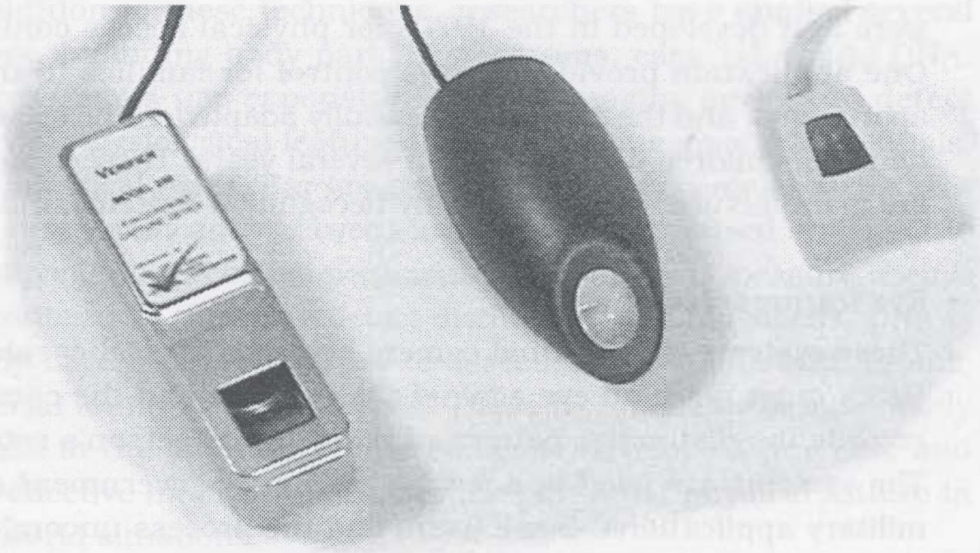


FIGURE 7.2: *Fingerprint scanners.* Like computer circuitry, their size and cost has dropped over time. Leftmost is the Model 250 Verifier from Cross Match Technologies. Next is the original Bio-Mouse from American Biometric. Rightmost is a recent fingerprint scanner built by Compaq.

However, a small portion of the population does not have good enough fingerprints for the scanners to recognize reliably. Numerous vendors produce hardware and software for fingerprint authentication, including the following: *see Note 6.*

- American Biometric
- AuthenTec
- Biometric Access
- Compaq
- CrossCheck
- DigitalPersona
- Identix
- Informer Systems
- Mytec Technologies
- Sony

- **Hand geometry**

  These systems use special readers that sense a user's finger size, thickness, and palm geometry. Mechanical systems of this type

were first developed in the 1970s for physical access control. One application provided access control for families to their apartments, and the system successfully adapted to the growing hands of children over a period of several years. These systems are manufactured commercially by Recognition Systems, Inc.

- **Eye features: retina**

  These systems use a retinal camera behind a special eyepiece. Users must place an eye against the eyepiece and the camera records the distinctive pattern of veins on the person's retina. The systems are used in a few high-security government and military applications. Some users find the process uncomfortable and even threatening, although the process cannot damage the eye.

- **Eye features: iris**

  These systems use a special camera that examine's the user's iris from a distance of a few feet and extracts its distinctive texture. Like fingerprints, individuals have unique irises even when they have similar or identical genes (that is, identical twins). Commercial iris authentication systems are produced by the following companies: *see Note 7.*

  - IriScan
  - Sensar

- Face

  These systems recognize people by examining images of faces. In authentication applications, the system has a camera that searches for a user's face and matches it against the face stored in the user record. The technology is widely used by law enforcement to search "mug shot" files. A shortcoming of face authentication is that it cannot distinguish between identical twins. Commercial face recognition systems are produced by the following companies: *see Note 8.*

  - Miros
  - Visionics

In addition to these techniques, researchers have studied several others, including body part thermograms, ears, odor, and DNA. Thermograms use expensive infrared imaging devices to detect internal physiological features that are unique among individuals, even identical twins. Ear imagery extracts fairly complex structures that vary significantly between individuals. Odor-based systems use sensitive (and rather temperamental) sensors that measure chemical contents of the air to detect distinctive personal smells. DNA is widely used in criminal trials to associate individuals with genetic material found at crime scenes. These techniques will most likely remain in the laboratory until someone develops a practical and cost-effective method to measure the particular physical feature in real-world situations.

MEASURING BEHAVIORAL TRAITS

These techniques measure some type of human behavior, and use that measurement to generate a unique biometric reading. Unlike physical traits, these biometrics do not have to measure the exact same phenomenon each time: a person could be directed to speak, write, or walk a certain way in order to reduce the risk of a replay attack.

- **Speaker recognition**

  These systems prompt a user for some spoken words, and authenticate the user based on distinctive speech patterns ("voiceprints"). Some systems rely on specific microphones or other restrictions on the speaking environment, while others try to use microphones included in modern "multimedia" personal computers. Speaker recognition can fail in noisy environments, or in some cases be fooled by tape recordings of the user's voice. Apple's Mac OS 9.0 included speaker recognition for authenticating different users of a shared Macintosh computer (Figure 7.3), which has made biometric authentication widely available to the general public. The following vendors produce speaker recognition systems for authentication: *see Note 9.*

  - Apple Computer

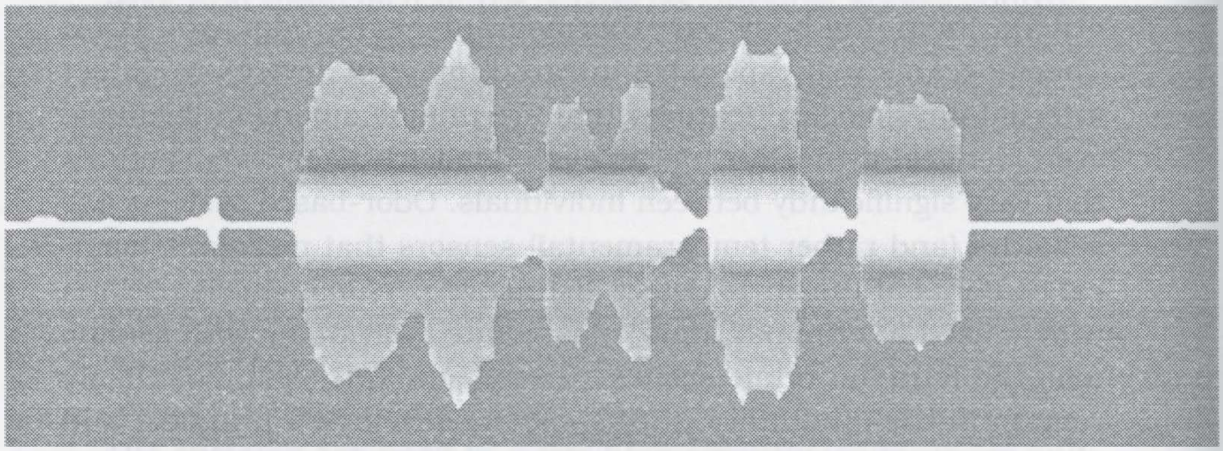


FIGURE 7.3: *Spoken passphrase collected by Apple's Mac OS 9.0.* This is a "voiceprint" for the passphrase "My voice is my password" collected by Apple's Mac OS 9.0. The system matches the way a user says the passphrase against previously collected patterns generated by the same phrase.

- Keyware Technologies
- T-NETIX
- Veritel

- **Written signature**

  These systems collect a written signature from the user and compare it against a record of that user's signature-writing behavior. In order to reduce the risk of forgery, the systems rarely rely on the signature's image alone: reliable systems measure the dynamics of pen motion as well. However, these systems are rarely used for authentication; they are more often used to capture a written signature in a digital form that makes forgery easier to detect. Software for signature authentication is produced by the following vendors:

  - Communication Intelligence Corporation
  - Cyber SIGN

- **Keystroke dynamics**

  These systems sense a user's behavior while typing at a computer keyboard and use unique features of that behavior to authenticate the user. The systems collect information about the time delay between hitting various pairs of keys and the amount of time a person takes to hit and release each key. Researchers have found that this information is distinctive among different individuals. A few recognition systems exist that use this technique. One obvious advantage is that the technique does not require additional equipment to sense user behavior.

Researchers have also built systems to collect information about how a user walks and use distinctive features of the user's gait for authentication. Such systems rely on special camera systems to collect the information. At present, no practical implementations have been demonstrated or used.

7.3 How Biometrics Work

Despite the range of sensors and measurements used for biometrics, all systems share the same fundamental design (Figure 7.4). All systems use a device that collects digitized measurements of a personal trait. The first step is to use the right equipment to measure the particular trait associated with the biometric. For example, an iris pattern system uses a specialized camera that focuses specifically on the iris and collects its image. Fingerprint systems typically use specially designed scanners to collect the fingerprint, and the scanners intentionally do poorly at detecting patterns in other types of surfaces.

Next, a biometric system performs *feature extraction* on the digitized data to identify the distinctive features associated with the particular biometric. In a fingerprint system, features consist of splits and junctions in the print's concentric ridges, combined with information about the features' relative locations. The extraction process yields a data item called the *biometric signature*. We use the term "signature" here in a general sense, and not just to refer to digitized versions of written signatures. Different biometric systems produce

their "signatures" from a user's fingerprint, voice, or other personal trait.

To authenticate a particular user, the system looks up the user's record, retrieves the *biometric pattern* associated with that user, and compares the derived signature against that pattern. The comparison process is generally designed to expect partial matches and to specify the degree to which the signature matches the pattern. If the two match closely enough, the system authenticates the user.

The system requires a biometric pattern for every recognizable user. The system constructs these patterns during an enrollment process that collects one or more biometric signatures from each user. The enrollment process uses additional signatures from a user to construct patterns that reflect likely variations in that user's signatures. This makes the system less likely to falsely reject a legitimate user.

Biometric systems designed for the uniqueness and identity applications can establish a user's identity without asking for the user's name. To do this, the system collects a biometric signature from the user and matches that signature against each biometric pattern in user database. The system identifies the user by finding the pattern with the closest match to the biometric signature.

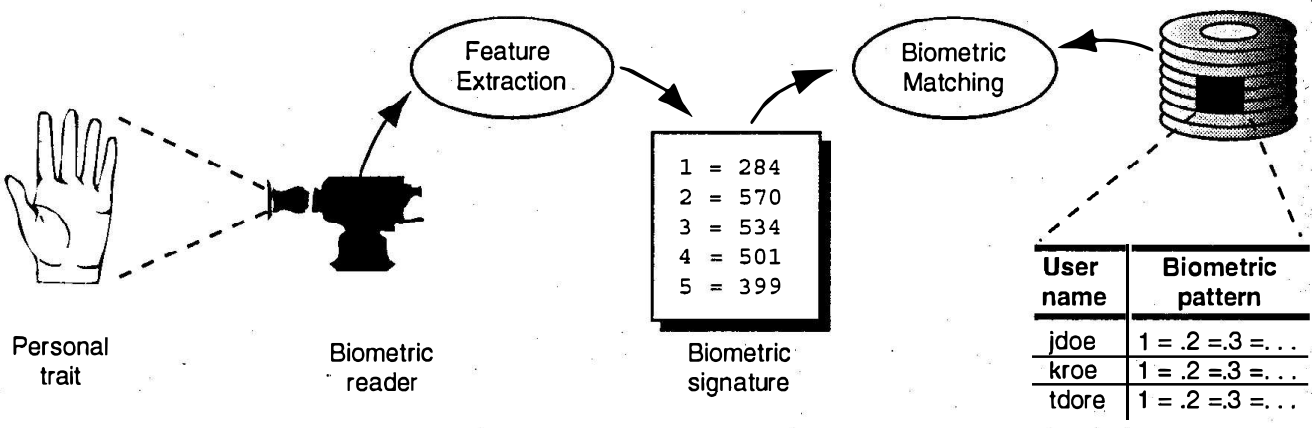


FIGURE 7.4: *Elements of a biometric system.* First, the biometric reader produces a digital repr sentation of the trait. Next, the feature extraction procedure generates a biometric signature co taining the individual's distinctive measurements. To authenticate someone, the system extrac the user's biometric pattern from the user record and compares how closely the person's biome ric signature matches the user's biometric pattern. Close matches are accepted.

7.4 TAKING A BIOMETRIC READING

The first step in any biometric procedure is to take a reading of the personal trait. The biometric reader in Figure 7.4 performs this task. The reader yields a digitized representation of the trait, which is converted into the biometric signature. This process is the first point where errors can work themselves into biometric authentication.

In some cases, the biometric reader might actually be a general-purpose device, like a flatbed image scanner or a computer's built-in microphone. For example, the speaker authentication process used by Apples's Mac OS uses a built-in microphone to collect the user's passphrase. Some face-recognition systems are also designed to use standard, off-the-shelf video cameras. While general-purpose devices can lower the cost of biometric authentication, they put more of a burden on the feature-extraction software. For example, different cameras might produce a "standard" video signal, but there will still be subtle differences between them that can affect the biometric pattern. The feature extraction process must be able to filter out the variations between different input devices.

Other systems use special-purpose biometric readers. In theory, fingerprint systems could use standard flatbed scanners, but in practice they tend to use specialized devices. Hand geometry readers collect three-dimensional information that a scanner cannot. In some cases, the reader may perform some or all of the biometric signature processing, so that the authentication system only has to perform a match on the resulting data.

FEEDBACK DURING BIOMETRIC INPUT

An interesting feature that shows up in many systems is a mechanism that displays the user's biometric image as the system tries to collect a reading. For example, fingerprint-recognition systems like the one provided with the BioMouse, from American Biometrics, displays a window with the image of the user's fingerprint. Users quickly learn how to place their finger correctly on the print reader, since correct placement has a distinctive appearance, and incorrect placement is generally rewarded with a false rejection.

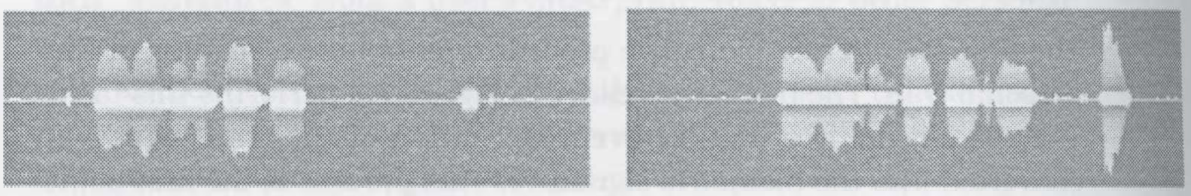


FIGURE 7.5: *Recognizing a passphrase in a voiceprint.* These two voiceprints represent a six-syllable passphrase, followed by noise. In the left example, a small burst of background noise followed the passphrase. In the right example, the user coughed after reciting the passphrase. Ideally, the system should isolate the phrase from other sounds, since that increases the likelihood of correctly matching the corresponding biometric pattern.

Apple's Mac OS 9 displays a waveform when someone uses its voiceprint authentication, as shown in Figure 7.5. Users can quickly learn to identify excessive background noise and other audio problems that may interfere with correct authentication. The figure shows two waveforms produced when the user uttered the phrase "My voice is my password." In both cases the system produced a group of waveforms that obviously correspond to the six-syllable phrase. In the left-hand example the phrase is followed by a small burst of noise, and in the right-hand example it is followed by a cough. A sophisticated recognition system can automatically isolate the passphrase from the noise. However, the display also helps out users of less-sophisticated systems. If the system rejects the passphrase, the user can "see" the noises by looking at the waveform. This at least gives the user the opportunity to try to prevent the troublesome noises when speaking a passphrase.

FORGING A PHYSICAL TRAIT

The biometric reader is an obvious point of attack on a biometric system. If the reader simply collects a digital image from a sensor and passes it on, then we can masquerade as someone by feeding the sensor the same data that the victim would provide. The strategy is obvious for fingerprint systems: we produce an image of the victim's fingerprint and present that to the system instead of the victim's finger. The strategy could work, too, if fingerprint systems typically used standard scanners to collect fingerprints.

A-41

As we noted earlier, however, fingerprint systems tend to use specialized scanners to collect fingerprints. These scanners are usually designed to detect the three-dimensional structure of fingerprint ridges and troughs as opposed to their simple two-dimensional appearance. While this may prevent masquerades using photocopies of fingerprints, it doesn't always prevent more sophisticated ones.

For example, the product-testing laboratory at *Network Computing*, a trade magazine, tried out two sophisticated fingerprint forgery techniques against several fingerprint readers. In one case, they collected "latent" fingerprints from a table using techniques similar to those of police detectives. Then they photocopied the latent prints onto transparency material and used the transparency to masquerade as a person's fingerprint. This technique actually succeeded in tricking a couple of fingerprint systems. This could be a serious problem since it suggests that an attacker can pick up a user's fingerprint left casually on a glass, doorknob, or other surface and then masquerade as that user.

see Note 10.

In another experiment, the lab took wax impressions of users' fingers and constructed fake fingers out of rubber. These copies succeeded in tricking most scanners. Fortunately, this technique requires the victim's cooperation.

The basic countermeasure for such attacks is to use more sophisticated sensors that measure additional properties of the physical trait. This is why many fingerprint readers try to detect the print in three dimensions, so that the ridges must actually be separate from the troughs to appear in the biometric reading. Some fingerprint readers sense skin capacitance or temperature in order to spot fakes.



D-41

A similar type of replay can take place in behavioral biometrics, notably speaker recognition. For example, an attacker can collect a high-quality recording of the victim's passphrase. In many systems, the microphone will filter out most differences between live speech and the playback of a high-quality recording. An extreme and fictional example of this appeared in the 1992 film *Sneakers*.



A-42

A countermeasure to prevent such replay is to require a different behavior for each authentication attempt. This is the same basic idea as the "challenge response" passwords described in Chapter 10. Some speaker-recognition systems achieve this by prompting the

user for a different passphrase each time. A relatively simple example is to have the user recite a sequence of digits and randomly choose a new sequence for each authentication. A more sophisticated approach would be to collect an arbitrary sentence of spoken words. However, any of these strategies will place a significant burden on the biometric system. While it may be tolerably convenient to teach a system to recognize the user when reciting a fixed phrase, varying phrases will take a lot more training. The system must be trained to recognize any words that might be used and to recognize any peculiarities about how the user might use pairs of words together. Thus, such systems might not be practical for some environments.



D-42

7.5 BUILDING AND MATCHING PATTERNS

Once the system has collected a biometric signature, the next problem is to match it reliably with the right user. If Cathy produces the biometric signature, then it should match the biometric pattern stored in her user record. It should not match Tim's record.

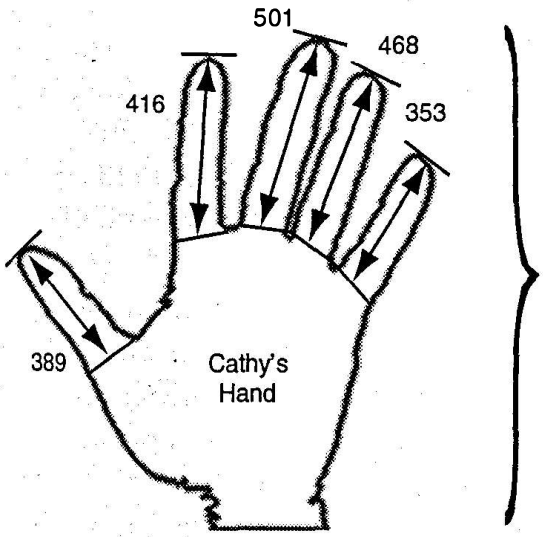
A physical trait makes a good biometric if it yields a set of unique measurements for each person. The trait's individual properties, like its overall size or basic arrangement, might not provide unique measurements by themselves. For example, we could try to use people's glove sizes as a biometric, except that lots of people share the same glove size. The biometric is vulnerable to successful masquerades if many people share the same measurement.

A practical biometric system must collect enough measurements of the personal trait so that their aggregate will vary from one person to another. A single spoken word or cough doesn't provide enough information for a speaker recognition system to distinguish between different people, so people often have to recite a phrase to such systems (Figure 7.5). Longer speech gives the system a better chance to detect the inevitable differences in the way that different people speak.

EXAMPLE: A TRIVIAL HAND GEOMETRY BIOMETRIC

A sophisticated hand geometry system might measure finger lengths, widths, thicknesses, and the relative location of crucial features like knuckles and other joints in order to construct a unique pattern for each person. However, we can easily illustrate the essential features of biometric matching by looking at the grossly simplified version in Figure 7.6: a biometric based on the lengths of the person's fingers. The biometric measuring process generates a signature that is an array of five numeric values, one for each finger in order, thumb first.

It should be obvious that biometric finger-length measurements will not be precise; that is, we can get subtly different measurements each time we measure the same finger lengths. Although adult fingers don't tend to lengthen or shorten significantly over time, the elastic nature of skin makes it impossible to get the same measurements consistently. To measure the length of a finger from a palm image, the system must be able to identify enough features of the hand to compute each finger's length. While fingertips may be easy to recognize, the location of each finger's base depends on folds



Today's Biometric Signature from Cathy:	Cathy's Stored Biometric Pattern:	Tim's Stored Biometric Pattern:
389	390	284
416	418	570
501	502	534
468	471	501
353	355	399
	Distance = 4 from that signature	Distance = 199 from that signature

FIGURE 7.6: *Biometric example using trivial hand geometry.* We construct the biometric signature of Cathy's hand by measuring her finger lengths as the features of interest. The signature consists of a five-element array, with the elements representing measured finger lengths. We match the signature against Cathy's pattern by computing the mathematical "distance" between them. Note that the distance is lower (4) than the distance between the signature and Tim's pattern (199).

of skin and other elastic hand features. The length of a given finger can vary with changes in the hand position and finger pressure, since these properties affect the folds of the skin between fingers.

In Figure 7.6, Cathy's hand yielded a particular set of lengths that are reflected in the signature. If Cathy submits to another biometric reading, the next signature might yield a thumb length of 392 instead of 389, a middle finger length of 503 instead of 501, and so on. Similar variations appear in any biometric measurement. Fingerprint features will be slightly farther apart or closer depending on how the finger is placed on the reader. Spoken phrases used for authentication will be said more quickly or with different emphasis regardless of how careful the speaker is (see Figure 7.5). A well-designed biometric system tries to factor out such variations as much as possible, but variations will inevitably crop up.

When we try to match Cathy's signature against her biometric pattern, we want the signature to match regardless of the imprecision. Naturally, this requires a completely different matching strategy from what we use for passwords or other more traditional base secrets, since we're matching approximate values instead of exact values. For the finger-length biometric, we can use the simple strategy of computing the mathematical "distance" between a given signature and a given pattern. In other words, we treat the signature's five values and the pattern's five values as representing geometric points in five-dimensional space, and we measure how far these points would be from each other. Thus, a close match yields a small number. In Figure 7.6, Cathy's own pattern is obviously a close match, and it yielded a distance of 4. Tim's pattern, on the other hand, yielded a distance of 199.

Real-world biometric systems tend to be more complex than the simple finger-length computation. Most biometric signatures require a lot more than five integers to represent their measurements. Some signatures include graph representations to indicate positional relationships between features as well as, or instead of, relative locations. The pattern-matching procedure likewise will be tailored so that results are accurate even if the signature's measurements aren't always precise. It is a delicate balancing act to design a biometric system. The system must pick up the features that make different people distinct but avoid picking up "noise" or other infor-

mation that confuses the reading. If the system rejects too much apparent noise, it may also reject the details that distinguish one user from another.

ENROLLING A USER

In the simplest case, biometric enrollment isn't particularly different from enrolling a user with a password. Instead of typing in a new password, however, the new user provides a biometric reading. The system uses the reading to construct a biometric pattern to store in the user record. However, this doesn't provide a reliable pattern in all cases. Some biometric systems require *training* so that the system can reliably recognize the person.

Training is necessary when the biometric trait, like a spoken phrase, is subject to too much variation from one instance to the next. In order to construct a reliable biometric pattern, the system must collect several biometric signatures from the user and combine them to construct the pattern. In some cases, the combining process simply takes an average of the different readings and perhaps retains information about the readings' variance. In other cases the system might analyze the signatures' differences in detail to figure out which of the user's measurements are distinctive, and then build a pattern to detect those measurements.

From a usability standpoint, the best systems require only a single biometric signature to enroll a user. Physical traits like fingerprints or hand geometry usually require a single reading. Behavioral traits like voice or written signature are more subject to variation and often require multiple readings to train the system to recognize each user. If fingerprint systems take multiple readings, the readings are usually taken from different fingers. The separate readings then allow the user to authenticate using different fingers, in case one or another is injured.

7.6 BIOMETRIC ACCURACY

For a biometric technique to work, it must be *accurate;* that is, it must be able to reliably associate the biometric signature taken from a particular person with that person's biometric pattern. A perfect system would never generate either false acceptances or false

rejections. We measure a biometric system's accuracy by estimating the percentage of false acceptances and false rejections the system should typically produce.

An accuracy estimate is usually based on experiments that test a large number of biometric signatures to see how closely they match a set of biometric patterns. In general, the experiments try to collect a large number of readings from a large number of randomly selected people, preferably from the population that is likely to use the biometric system.

The graphs in Figure 7.7 show what the results might look like, given that the matching procedure measures the difference (or distance) between a signature and a pattern. Ideally, whenever a user like Cathy uses the system, it should compute a difference of zero between the biometric signature it collects and Cathy's pattern. In practice, of course, the difference will vary with each signature. As more and more people use the system, we should find that most signatures yield a low difference, but users will occasionally produce signatures that yield a relatively high difference. This is shown on the left-hand graph. If one of Cathy's signatures yields a difference that is high enough, the system will conclude that someone other

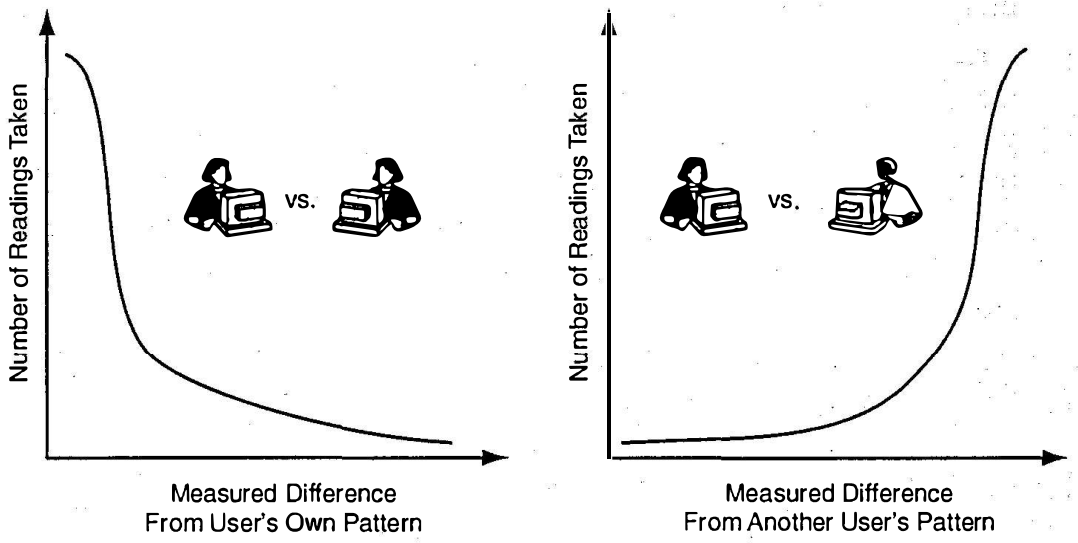


FIGURE 7.7: *A reasonably accurate biometric system.* These graphs show how a biometric system should match readings to individual users. On the left, we collect readings from authorized users and measure the difference between those readings and the users' own biometric patterns as collected by the system. On the right, we collect readings from arbitrary members of the population and measure their difference from randomly selected patterns belonging to authorized users.

than Cathy produced the signature and reject her authentication attempt.

The right-hand graph shows what we find if we calculate the difference between a signature produced by a random person and a biometric pattern belonging to someone else. Ideally, all such computations should yield a very large value. In practice, most such computations produce a large value, but a few produce smaller values, as shown on the graph. Thus, we will occasionally encounter successful masquerades if enough people try.

TRADING OFF USABILITY AND SECURITY

Figure 7.8 combines those two graphs to show how we can tune the system to trade off security and usability. The trick is to balance the likelihoods of false rejection and false acceptance, so the system rarely locks out legitimate users and it doesn't fall for masquerades. To some extent, this depends on the shape of the two curves, which depend on the design of the biometric signature and the matching procedure.

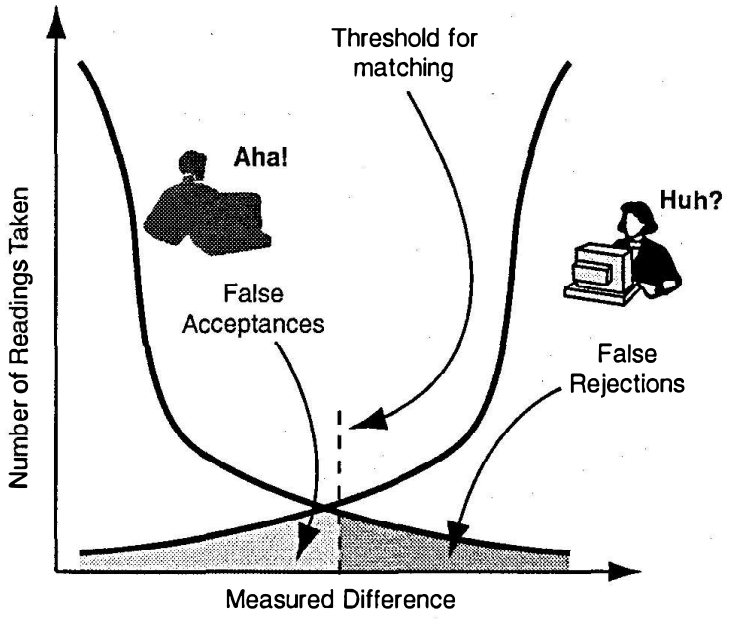


FIGURE 7.8: *Balancing usability and security.* This graph combines the two curves from Figure 7.7. Biometric readings from a particular user should closely match that user's biometric pattern; at worst, the difference should never exceed the matching threshold shown by the dashed line. We reduce the risk of masquerades (false acceptances) by moving the threshold to the left, which reduces the amount of gray area on the left. However, this reduces usability because it increases the likelihood of false rejections (the gray area on the right).

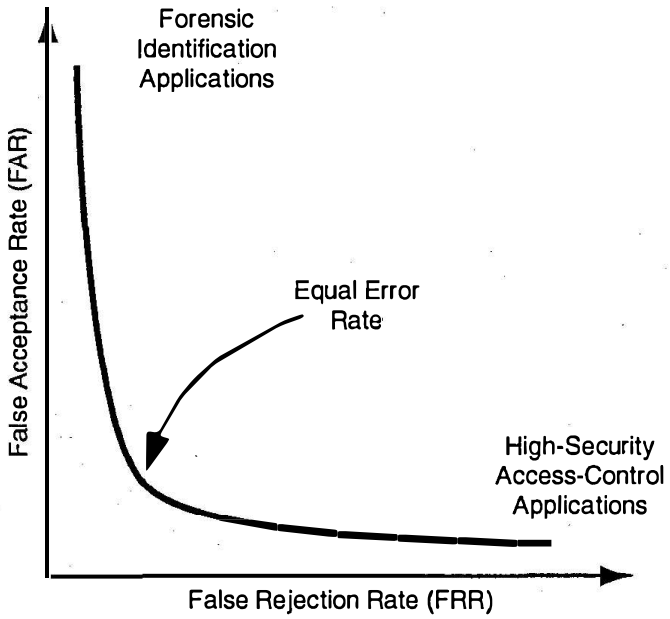


FIGURE 7.9: *ROC-curve and biometric applications.* Each point on this curve represents a different threshold value for biometric matching, and plots the resulting FAR against the resulting FRR. As noted in the upper left, forensic identification applications seek a low false rejection rate and can tolerate a high false acceptance rate. Access-control applications, on the other hand, will place a larger premium on a low false acceptance rate even at the cost of numerous false rejections.

In addition, the trade-off depends on the amount of difference the system is willing to tolerate when matching a signature. If a signature exceeds the system's tolerance threshold, the system will reject the match. The dashed vertical line in Figure 7.8 represents this threshold. The shaded area to the left of the threshold indicates how often random masquerades might succeed. We can generally decrease that risk by moving the threshold to the left, reducing its value. The shaded area to the right indicates how often the system might reject legitimate authentication attempts. As we reduce the threshold and reduce the risk of masquerade, we increase the risk of occasionally locking out a legitimate user.

Developers and analysts of biometrics typically use the results of these experiments to compute error rates based on different choices for the threshold value. This yields the *false acceptance rate*, called the FAR, and the *false rejection rate*, called the FRR, for different threshold values. These rates are scaled as percentages or fractions of the number of matches attempted. Since a particular threshold value will yield a particular FAR and FRR, we can plot one against

the other to produce a *receiver operating characteristics curve*, or ROC-curve, as shown in Figure 7.9. The curve will generally contain a point, called the *equal error rate*, at which the FAR and FRR are equal. A low equal error rate usually indicates that the system can achieve strong security without excessive numbers of false rejections.

AVERAGE ATTACK SPACE

We can also use the average attack space concept from Section 2.7 to roughly compare attacks on passwords with attacks on biometrics. In particular, we can use the FAR to compute an average attack space for a biometric system. This is because both computations are based on similar types of large-scale random attacks. To find the attack space (**V**) from the FAR ($\mathbf{A_{FAR}}$), we compute the following:

$$V = \log_2\left(\frac{2}{A_{\text{FAR}}}\right)$$

Some biometric systems, like those for fingerprint recognition, support a range of threshold values that allow each site to choose their own trade-off between safety and ease of use. In particular, some systems calibrate this threshold in terms of the false acceptance rate. Table 7.1 compares different FAR values with other average

TABLE 7.1: *Biometrics and Average Attack Spaces*

Example	Style of Attack	Average Attack Space
Biometric with a "one percent" FAR (1 in 100)	Interactive	6 bits
Random 10-letter English text	Off-line	16 bits
Biometric with a FAR of 1 in 100,000	Interactive	16 bits
Biometric with a FAR of 1 in 1,000,000	Interactive	19 bits
Personally chosen 8-character Unix password	Off-line	22.7 bits
Password from a FIPS 181 10-letter password generator	Off-line	39.5 bits
56-bit Data Encryption Standard (DES)	Off-line	54 bits
128-bit Advanced Encryption Standard (AES)	Off-line	127 bits

attack spaces. Biometrics can rely on relatively small average attack spaces because a well-designed system is not vulnerable to the off-line attacks that plague encrypted password databases. A well-built biometric system will accept only interactive authentication attempts, in which case a 10- to 20-bit average attack space should suffice.

7.7 BIOMETRIC ENCRYPTION

Cryptographic techniques like encryption provide a way to incorporate biometrics into the modern world of network computing. By themselves, biometric signatures can be easy to forge. Moreover, many people are worried about the privacy risks inherent in sharing their biometric readings. The most practical way of addressing these problems is to use cryptographic techniques to protect biometric information. Because most cryptographic techniques require base secrets, we must manage base secrets as well as biometric patterns.

PRESERVING SECRECY

Section 7.1 noted how many people are concerned about privacy risks in biometric systems. If an attacker can intercept a person's biometric data, then the attacker might use it to masquerade as the person, or perhaps simply to monitor that person's private activities. In any case, biometric systems are more likely to be accepted by a user community if the system takes steps to preserve the privacy of people's biometric information.



While systems built along the lines of the local authentication pattern can rely on physical security for privacy assurances, this obviously won't work for networked and distributed systems. They must use encryption. Some models of fingerprint readers even provide encryption to protect the fingerprint data as it travels from the reader to the computer.



Figure 7.10 sketches the basic strategy for cryptographically protecting biometric data. The biometric reader collects its data and cryptographically protects it before transmitting it to other devices in the system. The reader resides within a physically protected perimeter established by the reader's case or within a physically secure office environment. The device stores a base secret within

that perimeter, and uses the base secret to cryptographically protect the biometric data. The device that receives the biometric data must also reside within a physically safe environment, and it maintains its own copy of the base secret in order to decrypt and verify the biometric data.

The biometric data remains safe as long as the base secrets are kept safe and attackers can't breach the physical perimeters. Base secrets can complicate the process of managing biometrics, since they represent another sensitive data item to distribute and maintain. One way to simplify the handling of encryption keys is to use public key encryption, which is discussed further in Chapter 13. A shortcoming of public key encryption is that it can be easy for an attacker to masquerade as a legitimate biometric reader. However, there can be similar weaknesses in other encryption techniques, so a careful implementation may still consider the authenticity issues discussed later in this section. *see Note 11.*

In Section 1.3 we introduced hashing as a strategy to protect passwords from theft. There is no practical way to deduce the original data from the hashed information, so it might seem like a promising strategy to protect biometric privacy. Unfortunately, we cannot hash biometric data and use the resulting hash in the biometric matching procedure. An essential property of a safe hash function is that the hashes of two similar data items aren't at all likely to be

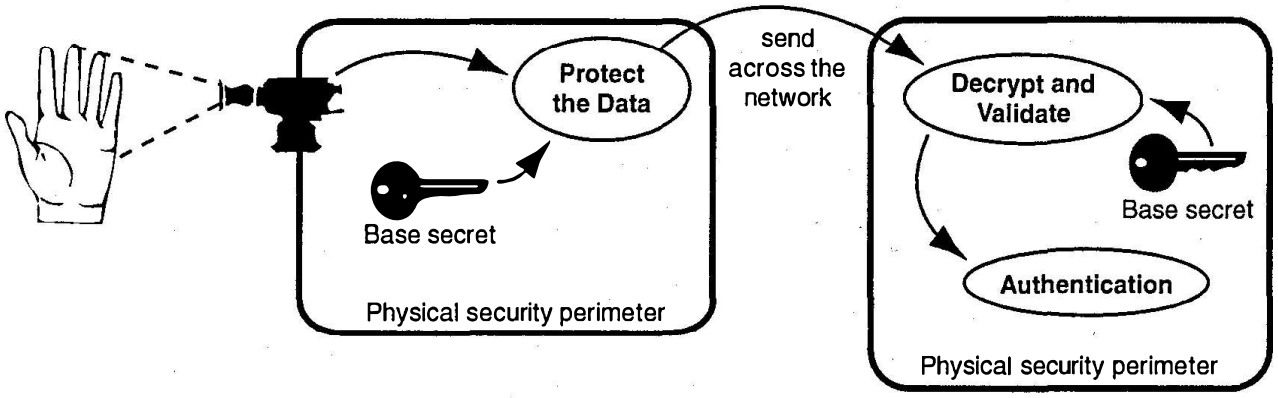


FIGURE 7.10: *Using cryptography to protect biometric data.* In distributed applications, we need to protect biometric data as it travels between devices. We can cryptographically protect it from eavesdropping and/or from modification. On the left, the biometric reader collects a reading. The reader uses a base secret to protect the reading as it is transmitted to another device. On the right, the recipient uses the base secret to decrypt and/or validate the biometric data.

similar. Thus, we can detect the match between a biometric signature to a similar biometric pattern, but we can't expect their hashes to match at all. This makes hashing an impractical strategy for protecting biometrics.

AUTHENTICITY OF BIOMETRIC DATA

The fundamental value of biometric authentication is that it relies on information that only a specific individual can provide. In practice, of course, anybody can provide any piece of information that they happen to have a copy of. While John Doe might not have the same thumbprint as Cathy Roe, he might have a digitized copy of it. If he does, then he might be able to send it to a computer and claim to be Cathy Roe. This is the essence of the biometric replay attack: the attacker intercepts the victim's biometric feature somehow, constructs a biometric signature that matches it, and transmits the signature to masquerade as the victim.


A-44

In a sense, this puts us back where we started. The computer needs to authenticate Cathy Roe. If she sends a biometric for authentication, the computer can't tell if someone forged her biometric signature. The solution is, of course, to authenticate the biometric authentication data. To do this, we must install a base secret in the biometric reader as shown in Figure 7.10. The base secret is used in conjunction with a cryptographic integrity check. If an attacker tries to forge a biometric signature, its integrity check will fail, because the integrity check depends on the base secret. The integrity check will also detect attempts to modify a legitimate biometric signature. Section 8.5 describes how these integrity checks work.


D-44

In other words, we can't trust biometric authentication on an untrustworthy network, unless we distribute base secrets to authenticate the biometric readers. This largely eliminates the perceived benefit of biometrics for some proprietors, since they still need to distribute and manage base secrets to make the system work. On the other hand, this is an acceptable expense in some applications. Many biometric uniqueness applications will require distributed biometric readers and a centralized biometric database. The benefits of the uniqueness check can outweigh the costs authenticating the biometric data.

THE PROBLEM OF BIOMETRIC EXPLOITATION

Even if a particular system isn't explicitly designed to support individual surveillance or law enforcement, some observers fear that the system might someday be exploited for those purposes. Some organizations with large-scale biometric systems have established privacy guidelines and even legislation that restricts the use of biometric data in order to address those concerns. Such systems may also incorporate special technical measures, like biometric encryption and tight security around the biometric database, to ensure the privacy of biometric data.

Nonetheless, some privacy advocates remain skeptical and even alarmed at the remaining privacy risks. The fundamental problem is that we cannot prevent "function creep." Even if today's policies, guidelines, and laws protect users' biometric privacy, there is nothing to prevent new legislation from stripping away the privacy protection. Some observers have suggested that a form of biometric encryption might be able to prevent an existing biometric database from being exploited in an unanticipated way. Such protections might, in theory, prevent a system from accepting any query except from one of its own live sensors. The protections might also prevent third parties from merging the system's databases with those of a large-scale identification system, like FBI's AFIS database. *see Note 12.*

Despite the hopes of some observers, there really is no technological fix to prevent misuse of biometric systems. The problem is a simple one: computing equipment is completely amoral and cannot tell whether it is being used for "good" or "bad" purposes. If a system can find an identity based on a biometric signature, then there is no purely technological way of controlling *why* a given search is performed. Given enough collusion among system operators and proprietors, there is always a way to fool the system into performing its function for unintended purposes.

The most promising technique for controlling the use of biometric systems is to place the system in the custody of a consortium of interested parties. The consortium becomes responsible for protecting users' privacy and for ensuring that the system operates efficiently and reliably. The consortium can establish procedures to monitor the system's integrity and satisfy all members that biometric data remains under the system's control. While it is not clear

whether technology exists to prevent a highly motivated conspiracy from subverting the system's privacy, the consortium can use multi-person control techniques and separation of duty to minimize the risks.

7.8 SUMMARY TABLES

TABLE 7.2: *Attack Summary*

Attack 	Security Problem	Prevalence	Attack Description
A-41. Forged biometric trait	Masquerade as someone else	Physical	Attacker makes a copy of the user's physical trait and presents the copy to the biometric sensor
A-42. Replay user behavior	Masquerade as someone else	Physical	Attacker records the user's behavioral trait and replays it for the biometric sensor
A-43. Biometric interception	Masquerade as someone else	Common	Attacker sniffs a victim's biometric sensor reading or signature while being transmitted between devices, invading the user's privacy
A-44. Biometric signature replay	Masquerade as someone else	Common	Attacker replays a biometric sensor reading or signature so that it is treated as having been collected from the actual person

TABLE 7.3: *Defense Summary*

Defense	Foils Attacks	Description
D-41. Detailed sensing	A-41. Forged biometric trait	Collect additional biometric measurements from the person since a simple copy of the biometric will not have all of the properties of a live person
D-42. Varied behavior	A-41. Forged biometric trait A-42. Replay user behavior	The system requests a different action from the person for each authentication attempt so that a simple recording and replay won't be accepted
D-43. Biometric data encryption	A-43. Biometric interception	Encrypt biometric data when it is transmitted between different devices
D-44. Biometric signature authentication	A-44. Biometric signature replay	Apply authentication measures to biometric data to ensure that it came from a trustworthy source

Networking/Security

Authentication



Computer access control is an ongoing challenge. Left to themselves, computers tend to treat one user no differently than another. Computers use authentication to confidently associate an identity with a person. ***Authentication: From Passwords to Public Keys*** gives readers a clear understanding of what an organization needs to reliably identify its users and how different techniques for verifying identity are executed.

Authentication is one of the basic building blocks of security. To allow a computer system to distinguish between legitimate users and others, most sites give passwords to authorized users. Unfortunately, just as car thieves have found ways to defeat sophisticated locks and alarms, computer hackers are always finding new ways to circumvent password systems. The good news is that organizations now have available to them a broad range of alternatives to passwords, and a variety of ways to make passwords safer. A well-designed authentication system allows users to prove their identities conveniently and gain access to the network without threatening the safety of the organization.

The first of its kind, ***Authentication*** describes the entire range of authentication methods used today. It examines situations in which certain techniques fail and points out ways to strengthen them. Network professionals, designers, developers, administrators, planners, and managers will find in these pages the authentication strategy to protect their valuable systems. Through diagrams and examples, the author thoroughly explains the technical concepts behind authentication, focusing on existing, off-the-shelf solutions to security problems.

Authentication highlights real products and solutions. If you are a network professional searching for the how and why of computer authentication, this is the book that will help you prevent unauthorized access on your network.

Richard E. Smith, Ph.D., CISSP has spent over a decade as a computer security professional. Author of *Internet Cryptography* (Addison-Wesley, 1997), he is a researcher and information security architect at Secure Computing Corporation, a leading computer security company.

http://www.informit.com/aw

Cover design by Katie Noyes
Cover illustration by Peter Steiner, The Cartoon Bank

Addison-Wesley
Pearson Education

ISBN-13: 978-0-201-61599-9
ISBN-10: 0-201-61599-1

9 780201 615999 55499

$54.99 US $59.99 CANADA