Exposing Infiltrators and Traitors

Canary Traps, Barium Meal Tests, Embedding, Trap Streets, Traitor Tracing and Fingerprinting.

How to evaluate new members...

Weed out informants and agent-provocateurs.

Assessing the risks.

It is imperative that you run tests to verify the reliability and integrity of new recruits who are applying to join your cell. Failure to evaluate recruits will result in your group being penetrated by your adversary.

Every time you admit a new recruit into your cell you are risking the security of your group. Yes, the recruit might be a *bona fide* supporter of your cause – or he might be an informant or an agent-provocateur.

The Informant. The informant is a cell member who is providing information to your adversary. He may betray you for money. She may betray you because she is being blackmailed. He may betray you because he is unethical, immoral, and weak-willed. She may betray you because she has a passive-aggressive personality disorder.

The Agent-provocateur. The agent-provocateur is someone who feigns enthusiastic support for your cause while enticing you to commit acts that are illegal. She is acting on the instructions of the Security Services – or she may actually be an MI5 agent. You are being set up for arrest, interrogation, and conviction.

The Mole. The mole is a cell member who quietly works to sabotage your operations. He may deliberately *forget* to do things that result in failed operations. He may intentionally *ruin* meetings with specious arguments and pointless debate, often introducing paranoia into the discussion. A typical mole is a long-time cell member who has been recruited by the Security Services, perhaps by blackmail. Less frequently the mole is an MI5 agent who has penetrated the organization at an early stage in its development.

The Counterintelligence Role. It is vital that your organization have a *counterintelligence member*. This is someone whose role is to detect and neutralize attempted penetrations by the enemies of your organization. Whether this is a formal position or an *ad hoc* role is not important. Someone in your group must take steps to systematically and conscientiously evaluate new recuits.

Uncover informants...

Here is how established resistance movements uncover informants.

Reveal some sensitive bogus information to the suspected informant, then wait for things to go wrong.

First, reveal some sensitive information to the recruit – and *only* to the recruit. For example, you might inform him of the existence of a (bogus) hidden cache of weapons.

Then wait and watch. If the cache is suddenly discovered by the authorities, you may be dealing with an informant. More tests may be required to confirm your suspicions.

In serious cases where you're playing by Big Boys' Rules, you might need to use live bait. If your adversary is sophisticated and experienced, you might need to reveal genuine secrets to the recruit you're evaluating. For example, you might reveal the name of a *whistleblower* who is leaking information to you about your adversary. If your recruit betrays your information to your adversary, you'll have lost your whistleblower – but you'll have unmasked an informant before he can do too much damage.

Unmask an agent-provocateur...

The most reliable method for unmasking an agent-provocateur is to ask him to be the first to commit to action.

Here is how any organization can unmask an agent-provocateur.

If the person is full of ideas for future operations, then *insist that he lead by example*. Make him commit himself first. Or, to put it another way, make him incriminate *himself* first before asking others to risk injury, exposure, or arrest.

If the person balks, then he may simply be "all talk". Or he may be a coward. Or he may be an agent-provocateur. In either case, you've called his bluff and now you know not to fall for his *jive-talk*.

Enforce compliance...

Here is how resistance movements enforce compliance with the counterintelligence functions.

If a trusted cell member brings an outsider into your group – or reveals sensitive information to an outsider – without performing any of these counterintelligence measures, then that cell member must be severely disciplined.

Depending on your situation, simply ostracizing the individual may suffice. Revoking his membership may be all it takes to remove the threat he poses. Or firmer steps may need to be taken.

SOME OTHER TECHNIQUES AND IDEAS FOR DISCUSSION BELOW:

Canary trap

A **canary trap** is a method for exposing an information leak, which involves giving different versions of a sensitive document to each of several suspects and seeing which version gets leaked.

The term was coined by Tom Clancy in his novel *Patriot Games*, though Clancy did not invent the technique. The actual method (usually referred to as a **Barium meal test** in espionage circles) has been used by intelligence agencies for many years. The fictional character Jack Ryan describes the technique he devised for identifying the sources of leaked classified documents:

Each summary paragraph has six different versions, and the mixture of those paragraphs is unique to each numbered copy of the paper. There are over a thousand possible permutations, but only ninety-six numbered copies of the actual document. The reason the summary paragraphs are so lurid is to entice a reporter to quote them verbatim in the public media. If he quotes something from two or three of those paragraphs, we know which copy he saw and, therefore, who leaked it.

A refinement of this technique uses a thesaurus program to shuffle through synonyms, thus making every copy of the document unique.

Barium meal test

According to the book *Spycatcher* by Peter Wright (published in 1987) the technique is standard practice which has been used by MI5 (and other intelligence agencies) for many years, under the name "Barium meal test". A Barium meal test is more sophisticated than a canary trap because it is flexible and may take many different forms. However, the basic premise is to reveal a secret to a suspected enemy (but nobody else) then monitor whether there is evidence of the fake information being utilised by the other side. For example, the double agent could be offered some tempting "bait" e.g. be told that important information was stored at a dead drop site. The fake dead drop site could then be periodically checked for signs of disturbance. If the site showed signs of being disturbed (in order to copy the microfilm stored there) then this would confirm that the suspected enemy really was an enemy e.g. a double agent.

Embedding information

The technique of embedding significant information in a hidden form in a medium has been used in many ways, which are usually classified according to intent:

- Watermarks are used to show that items are authentic and not forged.
- Steganography is used to hide a secret message in an apparently innocuous message, in order to escape detection.
- A canary trap hides information in a document that uniquely identifies it, so that copies of it can be traced.

Appearances in fiction

The canary trap was also used in Clancy's (chronologically) earlier novel, *Without Remorse*, when a CIA official alters a report given to a senator, revealing an internal leak who was giving information to the KGB.

Barium meals are also administered in Robert Littel's book The Company, and later in the TV short-series with same name.

The technique (not named) was used in the 1970s BBC television serial *1990*.

A variation of the canary trap was used in *Miami Vice*, with various rendezvous dates leaked to different groups.

Appearances in media

When distributing *Broken* to friends, Trent Reznor claims that he watermarked the tapes with dropouts at certain points so that he could identify if a leak would surface.

Screener versions of DVDs are often marked in some way so as to allow the tracking of unauthorised releases to their source.

- Fingerprinting gives a good overview of different kinds of canary trap techniques.
- EFF.org DocuColor Tracking Dot Decoding Guide How to read the date, time, and printer serial number from forensic tracking codes in a Xerox DocuColor color laser printout.

Trap street

A **trap street** is a fictitious entry in the form of a misrepresented street on a map, often outside the area the map nominally covers, for the purpose of "trapping" potential copyright violators of the map, who will be unable to justify the inclusion of the "trap street" on their map. On maps that are not of streets, other "copyright trap" features (such as non-existent towns or mountains with the wrong elevations) may be inserted or altered for the same purpose.[1]

Trap streets are often nonexistent streets; but sometimes, rather than actually depicting a street where none exists, a map will misrepresent the nature of a street in a fashion that can still be used to detect copyright violators but is less likely to interfere with navigation. For instance, a map might add nonexistent bends to a street, or depict a major street as a narrow lane, without changing its location or its connections to other streets.

Trap streets are routinely denied and rarely acknowledged by publishers. This is not always the case, however. A popular driver's atlas for the city of Athens, Greece, warns inside its front cover that potential copyright violators should beware of trap streets.[2]

In an edition of the BBC Two programme *Map Man*, first broadcast 17 October 2005, a spokesman for the Geographer's A-Z Street Atlas company claimed there are "about 100" trap streets included in the London edition of the street atlas. One such street, "Bartlett Place", a genuine but misnamed pedestrian walkway, was identified in the programme, and will appear in future editions under its real name, Broadway Walk.

Legal issues

Street traps appear not to be copyrightable, at least under the federal law of the United States. In *Nester's Map & Guide Corp. v. Hagstrom Map Co.*, 796 F.Supp. 729, E.D.N.Y., 1992, a United States federal court found that copyright traps are not themselves protectable by copyright. There, the court stated: "[t]o treat 'false' facts interspersed among actual facts and represented as actual facts as fiction would mean that no one could ever reproduce or copy actual facts without risk of reproducing a false fact and thereby violating a copyright . . . . If such were the law, information could never be reproduced or widely disseminated." (Id. at 733)

In a 2001 case, the Automobile Association in the United Kingdom agreed to settle a case for £20,000,000 when it was caught copying Ordnance Survey maps. In this case, the identifying "fingerprints" were not deliberate errors but rather stylistic features such as the width of roads.[3]

In another case, the Singapore Land Authority sued Virtual Map, an online publisher of maps, for infringing on their copyright. The Singapore Land Authority stated in their case that there were deliberate errors in maps they had provided to Virtual Map years earlier. Virtual Map denied this and insisted that they had done their own cartography

Fingerprinting

Neal R. Wagner.

Ordinary human fingerprints are often used for identification. This writeup extends the notion of *fingerprint* to include characteristics of any object that distinguish it from other objects. The word *fingerprinting* refers here to the process of adding fingerprints to an object and recording them, or of identifying and recording fingerprints that are already present.

People commonly confuse these fingerprints with digital signatures. Such a signature authenticates an electronic object to identify the object, perhaps through the individual who created it. Fingerprints are usually intrinsic to an object and not easily removed; in contrast, signatures cannot be forged but can easily be stripped off the object. Other techniques attempt to hide information inside objects, especially in images.

Fingerprints can either be inserted or discovered, and the insertions can take the form of additions, modifications, or even selected deletions. Fingerprints can occur on physical objects or on data. Examples of fingerprinting in action illustrate these concepts. Most consumer goods come with a unique identifying number, such as the vehicle ID number on a car. Detectives routinely match typed characters with a specific typewriter, or a fired bullet with a specific weapon. Businesses may place similar advertisements in different markets with slightly varying return addresses, to determine the market yielding the best response. Mapmakers insert slight deliberate variations from reality to identify copiers.

Any object that might be misused needs a fingerprint to identify the object's owner after misuse. For identification to succeed, an authority must record the fingerprint along with an ID of the owner. The recording might take place at the time of sale or of delivery. A method from the previous chapter would then identify the individual taking charge of the object. Imagine the uselessness of identifying the purchaser of dynamite employed in a crime as ``John Smith, address unknown."

Fingerprints should be hard or impossible to remove, as dictated by the particular application. For example, using different post office boxes for alternative return addresses provides a perfect fingerprint: the box used reveals the source of the address. In some cases one can have a perfect fingerprint like this, and in others one can at best make it difficult or expensive to remove the fingerprint. Thus a car with its vehicle ID number stamped onto half the parts and etched onto every pane of glass becomes more secure from theft.

Fingerprinting ought to be ubiquitous. Society can and should do a better job of tracking objects, especially stolen or valuable objects. Law enforcement agencies already keep lists of stolen goods or of items left at pawn shops. Sometimes the lists are computerized, and sometimes there is cross-checking of lists. Notice that the items need fingerprints to identify them; the lists record these fingerprints. The lists should be all-inclusive and coordinated. Initially, such measures might help recover what was stolen and help catch the thieves, but in time the use of these measures would be a powerful deterrent. A television set stolen in New York could not be pawned in California. Stolen goods taken across national boundaries pose another problem that cooperation between the involved countries can solve.

Fingerprints on Physical Objects

Bullets illustrate many issues about fingerprinting physical objects. When a bullet goes through a gun barrel, it acquires characteristic rifling marks from the barrel. These are fingerprints that can match a bullet with a gun. In this way one can associate the same unknown gun with more than one crime; with the gun in hand, one can tie this gun to various crimes.

As a first step, laws should require determination of the rifling marks of each gun before sale and require the recording of these marks along with the identity of the gun purchaser and the serial number of the gun. Then it would be natural to enhance and expand these rifling marks, to make them show up more prominently and to identify the gun uniquely. The ideal would provide on each fired bullet a fingerprint that identifies a unique gun, traceable to an individual. (Stolen guns present an additional problem discussed below.) It would take considerable research to determine how well this could work in practice, and the comparison of fired bullets is so inexact, depending on the condition of the bullet and other factors, that such fingerprints will never be completely reliable.

As a second step, manufacturers should fingerprint every bullet. The fingerprints could be the same for each batch or box of bullets sold as a unit to an individual. One method would add trace amounts of various elements to the bullet's material. It would then be feasible after the fact to analyze the bullet's composition and thereby read its fingerprints. These fingerprints would survive an impact that destroyed the shape and rifling marks on the bullet. As before, laws would require the presence of fingerprints and a record of the purchaser, holding the purchaser responsible for the use of these bullets.

As additional steps, one could fingerprint batches of lead or other materials used to make bullets to help trace those who make their own bullets, and one could add a volatile substance to bullets that could be detected in the air, say, at airports. Similarly, researchers could find ways to identify guns from a distance. The theft or illicit resale of guns and bullets creates another problem. A fine or even forfeiture of escrowed money would work in such cases. Eventually, society can manufacture high-tech guns that will not fire when stolen, as discussed later in this writeup.

Adding different mixtures of trace elements to the material used to make successive batches of bullets need not in principle be much additional cost. The record-keeping would be more significant and would need to be computerized. Note that much of the work and expense would only be necessary in case of an investigation into a crime.

Some readers, particularly ones outside the United States, might find this discussion wrongheaded. They might wonder why the proposal is not to regulate and limit the sale of guns and bullets themselves. Such regulations would be a benefit, but even then, bullets will still be sold, and the fingerprints would still be useful.

Pollution gives another example of fingerprinting in action. Laws should require fingerprints on all industrial waste. There would be requirements that suppliers of raw materials to industry add trace amounts of identifying elements or compounds to those raw materials. Thus a chemical company would have to supply solvents in fingerprinted form. There would be opportunities for cheating or bribery, so unannounced inspections and controls would be needed. A dishonest official might even insert another company's fingerprints, so companies would want to check for themselves that the proper fingerprints are present. At each stage of a complex process, the industry would add additional fingerprinting materials. In the face of environmental pollution, the pollutants themselves would indicate their source and even the percentage involvement of several industries. Pollution with no fingerprints would uncover cheating. Notice that these techniques attempt to catch polluters after the fact, to stop them and perhaps punish them. In another approach, agents could detect and halt pollution as it starts to occur --a better way.

Suppose a hit-and-run driver leaves part of his car and a paint sample at the scene. Then suppose an investigation reveals that only 100 cars of this type, with its special paint, had been sold in the U.S.~~ Authorities narrow the search to just a few cars registered near the accident and are able to find the offender. The public would welcome the diligence and luck of the investigators, but society could make this the norm by requiring coded particles (or another identifying residue) in all cars, particles that would remain after an accident to uniquely identify the car.

In 1996, the U.S. Congress passed the Antiterrorism and Effective Death Penalty Act, which called for the study of tagging materials to add to explosives to make them easier to detect before an explosion and to allow identification of the source of the bomb materials after an explosion. These methods are promising but need further research; progress is blocked in the U.S. at present by various groups such as the National Rifle Association. In addition there were proposals for additives to ammonium nitrate to neutralize its explosive properties -- methods which do not appear promising. Other examples of physical fingerprints include the serial number on currency bills in circulation.

Society should use the fingerprint to track all currency, eventually tracking all electronic money as well. Admittedly, tracking all money is more controversial than tracking bullets or dynamite, but money laundering is another crime that such tracking would address.

Endangered animals can also be fingerprinted, as is the case with badgers in the U.K., where the popular but illegal sport of badger-baiting faces badgers equipped with a waterproof coating containing a unique set of chemical tracers which can even identify those who handle a marked animal. The U.K. even has a database of shoe imprints.

These examples illustrate what one should do with every hazardous object or material, with anything valuable that might be stolen or destroyed or misused, and with many other objects as well: insert or identify fingerprints; record them; and keep centralized records and correlate the records. Often multiple fingerprints for the same object are desirable -- identifying several characteristics already present, and adding identifying features, including residue that would remain after misuse as well as a volatile residue that instruments could detect during misuse. In the case of goods for sale that might be shoplifted, some manufacturers now insert standard tags that will trigger an alarm when an item is taken from a store without deactivating the tag. These tags lie deep in the item itself and are more difficult to remove than common anti-theft devices. Laws should require such tags in all dangerous or valuable objects. In collaboration with the fingerprinting (or sometimes independently), software agents, monitors, and sensors should track objects, recording and saving this data.

Fingerprints on Data

Here ``data'' refers to anything machine-readable. Examples include English language texts, program source, executable files, files of raw data, database files, digital pictures, and digital video. All such objects allow inexpensive fingerprint insertion, which society should routinely require.

Suppose you are a staff member for a U.S. senator, working with one of the senator's committees. You have a confidential memorandum ready for distribution to the committee. Recent events indicate that a senator on the committee either leaks such documents himself or has a leak in his staff. You could fingerprint the memo by preparing a unique version for each senator. Each version would have tiny variations throughout, say in the typefaces used or in the spacing -- not noticeable unless one is looking for it. Now if a senator leaks a photocopy of a portion of the memo, an analysis would determine the leaker, assuming the fingerprints are throughout.

Once the word got out, any leaker would know that he must retype a memo before leaking it. You can foil this new strategy by making small *textual* changes in each version of the memo. It is easy to find places in a text that can be worded in several ways. Then one can employ different combinations of these variations for the different senators. As a bonus, this method can be automated to allow, under direct computer control, fingerprint insertion, recording of the memo version and the person to whom it is distributed, and determination of the version leaked in case a portion of the memo appears in the press. For readers familiar with computer jargon, the method could be the following: first determine individual points of variation in the text and then use a pseudo-random number generator, with the ID of the person receiving the memo as seed, to determine which individual variation is used at each stage.

A popular novel, *Patriot Games* by Tom Clancy, described exactly this strategy (referred to as the ``canary trap'' in the book).

> Each summary paragraph has six different versions, and the mixture of those paragraphs is unique to each numbered copy of the paper. There are over a thousand possible permutations, but only ninety-six numbered copies of the actual document. The reason the summary paragraphs are so -- well, lurid, I guess -- is to entice a reporter to quote them verbatim in the public media. If he quotes something from two or three of those paragraphs, we know which copy he saw and, therefore, who leaked it. They've got an even more refined version of the trap working now. You can do it by computer. You use a

thesaurus program to shuffle through synonyms, and you can make every copy of the document totally unique.

In time, potential leakers will discover this approach also, and realize that they must paraphrase any leaked memo. They fall back to leaking the information in the memo. Now what can one do in an attempt to fingerprint the memo? The method from the Clancy novel fails completely in this case. It sounds extreme, but one can change the *information* in the memo: altering facts slightly, adding pieces, leaving pieces out. The challenge is to find facts to alter without changing the thrust, meaning, and completeness of the memo, and without calling attention to the fingerprints. In this environment potential leakers know that they must alter the basic information and facts of any memo they leak in order to escape detection.

Assume one carries out alterations and sends fingerprinted data to a number of individuals. When the data returns after the leak, a statistical analysis can test the hypothesis that each individual is the source of the leak. The knowledgeable opponent will counter by altering values himself before leaking them, perhaps by rounding them. This tactic will not work indefinitely for the opponent, however. First, the opponent's values have already been altered within acceptable limits. If he alters them much more, the leaked data will be too inaccurate to be of use. More significantly, no matter how much the opponent further alters the data before leaking, given sufficiently many leaked values, a statistical procedure will correctly identify him as the opponent with any desired degree of confidence. An opponent who continues leaking cannot protect himself from eventual detection.

Consider a specific scenario. Suppose the U.S. is to build a new line of tanks and trucks for use by its allies in Europe. Early in the project, each country wants to know the width, height, and weight of the various vehicles. (They may wish to know which roads and bridges the vehicles can travel on.) Suppose further that one of these countries is the source of a leak to an opponent (``the enemy''). The U.S. could supply each country with data altered within an acceptable range, since one would want leeway in the measurements anyway. After a leak, if the data returns to the U.S. somehow, the U.S. could try to identify the leaker. If the returned data was not further altered, this data itself would identify the country of the leaker immediately. But even if the opponent further altered the data, beyond the initial fingerprints, the hypothesis testing mentioned above would eventually pinpoint this opponent. The smaller an opponent's alterations, the quicker he would be identified, but larger alterations make the data less valuable, since it is less accurate. The opponent faces a dilemma: the more valuable his data, the more quickly he is caught, and he cannot avoid eventual detection.

Subtle Fingerprints

The fingerprinting process sends multiple copies of data out into the world. If a copy comes back, even an altered copy, the fingerprints may allow one to deduce the source of the returned copy.

It may be possible to deduce from alterations in the returned copy something about the path through the world that the original data took. For example, if 500 miles is sent out, and 497 miles returns, one might suspect that the 500 miles was converted to 804.675 kilometers, rounded to 800 kilometers, converted back to 497.095 miles, and finally rounded to 497 miles. (Different agents are rounding by different amounts, and so leave the fingerprint.)

Along similar lines, a news agency recently reported that a giant floating iceberg was 656 feet 2 inches thick -- a precise-sounding measurement that in metric units is exactly 200 meters, the true approximate figure. The same report said that the iceberg was the result of a 36.5 Fahrenheit temperature rise since the 40s. But the actual rise is a 2.5 Centigrade increase. A reporter converted the temperature rather than the increase, which should have been given as 4.5 Fahrenheit.

A fingerprint left by the Unabomber gives a final example, where he wrote in his ``Manifesto":

185. As for the negative consequences of eliminating industrial society -- well, you can't eat your cake and have it too. To gain one thing you have to sacrifice another.

The phrase about eating and having cakes also appears in an early letter of the suspect in the case. American reporters termed this a ``twisted cliche'' and said it was ``turned around.'' Its presence in writings by the Unabomber and the suspect provided a link between the two. Current American usage expects to hear the words ``eat'' and ``have'' reversed, so it is surprising to find that the *Oxford English Dictionary* lists only the Unabomber's version of this saying. Other dictionaries of idioms (British and American) list both versions. It now seems likely that the Unabomber used this as part of his normal English, and not as a clever reversal of a standard phrase. He may have inadvertently left this subtle fingerprint because he was not familiar with the modern American preference -- after all, he seldom talked with people and had no electricity for radio or television. This fingerprint supports the *verification* of an identify after the fact. Imagine carrying out an earlier *identification* based on similar fingerprints, using an automated search through vast amounts of published materials. Such identifications will be increasingly feasible as more library materials become machine-readable. The same process occurs when a literary researcher tries to decide whether a ``lost'' play was written by a particular playwright or when searching for plagiarism in published material.

Similar techniques will check if computer students copy or exchange programs for an assignment, as well as checking for other academic plagiarism. Software is readily available to compare two programs in a variety of computer languages or even to compare two term papers in English. If a whole class hands in programs, the instructor can check all possible pairs for similarities. The plagiarism detection software is subtle and hard to deceive; it easily copes with the common tricks of students who copy programs: change program identifiers, rewrite all the comments, reorganize the program in a new style, and arrange elements in a different order. As for detection of plagiarism in ordinary English writing, the grand opera singers of detectors are two employees of the National Institutes of Health, Walter Stewart and Ned Feder, who started out looking for scientific fraud and ended up checking for English text plagiarism.

Crime-proof Hardware

At this point the discussion will move beyond fingerprinting, from methods that identify misuse, to those that will not allow misuse.

For example, if a thief steals a fancy radio/CD player from a car, he may find that it no longer works when removed. This is a simple case of a piece of hardware that does not permit successful theft.

Most consumer goods are getting electronic innards and are developing higher intelligence -- from cars to refrigerators these machines are capable of more sophisticated actions -- even of adaptive behavior. In time, there will be enough extra computing capacity in electronic objects in the home or workplace so that they can be programmed to work as intended and in the assigned environment, and not to work if there are any changes, such as removal from the environment. For example, appliances could repeatedly verify that they are still in the proper house, using cryptographic authentication techniques. Such verification can be made foolproof, but with current systems this would substantially drive up the price of the appliance. Future appliances will have computing power to spare for this task. Appliances may broadcast their position, as with some stolen laptop computers that now try to ``phone home'' at random times to give their current location.

It must not be inexpensive to replace this module that controls appliance operation. Many of these appliances of the future will consult their brain before doing anything, and these brains will be a significant part of the appliances' cost. Thus the problem of theft and re-engineering should lessen also.

Some software vendors require that the authorized user retrieve a special enabling password or code, needed to run the software. (They may also require a hardware device inserted in the back of the computer.) Such software can be copied and backed up, but it does not run without the special password. It is even possible to use an identifying hardware ID within a specific computer and supply a password that will only work with that specific copy of the software and that specific computer. Take the software and the password to a different machine, and it will not run. Cryptographic techniques can create passwords that users are not able to break.

In the same way, manufacturers of microprocessors may one day protect against theft by requiring a special password that is tied to the specific microprocessor chip and to the specific computer. When the *hardware* is started up by the user, it could first insist on accessing the microprocessor vendor by phone or over the Internet, to let this vendor verify that the chip was not stolen.

Society could use similar techniques to make automobile theft nearly impossible. If an unauthorized person tries to start or even enter the car, the car's computers could be programmed to lock up in a way that would require resetting by a dealer.

A cartoon image showed a parking meter spewing hot tar over the car of a hapless motorist who violated the time limit. But a serious Philadelphia inventor has a real parking meter which resets itself when a car leaves. It then demands fresh money from the next car. The meter, equipped with infrared sensors, does not add time for inserted money if the meter has expired and the car has not moved. The meter also keeps track of the expiration time, to counter claims that the meter had just run out. This prototype meter is an early example of the new line of intelligent autonomous machines. Whether or not this particular meter is successful, similar machines will soon be available in many application areas.

Now move the level of sophistication yet one notch further up, from hardware that will not allow theft, to hardware that directly disallows the commission of a crime. A simple first example illustrates the idea: In some societies, such as Singapore, laws require the flushing of toilets after use, with a stiff fine for not flushing. Many new public toilets in the U.S. sense that a user has departed and flush themselves automatically, making it impossible to carry out the ``crime" of ``failure to flush."

As another example, if the U.S. society is unwilling to restrict the sale and ownership of guns, it could create guns that only the owner would be able to fire. An implementation might involve verifying the owner's hand geometry or fingerprint before firing, or might use a special enabling ring the owner wears. Such a system is not much different from a reliable trigger lock, but an owner can leave the trigger unlocked, while the other systems would reset themselves after each use. Guns could also have disabling mechanisms that would prevent them from discharging in public areas, since a gun owner ought to buy a gun to protect himself in his home, not to shoot at someone in an airport or a store.

Traitor tracing

Traitor tracing is a copyright infringement detection system which works by tracing the source of leaked files rather than by direct copy protection. The method is that the distributor adds a unique value to each copy given out. When a copy of it is leaked to the public, the distributor can check the value on it and trace it back to the "leaker".