computers are bad

You are receiving this facsimile because you signed up for fax delivery of this newsletter. To stop delivery, contact Computers Are Bad by email or fax.

https://computer.rip - me@computer.rip - fax: +1 (505) 926-5492

2023-06-30 calling in the alarm

I currently find myself on vacation in the Canadian Rockies, where internet is hard to come by. But here's something short while I'm temporarily back in the warm embrace of 5G: more about burglar alarms. I recently gave a presentation on this topic and I'll probably record it for YouTube when I'm back home, but I think the time has finally come to write a post on a specific and niche element of intrusion alarms that I find particularly interesting: alarm reporting protocols.

Let's briefly recap the architecture of a typical intrusion alarm. An intrusion alarm system (and essentially the same goes for fire alarms as well) consists of a controller that monitors sensor zones. When the controller detects that a sensor has been violated while the system is alarmed, it enters the alarm state. Once this happens, it reports the alarm to a Central Alarm Receiver (CAR) at a Central Alarm Station (CAS). The CAS is then responsible for dispatching emergency services (or private security) to respond to the alarm.

Early intrusion alarms relied on dedicated wiring to report to the CAS. In a major city, the municipal government often operated alarm infrastructure for fire (most commonly a Gamewell fire telegraph system), but for security private companies were more common. One of the biggest names in alarms to this day, ADT, clearly reflects this heritage: the acronym stands for American District Telegraph. The company originally provided stock quotations over their private telegraph networks, but later made a lot of money using their telegraph infrastructure for alarm monitoring.

The oldest alarm systems reported to the CAS using the "polarity reversal" scheme, which could be used either over privately owned wiring or a leased telephone line specified as "dry" (meaning that the telephone exchange did not apply battery power or dial tone). The burglar alarm controller normally put a voltage on this pair. When the alarm was triggered, the polarity of the voltage was reversed. In the CAS, the change in polarity caused a metal flag held in place by an electromagnet coil to drop down, informing the CAS operator that an alarm had occurred. The major advantage of the polarity reversal scheme is that, with an appropriately designed system of coils around the flag, the CAS operator could tell whether the polarity reversed (an alarm) or the voltage went away entirely (a trouble). During an armed state both of these conditions merited a response, but knowing the difference was useful for troubleshooting the alarm and reporting infrastructure.

The concept of applying a voltage to the monitoring line at all times is a simple form of *supervision.* It means that any interruption in the connection can be detected. Supervision is one of the most important concepts in technical security [1]: supervision is a set of techniques that allow an alarm system to ensure that it has not malfunctioned, been damaged, or been intentionally tampered with. Supervision is the key thing that differentiates a "life safety system" from other types of electronics: fire alarms (for safety reasons) and intrusion alarms (for security reasons) have to be highly tamper- and failure-evident.

Polarity reversal is a simple and effective scheme for alarm reporting, and it's still used in some institutional environments where the CAS is on-site and the alarm wiring from each building to the CAS is already in place. It suffers a major limitation, though: there is no support for multiplexing. In other words, every alarm system has to be connected to the CAS by its own dedicated pair of wires, or the CAS will not be able to differentiate where an alarm came from.

In the early 20th century, telegraph technology was producing the first digital communications protocols. Pulse-based telegraph systems could be connected to mechanical receivers that used clockwork mechanisms to count pulses, differentiating which signal had been received. For example, while early fire telegraphs required station personnel to interpret the paper tape manually, by the 1930s some fire stations were equipped with telegraph receivers that could count out pulse patterns to identify certain box calls and sound a bell in the station automatically. Similarly, railroads began to use "coded" signal circuits where multiple signals were connected to a shared bus (wired in parallel) and counted pulses to recognize addressed commands. These geartrain telegraph receivers, steampunk by modern standards, were the genesis of digital communications. The same methods were applied to intrusion alarms.

While the Gamewell system was designed for fire reporting, some areas used Gamewell telegraphs for intrusion reporting as well, and many proprietary intrusion reporting systems were substantially similar to the Gamewell design. A Gamewell fire box had a "hook" that, when pulled, would compress a spring. The spring then released its force into a clockwork mechanism that rotated a notched wheel---emitting a pulse onto the telegraph line every time a notch passed a switch. The internal design was fairly similar to a telephone dial, but with a different interface since the pulse train that was sent was fixed for each box, determined by the position of the notches on the wheel.

Originally these Gamewell boxes were mounted streetside where passersby could pull the hook, but Gamewell systems proved surprisingly durable and were in service well into the era of electronic fire alarms. Gamewell sold boxes which were electrically activated, meaning that they could be wired to a fire alarm system so that the "hook dropped" automatically when the fire alarm sounded. Look around in back alleys of a city and it is not unusual to find lonely Gamewell boxes still mounted on the backs of buildings, often a legacy of when they functioned as the reporting system for the fire alarm.

This basic design became common in life safety and intrusion alarms. In modern terms, the electrically-activated Gamewell box functioned as a "communicator" to report the alarm to a CAS when activated. Many intrusion alarms used very similar designs, with an electromechanical telegraph communicator triggered by a voltage output from the alarm controller. During the era of telegraph systems, central alarm receivers (CARs)---the equipment that actually receives the alarm signal---became progressively more complex, producing paper tape logs of all calls in addition to activating appropriate signals to CAS operators based on the received code.

One might wonder how supervision worked with telegraph systems, since it's not possible for the CAS to simply monitor for the presence of voltage when multiple alarm controllers share the same lines. Telegraph systems introduced period supervision: the alarm periodically sent a signal, and the CAS interpreted the lack of any message over a certain time period as being indication of trouble. Early periodic supervision was actually very simple, as it was common for early alarm systems to report any entry (authorized or not), as well as arming, to the CAS. In a common bank vault alarm, for example, the CAS would be informed when the vault was closed at night, and when it was opened in the morning, regardless of armed/disarmed status [2]. Operators at these CASs often had a checklist of sorts, where they expected to see each bank vault being opened in the morning. If they didn't, it likely indicated trouble with the reporting system.

Later on, timers were used to send "supervision reports" at configurable intervals. For dedicated alarm wiring, supervision might be very frequent, perhaps multiple times per hour. As telegraphic alarm reporting systems matured, though, dedicated wiring started to fade away. This wasn't a total abandonment of dedicated alarm reporting infrastructure, which can still be found in commercial areas of some cities and is especially common on institutional campuses where dedicated fiber lines might be run for fire and burglar alarm monitoring. But during the '60s and '70s, burglar alarms caught on in the home, and for home users a leased telephone line (or the cost of building out private alarm wiring) would significantly drive up the cost. Just about every home had a telephone line, though, and telephones were already the most common way to reach emergency services.

The vast majority of home burglar alarms, until perhaps the last 15 years, were installed with telephone communicators. Like the old Gamewell boxes these were dedicated modular devices, but they were reduced over time to a single PCB mounted in the alarm controller cabinet. The telephone communicator would be connected to a phone jack so that, in the event of an alarm, it could dial a call to the CAS and report the event.

There's a surprising amount of nuance to telephone communicators, perhaps unsurprisingly since they were in common use for a period of some fifty years. First, many telephone communicators could be configured for use with a dedicated telephone line (with the advantage of much more frequent supervision reporting without tying up the phone) or for use on a line shared with telephones (a much lower-cost option, typical for residential installations).

Sharing a phone line with telephones posed a problem, though. Say a fire or breakin (especially the activation of a panic button) happened when a phone was off-hook. Fire alarm codes in particular required that central reporting still work in this situation. The solution is simple but also surprisingly obscure: a special telephone jack.

Many homes built in the late '60s through the '80s, the golden era of residential intrusion alarms, include an "RJ31X" or burglar alarm telephone jack. It's usually found in a strange place for a telephone jack, like the master bedroom closet [3]. The RJ31X jack (this is actually a technically correct description of the jack, unlike most modern uses of the "RJ" identifications) is an 8P8C modular connector (same as Ethernet) with two phone lines terminated to it. One phone line goes directly to the telco via the network interface device (demarc), while the other goes to the house's internal telephone wiring. An RJ31X jack has a special shorting bar inside the jack housing that bridges the outside and inside phone lines together. When a plug is inserted, it pushes the shorting bar out of the way, disconnecting the inside and outside phone lines and routing them to the alarm communicator instead.

The alarm communicator now has complete control of the household phone line. All telephones in the house are connected to the "inside" wiring. Normally the communicator connects the inside and outside lines together the same way that the jack's shorting bar had, but when an alarm occurs, the communicator's "line seizure relay" disconnects the internal phone wiring from the telco. After waiting a moment for the telephone exchange's line card to reset the line to its on-hook state (assuming a call might have just been cut off), the alarm communicator can go back off-hook and dial its own call.

RJ31X jacks are now mostly a thing of history, but like many parts of history they sometimes protrude into the present. It's not unusual for scratchy, intermittent phone lines to be tracked down to an RJ31X in a closet with a loose or dirty shorting bar. Today the jack is usually removed rather than fixed.

Once the alarm communicator has dialed a call and wait for the far end (a CAR) to pick up, it has to transmit the details of what has happened. There is a confusing range of different standards here. Older communicators often used DTMF, sending a series of digits. A common example is "Contact ID," developed by Ademco and standardized by the Security Industry Association (SIA). When the CAR picks up the phone line, it is expected to send a pulse of 1400Hz, silence, and then a pulse of 2300Hz. This informs the alarm communicator that it has indeed reached a Contact ID endpoint, an important issue since some residential alarms especially were also configured to call the homeowner directly and supported other methods like voice recordings for non-Contact ID endpoints. After hearing the pickup tones, the alarm communicator sends 16 digits by DTMF. After the 16 digits, the CAR responds with a 1400Hz tone to confirm receipt.

The 16 digits consist of a 4-digit account ID (used to identify the specific alarm), a 2-digit message type (typically 18 which identifies the SIA Contact ID standard), a one digit event type, and three digits of event detail that typically indicate the type of zone that was triggered. The message ends with a 2-digit partition ID (used for multi-partition alarms such as in multifamily housing), a 2-digit zone ID, and then a checksum digit.

The 4-digit account ID might be a bit surprising. Some of the older alarm protocols really didn't support that large of a namespace, and early on most CAS were local operations with relatively small customer counts. As CAS became increasingly monopolized, many CAS had to have a large number of incoming phone lines so that they could differentiate alarms by the phone number they were configured to dial as well. One motivating factor in the development of more advanced alarm reporting protocols was to increase the address space and make reporting calls shorter, both of which allowed for more alarms reporting to the same phone line.

Indeed, other telephone reporting protocols used more advanced digital methods similar to data modems. For example, another SIA-standardized protocol (often referred to just as "SIA") uses frequency-shift keying compatible with the Bell 103 modem. SIA messages are short ASCII sequences with a several-character preamble identifying the protocol and giving an address for the intended receiver (a way to allow multiple logical receivers to share phone numbers), an account ID of up to 16 characters of hexadecimal, and then event and zone IDs. Once again, the message ends with a checksum to confirm correct receipt, and the communicator will retransmit if it does not receive an acknowledge tone from the CAR.

The use of digital modems over telephone lines starts to sound a lot like dial-up internet, and you might wonder if intrusion alarms used similar techniques. The answer is yes, but in several different ways.

One of the most interesting innovations in alarm reporting was a system called DCX, for Derived Channel MultipleX. I think I've mentioned previously that "derived channel" is a common term in the telecommunications industry for the use of any technique to get an additional data channel out of a medium. The most widely known example of a derived channel on the telephone network is DSL, and indeed DCX works very similarly.

A DCX communicator is connected to the telephone line much like a conventional telephone communicator, but it doesn't dial. Instead, it sends high-frequency FSK messages regardless of the state of the telephone line. Much like DSL, the FSK is outside of the normal voice passband of the phone system, so telephone calls won't interfere with it. That said, DCX communicators weren't completely inaudible like DSL---they used low enough frequencies that, if a DCX communicator happened to send a message while you were on the phone, you would hear it. This was a much more likely situation because DCX took advantage of the lower connection setup cost from not having to dial a call: the biggest advantage of DCX was significantly more frequent supervision, with DCX communicators reporting status to the CAS as often as every twenty minutes.

DCX signals can't pass through the telephone network, so just like DSL, DCX requires the customer's telephone line to be directly connected to a DCX receiver. CAS that used the DCX technology arranged to install receivers in telephone exchanges, and these DCX receivers used leased lines to send real-time reports to the CAS. Altogether the system was fairly elegant, but the need for specialized phone exchange equipment meant that DCX was only ever available in certain cities. It was mostly popular with commercial customers, where insurance companies often required frequent supervision intervals that made it impractical for the alarm to share a phone line with the normal business phones.

DCX has fallen out of use, but I can't resist the urge to share a charming detail of the implementation. The DCX receiver system was implemented as software on a normal IBM-compatible PC, but UL standards for burglar alarms require a hardware failsafe on all central alarm reporting and receiving systems. DCX's solution is the kind of wonderful PC-era computer accessory you rarely see today: a small box that went inline with the computer power supply and connected to a serial port. The DCX software pulsed a line on the serial connection (I suspect, from experience with this kind of thing, not even a data pin but likely a control pin), and the box functioned as a watchdog timer, probably using a simple delay relay. If too long elapsed without a pulse, the accessory box cut power to the PC for a few seconds. This is, of course, mundane, and today we are still configuring switches to cycle PoE on ports when ping checks fail. What really delighted me about it is the fact that this device is the only custom hardware involved in the receiver system, so the manual really sells it as a major DCX innovation. It has a blinking LED, so you know it's working.

DCX's major limitation, and what killed it off today, is the fact that it functions much like an internet connection but without the benefit of providing IP transit---or even coexisting with it, since DCX and DSL cannot be combined on the same line due to near-overlapping frequency ranges. By the mid-2000s, alarm communicators were making the transition to the internet.

Most modern intrusion alarm systems use SIA DC-09, the SIA-standardized protocol for alarm reporting over either TCP or UDP. DC-09 optionally supports AES encryption, and it's reassuring to know that the connection is optionally secure. DC-09 is very similar to the SIA FSK protocol in terms of the message structure; no surprise since it was designed in part for easy implementation in existing communicator and receiver systems. DC-09 is extremely simple and not really all that interesting of a protocol. The alarm communicator sends a single packet containing the message and waits for an acknowledge packet from the receiver.

What is interesting about DC-09 is the substantial benefit that intrusion alarms get from our modern world of pervasive IP connectivity. One of the key threat models to intrusion alarms has always been isolation. Going back decades, films have depicted burglars cutting the phone line before entering a house. It's well known that alarm communicators can be disconnected from the CAS, and with the long supervision intervals used by many alarm systems the disappearance of the alarm isn't likely to be noticed before the burglars have finished their work. The best protection against isolation is redundancy.

Alarm communicators that support multiple redundant connections are called *dual-path*, and dual-path communicators weren't common in residential installations until the internet era... there weren't many options besides a second phone line, and that would connect to the house by the same drop cable anyway.

Many burglar alarm systems today will attempt reporting events both through the homeowner's internet service and a cellular carrier, which makes it a lot harder for an intruder to isolate the alarm system. It also provides valuable redundancy for life-safety purposes, making it more likely that a fire or flooding alarm will be reported even if there has already been damage to the building or infrastructure in the area. IoT cellular service has gotten so inexpensive that there's no reason not to have dual-path reporting in most alarm installations.

And the fact that internet connections are inherently multiplexed yields a big advantage when it comes to supervision. One of the big problems with consumer burglar alarms was their infrequent supervision intervals... since a supervision report tied up the household phone line, many alarms were configured to never send supervision reports at all. This meant that a cut phone line or, perhaps more likely, a malfunctioning communicator could go unnoticed until it was too late.

Unfortunately, central alarm reporting suffers a lot from its legacy. Modern alarm systems often report surprisingly few events and supervise surprisingly infrequently, considering that the IP connection is inexpensive and has very little contention. One of the issues that I most often complain about is the rarity of disarm supervision.

If a burglar enters a house and is able to locate and destroy the alarm controller before it reports the alarm (which is typically after the entry delay of 30-60 seconds and a post-alarm reporting delay of 20-60 seconds), the CAS may never know that there was a problem. Hiding the alarm controller and surrounding it with immediate zones is the traditional solution to this problem, but an obvious modern one (that has been used in high-security environments going back decades) is to have the alarm report that it has begun the entry delay immediately, and then report when it has been disarmed. If the CAS receives an entry delay message and then doesn't receive a disarm message within a minute or so, it can assume that the communicator has failed and treat the situation as an alarm. Many intrusion alarms and receivers support this functionality, but vendors have inconsistent names for it (when they advertise it at all) and it's not often enabled. Hesitation to use disarm supervision is understandable when each message requires dialing a phone call, but today we have the internet and packets are cheap.

Burglar alarms aren't, though, and that's part of the problem. Consumer interest in burglar alarms has decidedly moved to the low-end, with even cheap wireless systems failing to produce the sales volume of burglar alarms in the '70s. Consumers have little awareness of the practicalities of alarm reporting, and alarm vendors advertise their smartphone apps and home automation features but barely mention the actual security properties of the system. Unsurprisingly, those properties are usually poor, and a lot of burglar alarms today have limited value against an informed actor who has planned ahead. That's the thing, though---not many home burglars plan ahead, and so the most primitive of alarm systems can do a pretty good job.

[1] The terminology in security can be confusing, especially since "cyber security" has come into the landscape and coopted a lot of the existing terminology. "Technical security" tends to refer to more "old-school" forms of electronic security, particularly intrusion detection and technical counter-intelligence.

[2] This concept is still common in high-security institutional environments like

military installations, where many alarms have no sense of "disarming" and will report any entry at all to the CAS. Authorized users are expected to contact the CAS either before or immediately after entering to identify themselves and explain their purpose. This can be a much more secure arrangement since it allows the CAS to audit every entry against work orders or duty assignments, discouraging insider theft.

[3] For various reasons, mostly for "smash-resistance," it’s a good idea to install the alarm controller in a somewhat hidden location. Of course most home builders and alarm installers were not very creative, and so the hidden location is virtually always a closet, most often the closet of the master bedroom. Fortunately, due to the longstanding architectural principle of private-public space separation, the master bedroom closets is often one of the furthest points from an entrance to the house. Unfortunately, every burglar knew that’s where the alarm controller was most likely to be, which made the master bedroom window an appealing point of entry. This highlights the importance of "immediate" alarm zones on nonstandard entry points like windows and motion sensors in rooms without exterior doors.