

Taking Sensor Networks from the Lab to the Jungle

Andrew S. Tanenbaum, Chandana Gamage, and Bruno Crispo

Vrije Universiteit Amsterdam

STOPPING the flow of illegal immigrants, smugglers, and terrorists across national borders is a hot topic from the deserts of Arizona to the jungles of Sri Lanka. Not surprisingly, people have suggested using technology, in the form of sensor networks, to aid in detecting these unlawful intruders. Sensor networks are also candidates for solving a variety of other social, military, and environmental problems, including secretly monitoring enemy activities on a military battlefield located in inhospitable terrain; and detecting wildfires in a densely forested area. In all these cases and others, people have proposed using manned or unmanned aircraft to drop battery-powered sensors equipped with built-in radio transmitters to detect and report potential incidents.

Because these scenarios all involve security-sensitive applications, a great deal of attention has been paid to the cryptographic protocols used on the radio channels. Other research has focused on network issues such as MAC layer protocols, message routing, data aggregation and node localization. However, in our view, insufficient attention has been paid to the systems aspects of such networks, so it is these aspects that we will discuss in this column.

All of the above scenarios share some characteristics that make them different from traditional local-area and wide-area networks. In particular, electricity is absent (requiring the sensors to be battery powered), a wire-line network is absent (requiring the sensors to communicate using an ad hoc wireless network), roads are absent (requiring the sensors to be air dropped), and large numbers of sensors are needed (requiring the sensors to be low cost). The large areas to be covered and the need to do detection 24 hours per day usually preclude airborne or satellite surveillance, making ground-based sensor networks attractive.

THE SENSOR NETWORK SOLUTION

As most of the applications envisaged for these very large sensor networks are required to operate under hostile conditions, it was necessary for researchers to develop security mechanisms to provide authenticity, integrity, and confidentiality protection for messages transmitted over the wireless network. Due to the low speed of the sensors' CPUs, small amount of RAM, and large amount of power needed to transmit each bit, researchers have worked on developing tailor-made solutions. Many security papers propose cryptographic methods for countering attacks such as injection of false data, modification of data, eavesdropping on wireless communication, cloning of nodes and node capture. Other research on power efficient MAC layer algorithms and routing protocols, distributed and cooperative schemes for node localization and efficient data aggregation techniques have all shown remarkable progress.

However, achieving reliable and secure communication on the radio channel is not the whole story. Systems aspects must be considered as well. We began our investigation of sensor networks by actually acquiring some commercially available sensors and measuring their radio range, something that has not been widely reported in the literature. We measured the radio range of a low-cost sensor on the ground in an open field to be about 7 meters; at a height of 1 meter, the range increases to about 35 meters. Interestingly enough, in a forested area the range is comparable to an open field, primarily due to multipath signal propagation (Gamage et al., Security for the Mythical Air-dropped Sensor Network, Proc. ISCC '06, pp. 41-47, 2006). Finally, indoors in a long narrow corridor, the range on the ground increases dramatically from 7 meters to 35 meters (and 42 meters when the sensor is 1 meter high) as the walls act as an excellent

wave guide. Unfortunately, most applications of air-dropped sensors are outdoors. Our conclusion is that the sensors must be spaced rather closely to avoid gaps in connectivity.

Another important measurement is the operating range for the actual acoustic, thermal, humidity, infrared, or magnetic sensing element itself. Field experiments have shown the range of these sensing elements to be less than the radio range, necessitating even closer spacing of the nodes (Arora et al., ExScal: Elements of an Extreme Scale Wireless Sensor Network, Proc. RTCSA '05, pp. 102-108, 2005).

Starting with these measurements, we have examined three commonly-proposed scenarios and discuss some issues relating to each one below.

MONITORING A BORDER

The goal for border monitoring is to detect humans surreptitiously crossing on foot (especially at night). Along a land border of several thousand kilometers, illegal crossings could happen at almost any point and require that the full length be under surveillance, lest intruders discover the weak spots. Intrusion detection can be based on sound or vibration, which good quality sensors can easily monitor but only within a sensor range of less than 10 m for humans on foot. However, distinguishing between the illegal immigrant and the wily coyote moving in the underbrush is not so simple due to the limited CPU power and memory in the low-cost sensors.

The first problem is planting the sensors. Using our measured 7-m radio range, the 3100-km long U.S.-Mexican border would need over 440,000 sensors. Alternatively, the sensors could be placed on 1-m long weighted sticks, in which case only 88,000 would be needed, but getting all of them to land exactly 35 m apart and stick in the ground upright is not entirely trivial, especially in rocky terrain. To avoid breaks in the network,

many more sensors would be needed in practice.

Furthermore, since the battery life is typically 6 months, the air drop has to be repeated twice a year. This sensor forest will not be inconspicuous, although that may have some deterrence value. In our view, air dropping these sensors is infeasible and planting them manually twice a year is highly labor intensive, precisely what the air drop was supposed to eliminate. This brings into question one of the motivating factors for using large scale sensor networks: ease of deployment.

What happens when a sensor detects an intrusion? Intentionally setting off a land mine is illegal in most countries, so the only reasonable response is to send a wireless message to the nearest law-enforcement post so that an officer can speed to the scene and apprehend the intruder. Let us assume the border patrol posts are spaced at 20-km intervals. If the border patrol officer jumps in his vehicle instantly and travels at 100 km/hr to the sensor raising the alarm, it may take as much as 6 minutes to get there. By that time, an illegal migrant running at 10 km/hour could be anywhere in a semicircle 1 km in radius (assuming the migrant runs away from the border). Finding the person at night may not be easy, especially if the terrain offers places to hide.

For a supposedly unmanned operation, the costs are significant. To station an officer at each of the 155 posts along a 3100-km border 168 hours a week requires four shifts and thus a minimum of 620 people, plus 155 vehicles. When added to the cost of the semiannual sensor replacement (currently \$100 per unit, exclusive of deployment), the cost of monitoring the U.S.-Mexican border runs into many tens of millions of dollars per year. While not impossible, it is far from the vision of the people who think that a one-time drop of a few thousand sensors from an aircraft will do the job. With such costs, alternative approaches (such as building a long fence) have to be considered. Furthermore, the presence or absence of better cryptography is hardly the factor that will drive deployment, despite what some researchers appear to believe.

BATTLEFIELD OBSERVATION

Another widely cited application example in the sensor network research lit-

erature is the real-time monitoring of a battlefield. The objective is to secretly deploy thousands of sensor nodes onto a large geographic area and then have these sensors establish clusters of ad-hoc networks to sense and collect data on movement of personnel and vehicles. The first systems aspect that is not properly addressed for this application is the mechanism for secretly placing the sensors in a highly dangerous area. The often-stated solution is to air drop these sensors. If such air dropping of sensors were to be done using small UAVs flying at very low altitudes, the setting up of a large sensor network will require either a large number of UAVs operating simultaneously or few UAVs making multiple passes over the battlefield. Neither of these scenarios would provide the secrecy required for the battlefield monitoring application.

Here, too, radio range is a major issue. To monitor even a small area of 1 square km, a sensor network of over 10,000 nodes is required. In reality, this is such a small area that it can be effectively monitored by a single soldier from a safe distance using a good pair of binoculars, while the high density of the nodes on the ground makes it impossible to hide its existence from the enemy. For much larger areas, the cost of all the sensors becomes an issue. Consequently, the lack of attention to systems aspects of the problem can easily result in a solution that is completely unusable.

FOREST FIRE DETECTION

A sensor network designed as an early warning system for forest fires has systems aspects that are somewhat different from the two earlier examples. While border monitors need to respond within a few minutes, the response time for a fire detection sensor network using thermal or smoke detectors can be much longer. By careful placement of sensors closer to the areas of the forest most prone to fires such as hilltop areas subject to lightning strikes, the number of sensors required to cover a large geographic area can be substantially reduced. Our experiments have shown that multipath radio signal propagation that happen in an area with vegetation actually helps in providing a reasonable radio hop distance compared to open spaces.

However, the most important systems

aspect of a sensor network for forest fire detection is its lifetime. Unlike a border monitoring network with many detected events over a moderate lifetime or a battlefield observation network with a massive number of events over a relatively short lifetime, the fire detection network must operate for a very long period of time to detect its comparatively rare event. The nodes of a forest sensor network are subject to random failures due to battery exhaustion, as well as antennas being reoriented in the wrong direction by falling branches, curious animals, wind, etc. Since such networks relay messages hop by hop, failure of several closely spaced nodes may partition the network into noncommunicating sub-networks. As a result, when a fire is actually detected, the sensor discovering it may not be able to get its message to the nearest base station. Thus, overlooking the systems aspect of a forest fire being a randomly occurring event over a long time period can make the sensor network fail precisely when it is needed.

SO, WILL SENSOR NETWORKS GET TO THE JUNGLE?

Our tentative conclusion is that none of the commonly cited applications for air-dropped large scale sensor networks would actually work very well and certainly not be cost efficient with current sensors. More recent radio technology such as ZigBee or newer sensors (at various frequencies) with theoretical radio ranges 3-10x longer than ours are becoming available, but they are more sensitive to noise than ours and how well they actually work in practice is still an open issue. Still, many of the problems discussed above, such as the need for border patrol stations to be closely spaced, battery lifetime, etc. are not solved by increasing the radio range of the sensors. Furthermore, as we pointed out, the range of the motion, heat, etc. sensor itself is often the limiting factor, rather than the radio transmission range.

The real problem is that researchers have solved only those network problems easily simulated in the lab and have not addressed the system problems out there in the real jungle or desert. Nevertheless, there are applications of sensor networks that are clearly effective. The most obvious one is guarding the perimeter of a

fixed and valuable asset, such as a military base, oil refinery, or nuclear reactor. Another main application area for wireless sensor networks is as a replacement for wired SCADA systems that monitor industrial machinery where performance and safety are critical. Here the sensors can be accurately installed by hand, relatively expensive, and wired to the electricity grid, giving them greater range and eliminating the need for frequent replacements due to battery exhaustion. But even here, caution is required. When

the Soviet army had bases in Afghanistan during the 1980s, rebels threw live rabbits over the base fences to trigger the motion detectors. After several weeks of false alarms, the bases turned off the motion detectors, allowing them to be attacked unnoticed, a clear failure to take systems aspects into account in the design.

WIRELESS sensor networks also have applications, such as habitat observation of endangered species in

a small island or environmental monitoring on an active volcano. However, these applications tend not to require either a large scale network or military-grade security. In conclusion, while we believe that sensor networks have a bright future, we would suggest that researchers put more effort into dealing with the *systems aspects* of these networks, rather than only the *network aspects* of MAC protocols, message routing, data aggregation, node localization, security and cryptographic aspects.