Security for Autonomous and Unmanned Devices: Cryptography and its Limits

Martin Strand
Norwegian Defence Research Establishment (FFI)
NORWAY

martin.strand@ffi.no

Jan Henrik Wiik
Norwegian Defence Research Establishment (FFI)
NORWAY

jan-henrik.wiik@ffi.no

Keywords: anti-tamper, cryptography, UAV, AUV

ABSTRACT

We survey existing literature on protecting resource constrained autonomous and unmanned devices. The goal is to secure data both during transmission and at rest, also when facing highly skilled and motivated adversaries. Our main observation is that many works are lacking in terms of rigidity and maturity, and much work remains to be done in this field. Cryptography can help us to ensure confidentiality and integrity of communication, but our analysis confirms that tamper protection is crucial to protect data at rest and the keys used for communication. In particular, if tamper protection can ensure that an adversary cannot control a device immediately after capture, one can use cryptography to secure the rest of the network.

1.0 INTRODUCTION

Unmanned devices come in all formats and sizes, ranging from large vehicles and drones, to miniature sensors. They can be used for missions too dangerous, tedious or boring for humans. The former is in particular true for military applications, but this topic has many commonalities to civilian applications. For military missions, security is critical: be it confidentiality, integrity or availability. Every system have their own characteristics, but there are often some commonalities, for example (relative) low cost, resource constraints, or a need to securely store and transmit data [1].

The primary contribution of this paper is a survey of the available literature through the eyes of cryptographers. Our goals can be summarised in five points:

- Secure communications to and from the device so that unauthorised parties are unable to listen to sensitive information (confidentiality).
- Ensure that changes injected in the communication will be detected (integrity).
- Guarantee that the sender is the one it claims to be. The recipient should only accept the message if the sender can be verified (authenticity).
- Ensure that all information stored on the device remains secure if the device is lost.
- Securely process data on the device – e.g. map and sensor data.

Availability is often listed as an important goal for a system and as part of the Confidentiality—Integrity—Availability (CIA) triad. However, we view it as an *operational* goal rather than a security goal. In some sense, confidentiality and integrity are meaningless unless we also have some degree of availability. Any discussion about the trade-off between security and availability is therefore outside the scope of this paper.

We presuppose the existence of a strong adversary. Concretely, we assume the following about the adversary:

- All algorithms and protocols are known.
- All unencrypted information in the network can be read (and readily understood).
- The adversary can control the message flow on the network and alter packages at will.
- The adversary may choose to take over control over some devices, and use these to attack his target device.
- Once the adversary takes control over a device, then all of its internal state and stored keys are automatically revealed. The adversary may choose to operate the device in an honest way until he is ready to mount his attack against his real target.

Note that we apply these assumptions to an ideal mathematical model of the system, not the actual devices. The motivation is that the foundations should be sound, and then require that the implementation maintains the properties. While this approach has resulted in strong results for traditional computer networks, it remains an interesting problem whenever the devices are limited with regards to either processing power, transmission bandwidth or power supply.

At this point in the discussion, we are ignoring any anti-tamper protection. We acknowledge that this may feel counter-intuitive for an anti-tamper audience, but the result of this approach is that we may arrive at a more precise image of where the mathematics of cryptography may be sufficient, and those areas where strong anti-tamper is an absolute necessity.

These assumptions may seem unrealistically strong at first, in particular since they do not take the time of recovering the state into account. However, they are applied in the analysis of protocols like Transport Layer Security (TLS), which is used to secure the internet. If we achieve security while placing strong assumptions on the adversary, then it will also be secure against all weaker adversaries.

The body of published research within this field directly related to autonomous and unmanned devices turned out to be small, and so we extended the scope to include research on the internet of things (IoT), mobile ad-hoc networks and mobile sensor networks.

Our original research [1] divided our findings into three main groups:

- *Security models:* How other scientists have reasoned about players in a system, security requirements and trust assumptions.
- *Cryptographic algorithms for lightweight applications:* Which algorithms are best suited when used on devices with restrictions on bandwidth, power, storage or processing power.
- *Practical experiments:* Some attempts to secure small devices have been documented in academic research and showcase interesting considerations and results.

Ideally, we would like to use mathematics to secure unmanned devices, so that when they are captured, no security is lost. However, that may not be possible, and so a second goal is to identify cases where tamper protection is fundamentally necessary.

Autonomy can be viewed as a strictly stronger property than being unmanned, but the distinction has limited consequences for this paper. We will therefore colloquially use the term unmanned to refer to both, as we absorb difference regarding control and error signals into the model; everything is a message. Furthermore, different devices will have different security goals. One goal of this work was to identify whether this had been discussed and abstracted in a meaningful way in the existing literature.

1.1 Overview of the Paper

The body of this paper consists of three sections. In the following section, we survey security models by introducing the concept and assess a selection of the available literature. Section 3 mentions some experiments from other researchers. Then, in Section 4, we use the lessons learned in the previous parts to lay out which tasks we as cryptographers need to delegate to tamper protection.

2.0 SECURITY MODELS

Good security requires a clear picture of what one wants to defend and against whom. One therefore have to start by doing a thorough modelling work. It is common to require cryptography to provide confidentiality, integrity and authenticity to the data. That is still too coarse a requirement. For example, assume the communications between device A and B are compromised, but that should not lead to less security in the communications between device B and C. We therefore want more fine-tuned definitions for those exact properties we seek.

2.1 Preliminaries: Cryptographic Modelling

Before moving on to the survey, let us review some assumptions of working cryptographers. We introduce this by summarising a work by Do, Martini and Choo [2]. The authors have surveyed – from an outside perspective – the modelling of adversaries in cryptography, mobile phones and IoT. Inspired by the work done in cryptography, they give recommendations for researchers in the other fields.

The modelling tradition in cryptography can in part be traced back to Dolev and Yao [3], who stated that we should consider the network being under adversarial control. Their work was later continued and refined by Bellare and Rogaway [4], and in numerous papers in the last 30 years. Bellare and Rogaway introduced the idea that the adversary could send *queries* to players in the network, for instance *Send* [a message], *Reveal* [a key] or *Corrupt* to take complete control of a player. The last query can for instance model the effect of having an unfaithful servant in your organisation.

The cryptographic literature is normally less concerned with physical attacks, something that Do et al. note as a drawback. They still conclude that "[o]ther security-based research should look to cryptographic protocols as the gold standard for adversary models (...)." The same cannot be said about the state of the IoT domain: "IoT security, particularly, is a research field in its infancy."

The authors ended by providing three recommendations for IoT security:

- Make concrete assumptions about the environment the devices will be deployed to.
- Clearly define what causes the adversary to win.
- Specify precisely the capabilities of the adversary.

2.2 Literature Review

There is a limited amount of published academic research directly geared towards security in constrained unmanned devices. We therefore expanded our scope to include wireless sensor networks and IoT. What these fields have in common is the deployment of several cheap devices, potentially limited in regards to computational ability, bandwidth and power.

Many of the papers we discuss below seem to find inspiration from previous work on mobile ad-hoc networks (MANET). That is natural as such networks in principle are autonomous. Sauveron comments on this in a presentation of a work we will return to: it is not sufficient to inherit the security measures from MANETs. For example, one approach lets devices earn reputation by behaving according to protocol over time. This is in conflict with our way of thinking: we assume that any device under adversarial control will continue to follow protocol until it performs its (potentially devastating) attack. By then it is too late to act. One can also extrapolate the argument: a device should never be able to recover reputation after a breach of protocol. We therefore view this approach as unsuitable for military applications.

Let us now discuss selected works, paper for paper.

Marzi and Marzi [5] aims to secure wireless sensor networks and design a system where the nodes in the network together perform computations to decide which other nodes they trust. There are two approaches: the first is from biology, where “ants” leave “pheromones”. In total, these provide an ideal way though the network graph. The other approach is based on each node’s assessment of its neighbours.

Both models lack an analysis of which assumptions one can make about the adversary’s capabilities, and the model therefore implicitly assumes that the attacker will leave hints about his presence.

Pathak and Patil [6] have tried to describe how one can handle the roaming problem effectively: how one node can move from the coverage area of one station to another, and with minimal overhead for reauthorisation in the network. The authors list pseudo code to implement this functionality.

There is no security analysis of the algorithm, and it is not measured against pre-defined goals. Furthermore, the authors do not discuss how it works in existing networks like Wi-Fi and GSM and subsequent networks.

In a formally unpublished manuscript from 2013, Sen [7] discusses the security challenge in wireless sensor networks on several levels: inherent constraints, security requirements, vulnerabilities and potential solutions. The paper provides a comprehensive overview over problems and possible solutions.

Sen makes a point out of the fact that transmitting a single bit of information is as expensive as a large number of instructions. This implies that we must find a balance between bandwidth limitations and power constraints.

He goes on to list a number of security and functionality requirements.

- *Data confidentiality:* If a message is sent through a neighbouring node, then said neighbour should not be able to read it unless explicitly authorised. Sufficiently much of the metadata should be encrypted to protect against traffic analysis attacks.
- *Data integrity:* No message can be altered during transit without it being detected.
- *Availability:* The services should be available even in presence of an internal or external attack, including denial of service attacks (DoS).
- *Data freshness:* Received data should recently have been sent, and no adversary should be able to replay former messages.
- *Self-organisation:* The nodes should be able to self-heal and distribute keys. Sen notes this as a particular security challenge.

- *Secure localisation:* Sensors can be used to locate devices through their transmission strengths and similar. However, an adversary may be able to use this to make a device appear to be at a different location. The network should provide secure measures to locate its units.
- *Time synchronisation:* All devices should run by the same clock, despite adversarial efforts.
- *Authentication:* The communicating node is the one it claims to be.

Next, he lists several attacks on the various network layers. At the physical level, tampering is a major problem. Many of the potential attacks at higher levels can be avoided with suitable authentication mechanisms.

The security requirements by Sen are comprehensive but well arranged: they consider the guarantees one wishes to make, not how the inner workings of the nodes should be.

This leads us over to a paper by Akram et al. [8]. In what they describe as a position paper, they consider autonomous swarms of unmanned aerial vehicles (UAV). The paper assumes a strong adversary and presents a list of security requirements:

1. The UAV should have a secure element to ensure security and privacy of its missions.

2. Either all of the UAV or at least the secure element should be tamper resistant

3. The secure element should withstand an attacker taking over a functioning unit.

4. Each UAV should have a unique ID.

5. The UAV should have mechanisms for key management and provide confidentiality and integrity within the swarm.

6. Secure storage for sensitive data.

7. The UAV should work as a platform so that security is independent of the concrete task for which it is equipped.

Akram et al. propose a solution to satisfy these requirements by using a smart card connected to a radio. One can compare Sen’s requirements with these. The latter are more oriented towards the inner workings of the system, rather than the guarantees one wants *any* implementation to achieve. It is useful to make a note of Akram et al.’s list, although it lacks the level of abstraction a cryptographer would look for.

The last set of works we consider here is a series of publications from an EU funded research project RERUM (REliable, Resilient and secUre IoT for sMart city applications) (2013–2016). In what appears to be a program declaration at the outset, Pöhls et al. [9] state that “RERUM will work on the definition of an authentication process for heterogeneous objects with different computational and connection capabilities”. Their objectives align nicely with ours.

Fragkiadakis, Angelakis and Tragos [10] of the same project mention three families of attacks against wireless sensor networks. The first two families are based on the adversary being able inject messages into the network, and are reasonably simple to defend against using authentication. The third family contains DoS attacks, which is out of scope for this work. The mere observation that they dedicate so much space on attacks easily thwarted by straightforward measures hints to said measures not being routinely used in such networks.

The third RERUM contribution [11] we wish to mention tackles the following scenario: sensors collect data, sign and send to the network. Some of the data is only relevant for internal use; other should be gathered into a larger collection and forwarded outside the local network. The objective is to packet this data in such a way that an observer who reads the collection is unable to decide which sensor provided which part and if parts have been held back, while still maintaining the original signatures on the data.

The idea is to attach a label to each packet and sign both the message and the label. The signature can later be updated to include more messages and labels, or one can remove the same. The system is based on hash functions and the asymmetric cryptosystem RSA. However, RSA is vulnerable against a quantum computer and as such not applicable for long-term protection in military systems. We also note that this scheme may be unsuitable for constrained devices, due to RSA's large keys and relatively expensive computations.

2.3 Assessment of the Modelling Literature

While we are only citing a low number of papers here, they appear to be representative for the current state. The RERUM papers and others support the view that civilian sensor networks generally have weak security, and that strong authentication measures are not routinely used. This impression is further strengthened by Do et al.'s characterisation of the maturity of IoT security.

We make a note of the security requirements suggested by Sen and Akram et al., as well as the updateable signatures by Frädrich et al., as a basis for further research.

3.0 ACADEMIC EXPERIMENTS

We have so far found that the research literature is somewhat lacking. Despite this, some have carried out concrete experiments with seemingly good results. We now discuss such examples.

Dini and Duca [12] have described a solution for acoustic underwater communications. Their approach was tested in the Trondheim fjord [13] with promising results. The main challenges for submerged message exchanges are a high rate of packet loss and the severely limited bandwidth offered by acoustic channels.

The authors use a block cipher in ciphertext stealing (CTS) mode to provide confidentiality. This allows them not to pad the messages to fill complete ciphertext blocks, and so they save some bits of ciphertext length. Integrity and authenticity is maintained by truncating a hash function to mere 32 bits. This would be unacceptably low in almost any other scenario. The authors justify this choice by observing that the vessel can only receive three messages per second (given a bandwidth of 500 bit/s and average message length of 184 bits), and that it would then take 25 years to send sufficient many messages to brute-force message integrity with a reasonable probability.

The key distribution protocol S2RP [14] provides key authentication by first generating a chain of keys where each is the hash of the previous. They are then distributed in the opposite order, so that one can verify a key by computing its hash and checking that it is equal to the previous key. This gives a very efficient protocol, while sacrificing some security: it is impossible to exclude compromised units, and an adversary who records all messages and only picks up a key at a later stage will be able to decrypt all former messages.

Dini and Duca's system seems reasonable for its application, and worth remembering for further developments. A concrete assessment of the environment should be repeated also in other cases.

We conclude this section with a somewhat more esoteric example. Cheon et al. [15] have developed a homomorphic lightweight cipher that can compute the correct rotor response on a quadcopter, so that the drone can fly correctly along its route. The scenario assumes that the sensors sign their encrypted data and that the controller computes the response without decrypting. The rotors then verify the authenticity and decrypt the correct setting. An authentication failure will make the drone return to its initial location. This means that the on-board controller is no longer a trusted part of the setup, and does not need any particular tamper protection. Fascinatingly enough, this system has been implemented and demonstrated.

4.0 THE ROLE OF TAMPER PROTECTION

Cryptography plays a necessary role in securing unmanned devices, both regarding communication and in order to harden the device itself against hijacking attempts. For the latter, one can for instance encrypt any information not currently in use, so that it neither can be altered nor read. This partly reduces the problem to securing the keys. However, the preceding surveys also show that the current state of cryptographic research is unable to sufficiently secure a device that falls into adversarial hands. Failing to reach one of our stated goals in the introduction, it underscores the need for strong tamper protection of cryptographic keys in particular.

Cryptographers normally assume that if we lose a device, then that device immediately falls under the adversary's control, and that all communications keys are revealed instantly. This line of thinking is evident in the way cryptographers model systems. This presents us with a difficult problem: how can we detect a captured device? If one can provide protection that ensures a sufficiently long delay from when a device is lost, and to when its data is extracted, then one can use a simple timeout criterion to exclude devices. It would be valuable to combine the formality of cryptography with the practical results from tamper protection.

One of our goals is to ensure that data stored on the device remains secure after a complete loss. This is relatively easy when it comes to collected data that will only be read once it returns to base: just encrypt it with a shared key, and then delete the key. It is more difficult with sensitive data that might be necessary for the duration of the mission, e.g. maps or target signatures. We have worked on some theoretical solutions to this problem, but none will be acceptable for a real-life small device. Hence, we see no other options but to leave this problem in the hands of the physical protection. This issue also includes keys used for communication.

Computations can be analysed to reveal their underlying data. Concretely, the encryption circuit can leak the key through various side channels like timing, power consumption, radiation or faults. While it is expected that implementations use constant-time variants of the algorithms, that may not be sufficient. In particular, resourceful adversaries will try to modify the circuit to produce unexpected results that again might leak sensitive information. This should be made as hard as possible.

From a cryptographer's point of view, the tamper protection would be successful if even the designer is unable to remove the protection without also annihilating the protected part. If cryptographic keys are rotated sufficiently often, then a tampering protection that merely delays the adversary from getting full access to a device will be a successful part of the greater picture.

5.0 CONCLUSION

Our goal is to secure data storage and the communication to and between unmanned devices without using more bandwidth, battery or computational power than necessary. The messages should be unintelligible for unauthorised parties; there should be no doubt that sensor information really originates from that small helicopter far behind the enemy's line, and the unmanned car will not follow orders unless it comes from the right authority. The lone autonomous underwater vehicle (AUV) should not have to use unnecessarily much power and bandwidth on too large transmissions.

These goals were used as a filter to select literature to be surveyed.

Many of these challenges have solutions rooted in cryptography. Others require a solid contribution from tamper protection specialists. We conclude in three points:

- There exists initial work to model the security of unmanned systems. However, it is still not sufficiently formal and thorough for us to build upon.
- There have been successful trials. The experiences from those can be mixed with a solid theoretical foundation to find long-term solutions.
- Tamper protection and cryptography will play a crucial role in complementing each other.

REFERENCES

[1] M. Strand og J. H. Wiik, «Kryptografisk sikring av autonome og ubemannede enheter - eksisterende forskning,» Forsvarets Forskningsinstitutt, Kjeller, 2019.

[2] Q. Do, B. Martini og K.-K. R. Choo, «The Role of the Adversary Model in Applied Security Research,» *Computers & Security,* 2018.

[3] D. Dolev og A. C.-C. Yao, «On the Security of Public Key Protocols (Extended Abstract),» i *22nd Annual Symposium on Foundations of Computer Science, Nashville, Tennessee, USA, 28-30 October 1981*, 1981.

[4] M. Bellare og P. Rogaway, «Random Oracles are Practical: A Paradigm for Designing Efficient Protocols,» i *CCS '93, Proceedings of the 1st ACM Conference on Computer and Communications Security, Fairfax, Virginia, USA, November 3-5, 1993*, 1993.

[5] H. Marzi og A. Marzi, «A security model for wireless sensor networks,» i *2014 IEEE International Conference on Computational Intelligence and Virtual Environments for Measurement Systems and Applications (CIVEMSA)*, 2014.

[6] G. R. Pathak og S. H. Patil, «Mathematical Model of Security Framework for Routing Layer Protocol in Wireless Sensor Networks,» *Procedia Computer Science,* vol. 78, pp. 579-586, 2016.

[7] J. Sen, «Security in Wireless Sensor Networks,» *CoRR,* vol. abs/1301.5065, 2013.

[8] R. N. Akram, P.-F. Bonnefoi, S. Chaumette, K. Markantonakis og D. Sauveron, «Secure Autonomous UAVs Fleets by Using New Specific Embedded Secure Elements,» i *2016 IEEE Trustcom/BigDataSE/ISPA, Tianjin, China, August 23-26, 2016*, 2016.

[9] H. C. Pöhls, V. Angelakis, S. Suppan, K. Fischer, G. C. Oikonomou, E. Z. Tragos, R. D. Rodriguez og T. Mouroutis, «RERUM: Building a reliable IoT upon privacy- and security- enabled smart objects,» i *2014 IEEE Wireless Communications and Networking Conference Workshops, WCNC Workshops, Istanbul, Turkey, April 6-9, 2014*, 2014.

[10] A. Fragkiadakis, V. Angelakis og E. Z. Tragos, «Securing Cognitive Wireless Sensor Networks: A Survey,» *International Journal of Distributed Sensor Networks,* vol. 10, p. 393248, 2014.

[11] C. Frädrich, H. C. Pöhls, W. Popp, N. Rakotondravony og K. Samelin, «Integrity and Authenticity Protection with Selective Disclosure Control in the Cloud & IoT,» i *Information and Communications Security - 18th International Conference, ICICS 2016, Singapore, November 29 - December 2, 2016, Proceedings*, 2016.

[12] G. Dini og A. L. Duca, «A Secure Communication Suite for Underwater Acoustic Sensor Networks,» *Sensors,* vol. 12, p. 15133–15158, 2012.

[13] A. Caiti, V. Calabrò, A. Munafò, G. Dini og A. L. Duca, «Mobile Underwater Sensor Networks for Protection and Security: Field Experience at the UAN11 Experiment,» *J. Field Robotics,* vol. 30, p. 237–253, 2013.

[14] G. Dini og I. M. Savino, «S2RP: a Secure and Scalable Rekeying Protocol for Wireless Sensor Networks,» i *2006 IEEE International Conference on Mobile Ad Hoc and Sensor Systems*, 2006.

[15] J. H. Cheon, K. Han, S.-M. Hong, H. J. Kim, J. Kim, S. Kim, H. Seo, H. Shim og Y. Song, «Toward a Secure Drone System: Flying With Real-Time Homomorphic Authenticated Encryption,» *IEEE Access,* vol. 6, p. 24325–24339, 2018.