Threats to the Swarm: Security Considerations for Swarm Robotics

Fiona Higgins, Allan Tomlinson and Keith M. Martin

***Abstract*—Swarm robotics is a relatively new technology that is being explored for its potential use in a variety of different applications and environments. Previous emerging technologies have often overlooked security until later developmental stages, when security has had to be undesirably (and sometimes expensively) retrofitted. This paper compares swarm robotics with related technologies to identify their unique features where existing security mechanisms can not be applied. We then review some of the emerging applications where ineffective security could have significant impact. We conclude by discussing a number of security challenges for swarm robotics and argue that now is the right time to address these issues and seek solutions. We also identify several idiosyncrasies of swarm robotics that present some unique security challenges. In particular, swarms of robots potentially (i) employ different types of communication channels (ii) have special concepts of identity, and (iii) exhibit adaptive emergent behaviour which could be modified by an intruder. Addressing these issues now will prevent undesirable consequences for many applications of this type of technology.**

***Index Terms*—swarm robotics, security, autonomy, adaption, emergent behaviour**

I. Introduction

Swarm robotics is a relatively new area of research, and one which is growing rapidly. As with many emerging technologies, there is no formal definition of the field that engenders universal agreement, however comprehensive reviews of the state-of-the-art identify some characteristics that have been generally accepted [1]–[3]. These characteristics include robot autonomy; decentralised control; large numbers of member robots; collective emergent behaviour and local sensing and communication capabilities. Thus, from a security perspective, it is reasonable to consider swarm robotics as a special type of computer network with the aforementioned characteristics.

It has often been the case that the security of a new technology is an afterthought rather than an explicit design objective. This is not entirely surprising given the creative nature of research and the diversity of disciplines investigating the technology. Typically it is only as the technology matures, and begins to be deployed, that the security implications then become apparent. This was the case with, for example, mobile phone technology. The first generation of mobile phones were analogue, and easy to clone since they broadcast their identity clearly over the airwaves. It was also easy to eavesdrop on them by simply tuning a radio receiver to pick up conversations. Subsequently the underlying technology has been continuously modified in order to address threats that became apparent after deployment. The development of the Internet is another example of security being retrofitted to the technology.

In the case of swarm robotics, the particular security requirements of swarm robotic networks do not appear to have been investigated in any detail so far. Thus, for the above reasons, we believe that this is an opportune time to consider these issues, before any wide-scale deployment. Deferring security research until later in the technology's evolution could, depending on the application, be a risky strategy and may lead to undesirable consequences.

As far as we are aware, this is the first attempt to categorise security challenges for swarm robotics. Very little prior work appears to have been openly published. A notable exception to this is the work of Winfield and Nembrini [4] who identify several threats to a swarm of robots, which they classify as hazards. In identifying the main security challenges to swarm robotic networks it is our hope that this paper, and the work we presented at ICAS09 [5], will result in the development of robot swarm technology that is reliable and safe to deploy even in potentially hostile environments.

In Section II, we briefly review technologies that are similar to swarm robotics, highlighting the key differences and defining what we mean by a robotic swarm. In Section III we provide examples of applications that potentially will make use of this technology and show how vulnerabilities may exploited. In Section IV we discuss security, commencing with a short high level overview of security, and then cataloguing aspects of the swarm robotic environment which present challenges to security. Finally, in Section V we draw some conclusions.

II. Swarm Robotics and Related Technology

Before considering the security of swarm robotic networks it is necessary to establish the scope of the type of system we wish to secure. In other words, it is necessary to define what we mean by a swarm robotic network. There are many technologies which are similar to swarm robotic networks in some respects but differ in particular aspects. It is useful therefore to review how similar technologies, some of which have been subjected to a degree of security analysis, relate to robotic swarms. This will allow us to identify the unique features of robotic swarms that may benefit from closer scrutiny in terms of security. It is these unique features that we wish to focus on in order to identify vulnerabilities that are

The authors are with the Information Security Group, Royal Holloway, University of London, UK
f.l.higgins, allan.tomlinson, keith.martin@rhul.ac.uk

particularly pertinent to swarm robotic networks and perhaps to identify aspects of these systems that may be exploited to enhance security.

In this section we consider four technologies closely related to swarm robotic networks and then describe the unique distinguishing features of the latter.

A. Multi-Robot Systems

Like robotic swarms, multi-robot systems are a collection of robots, working together to achieve a common goal. To accomplish this, multi-robot systems are typically managed by a well–defined command and control structure. Swarm robotic systems differ from more traditional multi-robot systems in that their command and control structures are not hierarchical or centralised, but are fully distributed, self-organised and 'inspired by the collective behaviour of social insect colonies and other animal societies' [6].

Self-organisation means that sometimes the collective behaviour, even if unpredictable, may well result in solutions to problems that are superior to ones that could have been devised in advance. The parallel drawn with social societies in the animal world extends to communication – interactions between the robots can be indirect as well as direct. Fault tolerance, which is related to security, has already been extensively explored within the context of multi-robot systems with hierarchical command and control, notably in the work of Parker's ALLIANCE control architecture [7].

B. Mobile Sensor Networks

Sensor networks consist of collections of devices, or nodes, with sensors that typically communicate over a wireless network. A *mobile* sensor network is a sensor network where the nodes are either placed on objects which move [8] or where the nodes may move themselves [9]. In the latter case they are sometimes known as robotic sensor networks [10]. Hybrid systems also exist [11], where mobile robots work in conjunction with static sensors.

Although mobile sensor networks exhibit many similarities to swarm robotic networks, there are distinct differences. For example, robotic swarms may utilise a wider range of communications technologies, which extend to indirect communication such as stigmergy, as described in Section IV-D. Moreover, individual identity may be more significant in a sensor network if it is important to determine exactly where the sensed data originated.

Perhaps the most important difference is that a sensor network is not designed to have the collective emergent behaviour of a robotic swarm.

C. MANETs

Mobile Ad-hoc Networks (MANETs) consist of wireless mobile nodes that relay each others' traffic, with the nodes spontaneously forming the wireless network themselves. The special properties of MANETs, such as the lack of infrastructure, absence of trusted third parties, as well as possible resource constraints, make implementing security a very challenging task. MANETs can consists of many types of mobile devices and there is considerable existing work on their security [12], [13]. Although MANETs do not exhibit the emergent behaviour of swarms, some MANET security techniques could have relevance to swarm robotics depending on the communication method used by the swarm.

D. Software Agents

There is no universally agreed definition of a software agent, but we take one proposed by Wooldridge [14]: 'An *agent* is a computer system that is situated in some environment, and that is capable of autonomous action in this environment in order to meet its design objectives'. A *multi-agent system* (MAS) [14], [15] is a system composed of multiple autonomous agents, where each agent cannot solve a problem unaided; there is no global system control; data is decentralised; and computation is asynchronous. A *mobile* agent is a particular class of agent with the ability during execution to migrate from one host to another where it can resume its execution [15]. Thus *mobile multi-agent systems* may share many features with swarm robotic systems, but in a virtual world.

Corresponding to the active interest in mobile software agents and their rapid adoption, there has been much interest in their security [15]. However this does not always translate easily to robotic swarms because of the particular characteristics of robotic swarms which differentiate them, such as their physical nature, diverse communication mechanisms and control structure.

E. Swarm Robotics

From the brief discussion above it is clear that producing a well–defined taxonomy of mobile robotic networks will require careful consideration. Instead we now attempt to define what is meant by a swarm robotic network and show how this differs from the above related technologies.

The term 'swarm robots' generally refers to a large collection of mobile robots working on a single task [16], and the development of this technology is growing very rapidly [6]. Some of the reasons for this can be found in the perceived benefits in the characteristic properties of problem solving by social insects. The properties exhibited by social insects result in flexible, robust, decentralised and self-organised systems [6] and it is the desire to imitate these natural systems that is influencing research in swarm robotics.

Motivation

Social insects are regarded as highly effective and, some would argue, the most successful life-form on the planet. There are

many reasons why they are so successful, and the properties which make them so are highly desirable in a swarm of robots. Some of these desirable properties are:

1) Redundancy, Reliability, and Scalability: Each entity within a swarm is highly redundant. This redundancy means that the loss of individual entities has little impact on the success of the task at hand, unless all or the vast majority of them are lost. As early as 1989, Rodney Brooks at MIT proposed to NASA that teams of hundreds of inexpensive ant-like reactive robots be sent to Mars in an article entitled 'Fast, Cheap and Out of Control'. The rationale for this was, in part, to provide a degree of fault-tolerance. Having a large number robots meant that any robots damaged in transit or during landing would not have a real impact on the overall mission [17].
2) Decentralised Coordination: Coordination is completely distributed, and the task in hand will be carried out regardless of whether one or more of the individuals is lost - there is no central point of control in the swarm.
3) Multiplicity of Sensing: In a swarm, many individuals sense the same data. This means that the signal-noise ratio is greatly increased.
4) Dynamic Adaptability to the Working Environment: A swarm will adapt itself to the environment to meet the needs of the swarm.

In addition to imitating the above characteristics of social insects, there may also be practical reasons that a swarm of robots working together may be desirable. For example:

- Some tasks may be too difficult for a single robot, and may require robots working in a team to complete them.
- Using several robots may increase the speed of performing tasks.
- Designing, building, and using several simple robots may be easier, cheaper and more fault-tolerant than using a single robot.
- Theories of self-organisation show that, sometimes, the collective behaviour of a swarm results in patterns which are qualitatively different from those that could be obtained by a single entity. Randomness or fluctuations in individual behaviour, far from being harmful, may in fact greatly enhance the system's ability to explore new behaviours and find new solutions [6].

The foregoing has described some of the benefits that may be expected in systems that imitate swarms. However, as well as the advantages outlined above, there are also some potential challenges to be overcome. For example:

- Lack of global knowledge may mean that a swarm of robots does not have the information required to perform a task, and stagnates, unable to make any progress.
- Also, there is a 'too many cooks spoil the broth' effect. Having more robots working on a task or in a team increases the possibility that individual robots will unintentionally interfere with each other, lowering the overall productivity [16].
- Programming: The concept of a swarm robotic network is that individual entities are autonomous. Nevertheless their deployment implies that there is a task that they are required to do. Often the solution to the task is *emergent*, and it may be extremely complicated to program the robots to perform the task [6]. Swarm engineering is a new discipline proposed by Winfield [18] which aims to help solve this problem.
- Control and Mediation: Complex systems with swarm intelligence might be very difficult to control or mediate if they started to exhibit undesirable behaviour. Such systems would therefore need to be designed and validated for a high level of assurance that they exhibit intended behaviours, and equally importantly do not exhibit unintended behaviours [18].

Definition of Swarm Robotics

The preceding has identified the desirable properties of naturally occuring swarms thus motivating research into artificial swarms. The term 'swarm' as applied to robotics was coined by Gerardo Beni and Jing Wang in 1988 at a NATO robotics workshop in Italy [19]. At the time, the discussion was about 'cellular robots'. This term 'cellular robots' was applied to a group of robots that could work like cells of an organism to assemble more complex parts. Beni was discussing a class of cellular robots that could behave in an unpredictable way and move and interact dynamically. For this application it was generally agreed that the adjective 'cellular' was not particularly descriptive, and that 'swarm' was much better. Moreover, it more accurately portrayed the characteristics of the robots that were under discussion, which were seen to behave in a similar way to swarms occurring in nature..

When Beni and Wang introduced the term, the concept of swarm robotics was largely theoretical, but now it is a fast-evolving reality, with many research projects taking place worldwide – examples being the EU 'Guardians' project [20], 'Ultraswarm' at the University of Essex [21], Maxelbot at the University of Wyoming [22], Idaho National Laboratory projects [23] and SYMBRION [24] which is a project funded by the 7th Framework programme of the European Union.

At the 2004 Swarm Robotics workshop, Erol Şahin has proposed the following definition for swarm robotics, along with a set of distinguishing criteria to differentiate this technology from other multi-robot research:

> "Swarm robotics is the study of how large numbers of relatively simple physically embodied agents can be designed such that a desired collective behaviour emerges from the local interactions among agents and between the agents and the environment." [25]

Based on Şahin's definition a number of criteria may be described that identify a robotic swarm. These are:

1) Autonomous Robots: They should have a physical embodiment in the world, be situated and should be able to physically interact with the world.

2) Large Number of Robots: There should be a large numbers of robots (or the studies should be applicable to the control of large robotic swarms)
3) Few Homogeneous Groups of Robots: There should be relatively few groups containing large numbers of homogeneous robots.
4) Relatively Incapable or Inefficient: The robots should be relatively simple and incapable such that the tasks tackled require the co-operation of the individual robots.
5) Robots with Local Sensing and Communication Capabilities: The robots should only have localised and limited sensing and communication abilities. This constraint ensures that the coordination between the robots is distributed. However, it is acceptable to use global communication channels for a purpose such as to download a common program onto the swarm, but not for coordination among the robots.

This allows us to compare swarm robotic networks with the more established technologies described at the beginning of this section. Table I attempts to summarise and compare the characteristics of swarm robotic networks defined above with the aforementioned more mature technologies. Where there is no entry there is no clear answer to whether the technology fits the criteria or not Table II does the same for some of the implicit characteristics not explicitly included in Şahin's definition.

III. Use and Mis-use of Swarm Robotics

Section II described the scope of the systems we wish to study and showed how swarm robotic systems differ from similar, more mature technologies. Within this scope we may now begin to look at the threats that may be unique to swarm robotic systems, and perhaps identify unique features of swarm robotic systems that may help mitigate these threats.

In analysing the threats to swarm robotic systems it is useful to have an idea, first of all, of how the systems may be used; and then how the systems may be mis-used for malicious ends. In order to describe how systems may be mis-used some security terminology is introduced.

A. Basic Security Terminology

The International Organization for Standardization (ISO) has provided definitions for a number of high level security concepts and we will follow the nomenclature of ISO 13335-1 [26] in our discussion.

ISO 133335-1 defines a *threat* as 'a potential cause of an incident that may result in harm to a system or organization.' In our case, this can be interpreted as any potential incident that may adversely affect the intended objective of the swarm robotic network. The threat may be a threat to the swarm itself or to the information being processed by the swarm. Moreover, the threat can be the result of deliberate or accidental actions. The standard provides a number of examples of threats such as eavesdropping; information modification; malicious code; and physical accidents. In the following we will expand on these examples to illustrate the threats pertinent to swarm robotic networks.

Threats that are not mitigated leave *vulnerabilities* in the system. These threats may be then exploited, causing harm to the system. In other words, although all threats define a potential cause of harm, it is only the unmitigated threats that leave vulnerabilities. In the remainder of this section we consider the threats, rather than specific vulnerabilities.

We refer to the deliberate exploitation of a threat as an *attack* and those that initiate their execution as *attackers* or *adversaries*.

ISO 13335 also defines the notion of the *impact* of the exploitation of a threat, and the *risk* of a threat being exploited. While these are considerations that should be made in the deployment of any system, our objective is not focused on the specific details of a particular application and thus we will focus on the threats.

An example of a threat could be that an unauthorised person might see top secret information; a vulnerability could be that trust is misplaced in a courier delivering this information in a document. The courier may accidentally lose the document; or an attack could be that someone steals the document in transit and publishes it. The impact of a information loss will depend on the content of the document.

Security in any environment, including swarm robotics, is fundamentally about the provision of core *security services.* These services can be defined at a high level without binding the service provision to a technology specific *security mechanism*. The ISO standard for the security architecture of the OSI reference model [27] identifies a number of security services, of which the following are relevant to swarm robotic networks.

Confidentiality
: The confidentiality service protects data from unauthorised disclosure. It may protect all data in a message or selective fields. It may also be used to prevent traffic analysis.

Integrity
: An integrity service prevents prevents data from being altered in an unauthorised or unintended way; for example, by modification, insertion or deletion. As with confidentiality, it may be selective or apply to the entire message. An integrity service may also be used to detect data that has been replayed.

Authentication
: Authentication services may be classed as *peer entity authentication* services or *data origin authentication* services. The former provides assurance that the peer entity in the communication protocol is who they claim to be. The latter provides assurance that data came from its reputed source.

Availability
: Although, strictly speaking, not a security *service*, availability is defined in ISO 7498-2 as the *property*

Table I: Comparison of Explicit Characteristics

	Swarms	Multi-Robot	Mobile Sensor Networks	MANET	Multi Agent Systems
Autonomous	✓	✗			
Large Number	✓		✓		
Few Homogeneous Groups	✓		✓	✗	✓
Simple	✓		✓		✓
Local Sensing and Comms.	✓	✓	✓	✓	✓

Table II: Comparison of Implicit Characteristics

	Swarms	Multi-Robot	Mobile Sensor Networks	MANET	Multi Agent Systems
Self Organising	✓	✗	✗		
Emergent Behaviour	✓	✗	✗	✗	
Co-operate to accomplish task	✓	✓		✗	✓
Distributed Command/Control	✓	✗	✗		✓
Mobile	✓	✓	✓	✓	✓
No ID Necessary	✓		✗	✓	

of being accessible and useable upon demand by an authorised entity. The term *denial of service* is often used in reference to loss of availability.

ISO 7498-2 defines two more services: access control, and non–repudiation. Although these services are important in the context of the OSI reference model, they are of less relevance to swarm robotics. Entities in a robotic swarm are typically simple devices that do not provide access to a service, and which operate in a closed network where disclaiming a previous transaction is not a high risk.

Security mechanisms used to provide the above services include *encryption*, for confidentiality, and *digital signatures* and *message authentication codes* for integrity and data origin authentication. Entity authentication usually requires the completion of a security protocol. The Handbook of Applied Cryptography [28] provides a good introduction to these mechanisms. Mitigating denial of service attacks is more dependent on the particular application. In any system, the provision of security is a holistic process. This requires careful management processes that oversee the use of specific security technologies that can be applied to devices and networks. These include *firewalls, access control mechanisms* and *network security protocols*. At the heart of most security technologies is the deployment of specific *cryptographic primitives*, which are mathematical tools that can be applied to data to provide the core security services. These normally rely on the careful protection and maintenance of *cryptographic keys*, which are critical data items that must be stored securely.

With this background we may now review several scenarios where swarm robotic technology is being considered for use, and look at the potential threats to these systems.

B. Military Applications

Swarm robotic networks are of particular interest to the military, and in these applications the need for security is perhaps self-evident. There is currently a great deal of research taking place in the military use of robotic swarms. In the United Kingdom in August 2008, a challenge called 'The Grand Challenge' [29] took place, which was searching for the best ideas in defence technology to help solve some of the evolving threats facing front line troops. One prominent entrant to this was 'Swarm Systems' [30], which used swarms of micro air vehicles.

In the United States, US Army Research are funding and working with BAE Systems on the The Micro Autonomous Systems and Technology (MAST) project, [31] which will 'research and develop advanced robotic equipment for use in urban environments and complex terrain, such as mountains and caves. The project will create an autonomous, multifunctional collection of miniature intelligence-gathering robots that can operate in places too inaccessible or dangerous for humans'.

Mine clearance is another example of where robotic swarms may be deployed. Individual entities that constitute a swarm robotic system are dispensable, making the system suitable for domains that involve dangerous tasks. For instance, clearing a corridor on a mining field. Swarm systems would be better than a single more complex and expensive mine clearing robot because they can afford to be suicidal, and may be able to cover the area more quickly.

The major threat to military systems is from deliberate attacks on the robotic swarm. Such attacks may range from passive eavesdropping on communications, or monitoring traffic; to more sophisticated attacks where malicious robots may be injected into the swarm, much as viruses and Trojans are deployed in computer systems. Such sophisticated attacks may go un-noticed while the attacker manipulates data being processed by the swarm and possibly affects the emergent behaviour.

C. Monitoring

Robotic swarms are well–suited to environmental monitoring and, since they have motor as well as sensor capabilities, they could potentially provide solutions in the case of undesired environmental events. For example, the Elimination Units for Marine Oil Pollution (EU-MOP) project demonstrated that a robotic swarm could be used to detect environmental pollutants such as oil spillages, and subsequently clean them up [32].

At the Ecole Polytechnique Federale de Lausanne (EPFL) an investigation has been completed into how swarm robotics could be used for the autonomous inspection of complex engineered structures [33].

Accidental malfunction of the entities that make up the swarm is always a threat. The impact of this threat could be significant these applications if the swarm was the sole means of monitoring.

In addition to malfunctions or accidents, active threats to such robot swarms could arise from malicious organisations such as terrorists or criminals. Such groups could target the availability of the swarms, or the confidentiality or integrity of any information that they hold. For example, by injecting malicious code or malicious robots, they could physically or electronically hijack the system resulting in the loss of availability. As with current denial of service attacks on internet–based services, the threat of such an attack is itself sufficient to meet the adversary's requirements. As with internet–based services a malicious organisation could potentially use such a threat to extort money from the legitimate owners of the swarm. Many such attacks have already been launched against business websites on the Internet. The threat may even be exploited, with the robots only being returned to use after a ransom has been paid.

The data that a monitoring swarm holds could also be useful to an unauthorised third party. For example the location and extent of an oil spillage could be of interest to an environmental group; the location of faults in an engineering structure could be of great value to a competitor. Therefore, it is of importance that such data is kept confidential. Also, if such data could be corrupted accidentally or deliberately, it could lead to the swarm performing incorrectly, which could mean that monitoring is not taking place properly or the swarm is not trying to fix something that it is meant to be. Thus integrity protection would be a useful service to have in these applications.

D. Disaster Relief

The deployment of robot swarms during disaster relief operations is another application area that is considered for swarm robotic networks.

At the University of Utah, research has taken place into using swarms of robots to aid first responders in disaster situations [34] and the European Union 6th Framework GUARDIANS project [20] is addressing a similar application.

The GUARDIANS are a swarm of autonomous robots applied to navigate and search an urban space in situations which are dangerous and time-consuming for humans. The project's central example is an industrial warehouse in smoke, as proposed by the Fire and Rescue Service. The job is time consuming and dangerous since toxins may be released and humans senses can be severely impaired. The robots warn of toxic chemicals, provide and maintain mobile communication links, infer localisation information and assist in searching. They enhance operational safety and speed and thus indirectly save lives

In situations such as these, availability becomes a primary security requirement, as well as confidentiality, integrity and authentication/identification. Availability is necessary so that the swarm can respond as quickly as possible to the emergency at hand. If robots are unavailable due to malfunction, accident or because they have been hijacked either physically or electronically by an external agency, then they will be unable to perform their critical task. The motivation for such an attack may be difficult to comprehend, however, as discussed above, the threat of a denial of service attack may be sufficient for malicious groups to extort their demands.

Unauthorised access to data could also be a threat in this application. Eavesdropping on the robots communications may provide information to an attacker, for example about the location and extent of the damage and about any entities that the robots discover during the rescue operation. Such information may be highly sensitive and would require the protection of a confidentiality service.

Perhaps more importantly, there exists a threat of data manipulation. Integrity protection is necessary to ensure that the data being passed around the swarm is accurate, so that the robots respond correctly. In addition to integrity protection, data origin authentication may be required to provide assurance that information has come from a reliable source.

Threats arising from entity authentication failure are more subtle. It may be necessary to ensure that sensitive information obtained by the swarm is communicated only to legitimate parties e.g. the rescue service. If robots are communicating locally, problems could arise if the peer entity cannot authenticate itself. This scenario could arise if multiple swarms, perhaps with different goals, are operating in the same physical area.

Many relevant current technologies already provide full support for strong security. Communications between human personnel in emergency situations often use Terrestrial Trunked Radio (TETRA) [35] which is an open digital standard defined by the European Telecommunications Standard Institute (ETSI). Whether this technology is applicable to swarms, however, will depend on the particular implementation.

E. Healthcare

The use of swarm robotics has been considered for a wide range of healthcare services: from surgery and intrabody

diagnostics, to more routine tasks such as medication provision and patient monitoring. The European IWard project is proposing to use swarms of robots to provide assistance to healthcare workers [36].

Entity authentication will be very important for swarm robotics in healthcare situations. For example, it will be of vital importance that only legitimate robots are introduced into a human body, or sent to deliver patient medicines or read patient data from monitoring stations. Failure to authenticate could result in the introduction of swarm robots that would harm the patient surgically or whilst inside their body, by delivering incorrect medicines or by reporting medical data to unauthorised entities.

The confidentiality or privacy of patient data is paramount, and is protected by law in many countries. Apart from patients wanting to be able to choose who knows their personal medical history, it must be kept from organisations who may wish to have it for reasons such as pharmaceutical research, or to simply try and deny an individual access to insurance, employment or services.

Integrity of medical information must be ensured. Otherwise a swarm could damage a patient by responding to incorrect information such as wrong organ position, elevated blood pressure or blood sugar levels. Consequently, if a swarm were to respond incorrectly, this could seriously damage the patient's health, maybe fatally.

Availability of swarm robots in a healthcare situation is important, especially so where they are deployed in situations with critically ill patients. If they are not available and able to respond immediately then such patients could suffer greatly, and maybe die as a result.

F. Commercial Applications

As the technology develops, the hope is that robotic swarms will find commercial use. Commercial uses could include some of applications already discussed, for example monitoring or healthcare, as well as many other routine tasks that are 'dull, dirty, or dangerous' [16]. Unfortunately in any commercial environment, the motivation to gain competitive advantage will undoubtedly result in attempts to steal information, manipulate data, and disrupt services.

For example, if an organisation can interrupt their competitors service and make it unavailable or unstable, damaging its reputation, then they will become the organisation of choice. If they can steal information from their competitor then they may be able to find out their trade secrets for their own commercial gain. If they can amend their competitors data then they can make them operate unpredictably, again damaging their reputation, and making themselves appear preferable.

Thus consideration of confidentiality, integrity, authentication and availability will be required for commercial applications to be successfully adopted and deployed.

IV. Security Challenges in Swarm Robotic Environments

Section II described what we mean by swarm robotics and discussed how this differs from related technologies. In the preceding section we have described the threats that may arise in the deployment of swarm robotic systems and identified the security services that may be applied to mitigate these threats.

It is appropriate now to consider the challenges to providing these security services in swarm robotic networks. It is clear that some security problems are similar to those experienced by other related technologies, and that some solutions from these technologies may apply to swarm robotics. However, not all of these shared problems have been fully solved. Furthermore, the swarm robotic environment introduces particular security challenges that do not exist in other technologies.

A. Resource Constraints

According to our definition of swarm robotics, the robots should be relatively simple. The less complicated a device is, the greater the challenge in providing security becomes. This is due to resource constraints: storage for static and ephemeral data is restricted; communication bandwidth is constrained; and processing power is limited. Most importantly, where mobile devices are concerned, power consumption has to be minimised to preserve energy.

Resource constraints restrict the types of existing security technologies that can be deployed and special cryptographic mechanisms may be required to reduce the consumption of resources on such devices [37]–[39]. However, attacks on the provision of resources can still lead to the device becoming inoperable – permanently so if the resource is not renewable e.g. a battery. This would result in loss of availability of the device and potentially the swarm.

B. Physical Capture and Tampering

Robotic swarms are unique in their combination of physical entities with autonomous behaviour, mobility, and distributed control. Consequently, the owner of a swarm may not know the exact location of each device and what other entities may be in the vicinity. Thus individual swarm robots may be captured by an attacker.

Physical capture of a robot may lead to immediate loss of availability. The attacker may also use the device to manipulate data being reported, and may attack the device hardware to extract any secret data.

In the worst–case scenario an attacker could modify the device and re-introduce it to the swarm, enabling a number of other attacks to be carried out. Such a rogue device may continue to manipulate data as the swarm moves to new locations. It may eavesdrop on communications. It may even be able to introduce malicious code or commands to other devices. In the worst case it would be able to alter the behaviour of the swarm without the attack being detected. This capture and 're-introduction attack' is unique to swarm robotic technology.

C. Monitoring and Control

Systems employing swarm intelligence do not have a hierarchical structure with specific points of monitoring and control. Moreover, the individual entities within these systems take decisions autonomously, based on local sensing and communications. With such systems it is evident that there could be many risks if, for any reason, deliberate or accidental, they went 'out-of-control'. These risks include many security violations such as loss of confidentiality, integrity or availability. Monitoring and control presents an interesting challenge to security within swarm robotics.

D. Communication

Unlike the related technology discussed in Section II, robotic swarms may be designed to interact either explicitly, or implicitly [40].

Explicit communication can be achieved via broadcast or directed messages. Radio frequency (RF) and infra-red (IR) technologies have been widely used for explicit communications within swarms. Other technologies include coloured LED display, body-language or sign-language, colour patterns on a robot's body, coil induction, haptics, audible sounding, combination of LED display and audio signalling and acoustic signalling in an underwater environment.

Implicit communication is unique, amongst the technologies discussed in Section II, to swarm robotics. By considering implicit communication, we include interaction via sensing other robots and their behaviours, and interaction via the environment. The latter acts as a sort of shared memory and is known as *stigmergy* [6], [41], [42].

From a security perspective, any open implicit or explicit communication method can be jammed, intercepted or otherwise disturbed relatively easily by an attacker. The security of RF and IR has been well–researched, but the security of the remaining more 'exotic' interaction methods needs to be thoroughly investigated and presents a fascinating security challenge.

E. Swarm Mobility

Security is difficult to provide in any mobile environment, however the mobility of robot swarms, combined with the autonomous behaviour, is quite unusual. This has some interesting characteristics that might make some security services easier to implement than for related technologies.

One example is entity authentication discussed below. Swarm robots may be able to move towards the peer that they wish to authenticate. The authentication service could then be provided through visual sensing and physical data exchange. A similar example is key distribution: robots may move within the swarm to distribute shared keys.

However any constraint on the movement of swarm members, for example to remain in the 'bounds' of the swarm, could present additional security issues.

F. Entity Authentication and Identity

As discussed in Section III, in some applications it may be very important for a swarm robot to determine whether it is interacting with a legitimate entity or not. Data origin and entity authentication often require some notion of *identity*. This is a particular problem where *individual identity* within a swarm is undesirable [43]. However, other work on robotic swarms has used *group identity* [44], or individual identity which is broadcast regularly [45].

If identity can be assumed or changed, then attacks can be launched on entity authentication, confidentiality, integrity and availability. The notion of identity within a robotic swarm thus presents an interesting challenge from a security standpoint.

G. Key Management

Security services deployed in a robot swarm inevitably require the need to manage cryptographic keys [46]. These keys define which pairs (or groups) of robots can apply security services. As robots join and leave a swarm, it may be necessary to update this keying material. Thus the dynamic and interactive nature of a swarm presents sophisticated key management challenges although the intelligent mobility of the swarm may provide some novel solutions to this.

H. Intrusion Detection

When an unauthorised entity joins a network it is sometimes called *intrusion,* and the field of network intrusion detection is well–established [47]–[49]. However, the physical nature of the entities in a swarm robotic network means that intrusion in a robotic swarm robotic network is not the same as intrusion in a traditional data network. Deliberate intrusion was alluded to in Section IV-B and accidental intrusion in Section III-D, where several swarms may operate in the same geographical location. In these cases the intrusion mechanism is the physical insertion of a rogue agent into the swarm, which is not addressed by the established network intrusion detection systems.

Intrusion detection systems typically attempt to detect anomalous behaviour in the network. In a robotic swarm this behaviour may extend to the physical behaviour of individual robots. New mechanisms will be required to detect anomalous physical behaviour. Moreover, the autonomous nature of robots and collective emergent nature of the behaviour of the swarm will make any anomalous behaviour difficult to detect.

If undetected, one or more foreign robots could infiltrate the swarm, either maliciously or accidentally, and ultimately affect the desired emergent behaviour. The situation is illustrated in figure 1 where the shaded nodes represent foreign robots infiltrating the swarm. In figure 1a a single intruder has entered the swarm. It is not unreasonable to expect the swarm to be able to detect this intruder. If several foreign robots can infiltrate the swarm, as shown in figure 1b, it may be easier for the intruder to affect the behaviour of the swarm. As

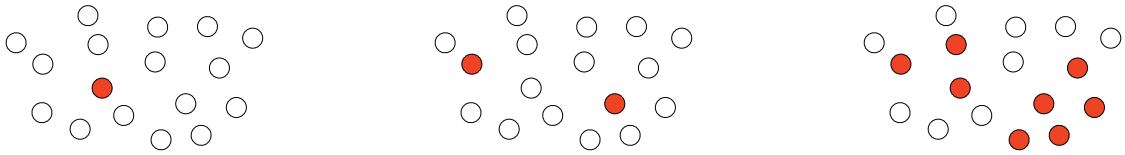

(a) Single Intruder (b) Multiple Intruders (c) Overlapping Swarms

Figure 1: Intrusion of rogue robots into a swarm

more foreign robots infiltrate the swarm, 1c, it will become more difficult to distinguish the intruders from the original swarm, particularly if the notion of identity within a swarm is forbidden.

Once an intruder is detected, an appropriate response will need to be formulated according to an *Intrusion Protection System* [49]. Depending on the application the response could be to simply ignore the rogue device, to monitor its behaviour, or to find a way to either disable it or remove it from the system. In some scenarios it may even be desirable to manipulate the behaviour of the intruder in a counter-attack.

Intrusion detection and protection looks to be particularly challenging in a swarm of robots, and will need a specifically tailored approach.

I. Managing Learning

Robotic swarms are designed to learn and react to environmental changes by means of adaption. A malicious entity might present changes to the *environment* which will cause a swarm to adapt in an undesired way. For example, if anomaly detection is used to detect intrusion based on learning and monitoring typical behaviour, then a malicious entity could manipulate the environment to change the pattern of 'typical' behaviour in order to gain entry to the network.

V. Conclusion

The development of swarm robotic technology has reached a point where many new applications beginning to emerge. Therefore, we believe that this is an opportune moment to take a closer look at the security of swarm robotic systems - before widespread deployment.

To that end we have identified several unique features of swarm robotic networks that distinguish them from related technology and consequently justify further study from a security perspective. Most notable of these characteristics are the autonomous behaviour of the swarm and the emergent behaviour. Although much has already been accomplished to provide security for related technologies, the characteristics of autonomy and emergent behaviour, combined with mobility and distributed control, make robotic swarms significantly different from these technologies to raise a number of new security problems. These new security problems are not only of theoretical interest but will have implications on many practical applications of swarm robotic technology.

Bearing this in mind, we described a number of challenges to robotic swarm security, many of which are unique to this technology. For example, the application of stigmergic communications may provide a new attack surface that will require the development of new security mechanisms. The notion of identity within a swarm may also necessitate research into the provision of entity authentication within a swarm. And finally the potential to modification of emergent behaviour if a malicious entity manages to infiltrate the swarm may require further investigation into intrusion detection, especially where the intruder is a physical mobile agent. We therefore believe that an investigation of the above areas is timely.

References

[1] E. Şahin and W. M. Spears, Eds., *Swarm Robotics: SAB 2004 International Workshop*, ser. LNCS. Santa Monica, CA, USA: Springer Berlin / Heidelberg, Jul 2005, vol. 3342.

[2] E. Şahin, W. Spears, and A. Winfield, Eds., *Swarm Robotics: Second International Workshop, SAB 2006*, ser. LNCS. Rome, Italy: Springer Berlin / Heidelberg, Sep 2007, vol. 4433.

[3] L. Bayindir and E. Şahin, "A review of studies in swarm robotics," *Turkish Journal of Electrical Engineering*, vol. 15, pp. 115–147, 2007.

[4] A. F. T. Winfield and J. Nembrini, "Safety in numbers: fault–tolerance in robot swarms," *International Journal of Modelling, Identification and Control*, vol. 1, no. 1, pp. 30–37, 2006. [Online]. Available: http://inderscience.metapress.com/link.asp?id=byu8g6fqm5g55v9x

[5] F. Higgins, A. Tomlinson, and K. Martin, "Survey on security challenges for swarm robotics," in *Fifth International Conference on Autonomic and Autonomous Systems (ICAS 2009)*, R. Calinescu, F. Liberal, M. Marin, L. P. Herrero, C. Turro, and M. Popescu, Eds. Valencia: IEEE Computer Society, April 2009, pp. 307– 312.

[6] E. Bonabeau, M. Dorigo, and G. Theraulaz, *Swarm intelligence: from natural to artificial systems*. Oxford University Press US, 1999.

[7] L. E. Parker, "Alliance: an architecture for fault tolerant multirobot cooperation," *IEEE Transactions on Robotics and Automation*, vol. 14, no. 2, pp. 220–240, Apr 1998.

[8] T. Wark, C. Crossman, W. Hu, Y. Guo, P. Valencia, P. Sikka, P. Corke, C. Lee, J. Henshall, K. Prayaga, J. O'Grady, M. Reed, and A. Fisher, "The design and evaluation of a mobile sensor/actuator network for autonomous animal control," in *IPSN '07: Proceedings of the 6th international conference on Information processing in sensor networks*. New York, NY, USA: ACM, 2007, pp. 206–215.

[9] K. Dantu, M. Rahimi, H. Shah, S. Babel, A. Dhariwal, and G. Sukhatme, "Robomote: Enabling mobility in sensor networks," in *Proceedings of Fourth International Symposium on Information Processing in Sensor Networks*, 2005, pp. 404–409.

[10] (2008) Robotic sensor networks. Minnesota University. [Online]. Available: http://rsn.cs.umn.edu

[11] J. Reich and E. Sklar, "Toward automatic reconfiguration of robot-sensor networks for urban search and rescue," in *Proceedings of the 1st International Workshop on Agent Technology for Disaster Management*, 2006, pp. 18–23. [Online]. Available: http://users.ecs.soton.ac.uk/sdr/atdm/ws34atdm.pdf

[12] L. Buttyán and J.-P. Hubaux, *Security and cooperation in wireless networks: thwarting malicious and selfish behavior in the age of ubiquitous computing*. Cambridge University Press, 2007.

[13] E. Hansson, A. Bengtsson, and A. Vidström, "Security solutions for mobile ad hoc networks," Swedish MOD, FOI Defence Research Agency Command and Control Systems, Tech. Rep. FOIR-1694-SE ISSN 1650-1942, Aug 2005.

[14] M. Wooldridge, *An Introduction to MultiAgent Systems*. Wiley, 2002.

[15] N. Borselius, "Multi-agent system security for mobile communication," Ph.D. dissertation, Department of Mathematics, Royal Holloway, University of London, 2003.

[16] R. R. Murphy, *Introduction to AI Robotics*. MIT Press, 2000.

[17] R. A. Brooks and A. Flynn, "Fast, cheap and out of control - a robot invasion of the solar system," *Journal of the British Interplanetary Society*, vol. 42, pp. 478–485, 1989.

[18] A. F. T. Winfield, C. Harper, and J. Nembrini, "Towards dependable swarms and a new discipline of swarm engineering," in *Swarm Robotics: SAB 2004 International Workshop*, ser. LNCS, E. Şahin and W. M. Spears, Eds., vol. 3342. Santa Monica, CA, USA: Springer Berlin / Heidelberg, Jul 2005, pp. 126–142.

[19] Beni and Wang, "Swarm intelligence in cellular robotic systems," in *NATO Advanced Workshop on Robotics and Biological Systems*, Il Ciocco, Tuscany, Italy, 1989.

[20] EU. (2007) 6th framework programme: Guardians project. European Union. [Online]. Available: http://www.guardians-project.eu/

[21] R. D. Nardi and O. Holland, "Ultraswarm: A further step towards a flock of miniature helicopters." in *Swarm Robotics: Second International Workshop, SAB 2006*, ser. LNCS, E. Şahin, W. Spears, and A. Winfield, Eds., vol. 4433. Rome, Italy: Springer Berlin / Heidelberg, 2006. [Online]. Available: http://gridswarms.essex.ac.uk/publications/DeNardi2006UltraswarmFurther.pdf

[22] W. M. Spears and D. F. Spears. Maxelbot project. [Online]. Available: http://www.cs.uwyo.edu/~wspears/ap.html

[23] Idaho National Laboratory. [Online]. Available: http://www.inl.gov/adaptiverobotics/robotswarm/index.shtml

[24] [Online]. Available: www.symbrion.eu

[25] E. Şahin, "Swarm robotics: From sources of inspiration to domains of application," in *Swarm Robotics: SAB 2004 International Workshop*, ser. LNCS, E. Şahin and W. M. Spears, Eds., vol. 3342. Santa Monica, CA, USA: Springer Berlin / Heidelberg, Jul 2005.

[26] *Information technology – Security techniques – Management of information and communications technology security – Part 1: Concepts and models for information and communications technology security management*, International Organization for Standardization (ISO) Std. 13 335, Rev. 1, 2004.

[27] *Information processing systems – Open Systems Interconnection – Basic Reference Model – Part 2: Security Architecture*, International Organization for Standardization (ISO) Std. 7498-2, Rev. 1, 1989.

[28] A. J. Menezes, P. C. van Oorschot, and S. A. Vanstone, *The Handbook of Applied Cryptography*, 5th ed. CRC Press, 1996. [Online]. Available: http://www.cacr.math.uwaterloo.ca/hac/

[29] (2008) The grand challenge. United kingdom ministry of defence. [Online]. Available: http://www.challenge.mod.uk

[30] (2008, Aug) Swarm systems: providing swarms of micro air vehicles. [Online]. Available: http://www.swarmsys.com

[31] BAE-Systems. (2008, April) Mast project. [Online]. Available: http://www.baesystems.com

[32] K. Gkonis, T. Pavlidis, N. Kakalis, and N. Ventikos, "Final project report for the elimination units for marine oil pollution (eu-mop) project," European Union, 6th Framework Programme, Tech. Rep., 2008.

[33] N. Correll and A. Martinoli, "A challenging application in swarm robotics: The autonomous inspection of complex engineered structures," *Bulletin of the Swiss Society for Automatic Control*, vol. 46, pp. 15–19, 2007.

[34] D. P. Stormont, "Autonomous rescue robot swarms for first responders," in *IEEE International Conference on Computational Intelligence for Homeland Security and Personal Safety (CIHSPS 2005)*, Orlando, FL, USA, Apr 2005, pp. 151–157.

[35] (2008) Terrestrial trunked radio. Tetra-Association. [Online]. Available: http://www.tetra-association.com

[36] EU. (2007) 6th framework programme: Iward: Intelligent robot swarm for attendance, recognition, cleaning and delivery. European Union. [Online]. Available: http://www.iward.eu/

[37] S. Kumar and C. Paar, "Reconfigurable instruction set extension for enabling ecc on an 8-bit processor," in *International Conference on Field-Programmable Logic and Applications (FPL) 2004,*, Antwerp, Belgium, 2004.

[38] L. Batina, J. Guajardo, T. Kerins, N. Mentens, P. Tuyls, and I. Verbauwhede, "Public-key cryptography for rfid-tags," in *IEEE International Workshop on Pervasive Computing and Communication Security*, New York, USA, 2007.

[39] A. Bogdanov, L. Knudsen, G. Leander, C. Paar, A. Poschmann, M. Robshaw, Y. Seurin, and C. Vikkelsoe, "Present: An ultra-lightweight block cipher," in *Cryptographic Hardware and Embedded Systems - CHES 2007*, 2007, pp. 450–466.

[40] L. E. Parker, "Current state of the art in distributed robot systems," in *Distributed Autonomous Robotic Systems 4*, L. E. Parker, G. Bekey, and J. Barhen, Eds. Springer, 2002, pp. 3–12. [Online]. Available: http://www.cs.utk.edu/~parker/publications/DARS_2000_overview.pdf

[41] P.-P. Grassé, "La reconstruction du nid et les coordinations interindividuelles chez bellicositermes natalensis et cubitermes sp. la théorie de la stigmergie: Essai d'interprétation du comportement des termites constructeurs," *Insectes Sociaux*, vol. 6, no. 1, pp. 41–80, 1959.

[42] T. White, "Expert assessment of stigmergy: A report for the department of national defence," School of Computer Science, Carleton University, Ottawa, Ontario, Canada, Tech. Rep., 2005. [Online]. Available: http://www.scs.carleton.ca/~arpwhite/stigmergy-report.pdf

[43] P. Flocchini, G. Prencipe, N. Santoro, and P. Widmayer, "Gathering of asynchronous robots with limited visibility," *Theoretical Computer Science*, vol. 337, no. 1-3, pp. 147–168, 2005. [Online]. Available: http://www.sciencedirect.com/science/article/B6V1G-4FC37VR-1/2/119f08be923e94ddbdb49819df053e63

[44] R. A. Russell, "Visual recognition of conspecifics by swarm robots," in *Australasian Conference on Robotics and Automation*, 2004. [Online]. Available: http://www.araa.asn.au/acra/acra2004/papers/russell.pdf

[45] J. Fredslund and M. Matarić, "A general algorithm for robot formations using local sensing and minimal communication," *IEEE Transactions on Robotics and Automation*, vol. 18, pp. 837–846, 2002.

[46] S. Dolev, L. Lahiani, and M. Yung, "Secret swarm unit reactive k – secret sharing," in *Progress in Cryptology – INDOCRYPT 2007*, ser. LNCS, vol. 4859. Chennai, India: Springer Berlin / Heidelberg, Dec 2007, pp. 123–137.

[47] J. P. Anderson, "Computer security threat monitoring and surveillance," James P. Anderson Co., Tech. Rep., 1980.

[48] D. E. Denning, "An intrusion detection model," in *Proceedings of the Seventh IEEE Symposium on Security and Privacy*, May, 1986, pp. 119–131.

[49] K. Scarfone and P. Mell, "Guide to intrusion detection and prevention systems (idps)," National Institute of Standards and Technology (NIST), Gaithersburg, MD, USA, Tech. Rep. SP 800–94, Feb 2007. [Online]. Available: http://csrc.nist.gov/publications/nistpubs/800-94/SP800-94.pdf