The Layered Security Model and its Representation using Bigraphs to Analyse Critical Infrastructure

Clive Blackwell
Information Security Group, Royal Holloway, University of London
Egham, Surrey. TW20 0EX. United Kingdom.
C.Blackwell@rhul.ac.uk

1 Introduction

There is a widening gap between our understanding of systems and their ever increasing complexity, functionality and connectivity. We require more sophisticated functionality for novel applications, and systems to interoperate with each other dynamically and autonomously to meet their different objectives. Piecemeal defences address limited technical problems, rather than tackle system requirements comprehensively. This leads to brittle systems with single points of failure that break easily with unpredictable consequences. We still often rely on the binary distinction between insider and outsider, whereas we need more fine grained measures to cope with a continuum of access rights and to manage the effects of successful attacks.

Some security issues that need to be seriously addressed include emergent system behaviour, effects at a distance, unexpected changes to a system and its environment, and new methods of attack. As Einstein said, “we can’t solve problems by using the same kind of thinking we used when we created them”. We provide an informal architectural model that can be formalised, which can analyse systems that have multiple independent mechanisms operating at different layers and locations with different protective characteristics. This helps to plan, design and build systems to help provide comprehensive protection and assurance that they will complete their missions in the presence of security vulnerabilities and functional defects rather than respond tactically to every little problem.

Only a few systematic models of security can represent systems in their entirety, rather than as technical systems alone. There is the longstanding effort in classifying the important aspects of dependability, including security, which offers a comprehensive taxonomy of the different types of fault and methods to manage them [1]. Neumann [2], and originally Neumann and Parker [3], organised systems into eight layers for security analysis, which are listed starting from the highest layer as the external environment, user, application, middleware, networking, operating system, hardware and internal environment. We consider this as a logical and useful aid to understanding systems, but we have introduced some new organisational criteria to give a simplified model that has only three layers. Howard and Longstaff [4] present a classification system for network security incidents, which shows the different types of entity involved in an attack and their interrelationships. The classification can be extended with dual concepts to model the defence, and by explicitly including the semantic and physical aspects of systems as well as computer and networks.

2 The Layered Security Model

2.1 The Layers

We model systems and their interactions in a three-layer hierarchy, where each layer can have sub-layers when required for detailed analysis (figure 1). The *semantic* or *conceptual layer* is the top layer that includes people and the abstract representation of systems including their requirements. The *logical layer* is an intermediate layer containing entities in an intangible form including data and software that are stored and processed on computers and transmitted between them. The purpose of this layer is to carry out the objectives of semantic layer entities, as they cannot interact directly with logical entities. For example, people are represented by a logical proxy such as an account, a process or a cryptographic key to act on their behalf. In addition, the logical layer contains helper entities such as network routers that aid other logical entities to fulfil their requirements. The *physical layer* is the bottom layer that represents the physical or basic existence that all entities have in the real world. The physical layer includes the physical components of systems and the environment including both tangible objects and electromagnetic radiation.

Every subject, object, relationship or piece of information, except abstract concepts, has a physical representation in addition to its existence at higher layers. Any activity carried out by a system is ultimately on behalf of a semantic subject, but the work must take place at the physical layer. However, higher layer entities cannot be understood at the physical layer. For example, information is ultimately stored physically, but understanding involves knowing its meaning, purpose, and maybe other attributes such as its origin and correctness that can only be fully appreciated at a higher layer.

This is much simpler than Neumann's eight-layered model, but it can still provide detailed analysis of systems by allowing each layer to have sublayers. For logical network communication, the best-known model is the seven-layer OSI network model [5]. We would use Tanenbaum's simplified five-layer network model [6] as sub-layers of our logical layer with the link, network and transport layers as intermediate sub-layers and the upper application and lower physical sub-layers interfacing to the social and physical layers of our model respectively.

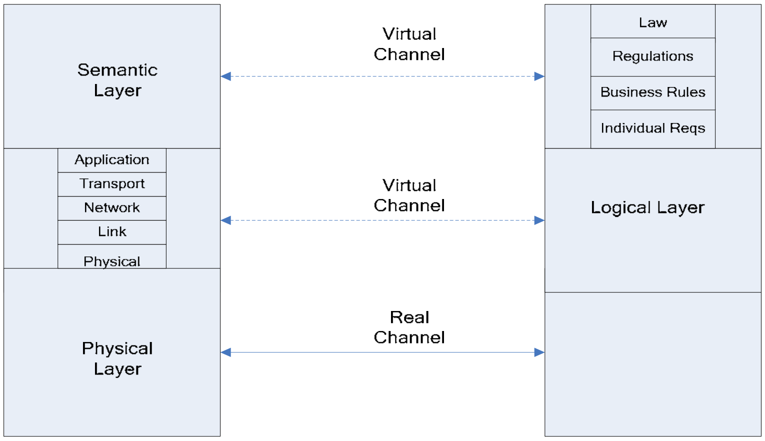

Figure 1-The three-layer model showing the sub-layers of the semantic and logical layers

Each layer has a separate concept of location and distance between entities. Higher layer entities also have a physical existence, so they are represented by different locations at each layer. The entities at each layer have different extents, dependencies and interactions and should meet the system requirements for that layer. A *channel* is an entity that carries flows of information and objects from one location to another at the same layer. Entities in different locations at the same layer use a channel to communicate. The channels at the higher layers are virtual, and must use a physical channel to communicate analogously to communication in the OSI network model (figure 1). The introduction of horizontal scope allows us to remove some of Neumann's layers [2] such as the networking and middleware layers and represent them as horizontal communication channels for computer and application entities respectively.

A *virtual entity* is a logical entity that uses some controls to limit access, so that it can only be understood, accessed or used with special methods or knowledge such as using cryptographic keys. Application layer resources such as data and services can be virtualised by replicating them, which removes the reliance on protecting single locations and thereby avoids single points of failure. In addition, many systematic controls can be represented using virtual entities including Trusted Platform Modules (TPMs) and virtual machines.

2.2 Protection

All systems have a horizontal scope at every layer. Neumann [2] considers four conceptual locations for compromise at every layer; from outside, above, within and below. Protection from an external entity at the same layer requires horizontal controls, whilst protection from a higher layer entity requires a vertical boundary between the layers. Insiders should be constrained by partitioning the system with additional internal system boundaries they should not be able to breach. However, complete protection from insiders may not be possible, so these controls may use detection and recovery mechanisms such as auditing and redundancy, so that misuse of the system is detected and recovered from, rather than prevented. Replication is an effective defence against insider attacks by backing up data or providing standby services in separate locations that insiders cannot access. Some components that control the system must be trusted and so they should be made simple enough to assure. They should reside in an inaccessible location at the bottom of the system or use a secure control network to stop external interference. In conclusion, all entities, apart from unconditionally trusted entities, should be outsiders relative to one or more controls that moderate their use of the system.

The boundaries can be annotated with their access controls represented by logical predicates or probabilistic estimates of successful defence. All the controls should be analysed together to show that they meet the system policy. For example, a firewall could be annotated by the port numbers of protocols it blocks, and anti-virus software by the signatures of malicious code it detects. This is intended to deal with attacks by partitioning the defence at the network and application layers, but a new virus using an allowed protocol would breach the defence, unless there were additional controls.

3 Modelling Multi-layer Systems

3.1 Coordinate Representation

We can represent the location of logical entities in multiple dimensions as for physical entities. For example, physically separate entities can communicate securely by the creation of a virtual tunnel that can be represented in a higher dimensional space. The location of real entities is represented in real coordinates (x, y …), whereas cryptographic entities have additional dimensions as well, which are represented by complex coordinates (z, w …). The coordinates indicate the location where an entity can access the data, which is only possible if it has access to the correct keys, indicated by access to the right virtual coordinates, and it can reach the real location represented by the real coordinates. A specific example is an SSL connection, whose real components are the Internet addresses of the path taken, and the virtual component might be an integer identifying the cryptographic channel uniquely amongst current cryptographic sessions. Cryptographic protection offers weaker protection semantics than physical security, as data can be deleted or altered with access to the physical location or communication path alone, without having access to any keys.

The coordinate representation has many applications such as reasoning about possible breaches of security by attackers in different logical locations with different knowledge and abilities. Bigraphs are a more abstract topological representation of this idea that only retains the shapes of entities by discarding the location coordinates, which simplifies the analysis and, in addition have executable semantics. Both methods can represent Neumann's four conceptual locations for compromise more formally [2].

3.2 Bigraphs

Our model can be formalised using multi-level graphs with one level for each layer of our model. We propose formalisation using Milner's bigraph model [7] that can represent physical and logical levels and the interaction between them. It represents the semantic layer indirectly through its actions and effects at lower layers. The model is quite flexible having its origin as a unification theory for process calculi such as Petri nets and the pi-calculus that model communication, together with models such as the ambient calculus that handle physical movement. The model natively incorporates the structure and organisation of the physical and logical layers including their interaction and dependence on each other, which is not considered in most security models. We propose a novel use for modelling security architecture and apply it to critical infrastructure protection.

The bigraphical structure composes two graphs with one to represent logical communication and the other physical mobility. A bigraph can model the security architecture of systems, as security mechanisms can be represented as graph rewriting or reaction rules. The application of a rule in one direction represents the defence, which is reversed by the user to access the system. The defender's objectives can be defined by invariants of the bigraph that hold in secure states of the system, and certain reaction rules representing actions that should only be performed by the defender.

The system is represented by a bigraph with security requirements represented as constraints in the bigraph. Different types of attacker with various powers and locations can be represented using an attacker bigraph occupying a particular kind of node, having particular communication possibilities and accessing certain reaction rules. The model is executed to discover if the attacker can breach or inactivate security controls, access critical assets or reduce system functionality.

4 Critical Infrastructure Protection

There are many applications of our three-layer model as computer systems and physical objects must always be used to meet higher-layer organisational and personal requirements. One important application is modelling critical infrastructure, which are very complex systems that are impossible to analyse manually and the effects of failure could be devastating. These systems have large numbers of people with various degrees of physical or logical access to parts of the system such as buildings, equipment, computers and control systems. The large horizontal extent at all layers may allow unauthorised remote access to computers, and access to insecure physical components. They can function in unexpected ways with remote effects and complex interaction between the layers with unpredictable consequences.

The model can faithfully represent both physical and logical attacks on control systems and communication links in critical infrastructure. We can formalise the representation using bigraphs and analyse the resulting model for vulnerabilities. It can model dependencies between components, attackers in various locations with various powers, and handle effects in remote locations and other layers. It can model hybrid attacks that use several layers and transitive attacks that operate in several stages. Controls can be proposed to avoid or mitigate undesirable effects by partitioning and isolating systems horizontally, and restricting changes of layer vertically.

In figure 2, the ellipses represent physical locations, whereas the arcs represent logical communication. The locations can be of different kinds, which may have different modes of interaction and communication. The outermost ellipse might be a building, the small circle might represent a computer or machine, and an intermediate size ellipse may represent a room (or possibly a network). The flexibility of the model is demonstrated by the example, as these ellipses could equally represent networks, computers and applications instead.

The physical movement of people and tangible objects is modelled by movement between ellipses, which is controlled by the kinds of the nodes and the available reaction rules. The arcs represent different types of communication, interaction and control including computer, telecommunication and power networks. Communication is controlled by the kinds of communicating entities, the type of the channels, and the available reaction rules. Entities are allowed to move over communication channels as well. A logical entity such as a user account (acting as a proxy for a person) may move over the logical link to the remote computer if it can log on. All links should be protected cryptographically or by enclosure within a secure physical boundary. This interaction between physical location and logical communication is represented natively using bigraphs, but not most formal models that discard location information.

The lower large ellipse represents a building containing a room holding a telecoms switch S that can be accessed and controlled through the administrator's workstation. A defence objective is that there should be no path from the outside to the switch S by either physical or logical means from unauthorised people or malware.

A user can access the administrator's computer remotely if he can use the correct key K required for authentication, which is indicated by the graph reaction rule as shown. The defence initially set up the requirement for external authentication using the reverse reaction rule. Figure 2 shows that the switch S can be accessed in multiple ways from outside both logically through transitive access to S via the administrator's workstation A, and physically by entering the building and then the room. It is also possible to represent hybrid attacks where both logical and physical accesses are combined. The link to the room containing S could be used to send a command to turn off the power supply for example. All these controls can be represented by the kinds of node, types of channel and available reaction rules, so the model can be executed to determine if an external attacker can breach any of the controls to interfere with S.

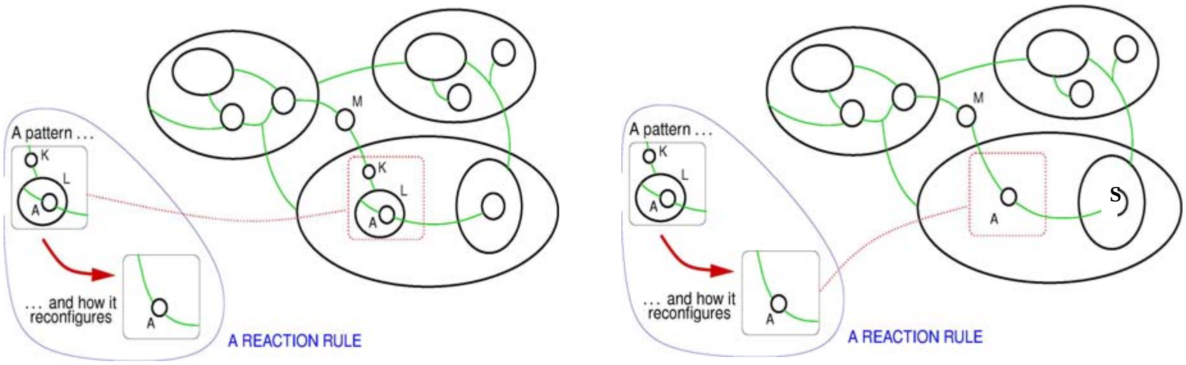

Figure 2 – Use of an authentication key K to remotely access an administrator's workstation A through the boundary L represented by a reaction rule (© Milner (2005) [7])

When compared to our coordinate representation, the diagram can be considered a flattening of the multidimensional structure into two dimensions. The physical entities such as buildings have height one, whereas the logical entities such as computers have height two. Communication is considered to take place at the highest layer of the communicating peers. The cryptographically protected communication channel is considered to occupy a fourth orthogonal dimension not accessible to real entities.

5 Conclusion and Further Work

Comprehensive protection against the different types of threat can be provided by multiple defensive controls that each create a boundary meeting different protection requirements. For example, it was suggested that further internal controls be used to protect against insiders, which allows the treatment of insiders and outsiders to be unified.

We demonstrated an informal three-layer model for modelling security architecture that allows us to reason about the structure and organisation of systems components and their interaction. A coordinate system was provided to represent the location of logical entities that allows the modelling of Neumann's four conceptual locations of attacks at all layers. We formalised the three-layer model using bigraphs and used it in the critical infrastructure example to show how systems can be compromised at the physical and logical layers including multi-layer attacks that use both.

We have used bigraphs to represent cryptographic primitives such as hash functions and digital signatures [8] and intend to use it to analyse Kerberos, which is a complex network security protocol that takes account of physical vulnerabilities such as insecure workstations as well as logical vulnerabilities.

We propose some extensions to the bigraph model to broaden its applicability and to model the security requirements of systems more faithfully. Intermediate layers can be introduced to model different layers of the network stack such as the network and application layer. Additional layers can also model the virtualisation of hardware or the operating system. Users interact with the virtual layer, which is translated to the real activity performed by the layer underneath. This sandwich layer allows, among other things, policy enforcement with additional security checks, or the virtual system acting differently to the underlying layer.

An important application is modelling the interaction between the control elements of a system and its functional components. The control space must interact with the rest of the system through physical proximity or at a distance through logical communication channels. For example, a hardware-based Trusted Platform Module (TPM) has its own separate components in a secure location performing computation with its own processor, communicating with its own dedicated buses and using its own physical storage. The TPM can be represented by a separate bigraph encapsulated within the complete system, which communicates through dedicated control channels to control access to the system resources.

References

1 Avizienis, A, Laprie JC, Randell B and Landwehr C, "Basic Concepts and Taxonomy of Dependable and Secure Computing", *IEEE Transactions on Dependable and Secure Computing*, vol 1, no 1 (2004).

2 Neumann, PG, "Practical Architectures for Survivable Systems and Networks" (2000), online at www.csl.sri.com/neumann.

3 Neumann PG, Parker D, "A Summary of Computer Misuse Techniques", *Proceedings of the 12th National Computer Security Conference*, (1989).

4 Howard JD and Longstaff TA, "A Common Language for Computer Security Incidents" (1998), Sandia National Laboratories, online at www.sandia.gov.

5 Day JD and Zimmermann H, "The OSI Reference Model", *Proceedings of the IEEE*, vol 71, pp 1334-1340, (1983).

6 Tanenbaum, AS, *Computer Networks (4th edition)*, Prentice-Hall, Upper Saddle River, New Jersey, (2003).

7 Milner R, "Bigraphs, a Tutorial", (2005), online at www.cl.cam.ac.uk/users/rm135.

8 Blackwell C, "Using bigraphs to represent cryptographic primitives", submitted to *ACM New Security Paradigms Workshop* (2007).