Anti-Social Networking:

Crowdsourcing and the CyberDefence of National Critical Infrastructures

Chris W. Johnson,

School of Computing Science, University of Glasgow, Glasgow, Scotland, G12 8RZ.

Christopher.johnson@glasgow.ac.uk, http://www.dcs.gla.ac.uk/~johnson

Abstract

The last decade has seen a growing number of cyber-attacks, for instance on Estonia, Belarus, Lithuania, Georgia, Pakistan and India. It has been difficult to determine whether or not these incidents were state-sponsored. This paper identifies three different roles that social networking and social media have played in this 'attribution problem'. Firstly, social networks have motivated individuals to participate in mass Denial of Service (DoS) attacks. They have disseminated information and provided access to resources, including botnets that were originally developed by cyber-criminal groups. Secondly, we show how information about an individual's social networks has supported targeted attacks, such as spear phishing, on opposition groups. Malware is, typically, disguised in a document that was intercepted from a colleague or friend. The recipient is more likely to open an attachment or link if it has been sent from a trusted source. Thirdly, we show how the development of Cloud infrastructures to support social networking applications has created disposable architectures for the Command and Control servers that coordinate malware attacks. The ubiquitous and distributed nature of these architectures makes it increasingly difficult to determine who owns and operates these systems. The closing sections of the paper identify a roadmap for the defensive measures that might be used to minimise the future threats from the 'dark side' of social networking[1].

Keywords: Cyber-defence, National Critical Infrastructures, Software Security.

1. Introduction

A small number of 'mass market' software infrastructures now support a broad range of critical systems. Linux variants have been used in applications that range from power distribution, to air traffic management and fire/rescue dispatch. The Internet Protocol now supports the

[1] This is an initial draft, comments and criticisms are very welcome.

dissemination of flood warnings, of electronic patient records and of food dispatch requests. The ubiquitous nature of these information infrastructures creates common failure modes. Attacks on critical software components have consequences that cross the borders between nation and industries borders.

Social Networking and the 'Attribution Problem': The vulnerabilities created by common, critical software infrastructures create particular concerns because we have seen a growing number of cyber-attacks. The following pages focus on seven different case studies, including:

1. Estonia, April-May 2007
2. Radio Free Europe and Radio Liberty, Belorusia, April 2008
3. Lithuania, June 2008.
4. Georgia, August 2008
5. China, GhostNet and the Shadow Networks, March 2009
6. Pakistan and Indian Cyber Armies, November-December 2010
7. W32.STUXNET, March 2010

In most cases, it is difficult to determine whether or not these incidents were state-sponsored. The following pages consider the role that social networking has played in this 'attribution problem'. Social networks describe the interconnections that exist between individuals and organizations over time. These relationships may be based on family ties, sexual relationships, employment etc. However, these interconnections can also represent the common beliefs and concerns that motivate groups to launch mass cyber-attacks in support of state policy. The following pages also argue that 'social media' increasingly provides state agencies with mechanisms for encouraging and sustaining these social networks. Social media are Internet-based applications that allow for the generation and dissemination of user-generated content. As we shall see, these applications have been used to encourage groups to launch cyber-attacks and have provided mechanisms for the dissemination and control of malware. We also show how these techniques support forms of indirection that make it difficult to identify explicit state responsibility for many cyber-attacks over the last decade.

Social Networking and Tailored Attacks: The analysis of social networking in cyber-attacks goes beyond the use of on-line forums to focus 'patriotic' support for state policy or their use in the dissemination of malware. In previous incidents, attackers have also used information about an individual's social networks to gain access to their accounts. Malware is, typically, disguised in a document that was intercepted from a colleague or friend. The recipient is more likely to open an attachment or link if it has been sent from a trusted source. The use of these techniques against several individuals in the same organisation does not, however, imply state sponsorship. Opportunity and motivation do not imply causation without more direct forms of evidence. As we shall see, implicit state support for cyber-criminal activities helps to sustain communities that can then be 'encouraged' to act against political rather than purely economic targets. However, the level of indirection supports 'plausible deniability' by state agencies even though they benefit from the intelligence derived in this way.

Social Networking and the Threat from Cloud Infrastructures: The third area of concern in this paper is that social media often depends upon Cloud architectures that support increasing levels of indirection in cyber-attacks. Cloud computing refers to the distribution of a computational service across multiple servers using common network infrastructures. Typically, the end use of the service does not download or run an application on their local machine but instead depends upon the execution of a program on a remote server using data that is stored across the Cloud infrastructure. Consequently, it may be difficult for the end user and even the host organisation to determine exactly where an individual's data is stored. These systems create significant opportunities for cyber-attacks, for instance when malware might send back data from a compromised machine to be stored in the Cloud for later collation. The low overheads associated with setting up an account also make these infrastructures ideal for 'disposable' Command and Control servers that can then be abandoned once an attack has been detected.

The closing sections of the paper identify a roadmap for the defensive measures that might be used to detect and dissuade future attacks of this nature.

2. Estonia, April-May 2007

The first case study focuses on a series of cyber-attacks that were directed against Estonia during April and May 2007. This was not the first example of a coordinated cyber-attack; for instance the US government identified the 'Titan Rain' attacks on companies, such as Lockheed Martin, national laboratories, including Sandia, and military targets, such as the Redstone Arsenal, well before 2003. Attackers from the Netherlands successfully compromised some 34 US defence sites in 1990. However, the political impact of the Estonian attacks justifies their inclusion.

The probable causes of this incident can be traced back to long running tensions between ethnic Estonians and Russians. These had been exacerbated by the agreement between Hitler and Stalin that led to the annexation of the country by Soviet troops in June 1940. After independence, many ethnic Russians were left without any form of citizenship following a series of disputes between the two governments. More immediate triggers included a decision to move 14 Soviet war graves and the Bronze Soldier of Tallinn war memorial. One justification for moving the statue was that it had become a focus for clashes between Estonians and Russian nationalists. The subsequent riots left one person dead and 150 injured. Some argued that the statue was a symbol of the Soviet occupation of Estonia while others argued that it was a memorial to the Russians who fought against the Nazi occupation (Tikk, 2009).

Estonia was particularly vulnerable to cyber-attacks because the country had embraced digital communications during the moves towards independence between 1988 and 1991. By the mid-1990s, Estonia had developed a comprehensive legal framework governing the IT industries. This was in strong contrast to the piecemeal legislation in many other European states. The development of Internet based infrastructures enabled Estonia to modernise national services whilst at the same time retaining the lowest ratio of government debt to GDP in Europe. The government supported the development of infrastructures that enabled the first on-line elections to be held; 80% of the population had national identity cards that supported digital signatures. These innovations had knock-on effects in supporting the development of on-line banking and retail systems. At the time of the attack some 97% of all Estonian bank transactions were made on-line and 65% of the population were classified as 'active users' of the Internet.

It is possible to distinguish several different phases in this incident. An initial emotional/physical stage was closely tied to the riots around the Bronze Soldier. A series of web-site defacements then took place but these were then superseded by a third phase in which Denial of Service (DoS) attacks were launched (Tikk, 2009). It is difficult to assess the impact of the DoS against Estonian infrastructures. One reason for this is that many different mechanisms were used. These included single individuals using primitive Internet Control Message Protocol (ICMP) floods. An ICMP flood (Ping flood or Smurf attack) sends large amounts of ICMP packets. The aim is to crash the target's TCP/IP stack to stop it responding to legitimate TCP/IP requests. The Estonian attacks also included more coordinated methods using botnets that would otherwise support Spam farms.

Arbor Networks issued reports on the extent of the attacks as they occurred using their Active Threat Level Analysis System (ATLAS). The company identified more than 120 different attacks over a three week period. Approximately 100 of these involved simple ICMP floods. 4 involved more sophisticated TCP SYN floods, where the attacker sends TCP connection requests faster than the recipient can process them. The analysis also traced the targets of these attacks. Most were aimed at government web sites, especially the Ministry of Finance (35 out of 128). The attacks gradually increased in number –from 21 on the 3rd May to 58 on the 9th May and then down to only one attack on the 11th. Most attacks lasted less than an hour but 7 went on for 10 hours or more. It is important to remember, however, that the bottlenecks and wider infrastructure problems created by massive network congestion would take far longer to resolve. The Arbor Networks study identified "a decent sized botnet behind the attack, with aggregate bandwidth at our points of measurement maxing out at nearly 100 Mbps… All in all, someone is very, very deliberate in putting the hurt on Estonia, and this kind of thing is only going to get more severe in the coming years"(Nazario, 2007).

The impact of these DoS attacks has, arguably, been over estimated. For instance, a Chatham House report argued: "The severity and length of the attacks directed at one of NATO's most electronically connected members put the alliance on guard. If a highly wired small state could be brought to its knees by a well-orchestrated Internet-based attack, then what type of havoc could be wrought upon larger states with more heterogeneous systems and critical infrastructure open to attack?" (Hughes, 2008). It is important to recall that the two largest banks, Hansapank and SEB Eesti Ühispank, together accounted for around 80% of the market. These were attacked several times between the 9th and the 15th May. However, the DoS attacks had limited effects; Hansapank's e‑banking service had to be shut down for up to 2 hours on 9th and the 10th May 9. SEB Eesti Ühispank's online banking service was offline for a similar period on the 15th May 15. The banks were able to limit the impact of the attacks by, for instance, limiting access to customers outside of Estonia. Telecommunications companies were also targeted. Three major Internet Service Providers – Elion Ettevõtted, Elisa Andmesideteenused, and Starman came under DoS attacks. Tikka et al also describe how mobile service operators experienced disruption (Tikk, Kaska and Vihul, 2010). The Zone.ee web hosting service provider and directory service provider (ee.ee) were attacked but in each case their services were restored relatively quickly. The US Computer Emergency Readiness Team (CERT) assessment concluded that although this was a 'watershed' attack, it was not 'revolutionary' in terms of the scale or sophistication of the techniques that were used (Waterman, 2007).

The Estonian Foreign Minister Urmas Paet accused the Kremlin of direct involvement in the cyber-attacks. However, there was no direct evidence to link the origins of the DoS with any state sponsored agencies. These attacks provide an example of the wider 'attribution problem' that stems from the attackers' ability to launch 'anonymous' attacks using features of the Internet Protocol. The same methods that provide resilience in the face of network failures also support subterfuge – for example, using different forms of server indirection.

The attribution problem is complicated because the Estonian attacks may provide an early example of ‘crowdsourcing’ in cyber-security. In much the same way that state agencies provided implicit support for the blockage of the Estonian embassy in Moscow, the same informal mechanisms were used to encourage distributed attacks from many different groups and individuals. This interpretation is supported by subsequent legal action against a student living in Talinn who was eventually fined around $1,000 for attacking the website of the Estonian Reform Party (Tikk and Kaska, 2010). Other attacks have been linked to the Nashi; a Moldovian youth movement. The incidents have, therefore, been likened to a ‘cyber-riot’ rather than a well-coordinated attack (The Economist, 2007). The terms ‘patriotic hacking’ and ‘hacktivism’ have also been used to describe the attacks “performed by a group of people who take action “pro patria” in cases where they believe that this is the right thing for their government to do or where they perceive the government as unable to do “the right thing” (Tikk, Kaska and Vihul, 2010). Political activists protested by engaging in coordinated cyber-attacks against the online presence and Internet infrastructure of Estonia.

Debate continues about whether the distributed DoS attacks from late April show more evidence of coordination than the earlier more random and arguably less effective attempts. In any event, the activities of disaffected individuals provided opportunities for ‘plausible deniability’ by a range of state agencies that might otherwise stand accused of coordinating the attacks.

Although the attacks had a significant impact, Estonia was not ‘brought to its knees’. In contrast, the Estonian response to the attacks reinforced their status as a leader in cyber-defence. The Chatham House report is, however, accurate when it argues that the political consequences outweighed the effects on economic infrastructures. Governments around the globe began to reinvigorate initiatives and infrastructures to counter potential cyber-attacks. NATO members held a meeting to consider the wider implications of the attacks on Estonia, leading to the development of NATO’s Cooperative Cyber Defence Centre of Excellence (CCDCOE) in 2008. The political significance of the attacks are illustrated by the decision to site CCDCOE in Tallinn. This organisation is not funded directly by NATO but from the supporting member states. CCDCOE has a wide set of objectives. It provides a meeting place for military experts, scientists, lawyers and politicians, albeit to a limited extent. These teams play a significant role in supporting a range of recent simulation exercises across member states. It also publishes technical reports on emerging cyber threats and provides ‘white papers’ on policy decisions including the development of international legal frameworks to clarify the response to future attacks, in particular when it is not possible to identify whether a state organisation is responsible for a threat.

The CCDCOE also acts as a ‘think tank’ identifying different ways of mitigating the threat from cyber-attacks. There are strong differences in policy between the supporting nations. For instance, the United States and the UK have significantly increased the funding of cyber-defence programmes. They have also created an array of over-lapping initiatives and agencies to help protect critical infrastructures. Some states have also proposed offensive capabilities that might deter other nations from supporting or encouraging future attacks. At the other end of this spectrum are nations whose cyber-policy focuses more on public awareness and the promotion of distributed resilience through public and commercial contingency planning.

The attacks on Estonia proved to be precursors for subsequent attacks against other NATO states. There were strong similarities between the events described in previous paragraphs and a series of attacks launched against Lithuanian sites in June 2008, see later sections. This was triggered by their suspension of discussions with Russia about an energy partnership with the European Union. More than 300 web sites were defaced in an attack that mirrored the first phase of the Estonian experience; in this case the content was replaced by the hammer and sickle symbol. All of the sites shared the same Internet Service Provider (Hostex); increasing suspicions that this was a coordinated attack. The attack originated in compromised machines located in France and Sweden (Tikk, 2009).

Some attacks were less public; including intrusions on government systems. In August 2007, Chancellor Merkel had to field reports of Chinese attacks on German ministries during a visit to China. In May 2008, the Belgian justice minister warned that similar attempts had been made against their ministries. In contrast, further attempts were made to compromise commercial targets. During November 2007, the head of MI5 issued a public warning that the People’s Liberation Army was supporting commercial espionage in the UK. In the United States, a series of working groups were created to consider the potential impact of cyber-attacks on key industries; with particular attention focussed on the banking and financial sectors.

3. Radio Free Europe and Radio Liberty, April 2008

Radio Free Europe/Radio Liberty grew out of the Cold War providing political counterpart to the emerging military strategies in the years after World War II. They had their beginnings in the work of the anti-communist National Committee for a Free Europe and the American Committee for the Liberation of the Peoples of Russia (Amcomlib) but were supported by the CIA. They were intended to provide a means of broadcasting information across the Iron Curtain. Radio Liberty focussed on the Soviet Union. Radio Free Europe broadcasts were intended for the satellite Soviet nations.

In 1972, CIA funding was halted and the broadcasters were then made more accountable through indirect financial support from the US Congress. They were also merged to form Radio Free Europe/Radio Liberty (RFE/RL). With the end of the Cold War, their mission extended to provide information from Eastern Europe, to Asia and the Middle East. RFE/RL were initially based in Munich, Germany, (1949-1995) before moving to Prague, Czech Republic. However, they also maintain a number of local bureaus that support broadcasts in more than 20 countries including Afghanistan, Iran, Iraq, Pakistan and Russia.

The broadcasters have always faced considerable technical difficulties in transmitting their signals. Some of these were technical; for instance, transmitters had to be located in Taiwan to reach the Eastern provinces of the Soviet Union. Other difficulties stemmed from the opposition of many state authorities, which introduced jamming techniques to prevent the populace from listening to the transmissions. There have also been more active threats to their work. For instance, several communist regimes deliberately sent agents to try and infiltrate the RFL/RE organisation. In most cases, the intention was explicitly not to interrupt the broadcasts but to gather information about their activities. However, there were direct attacks on RFE/RL staff. These included a Czechoslovak Intelligence Service attempt to poison food in the canteen as well as a car bomb that injured several employees and caused significant damage to the Munich headquarters in 1981. This was later linked back to the Romanian intelligence services.

The attacks associated with April 2008 were focused on Radio Svaboda, the RFE/RL service in Belarus. Radio Svaboda had been established in 1954 to broadcast each day in the native language. In 1988, their work was extended to support a web site on www.svaboda.org. Views differ over the political impact of these transmissions; however, Svaboda remained one of the few sources of external analysis even after Belarus gained its independence in 1991 (Tikk, Kaska and Vihul, 2010).

As a former Soviet republic, Belarus has arguably maintained the closest links with the Russian Federation. A situation which contrasts with that in Estonia, presented in the previous section, and in Georgia, discussed in the next section. The catalyst for the attacks stemmed from successive re-elections of the President Alexander Lukashenka. He was initially voted into power in 1994 then used referendums in 1996 and 2004 to extend his period in office. Many independent observers criticised the electoral processes that were used in these plebiscites. Criticisms have also been voiced about human rights issues in Belarus with Lukashenka being compared to a dictator from the Cold War era. In particular, the state maintained close control over the broadcast media.

During 2008, a new Statute on Mass Media extended previous regulations to cover the Internet information sources. This provided legislative support for the technical restrictions that had been implemented through the Beltelekom ISP which is run as a state monopoly. They were able to use their market position to restrict or deny access to sites that hosted material which was felt to be critical of the administration, including the Radio Svaboda site mentioned earlier.

The attacks began during the morning of the 26th April 2008. The servers supporting the Radio Savboda web site on behalf of RFE/RL were subjected to a distributed DoS; they were "inundated with about 50,000 fake pings every second, which the organisation reported as unprecedented in the history of cyber assaults against them" (Tikk, Kaska and Vihul, 2010). In the hours that followed the attacks extended to several other RFE/RL websites including those for Kosovo, Azerbaijan, Tatar-Bashkir. The attacks also targeted the web sites associated with Radio Farda in Iran as well as the South Slavic and Tajik services. Again, these attacks did not bring the service 'to its knees'. Several local web sites in Belarus offered to carry content for RFE/RL before they could recover even though this might have compromised their relationships with the state ISP. In most cases, the effects of the DoS were limited to 2-3 days. By 28th April most of the RFL/RE sites had been restored with significant improvements in their security infrastructures.

The RFE/RL attacks form a significant contrast to those in Estonia because their timing suggests a more prominent role for state coordination. The attacks took place on the 22nd anniversary of Chernobyl; which remains the world's worst civil nuclear incident. This accident occurred in the Ukraine on the borders of Belarus and caused widespread radiological damage leading to increases in cancer rates and birth defects for the local population. The anniversary of the accident had served as a focus for opposition supporters; many were unhappy about the lack of compensation for those affected by the accident. Others were angered by the administration's support for the Belarus civil nuclear program. Tikk et al (2010) argue that one motivation for these attacks was to disrupt Radio Svaboda's scheduled coverage of the protest rally; "Jeffrey Gedmin, the director of the RFE/RL, believed that the Belarusian government was most likely behind the attacks, describing the cyber-attacks as a weapon of dictators who were trying to prevent unfiltered news and information from reaching their people". However, this does not explain why the scope of the attack was then broadened to affect the other RFE/RL servers. Neither does this interpretation explain why Svaboda had experienced a similar minor attack on the previous anniversary of the Chernobyl accident. Again the attribution problem was complicated by official denials of any involvement in the DoS attacks. Motivation and opportunity cannot be used as surrogates for more direct evidence.

Arbor Networks argued that "a Russian language DDoS botnet" was involved in the attacks (Nazario, 2008). They also showed that the same farms were behind attacks on opposition media sites, including the Charter 97 news organisation. This is an English language news site about political activities in Belarus. Their www.charter97.org site remains one of the most popular portals in Belarus. It has attracted numerous cyber-attacks in recent years, hence the involvement of the same botnets that were used on the RFE/RL Belarus attacks is less of a coincidence that it might appear. Further attacks were focused on another popular news web site; BelPartyzan as well as Russian language new or information sites legis-group.ru and compromat.net.

The RFE/RL attacks are significant for a number of reasons. Firstly, they again illustrate the complexity of attribution both during and after cyber-attacks. Even though there are clear-cut motivations for state influence behind the attacks, this is insufficient to establish a causal relationship. There is also contradictory evidence, in particular any 'narrow' Belarussian focus cannot explain the other targets that were also affected in other states. Secondly, the Arbor analysis of the attacks on RFE/RL establishes a pattern to be repeated many times in this paper. The use of botnets in DoS attacks that would otherwise be used for criminal activities shows the blurred boundaries that exist between cyber-crime and cyber-defence. Some states may provide implicit support to these wider activities in the hope that that they might 'influence' their use as an extension of state policy. In this interpretation, DoS attacks can be triggered by informal links between the administration and local crackers. These groups might then act for patriotic reasons or for to ensure continued implicit state approval; in either case the relationships between state and cyber-attacker are implicit. This level of indirection has particular advantages for 'plausible deniability'.

The RFE/RL attacks also provide significant insights into the resilience of broadcast organisations and social media. In this case, sustained DoS attacks triggered rapid repair actions from the Belarusian site administration team. They were able to call on significant support not just from their parent organisation but also from a wide range of sympathisers both within the country but also from communities interested in protecting freedom of expression across the Internet. This response is a prototype for the collective action that has been seen during 2011 in what has been termed the 'Arab Spring'. The uprisings and unrest in Algeria, Bahrain, Egypt, Iraq, Jordan, Libya, Morocco, Oman, Syria, Tunisia and Yemen etc have all been influenced by the use of social media including Facebook, Twitter and YouTube. Attempts by many states to disrupt Internet communications have had significant success. Just as in Belarus, however, informal networks have helped opposition groups to continue disseminating their message, for instance using externally hosted sites. In some cases, this support also provided store and forward support for text messaging/voice communication linked to Twitter feeds to avoid Internet blackouts. We cannot be sure that this level of ingenuity will provide resilience against future attacks, however, the events of 2008 and 2011 illustrate the limited effectiveness of existing cyber-attack methods against the 'social resilience' of existing communications infrastructures.

4. Lithuania, June 2008.

The third case study occurred some two months after the attacks on the Radio Free Europe servers and focused on Lithuania. There are strong similarities between Estonia and Lithuania in terms of the impact of the attacks on their developing national economies. Lithuania had been pursuing a long-term development strategy based on the 'knowledge society' together with public safety and economic prosperity. The information infrastructure was seen as a key enabler of the knowledge society and so a 3-year Information Society Development Program was established in 2006. As in Estonia, this included the development of electronic interfaces between the civil population and local as well as national government. The development of the electronic infrastructure had been enhanced through legislative support for the use of electronic signatures, data protection, information security etc. The national objectives were supported by agencies including the Electronic Signature Monitoring Authority and an Information Society Development Committee.

These initiatives had resulted in significant growth in IT usage across Lithuanian society; around 30% of government organisations had a web presence. Almost half of all households possessed a computer prior to the attacks. Almost one third of households had an Internet connection. These statistics need to be set in context of a relatively low initial uptake. In other words, at the time of the attack Lithuanian was experiencing a rapid increase in the use of computational infrastructures although it had not yet reached the levels seen across many other European states.

Lithuania's recent history forms a strong contrast with that of its neighbour Belarus subsequent sections. Lithuania was the first former Soviet republic to declare independence. As in previous case studies, the longer term causes can be traced back before this declaration of March 1990. Between 1918 and 1940, Lithuania enjoyed a brief period of independence. However, they were first 'invaded' by the Soviet Union and then by Nazi Germany and then again by Soviet forces as the German's retreated during 1944.

It is against this background that a right-wing opposition party introduced a series of amendments in the Seimas, or Lithuanian Parliament in June 2008. The intention was to prohibit the public display of Soviet or Nazi insignia. It was also made illegal to play anthems associated with the former occupiers. The wide popularity of these measures contributed to the subsequent return of the proponents, the Homeland Union-Lithuanian Christian Democrats, as the largest party in the Seimas following the elections in November 2008. These events should also be seen as part of the wider Baltic tensions that have been described in the previous sections on the Estonian Cyber-attacks. In particular, local memories of the Soviet occupation created considerable dissonance with the rising self-confidence of the Russian Federation as they began to reassert a significant influence on their bordering states in the Baltic (Pääbo, 2008).

The immediate passage of the amendments did not seem to trigger the same level of hostility in the local Russian population, as had been observed amongst ethnic Russians in Estonia (Tikk, Kaska and Vihul, 2010). However, President Medvedev joined with the Belarusian President Lukashenka to express the Russian Federations concern over the legislation. The Russian Duma passed resolutions denouncing what they saw as the Lithuanian attempts to re-write the Soviet forces role in freeing the local population from Nazi oppression. There were simultaneous protests outside the Lithuanian embassy in Moscow.

The physical protests were accompanied by a significant number of cyber-attacks on web servers. Tikka et al estimate that some 300 sites were defaced by anti-Lithuanian slogans and by pro-Soviet symbols that would have been banned under the amendments being enacted by the Seimas. The attacks were first detected on the 28th June. The peak for attacks occurred between 17:00-18:00 EET on Sunday 29th June. Companies and government web sites were targeted ranging from the Lithuanian equivalent of the Securities and Exchange Commission to the web portal for the Lithuanian Social Democratic Party. By Monday 30th June 30, Rytis Rainys head of information security in the Lithuanian communications regulatory authority (Ryšiu Reguliavimo Tarnyba) warned that "network administrators should not relax" even if the scale of the attacks seemed to be falling (Tikk, Kaska and Vihul, 2010). However, by 2nd July most of the sites were able to resume normal operation.

The vulnerabilities that were exploited in the attacks helped to focus the impact on commercial organisations. Part of the reason for this is that Lithuanian government agencies had been warned to expect an attack after an Estonian TV channel and journal had reported the possibility of a cyber-attack on Baltic States on the 26th July. The Lithuanian Computer Emergency Response Team (CERT-LT) responded by warning the government agencies on its list of 'at risk' organisations. Commercial organisations did not receive the warning and hence remained at risk. CERTs have been established by commercial and research organisations across many nations. Initially, information was exchanged between these private bodies through the European Task Force for Collaboration between Security Incident Response Teams (TF-CSIRT). Increasingly, however, it has been recognized that there is a need for states to support national CERT/CSIRTs to act as a hub for national policy formation and for the coordination of any response to adverse events including cyber-attacks. The Lithuanian CERT was established just over eighteen months before the web servers were defaced. As with many of these bodies, it was initially set up to support the education sector, however, its role had steadily expanded with the development of national strategic plans for ICT. The lack of warning and consequent impact on commercial organisations from the Lithuanian cyber-attack illustrates the need to create and maintain social media that support potential victims of future incidents.

Commercial organisations criticised the lack of warning in the aftermath of these incidents. However, the Lithuanian attacks raise deeper questions about the relationship between state and private organisations during cyber-attacks. .For instance, it may be difficult to issue future warnings if they compromise intelligence sources. Conversely, these warnings may become less effective over time if there are false alarms. Other questions relate to financial liability for future attacks in which the state does not disclose relevant information or conversely for warnings that turn out not to be justified.

In this case, the advanced warning to government agencies provided by CERT-LT helped to mitigate the impact of these attacks. There were “no real danger to the private sector or strategic State administration was determined, although the cyber-attack did disturb the Hostex-operated servers which contained the Internet sites of several hundred companies” (Tikk, Kaska and Vihul, 2010). However, there was recognition that Lithuania has been fortunate. In other circumstance, they might not rely on receiving information from media outlets in other Baltic states. A key lesson was the need to improve their monitoring of Internet sites for better intelligence about potential attacks.

CERT-LT also played a prominent role in the forensic investigations that identified the attack mechanisms. They found that most of the web sites that were attacked had been hosted by Hostex web servers. This company was one of the largest web hosting services in Lithuania. The pattern of attack suggested that once the vulnerability had been identified, all of the Hostex users were targeted. There does not seem to have been a premeditated list of specific sites against which the attacks were directed. CERT-LT provided specific assistance to Hostex engineers as they worked to mitigate the effects of the attacks.

Initial reports focused on vulnerabilities in PHP servers that were then used to run a series of scripts distributed amongst a small group of attackers. However, other modes of attack included the generation of large amounts of spam email linking to "Hackers United against External Threats to Russia". This urged that an initial attack should be extended across the Baltic and to the Ukraine (Tikk, Kaska and Vihul, 2010). Further attacks were, however, focussed on Lithuania. These included what appears to have been a DoS incident involving the Lithuanian Tax Office almost a month after the original attacks. This incident illustrates the *uncertainty principle* in cyber-attacks. The damage done in an initial attack cannot simply be quantified in terms of the costs of restoring operations. There is also a continuing overhead in terms of the uncertainty over future attacks. In consequence, many potential e-commerce and government sites observed considerable traffic reductions as end users became increasingly concerned over the vulnerability of their transactions.

As in the majority of attacks described in this paper, it was extremely difficult to identify the source of the attacks. CERT-LT and RRT, the Lithuanian communications regulatory authority, identified the use of proxy servers to the East. It seemed "likely that the attacks were organised in advance, considering the fact that signals, invitations, and agitation were spread on the Internet prior to the attacks. However, the RRT refused to speak of any particular countries as the initiator of the attack" (Tikk, Kaska and Vihul, 2010). Subsequent analysis identified compromised hosts in France and Sweden as possible candidates for the servers used to launch the defacement attacks. The difficulty of recreating the attack vectors has also been used to argue that these attacks were carefully planned some time in advance of the 28th June. Alternatively, it might be that these attacks were using mechanisms that had already been well rehearsed for other purposes, including lower profile cyber-crime activities.

Attempts to identify the source of the attacks have focussed on Russian 'cracking' groups. This is justified given the convergence of national sentiment with the amendments in the Seimas. Circumstantial evidence is provided by discussions on forums including hack-wars.ru. However, direct proof was difficult to obtain. For example, the Cyber Police Unit of the Lithuanian Police argued that the attackers had used the Tor "onion routing" protocols. Tor is used by journalists, whistle-blowers, law enforcement agencies to prevent monitoring of Internet traffic. It relies on 'onion routing', where 'onion' refers to the use of encryption at multiple layers of a communication network. The system also relies on a volunteer network of sensors to help 'disguise' communications through the various intermediaries. As we shall see, this infrastructure was also implicated in the GhostNet attacks on the communications infrastructure of the Tibetan movement. The presumed involvement of the cracking groups in these attacks rather than state agencies is also illustrated by the Lithuanian response. The Police treated the incident as a cyber-crime conducted against Hostex.

The Lithuanian attacks illustrate a number of common themes across many of the incidents in this paper. The problems of attribution were complicated by the lack of direct evidence for state involvement even if there is evidence of significant planning and some coordination. The involvement of 'cracking groups' communicating over Internet forums and social media partly explains these observations. Implicit state support for the development of cyber-criminal communities can help to create or prepare the infrastructures that can eventually be used to launch cyber-attacks of the form experiences in Lithuania and Estonia.

The failure by CERT-LT to pass on warnings to commercial organisations also demonstrate the consequences that arise when state agencies fail to fully exploit the same communication channels that were employed in launching an attack. In other words, the exclusion of particular groups from the network of dissemination created vulnerabilities that exacerbated the consequences of the defacements. The lack of warning and consequent impact on commercial organisations from the Lithuanian cyber-attack illustrates the need to create and maintain social media that support a wide range of potential victims during future incidents.

5. Georgia, August 2008

The fourth case study forms part of the armed conflict that took place between Georgia and the Russian Federation. The conflict centred on Ossetia, a region that straddles the Caucasian mountains but whose population share ethnic and linguistic ties. The geographical features were used to justify a division of the territory in 1922 when North Ossetia was ceded to the Soviet Union while the area South of the mountains became part of the Georgian Soviet Socialist Republic (GSSR). The GSSR was established after Bolshevik forces had suppressed the Democratic Republic of Georgia in 1921. Ethnic tensions continued between Georgians and Ossetians, especially during the dissolution of the Soviet Union. By 1990 Southern Ossetia achieved a form of 'de facto' independence. This did little to reduce the tensions, while the international community continued to recognise the region as formally part of Georgia.

The Georgian declaration of independence in April 1991 provided a further catalyst for wider ethnic disputes. During transition to independence, Georgia´s first President, Zviad Gamsakhurdia, alienated the populations of Abkhazia and South Ossetia by using nationalist slogans such as "Georgia for the Georgians". Fighting borke out between Georgian forces and separatists in South Ossetia (1991 – 1992) and then in Abkhazia (1992 -1994 ended with Georgia losing control of large parts of both territories. In consequence, the Organization for Security and Co-operation in Europe (OSCE) formed a peacekeeping force in 1992 combining troops from Russia, Georgia and South

Ossetia. The OSCE is a security oriented inter-governmental organisation supported by almost 60 states formed under the United Nations Charter. It was created during the Cold War era as an East-West forum but its role has evolved since the end of the Soviet Union. However, the composition of the peacekeeping force did little to reduce animosity between the various factions (European Commission, 2009).

There was growing concern in Moscow about US support for Georgia under President Saakashvili; military aid was provided by the Bush administration to help Georgian forces clear Chechen fighters from the Pankisi valley. Tensions were also increased when the 2002 Russian Law on Citizenship opened ways for South Ossetians to obtain Russian passports. By April 2008, there was a significant worsening in relationships between Russia and Georgia as fighter aircraft and UAVs began to patrol the supposedly 'demilitarised' areas. Russian engineers also began work on improving rail infrastructures for 'humanitarian purposes'. Large scale military exercises were conducted on both sides of the border. Georgia also argued that there were a rising number of attacks on their villages carried out by members of the South Ossetia militia,

On 7th August 2008, Georgian forces began an artillery bombardment and then attacked the town of Tskhinvali and surrounding areas. The attack was justified by various figures inside the Georgian government as an attempt to restore the proper 'constitutional order' and then later as a response to an alleged Russian invasion. By 8th August, Russian forces had responded by moving into Georgian territory. Their Air Force began to attack central Georgia, extending the conflict to the Senaki military base and the naval port of Poti. Dual purpose infrastructures were targeted, including Tbilisi airport radar, railroad junctions and telecommunications facilities.

The Russian and Ossetian advance continued and many thousands of civilians were displaced by the fighting. By 10th August, the Georgian Government declared a unilateral ceasefire and announced its intention to withdraw Georgian forces from South Ossetia. Fighting continued so that by the 10th-11th August, Russian troops had crossed the boundaries of Ossetia and entered a number of Georgian towns. On 12 August, French President Sarkozy travelled to Moscow and Tiblisi, representing the European Council. He proposed a six point peace plan that ultimately led to a ceasefire. By 22nd August, Russian troops had withdrawn from their positions beyond the former boundaries of South Ossetia and Abkhazia. Georgian sources claimed that 170 servicemen, 14 policemen and 228 civilians were killed with 1,747 wounded. The Russian Federation claimed that 67 of their personnel had been killed with 283 wounded. In total around 850 people, including civilians were killed in the conflict.

The initial wave of cyber-attacks on the Georgian government infrastructures occurred well before the armed conflict with ICMP floods and large volumes of spurious HTTP 'GET' requests being used in July 2008 (Nazario, 2008a). However, the main effects began to be felt late on the 7th August, before the Russian and Ossetian forces had fully mobilised. There are significant differences between the DoS and web-defacements of August 2008 and those experienced by Estonia just over twelve months earlier. In this case, it was less easy to disguise the links between and cyber-attacks and external state intervention. There are further differences. The previous sections have described how Estonia had made rapid advances in terms of Internet uptake since its independence. More than 50% of the population were described as active users. In contrast, this figure was close to 30% in Lithuania but only 10% in Georgia. However, this percentage had been growing rapidly in the months prior to the conflict.

The resilience afforded by the relatively low levels of Internet usage has to be contrasted with the vulnerabilities that were created by physical constraints on Internet routing into Georgia. More than half of the 13 interconnections passed through Russia (Zmijewski, 2008). However, the level of traffic did not necessarily reflect the majority physical routing. Most of Georgia's Internet prefixed were routed via Turkish or Azerbaijan service providers. This can partly be explained by the influence of increasing tensions with Russia and growing trade opportunities with the EU prior to 2008. Even so, the limited routing options open to Georgian telecommunications providers formed a strong contrast with the high-capacity links between Estonia and Finland, Sweden, Latvia and Russia. The Georgian telecommunications market was also significantly smaller than its Estonian counterpart with only five companies offering Internet access and more than three quarters of the traffic being met by Caucasus Network Tbilisi. Estonia exploited her more diverse infrastructure to implement a range of contingency plans that could be used to spread the load when DoS attacks began to overwhelm particular service providers. In contrast, Georgia was aware of the communications vulnerabilities prior to the conflict. Work had almost been completed on a high capacity link with Bulgaria before the invasions began.

The attacks on Georgia were remarkably similar to those used against Estonia. They included the defacement of public web sites and DoS attacks. The targets included www.president.gov.ge, where images of Adolf Hitler were substituted for pictures of President Saakashvili. A collage of dictators was also placed on the National Bank of the Republic of Georgia's site. Other targets included the portal provided by the Ministry of Foreign Affairs as well as newspaper and media agencies both in Georgia and Azerbaijan. Meanwhile, the DoS attacks were aimed at government sites, including the Parliament and Ministry of Education and Science. They also attacked news and media sites. These included on-line communities, the web-sites of English language newspapers and the Associated Press in Georgia. DoS techniques were also used against commercial banks. A sustained attack was also launched on www.hacking.ge; arguably undermining a potential source of information about the nature and extent of the attacks. Russian sites also seem to have been the target of secondary attacks, such as that operated by the RIA Novosti news outlet (Nazario, 2008b).

The DoS statistics gathered during the Georgian attacks were similar to those witnessed during the Estonian case study. Most attacks lasted around two hours with the longest sustained attempts reaching some six hours in duration. There is no conclusive proof about the origin of the attacks. However, many intelligence agencies had learned from the experience of the previous year and were better equipped to monitor the attack patterns. The widespread distribution of the coordinated attacks suggested that several botnets must be involved. One line of analysis identified the use of a MachBot controller for an HTTP-based botnet attack, a signature of Russian herders. However, it can be dangerous to assume that a well-known pattern of attack need necessarily provide reliable evidence about the actual identity of any coordinator (Tikk, Kaska, Rünnimeri, Kert, Talihärm and Vihul , 2008). At least one of the botnet Command and Control servers was physically located in the United States.

The evidence of organisation cannot easily be distinguished from the impact of 'crowdsourcing' that we have identified in the Estonian attacks. In this instance, a batch file was distributed on a number of web sites that executed DoS attacks on Georgian servers. This was then passed on between sympathisers for the Ossetian cause. Other sites, including stopgeorgia.ru and stopgeorgia.info provided attack tools and a list of 36 Georgian sites that were identified as priotiry targets. The net effect was similar to what Tikk et al describe as the 'emotional phase' of the Estonian attacks where 'a downloadable script to ping flood Estonian websites (both DNS and IPs) was shared on various Russian language message boards. Instructions on how to ping flood Georgian government web sites were also distributed on Russian language websites and message boards, as well as lists of Georgian sites vulnerable to remote SQL injections, facilitating automatic defacement of them' (Tikk Kaska, Rünnimeri, Kert, Talihärm and Vihul, 2008). SQL injections occur because commands written in one language may be hidden inside the instructions of another. If the input is not correctly filtered for embedded command sequences then these can be inadvertently executed by higher layers of an application.

The Georgian conflict extended beyond the defacing of web sites and the use of DoS attacks. Several sources have suggested that the email addresses of prominent Georgian politicians, academics and industrialists were distributed to encourage spamming and more targeted attacks (Tikk Kaska, Rünnimeri, Kert, Talihärm and Vihul, 2008). These tactics were also used in Estonia; where every individual has the right to communicate with the government using electronic means. The effect was to force major figures to create new email accounts and re-establish contacts in a piecemeal fashion, at a time when their attention was being demanded by many other priorities. Further disruption may have been caused by attempts to alter Internet routing tables; it has been alleged that route requests from the United States on Georgian government sites were showing that access was blocked via TTNEt (Turkey). Many aspects of the attacks in Estonia and Georgia have been linked to the Russian Business Network (RBN) criminal organisation and associated botnets. The accuracy of this claim is understandably difficult to establish, as is the nexus between the founder of the RBN and Russian government agencies. It may be that the RBN provided access to botnet resources without mounting the DoS attacks; following a 'business model' that supported their wider activities. 5 stages can be identified in the coordination of the Georgian attacks:

1. Spreading encouragement to get involved in the cyber war against Georgia;
2. Publishing a target list of Georgian government Web sites which have been tested for access;
3. Selecting types of malware to use against the target Web site;
4. Launching the attack and optionally,
5. Evaluating the results and iterating previous stages. (Grey Goose, 2008)

It can be argued that this level of coordination is entirely different from the cyber 'riots' seen in Estonia. However, it does not automatically follow that a government agency should be explicitly involved in leading any one of these different phases. The lack of explicit evidence to connect Russian state agencies with either the Estonian or the Georgian attacks is very significant and too easily overlooked. Too often commentators use these incidents to illustrate elaborate theories about the future of 'cyber-warfare'. However, they often over-emphasise the explicit influence of state agencies and arguably neglect the importance of small groups and individuals during these attacks. There seems to have been a widespread understanding that the DoS attacks and defacement activities were at the very least ‘tolerated’ by state agencies even though there do not seem to have been any official announcements to that effect. Some states might deliberately choose to show a degree of tolerance for a wider range of more conventional cybercrimes; phishing attacks, identity fraud etc. This helps to establish the ‘anti-social networks’, that provide the infrastructures, such as those operated by the RBN, and the agents who might then initiate strategic cyber-attacks. In contrast, the response to the attacks was organised in a more centralised fashion by the Computer Emergency Readiness Team (CERT Georgia) that normally served the Higher Education sector. They assumed national responsibilities and began to coordinate with their Polish counterparts to analyse IP data and with CERT France on data collection. Further assistance was provided by technical experts from CERT Estonia.

Many of the victims of the cyber-attacks simply did nothing except to wait for them to subside. Other sites temporarily changed their IP addresses while others changed hosts, sometimes with support from international hosting services. There was considerable sympathy for many of the organisations that were attacked; some were seen as the victims of an attack on freedom of expression over the Internet. In consequence, accounts were provided by blogspot.com for both Georgian news agencies and for government ministries. The websites of the Georgian Ministry of Defence and the President were relocated to a company in Atlanta, Georgia, USA. Further symbolism was provided when the Ministry of Foreign Affairs moved their web site to an Estonian server. These counter-measures had limited success for users inside Georgia. Several key routers went down under the volume of traffic. In particular, the flood attacks on Caucasus Network Tbilisi led to rerouting that had knock-on effects for the smaller service providers. The impact of the DoS attacks was exacerbated by the physical damage to network infrastructures that ran through the war zone, recall that many of the Georgian external connections ran North into Russia.

It is clear that these attacks had different effects on Estonia and Georgia, however, neither country was ‘brought to their knees’. It has been argued that these two case studies illustrate a correlation between ICT maturity and the vulnerability to cyber-attacks. The more advanced nations will suffer most from the loss of these infrastructures because other countries retain alternate mans of communication and coordination. In the Estonian attacks, key government and financial services were affected for a wide proportion of the population. This was not the case in Georgia where Internet usage was more limited. Alternatively, it can be argued that countries with a higher level of ICT maturity are better placed to develop alternate network infrastructures/routing when an attack has occurred. Less mature nations find it harder to bring their electronic infrastructures back on line. In Georgia, the attacks forced the suspension of electronic banking services for ten days from the 9th August.

There were significant psychological consequences for both attacks. These psychological effects contrast with recent events during the 'Arab Spring' when governments have sought to disrupt communication between its citizens and the outside world. In Georgia, they centred on the loss of key communication channels between the government and its citizens. The psychological impact was exacerbated by the presence of Russian troops and Ossetian militias inside the country. In both cases, the attacks also severed important communications channels between the governments and international organisations. This made it difficult to pass information to broadcast agencies. Tikk Kaska, Rünnimeri, Kert, Talihärm and Vihul (2008) observe that "while the attacks did not have a permanent or even a long-run devastating effect on the Georgian Internet infrastructure, the damage caused by the attacks was most acutely experienced at the time when Georgia was the most dependent on the availability of their information channels". The timing of the attack was ideal; showing the vulnerability of the Georgian government to external forces.

6. China, GhostNet and the Shadow Networks, March 2009

This paper takes a deliberately broad view of CyberDefence. The Estonia case study illustrates the manner in which political tensions can trigger the use of DoS attacks in a relatively uncoordinated way. These tensions establish and reinforce the 'patriotic' social networks that facilitate further cyber-attacks. The RFE/RL DoS incidents provide further examples of the attribution problem. Even though the Belarusian state authorities had motivation and opportunity to coordinate the attacks; causal connections cannot be established without more direct evidence. The Georgian incidents also show how DoS attacks can be supported by social media and, to a certain extent, can be directed through implicit means as an arm of state policy. In contrast, the fourth case study focuses on the use of cyber techniques to support longer term espionage that suggests greater forms of coordination and planning but which, nevertheless, illustrates the importance of social media in the creation of vulnerabilities and new modes of arrack.

The name 'GhostNet' is used to describe a long term and large scale programme of cyber espionage uncovered by a group of Canadian security specialists in March 2009 (IWM, 2009). The extent of the network is believed to have affected more than 100 countries. Targets included foreign ministries in Iran, Bangladesh, Latvia, Indonesia, Philippines, Brunei, Barbados and Bhutan; embassies in India, South Korea, Indonesia, Romania, Cyprus, Malta, Thailand, Taiwan, Portugal, Germany and Pakistan; the ASEAN (Association of Southeast Asian Nations) Secretariat, SAARC (South Asian Association for

Regional Cooperation), and the Asian Development Bank. The attacks also compromised computers in news organizations and a range of commercial organisations. There was no evidence to suggest that machines in the United States or the United Kingdom had been targeted, however as mentioned in previous sections, both governments had issues warnings about previous Chinese cyber-attacks.

There are reasons to believe that GhostNet was coordinated and controlled by the People's Republic of China as part of the wider policy of active defence. The Information Warfare Monitor report that initially identified the threats associated with GhostNet observed that "Chinese authorities have made it clear that they consider cyberspace a strategic domain, one which helps redress the military imbalance between China and the rest of the world (particularly the United States). They have correctly identified cyberspace as the strategic fulcrum upon which U.S. military and economic dominance depends" (IWM, 2009). The Information Warfare Monitor is a public-private partnership between the University of Toronto and the Secdev group. They reiterated comments made in previous sections of this paper that state coordination cannot be assumed behind all of the malware that originates from a particular nation. In the case of China, the sheer numbers of Internet users and the economic potential from cyber-crime make it likely that many forms of attack will stem from criminal activities. Also in common with the previous there is no definitive proof to support Chinese state involvement. There are also significant differences with the Estonian and Georgian attacks. In the previous incidents, social networks in the broader sense helped to multiple the impact of DoS attacks and web site defacement. In contrast, GhostNet provides significant insights into the vulnerabilities that are created by the rise of social networking. Information about an individual's on-line contacts can be exploited by 'socially engineered malware'.

The starting point for the forensic investigations that eventually led to the identification of GhostNet was the concern that a modus operandi had emerged in which a Trojan horse was introduced onto a victim's machine that would then disclose information to control servers, ultimately reporting back to Chinese sources. It was difficult to find evidence to support these concerns; however, obvious targets included the Tibetan community, the Falun Gong, the US and Taiwanese Governments and a number of large corporations engaged in key areas of Africa and the Far East.

One of the employees at the Information Warfare Monitor had personal links into the Tibetan community and was, therefore, ideally placed to coordinate a more detailed investigation of any potential attacks. Tibetan groups had raised concerns about information security as early as 2002. The Office of the Dalai Lama does not have any secrets, however, it coordinated many aspects of the Tibetan movement since the Chinese invasion. It, therefore, receives many time-critical communications of a sensitive nature. In particular, it handled diplomatic communications with the Tibetan Government in Exile and correspondence with their supporters. The IWM analysis began by collecting evidence that confidential information had been leaked from within the Tibetan Government in Exile and the private office of the Dalai Lama. In this first phase, network monitoring software was also installed in a number of machines to determine whether there was evidence of malware. A real-time packet capturing program captured all traffic from each machine. It was eventually observed that the compromised systems tried to connect to control servers to execute further instructions, including commands to access instructions on other servers.

An initial analysis was conducted on the machines in Dharamsala, India and at Tibetan missions. However, the bulk of the analysis was conducted during a subsequent second phase. This "led to the discovery of four control servers and six command servers. These control servers were identified and geo-located from the captured using a simple IP lookup. The control servers were then probed and web-based control interfaces were identified on four control servers, which allowed us to view and control the network. The system was actively monitored for two weeks, which allowed us to derive an extensive list of infected systems, and to also monitor the systems operator(s) as the operator(s) specifically instructed target computers" (IWM, 2009). This work was supported by further analysis and site visits by researchers from the University of Cambridge, documented in Nagaraja and Anderson (2009).

The IWM investigators were unable to establish the mechanisms that were used to infect machines in the Dalai Lama's office. However, the subsequent report argued that the attacks made extensive use of 'social malware' (Nagaraja and Anderson , 2009). A tailored email would be sent to the target to increase the likelihood that an attachment would be opened – the example given is a message from campaigns@freetibet.org. The success of the approach depended upon convincing the recipient to open the attached document containing the malware. Hence the terms 'social malware', 'social engineering' and 'spear phishing' describe the use of information about the target's network of contacts. The aim is to increase the plausibility of the email vector. In some cases, legitimate messages were simply reused. In other cases, documents were taken from other compromised machines and then passed to new targets with the embedded malware to increase the probability that they would, in turn, infect their machines. The Cambridge group also argued that legitimate emails were intercepted and legitimate attachments were then infected before the message was forwarded across the network (Nagaraja and Anderson , 2009).

Many of these were coordinated by the 'ghost RAT' (Remote Access Tool), this extended the scope of the control that was possible be enabling limited forms of remote dynamic reconfiguration. Only one third of the thirty anti-virus programs that Information Warfare Monitor tested were able to recognise the malware embedded in a sample Word document. The Word Trojan is a useful example because it illustrates some of the common mechanisms that were embedded in these different variants. Once executed, the program performed a DNS lookup to locate and then connect to its control server. Most of these were associated with IP addresses that had been assigned to China, although the IMW report identified others in the USA, Sweden, South Korea and even Taiwan – again increasing the difficulty of attribution and allowing for plausible deniability. This example was only one of eight different varieties of Trojan horse that were deployed against the Tibetan communities. Others used up to four different IP addresses using slightly different methods to establish communication with control servers. Most of these different variants of malware used the Hypertext Transfer Protocol (HTTP) to disguise communications as 'normal' activity. In one case, commands appear to have been embedded within JPEG image files that were exchanged using a PHP server. In another case, the malware used the regular HTTP POST command to send information to a CGI script hosted on the control server. IWM argued that one type of connection was used for coordination while the other was used for the transfer of information.

The initial GhostNet report provides the following example of the ways in which information might have been used against Tibetan sympathisers. Drewla relies on young Tibetans with Chinese language skills to spread information about the plight of their country using Internet social media to people in mainland China; "A member of Drewla, a young woman, decided to return to her family village in Tibet after working for two years for Drewla. She was arrested at the Nepalese-Tibetan border and taken to a detention facility, where she was held incommunicado for two months. She was interrogated by Chinese intelligence personnel about her employment in Dharamsala. She denied having been politically active and insisted that she had gone to Dharamsala for studies. In response to this, the intelligence officers pulled out a dossier on her activities and presented her with full transcripts of her Internet chats over the years. They indicated that they were fully aware of, and were monitoring, the Drewla outreach initiative and that her colleagues were not welcome to return to Tibet. They then released her and she returned to her village" (IWM, 2009). The Cambridge study uses similar accounts to suggest that the Chinese were behind the Ghostnet. IWM found that many of the control servers are associated with IP addresses that can be located on the island of Hainan, which is the home to the Lingshui signals intelligence facility and the Third Technical Department of the People's Liberation Army. Unfortunately, the are very few simple solutions to the attribution problem in Cyberdefence. Many members of Drelwa use Tencent QQ; this is the most widespread instant messaging service in mainland China, to spread information about the Tibetan cause. Other sympathisers had used a Chinese variant of Skype. Both of these infrastructures had attracted the interest of the authorities. For example, pressure continues to be placed on QQ operators to disclose the identities of individuals who set up chat rooms and groups. Police interest is justified as part of national attempts to fight pornography. However, it is also possible that Tibetan supporters have been compromised through the authorities observation of social media rather than through the Trojans that formed part of the GhostNet.

The attribution problem is further complicated by investigations that were triggered by the IWM report. Scott Henderson (2009) maintains a web site and blog about Chinese 'crackers'. The term 'cracker' is used rather than hacker – to avoid any suggestion that these individuals are poor coders. He initiated a search on the emails associated with some of the web sites involved in the coordination of the GhostNet Trojans. Rather than linking back to government agencies; these seems to be associated with individual members of the Chinese 'cracker' community that could be identified using QQ numbers. Nart Villeneuve, a member of the IWM teams and Chief Security Officer at the SecDev Group, responded to Scott's work in his blog; "Questions regarding those who are ultimately responsible for this cyber-espionage network remain unanswered. We were, however, able to benefit from a great investigation by The Dark Visitor who tracked down lost33, the person who registered some of the Shadow[2] network's domain names that we published in the GhostNet report and his connections to the underground cracking community in China. Based on the IP and email addresses used by the attackers we were able to link the attackers to several posts on apartment rental sites in Chengdu. This, of course, does not reveal the role of these specific individuals nor the motivation behind the attacks. However, the connection that The Dark Visitor drew between lost33 and the underground hacking community in China does indicate that motivations such as patriotic hacking and cybercrime may have played a role. Finally, the nature of the data stolen by the attackers does indicate correlations with the strategic interests of the Chinese state. But, we were unable to determine any direct connection between these attackers and elements of the Chinese state. However, it would not be implausible to suggest that the stolen data may have ended up in the possession of some entity of the Chinese government" (Villeneuve, 2010). A more sustained analysis of Henderson's work was published in the IWM's follow-up report to the original GhostNet analysis (IWM, 2010).

The subsequent analysis of the IWM report is significant because it draws clear parallels with the investigations that followed on from the attacks against Georgia. In both cases, it was impossible to identify direct state involvement. Links could, however, be identified with members of the Russian

[2] The term 'Shadow' was coined by IWM in a follow-up investigation discussed in the closing paragraphs of this section, it is significant because it refers to the influence of social media and Cloud techniques in cyber attacks.

and the Chinese ‘cracking’ communities. The subsequent discussions also illustrate the ways in which blogs and other forms of social media support the analysis of Cyber-attacks; Henderson picked up on elements in the IWM report and was able to use his Chinese linguistic skills to extend the analysis, reinforcing key themes in the initial study. Very similar dialogues can be seen in the analysis of the Georgian attacks, mentioned earlier.

There are further parallels between the discovery of GhostNet and the previous attacks described in this paper. For instance, the Estonian and Georgian attacks have motivated a recent study by Tikk, Kaska and Vihul of the legal implications of cyber-defence (Tikk, Kaska and Vihul, 2010). They point out a number of areas of concern:

1. There is a need to cover the spectrum of cyber incidents in national laws and also to ensure international cooperation and coordination;
2. There is a need to clarify the freedom of expression and state responsibility to support it;
3. There is a need to consider the threat of politically motivated attacks by non-state actors, as well as governmental duties toward the private sector regarding threat warnings; and
4. There is a need to reflect the importance of cyber defence in different fields of law (law of armed conflict, criminal law, and legal regulations for the ICT industry).

The individuals responsible for the IWM report faced particular dilemmas that provide examples of several of the broader concerns in the report by Tikka et al. They were unsure how to disclose information about the extent of the attacks that they had uncovered. For instance, it was possible that legal processes might have been used to force them to disclose the extent of their discovered (e.g. IP addresses of compromised machines) to groups that could them have abused the information. In particular, this information would help any agencies involved in the attack to determine the extent of the investigation. IWM, therefore, contacted the national Canadian Cyber Incident Response Center (CCIRC), who then assumed responsibility for contacting the organisations that had been compromised by the GhostNet attacks.

The GhostNet investigations were subsequent extended by the IWM in directions that are particularly important for the topics in this report. These ‘Shadow 2.0’ studies described how malware was “to a large degree organized and operated through the misuse of social networking and cloud computing platforms, including Google, Baidu, Yahoo!, and Twitter, in addition to traditional command and control servers. Second, although we are able to piece together circumstantial evidence that provides the location and possible associations of the attackers, their actual identities and motivations remain illusory. We catch a glimpse of a shadow of attribution in the cloud, in other words, but have no positive identification. The 2.0 designation also contains a double entendre: it refers to a generational shift we believe is unfolding in malware networks in multiple dimensions, from what were once primarily simple to increasingly complex, adaptive systems spread across redundant services and platforms, and from criminal and industrial-based exploitation to political, military, and intelligence-focused espionage" (IWM, 2010). The follow-on report describes how many of the domain names that were previously used to control the GhostNet Trojans were abandoned within days of the first study being published. IWM working closely with the Shadowserver foundation, therefore, acquired them to monitor any further attempts to connect with them by any other compromised machines. Shadowserver is a volunteer led site for the investigation and exchange of information on cybercrime. With their additional expertise it was possible for the IWM investigators to identify two different forms of attack on the Office of His Holiness the Dalai Lama's computer systems. The subsequent analysis indicated the involvement of two and possibly three different groups using different command and control infrastructures for these different forms of malware. These infrastructures are critical for the success of any sustained attack because they provide the mechanisms for directing and coordinating the exfiltration of information from the compromised systems.

The second IWM report focused on one of these infrastructures that they termed the Shadow network, in contrast to the GhostNet described in previous sections. Shadow "is a complex network that leveraged social networking websites, webmail providers, free hosting providers and services from some of the largest companies on the Internet as disposable command and control locations" (IWM, 2010). Blogs, newsgroups and social networking sites were used because they helped to mask the malicious nature of the information transfers. If the connections had been detected then it is likely that they would have been dismissed as legitimate and benign activities; consistent with the activities of the staff using the systems to coordinate the Tibetan movements. These Cloud based services offered further benefits; they were disposable. Free or low cost hosting services with minimal background checks offered huge advantages for locating command and control services for malware operations. If the free hosting sites had to be abandoned then malware could check for a new location using social networking sites. Twitter, Google Groups, Blogspot, Baidu Blogs, and blog.com accounts were used to post a list of updated links to more free hosting accounts. Alternatively, if this mechanism failed the malware could refer back to more stable servers in China; these would then point the malware back to an alternate benign and disposable hosting service reducing the chances of further detection. 27 different malicious binaries were identified during this second phase of the investigation. To illustrate the general approach to command and control, two of these used Yahoo! Mail accounts. Periodic checks were sent by the malware to these addresses to indicate that they were still functioning. The same accounts were also used to send additional malicious binaries to the infected machines. Google Pages was also (ab)used to host malware; "the use of social networking platforms, blogs and other services offered by trusted companies allows the attackers to maintain control of compromised computers even if direct connections to the command and control servers are blocked at the firewall level. The compromised computers can simply be updated through these unblocked intermediaries to point to a new, as yet unknown, control server" (IWM, 2010). None of these techniques were novel. Arbor Networks' Jose Nazario (2009, 2009a) had previously found both Google Ap pages and Twitter being used for command and control.

The use of the Tor anonymity network provides a further illustration of the techniques uncovered in the subsequent study, indicating the potential involvement of more than one group in the attacks. This is the same system that was identified by the Lithuanian Police Cyver-Crime Unit in the investigation of the attacks launched during June 2008. As mentioned, Tor is used by journalists, whistle-blowers, law enforcement agencies etc to prevent monitoring of Internet traffic. It relies on 'onion routing', where 'onion' refers to the use of encryption at multiple layers of a communication network. The system also relies on a volunteer network of sensors to help 'disguise' communications through the various intermediaries. Dan Egerstad monitored the data exiting the Tor network and found some of it related to the Dalai Lama's office. If the exit point of a Tor network is not encrypted then it can be accessed in plain text – and anyone can, in principle, operate a Tor exit node. He argued that Tor was unlikely to have been used by the Dalai Lama's staff and that it was most probably being forwarded as part of the indirection employed by a group attacking their servers.

IWM also identified a further attack mechanism associated with the Enfal Trojan. This provided important forensic information. It is not widely available and suggested clear links with a number of crackers who either initiated the attack or who provided expertise to the groups responsible for deploying the malware. The subsequent report argued that such exchanges were characteristic of the vibrant "cracker community' in the People's Republic of China with informal links to various government agencies.

One common theme regarding attribution relating to attacks emerging from the PRC concerns variations of a privateering model, in which the state authorizes private persons to perform attacks against enemies of the state. This model emerged because studies have shown that there is no direct government control over the loosely connected groups of hackers in the PRC" (IWM, 2010). If these groups did infiltrate the Tibetan community then it is likely that information was passed to one of the (rival) state agencies in China. It is ironic that the crackers probably face many of the same dilemmas that arose when IMW passed their findings to the Canadian Cyber Incident Response Center (CCIRC). Both arguably acted to support national self-interest without specific coordinateing instructions from state agencies.

The impact of GhostNet and Shadow can be summarised in the findings of the second report. They identified the enormous benefits derived from Cloud computing, peer to peer techniques and from social networking. However, as these techniques become more deeply engrained in personal, business and government computing it is clear that "these new platforms are also being used as vectors of malware propagation and command and control" (Johnson, in press).

7. Pakistan and Indian Cyber Armies, November-December 2010

It is difficult to provide detailed insights into the mechanisms used by more recent attacks because forensic investigations must still be completed. With these caveats in mind, it is possible to sketch some of the mechanisms that have been used in a number of recent attacks between groups in India and Pakistan. As with the other examples in this report, it is important to understand the political context against which these disputes arose. Pakistan is located at the intersection of land routes between the Middle East, South and Central Asia. By the time that the British East India Company came to increase their dominance of the region during the C18th, they found a culturally diverse but predominantly Muslim population. The failure of successive armed uprisings against the British led to the non-violent protests led by the Indian National Congress. However, this was mirrored by the development of an All India Muslim League who were concerned about their lack of representation.

By 1940 there were calls to create two nations from the Indian sub-continent; one predominantly Hindu and the other predominantly Muslim. By 1947, nationalist leaders including Nehru, together with representatives of the Muslim League, such as Jinnah, as well as representatives of the Sikh community had agreed to the shape of modern India following independence. This led to the formation of Pakistan in August 1947. However, almost immediately conflicts arose over the former provinces of Punjab and Bengal. Many Muslim were forced to flee across the border into Pakistan while Hindus moved in the opposite direction. Further conflict focussed on the Kashmir region where the local Hindu ruler had seceded to India. Since that time, India and Pakistan have fought four major wars (1947, 1965, 1971, 1999). There have also been numerous border clashes. Both sides have been accused of supplying arms to local insurgent groups.

These conflicts have taken place during a period that has seen a significant expansion of communications infrastructures across both India and Pakistan. However, the distribution of Internet access has been very uneven; reflecting significant geographical and socio-economic divisions. India has the world's largest population living below the World Bank poverty line and the most rapid increase in telecommunications subscribers. Both countries have addressed these imbalances – for instance, Pakistan's National ICT R&D Fund offers incentives for undergraduate students from rural areas. In consequence, both countries have a large, comparatively young population of well-trained engineers.

Previous case studies have focused on the attribution problem that complicates attempts to link explicit state intervention with cyber-attacks. However, more direct forms of intervention have been used to control domestic access to Internet resources. In particular, the Pakistan government has restricted access to anti-Muslim material. Almost all domestic traffic is routed through the Pakistan Internet Exchange, which is operated by a state-owned company. Attempts to restrict access within the country to blasphemous YouTube material created routing problems that denied access to the video servers across the globe (Johnson, in press). It was only after the event that it became clear this was the result of a mistake rather than a deliberate attack on the address space of the video site.

Internet censorship has also been enforced in India. In previous years, this has arguably been less organised and sustained than in Pakistan. However, the Mumbai attacks of 2008 helped to focus attention on the role of Internet sites that were perceived to have played a role in the early planning of these atrocities. Subsequent amendments to the Indian Information Technology Act supported the State monitoring and blocking of Internet forums and social networks. The mechanisms for doing this are largely centred on licensing terms for Indian ISPs; requiring them to filter requests to sites that are banned. The same Act also made it an offence to attack ICT infrastructures both inside India or overseas; punishable with up to three years imprisonment or substantial fines or both. These provisions have subsequently created problems for the Indian National technical research organisation and Defence Intelligence Agency as they have worked to establish strategic, cyber-offensive capabilities.

The Indian CERT (CERT-IN) was established in 2003, as the government realised the growing strategic importance of their domestic ICT industries. There were particular concerns about the threats posed by cyber-attacks from Pakistan. CERT-IN has considerable powers – certainly greater than those associated with similar bodies identified in the previous sections of this paper. In particular, they can require ISPs to take action to protect critical infrastructures, whereas organisations such as CERT-LT have a far more advisory role. In Pakistan there are several organisations providing CERT services including PAK-CERT and CERT-Pakistan. The former is a commercial organisation whereas the latter is a non-profit group sponsored by a parent company. Both promote the sharing of information about previous attacks and offer ‘members services’ to disseminate warnings about potential incidents.

In both India and Pakistan, military organisations have also taken an increasing interest in cyber warfare. In 2008, the Indian army organised an internal review of cyber-security. This led to the creation of specialist units that were attached to each division; their responsibilities include monitoring conformance to digital security operating procedures. At a more strategic level, the Army Cyber Security Establishment (ACSE) was directed to conduct regular ISO 27001 audits to ensure that the armed forces could anticipate new threats and vulnerabilities over time. They were also tasked to support the tri-service integrated defence staff where information infrastructures extended across naval and airborne operations. These initiatives were justified by the investments into cyber-defence being made by both China and Pakistan. However, successive internal reports have stressed the slow pace of progress within the Indian army compared to the rapid changes that have characterised the civilian ICT industries. In other words, there is a concern that military systems might continue to be vulnerable to forms of attack that have already been identified in civil systems. Partly in consequence, the Indian army has established their own internal CERT building on the precedents created in civilian organisations. As might be expected, these initiatives have triggered similar responses from inside the Pakistan military. In particular, the Army has created close links with the Cybercrime division of FIA (Federal Investigation Authority) to address growing concerns over potential vulnerabilities to future attacks.

It is against this background of continuing conflict between Pakistan and India that a series of cyber-attacks have taken place. These again follow the pattern of ‘patriotic hacking’ and ‘hacktivism’ described in the previous case studies. However, the initial incidents triggered successive reprisals; demonstrating the future potential for the escalation of cyber conflicts as attackers target different areas of their ‘enemies’ national infrastructures. The cyber-attacks can be traced at least as far back as 1998. They were triggered by a series of nuclear weapons tests conducted by both India (Operation Shaki) and Pakistan (Changai-I). Shortly after India had announced the Shaki test, a group of Pakistani activists known as milw0rm managed to deface the web site of the Indian Bhabha Atomic Research Center.

These sporadic attacks continued into the Kargil War, including an infamous defacement of http://www.armyinkashmir.com. This illustrated the potential military and political impact. The web site provided information and news about the Indian army in Kashmir. However, it was altered to show what appeared to be Indian forces killing Kashmiri militants. This led to a series of reciprocal attacks on Pakistani sites. Some of these attacks used spoofing. For example, the

‘Patriotic Indian’ group registered pakgov.org in opposition to the official government site on pakistan.gov.pk. Other notable Indian groups active from the late 1990s include the Hindustan Hackers Organization (H20). Their attacks have been countered by groups including the Pakistan Hackers Club (PHC) and G-Force.

These reciprocal attacks have continued at varying levels of intensity for almost 15 years. In May 2010, the Indian Cyber Army launched a further series of ‘hacktivist’ defacement attacks on Pakistani sites. This triggered attempts to deface more than 1,000 Indian sites by a range of different groups including PakHaxors, TeaMp0isoN, ZCompany Hacking Crew etc. Their targets included the web sites of government agencies, including the Indian Criminal Investigation Department (CID), as well as commercial sites such as the Box Office of Indian. They also attacked the public portals of military organizations; exposing limitations in the audits and doctrine described in previous paragraphs.

In November 2010, the Indian Cyber Army attacked a series of high-profile web sites, including those belonging to the Pakistan Army. Other targets included the Council of Islamic Ideology as well as a number of agencies associated with the Pakistani ICT infrastructures, Although these attacks were partly in response to the previous wave of defacements, mentioned above, they were also influenced by popular unrest in India following renewed claims that the 2008 Mumbai attacks were organised by Pakistani groups. This led to further defacements by the rival Pakistan Cyber Army in December 2010. They followed up the earlier attacks on the Indian CID by targeting the Central Bureau of Investigation (CBI). Subsequent enquiries revealed that the site had not had a full security audit in the twelve months since it had been brought on-line. This illustrates the difficulty of ensure the security of the hundreds of potential targets that could be attacked across dozens of government agencies.

These attacks have affected sites ranging from the Pakistan’s Oil and Gas Regulatory Authority to the website of the Jadavpur University Department of Economics. In many cases, the targets seem to have little political or strategic importance. In most cases, service was restored within hours or days.

These reciprocal attacks have extended beyond Pakistani and Indian sites. In particular, groups in Pakistan have been implicated in attacks on US-Israeli sites to support the Palestinian cause. Just as in previous attacks, it can also be difficult to distinguish clearly between examples of patriotic hacking and wider forms of cyber-attack. For instance, it has been alleged that the Pakistan Cyber Army obtained customer data from a number of commercial web sites. Subsequent statements by the group made it clear that this was not done for criminal reasons but to illustrate the potential vulnerability of European servers and that the personal information was destroyed after the companies were informed.

8. W32.STUXNET, March 2010

W32.Stuxnet is a malware program with many different components, intended to reprogram Programmable Logic Controllers (PLCs) within industrial control systems (Falliere, Murchu and Chien, 2011). It, therefore, forms a considerable contrast to the DoS attacks and Trojan horse techniques identified in previous sections of this report. The majority of previous case studies have focused on private email systems and broadcast web sites. In these applications there has been a continuing concern to identify security threats and patch known vulnerabilities; even if some end users have not used this information. In contrast, Stuxnet was intended to target control systems that have often relied upon security through anonymity. PLCs have not been significant targets for wider forms of cybercrime; their specialist nature also implies that the potential attackers must understand how these devices work. It is important to recognise that Stuxnet is not the first malware to target industrial control systems. However, the sophistication of the combined methods of attack together with the potential target of the malware makes it a significant landmark in cyber-defence. The worm was first reported in mid-June 2010; its name was derived from keywords discovered in the software.

The complex nature of the malware together with the generic problems created by attribution for cyber-attacks makes it difficult to derive conclusive proof about the target of the attack. However, it is widely assumed that Stuxnet was created to inflict damage on the Iranian nuclear rporgramme. This began in the1950s with assistance from the United States. However, Western support ended with the 1979 Iranian Revolution. This created a hiatus in the civil energy programme until the commissioning of the Bushehr I reactor, which was launched in August 2010. In the interim, however, the UN Security Council became increasingly concerned over the existence of an Iranian military nuclear programme. They passed six resolutions between 2006 and 2010. One aim behind these was to impair the enrichment of uranium to the concentrations required for military applications. For instance, resolution 1803 (March 2008) banned the export of dual use technologies with nuclear applications to Iran.

The enrichment process supports isotope separation, for example separating natural uranium into enriched uranium and depleted uranium. The enrichment process is essential to develop uranium fuel for civil energy generation but also for nuclear weapons development. Typically, the process relies on cascades in which each stage produces successively higher concentrations. One cascade approach relies on strings of centrifuges; heavier isotopes are displaced towards an outer radial surface. This approach is believed to have been employed by several countries, including Pakistan, to support the accelerated development of nuclear weapons.

Iran had already announced its ability to enrich uranium up to 3.5 % using a string of more than 100 centrifuges in 2006. The International Atomic Energy Agency (IAEA) found evidence that Iran had created at least two strings capable of enriching uranium up to 20% by August 2010. They also argued that Iran had more than 2.5 tons of low-enriched uranium which would be sufficient to support the development of nuclear weapons using these strings of centrifuges.

Stuxnet contains a number of different forms of attack. These include up to four zero-day exploits; vulnerabilities that exist between the time at which they are first exploited and the time when a fix is applied. The interval during which any vulnerabilities remains will depend in part upon the comparative effectiveness of the social and professional networks that have been developed between attackers and defenders. If information about potential vulnerabilities is not effectively distributed to users then zero day vulnerabilities will persist, increasing the exposure to any attack as time goes on. The number of zero day attacks in Stuxnet is highly unusual. Each vulnerability has a considerable 'market value' to attackers while it remains unpatched. To use four in the same malware suggests a significant investment of time and effort in ensuring the success of the attack. Other unusual aspects of the malware include the range of programming languages that seem to have been used during its development, including C and C++, suggesting the integration of ideas from a team of programmers working together in a coordinated manner.

Stuxnet also included a Windows rootkit, providing privileged access to the underlying operating system while at the same time masking the intrusion. There is also evidence that Stuxnet successfully exploited the first rootkit for PLCs. Part of the attack included hooking code that was designed to intercept operating system calls. Other elements included process injection, network dissemination and replication techniques as well as a command and control interface with superficial similarities to the approaches used to monitor and control the GhostNet vectors.

The diversity and complexity of the attack can be illustrated by the different techniques that were used to spread the malware – these included the use of vulnerabilities in the auto-execution mechanisms associated with removable drives, such as USB sticks. Stuxnet also exploited vulnerabilities in the Windows Print Spooler as well as Windows RPC mechanisms to replicate over local area networks building on vulnerabilities already exploited by the Conficker worm.

These different approaches are necessary because PLCs are not, typically, connected to networks. Instead they are periodically attached to PCs and code is downloaded directly on the controllers. In consequence, attackers must first infect the machines that are used to develop the PLC application. In the case of Stuxnet, the malware focussed on PLCs that were programmed using Siemen's Step 7 Industrial Control Software, for reasons that will be discussed in later sections. If the Step 7 environment was not detected on the PC then no actions were taken; hence the malware was aimed at a very specific call of systems even though the transmission vectors were intended to support widespread infection.

Stuxnet enabled additional instructions to be inserted into the relevant code before it is downloaded onto the PLC. This is a non-trivial task – for instance, it is likely that the developers of Stuxnet would require access to the target hardware so that they could test early versions of their malware. This, in turn, would require knowledge of the systems architectures, either via inside sources or through information obtained by earlier versions of the malware. In addition, the attackers used two different digital signatures to authenticate drivers in the malicious binaries. These signatures belonged to two Taiwanese companies, JMicron and Realtek. Some have argued that these digital signatures could only have been obtained through other forms of industrial espionage (Matrosov, Rodionov, Harley and Malcho, 2010).

Three different variants of the malware have been identified (June 2009, March 2010, April 2010). Each modification introduced new infection vectors to increase transmission of the virus. Each variant took measures to disguise the infection. One approach that was used to hide the source of the infection was for Stuxnet to delete any malware left on a USB stick after the third infection. Variants of the malware were also programmed to delete themselves on the 24th June 2012. Another technique involved scanning the registry for indicators that a range of security software had been installed. In some cases, the malware would inject code into trusted processes, including anti-virus software. However, if the security product was considered to be resilient to the attack then no injection would be attempted. Where necessary, it would employ one of two different techniques to escalate process privileges in order to acquire sufficient rights to complete the injection attack.

The Sematec team that monitored the Stuxnet command and control servers found that the initial infections were clustered in Iran, Indonesia and India. However, they ventually found connection requests from over 40,000 unique IP addresses in more than 155 countries. Of these, approximately 60% were located in Iran, of this 60% of total infections it was found that 68% involved machines operating Step 7. These command and control servers exploited http and were located in Malaysia and Denmark on www.mypremierfutball.com and www.todaysfutball.com. However, various countries have since taken measures to block these servers. In consequence, there have been no notifications from Iran after 22nd August 2011. Considerable sophistication can again be seen in the engineering of these interactions. For instance, the malware would first attempt to communicate with a valid address, including www.msn.com, to test connectivity. It was also possible for additional code to be retrieved from the C&C servers for subsequent execution on the infected machines.

Stuxnet exploited a zero-day exploit in the Step 7 SCADA database software in the form of a hard-coded database password. Once the malware had infected a host, it focused on Step 7 project files used to program the PLC devices, mentioned above. It used this by hooking operating system calls to open project files associated with this application. However, it did not attempt to infect all data. It was only interested in projects that had been used or accessed within the previous 3.5 years, it ignored the Step 7 example projects and those that did not contain any meaningful code. The virus will also check if a project has already been infected. If this is the case then it will update the infection if the present version of the virus is newer than that associated with the existing project.

Once a target had been infected, the malware would have been downloaded onto the PLC. The objectives for the infection are to enable Stuxnet to monitor, replace and disguise blocks of code being written to the PLCs. In particular, the code looks for PLCs with 6ES7-315-2 (series 300) CPUs. Further steps are taken to monitor Profibus communications processors looking for data associated with more than 33 frequency converter drives manufactured by Fararo Paya in Teheran, Iran and Vacon in Finland. These devices are used to control the speed of motors and other industrial equipment. The changes to the code would have incurred significant physical damage on the equipment being controlled by these devices.

Stuxnet monitors the frequency of the attached motors and only attacks systems that normally operate at a frequency between 807 Hz and 1210 Hz. This covers a diverse range of motors including those used in industrial pumps and gas centrifuges. If the intended target is detected then the malware triggers a state machine to control the 'sabotage'. After an initial delay of some 13 days, the state machine in Stuxnet will set the maximum frequency to 1410 Hz, after a further delay of some 27 days this is changed to 2 Hz and then 1064 Hz before the sequence is repeated. This helps to disguise the infection as the converters are slowed and then speeded up beyond normal operating limits. Falliere, Murchu and Chien (2011) argue that this could also have resulted in considerable 'colateral damage' beyond the specific devices that were the target of the Stuxnet attacks. The rootkit that is installed on the PLC masks the operation of the malware to disguise the reasons for changes in rotational speed from monitoring systems.

As mentioned before, Stuxnet was designed to focus on PC's using Siemen's Step 7 environment. The company has reported that the malware did not cause any damage to any of its registered customers. They responded to forensic studies of the malware by distributing a detection and removal tool; some of these affected the password vulnerabilities mentioned in previous paragraphs. Siemens also advised the installation of various Microsoft security patches and the enforcement of polies banning the use of unauthorised USB drives. The development and validation of these changes was complicated by the specialist nature of the PLCs and the possibility of cross-infection.

Even though Stuxnet did not damage the equipment operated by Siemen's registered customers, suspicions remain that the intended target was part of the Iranian nuclear programme. This uses centrifuges that require the control devices, which were the focus of the infection. However, this equipment falls under the UN Security Council embargoes that have been described in the opening paragraphs of this section. Hence, it is likely that the developers of Stuxnet had reason to believe the Iranian programme relied on unauthorised applications of the Siemens technology. This, together with the sophistication of the attack methods, raises suspicions that Stuxnet is the product of state programmes. These suspicions have focused on the relationships between Iran, the United States and Israel. For example, there is circumstantial evidence of joint work by Siemens and the Idaho National Laboratory to identify vulnerabilities in the Step 7 suite that was then exploited by Stuxnet. In particular, it is likely that the malware was intended to disrupt the Natanz enrichment facility, which suffered a series of 'technical difficulties' during 2009-10. The Iranian President Mahmoud Ahmadinejad subsequently confirmed that these problems were related to the operation of the centrifuge strings.

It has been argued that Stuxnet is only one part of a wider state-sponsored campaign to disrupt the Iranian nuclear programme. This view is supported by a series of car bomb attacks on senior Iranian scientists and defence officials. Symantec has estimated that Stuxnet would have taken from five to thirty people around six months to prepare (Falliere, Murchu and Chien, 2011). However, this does not preclude the possibility that a small group of committed individuals acted together to create the malware with or without state encouragement. Many of the subsequent studies that identify the necessity of state involvement rely almost entirely on circumstantial evidence and inference. They also underestimate the importance of social networking and state endorsed forms of 'patriotic hacking' in the previous attacks studied in this report.

Previous sections have argued that cyber-attacks seldom have the medium and long term effects that many commentators have feared. The same caveats can also be raised about the impact of the Stuxnet malware. The IAEA reported signs of damage to around 900 centrifuges during the period when Stuxnet was active. However, this seems to have been dismantled and replaced over a period of months rather than years. As mentioned previously, the Iranian authorities acted to block access to the command and control servers. Anti-virus software was recommended, however, the Iranian

Technology Council became concerned that this might itself harbour the Stuxnet virus. As we have seen, these concerns were well founded. However, independent studies confirmed by the IAEA suggest that Iran was able to increase its ability to enrich uranium even in spite of the problems that may or may not have been caused by Stuxnet (Amarello, 2011). One interpretation of this conflicting evidence is that the Iranian authorities deliberately exaggerated the impact of the malware to create the impression that greater damage had been inflicted on their nuclear programs. Other commentators have argued that the preoccupation with Iranian centrifuges has masked the real purpose of the Stuxnet virus (Gaycken, 2010). In this interpretation, the attacks represented a limited field test of a targeted cyber-attack. The widespread dissemination of the virus, with only limited targets was used to determine how easy it would be to transmit the virus across a range of different systems/security cultures. The real purpose was to use this information so that future weapons might have more widespread strategic targets during a potential cyber-war.

9. Roadmap for Social Media in Cyber-Defence

The previous sections in this report have provided an overview of seven different attacks. In particular, the intention has been to highlight the different roles that social media and social networking have played in many different incidents, including:

1. Estonia, April-May 2007
2. Radio Free Europe and Radio Liberty, Belorusia, April 2008
3. Lithuania, June 2008.
4. Georgia, August 2008
5. China, GhostNet and the Shadow Networks, March 2009
6. Pakistan and Indian Cyber Armies, November-December 2010
7. W32.STUXNET, March 2010

However, this analysis is of little benefit in itself unless it can be used to develop a roadmap for future intervention. The following sections identify a number of different ways in which potential targets could increase their resilience to these diverse forms of attack.

Social networking has played many different roles in the incidents that we have studied. The opening sections of this paper identified three primary mechanisms – the use of social networking to inspire patriotic attacks that complicate the ‘Attribution Problem’; the use of social network to infer inter-personal relationships that can then be exploited by malware, for example using spear phishing techniques and the threats posed by Cloud Infrastructures that often underlie social networking applications, for example to provide ‘disposable’ command and control servers.

The previous case studies have provided numerous examples of these concerns. For example, both Georgia and Estonia show how social networking provided patriotic groups with both the motivation and technical direction to implement distributed DoS attacks. In both of these incidents, the attribution problem was complicated because it is impossible to determine whether or not direct state intervention played any role in sustaining the groups who ultimately launched the attacks. In both these cases, it is clear however that the lack of state intervention to tackle other forms of cyber-crime helped to establish the resources, such as the bot-nets, and the knowledge necessary to launch these attacks.

Previous sections have also provided examples of the use of social media to obtain personal information that can inform further attacks. It seems clear that many different intelligence agencies are now using advanced information retrieval techniques to monitor interaction over existing sites, just as the QQ feeds may have been used to track the activities of the Drelwa. Groups of crackers may also be using similar techniques to target attacks on individuals using malware similar to that employed during the GhostNet and Shadow attacks. In both cases, politically motivated cyber-attacks have exploited the same targeted techniques used for identity theft in wider areas of cyber-crime.

Similarly, our case studies have identified the use of Cloud infrastructures that often support social networking as an important mechanism in the attack vectors that have been used against nation states. . The most straightforward approach of using the web servers to coordinate elements of the Stunext attacks can be contrasted with the use of ‘disposable’ command and control servers to support indirection and dynamic reconfiguration by both the GhostNet and Shadow attacks. Similarly, the role played by anonymity services, such as the Tor system, in both the Lithuanian attacks and the GhostNet malware, provides a further area of concern.

These observations can be brought together to develop an overview of the roles that social media have played in previous attacks, illustrated in Figure 1. For instance, we have seen how groups of ‘crackers’ have exploited social networking in at least five different ways to coordinate attacks both across and within national borders. Firstly, these groups have used social media to encourage others to get involved in cyber conflicts. Secondly, social media has been used to disseminate lists of potential targets, including government sites that have already been probed for potential vulnerabilities. These target lists may also include the identities of individuals vulnerable to spear-phishing, mentioned in previous sections. Thirdly, community forums have been used to discuss, select and tailor specific malware applications to use against these selected targets. Fourthly, social media have been used to launch these attacks, most notably through C&C and anonymity servers. Finally, social networking has helped to debrief participants in an attack, to evaluate lessons learned and to support iterations through each of these previous stages.

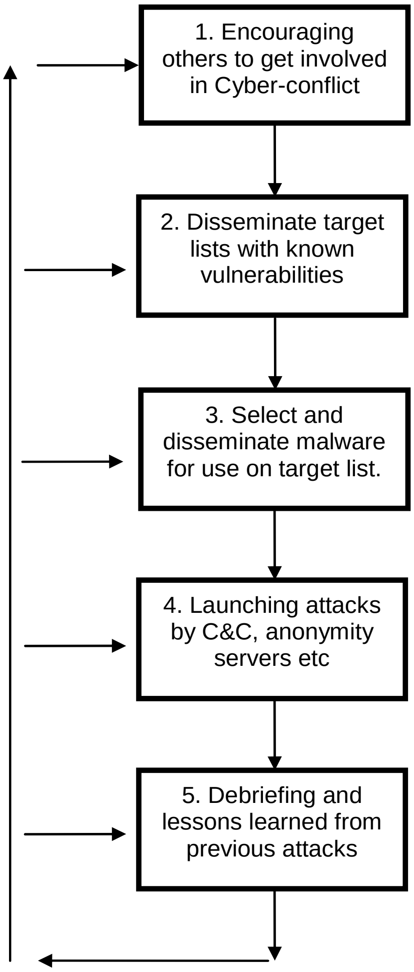
1. Encouraging others to get involved in Cyber-conflict
2. Disseminate target lists with known vulnerabilities
3. Select and disseminate malware for use on target list.
4. Launching attacks by C&C, anonymity servers etc
5. Debriefing and lessons learned from previous attacks

Figure 1: Different Stages in the Use of Social Networking for Cyber Attacks

The high-level overview of the use of social media in cyber-attacks can be used to construct a roadmap for future engagement in social media by security and intelligence services. For instance, we have seen how the Lithuanian CERT was alerted to the possibility of an attack on a Baltic state by an Estonian television channel and regional journal. These sources, in turn, were aware of the use of social networking to recruit potential participants in a cyber-attack during stages 1 and 2 in Figure 1. The Lithuanian CERT arguably was fortunate to receive the warning from the media so that they could warn government agencies of the increased risk. However, in the future more might be gained from a closer monitoring of these groups directly rather than relying on reporters to gather the source data that is then used to increase national resilience. Similarly, as Cloud computing techniques become more and more prevalent it seems critical that state agencies become more sophisticated in their appreciation of the risks; especially considering the role that some of these infrastructures have played in previous attacks. This implies a more detailed analysis of stages 4 and 5 of Figure 1.

Previous paragraphs have argued that the case studies in this report can be used to provide generic insights into the role that social media has played in previous cyber-attacks. It has also been argued that these insights can be used to develop a more coherent roadmap for security and intelligence services. Such an outline can help both to improve our understanding of the threats posed by the use of social media and also to consider the opportunities that social networking tools provide to increase our resilience against future incidents.

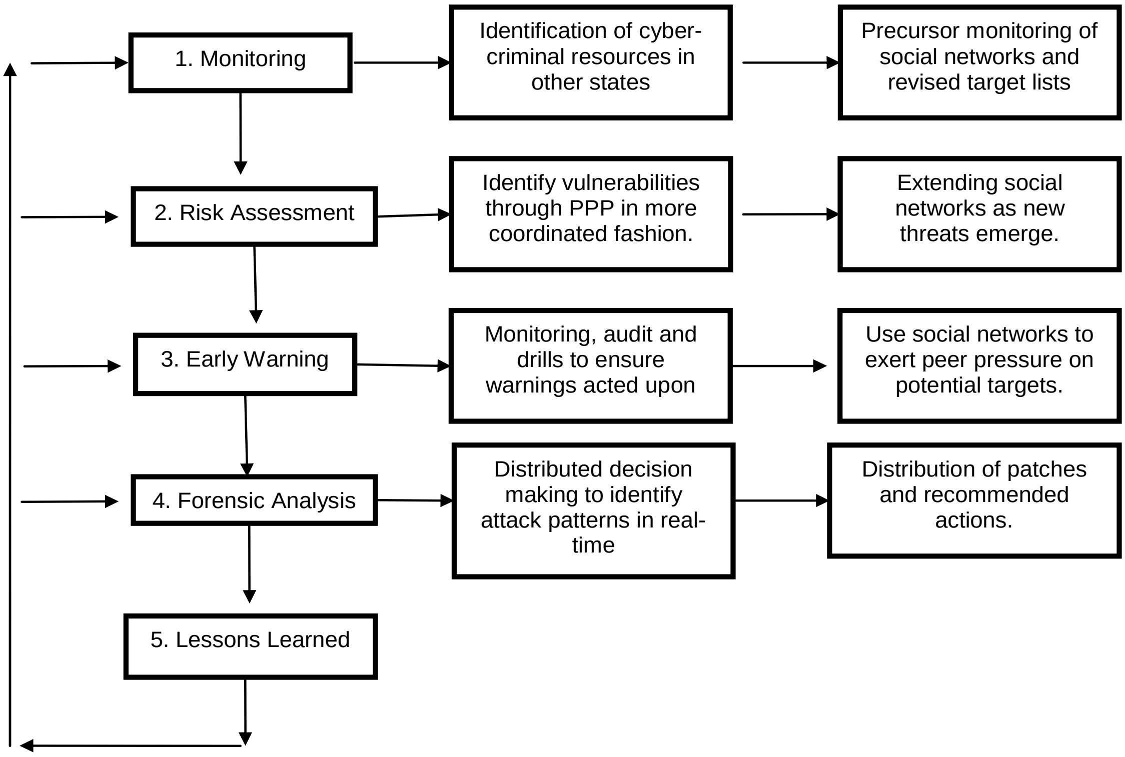
1. Monitoring
Identification of cyber-criminal resources in other states
Precursor monitoring of social networks and revised target lists
2. Risk Assessment
Identify vulnerabilities through PPP in more coordinated fashion.
Extending social networks as new threats emerge.
3. Early Warning
Monitoring, audit and drills to ensure warnings acted upon
Use social networks to exert peer pressure on potential targets.
4. Forensic Analysis
Distributed decision making to identify attack patterns in real-time
Distribution of patches and recommended actions.
5. Lessons Learned

Figure 2: Roadmap for the Use of Social Networking for Cyber Defence

Figure 2 provides an initial sketch of a roadmap for the use of social media to help mitigate the threats posed by cyber-attacks. It is more complex than the overview of social media in previous attacks, illustrated in Figure 1. This is because the first diagram was based on evidence drawn from the sustained case studies presented in this paper while the elements of Figure 2 are hypothesised interventions to mitigate the risks of future attacks. Hence they require further analysis and justification. It should also be stressed that this is a starting point for future work in this area.

As can be seen, the first element of the roadmap focuses on the monitoring of social media. This can be broken down into a number of concerns. For example, previous sections have described how the majority of 'politically motivated' cyber-attacks exploited techniques and resources that were initially developed to support more general forms of cyber-crime, including identity theft. These range from the bot-nets that used to attacks Estonia and Belarus through to the day-zero techniques and root kits that were employed by Stuxnet. All of these resources and techniques had been discussed on a range of social media well before the attacks even if these for a did not explicitly discuss the exact mechanisms that were used. Initial work has already begun to use automated information retrieval techniques to identify social networks sharing an interest in these resources – the aim is to automatically identify intersections between these groups and those who are politically motivated to launch an attack. This is just one approach to the use of advanced monitoring techniques to identify the precursors to future attacks in a more systematic way than was done in all of the case studies discussed here.

The second stage of the road map, presented in Figure 2, focuses on risk assessment for cyber defence. This reinforces links between the road map and existing security management systems; however, our focus on the use of social media in cyber-attacks distinguishes our approach from this more general work. Risk assessment is important because, at present, it is not performed in an organised or systematic way by many of the potential targets of cyber-attacks across NATO countries. One reason for this is that Public-Private-Partnerships have bneen used by many states as a cost-effective means of encouraging commercial and public bodies to increase their resilience to future attacks. Unfortunately, this approach has not worked very well. Many companies have been enthusiastic participants in the first round of these initiatives. However, many others have not participated at all. Even those that were initially enthusiastic have not sustained their work in the monitoring and audit of the implementation of security policies across all areas of the organisation – evidence for this is provided by the widespread impact of the attacks described in this report. Social networking provides one means of addressing these concerns – for instance, following the development of social media as a means of developing common communities across the safety industries using resources such as EUROCONTROL's Skybrary resource. In particular, it is essential that any future use of social media to increase resilience by extending and revising risk assessments should extend the range of participants within existing PPP security initiatives. For instance, Europe has recently implement Safety of Life services using their EGNOS satellite based augmentation system. For the first time, this enables end users to integrate navigation and timing information derived from the GPS or GLONAS constellations into safety-critical systems including a range of flight-deck applications. At present most of the attention has been on demonstrating that these infrastructures meet safety requirements, however, this is one example of an area where NATO and European member states urgently needs to develop social networks that can consider the security implications of future attacks on this emerging infrastructure (Johnson and Atencia-Yepez, 2011).

The third element of the road map focuses on the use of social media to disseminate early warnings about potential cyber-attacks. The importance of this should not be underestimated, the consequences of failing to communicate with all stake holders was seen in the Lithuanian attacks when the CERT did not extend the warnings it provided to commercial and private organisations. Hence, it is necessary to use drills and exercises to determine the effectiveness of communications mechanisms that might be used during a future attack. Social networks can provide critical support during the planning of these exercises. Too often these drills assume perfect conditions for CERTs and other state agencies to respond to potential attacks. However, the Georgian cyber- attacks during the 'physical' occupation of Ossetia demonstrates the limitations of plans that are based on such 'perfect' assumptions. We have also seen new patterns of attack where different groups engage in mutual hostilities, for example in the conflicts between the Indian and Pakistan Cyber Armies. Such scenarios are seldom considered in existing exercises. By developing wider communities of interest in cyber-security, by encouraging participation in drills and exercises, we can also obtain evidence of vulnerabilities and the consequences of future attacks that provide persuasive data to encourage wider participation in these initiatives.

The penultimate stage in the roadmap focuses on the use of social networks in forensic analysis after an attack has occurred. At present, it is extremely difficult to detect the early stages of an attack. A range of network monitoring companies, mostly located in the United States, provide detailed information on patterns of connectivity failures that can provide real-time evidence of the impact of a DoS attack. However, the information that they provide is not coordinated in any clear way. In consequence, CERTs trigger their response after an attack is well underway and significant damage has already been sustained. Greater benefits could be derived from research into the particular patterns that have characterised the initial stages of previous attacks, for example using abnormal Border Gateway Protocol events. The real-time identification of an attack is not simply a technical exercise but typically requires input from a range of different experts using social and professional networks that should be developed in a more coordinated manner than is done at present. Similar comments can be made about the distribution of patches and recommended interventions in the aftermath of an attack. In each of the case studies presented in this paper, many critical infrastructures remained vulnerable long after forensic reports had been published into the causes and mitigations for previous attacks.

The final stage of the roadmap for using social media to support cyber defence focuses on the identification and dissemination of lessons learned. At present this remains an ad hoc process. Many stakeholders remain ignorant of the extent of previous attacks, especially outside their own national borders. Further problems arise when vested interests exaggerate the impact of cyber-attacks; this alienates engineering staff who play a vital role in implementing potential mitigations. As we have seen, reports that cyber-attacks will 'bring a nation to its knees' cannot be justified on the evidence of previous incidents although this should not be used to justify future complacency. One aspect of dissemination is the difficulty that managers and operators have in obtaining clear information about previous attacks. This has been a motivation in the compilation of this report. Strong contrasts can be drawn between the lack of information in this area that might be used to motivate increased vigilance and the plethora of accident and incident reports published across safety-related industries. Of course, it can be argued that the wider dissemination of this information might only encourage future attacks. However, this argument ignores the anti-social networks that have already sprung up to encourage and sustain previous attacks.

10. Conclusions

The last decade has seen a growing number of cyber-attacks, for instance on Estonia, Belarus, Lithuania, Georgia, Pakistan and India. It has been difficult to determine whether or not these incidents were state-sponsored. This paper has identified three different roles that social networking and social media have played in this 'attribution problem'. Firstly, social networks have motivated individuals to participate in mass Denial of Service (DoS) attacks. They have disseminated information and provided access to resources, including botnets that were originally developed by cyber-criminal groups. Secondly, we show how information about an individual's social networks has supported targeted attacks, such as spear phishing, on opposition groups. Malware is, typically, disguised in a document that was intercepted from a colleague or friend. The recipient is more likely to open an attachment or link if it has been sent from a trusted source. Thirdly, we show how the development of Cloud infrastructures to support social networking applications has created disposable architectures for the Command and Control servers that coordinate malware attacks. The ubiquitous and distributed nature of these architectures makes it increasingly difficult to determine who owns and operates these systems. The closing sections of the paper have identified a roadmap for the defensive measures that might be used to minimise the future threats from the 'dark side' of social networking.

11. References

M. Amarelo , New Federation of American Scientists Report Demonstrates Iran Improved Enrichment in 2010. Federation of American Scientists, Washington DC, USA, 21 January 2011. Available at: http://www.fas.org/press/news/2011/issuebrief_iran.html, last accessed June 2011.

S. Gaycken, Stuxnet: Wer War's? Und Wo Zu? Die ZEIT. 26 November 2010. Available at: http://www.zeit.de/2010/48/Computerwurm-Stuxnet/seite-1, last accessed June 2011.

The Economist, A Cyber-Riot: Estonia Has Faced Down Russian Rioters. But Its Websites Are Still Under Attack, May 10th 2007, Available at: http://www.economist.com/node/9163598, last accessed June 2011.

European Commission, Independent International Fact-Finding Mission on the Conflict in Georgia, Council of the European Union, Brussels, Belgium, September 2009. Available at: http://www.consilium.europa.eu/uedocs/cms_data/docs/pressdata/en/er/110370.pdf, last accessed June 2011.

N. Falliere, L. O. Murchu and E. Chien, W32.Stuxnet Dossier, Symantec, Cupertino, California, USA, February 2011. Available at: http://www.symantec.com/connect/blogs/w32stuxnet-dossier, last accessed June 2011.

Grey Goose Project, Russia/Georgia Cyber-War: Findings and Analysis, Phase I Report, October 2008. Available on: http://www.scribd.com/doc/6967393/Project-Grey-Goose-Phase-I-Report, last accessed June 2011.

S. Henderson, Hunting the GhostNet Hacker. The Dark Visitor. 2nd April 2009. Available on: http://www.thedarkvisitor.com/2009/04/hunting-the-ghostnet-hacker/, last accessed June 2011.

R.B. Hughes, NATO and Global Cyber Defence, Ed. R. Shepherd, The Bucharest Conference Papers, Royal Institute of International Affairs, Chatham House, London, 41-55, 2008.

Information Warfare Monitor, Tracking GhostNet: Investigating a Cyber Espionage Network, Munk Center for International Studies, University of Toronto, The SecDev Group, Ottawa, Joint Technical Report JR02-2009, March 2009, Available on: http://www.scribd.com/doc/13731776/Tracking-GhostNet-Investigating-a-Cyber-Espionage-Network, last accessed June 2011.

Information Warfare Monitor, SHADOWS IN THE CLOUD: Investigating Cyber Espionage 2.0, Munk Center for International Studies, University of Toronto, The SecDev Group, Shadowserver Foundation, Ottawa, Joint Technical Report JR03-2010, April 2010. Available on: http://www.scribd.com/doc/29435784/SHADOWS-IN-THE-CLOUD-Investigating-Cyber-Espionage-2-0, last accessed June 2011.

C.W. Johnson, Case Studies in Major Failures of Telecommunications Infrastructures. In P. Theron and S. Bologna (eds), Improving the resilience of European telecommunications (in press), 2012.

C.W. Johnson and A. Atencia Yepez, Mapping the Impact of Security Threats on Safety-Critical Global Navigation Satellite Systems. In Proceedings of the 29th International Systems Safety Society, Las Vegas, USA 2011, International Systems Safety Society, Unionville, VA, USA, 2011.

A. Matrosov, E. Rodionov, D. Harley, and J. Malcho. Stuxnet Under the Microscope, ESET, Bratislava, Slovakia, 2010.
Available at: http://www.eset.com/us/resources/white-papers/Stuxnet_Under_the_Microscope.pdf , last accessed June 2011.

S. Nagaraja and R. Anderson, The Snooping Dragon: Social-Malware Surveillance of the Tibetan Movement, Computer Laboratory, University of Cambridge, UK, Technical Report Number 746, March 2009.
Available on http://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-746.pdf, last accessed June 2011.

J. Nazario, Estonian DDoS Attacks – A summary to date, Arbor Networks, Chelmsford, MA, USA, 17th May 2007, Available on: http://asert.arbornetworks.com/2007/05/estonian-ddos-attacks-a-summary-to-date/, last accessed June 2011.

J. Nazario, Radio Free Europe DDoS. Arbor Networks, Chelmsford, MA, USA, 29 April 2008, Available at: http://asert.arbornetworks.com/2008/04/radio-free-europe-ddos/, last accessed June 2011.

J. Nazario, Georgia On My Mind – Political DDoS, Arbor Networks, Chelmsford, MA, USA, 20th July 2008a. Available on: http://asert.arbornetworks.com/2008/07/georgia-on-my-mind-political-ddos/, last accessed June 2011.

J. Nazario, Georgia DDoS Attacks – A Quick Summary of Observations, Arbor Networks, Chelmsford, MA, USA, 12th August 2008b.
Available at: http://asert.arbornetworks.com/2008/08/georgia-ddos-attacks-a-quick-summary-of-observations/, last accessed June 2011.

J. Nazario, Twitter Based Botnet Command Channel, Arbor Networks, Chelmsford, MA, USA, 2009. Available on: http://asert.arbornetworks.com/2009/08/twitter-based-botnet-command-channel/ , last accessed June 2011.

J. Nazario, Malicious Google AppEngine Used as a CnC, Arbor Networks, Chelmsford, MA, USA, 2009a.
Available at: http://asert.arbornetworks.com/2009/11/malicious-google-appengine-used-as-a-cnc/ , last accessed June 2011.

H. Pääbo, War of Memories: Explaining 'Memorials War' in Estonia, Baltic Security & Defence Review, (10)5-28, 2008.

E. Tikk, Frameworks for International Cyber Security, NATO Cooperative Cyber Defence, Centre of Excellence, Tallinn, Estonia, 2009.
Available on:
https://www.nsm.stat.no/upload/Konferanser%2009/05_Framework%20of%20cyber%20incident_Tikk.pdf, last accessed June 2011.

E. Tikk and K. Kaska, Legal Cooperation to Investigate Cyber Incidents: Estonian Case Study and Lessons. In the Proceedings of the 9th European Conference on Information Warfare and Security, Thessaloniki, Greece. Academic Press, Reading, MA, USA, 288-294, 2010.

E. Tikk, K. Kaska and L. Vihul, International Cyber Incidents: Legal Considerations, NATO Cooperative Cyber Defence, Centre of Excellence, Tallinn, Estonia, 2010.
Available at: http://www.ccdcoe.org/publications/books/legalconsiderations.pdf, last accessed June 2011.

E. Tikk, K. Kaska, K. Rünnimeri, M. Kert, AM. Talihärm, L. Vihul, Cyber Attacks Against Georgia: Legal Lessons Identified, NATO Cooperative Cyber Defence, Centre of Excellence, Tallinn, Estonia, November 2008.
Available on: http://www.carlisle.army.mil/DIME/documents/Georgia%201%200.pdf, last accessed June 2011.

N. Villeneuve, Shadows in the Cloud, Personal Blog, April 2010.
Available on: http://www.nartv.org/2010/04/05/shadows-in-the-cloud/, 2010, last accessed June 2011.

S. Waterman, Security Industry Analysis: Who Cyber Smacked Estonia? United Press International. 11th June 2007.
Available at: http://www.upi.com/Business_News/Security-Industry/2007/06/11/Analysis-Who-cyber-smacked-Estonia/UPI-26831181580439/#ixzz1OhCASAE9, last accessed June 2011.

E. Zmijewski, Georgia Clings to the 'Net, Renesys, Manchester, New Hampshire, USA, 10th August 2008.
Available on: http://www.renesys.com/blog/2008/08/georgia_clings_to_the_net.shtml, last accessed June 2011.