Applying Robert A. Pape's Denial Strategy to Computer Warfare

Trevor Sutherland

The aim of this study is to use Robert Pape's Bombing to Win as a stepping off point. This article analyzes the utility of computer warfare as a method to coerce states. It then continues on to consider other views on cyber coercion and how actors and targets can be classified. Actors are mapped to one of four quadrants while six separate types of targets are created.

Keywords: *Cyber war, computer network attack, strategic bombing, malware, computers*

In the aftermath of World War I, there was much debate over the future role of the airplane in the military conflict. During the war, it had shown itself to be a platform with much potential, filling a number of roles and assisting the traditional land and sea forces with their missions. At the same time, strategists realized that its incorporation into warfare also enabled forces to undertake missions unlike those that were previously imagined. Rather than simply reconnoitering and spotting for artillery, theorists saw that the airplane could be turned into a flying piece of artillery, enabling the birth of what we now think of as strategic bombing.

After World War I, several thinkers became active in the debate over the use of Airpower. Of the major thinkers in the early interwar period, Giulio Douhet remained today the most widely referenced. Douhet focused mainly on strategies that involved attacking non-military targets such as industry, transportation, and government centers (Ferrari 1942, 179). By attacking targets that would conceivably disadvantage the civilian populations of home countries, Douhet felt that the civilian populations would pressure their governments into submitting to the will of those doing the bombing (Pape 1996, 60). This view was further popularized by the first commander of the British Royal Air Force (RAF), Hugh Trenchard, who began a policy of directly bombing civilian populations based on his experiences policing the British colonies. Both Douhet and Trenchard exemplify the feeling of many early airpower theorists that the airplane is an inherently offensive weapon. They felt that an effective defense would at best be attritional to the attackers and that some forces would inevitably survive to attack their target (Bradbeer 2004, 125).

A similar line of thought dominates the discussions of cyber war that are occurring today. As our daily lives become more interconnected and networked there are simply too many potential vectors for attack. Moreover, as this interconnectedness fosters increased efficiency it stands to reason that interdependency will only increase over time, ultimately putting more systems at greater risk. Most books devoted to cyber security focus on this problem in the civilian sense, either speculating on the

doi: 10.18278/gsis.1.2.5

effects of an attack made against element of infrastructure, like electricity distribution or economic institutions. While there is some evidence that such tactics might have been used to coerce individuals for monetary reasons (Brenner 2011, 557), there is no basis for judging the costs and effects of an attack made on a population as a whole. The general assumption put forward in most books is that the victimized population will cease to be able to effectively function and pressure their government to capitulate to the attacking force. This is very similar to the Douhet model of strategic bombing, in which a civilian population's morale is broken and rendered incapable of effective resistance.

If populations exposed to direct bombing in the past are used as an example, it stands to reason that the expected outcome of civilian capitulation would likely not occur. This strategy, employed extensively during World War II against population in Britain, Germany, and Japan, has since been found to be only moderately useful. While population-targeted coercion can be effective in situations when nuclear weapons are expected to be used, attacking civilian populations more often has an opposite effect. By bringing civilians into the war fighting process, a rally-around-the-flag effect is usually seen increasing resistance rather than undermining it, as was seen in both England during World War II and North Vietnam during the Vietnam War.

Instead, most effective strategies for coercive bombing focus on limiting the military effectiveness of an opponent. This is accomplished through the use of airpower to complicate the manufacture of arms, interdict their transportation to the battlefield, and disrupt communications on the battlefield and within the theater (Pape 1996, 69). To better understand what a coercive action involving the cyber sphere would look like, cyber capabilities should be analyzed in terms of how they can fulfill these goals without effecting the population as a whole.

Is Cyber Coercion Viable?

In *Bombing to Win*, Robert A. Pape examines the use of strategic bombing for coercion. He compares its usefulness for this task against land and sea-based measures and finds that airpower is ideal for coercion for a number of reasons. First, it is flexible and precise, allowing those that use it to better separate actions taken against the military from those taken against the population. Second, it allows greater amounts of ordnance to be put on target with more precision and over a greater area than either land- or sea-based measures. Lastly, unlike land-based coercion, strategic bombing does not require a decisive ground victory to be successful.

The advantages of cyber coercion are similar to those of airpower. Flexibility and precision can be achieved through defining the attributes that are present in a given environment before an attack can begin, as was evidenced in the outbreaks of both the Conficker and Stuxnet worms (Bowden 2011, 56; Falliere, Murchu, and Chien 2011, 7). This would further reduce the risk of collateral damage through misidentification of buildings or the location of any non-military buildings nearby. Payload delivery would also be more efficient in a cyber-coercive campaign, as physical distance and

munitions weight are not factors. Lastly, just as aerial coercion does not require ground superiority but only a measure of air superiority, cyber coercion needs neither. What is needed in some situations is "network superiority," the ability to function in an opponent's networks with complete freedom.

The Limitations of Denial

Pape also outlines the limitations of denial, his term for coercion carried out against military targets, in *Bombing to Win*. These are (a) effective denial within the area over which control is sought, (b) constant maintenance of pressure until concession is given, and (c) the ability to control the territory by force (Pape 1996, 32). These limitations present five major problems for cyber operations as they are typically understood today.

First, they require significant tangible effects. While there is proof that the manipulation of industrial control systems and Supervisory Control and Data Acquisition (SCADA) systems can yield spectacular results (Brenner 2011, 1497), these effects are often achieved through negligence on the part of the system operator and can be classified as targets of opportunity. This lack of diligence is seen in the utility of websites such as Shodan.com that act as a "google for SCADA." It is reasonable to assume that a government, especially one that is in a state of heightened conflict, would secure phone, power, and other essential systems quickly. Also, most SCADA systems that are accessible through the Internet are set up for the convenience of those maintaining the system (Bentek Systems 2012). If the networked controllers ever became a large-scale problem, most industrial and manufacturing systems could simply be "unplugged," from the Internet with few consequences.

Other systems are networked through their very nature and these will likely be the most vulnerable over the course of a potential coercion campaign. Such systems include the Internet, air traffic control systems, and the systems that control road and rail travel. For most of these systems there are non-networked alternatives, though these alternatives are less efficient, even after implementation. This loss of efficiency can be seen as a limited form of interdiction, though one not likely to be successful without other factors utilized as well.

The second problem presented by Pape's limitations is the need for persistent pressure. While assets and systems can be rendered inoperable through a cyber attack, once an alternative method of providing the same service has been established the coercive attempt has effectively failed. For a successful outcome, pressure must be maintained over a potentially long period of time. This can be achieved by consolidating control over the target computer networks, allowing for the pacing of operations to slowly degrade systems, which is similar to the actions Stuxnet took against Iranian nuclear processing centrifuges.

Once network superiority is achieved, another option for persistent capability is falsifying and modifying data rather than destroying it. By not deleting the data, the hope is that the target will not notice the extent of the infection, thus enabling

the coercer to influence the target over a longer period of time. Of course, it should be noticed that this strategy can cause significant attribution problems if not enacted properly. If actions are unattributable to the coercer through overuse of this subterfuge, then they constitute only wasted effort on the part of the coercer.

Third, the attempt at cyber coercion must be able to resist the target's attempts to undo the long term efforts of the coercer. The basis of cyber security is the realization that every program potential for a critical flaw that can be exploited for disastrous results. This theory applies equally in this case to the coercer and the target. By consolidating control over the target networks, the coercer is opening himself to the potential that his efforts will be flawed in a way that enables the target to take control of his network. While there are ways to minimize this risk, the potential consequences of such an action are disastrous.

There is also the risk that the target will notice the vulnerability that is being exploited and simply fix it. Many modern malwares, for example, act toward isolating a system from updates, as was seen with the Confikr worm (Bowden 2011, 54; Dhanjani, Rios, and Hardin 2009, 3189) before carrying out their ultimate end. By doing this, they seek to prohibit any actions that might either detect their presence or correct a vulnerability that they might be dependent on. If an attempt at coercion can be turned aside by simply updating a system to the most recent version, then it will most likely be found to be ineffective.

Fourth, effective coercion ultimately relies on the coercer possessing significantly greater military power than the target to have a chance at success (Pape 1996, 45). In land-, sea-, and air-based operations this superiority is needed to create the relative freedom of action that enables effective denial strategies. While there may be some disagreements over the finer points of relative power, a general consensus does exist over which countries are militarily stronger than others in conventional terms. However, as no true cyber competition has occurred with two attributable actors, no equivalent scale exists for computer-based capabilities. Moreover, a coercive strategy through networked infrastructure will only be effective against very highly developed countries. Given these facts, very few countries are susceptible to cyber coercion.

The last problem with cyber coercion qua Pape's limitations of denial is one of attribution. Even if an effective strategy is implemented and carried out to its fullest extent, it is wasted effort unless the target knows with certainty who is carrying out the actions and what demands they are making. When combined with the points listed earlier, a delicate situation comes into being: if a coercer acts too strongly, he risks his ability to maintain coercive pressure. However, if he acts too stealthily, he risks non-attribution and failure.

Analysis of Problems Posed by Limitations

While there does exist the possibility for a successful network-based coercion campaign, it is remote. Most literature focuses on coercive attempts against civilian populations for the good reason that that segment of the population

is far more vulnerable to attacks of this type. While some targeted attacks could be used to contribute to a larger coercive campaign, most possible actions fall short in one of two areas.

First, actions effecting general infrastructure such as traffic control, oil and electricity distribution, or economic structures harm the civilian population disproportionately over the military. As Pape points out, military forces tend to have auxiliary capabilities to provide for most of their needs (Pape 1996, 75). This reduces the logistical tail that is vulnerable to attack. Furthermore, in times of scarce resources, militaries usually have priority access to what resources are available. This results in most actions victimizing the civilian population while having negligible effects on the military. As explained earlier, population-centered coercion is rarely effective.

Second, the problem of creating a reliable, persistent, and effective framework through which to continue coercion poses a significant challenge. Aside from the attribution challenge explained earlier, it is hard to think of a way that a computer problem could elicit devastation similar to a bomb without rendering the system that it located on inoperative or exposing itself in such a way that it is allowed to remain a threat. Though some examples exist of malware that is capable of reaching non-networked systems, these remain costly and time consuming to create and ultimately rely on bad implementation of security best practices (Fallieret al. 2011, 3).

This is not to say that computer operations could not be used to augment the efforts of a larger coercive effort. When used in this way, actions could be taken against targets of opportunity while a greater effort could be put into sabotaging critical components (Pape 1996, 71) and decreasing the efficiency of the overall manufacturing system. Though this is a costly and time-consuming procedure, limiting the target list to several facilities could allow teams enough time to conceivably hinder the production of needed military goods.

Further Limitations

While the lens of Pape's thinking can serve as a tool to evaluate network attack, other authors have tackled this topic in a more head-on fashion. Most significantly, Thomas Rid and Brandon Valeriano and Ryan Maness seat computer operations firmly within greater frameworks of international relations and war. Jason Healey also offers useful advice, drawing from a well of knowledge gained from studying incidents of cyber conflict dating back to the 1980s. Lastly, there is a huge body of technical knowledge that seems largely absent from academic policy writings, yet can add significant depth and texture to any discussion.

In *Cyber War Will Not Take Place*, Thomas Rid makes a compelling case that war as defined by Carl von Clausewitz is unlikely to be waged solely with computers. He also enumerates, based on his analysis, the three avenues that computer operations could potentially be useful. Finally, he also creates a continuum for classifying what he calls "cyber weapons."

Referencing Clausewitz, Rid focuses on the three qualities that separate war from simple violence or contention: war is violent, war is instrumental, and war is political (Rid 2013, 2). While cyber operations can easily fit the second and third criteria, true violence is difficult to create reliably. Though SCADA vulnerabilities can lead to destructive malfunctions and even possibly explosions, it is hard to consistently construct threats that have a serious chance of causing death or injury to civilian or military personnel (Rid 2013, 66).

Continuing to reference Clausewitz, Rid moves on to the structural aims of war, namely to disrupt trust between a population and its government and military (Rid 2013, 22). In normal circumstances, this trust is attacked through violence; as the population loses faith in its institutions they begin to lose trust in each other until either the state capitulates or order breaks down completely. Cyber attacks can—in theory—facilitate this lack of trust through non-violent means.

Rid gives three ways that this can happen: through espionage, through sabotage, and through subversion (Rid 2013, 10). Espionage erodes trust by showing that a government is incapable of protecting its citizens' digital assets. This is, as Rid points out, not only non-violent but also of questionable instrumentality as theft is a clandestine activity and likely would not be publicized (Rid 2013, 81). Sabotage is similarly hamstrung as a vector of war as it aims to disrupt the trust of groups in their equipment, as was seen in the Stuxnet attacks. Once the source of the sabotage is presented, or even the existence of sabotage, this trust is restored (Rid 2013, 32). This leaves subversion as the best avenue for a true "cyber war," though the least likely to be classified as such because any resulting contention or violence is likely to be seen as internal struggle instead of the result of malicious code (Rid 2013, 114).

Finally, Rid provides a useful continuum for classifying the tools of cyber conflict. On one side are weapons that are broadly effective yet produce minimal results. Denial of service attacks (DoS) and website vandalism fall in this category as they are effective to some degree on all Internet-connected devices yet cause little lasting damage. The continuum's other side are tools like Stuxnet, which are deeply impactful yet highly specific. These are the equivalent of a sniper's bullet that can cause massive amounts of damage yet must be tailored very specifically to a given target (Rid 2013, 35).

In contrast to Rid's forecasting, Valeriano and Maness attempt to empirically quantify what we already know about cyber conflict. Significantly, they approach the issue from the assumption that cyber conflict is a sphere of diplomacy rather than warfare, neatly sidestepping Rid's questions of the need for violence. From this vantage point, they see several interesting trends: first, that almost all cyber attacks are rooted in a pre-existing rivalry. Second, those attempting to use cyber weapons are very likely to be restrained in their use. Lastly, states will sometimes actively support cyber terrorism, though only in very specific situations.

The bedrock observation of Maness and Valeriano is that cyber contention stems from traditional interstate rivalries, similar to economic and military contention (Valeriano and Maness 2015, 8). This means that the process of attribution, usually

held up as a key problem in cyber warfare, is simplified as the list of potential suspects is greatly reduced. Also, it signifies that the vast majority of conflicts will occur between neighbors, as they are more likely to be in contention with one another than countries with no shared borders. Moreover, the exceptions to this rule will be more significant and more constrained than non-neighbors.

Next, the nature of cyber conflict incentivizes restraint in its use. Because cyber tools are less predictable than conventional munitions, they are more likely to go awry in several ways: most significantly, they can be difficult to control, meaning that unintentional overreach is a possibility (Valeriano and Maness 2015, 4). Second, as cyber munitions are not expended when used, there is the possibility that victims or even third parties will reuse the tools for their own ends. This restraint will likely manifest itself in the use of cyber tools for primarily low-level actions, such as espionage, or to exploit obvious weaknesses (Valeriano and Maness 2015, 72).

Lastly, Valeriano and Maness claim that states have and will resort to cyber terrorism in certain situations. Primarily, state-sponsored cyber terrorism will allow less powerful nations to act with greater effect against more powerful foes. Second, states will resort to terrorism when they wish to distance themselves from their actions. Finally states will resort to terrorism when they want to quickly amplify their power for very simple purposes, as was seen in the Russian–Estonian War of 2007 (Valeriano and Maness 2015, 70). Interestingly, Valeriano and Maness do not seem confident that cyber terrorism can effect change.

Jason Healey's sweeping recap of conflict in the cyber arena, *A Fierce Domain*, covers a whole range of what could be termed "cyber attacks," between 1980 and 2012. While many of the attacks that he discusses are purely civilian in nature, the lessons learned from aggregation show that relatively little has changed in the past 35 years. Key among these lessons is the importance of public–private cooperation in computing crises, the nebulous nature of US cyber command and control, and that, in general, the more significant a cyber conflict is, the more similar it is to other conflicts.

From the beginning of cyber conflict, the importance of sharing of information between the public and private sectors has been crucial to both defense and recovery from attacks. This is seen as early as 1986's Lawrence Berkley Labs intrusions (Healey 2013, 2117) to as recently as the recovery from 2007's Estonian "cyber war" (Healey 2013, 1691). Similarly, in offense some states often employ or allow non-state actors to contribute to state-led efforts, as is seen in Healey's chart, Spectrum of State Responsibility (Healey 2013, 1218), and was also demonstrated in the 2007 Estonian event.

Contrasting this need for cooperation is the way in which the United States has handled the increasing militarization of cyber security. Healey quotes General Dusty Rhodes, former head of the 609th Information Warfare Squadron as saying that it was a great detriment to the cause of information security that all of the 609th's offensive operations remained classified, as well as many of their defensive actions (Healey 2013, 807). This is expounded upon by other statements made by other military commanders and policymakers (Healey 2013, 1126, 1327) indicating that many neglected cyber capabilities are due to their poor integration within the military structure. Put together,

these sentiments point to a system that was and largely still is isolated from much of the day-to-day operations of the groups that it claims to protect.

Lastly, similar to Valeriano and Maness, Healey makes the argument that cyber conflicts behave more similarly to non-cyber crises as they grow in significance, with one exception. That exception, the increased presence and ability of non-state actors, is squarely at odds with arguments made by Valeriano and Maness (Healey 2013, 494). This differentiation is important, as it substantially broadens the field of potential attackers, introduces more variance in how they act, and complicates attribution.

Technical Writings

While many policymakers and thinkers are familiar with the sources cited above, there is another category of sources that are rarely referenced. There is a large and ever-growing trove of books, blogs, and media published both formally and informally that shows the conflict over security from the tactical side rather than the political. While much of this information is of little value to decision makers, there are certain fundamentals that can give those at the strategic and political levels of decision-making valuable insights.

First, there are the core goals of security. According to Harris, an expert on information security, the goal of a computer security is to maintain data in a way that it is always available, accurate, and confidential (Harris 2012, 1212). At first glance, this maps very closely to Rid's sabotage, subversion, and espionage, as each attack targets its respective value (sabotage attacks availability, subversion targets accuracy, and espionage targets confidentiality). Still, as Harris goes on to point out most businesses (and governments) do not exist to be secure, that is security is a secondary goal that has no value if the company is not successful in its primary endeavor. Contrary to kinetic warfare where a state by definition must have a monopoly over force within its borders, computer security will always be less important than physical security and day-to-day government operations.

Closely tied to these goals are the concepts of risk, threat, and vulnerability (Figure 1). While these terms are used in both the kinetic and computer arenas, in computer security, these words have very specific meanings: a threat is a possible danger. A threat agent is something that actually uses a threat to cause damage, assuming that vulnerability can be found. Risk is the probability that one's assets take damage from a threat agent given once exposure and safeguards (Harris 2012, 1312)

This model warrants a moment's consideration, as most of the hyperbole that Rid, Valeriano and Maness, and Healey frequently reference can be seen in this model. In the arguments of thinkers like Richard Clarke and Winn Schwartau, both noted for their contributions to information security policy, there are effectively an infinite number of threats and vulnerabilities meaning that there is a near infinite risk. This has clearly not been the case as observed by the authors profiled in the previous section and is due largely to difficultly of constructing treat agents that can act in concert with a states' wishes while limiting the probability of blowback.

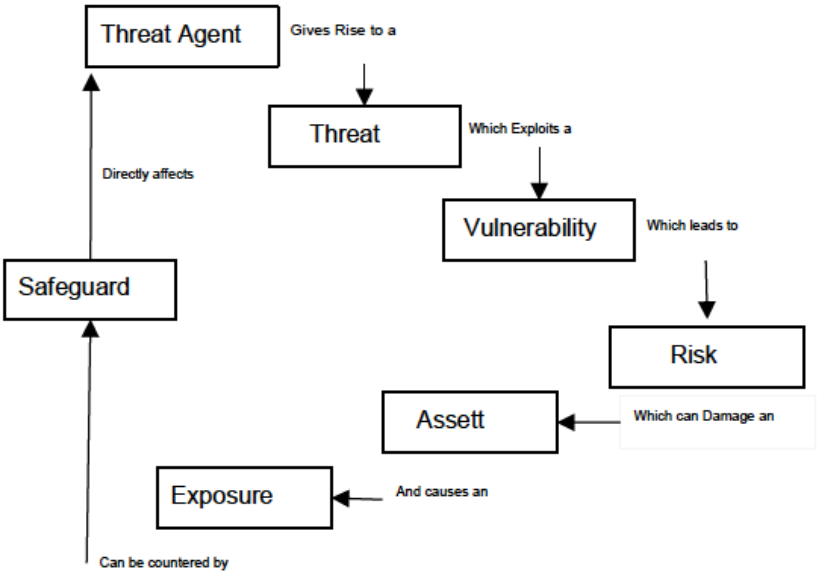

Figure 1: Relationships of Different Security Concepts (Shon 2012, 1312)

Next is the concept of defense-in-depth, wherein layers of protective measures are used to provide overlapping means of security. According to Harris, these layers can each have one of six functions (Harris 2012, 1343) as follows.

Deterrent	Discourages a potential attacker
Preventive	Intended to avoid an accident
Corrective	Fixes a component/system after an incident
Recovery	Intended to bring environment back to normal
Detective	Watches a system's activities for signs of an incident
Compensating	Provides an alternate measure of control

While not all of these functions can realistically scale up to a strategic or political level, it is very clear that some of these have been favored above others in the discussions on security. For example, Valeriano and Maness extensively discuss the lack of utility of a deterrent strategy and the NSA's mandate to secure the American military space is very akin to a strategic-level detective function. Still, rather than focusing on preventative measures as the only solution, other options such as compensation and correction might well have a place in a strategic computer security solution.

Further Analysis

The various insights of the aforementioned authors paint a much more complete picture of cyber conflict as it exists today. Rid's integration of cyber war into the Clausewitzian understanding of war is significant, though by dropping the traditional requirement of violence, more understanding can be gained. Similarly,

Valeriano and Maness' framing of cyber conflict within the greater international system yields great possibility for moving thought forward on this topic.

Rid's assertion that trust is key in cyber attacks is crucial to understanding the landscape; however, there are more functions that attack trust than just espionage, sabotage, and subversion. An additional three "hybrid attacks" are conceivable, each a combination of two basic types (Figure 2). These are insurgency, transparency, and APT (advanced persistent threat).

APT is a term that used to be reserved for high-level state threats, but is now used to denote any threat that has the long-term capability to enter a computer system at will and take information. APT is usually achieved by sabotaging the computer's security system crippling its abilities to detect the intrusion and subsequently exfiltrate information. This has been seen in a number of incidents—usually attributed to the Chinese state—including the Moonlight Maze and Night Dragon attacks (Healey 2013, 1138, 1611).

Insurgency, much like its kinetic counterpart, is a combination of subversion and sabotage. This type of attack is what Valeriano and Maness refer to as state-sponsored cyber-terrorism. In these incidents, non-state forces are coerced into acting on the state's behalf to publicly sabotage a target of the state's choosing, undermining trust in the state. Examples of cyber insurgency are the 2007 Russian–Estonian conflict and the computer component of 2008's Russian–Georgian War.

Transparency is arguably the most powerful attack methodology, a combination of espionage and subversion. Though the results of this vector are rarely considered a cyber incident, well-timed revelations of stolen information like the Edward Snowden leaks or WikiLeaks can have a huge subversive effect on a population.

All three of these hybrid attack methodologies are unique to their fundamental counterparts in scope. APT is greater than simple espionage in that it creates a paranoia that undercuts trust more thoroughly once it is detected and exfiltrates more information while it is being used. Similarly, transparency is more effective than simple espionage because it is noticed more and can undercut not only institutional trust but also faith in institutional motives. Lastly, infiltration can create an ad hoc strategic weapon capable of surpassing the tactical capabilities of most cyber tools.

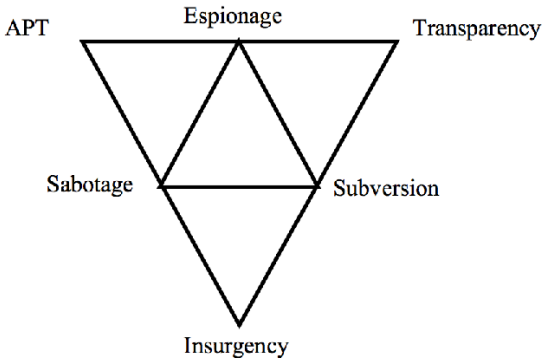

Figure 2: A Riddian Triad Modified with Hybrid Attack Methods

Worth noting in the examples above is the presence of non-state actors. Healey writes that as cyber crises become more significant, they become more like traditional crises save the greater number of non-state actors. This is in contention with Valeriano and Maness (2015, 165), who claim that non-state actors have little effect on states in a cyber conflict other than that they may be co-opted into a state's official plans. This is an unrealistic expectation for several reasons.

First, if the true aim of an attack is to damage a population's trust in its institutions, it stands to reason that portions of a population will become involved in the conflict if the barriers to entry are low enough. This is related to Pape's "rally around the flag," and is similar to resistance movements around the world. The Internet's low entry criteria and huge potential for individual anonymity create the perfect situation for individual involvement.

Second, as the Estonia–Russia conflict of 2007 shows, in times of extreme crisis, it is civilians, not state forces that are available to assist (Healey 2013, 1693). The reasons for this are twofold: primarily, it is because the Internet is run by nongovernmental groups and these are the groups with the expertise to aid in large-scale restructuring of network infrastructure. Additionally, because state assets are many times bound in nebulous command structures that prioritize secrecy, they are unable to help directly. This non-state, ad-hoc cooperation was also seen in the public response to the Windows Confikr worm (Bowden 2011), as well in 1986's Lawrence Berkley Labs (Cookoo's Egg) intrusions (Healey 2013, 2125).

Lastly, as the Internet and its users mature, non-state actors' abilities are constantly growing. Far from the website vandalism and DDoS attacks of the early 2000s, individuals, such as the social activist "The Jester", are able to cause significant real damage by themselves. When individuals with these skills lend or sell them to groups with political agendas, the result is semisophisticated incidents like 2012's Shamoon virus which targeted Saudi Aramco computers (Rid 2013, 56). Also, as the 2015 breakup of Italy's Hacking Team shows, non-state groups are actively developing a sophisticated arsenal of digital tools (Security Week 2015) that are being used by states.

While the presence of these non-state entities does complicate Valeriano and Maness' ease of attribution, it does not create undue confusion. Most of these non-state groups are interested primarily in crime and will only become involved in political events when they are tapped by their respective states (Carr 2009, 28). Other groups, like those that administer and run the Internet, have little interest in causing unrest and will only become involved in an incident to mitigate and diagnose problems.

These groups can, however, be classified in a way that can give clue to their intentions (Figure 3). When classified as either legal or illegal and further subdivided by level of organization a pattern emerges

	Organized	Illegal
Legal	• Infrastructure groups (IETF, ICANN, W3C) • Antivirus companies • University researchers • Professional groups (SANS)	• Working groups (NSC-SEC, Confikr Working Group) • Individual researchers
Individual / Ad Hoc	• Individual hackers • "Hacktivists" • Disgruntled insiders	• Individual hackers • "Hacktivists" • Disgruntled insiders

Figure 3: Types of Non-State Actors in Cyber Events and Examples

In general, each quadrant has common goals. Entities in the top left quadrant (Legal, Organized), have a vested interest in maintaining order online and understanding and mitigating systemic threats. While they are very unlikely to instigate problems, they are usually in the forefront of solving most major problems. The top right square (Legal, Individual/Ad Hoc) contains many of those that actually solve systemic crises. A common model, as seen in the reaction to the Estonian crisis and the Confikr virus, is for individuals associated with the groups in the first quadrant to come together to form and execute a solution.

The bottom row contains the groups that are likely to have a part in instigating crises, though only in well-defined contexts. The bottom left quadrant (Illegal, Organized) is traditionally the domain of organized crime groups, such as the infamous Russian Business Group. Though usually interested in traditional criminal enterprises, there is evidence that these groups have engaged in political activities with state sponsorship. Similarly, online wings of terrorist groups are an emerging phenomenon, but, thus far, organized activities have been limited to propaganda and local activities.

These groups' capabilities can be highly sophisticated and targeted, as seen in the 2013 intrusion on the US shopping chain Target and The Cutting Sword of Justice's re-weaponization of the Wiper malware (CNET 2012). Still, most of these groups' activities are less spectacular and usually comprise identity theft and website defacement.

Lastly, the bottom-right corner (Illegal, Individual/Ad Hoc) is the realm of so-called "black hats," (malicious individuals) and hacktivists. These are groups and individuals usually focused on short-term goals and causes. While they can be technologically proficient, more often than not, they use tools that fall on the lower end of Rid's continuum (broadly targeted and lightly damaging). There is no evidence of these groups creating strategic-level incidents without state support.

The last entry in this quadrant, the disgruntled insider has proved to be the most significant cyber foe of states (Andress and Winterfeld 2011, 1028). This is the group that individuals like Edward Snowden and Chelsea Manning belong to. Their methods are rarely sophisticated, though their access means that they have little need of sophisticated methods.

To return to Rid's Clausewitian argument, what usually separates malicious state actors from non-state ones is not sophistication but scale. While many of these groups are quite technologically able, they have not yet moved beyond what would be termed an operational level in traditional military terms. All cyber weapons are tactical, as is pointed out by Rid (2013, 35) and they have shown themselves capable of creating focused campaigns to attack a given target from multiple angles. What is yet to be seen is these groups mounting a true strategic campaign, such as an attack across a whole industry or geographic area.

This is likely for several reasons: first, being interested in profit alone, they have little incentive to invest substantial resources in difficult targets so long as easier ones exist. Second, the range of vulnerabilities needed to threaten multiple systems and network architectures is substantially greater than those needed to threaten one.

Businesses, as Harris explains, exist to make money, not to be secure. This leads to a situation wherein there will always be "easy targets" due to competing priorities within a company and industry as a whole. This same logic, incidentally, also applies to state groups acting under restraint as described by Valeriano and Maness: so long as targets exist that are easily attacked with little chance of bleeding into other sectors, there is little reason to devote the substantial resources needed to develop a Stuxnet-like threat (2015, 63). Overspill is of concern due to unintended consequences that might accompany losing control of a tool. Concern for this overreach can be seen in several tools: Confikr contained a test wherein if a computer was using a Ukrainian keyboard it would not be effected, likely for legal reasons (Bowden 2011, 56). Similarly, Sutxnet, the poster child of high-end malware, was carefully written to only effect the very particular combinations of hardware and software that were used in a certain Iranian nuclear refinery.

Second, as the Stuxnet dossier shows, creating highly targeted, hard hitting software is difficult. Sophisticated though it was, Stuxnet was only a tactical tool. If one compares it to its brethren in the Olympic Games campaign, it quickly becomes apparent that Rid's tradeoffs between complexity and range are very real. The ability to create a series of these tools to be used in coordination required not one but at least four separate tools, each with its own purposes. To return to the air power analogies that were made earlier in this paper, to create a one-size-fits-all cyber weapon would be similar to creating a single airplane that could simultaneously act in all the roles needed by a modern air force.

As Harris states, defense-in-depth is the standard when protecting digital assets. In general, the more valuable the target, the greater number of layers of defense it utilizes. While each layer may have one or multiple vulnerabilities, finding the correct threats to exploit these vulnerabilities in sequence is time consuming and difficult, especially if any failed attempts will result in patches to the system, thus negating previous work. Moreover, there are communities of professionals in place that frequently communicate known vulnerabilities, meaning that a failed attempt on one target might result in other targets becoming aware of their vulnerabilities.

Targets can be classified into five tiers, numbered one to five, with an additional level (numbered zero) for specific cases. These tiers are clusters of traits that should coalesce given a network's assets, contents, and security goals and tolerances. While they are not hard-edged cases, they are useful when thinking about a network or company's security stance.

Type one targets are typical of military installations. Given their highly secret nature and zero tolerance for breaches, they are highly controlled and have very complex security protocols that control both digital and physical access (NAVFAC 2015). Typically, they are secured beyond a level necessitated by normal compliance standards and are quick to react to any perceived threat. Due to their nature, however, they are slow to disclose any known vulnerabilities, meaning that attacks made on these systems can possibly be repeated elsewhere against non-military targets.

Tier two targets are less complex but still in compliance with a rigid set of industry standards. A typical type two target is a financial institution, multinational corporation, or non-military government facility. Their need to be secure is balanced by their need to be accessible to a wide number of users, creating inherent tradeoffs in security versus usability. While security is a priority for these groups, often times there is an element of calculated risk, balancing money spent on security against the costs associated with being exploited. They can be slow or fast to react to threats, with private sector entities tending to react faster due to the need to be seen as secure. If exploited, they are typically quick to disclose the attack.

The third type of target is typical of a business of medium size. It is generally in compliance with industry standards, though these may be allowed to slip between evaluation periods. These industry standards, like PCI (for taking credit cards), are usually the extent of security procedures. If they are attacked, they may be less quick to react and will only disclose the attack if it is deemed economically sound, especially if they are not in an industry that is especially security conscious.

Fourth is the level of security seen in most homes and small businesses. There is little, if any regulated security structures, and most of these are protocol driven and standardized. As a group, type four targets are slow to react to vulnerabilities and are often unaware of their risk. If an attack or vulnerability is discovered, it is likely reported quickly through professional groups though updates to mitigate the vulnerabilities are often slow to be applied.

The last level, zero, is reserved for hardware and embedded systems. These are typical of consumer level "Internet of Things," devices, though other devices like rolling-code garage doors (the most common type of remotely operated door) and Bluetooth devices also fall into this category. If vulnerability is discovered, it is often hard or impossible to mitigate without replacing hardware.

These levels are, as stated before, not clear-cut distinctions but rather common clusters of traits based on priorities. They become significant in state-level security because they offer a series of tradeoffs that a potential attacker or defender must consider. On one hand, the simplest targets are easy to attack and control but are of limited instrumentality and clearly within the civilian sphere, meaning that blowback

will likely be significant. To balance that, the more instrumental targets (like military systems) are more difficult to attack and also less likely to be disclosed publicly if successful.

The best target for an attacker is the one that is likely to be disclosed publically but does not directly affect the population at large and must balance the ability to rapidly change with the ability to be accessible to a large number of people. This means that type two systems that do not directly affect the public are the ideal target to undermine trust in institutions. Because they are quick to disclose attacks and slow to be able to change, the likelihood of creating great effects is significant.

Still, within the past year, there have been two attacks that fit this description, one on the US Office of Personnel Management and another one on the American offices of Sony Pictures. Neither, however, seems to have gained the traction that would be needed to be deemed effective.

Conclusion

While the term "cyber warfare," is used often, it is difficult to foresee a true cyberwar. Though network-based attacks are capable of producing some impact on physical systems, these are mostly one-off in nature and require a significant amount of time and planning per attack, making them an unlikely choice as a primary coercive measure. Similarly, though much time and effort has been put into planning for a limitless cyber attack on the US population, there is evidence that this type of attack will have little chance of success. The most likely role for cyber attacks in the near future will likely be variations on how they are used today: espionage, sabotage, and subversion. Still, there are models that can be used to classify targets, attackers, their tools, and their aims. By using these frameworks we can get a better understanding of the computer battlefield. In the near future, analytic value will come from seeing if and how these frameworks continue to be relevant.

References

Andress, Jason, and Steve Winterfeld. 2011. *Cyber Warfare: Techniques, Tactics and Tools for Security Practitioners*. Kindle Edition. New York: Syngress Press.

Bentek Systems. 2012. "Internet and Web SCADA." http://www.scadalink.com/support/web-based-scada.html (accessed September 14, 2015).

Bowden, Mark. 2011. *Worm: The First Digital World War. Kindle Edition*. New York: Atlantic Monthly Press.

Bradbeer, Thomas G. 2004. "Battle of Air Supremacy Over the Somme: 1 June–30 November 1916." MS Thesis, US Army Command and General Staff College.

Brenner, Joel. 2011. *America the Vulnerable: Inside the New Threat Matrix of Digital Espionage, Crime, and Warfare*. Kindle Edition. New York, New York: The Penguin Press.

Carr, Jeffery. 2009. *Inside Cyber Warfare: Mapping the Cyber Underworld*. Kindle Edition. Boston, Massachusetts:O'Riley.

CNET. 2012. "A Who's Who of Mideast-Targeted Malware." http://www.cnet.com/news/a-whos-who-of-mideast-targeted-malware/ (accessed September 14, 2015).

Dhanjani, Nitesh, Billy Rios, and Brett Hardin. 2009. *Hacking: The Next Generation*. Kindle Edition. Boston, Massachusetts:O'Riley.

Falliere, Nicolas, Liam O. Murchu, and Eric Chien. 2011. "W32.Stuxnet Dossier." http:/www.symantec.com/content/en/us/enterprise/media/security_response/whitepapersw32_stuxnet_dossier.pdf

Ferrari, Dino, Trans. 1942. *Giulio Douhet Command of the Air*. New York: Coward McCann.

Harris, Shon. 2012. *CISSP All-in-One Exam Guide*, Sixth Edition, Kindle Edition. New York, New York:McGraw-Hill Education.

Healey, Jason. 2013. *A Fierce Domain: Conflict in Cyberspace, 1986 to 2012*. Kindle Edition. Vienna, Virginia:Cyber Conflict Studies Association.

Pape, Robert A. 1996. *Bombing to Win: Air Power and Coercion in War*. Ithica, New York: Cornell University Press.

NAVFAC. "Physical Security of Sensitive Compartmented Information Facilities." https://www.wbdg.org/pdfs/dod_at/navfac_scif_ho.pdf (accessed September 14, 2015).

Rid, Thomas. 2013. *Cyber War Will Not Take Place*. Kindle Edition. New York: Oxford University Press.

Security Week. "Zero Day Exploits Released in Hacking Team Leak." http://www.securityweek.com/zero-day-exploits-leaked-hacking-team-breach (accessed September 14, 2015).

Valeriano, Brandon, and Ryan C. Maness. 2015. *Cyber War versus Cyber Realities: Cyber Conflict in the International System*. Kindle Edition. New York: Oxford University Press.