
Towards a Cyber Conflict Taxonomy

Scott D. Applegate

Center for Secure Information Systems
George Mason University
Fairfax, Virginia
sapplega@gmu.edu

Angelos Stavrou

Center for Secure Information Systems
George Mason University
Fairfax, Virginia
astavrou@gmu.edu

Abstract: This paper seeks to create a practical taxonomy to describe cyber conflict events and the actors involved in them in a manner that is useful to security practitioners and researchers working in the domain of cyber operations. The proposed Cyber Conflict Taxonomy is an extensible network taxonomy organized as a plex data structure. Subjects of the taxonomy are entered as either Events or Entities and are then categorized using the categories and subcategories of Actions or Actors. Each of these categories is further subdivided into increasingly specific subcategories used to describe the defining characteristics of each subject and labeled lateral linkages are used to illustrate the associative relationships between Entities and Events. The categories are organized in both a hierarchical and associative manner to illustrate the relationships between subjects and categories. A prototype of this taxonomy was developed and tested using a test set of recent cyber conflict events and used to explore the relationship and connections between these events and the states, groups or individuals that participated in them. Furthermore, this taxonomy can potentially identify actors across different events based on their similar method of operation, toolsets and target sets.

Keywords: *Cyber Conflict, Cyber Operations, Taxonomy*

1. INTRODUCTION

This paper seeks to construct a practical and comprehensive taxonomy to describe cyber conflict events and the actors involved in them in a manner that is useful to security practitioners and researchers working in the domain of cyber operations. Our aim is to provide an organized formal model that can be used to measure the impact of attacks and different defense strategies both in specific scenarios and in large-scale cyber conflicts. To study a subject effectively, one must have some means of organizing the knowledge related to that subject. A taxonomy provides a logical organizational framework for doing this and can act as a tool to assist users in visualizing relationships and classifying data in a useful manner. The military strategist Carl von Clausewitz discussed the importance of the “coup d’oeil” which he roughly described the ability for a military leader to be able to see and immediately grasp the implications of a military situation with one “cast of the eye” [1]. With this in mind, this project attempts to create a Cyber Conflict Taxonomy that will give the security practitioner a coup d’oeil of cyber conflict related events.

The use of the term Cyber Conflict Taxonomy versus a Cyber Warfare Taxonomy in this project seeks to recognize the fact that other entities beyond states, such as non-state actors, hacktivists groups and even private individuals, are playing a role in the ongoing hostile, politically motivated actions that are taking place in cyberspace. It is therefore important that a taxonomy designed to describe these events and actors take that fact into account, hence, the proposed taxonomy will attempt to describe not just events that take place solely between nation-states, but also events undertaken by non-state entities directed at other competitor states for political, nationalistic or ideological purposes.

To further this effort, a review of previously developed taxonomies was undertaken to give the paper a logical starting point and to determine what previous works were relevant to this work. To date, no one has undertaken a taxonomy specifically geared towards classifying and understanding cyber conflict, but numerous taxonomies have been created that address cyber threats and other aspects of cyber security.

2. SURVEY OF PREVIOUS RELEVANT TAXONOMIES

A great deal of previous work has been done in the area of classifying threats and vulnerabilities. Early taxonomies such as Bishop’s 1995 work focused on categorizing security vulnerabilities in software to assist security practitioners in maintaining more secure systems through an understanding of these vulnerabilities [2]. John Howard extended this idea in his 1997 work in which he analysed and

classified 4299 security related incidents on the internet. Howard's work was notable because he included attackers, results and objectives as classification categories expanding threat taxonomies beyond the technical details of an attack to include more intangible factors such as an attacker's motivation for conducting an attack [3]. Hansman and Hunt created a unique taxonomy in 2004 which was designed to be used by information bodies to classify new attacks. This taxonomy was based on four dimensions but was also designed to be extensible in that additional dimensions, some of which the authors suggested, could be added to the taxonomy as needed [4].

The vast majority of threat taxonomies are designed as attacker-centric frameworks which categorize attacks from the perspective of an attacker's tools, motivations and objectives. Killouri, Maxion and Tan created a taxonomy in 2004 designed to be defense-centric based on how an attack manifested itself in the target systems. Based on a test set of 25 attacks, this taxonomy was able to predict whether or not the defenders detection systems would be able to detect a given type of an attack [5]. In a similar effort, Mirkovic and Reihner created a taxonomy of Distributed Denial of Service (DDoS) Defenses which categorized DDoS defense mechanisms based on activity level, degree of cooperation and deployment location [6]. These two taxonomies are among the few that classify threats or security incidents from a defensive viewpoint and show the importance of addressing such issues from different perspectives to gain a more holistic view of security issues.

Another approach towards classifying cyber-attacks is to look at the actors involved versus the actual attacks. Kjaerland's 2005 study categorized cyber intrusions based on four categories; (1) method of operations, (2) impact of the intrusion, (3) source of the intrusion and, (4) target [7]. This study examined the likelihood of attacks against different kinds of targets and the likelihood of various kinds of attacks occurring together on a given target. It proved very valuable to this project in that it examined relationships between targets and the impact of attacks on those targets. In 2005, Rogers was one of a number of researchers who attempted to classify the actual attackers themselves. The Rogers' study modeled its taxonomy using a modified circular order circumplex which classified eight levels of hackers across two principal dimensions of skill and motivation [8].

Researchers at the University of Memphis created a cyber-attack taxonomy called AVOIDIT in 2009 which described attacks using five, extensible classifications: Attack Vector, Operational Impact, Defense, Informational Impact, and Target [9]. This taxonomy was created as a network plex taxonomy which, unlike previous efforts, allowed the classification of blended attacks. Additionally, it also allowed for the classification of attacks by both operational and informational impacts and was designed to help educate defenders by looking at attacks' various impacts,

vectors or target types. While this taxonomy focused exclusively on cyber-attacks, its structure and style were very useful in designing the proposed taxonomy in this paper, especially the ability to view and categorize attacks from different taxonomic perspectives.

In recent years, a number of researchers have begun to look at creating taxonomies specifically addressing SCADA systems. In 2010 Fovino, Coletta & Masera created a comprehensive taxonomy describing SCADA architecture, vulnerabilities, attacks and countermeasures [10]. In 2011 Zhu, Joseph, & Sastry highlighted the difference between what they termed standard information technology (IT) systems versus SCADA systems and focused on systematically identifying and classifying attacks against SCADA systems [11]. Neither of the papers presented a taxonomic view describing relationships between the areas they addressed and both focused on attacks while excluding many other relevant details such as actors, impact of the attacks or characteristics of the attacks such as attack vectors.

Moving outside the realm of traditional IT threat taxonomies, Cebula & Young created taxonomy of operational cyber security risks in 2010 which categorized risks into four classes: (1) actions of people, (2) systems and technology failures, (3) failed internal processes, and (4) external events. A valuable aspect of this taxonomy was its insight into the fact that risks can cascade and “that risks in one class can trigger risks in another class” [12]. This insight demonstrated the difficulty in trying to quantify events in a mutually exclusive manner when dealing with complex interactions in cyber security risk. This insight also holds true when trying to identify and classify the complex interactions involved in cyber conflict and was a contributing factor to the development of a network plex topology for the proposed taxonomy in this paper

3. REASONS TO CREATE A CYBER CONFLICT TAXONOMY

As the preceding section demonstrates, there are a number of previously developed taxonomies that address various aspects of cyber threats. While almost any cyber-attack can be categorized and described using these taxonomic frameworks, none of these previous frameworks are capable of illustrating the complex interactions between attacks, actors and other potentially related events and connecting them through logical links that formally describe their relationships. Previous taxonomies are valuable in classifying technical threats and vulnerabilities, but will fall short when it comes to linking actors with different methodologies, goals and patterns of behavior. For security practitioners operating in the realm of cyber conflict, understanding these interactions and the relationships between various aspects

of cyber conflict events can be critical in developing strategy and doctrine. For cyber operations practitioners who must develop doctrine and strategy, the ability to classify and study conflict related events from various taxonomic perspectives can give them unique insights that are not supported by previous works.

To address these issues, the proposed taxonomy has been developed to give users the ability to classify events and expose logical connections and links between different actors, types of attacks and vectors used and various types of impacts associated with each event. Once data is entered into the taxonomy, users can also look at cyber conflict events from discrete taxonomic perspectives such as looking at all events related to a particular actor or all attacks which use a social engineering vector, etc. and then explore the relationships between events and actors to look for commonalities that an operator could act upon.

4. PROPOSED TAXONOMY

The proposed Cyber Conflict Taxonomy is an extensible network taxonomy organized as a plex data structure. Each node in the taxonomy below the four primary category and subject headings can have more than one parent and any secondary or below level item in the plex structure can be linked to any other item based on defined relationships and classifications. This serves to organize the taxonomy into both hierarchical and associative categories which are useful in illustrating the many relationships that can exist between various nodes. The taxonomy is divided into categories and subjects. Categories are the taxonomic classifications that are applied to subjects and are further subdivided into subcategories. Subjects represent the real world events classified as cyber conflict and the real world entities such as individuals, groups or governments that participate in these events. Because cyber conflict involves interactions between states, non-state actors, and other competing entities, it is necessary to have a taxonomy that incorporates both events and entities and applies taxonomic classifications to them both in order to properly understand the complex relationships involved. The initial categories and subjects used in this taxonomy are defined below, however, since this taxonomy is designed to be extensible, additional categories and subjects may be added in the future as necessary.

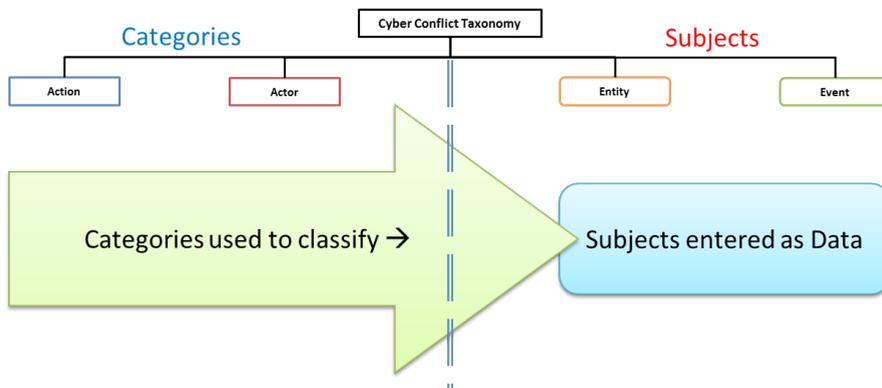


Figure 1. Cyber Conflict Taxonomy

A. SUBJECTS

Subjects are the actual real world cyber conflict related events and the individuals, organizations or states that participated in those events. Subjects represent the data objects that this taxonomy was meant to classify and are divided into Events and Entities. Subjects will always be linked to at least one category or subcategory and more than likely will be linked to multiple subcategories in order to provide accurate and discrete classification of the characteristics of the subject in question. Further subdivision of subjects, beyond Events and Entities, is not necessary for the taxonomy although specifications of subjects can be employed by the user to create logical groupings that may be useful when users wish to create groupings not covered by the actual classification scheme of the taxonomy.

Entities. The Entities subject heading is used to organize and list the actual, real world individuals, groups, organizations or governments that initiated, were targeted or took part in cyber conflict events. Entities will be classified using the Actors category of the taxonomy and will also be laterally linked to the specific Events in which they participated or in which they have suspected involvement. Entities can also be laterally linked to other entities with which they have a defined relationship. An example would be two entities which are directly politically opposed to each other.

Events. The Events subject heading is used to organize and list the actual, real world cyber conflict incidents which will be described in this taxonomy. Events will be hierarchically classified using the Actions category and subcategories of the taxonomy and will also be laterally linked to the specific Entities that participated in these events. Currently, Events are only organized by the specification Year in the

prototype, but no subdivision of Events is actually required by the taxonomy and this specification was added for the author’s purposes.

- Year. The Year specification is an optional subdivision used in the prototype that allows a user to organize events temporally by the year or years in which they occurred. Many events related to cyber conflict span multiple years and it may be valuable for a user to be able to view events from this perspective

B. CATEGORIES

Categories represent the various forms of taxonomic classification used to describe the subjects of this taxonomy. The two primary parent categories in this taxonomy are Actions and Actors which are divided into subcategories as necessary to provide discrete and accurate descriptions of subjects. Subcategories are arranged hierarchically but are applied associatively to subjects so that any given subject will be described by multiple subcategories.

Actions. The Actions category is used to describe cyber conflict events and the characteristics of those events in a manner that is useful for researchers and operators. Actions are subdivided into attack and defense related subcategories.

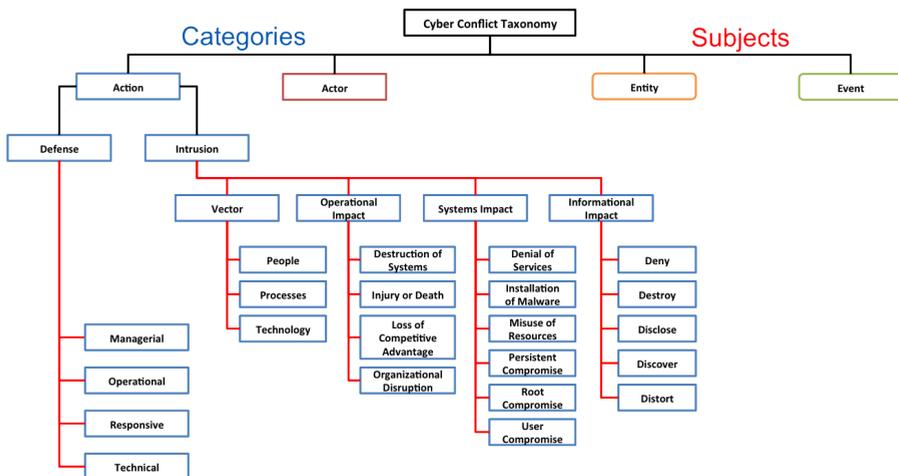


Figure 2. Actions Category of Cyber Conflict Taxonomy

- Intrusion. The Intrusion subcategory describes aggressive actions taken by one actor to affect other actors. Intrusions can be further divided into as many descriptive subcategories as necessary to describe said aggressive action. A

single intrusion may have many characteristics that must be classified in order to accurately classify the event in a complete and useful manner.

- Vector: This subcategory describes the path or means by which an attacker attempts to gain access to information resources or systems. This subcategory has been further divided into vectors which target people, processes or technology. Each of these subdivisions could be further subdivided into increasingly specific and discrete vectors as well.
 - People: This subcategory describes a vector based on the manipulation of people. An example would be the use of social engineering to gain credentials.
 - Process: This subcategory describes a vector based on the manipulation of flawed organizational processes. An example would be an organization that allows a visitor to hand carry their security credentials rather than mandating that the credentials be verified directly with the issuing source. An attacker might exploit this flawed process to illegitimately gain legitimate credentials to a system.
 - Technology: This subcategory describes a vector based on the manipulation of technology and technical processes. An example would be exploiting a vulnerability in a software program.
- Informational Impact: This subcategory describes the impact an intrusion has directly on the victim's information. This subcategory has been further divided into five additional child subcategories.
 - Deny: Denying legitimate users access to information within their own systems or networks.
 - Destroy: Destruction of information, usually through the permanent deletion of files, on a target system or network.
 - Disclose: Illegitimate access to or disclosure of sensitive, confidential or classified information.
 - Discover: Discovery of information previously unknown to an attacker which could potentially give the attacker additional advantages during follow on operations.
 - Distort: Distorting or changing information in a target system in a way that disadvantages the legitimate users of that information and provides advantages for the attacker.
- Operational Impact: This subcategory describes the impact of an intrusion on the victim's operations. The term operational should not be misconstrued

to mean the operational level of war; it is used in this context to indicate the effects of an intrusion on the personnel, business processes and operations of the victim or victim organization.

- Destruction of Systems: Impact of an intrusion, which results in actual physical damage or the destruction of systems. The systems in question may be the actual information systems or other types of systems attached to or controlled by information systems. An example of this would be the damage to centrifuges that resulted from the Stuxnet attack.
- Injury or Death: Impact of an intrusion, which results in actual physical injury or death. This subclass could be further subdivided to differentiate between injury or death to human beings versus injury or death to non-human life. For example, a cyber attack which causes the injury or death of wildlife or livestock.
- Loss of Competitive Advantage: Impact of an intrusion which results in a victim organization losing its competitive edge due most likely to disclosure of plans, proprietary information, classified information or confidential technical data. An example would be a competitor state stealing data from a defense contractor related to a classified technology which enables it to reverse engineer this technology for its own use.
- Organizational Disruption: Impact of an intrusion, which causes the disruption of operations within an organization. An example would be altering information in a supplier database system to reroute critical supplies to the wrong destinations.
- Systems Impact: This subcategory describes the impact of an intrusion on the actual information systems of the victim organization.
 - Denial of Service: Denying a victim access to information resources or system services.
 - Installation of Malware: The installation of malicious software onto the target host or system beyond what is required for the initial compromise of the system in question.
 - Misuse of Resources: An unauthorized use of system resources. This may consist of any system related function that requires certain elevated privileges and those privileges are then converted into abusive action [9].
 - Persistent Compromise: Gaining a persistent foothold on a particular host or within a particular network that goes undetected for an extended period of time. This type of compromise may remain undetected for months or even years and is usually used to facilitate other actions.

- Root Compromise: Gaining unauthorized root or administrative privileges on a particular host or system.
- User Compromise: Gaining unauthorized use of a non-administrator's user privileges on a particular host or system.
- Defense. The Defense subcategory describes actions taken by an actor to protect their information systems from attacks. Defense is divided into Managerial, Operational, Responsive and Technical subcategories, which can be further subdivided into more specific subcategories as is necessary for the user. Three of these subcategories roughly align to the security controls advocated by the National Institute of Standards and Technology [13]. The fourth subcategory, Responsive defenses, expands on the NIST standard to account for more active responses such as counter-attacks which would not be seen in a commercial setting but which could certainly be used in a cyber conflict setting.
 - Managerial: Defensive techniques and methods, normally addressed by management, regarding an organization's computer security strategy.
 - Operational: Defensive strategies based on policies and procedures implemented and executed by people, as opposed to systems, to improve the security of a system or group of systems.
 - Responsive: Direct responses to a malicious intrusion targeting the source of the intrusion. Examples could include counter-attack or counter-reconnaissance.
 - Technical: Defensive tools or strategies executed by automated systems to improve the security of individual systems or a group of systems.

Actors. The Actors category classifies the entities participating in cyber conflict by type. Currently, this category is divided into two subcategories; Non-State Actors and State Actors. These subcategories may be further divided down as needed.

- Non-State Actors: The Non-State Actors subcategory describes entities participating in cyber conflict events, which have no known ties to government entities.
- State Actors: The State Actors subcategory describes governments, government organizations or government sponsored entities that participate in cyber conflict events.

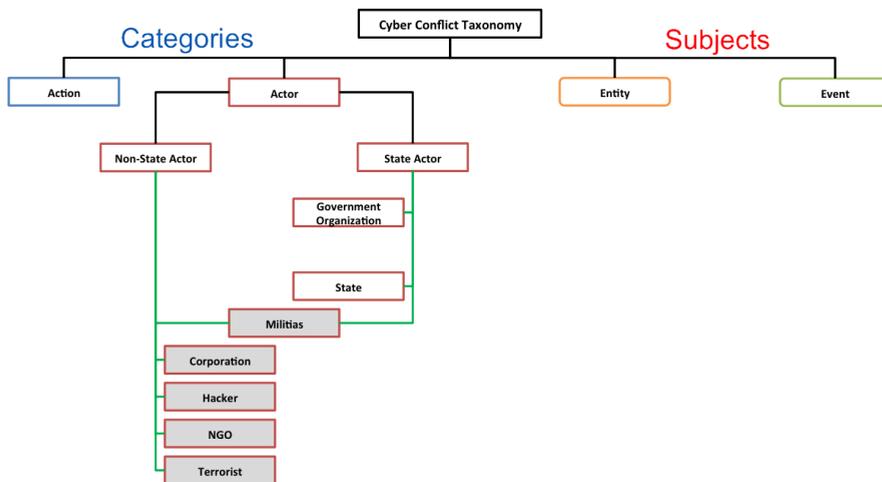


Figure 3. Actors Category of Cyber Conflict Taxonomy

C. TYING IT ALL TOGETHER

In order to begin testing the usefulness of the proposed taxonomy, two prototypes were developed. The first prototype was modelled using mind-mapping software called *The Brain*. Version 7 of this software was used for the development of the initial prototype. This software was used to rapidly build and visualize the proposed taxonomy. This first prototype provides the ability to show multiple child- and parent-relationships hierarchically in a network plex and to laterally link related entities and events together depicting the causal relationship between various subjects. The prototype also allows the user to define the different types of relationships that link nodes together throughout the taxonomy and to color-code, tag and categorize both nodes and links. This allows the user to search or filter the taxonomy based on key words, node types or even relationship types.

A sample set of a ten real world events was entered into the taxonomy as Events and then classified using the categories and subcategories previously described. Additionally, more than fifty entities were additionally entered into the taxonomy based on their relationship to the previously entered events. These entities represented the actors involved in these events, including those suspected of involvement in cases where definitive attribution (i.e. most cases) could not be established. This prototype proved to be very useful in developing classification categories and in visualizing the data entered into the taxonomy. The main limitation of this prototype, based primarily on the software package used to develop it, was the need to manually link each subject entered into the taxonomy to the various categories

and subcategories that would apply to it. For a large data set, this would be a very tedious task prone to omissions and errors. Ideally, a fully automated and polished version of the taxonomy would include simple drop lists with all the categories from which the user could select multiple classifications simultaneously to describe the subject. Additionally a similar list of subjects would be available to simultaneous select related or causal subjects as well.

A second prototype of the proposed taxonomy was modelled using *Protégé* version 4.1. *Protégé* is a free, open-source platform that provides a suite of tools to construct domain models and knowledge-based applications with ontologies using Web Ontology Language. Use of *Protégé* for the second prototype allowed for more formal and rigorous definitions of the relationships between entities and categories and provided a platform capable of more easily identifying trends in the knowledge base. In defining relationships, *Protégé* allows for the specification of domains and ranges for each relationship. It allows additional facets of such relationships to be specified such as transitive, functional, symmetric, asymmetric and reflexive properties. Additionally, due to the open source nature of the software, it would be easier to alter this platform to provide for easier data entry due to the availability of the original source code.

5. APPLICATION EXAMPLES

To demonstrate the use of the proposed taxonomy three examples are shown below all related to the same event, Operation Shady RAT. This event is shown from three different taxonomic perspectives; one view with the event as the central node in the taxonomy, one from the perspective of one of the event's systems impacts, and finally, a view from the perspective of its suspected initiator. Each view shows different characteristics of the event and illustrates the potential relationships between this event and other entities or events. It should be remembered that in the examples below, only a limited data set of ten events was entered into the prototype.

A. OPERATION SHADY RAT – TAXONOMIC VIEW OF AN EVENT

Operation Shady RAT was a targeted set of intrusions into more than 70 global companies, governments and non-profit organizations that took place from 2006 to 2011 [14]. When entered into the prototype taxonomy (see Fig. 4), the result shows links to the actors which were targeted, the suspected initiating actor, the years over which the event took place, and the various types of impacts. Additionally, other events are shown which took place during the same time frame, which had similar types of impacts, or which were related to the actors listed.

This initial view gives an operator a starting point to begin studying related events in order to look for trends or patterns in the data such as, for example, looking at other events which involved the installation of malware on targeted systems.

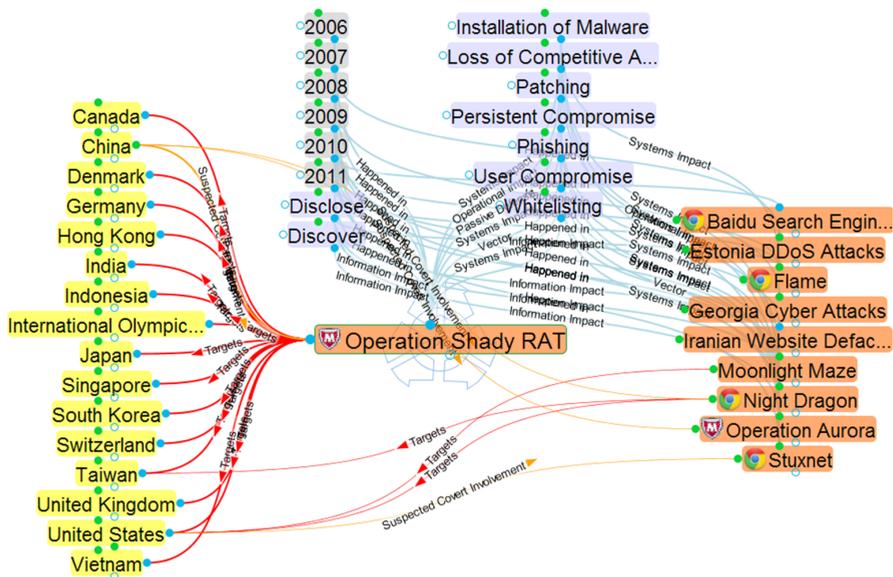


Figure 4. Taxonomic View of Operation Shady RAT

B. INSTALLATION OF MALWARE – TAXONOMIC VIEW OF A SYSTEMS IMPACT

To view this event from a different taxonomic perspective, an operator can simply select one of the categories by which the event was characterized such as the Systems Impact – Installation of Malware. As can be seen in Fig. 5, this view shows the user other events which shared this same systems impact. Additionally it shows links from these other events to additional systems impacts they exhibited allowing the operator to compare impacts of similar events.

C. CHINA – TAXONOMIC VIEW OF AN ENTITY

To view Operation Shady RAT from the perspective of the suspected initiating actor, the operator can select the State – China (see Fig. 6). This perspective shows other events in which China is suspected to have been involved and also displays which other actors were targeted by these events.

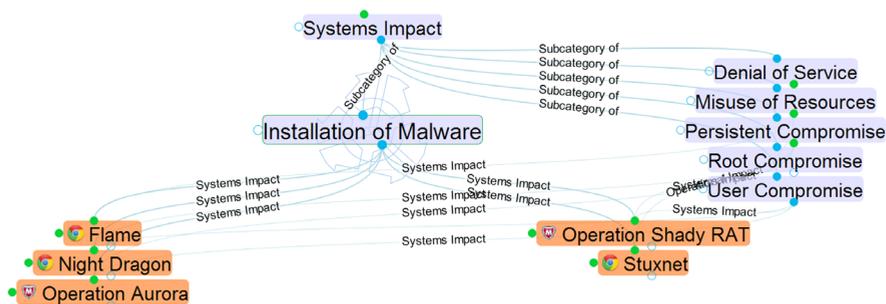


Figure 5. Taxonomic View of Systems Impact: Installation of Malware

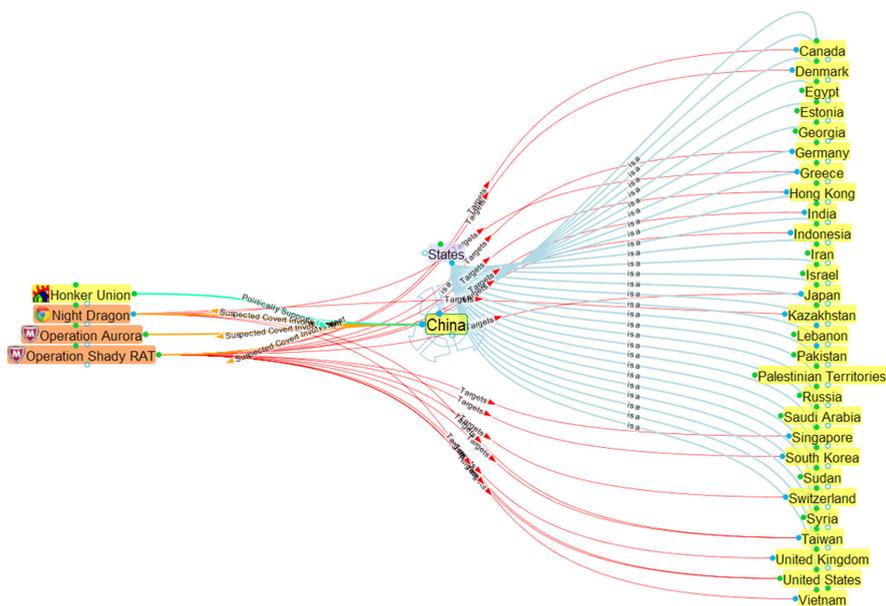


Figure 6. Taxonomic View of Actor: China

If the network plex is expanded by one additional level of connectivity the complexity of the events and interactions related to China becomes apparent. Relationships that are separated by 3 or 4 degrees of separation can now be illustrated and users can look for insightful patterns of behavior or similar methodologies. This expanded plex shows other state and non-state actors involved in similar events, the targets of these events, the time frame of these events and other related information such as the political allegiances of various non-state entities illustrated by the extended connectivity.

D. COMPARISON OF OTHER RELAVANT TAXONOMIES

Using Operation Shady RAT as a case study, the proposed taxonomy in this paper was studied in a side-by-side comparison with two other taxonomic systems previously discussed above. Howard's Computer & Network Attack Taxonomy classifies attacks using five classification categories: Attacker, Tools, Access, Results and Objective [3]. Table I shows the result of classifying Operation Shady RAT using this Taxonomy. While this taxonomy does provide some important information about this attack, it lacks a couple of important characteristics such as vector, defensive actions and the specific actors involved.

Table I. Classification of Operation Shady RAT using Howard's Taxonomy

Howard's Taxonomy						
Name	Attacker	Tools	Access		Results	Objective
Shady RAT	Spies	Toolkit	Design & Config Vulnerabilities	Unauthorized Use Unauthorized Access	Files Compromise of Information Disclosure of Information	Polioitical & Finaicial Gain

The AVOIDIT Taxonomy also classifies attacks using five classification categories: Attack Vector, Operational Impact, Informational Impact, Defense and Target. Table II shows the result of this classifying this attack using the AVOIDIT Taxonomy. While this taxonomy does improve on Howard's in some key areas such as attack vector and defensive strategy, it still lacks specificity when it comes to identifying actors involved in this attack.

Table II. Classification of Operation Shady RAT using AVOIDIT Taxonomy

AVOIDIT Taxonomy					
Name	Attack Vector	Operational Impact	Informational Impact	Defense	Target
Shady RAT	Spear Phishing	Installed Malware: Trojan	Discovery Disclosure	Remediation: Patch System, Whitelisting	Network

Classifying Operation Shady RAT using the proposed taxonomy, the first thing that becomes apparent is the inclusion of all the actors involved in this event (see Table III.). A compressed list was used for this paper as the original attack targeted more than 70 organizations across 14 nation-states. This taxonomy also differentiates between Systems Impact and Operational Impact while the AVOIDIT Taxonomy only highlights the technical impact of attacks on systems and excludes the impact of attacks on the target's operations. All information from the AVOIDIT Taxonomy is accurately captured in the proposed taxonomy and all information from Howard's taxonomy, with the possible exception of the vulnerability portion of Access, are also captured.

Table III. Classification of Operation Shady RAT using Cyber Conflict Taxonomy

Cyber Conflict Taxonomy								
Name	Vector	Informational Impact	Operational Impact	Systems Impact	Defense	Actors		
Shady RAT	Spear Phishing	Discover	Loss of	Installation of	Passive -	Targets:	Canada	Source: China (Suspected)
		Disclose	Competitive Advantage	Malware; Persistent Compromise	Whitelisting, Remediation - Patching	Denmark	Germany	
						Hong Kong	India	
						Indonesia	IOC	
						Japan	Singapore	
						South Korea	Switzerland	
						Taiwan	United Kingdom	
						Unites States	Vietnam	

An important feature of the proposed taxonomy that is not addressed in all of the previous taxonomies is the ability of this taxonomy to identify related subjects (both entities and events). Looking back at Fig. 4, a group of related events appears on the right hand side of the image (the 9 items which are circled). These events all share some of the characteristics of Operation Shady RAT. They may use the same vector, target the same states or organizations, or may have just happened in the same timeframe. Three of the nine events identified share a high degree of similarity with Operation Shady RAT and could potentially be related to this event. Given that this prototype had a very limited test-set, it is easy to see how this capability would be useful for researchers and planners working in the cyber operations domain. This capability can assist a researcher in attributing an anonymous event to a specific actor based on similarities in methodology, impacts and target sets.

Each of the above taxonomic frameworks can provide useful information; however, the proposed taxonomy provides the most robust classification scheme and provides the ability to identify related subjects. This improvement on previous taxonomic frameworks and the focus on cyber conflict events at an operational level make this proposed taxonomy a useful tool for both security researchers studying cyber conflict and for planners and operators working in the domain of cyber operations.

6. LIMITATIONS AND FUTURE RESEARCH

Over the course of this research, a number of limitations were identified in relation to the use of a taxonomy to evaluate cyber conflict events. Introducing such a taxonomy to classify the events and entities involved in cyber-conflict is important and offers a good first approximation of what a security analyst can derive and potentially plan for when it comes to cyber operations. However, there are inherent limitations that stem from the use of a taxonomy, which is a hierarchical categorization of entities within a domain. A taxonomy does not allow for any formal or empirical

relationships among the entities beyond parent-child relationships. To capture most, if not all possible relationships and characteristics between different actors and events, a more formal mechanism such as an ontology is needed. Unlike a taxonomy, an ontology allows for the formal description of multiple relationships between entities in an empirical manner. The creation of the second model using Protégé and OWL constituted the first step in this process and will be used in future research to expand the scope of this project. Once this second model has been more extensively defined and tested, a larger data set will be used to validate the model's ability to identify commonalities between related events.

7. CONCLUSION

This paper presents a taxonomy for classifying cyber conflict events and the entities involved in these events. All data are entered into this taxonomy as subjects and then classified according to the categories and subcategories used to describe the characteristics of these subjects. A prototype was developed which demonstrated that the proposed Cyber Conflict Taxonomy is useful in categorizing and describing events and entities involved in cyber conflict in a manner that would be beneficial to researchers and operators. All events and actors entered into the prototype were fully describable using the proposed categories. Even with a limited data set, the ability to study linkages between related subjects demonstrates patterns and provides researchers with insights into commonalities between different events and entities and would be useful when developing doctrine and strategy. This feature is unique to this taxonomic model and is an improvement on previous frameworks. It can potentially allow an operator to identify actors across different events based on their similar method of operation, toolsets and target sets.

Finally, this taxonomy is designed to be extensible so that users can categorize the characteristics of cyber events or entities using increasingly discrete descriptions. This allows this framework to be as specific as necessary for various purposes. For future work, a much larger data set should be created and empirical studies undertaken to validate the taxonomy's ability to identify commonalities between related events.

Acknowledgements

The authors would like to gratefully acknowledge the efforts of LTC André Abadie, COL Jody Prescott (Ret.), and Dr. Duminda Wijesekera who assisted in the editorial review of this paper. Portions of this project were conducted using the Protégé resource, which is supported by grant LM007885 from the United States National Library of Medicine.

REFERENCES

- [1] Lambe, P. (2006, April 18). Defining Taxonomy. Retrieved from Green Chameleon: http://www.greenchameleon.com/gc/blog_detail/defining_taxonomy/
- [2] Bishop, M. (1995). A Taxonomy of UNIX System and Network Vulnerabilities (University of California at Davis No. Report CSE-95-10). Retrieved from <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.33.5712>
- [3] Howard, J. D. (1997). An Analysis of Security Incidents on the Internet 1989-1995 (Doctoral dissertation, Carnegie Mellon University, Pittsburgh, PA, 1997). Retrieved from www.cert.org/archive/pdf/JHThesis.pdf.
- [4] Hansman, S., & Hunt, R. (2004). A taxonomy of network and computer attacks. *Computers & Security*, 24(1), 31-43. <http://dx.doi.org/10.1016/j.cose.2004.06.011>
- [5] Killourhy, K. S., Maxion, R. A., & Tan, K. M. C. (2004). A Defense-Centric Taxonomy Based on Attack Manifestation. Presented at the International Conference on Dependable Systems & Networks, Florence, Italy.
- [6] Mirkovic, J., & Reiher, P. (2004). A Taxonomy of DDoS attack and DDoS defense mechanisms. *ACM SIGCOMM Computer Communication Review*, 34(2), 39-53. <http://dx.doi.org/10.1145/997150.997156>
- [7] Kjaerland, M. (2006). A taxonomy and comparison of computer security incidents from the commercial and government sectors. *Computers & Security*, 25(7), 522-538. Retrieved from <http://dx.doi.org/10.1016/j.cose.2006.08.004>
- [8] Rogers, M. K. (2006). A two-dimensional circumplex approach to the development of a hacker taxonomy. *Digital Investigation*, 3(2), 97-102. Retrieved from <http://dx.doi.org/10.1016/j.diin.2006.03.001>
- [9] Simmons, C., Ellis, C., Shiva, S., Dasgupta, D., & Wu, Q. (2009). AVOIDIT: A Cyber Attack Taxonomy. Retrieved from http://issrl.cs.memphis.edu/files/papers/CyberAttackTaxonomy_IEEE_Mag.pdf
- [10] Fovino, I. N., Coletta, A., & Masera, M. (2010, March). Taxonomy of security solutions for the SCADA Sector, Deliverable: D 2.2, Version: 1.1. A European Network For The Security Of Control And Real Time Systems.
- [11] Zhu, B., Joseph, A., & Sastry, S. (2011). A Taxonomy of Cyber Attacks on SCADA Systems. *IEEE International Conferences on Internet of Things, and Cyber, Physical and Social Computing*. DOI 10.1109/iThings/CPSCCom.2011.34
- [12] Cebula, J. J., & Lisa, R. Y. (2010). A Taxonomy of Operational Cyber Security Risks (Carnegie Mellon University / Software Engineering Institute No. CMU/SEI-2010-TN-028). Retrieved from <http://www.sei.cmu.edu/library/abstracts/reports/10tn028.cfm>
- [13] National Institute of Standards and Technology (2009). NIST Special Publication 800-53 Revision 3: Recommended Security Controls for Federal Information Systems and Organizations. National Institute of Standards and Technology. United States Department of Commerce. Gaithersburg, MD.
- [14] Alperovitch, D. (Vice President, Threat Research, McAfee). (2011). Revealed: Operation Shady RAT. McAfee. Retrieved from <http://www.mcafee.com/us/resources/white-papers/wp-operation-shady-rat.pdf>.