To appear in *Proc. Twenty-seventh Annual Int. Computer Software and Applications Conf.*, IEEE (Dallas, Tex., Nov. 2003).

Measured Responses to Cyber Attacks Using Schmitt Analysis: A Case Study of Attack Scenarios for a Software-Intensive System

[1]James B. Michael, [2]Thomas C. Wingfield and [3]Duminda Wijesekera
[1]Department of Computer Science, Naval Postgraduate School, Monterey, California 93943-5118
[2]The Potomac Institute for Policy Studies, 901 North Stuart Street, Suite 200, Arlington, VA 22203
[3]Center for Secure Information Systems, George Mason University, Fairfax, VA 22030
[1]bmichael@nps.navy.mil, [2] twingfield@potomacinstitute.org, [3]diwjesek@gmu.edu

Abstract

In this paper we address the development of measured responses to coercive actions. We demonstrate, via a case study of kinetic and cyber attacks on a safety-critical software-intensive system, the application of the Schmitt Analysis to the question of whether the attacks have risen to the level of a "use of force" under international law, taking into account both the quantitative and qualitative aspects of the attacks.

1. Introduction

Subway systems have been the target of many terrorist acts, partly because they are relatively easy targets and such attacks can instill fear across a wide spectrum of the populace. These attacks have been primarily kinetic in nature, in addition to involving the physical presence of someone on the subway system to perform the terrorist act, such as to either place or detonate a bomb.

The control systems embedded in modern subways are software intensive, meaning that the majority of the control system's functionality is implemented in software rather than hardware. Such software is deemed to be safety critical in that it controls the forces exerted by the subway system (*e.g.*, acceleration of trains). Thus, one could envision that a well-trained terrorist cell with a high level of competency in applying information technology could remotely initiate a cyber-based attack on a subway system, with the effects being on par with that of a purely kinetic-based attack, for instance, causing a high delta-velocity collision between two trains by, in part, disabling the automatic train protection (ATP) system.[1] There have been reports that some members of the AlQaida terrorist network, for example, have probed for vulnerabilities in distributed control systems (DCS) and supervisory control and data acquisition (SCADA) systems [2].

The law enforcement, military, and intelligence communities obviously need to abide by U.S. domestic law while fashioning responses to terrorist acts. Less well understood by laymen is the need for the U.S. to abide by those portions of international law that the U.S. recognizes.[2]

Pagni [5] conducted a case study of the infamous sarin nerve gas[3] attack on the Tokyo subway by the Aum Shinri-kyo religious sect, with the aim of identifying key challenges for improving U.S. domestic preparedness in the area of consequence management: mitigating or ameliorating the effects of an attack (*e.g.*, rescue and treatment of wounded, decontamination). Pagni observes that "legal preparation is a vital but often overlooked aspect of domestic preparedness" and that such preparedness "…affords law enforcement the necessary powers to investigate and prosecute those who possess or attempt to use…" weapons of mass destruction (*e.g.*, sarin).

In this article, we argue that the novel and amorphous nature of legal reasoning in the area of responding to terrorist acts calls for a disciplined, principled analysis. While a full analysis may not be reduced to a simple mathematical algorithm, it is possible to objectively frame those aspects that are more concrete. This narrows the "gray area" of uncertainty, and provides a framework for evaluating differences in interpretation of the law. The Schmitt Analysis, then, is useful as a legal algorithm, but it is even more useful as a method for highlighting areas of uncertainty or disagreement in multiple legal analyses, and for providing a principled means by which to address all relevant aspects of a use of force against software-intensive systems that are part of the critical infrastructure.

[1] The train protection system is one of four subsystems of a train control system. It is responsible for "assuring safe train movement by a combination of train detection, separation of trains running on the same track or over interlocked routes, overspeed prevention, and route interlocking." [7]

[2] The following is of the specific priorities listed in the National Strategy to Secure Cyberspace [8]:

> "When a nation, terrorist group, or other adversary attacks the United States through cyberspace, the U.S. response need not be limited to criminal prosecution. The United States reserves the right to respond in an appropriate manner. The United States will be prepared for such contingencies."

[3] Sarin is a nonpersistent organophosphorus compound developed for use as a chemical weapon. Sarin, in addition to its other effects, paralyzes the muscles around the human lungs, causing the victim to suffocate.

2. Schmitt Analysis

The first question faced by lawyers evaluating attacks on critical infrastructure is, “When does the attack rise to the level of a ‘use of force’ under international law?” Military operators and their legal advisors have ample precedent to give finely-calibrated answers for attacks involving more traditional, kinetic means (*e.g.*, bombs and bullets), but the novelty of information operations, with its new digital weapons, modes of attack, and novel target lists, is more problematic.

Until recently, there were two broad schools of thought. The first, or “common sense” approach, was the more popular with laymen. It (very reasonably) postulated that the international legal regime was ostensibly in place to keep sub-war level operations from mushrooming into full-blown wars (and, conversely, to identify war-level activities as soon as possible to attach the appropriate legal protections to the participants as soon as possible). This being the case, the simplest and most sensible approach to applying the kinetic legal regime to the digital battlefield was to disregard the means of attack, and concentrate solely on the quantum of damage done. Put another way, it should be immaterial whether a refinery was destroyed by a 2000-lb bomb or a line of malicious code in its pressure-regulation subroutine; what did matter was the size of the smoking hole left in the ground after the attack. This quantitative approach had the benefit of simplicity, clarity, and logic. Unfortunately, it had the inescapable fault of being out of sync with the prevailing structure of international law, the UN Charter paradigm.

This second approach, more popular in academic circles, followed the logic of the Charter’s framers to its literal conclusion: that anything other than an armed attack (something very much like the tanks-across-the-border threat the Charter was written to obviate) was permissible. Or, in other terms, the quantity of force was less important than the quality of force: military coercion was to be discouraged (with a very low threshold of permissible activity), while diplomatic, economic, and political coercion were to be encouraged, or at least less discouraged, as a peaceful alternative to blitzkrieg. This approach had the advantage of academic purity and consonance with mainstream international legal thinking; unfortunately, by deploying a half-century old legal theory, it failed to take into account the newly-destructive capacities of what had been mere messages and signals.

This impasse lasted until the very end of the Twentieth Century, when Schmitt [6] proposed a reconciliation of the apparent conflict. In his words, “… as the nature of a hostile act becomes less determinative of its consequences, current notions of ‘lawful’ coercive behavior by states, and the appropriate responses thereto, are likely to evolve accordingly.”[4] That evolution will be driven by a return to first principles, and an expression of quantitative determinations with intellectually honest qualitative descriptions. Specifically, Schmitt examined why the framers of the Charter chose to characterize each type of coercion as they did. By applying a quantitative scale to each of the seven factors he identified, any given operation could be described in qualitative terms as being closer to one end of a spectrum or the other. In other words, an action’s qualitative nature (in seven more or less binary areas) could be determined by applying any fixed quantitative figure (say, a one-to-ten scale). Schmitt’s contribution in translating the qualitative Charter paradigm into its quantitative components—the legal equivalent of going from analog to digital—provides a framework for scholars and practitioners to organize analysis in something other than a quantum cloud of subjective uncertainty.

3. Scenario I: A Kinetic Attack

In our first scenario of events for which we will access our legal options for responding to an attack, terrorists after crossing an international border, release sarin gas on the Washington Metro—the subway system in Washington, D.C.—during rush hour.[5] The terrorists are citizens of countries with which the U.S., at the time of the attack, is nominally at peace.

3.1. Analysis

Severity: The release of the chemical agent on the Metro injured approximately 1,000 travelers and caused ten deaths. Fifteen minutes after the attack, law enforcement authorities

[4] Schmitt explains:

> “In the current normative scheme the consequences of an act are often less important than its nature. For instance, a devastating economic embargo is not a “use of force” nor an “armed attack” justifying forcible self-defense, even though the embargo may result in enormous suffering. [footnote omitted] On the other hand, a relatively minor, armed incursion across a border is both a use of force and an armed attack. [footnote omitted]. This contrary result derives from the law’s use of “acts” as cognitive shorthand for what really matters—consequences. Acts are more easily expressed (to “use force” versus to cause a certain quantum and quality of harm) and more easily discerned than an effects-based standard, on the harm suffered. This cognitive shorthand does not work well in the age of information operations because information attacks, albeit potentially disastrous, may be physically imperceptible.”

[5] The Washington Metro was designed from the start to be fully automated [3]. Although there is a conductor on each train, that person’s primary responsibilities are to assist in physical security (*e.g.*, detect obstacles on the tracks or platform) and take emergency measures if the automated system fails to operate correctly.

began evacuating all passengers, except emergency workers, from the entire Metro system. Physical property damage was negligible. However, there was a loss of intangible property, such as opportunity costs and lost productivity (*e.g.*, extra time required to travel to the workplace, which could otherwise been used to produce something or provide a service).

Immediacy: The immediate attack took a matter of minutes. In contrast, the second- and third-order effects will be long-lived. In particular, the short-term consequence-management tasks (*e.g.*, treating the wounded, decontaminating the trains and stations) will take weeks to complete, while other tasks such as those that address political, psychological, and sociological effects will continue over many months or years.

Directness: We can tell which cause produced the effect. The proximate cause was different terrorists on different trains releasing sarin.

Invasiveness: The terrorists physically crossed into the U.S. from other countries. The locus of the attack was the Washington, D.C., metropolitan area.

Measurability: We can count the dead and injured, assess the severity of the injuries, and assign a monetary cost to the deaths and injuries. However, there are effects of the attack that are difficult or impossible to measure due to the difficulty of observing these particular effects, such as psychological trauma experienced by the passengers and frequent users of the Metro, family of the passengers, and others. However, we can quantify, in this scenario, the immediate aspects of the attack.

Presumptive legitimacy: No one can launch this type of attack—not even nation states—against noncombatants. Thus, there is no presumptive legitimacy.

Responsibility: During the attack, the actors made no claim of national responsibility (*e.g.*, the terrorists did not wear uniforms). We can classify the responsibility as being "medium" because the actors know that the U.S. Government will find out who was responsible for attack, in addition to discovering links between the actors and terrorist training bases in the offending country.

3.2. Results

Each of the seven factors in the foregoing analysis is graphically reproduced in the addendum. Each diagram contains a brief description of the importance or distinctiveness of the factor, formulation of questions that would satisfy the requirements of the factor, and a vertical scale of the factor itself with one qualitative choice located at the bottom and the other located at the top. Schmitt divided the spectrum into three broad bands, one each for relatively clear cases of each qualitative choice, and a central "gray" area for factually uncertain determinations.

Obviously, a one-to-ten quantitative determination would allow for subdistinctions within each of these bands. More precision than this would be chimerical, in that it would present the appearance of more precision than actually exists. The authors have chosen to present Schmitt's work in this manner to provide clear structure for discussion, but not as an absolute algorithm for producing the "right" answer given any input. The results for Scenario I are given in Table 1.

Table 1. Consequences ascribed to Scenario I

	Numeric rating
Severity	8
Immediacy	8
Directness	8
Invasiveness	9
Measurability	8
Presumptive legitimacy	8
Responsibility	5
Total	54
Simple average	7.7

4. Scenario II: A Cyber Attack

Our second scenario, as in the first, involves an attack on the Washington Metro at rush hour. In this scenario, however, the terrorists use malicious code to strike the software-intensive automatic train protection (ATP) system of the Metro. The attack was orchestrated from outside the U.S. by using compromised administrative computers that are used by Metro officials to monitor operations.

4.1. Analysis

Severity: The changes made to the ATP system permitted the train control system to allow two head-on and three rear-end collisions to occur among trains, before the legitimate operators of the system could regain control by halting train traffic system-wide, evacuating the stations, and redirecting passenger traffic from other modes of transportation (*e.g.*, transit buses). As a result of the crash, thirty passengers were killed and approximately two hundred passengers were physically injured. An undetermined number of people experienced psychological effects—even people who were not proximate to the crash scenes. Property damage was extensive. There was also a loss of intangible property: it took considerable resources to track down and remove vulnerabilities within the software systems that were exploited by the terrorists, and to repair the integrity of the software (*i.e.*, to remove the modifications that had been made by the terrorists to the software).

Immediacy: The attack was executed in under two minutes. The effects of the attacks are tiered: there were instantaneous effects—the crashes themselves; the system was shut down after ten minutes; and many people avoided using the

Metro and other subway systems throughout the nation. It took days to clear the debris and repair the physical damage to the system, along with removing the vulnerabilities in the software comprising the command and control system. Shortly thereafter, regulatory and law enforcement officials gave permission to the transit operator to restart limited passenger service. However, as noted above, passenger confidence was not restored immediately.

Directness: The attackers used code to cause the disturbance in the Metro system. This attack represents a specific break-in: one act had one effect.

Invasiveness: The locus of the act was solely in the U.S.

Measurability: The effect of the attack can be quantified to some extent, as noted in our analysis of severity. However, the nonphysical effects, such as the loss of public confidence in the safety of the Metro system, are difficult to quantify.

Presumptive legitimacy: As in the first scenario, no one can launch this type of attack—not even nation states—against noncombatants. Thus, there is no presumptive legitimacy.

Responsibility: No countries claimed responsibility for this attack. However, although no one saw the terrorists modify the software, we can apply the legal principle of *res ipsa loquitur* to assume that the injury to the passengers was caused by the negligent action of another party because the train collisions are of the sort that would not occur unless some party acted in a negligent manner.

4.2. Results

The attack represents an “8” in terms of severity relative to the September 11, 2001, attack on the World Trade Center. The attack is extreme in both aspects of invasiveness, but lower for the intangible aspects and distance from the target, so we rated invasiveness as a “5.”

Table 2. Consequences ascribed to Scenario II

	Numeric rating
Severity	8
Immediacy	9
Directness	9
Invasiveness	5
Measurability	9
Presumptive legitimacy	5
Responsibility	5
Total	50
Simple average	7.1

5. Discussion

The results of both analyses place the consequences of the attacks in the low end of the high range on the Schmitt scale. The kinetic attack has military aspects to it, such as the crossing of national borders. Although the second scenario does not share the military aspects of the first, the immediacy of the attack in the second scenario is much greater than that of the attack in the first scenario. In both scenarios, however, we can say that armed attacks occurred. More specifically, we can reason that both the U.N. Article 2(4) (*i.e.*, use of armed force) and U.N. Article 51 (*i.e.*, immediacy of threat) thresholds were crossed. These circumstances portend a movement toward conflict between the aggressor and the U.S.

6. Conclusion

We have demonstrated how the Schmitt Analysis can be used to perform a more academically rigorous evaluation of the factors affecting a lawful response to a terrorist attack. Schmitt himself never intended his experiment to provide a mechanical algorithm, for solving what are some of the most technically and legally challenging questions we face; instead, he sees it to be a useful framework for analyzing the effect of key factors on the legal nature of an attack and the appropriate response. As such, it provides an invaluable tool for clarifying thought and highlighting areas of misunderstanding or disagreement. Further, it is an excellent basis for training lawyers, technologists, and decision makers in government. Finally, Schmitt’s methodology shows the way for parallel efforts to make more rigorous and more transparent legal analyses in neighboring areas.

The authors of this article have been working from two different but complimentary areas toward the joint effort represented here—lawful use of software-based deception[6] [4], and cyber law [9]. To this end, we are further refining the scenarios to use them, with the aid of some automated tools, to teach officials from the law enforcement, intelligence, and military communities how to reason about the legality of responses to terrorist attacks. With appropriate training, information, and analysis (both automated and with man-in-the-loop), it will be possible to reduce the “gray area” of legal uncertainty to an absolute minimum, and allow the most complete range of effective responses against those who attack a nation’s critical infrastructure.

Acknowledgements

Conducted under the auspices of the Naval Postgraduate School’s Homeland Security Leadership Development Program, this research is supported by the U.S. Department of Homeland Security. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the U.S. Government. The U.S.

[6] This is one of the many types of weapons that one might expect will be included in the U.S. arsenal for conducting cyber warfare. The U.S. Government is currently developing doctrine and policy for using such weapons [1].

To appear in *Proc. Twenty-seventh Annual Int. Computer Software and Applications Conf.*, IEEE (Dallas, Tex., Nov. 2003).

Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright annotations thereon.

References

[1] "Bush orders guidelines for cyber-warfare." *Washington Post*, 7 Feb. 2003, p. A01.

[2] "Cyber attacks by al Qaeda feared." *Washington Post*, 27 June 2002, p. A01.

[3] Greenway, J. P. and Sheldon, R. H. Automatic train control and communications for Washington Metro. *Communications Society: A Digest of News and Events of Interest to Communications Engineers* 12, 6 (Nov. 1974), pp. 14-21.

[4] Michael, J. B. and Wingfield, T. C. Lawful cyber decoy policy. In Gritzalis, D., Vimercati, S. C., Samarati, P., and Sokratis, K., eds., *Security and Privacy in the Age of Uncertainty*. Boston, Mass.: Kluwer Academic Publ., 2003, pp. 483-488.

[5] Pangi, R. Consequence management in the 1995 sarin attacks on the Japanese subway system. *Studies in Conflict and Terrorism* 25, 6 (2002), pp. 421-448.

[6] Schmitt, M. N. *Bellum Americanum*: The US view of Twenty-first Century war and its possible implications for the law of armed conflict. *Mich. J. Int. Law* 19, 4 (1998), pp. 1051-1090.

[7] U.S. Congress. Office of Technology Assessment. Automatic Train Control in Rail Rapid Transit. Washington, D.C.: Government Printing Office, 1976.

[8] U.S. President. Critical Infrastructure Protection Board. The National Strategy to Secure Cyberspace. Washington, D.C.: Government Printing Office, Feb. 2003.

[9] Wingfield, T. C. *The Law of Information Conflict: National Security Law in Cyberspace*. Falls Church, Va.: Aegis Research Corp., 2000.

Addendum

Severity

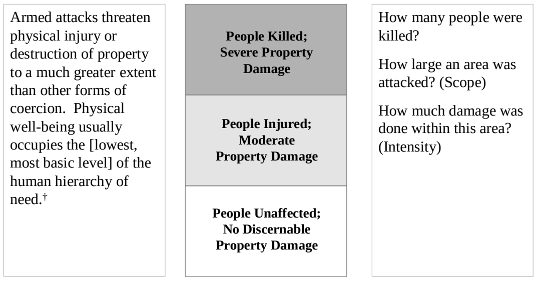

Armed attacks threaten physical injury or destruction of property to a much greater extent than other forms of coercion. Physical well-being usually occupies the [lowest, most basic level] of the human hierarchy of need.†

People Killed; Severe Property Damage

People Injured; Moderate Property Damage

People Unaffected; No Discernable Property Damage

How many people were killed?

How large an area was attacked? (Scope)

How much damage was done within this area? (Intensity)

† Michael N. Schmitt, Computer Network Attack and the Use of Force in International Law: Thoughts on a Normative Framework, 37 COLUM. J. TRANSNAT'L L. 887 (1999) at 914-15.

Immediacy

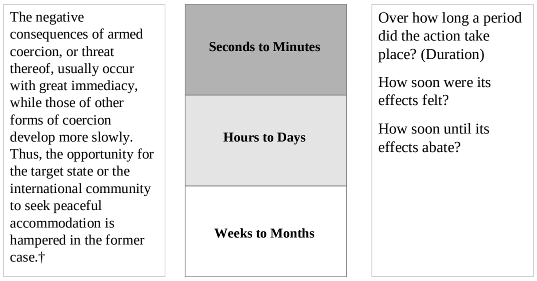

The negative consequences of armed coercion, or threat thereof, usually occur with great immediacy, while those of other forms of coercion develop more slowly. Thus, the opportunity for the target state or the international community to seek peaceful accommodation is hampered in the former case.†

Seconds to Minutes

Hours to Days

Weeks to Months

Over how long a period did the action take place? (Duration)

How soon were its effects felt?

How soon until its effects abate?

† Michael N. Schmitt, Computer Network Attack and the Use of Force in International Law: Thoughts on a Normative Framework, 37 COLUM. J. TRANSNAT'L L. 887 (1999) at 914-15.

Directness

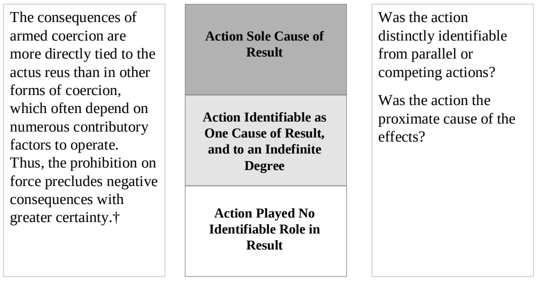

The consequences of armed coercion are more directly tied to the actus reus than in other forms of coercion, which often depend on numerous contributory factors to operate. Thus, the prohibition on force precludes negative consequences with greater certainty.†

Action Sole Cause of Result

Action Identifiable as One Cause of Result, and to an Indefinite Degree

Action Played No Identifiable Role in Result

Was the action distinctly identifiable from parallel or competing actions?

Was the action the proximate cause of the effects?

† Michael N. Schmitt, Computer Network Attack and the Use of Force in International Law: Thoughts on a Normative Framework, 37 COLUM. J. TRANSNAT'L L. 887 (1999) at 914-15.

Invasiveness

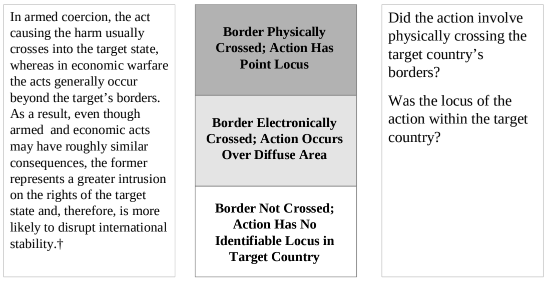

In armed coercion, the act causing the harm usually crosses into the target state, whereas in economic warfare the acts generally occur beyond the target's borders. As a result, even though armed and economic acts may have roughly similar consequences, the former represents a greater intrusion on the rights of the target state and, therefore, is more likely to disrupt international stability.†

Border Physically Crossed; Action Has Point Locus

Border Electronically Crossed; Action Occurs Over Diffuse Area

Border Not Crossed; Action Has No Identifiable Locus in Target Country

Did the action involve physically crossing the target country's borders?

Was the locus of the action within the target country?

† Michael N. Schmitt, Computer Network Attack and the Use of Force in International Law: Thoughts on a Normative Framework, 37 COLUM. J. TRANSNAT'L L. 887 (1999) at 914-15.

To appear in *Proc. Twenty-seventh Annual Int. Computer Software and Applications Conf.*, IEEE (Dallas, Tex., Nov. 2003).

Measurability

While the consequences of armed coercion are usually easy to ascertain (e.g., a certain level of destruction), the actual negative consequences of other forms of coercion are harder to measure. This fact renders the appropriateness of community condemnation, and the degree of vehemence contained therein, less suspect in the case of armed force.†

Effects Can Be Quantified Immediately by Transitional Means (BDA, etc.) with High Degree of Certainty

Effects Can Be Estimated by Rough Order of Magnitude with Moderate Certainty

Effects Cannot be Separated from Those of Other Actions; Overall Certainty is Low

How can the effects of the action be quantified?

Are the effects of the action distinct from the results of parallel or competing actions?

What was the level of certainty?

† Michael N. Schmitt, Computer Network Attack and the Use of Force in International Law: Thoughts on a Normative Framework, 37 COLUM. J. TRANSNAT'L L. 887 (1999) at 914-15.

Presumptive Legitimacy

In most cases, whether under domestic or international law, the application of violence is deemed illegitimate absent some specific exception such as self-defense. The cognitive approach is prohibitory. By contrast, most other forms of coercion—again in the domestic and international sphere—are presumptively lawful, absent a prohibition to the contrary. The cognitive approach is permissive. Thus, the consequences of armed coercion are presumptively impermissible, whereas those of other coercive acts are not (as a very generalized rule).†

Action Accomplished by Means of Kinetic Attack

Action Accomplished in Cyberspace but Manifested by a "Smoking Hole" in Physical Space

Action Accomplished in Cyberspace and Effects Not Apparent in Physical World

Has this type of action achieved a customary acceptance within the international community?

Is the means qualitatively similar to others presumed legitimate under international law?

† Michael N. Schmitt, Computer Network Attack and the Use of Force in International Law: Thoughts on a Normative Framework, 37 COLUM. J. TRANSNAT'L L. 887 (1999) at 914-15.

Responsibility

Armed coercion is the exclusive province of states; only they may generally engage in uses of force across borders, and in most cases only they have the ability to do so with any meaningful impact. By contrast, non-governmental entities are often capable of engaging in other forms of coercion (propaganda, boycotts, etc.). Therefore with armed coercion the likelihood of blurring the relative responsibility of the State, a traditional object of international prescription, and private entities, usually only the object of international administration, narrows. In sum, the consequences of armed coercion are more susceptible to being charged to the State actor than in the case of other forms of coercion.†

Responsibility for Action Acknowledged by Acting State; Degree of Involvement Large

Target State Government Aware of Acting State's Responsibility; Public Role Unacknowledged; Degree of Involvement Low

Action Unattributable to Acting State; Degree of Involvement Low

Is the action directly or indirectly attributable to the acting state?

But for the acting state's sake, would the action have occurred?

† Michael N. Schmitt, Computer Network Attack and the Use of Force in International Law: Thoughts on a Normative Framework, 37 COLUM. J. TRANSNAT'L L. 887 (1999) at 914-15.

Overall Analysis

Use of Force Under Article 2(4)

Arguably Use of Force or Not

Not a Use of Force Under Article 2(4)

Have enough of the qualities of a use of force been identified to characterize the information operation as a use of force?

† Michael N. Schmitt, Computer Network Attack and the Use of Force in International Law: Thoughts on a Normative Framework, 37 COLUM. J. TRANSNAT'L L. 887 (1999) at 914-15.

Calibrating the Factors: Primary Schmitt Analysis

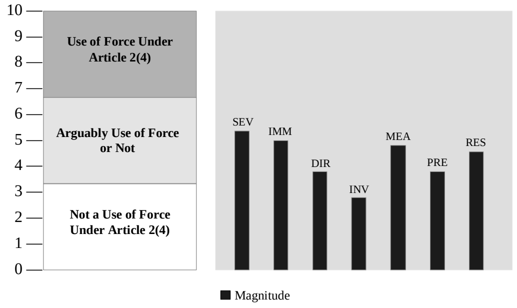

Weighting the Factors: Secondary Schmitt Analysis

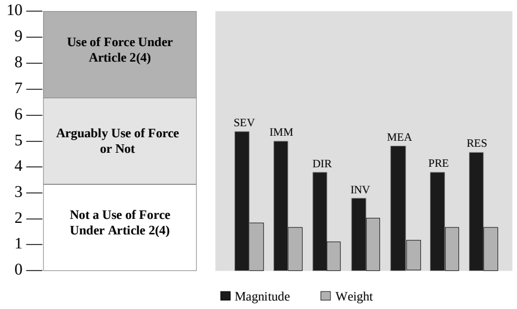

Averaging the Weighted Factors

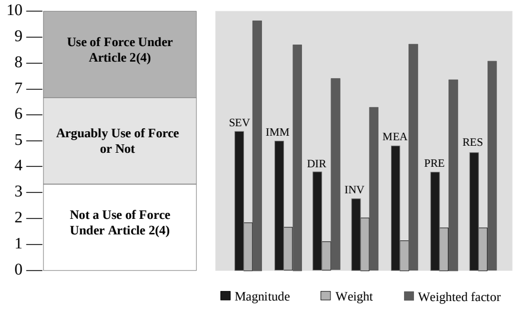
Use of Force Under Article 2(4)
Arguably Use of Force or Not
Not a Use of Force Under Article 2(4)
SEV
IMM
DIR
INV
MEA
PRE
RES
Magnitude
Weight
Weighted factor