Modifying Lanchester's Equations for Modeling and Evaluating Malicious Domain Name Take-down

Jonathan M. Spring
CERT Network Situational Awareness Group
netsa-contact@cert.org
Publication NetSA-2011-25

September 2011

Executive Summary

Domain names drive the ubiquitous use of the Internet. Criminals and adversaries also use domain names for their enterprise. Defenders compete to remove or block such malicious domains. This paper models this competition on large, decentralized networks using a modification of Lanchester's equations for combat. The model is applied to what is known of the current state of malicious domain activity on the Internet. The approach demonstrates limitations based on the general dynamics of the model.

When taken with the economic and physical laws to which the Internet is bound, the model demonstrates that the current approach to removing malicious domain names is unsustainable and destined for obsolescence. However, there are technical, policy, and legal modifications to the current approach that would be effective, such as preemptively populating watch lists, limits on a registrant's registrations, and international cooperation. The results indicate that the defenders should not expect to eliminate or significantly reduce malicious domain name usage without employing new digital tactics and deploying new rules in the physical world.

Software Engineering Institute | Carnegie Mellon

Contents

1   Introduction

Malicious domain names cause significant trouble on the Internet, and defenders can and should resist their damage. In deciding the best course of action in eliminating malicious domain names from the Internet, a model of the potential success of various approaches would be a powerful tool. Lanchester's equations model military combat, but can be modified for this purpose. Lanchester's equations are themselves a modification of the Lotka-Volterra equations, which model predator-prey interaction. Lanchester's equations have been critical to the modeling of warfare since their introduction in 1916 (Fowler, 2006).

The idealized conception of the equations has many assumptions and they have been modified many times using various assumptions to better accommodate various types of warfare, such as by Dolanský (1964), Lauren (2001), and Helmbold (1965). Following Lauren (2001), the basic combat interaction between a red force R and a blue force B over time, the Lanchester equations are:

$$\begin{aligned} dr/dt &= -K_b b(t); \; r(0) = R_0 \\ db/dt &= -K_r r(t); \; b(0) = B_0 \end{aligned} \tag{1}$$

Here, K_b and K_r are non-negative scalars representing the effectiveness of the two forces against each other. The number of active blue soldiers (or airplanes, tanks, etc.), b, in the battle will decrease proportional to the number of soldiers with which red, r, is opposing blue. The number of units changes as a function of time over the duration of the battle. The larger initial force will win, if K_b and K_r are equal. However, if R_0 were twice B_0, blue could try to compensate by being more effective at destroying red, i.e. increase K_b. But to overcome this 2:1 disadvantage in numbers, K_b would have to be 4 times more effective than K_r. That is, the basic Lanchester equations are second order.

In a military conflict in which both unit types have the same destructive effectiveness, the side with a numerical advantage will have significantly fewer total casualties by the end of the conflict. The effect is demonstrated in the following contrived examples. The Spartans, with superior numbers, achieve a lopsided victory just by maintaining parity with the destructive effectiveness of the smaller force (Figure 1). In order to overcome the Spartans superior numbers, the Athenians must be 4 times more effective to battle them to a draw (Figure 2).

Figure 1: Contrived example in which the Spartans have double the initial force and same destructiveness.
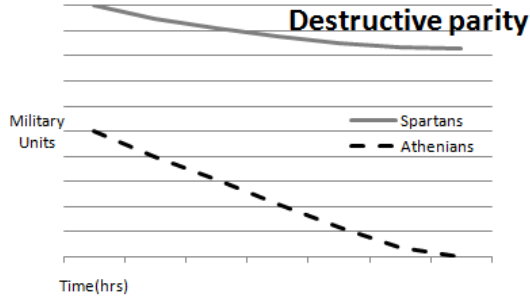

Figure 2: Contrived example in which the Spartans have double initial force but the Athenians have 4-times their destructive effectiveness.

There are many modifications of the Lanchester models. For example, the equations can be modified to take in to account heterogeneous forces (riflemen, infantry, tanks, etc. on each side) each with a different success rate against each type of enemy unit. This can be accommodated by constructing matrices analogous to (1) (Dolanskỳ, 1964).

Lanchester's basic equations are derived from the Lotka-Volterra model of predator and prey. Lanchester removed elements that accounted for the birth of prey and the natural death of predators. This is sensible, as the time scale of battles does not allow for the production of new soldiers during the battle. The basic predator-prey equations incorporate this concept as follows, where y is the number of predators and x is the number of prey:

$$\frac{dx}{dt} = x\,(A - By) \tag{2}$$
$$\frac{dy}{dt} = -y(C - Dx)$$

Here, the terms $x(t)$ and $y(t)$ are abbreviated as x and y, respectively. This

convention will continue for future variables which are a function of time. The symbols A, B, C, and D are non-negative scalars. The terms A and C account for the birthrate of prey and the natural death rate of predators, respectively. The terms B and D account for the rate at which the predators kill prey and use that energy to reproduce, respectively. Birthrates make (2) conceptually different from (1), as does the concept of consumption for reproduction.

Yet like (1), the Lotke-Volterra equations model entities that both destroy and are destroyed. This is not true of domain names. Therefore equations to model domain dynamics modify both the Lanchester and Lotke-Volterra equations significantly, while still taking their inspiration from the earlier equations.

2   Modeling Domain Name Take-Down

The equations introduced can be modified to suitably describe actors competing in the digital world, given the appropriate assumptions are extracted from (1) and (2). For example, the assumption that every ranged unit can target every enemy ranged unit is not physically feasible for a large force. On the Internet, however, this is reasonable. In usual day to day operations, every part of the Internet is supposed to be reachable by every other part of the Internet (Hall et al., 2011, p. 22). In this regard, the basic Lanchester equation format should apply to competition interactions that occur over the Internet.

The competition of domain name take-down is quite different from that of armed combat. Notably, the entities being destroyed are not the entities doing the destruction. This exception essentially violates the assumptions that entities in x and y are commensurate. Additionally, the competition is inherently one-sided. This is more similar to a predator-prey relationship than an armed conflict. The malicious domains do not seek to take down benign domains, but to steal information (intellectual property, personal information, credit cards) and resources (money, CPU time for botnets) from the target. For this reason, malicious domain activity or numbers do not directly affect malicious actors' opponents, even though a competition for resources may have indirect effects.

Like prey, which are born and die during the scope of the competition described by the above equations, domain names are definitely registered and taken down during the competition. On the other hand, digital competition contains asymmetries unlike either war or predation. These asymmetries complicate the connected equations necessary to describe the competition. The following variables are proposed to represent the salient aspects of the competition interaction among domain names, regardless of their particular malicious use.

Variables that are a function of time:

x_δ :=number of active malicious domains

x_ν :=number of malicious domains newly registered during the interval that have not been activated

x_m :=resources (either in time, person-hours, or money) malicious actors have available for registering and maintaining domain names

x_c :=resources lost by malicious actors due to non-technical socio-political and criminal penalties.

y_m :=resources the community or organization makes available to defensive actions, such as taking down or blocking malicious domains.

y_r :=resources (time, person-hours, money, intelligence, intellectual property, etc.) lost by the community or organization as a result of fraud, etc., enabled by malicious domains.

Scalars that represent effectiveness or efficiency:

$D_\sim :=$ various; convert units of variable to domain units; must be ≥ 0.

$C_\sim :=$ various; converts units of variable to monetary units; must be ≥ 0.

Values which are modeled as constants:

$N :=$ rate at which new domains are registered by malicious actors

$E :=$ engineering and development costs.

With these parameters, we can propose several equations following the spirit of the Lanchester and Lotka-Volterra algorithms. Variables in x represent aspects of the attacker, while variables in y represent aspects of the defender.

brief variables:
x_δ: active mal. domains
x_ν: registered but inactive
x_m: mal. actor resources
x_c: crim. penalty
y_m: defensive resources
y_r: lost resources
E: dev costs
N: new registrations
$D_\square$: scalar to domains
$C_\square$: scalar to money

$$\frac{dx_\delta}{dt} = D_{x_m}x_m + (D_{x_\nu}x_\nu) - (D_{y_m}y_m) \qquad (3)$$

$$\frac{dx_\nu}{dt} = -D_{x_\nu}x_\nu - (D_{m2}y_m) + N \qquad (4)$$

$$\frac{dx_m}{dt} = C_{y2}y_r - (C_{\delta2}x_\delta) - (C_{x_\nu}x_\nu) - (C_{x_c}x_c) - E \qquad (5)$$

$$\frac{dy_m}{dt} \approx -C_{y1}y_r \qquad (6)$$

$$\frac{dy_r}{dt} = C_{\delta1}x_\delta \qquad (7)$$

Equation (3) models the rate of change of the population of active malicious domains. The scalar D_{y_m} represents the effectiveness of take-down measures per unit of resources devoted. It is estimable by observation, in principle. Block listing has been observed to be reasonably effective (Moore and Clayton, 2007). The scalar D_{x_m} represents the effectiveness of efforts to maintain active malicious names and their infrastructure. Since not all newly-registered domains are activated right away, some percentage of the registered domains will be activated

over time, which is represented by D_{x_ν}. Measurements would be necessary to determine this percentage and if it is constant. Studies to this effect are not known, so for the time being it will be approximated as such.

To balance the equations, they must use the same units. The units of x_m and y_m (resources) are not commensurable with those of x_δ and x_ν (domains). However, the important aspect of the equation is that efforts to take down the malicious domains are offset by both efforts to keep them alive and the number of new domains that are activated. The units of D_{y_m}, D_{x_m}, etc., could simply be such that they convert units appropriately. This would not change the general dynamics of the equation. Similar conversions will be assumed for all the equations.

Equation (4) models the change in the number of registered but inactive domains. Since x_ν represents domains that have been registered, but not used, it is decremented by the number of domains that are activated in (3). For simplicity, domains are modeled to be registered at a rate independent of the other variables by the scalar N. Community take-down efforts could also reduce these domains, but with a different effectiveness coefficient than these efforts effect live malicious domains in (3).

Equation (5) models the resources available to the malicious actors. The scalar C_{y_r} describes a percentage of those resources stolen that can be incorporated into the malicious actors' resources. Scalars $C_{\delta 2}$ and C_{x_ν} essentially represent the cost of maintaining and registering a domain name, respectively. The scalar E is a cost independent of the number of domain names active. It represents various engineering, setup, human, and organizational costs.

The resources available to the defender must constrain the defensive resources allocated, y_m. Resources allocated to defense is related to resources lost because it is natural to devote more resources to a bigger problem. However, in practice there are many social, political, and economic factors that alter what resources are allocated for network defense, and such non-technical features are not modeled in (6). The lack of non-technical aspects would be most important to the model in (6), so here this modeling choice is most acutely felt.

In principle, the defender's losses may be reduced by legal action or insurance payments, however these recuperations will not, for the community as a whole, exceed the costs of providing them. For example, any insurance provider will have to charge more for premiums than they give out, or else that provider will become insolvent. Given losses of this nature it is sensible to assume that y_r is monotonically increasing, as in (7).

The change is positive in (7) because it represents increasing losses, rather than decreasing resources. The variable y_r is also presumed to have no limit in (7). So far, fraud losses have not been so great as to overwhelm the economy or resources of whole communities or organizations, but there is certainly some threshold that y_r could reach for which the defrauded entity would cease to be able to function. Such catastrophes are not considered in this model.

The starting resources available to malicious actors are non-zero, since there are certainly some initial resources rolled over from previous crime, digital or not. This starting funding is the value of x_m at $t = 0$. Evaluating the values of

the initial conditions is difficult, but for each variable it is greater than zero.

In theory, there are five costs to model in (5). In reality, for international cybercrime at least, C_{x_c} is essentially zero.[1] This is because there are few effective penalties. Therefore, the current realistic model for (5) is:

$$\frac{dx_m}{dt} = C_{y2}y_r - (C_{\delta2}x_\delta) - (C_{x_\nu}x_\nu) - E \tag{8}$$

From these differential equations relationships between the resources expended by malicious actors and the community can be derived. Malicious actors will gain resources and capabilities, i.e. $\frac{dx_m}{dt}$ will be greater than zero, as long as the following inequality holds:

$$C_{y2}y_r > C_{\delta2}x_\delta + C_{x_\nu}x_\nu + E \tag{9}$$

That is, if their profits exceed their expenditures. Furthermore, if the number of active malicious domains is to decrease, i.e. $\frac{dx_\delta}{dt}$ in (3) is negative, then the following must hold:

$$D_{y_m}y_m > D_{x_m}x_m + D_{x_\nu}x_\nu \tag{10}$$

The costs represented by x_m and C_{x_ν} are essentially the variable costs of a domain name to the malicious actor. Variable costs are opposed to fixed costs, which are initial investments. Total costs are variable costs plus fixed costs. Marginal cost is the change in total cost per one more unit of output, i.e. one more domain (Henderson and Quandt, 1980, p 84ff). In an unconstrained digital economy, reproduction is reduced to copying patterns of bits, which has a variable cost of essentially zero. As more units are produced, initial fixed costs are averaged out over more units produced and so marginal cost is negative and production costs asymptote towards the cost of a new unit. Therefore, unopposed, domains will approach being free to the malicious actor.

3   Current Dynamics

There are simple observations about the current dynamics of the digital environment that can further inform the models. Malicious actors can and do automate fixed costs, externalize fixed costs, reduce variable costs, and utilize existent infrastructure.

Both take down resistance and activation of domains are actions that can be automated. Automation increases D_{x_m} and D_{x_ν}, which is not beneficial for the

brief variables:
x_δ: active mal. domains
x_ν: registered but inactive
x_m: mal. actor resources
x_c: crim. penalty
y_m: defensive resources
y_r: lost resources
E: dev costs
N: new registrations
$D_\square$: scalar to domains
$C_\square$: scalar to money

---

[1]There are no international treaties to account for international cybercrime. The few bilateral treaties that exist are avoided by the criminals. INTERPOL cannot press charges, and so the lack of international agreement on what constitutes a crime renders the organization ineffective in this arena. The International Criminal Court (ICC) has not been approved by sufficient nations to be considered effective, especially lacking the support of the United States. The main purpose of the ICC is also war crimes, not cyber crimes, and so would require a significant increase in scope before it would be helpful to this particular problem. As such, the term for criminal penalties and costs is effectively zero. Implementing effective international criminal penalties is a necessary long-term solution.

defender, considering (10). Automation converts these activities into variable costs, rather than fixed costs. The engineering to automate the operations is the fixed cost. Attackers reduce E in (5) in this way, thereby increasing their profits. Digital, automated costs to the malicious actor will approach the cost of copying bits, i.e. zero, unless non-digital costs are imposed by the defenders. In a purely digital competition, the number of domains available to the malicious actor should always exceed the defender's ability to take them down; the variable cost of detecting domains is not zero. Increasing digital costs is an important and necessary defense, but it is not a sufficient defense. Some of the larger respites from malicious activity have been not due to defensive action, but due to the decisions of the attackers. The group of malicious actors known as Avalanche abandoned their phishing exploits on their own, although presumably to pursue other endeavors (Rasmussen and Aaron, 2010).

Malicious actors' activities can be modeled as having a negative marginal cost and low variable cost due to the dynamics of digital economics. Additionally, many initial costs are born by other organizations, such as registrars, who sell domain name services. Initial costs may also be borne by previous criminals, who have gone through the trouble to establish botnets usable as name servers or other services. This infrastructure further reduces set up costs, i.e. E in (5).

One might expect there to be switching costs involved in new domain names, i.e. that C_{x_ν} would incorporate some component of a switching cost for each new domain. However this is not clear. Since the DNS is one ubiquitous protocol designed to minimize switching costs, the costs are low as long as the malicious code is capable of asking for the correct names. This engineering is non-trivial, but not actually related to the domain names themselves. Further, malicious code has demonstrated the ability to incorporate both updates and outside data (Rodionov and Matrosov, 2011), and so the process of what names to look up is also automatable, thus providing an avenue to significantly reduce cost.

In practice, the cost of domain names to malicious actors is nearly zero. Dozens of dynamic DNS services provide free domain name registration. Furthermore, many registrars and registries permit "domain tasting," in which a registrant is permitted to use a domain for a few days to get a sense of the traffic available to it before paying. Even though ICANN successfully implemented policies to eliminate this practice within generic TLDs, the country code TLDs aren't bound by the same policies (ICANN, 2009). The current useful lifetime of a malicious domain is already below a few days (Cyveillance, 2008; Rasmussen and Aaron, 2011). So the attack patterns are already adapted to making use of a domain well within the time frame afforded by this free domain tasting.

If the malicious actor would actually purchase the domain, the cost of the domains could be charged to fraudulent or stolen credit cards, perhaps those obtained by previous attacks. Even if stolen credit cards are purchased the cost is minimal. The market for credit card credentials is flooded — price is based on the availability of processing time, not on the supply of stolen credentials (Hackworth, 2011). It is reasonable to imagine that some registrars are established by malicious actors for easier dealing in such stolen credentials. Such malicious establishments have been repeatedly observed for other functions

on the Internet, such as the infamous Russian Business Network and many others (Krebs, 2007; Wikipedia, 2011). This behavior also follows a digital information economy. However the fixed costs of establishing such a business are non-digital, and therefore a potentially useful target for defensive actions.

Given this rationale, it is reasonable to estimate that $C_{\delta 2}$ and C_{x_ν}, in (5), (8) and (9), that the cost of maintaining a DNS structure and registering new domain names each are near zero and getting nearer. The cost of engineering these solutions, E, is also slowly approaching zero as code is reused and existing infrastructure is leveraged. In this purely digital competition, the terms x_δ and x_ν effectively drop out of (8). It then simplifies to a depressing expression about the profits of the malicious actors:

$$\frac{dx_m}{dt} \approx C_{y2} y_r \tag{11}$$

Equation (11) states that the resources of the malicious actors will only increase in a purely digital competition with the defenders. Additionally, as long as malicious actors control malicious domains y_r should increase as per (9), and so the malicious actors' profits will increase ever faster.

4 Ramifications

These models make certain statements. From (9) one can surmise that a defensive tactic is to make x_ν have to be very high to increase the cost to the malicious actor. A defender could do so by taking down many domains, forcing many new ones to be registered. The take-down rate is (hopefully) increased by community expenditures, i.e. y_m, as scaled by D_{y_m} in (3). Equation (10) permits a simple evaluation of the success of community expenditures based on whether they take down more domains than the malicious actors can maintain and activate.

Equation (10) also demonstrates that community take-down efforts could be resisted if adversaries create domain structures more resilient to take-down or register and activate many new domains. In practice, both tactics are used.[2] In light of the dynamics of the Lanchester equations, the defender would seem to have an advantage. If defender resources, y, are much larger than adversary resources, x, then x would have to compensate by a geometric advantage in D_{x_m} (3) and C_{x_ν} (5). Therefore defense should be tenable for the defender if D_{x_m} and C_{x_ν} are relatively close in value to D_{y_m} (3). In the physical world it is nearly assured that large, technologically-advanced forces will not be grossly outgunned by ragtag criminals. Yet, the economics of digital information change the landscape significantly.

---

[2]For example, the Conficker C virus registered 50,000 new random domains per day, and each infected host would attempt to contact a pseudo-random selection of 500 of these. This randomness is time and effort in algorithm design, i.e. increased x_m, and the high domain volume is a large x_ν.

It is clear from equations (10) and (11) and the realities in Section 3 that an approach which attempts to limit criminal activity solely by removing domains used maliciously is ill-fated. Even if the lifetime of a malicious domain were forced towards its cost of production, i.e. zero, any reactive approach cannot actually eliminate the domain before its use. Therefore, reactive block list services alone[3] cannot reduce the domain lifetime to its cost, no matter how efficient or well-conceived. In order to react, the domain must have been used, which means profit can be generated in that one use on something that was free to obtain. To reduce the domain lifetime to less than one use, a domain-name-based block list must be predictive, i.e. remove domains before they are used. This ability to eliminate registered domains before they are activated is modeled in (4) by the term $D_{m2}y_m$. In the current landscape, D_{m2} is equal to zero.[4]

A predictive block list could force the average domain lifetime below one use, since if many maliciously registered domains are blocked before they are used their lifetimes are zero. There is only one example of such a method, which utilizes intelligence from TLD zone files. Even though it explores a useful direction, the false positive rate as proposed is too high to be used (Felegyhazi et al., 2010). Defenders should concentrate efforts on refining such approaches. Reactive blocking fails because the revenue derived from malicious domains, $C_{\delta 1}x_\delta$ (7), will exceed the cost of new domains, $C_{x_\nu}x_\nu$ (5). The value derived from activated domains is small, but it only needs to exceed a minuscule cost. A predictive block list would institute a non-zero D_{m2} scalar in (4), and alter the landscape significantly by preventing some x_ν from becoming x_δ.

Another defensive option would be to implement policies or actions that would increase the switching or registration costs, including time, for domain names. There are currently no realistic barriers to prevent one actor from registering tens of thousands of domain names in a week, and maintaining them on standby indefinitely, to be activated when necessary. This is one reason why $C_{\delta 2}$, the cost of maintaining domain names in (5), is effectively zero.

Investigation of registration activity would have precedent. Banks investigate suspicious withdrawals or deposits of large amounts of cash (Levitt and Dubner, 2010, pp. 88-96). Similarly, registrars and registries would be justified in investigating such anomalous behavior; there is evidence that malicious domains behave differently in the data they handle (Spring et al., 2011). Smart anomaly investigation should be feasible to design and implement. Following the example of banking anomaly detection, it is possible to be fruitful without hampering daily users. Simply capping batch registrations at a low number would be a start. Coordination from registrars and registries would make it more difficult to register a domain (C_{x_ν}), to maintain it (C_{x_δ}), and increase

---

[3] All block lists currently in use (McAfee RBL, Spamhaus, PhishTank, Google safebrowsing, etc.) and receiving academic acclaim (EXPOSURE (Bilge et al., 2011), Notos (Antonakakis et al., 2010), Kopis (Antonakakis et al., 2011)) are reactive.

[4] In the case of Conficker the community made a concerted effort to block domains before they were used. Conficker variant C eventually overwhelmed these efforts. There has been little effort to work on this in the general case.

defenders' ability to take down both active (D_{y_m}) and registered but unused (D_{m2}) names.

Another potential point at which to apply pressure would be on registrars. Certainly, some registrars do a better or worse job at preventing abuse than others. If such statistics were available to the community, if not publicly, pressure could be applied to those registrars to improve. Offending registrars should eventually lose their authority to register names. Without any such censure process in place, there is no way to prevent rogue registrars from aiding and abetting the criminals. Essentially, someone needs to audit and verify the necessary registrar anti-abuse measures.

Other brick-and-mortar institutions that serve to abet cyber criminals should also be sought out and censured or terminated. The criminal process is not an exclusively digital phenomenon, and traditional countermeasures must not be abandoned. For example, services that launder money such as banks would be a valuable target. Such institutions have much higher costs, both to replace, switch, and operate, than simple domain names. Finding, arresting, and incarcerating the criminals would be a deterrent. Unfortunately, this is not optional in the case of cyber crime. The models in Section 2 and the realities in Section 3 demonstrate that cyber crime cannot be effectively combated by digital means alone. Yet due to ineffective international coordination, current approaches to the problem are almost exclusively digital. This approach is not sustainable. Effective political changes need to occur to make criminal penalties for international cyber crime a reality.

5   Future Work

The line of inquiry described above still leaves much to be explored. One important aspect of the model is the values for the variables at the initial conditions. Estimates of some may exist, but there is significant room for improvement. Another possible area for exploration could incorporate the interaction of multiple types of forces within each competing entity, as Dolansky does for Lanchester's original equations (Dolanský, 1964). Incorporating such heterogeneity may prove useful in the discussion of domain name take down and competition as well. Each take-down technique will almost certainly have varying effectiveness on the various deployments of malicious domains. Additionally, these models may generalize to other types of network behavior. That determination will have to be made for several large classes of network behavior independently.

In regards to (7), while money may be recouped by the defrauded in a legal case, there is still increasing cost to the community in general in order to pay for the police, insurance, and legal activities. This paper does not introduce such players into the model, and such inputs are left to future work. Such additions would likely be an extension of introducing heterogeneity into the model. There is also some point at which a critical point is reached and the defender begins to collapse due to resources lost. To model this a bounded formula for the elements of (7) could be implemented.

Equation (6) is recognized to be incomplete. This equation oversimplifies the rate at which the defender's resources change. While defensive resources are certainly diminished by resources lost to the attacker, there are myriad other influences. These include resources allocated due to social pressure, legal requirements, ideological values, and others. It is also probable that a community's sense of urgency could increase defensive resources related to the number of malicious domains (x_δ), which would further complicate the interaction. Modeling such interactions is left for future work.

Equation (4) is also incomplete. The rate at which new domains are made available to the malicious actor is not constant, so the dynamics of N must also be modeled. The process is probably related to the agent's needs, opportunities, available resources, and other chance factors. The scalar E in (5) is similarly oversimplified, and would benefit from a more thorough analysis.

Identifying actual resources available to malicious actors is another realm of study. This empirical work is ongoing in certain niches, such as the spam economy (Kanich et al., 2011). In general, specifying the values of the scalars in the equations and the initial funding and existing malicious domains is future work for empirical study.

Section 4 recommends identifying and removing rogue registrars. Successfully doing so would not end the competition, but would likely push the battle to dynamic DNS services. Such services provide free domain names, name server access, and services for user-defined subdomains under domains that the dynamic DNS service provider controls. Essentially, they serve as an informal registrar, registry, and name server operator. Dynamic DNS services could be considered rogue registrars which would require censure; however they do not operate under the authority of ICANN or another central authority and so would be more difficult to find and censure.

6   Conclusion

Digital countermeasures to malicious domains are still necessary. Their effectiveness has been documented. However, as the competition continues, the malicious actors will continue to adapt around digital countermeasures. The models presented demonstrate that malicious actors should be expected to always be able to adapt around digital countermeasures and still profit. Given the necessary features of a digital economy and reactive blocking, the malicious actors will still have revenues exceeding their costs. These digital methods must be accompanied by physical and policy countermeasures to cyber crime and malicious domain name usage. Malicious domains serve as a means to a human economic end. Criminals will operate in the space where they will not be caught or punished. Without effective penalties equivalent to those for traditional crime, one cannot expect cyber crime to cease of its own accord or by digital countermeasures alone.

References

Antonakakis, M., Perdisci, R., Dagon, D., Lee, W., Feamster, N., 2010. Building a dynamic reputation system for DNS. In: 19th Usenix Security Symposium.

Antonakakis, M., Perdisci, R., Lee, W., Vasiloglou II, N., Dagon, D., 2011. Detecting malware domains at the upper DNS hierarchy. In: 20th Usenix Security Symposium. San Francisco, CA.

Bilge, L., Kirda, E., Kruegel, C., Balduzzi, M., February 2011. EXPOSURE: Finding malicious domains using passive DNS analysis. Proceedings of the Annual Network and Distributed System Security (NDSS).

Cyveillance, 2008. The cost of phishing: Understanding the true cost dynamics behind phishing attacks. Tech. rep., Cyveillance, Inc., Arlington, VA.

Dolanskỳ, L., 1964. Present state of the lanchester theory of combat. Operations Research, 344–358.

Felegyhazi, M., Kreibich, C., Paxson, V., 2010. On the potential of proactive domain blacklisting. In: Proceedings of the 3rd USENIX conference on Large-scale exploits and emergent threats. USENIX Association, pp. 6–6.

Fowler, C. A., March 2006. Asymmetric warfare: A primer. IEEE Spectrum.

Hackworth, A., July 2011. personal communication.

Hall, C., Anderson, R., Clayton, R., Ouzounis, E., Trimintzios, P., 2011. Resilience of the Internet interconnection ecosystem. Tech. rep.

Helmbold, R., 1965. A modification of lanchester's equations. Operations Research, 857–859.

Henderson, J., Quandt, R., 1980. Microeconomic theory: A mathematical approach, 3rd Edition. McGraw-Hill New York.

ICANN, 2009. The end of domain tasting: Status report on AGP measures. Tech. rep., ICANN.

Kanich, C., Weaver, N., McCoy, D., Halvorson, T., Kreibich, C., Levchenko, K., Paxson, V., Voelker, G., Savage, S., 2011. Show me the money: Characterizing spam-advertised revenue. In: 20th USENIX Security Symposium. San Francisco, CA.

Krebs, B., November 7, 2007. Russian business network: Down, but not out. Washington Post.

Lauren, M., 2001. Describing rates of interaction between multiple autonomous entities: An example using combat modelling. Tech. rep., New Zealand Defense Force, Defence Technology Agency.

Levitt, S., Dubner, S., 2010. SuperFreakonomics: Global Cooling, Patriotic Prostitutes, and Why Suicide Bombers Should Buy Life Insurance. Harper Collins.

Moore, T., Clayton, R., 2007. Examining the impact of website take-down on phishing. In: Proceedings of the anti-phishing working groups 2nd annual eCrime researchers summit. ACM, pp. 1–13.

Rasmussen, R., Aaron, G., October 2010. Global phishing survey: trends and domain name use in 1H2010. Tech. rep., Anti-Phishing Working Group.

Rasmussen, R., Aaron, G., March 2011. Global phishing survey: trends and domain name use in 2H2010. Tech. rep., Anti-Phishing Working Group.

Rodionov, E., Matrosov, A., June 2011. The evolution of TDL: Conquering x64, revision 1.1. Tech. rep., ESET, Bratislave, Slovakia.

Spring, J., Metcalf, L., Stoner, E., 2011. Correlating domain registrations and DNS first activity in general and for malware. In: Securing and Trusting Internet Names 2011.

Wikipedia, September 21, 2011. Bulletproof hosting. `http://en.wikipedia.org/wiki/Bulletproof_hosting`.