Nonlinear Resource Allocation in Restoration of Compromised Systems

Qunwei Zheng*, Xiaoyan Hong*, Sibabrata Ray§
* Computer Science Department, University of Alabama, Tuscaloosa, AL 35487
§ Google Inc. 604 Arizona Blvd. Santa Monica, CA 90401

Abstract— **Under security threats, today's networks are being made to be intrusion tolerant. In a large scale, services are continuing (at a degraded level) while compromising and recovering are both progressing. One of the key problems in the restoration procedure regards to the resource allocation strategies, typically a minimized total cost concerning both service loss and resource expense. In this paper, we investigate the achievable minimal total cost and corresponding resource allocation strategy for different situations. The situations include nonlinear relationship between resource allocation rate and the restoration rate, and its variant when time factor is concerned. We present cost models and numerical results. The results show the impact from various system parameters on the critical conditions for a successful system restoration and the minimal cost. An important result of our study suggests that tight operational region exists under certain conditions.**

I. Introduction

Despite the large amount of network security methods deployed over the world, it is highly desired that today's networks, be it Internet, Ad hoc networks, or sensor networks, are intrusion tolerant [3] [8], i.e., the systems are able to continue its operation even under attacks. When this happens, the performance of the system might degrade due to intrusion but the service it provides must never stop. In the meantime, the system's owners will typically use large amount of resources to restore the compromised system. Let us take an Internet worm attack for example. Fig. 1 shows the spread of Internet worm named Code Red in 2001. It infected hundreds of thousands of machines and caused a huge damage [6]. The figure shows that the infection rate was low at the beginning. Then during a certain period the rate became very high. The number of infected machines increased quickly. Eventually, the infection stopped when people learned how to defend the worm, and more and more machines were restored. The worm finally died out and all machines became normal again. Clearly, restoration of the system (here all the machines that could be infected) started at a certain time and the effort leaded to final recovery of the whole system.

One of the key problems in the restoration procedure regards to the resource allocation strategies, typically a minimized total cost concerning both service loss and resource expense. In the area of network security, cost has been used as a factor in analyzing the damage resulting from denial of service attacks [5], in suggesting investment strategies of security solutions [2], and in server replication strategies [7]. But none of these work nor other publications have studied resource allocation problem in restoration procedure and the cost incurred in that procedure.

In this paper, we present cost analysis on resource allocation for the restoration procedure. We focus ourselves on abstract systems instead of studying a particular attacking event like Code Red. We assume a system is comprised of a large number of machines (nodes) which have same restoration properties, e.g., recovery speed and restoration cost. And the owner has only limited resource. We will investigate the achievable minimal total cost and corresponding resource allocation strategy for different situations. These situations include nonlinear relationship between resource allocation and the outcome - restoration rate, and its variant when time factor is concerned. In the analysis, we model the total cost as a sum of service loss and resource expense. In lacking of closed form solutions, we present numerical results showing impact from various system parameters, e.g, the compromise rate, the initial system damage percentage, etc., on the critical conditions for successful system restoration, the time achieving total restoration, the minimum cost, etc.. An important result of our study suggests that tight operational region exists under certain conditions. One unsolved problem, though, is the validation of the model. Currently we are searching for real resource usage data from various sources, e.g., worm spread study. Nevertheless, the analysis and the results presented in the paper sheds a light on optimal usage of resources in combating network security breaches.

The rest of the paper is organized as follows. Section II describes the system model and the cost model. Section III introduces nonlinear functions for the restoration rates in terms of the resource allocation rate, with time-variable and time-invariable features. We then present cost analysis based on the time-variable nonlinear function and time-invariable nonlinear function in Section IV. Numerical results are given in Section V. Section VI concludes the paper.

II. The System Model

When a system is in a restoration procedure, the state of a system is decided by the rates of two opposite parties: the compromise rate v of the attacker and the restoration rate u of the owner. If the compromise rate v of the attacker is larger than the restoration rate u of the owner, more and more nodes in the system will be compromised. On the contrary, if restoration rate u is larger, increasing number of nodes will turn back to normal. The system is illustrated in Fig. 2 as a rectangle, where the shaded area represents the uncompromised part of the system, i.e., the service is still available. Particularly, the left rectangle in the figure is a normal system and the right one is a partially compromised system. In Fig. 2, l is the full service ability of the system, for example, the total number of nodes in the system. c is the portion of the system that has been compromised, e.g., the number of infected nodes. The attacker and the owner compete with each other in order to have full control over the system. The rate of compromise v is decided by the attacker. But the owner can adjust her restoration rate u which in turn is constrained by the amount of resource that she allocates, for example, budget, or manpower. However, if $u < v$, the system will eventually be totally compromised by the attacker.

We identify two types of cost in the restoration procedure. The first type is a service *loss* cost C_1. It comes from the fact that the compromised system can only provide a degraded service to its clients. The second type is a restoration *expense* cost C_2. It comes from the fact that the system owner must utilize all possible resources she possesses to put the system back to work. These include manpower, investment on new security software and hardware, the cost of recourse to security companies, etc. The total cost C is the sum of the loss and the expense: $C = C_1 + C_2$.

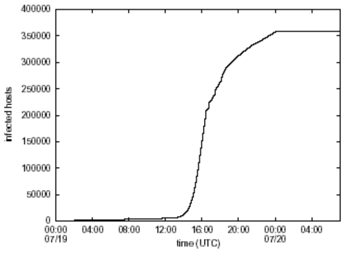

Fig. 1. Cumulative total machines infected by Code Red worm (from www.caida.org)

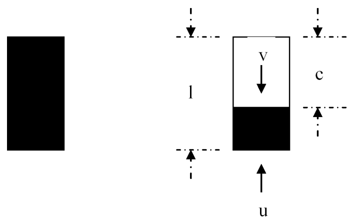

Fig. 2. Normal (left) and partially compromised system (right)

The total cost is affected by the two closely related components. If the system owner consumes more resources, that is, increases the expense C_2, the system will be restored faster and the loss C_1 will be smaller. If, on the contrary, the system owner consumes less resources, that is, the expense C_2 is small, the system will be restored slower and the loss C_1 will be large. Complexities come from many aspects: compromise rate v could be constant or time-varying, the usage of resources is not necessarily contributing to the restoration linearly, or the rate of restoration u could also be constant or time-varying.

In this research, we will investigate the achievable minimal cost for the situations that the usage of resources contributes to the restoration not linearly. In addition, the nonlinear relationship could be constant or variable over the time. For simplicity, we assume the compromise rate v is constant over time. This implies that the restoration is quick in the sense that it recovers the whole system before it is fully compromised. Or, it implies that the system is large so that the attacker could not attack all nodes before the total recovery. In our analysis, we will use a continuous model to approximate the discrete number of computers compromised, given that the accuracy is achievable when the system contains a large number of computers. Such assumption is common in worm spread research, for example, the modelling of Code Red worm [9]. We also assume that resource can be spent at arbitrary granularity and thus is regarded as a continuous variable. A good approximation is achievable if the amount of resource is large. Otherwise, once the optimal allocation is found in continuous case, a brute force search in the neighborhood will give optimal or near optimal integer solution.

III. Resource Allocation

In this section we define the nonlinear relationship between the resource usage and the restoration rate. Let x be the rate of resource spent for restoration, the restoration rate u is thus a function of x, denoted as $u(x)$. The relationship should satisfy *the law of diminishing marginal utility* [1]. In addition, the relationship could be constant or variable over the time. We present functions capturing these features in this section. Analysis on minimal cost will be presented in the next section (Section IV) for both time-invariable and time-variable cases.

A. Necessary Conditions

Spending resources has several limitations, so to the restoration. There are several conditions that the relationship of the two, $u(x)$ and x, should satisfy. First, the restoration rate u should be 0 when no resource is allocated. That is $u(0) = 0$. Second, the relation between the restoration rate and the amount of resources allocated should obey a general economics rule called *the law of diminishing marginal utility* [1]. The law of diminishing marginal utility states that as the individual's consumption increases, the marginal utility of each additional unit declines. In our cases, as the amount of allocated resources increases, the restoration rate increases but the amount of this increase declines. According to definition, the law of diminishing marginal utility requires that

$$\frac{du(x)}{dx} > 0, \qquad \frac{d^2u(x)}{dx^2} < 0$$

Third, there is an upper bound for the restoration rate, i.e., the rate of restoration can not be arbitrarily large even if the resource allocation could be infinitely large. This is true that a single recovery action has to spend a minimum amount of time. For example, when a machine is compromised by an Internet worm, a minimum time must be spent before it can be recovered. The time is needed for figuring out the problem, installing a solution, etc. Thus we have $\lim_{x\to\infty} u(x) = A$, where A is the upper bound.

B. Time-invariable Restoration Rate

One class of function that satisfy the above three necessary conditions is:

$$u(x) = A(1 - \rho^{\delta x}), 0 < \delta \leq 1, 0 < \rho < 1 \tag{1}$$

where ρ and δ are parameters that control the effectiveness in using resource x. This class of function is widely used by economists [2] in modeling the relationship between the investment and the outcome. Fig. 3 shows how the restoration rate changes when allocation x increases. Clearly, the three conditions are all satisfied. The time-invariable feature says that the restoration rate never changes during the whole period of restoration, given a specific x.

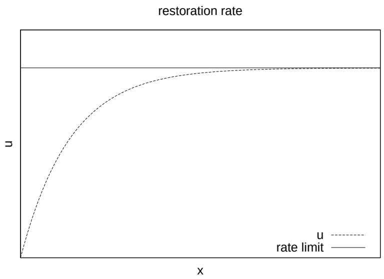

Fig. 3. Function for Time-invariable Restoration Rate

C. Time-variable Restoration Rate

In addition to nonlinearly relating to the resource allocation, the restoration rate u can also change over the time in some scenarios. Such a function is denoted as $u(x,t)$, where parameter t is for time. For example, given a group of persons fixing a worm attack, when time goes by, the increased experience would result in quicker restoration rate. In addition to the three necessary conditions, the time-variable restoration rate should satisfy $u(x,0)=0$. Here we propose to use the following time-variable function:

$$u(x) = A(1-\rho^{(1-r^{\theta t})x}), 0<\theta\le 1, 0<\rho<1, 0<r<1 \quad (2)$$

where ρ, θ and r are parameters that control the influence of resource x and time t. Fig. 4 shows the relationship. Clearly, the same upper bound exists since each restoration action needs at least a minimum time. The function also captures the feature that the more the resources allocated, the higher the restoration rate. But there exist an efficient zone where increasing the resource usage contributes greatly towards the restoration rate.

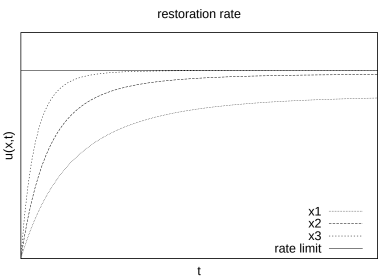

Fig. 4. Function for Time-variable Restoration Rate, here $x1 < x2 < x3$.

IV. Cost Analysis

In this section we present analysis on minimum cost for systems that the restoration rate u has the aforementioned function in relating to the resource usage and time. In the general discussion, we use $u(x)$ for simplicity. We assume the time that the system starts recovery is time $t=0$. The system state is illustrated at the right side of Fig. 2. Thus, total restoration of the system needs time:

$$\frac{c}{u(x)-v} \quad (3)$$

Let the maximum *loss rate* to be m when the system is totally compromised. When the system is partially compromised, the loss rate is proportional to the percentage of system compromised. Thus at time 0, the loss rate is $\frac{c}{l}m$.

It is easy to see that the loss rate at time t is

$$\frac{c-(u(x)-v)t}{l}m$$

So the loss in this period is

$$\begin{aligned} C_1 &= \int_0^{\frac{c}{u(x)-v}} \frac{c-(u(x)-v)t}{l} m\,dt \\ &= \frac{mc^2}{2l(u(x)-v)} \end{aligned} \quad (4)$$

The expense for restoring the system during the time (3), is

$$C_2 = \frac{c}{u(x)-v}x \quad (5)$$

Then the total cost for the restoration is

$$\begin{aligned} C &= C_1 + C_2 \\ &= \frac{mc^2}{2l(u(x)-v)} + \frac{c}{u(x)-v}x \end{aligned} \quad (6)$$

The constraints are: first, $u(x) > v$ must be satisfied, since this is necessary for a successful restoration; second, the owner's resources are limited, we must have

$$\frac{c}{u(x)-v}x \le R$$

where R is the upper bound of owner's resources.

The problem of how to allocate resources so that the total cost is minimum can be formulated as:

minimize

$$f(x) = \frac{mc^2}{2l(u(x)-v)} + \frac{c}{u(x)-v}x \quad (7)$$

subject to

$$\frac{c}{u(x)-v}x \le R \quad (8)$$

$$u(x) > v \quad (9)$$

A. Cost Analysis for Time-invariable Restoration Rate

For time-invariable restoration rate $u(x)$ with function (1), the minimum cost problem can be rewritten to

minimize

$$\begin{aligned} f(x) &= \frac{(mc^2)/(2l)}{(A(1-\rho^{\delta x})-v)} \\ &+ \frac{cx}{A(1-\rho^{\delta x})-v} \end{aligned} \quad (10)$$

subject to

$$\frac{cx}{A(1-\rho^{\delta x})-v} \le R \quad (11)$$

$$A(1-\rho^{\delta x})-v > 0 \quad (12)$$

No closed form solutions can be obtained directly. Numerical results will be presented in the next section. On the other hand, we give further analysis on constraint (11), which is critical for a successful restoration. The constraint (11) can be rewritten as

$$\frac{c}{R}x + v \le A - A\rho^{\delta x} \quad (13)$$

Let $\theta(x) = \frac{c}{R}x + v$ and $u(x) = A - A\rho^{\delta x}$. The two curves could have zero, one, or two intersecting points (Fig.

5) depending on the values of c, v, and R. When c/R or v is too large, there is no intersection. For $\theta(x) \leq u(x)$ (so to satisfy constraint 11), these two curves must have at least one intersecting point. The interval between the two intersection points a and b gives the operational region for resource allocation for a specific scenario where the system has given v, c and R.

On the other hand, the least condition for a possible restoration operation occurs when the two curves are tangent. At that point, we are able to study the relations among system conditions c, v, and R. The results are critical which directly indicate whether a restoration is feasible before the system is totally compromised. We solve the problem by taking derivative of both sides of (13)

$$\frac{c}{R} = -A\delta\rho^{\delta x} \ln \rho$$

Then we get the boundary condition for constraint (11):

$$\frac{\delta R(A-v)}{c} \ln \rho + 1 = \ln c - \ln(-AR\delta \ln \rho) \tag{14}$$

Studying the relations of c, v, and R depicted by Condition (14) leads to important results. For examples, we can obtain the minimum resource required for a success restoration given v and c; We can also calculate the latest time (the maximum portion c) that the owner has to start restoration given the v and R. These results are shown in Section V.

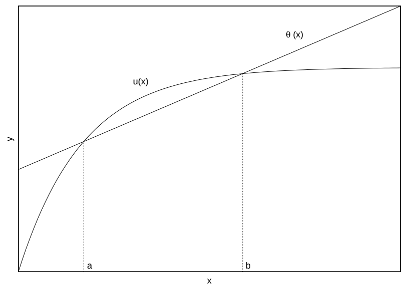

Fig. 5. Constraint $\theta(x) \leq u(x)$

B. Cost Analysis for Time-variable Restoration Rate

The restoration rate u is now a function of both x and time t as represented by Formula(2), i.e., $u(x,t) = A(1-\rho^{(1-r^{\theta t})x})$. In solving the minimum cost problem, we use differential equations. Fig. 6 shows a partially compromised system at time t (time 0 is the time that the restoration begins, with c portion of the system compromised). Thus the light shaded part represents the restoration effort since $t = 0$. This portion is denoted as $s(x,t)$, given the resource allocation rate x. Note that $s(x,t)$ could be positive or negative. If $s(x,t)$ is negative, it means that after restoration begins, more machines are compromised (since the initial restoration rate could be small). Let $f(x,t)$ be the total cost incurred till time t, $h(x,t)$ be the expense until time t, then we have the following derivative equations:

$$\frac{ds(x,t)}{dt} = u(x,t) - v \tag{15}$$

$$\frac{df(x,t)}{dt} = \frac{c - s(x,t)}{l} m + x \tag{16}$$

$$\frac{dh(x,t)}{dt} = x \tag{17}$$

$$s(x,0) = 0, f(x,0) = 0, h(x,0) = 0 \tag{18}$$

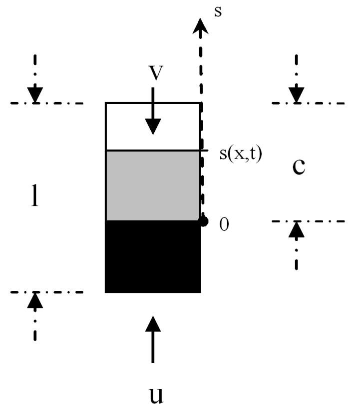

Fig. 6. System with $u(x,t)$

Let the time that the system is totally restored be $\tau(x)$, the total cost incurred during this restoration procedure is then $f(x, \tau(x))$. For each possible x, there is a corresponding cost f. Our task is to find the x (denoted as x^*) where the corresponding cost f is minimum. The problem is formulated as follows:

$$\begin{aligned} \text{minimize} \quad & f(x, \tau(x)) \\ \text{subject to} \quad & h(x, \tau(x)) \leq R \end{aligned} \tag{19}$$

We will give numerical solutions in the next section.

V. Numerical Results

In previous sections we have introduced our cost models for two restoration rates, one is resource-bounded and the other has a time constraint in addition. The minimal cost solutions for the two cases can not be solved in the closed forms. In this section we provide numerical results. All the numerical solutions are obtained through Maple [4].

For the purpose of easy illustration, we use and present integer numbers in the results. These numbers allow us to show the relations among the variables in question and the scale of changing trends. However these numbers by no means reflect any real scenarios in practice. To make the model useful for a particular real world event, estimations of the parameters must be conducted. The specific system parameters are: The network is composed of $l = 5,000$ nodes; the maximum loss rate is $m = 100,000$ if the system is totally compromised; The maximum possible restoration rate is $A = 100$. The following parameters are varied for the different results we show. But when not varied, they take the following default values: The maximum resources is $R = 50,000$; The compromise rate is $v = 50$; Initial compromised system is $c = 500$.

A. Time-invariable Restoration Rate

For time-invariable restoration rate $u(x)$ with function (1), we use parameters $\rho = 0.996$ and $\delta = 0.9$. The curve is illustrated in Fig. 3.

1) restoration vs. resource allocation: Fig. 7 shows the increase and decrease trends of the total cost, expense and loss when the resource allocation rate (x) increases. The figure shows that as x increases, expense increases but loss decreases. As a result, total cost first decreases then increases. At one specific point, total cost reaches its minimum value.

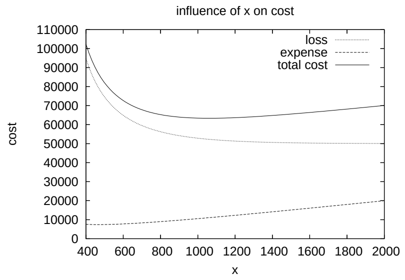

Fig. 7. Total cost changes as x increases

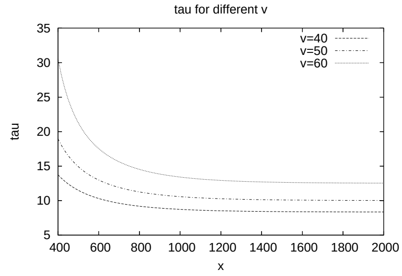

Fig. 8. τ changes as x increases.

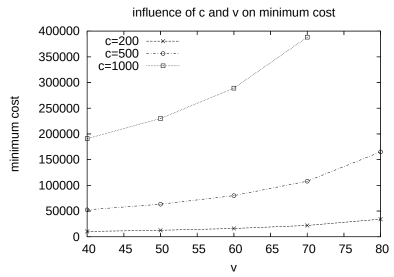

Fig. 9. Influences of c and v on minimum cost.

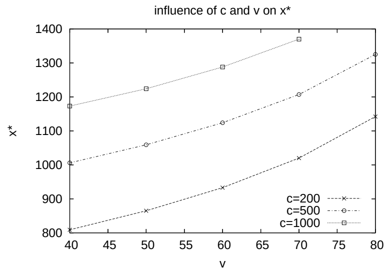

Fig. 10. Influences of c and v on x^*.

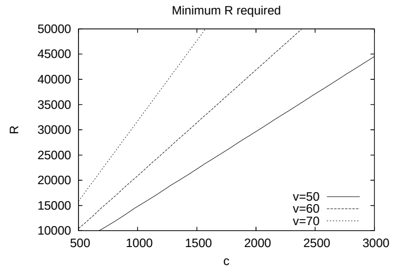

Fig. 11. Resource for a successful restoration

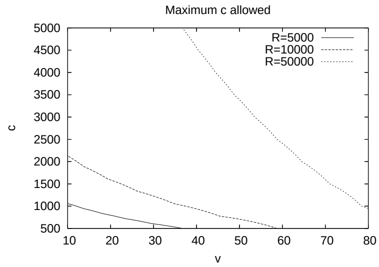

Fig. 12. Restoration constrained by R.

Our other results also indicate that a larger compromise rate will incur larger total cost, but minimum cost is achievable. On the other hand, when resource allocation rate x increases, time spent for recovery decreases. We have used τ to denote the time needed for total restoration. The result is shown in Fig. 8. Clearly, when more resource can be allocated, it is not necessary that the time needed to restore the whole system reduces proportionally, since the usage of the resource becomes less effective. When the compromise rate increases, τ increases as expected.

2) Achieving minimum cost: We have seen that for each specific c and v, there exists a x^* at which the cost is minimal. The Figures 9 and 10 show how the minimum cost C^* and x^* are influenced by various c and v. The observation is that as c or v increases, minimum cost C^* and the associated x^* increase. The figures also show that when v and c are too large (here when $v > 70$ at $c = 1000$), the minimum cost is not achievable.

3) The critical condition: The relationship among c, v, and R is expressed in Equation (14). It provides us opportunities to illustrate constraints from system parameters on a successful restoration. For example, Fig. 11 shows the minimum total resource required for a successful restoration under various c and v. A later start of restoration (larger c) or a stronger attack (larger v) all require larger amount of available resources. Less total resource will lead to a failure of restoration. Fig. 12 shows the operational region when the total resource is given. The figure suggests that for a specific v, the restoration must start no later than the presented value of c. Otherwise, no matter how the allocation rate x will be, total restoration is impossible. The figure also shows that the larger the R, the wider the operational region, which tolerates larger ranges of v and c.

B. Time-variable Restoration Rate

The time-variable restoration rate $u(x, t)$ has a function in the form of Equation (2). We use parameters $\rho = 0.996$, $\theta = 0.993$ and $r = 0.9$ here. The curve is illustrated in Fig. 4. The solutions towards the minimum cost rely on solving equations (15) to (18) (Section IV-B).

With the help from Maple, we are able to get the values of costs for various system variables, and to draw the various cost curves. Fig. 13 shows how a cost value can be calculated. Given an x, curves $s(x, t)$, $h(x, t)$ and $f(x, t)$ are generated by Maple. At the time of $s(x, t) = c$, the whole system is restored. Recall the time is denoted as τ. Thus the value of $f(x, t)$ at $t = \tau$ is the total cost spent for recovery. The value of $h(x, t)$ at τ is the expense. We use this method to compute the cost for each x and the minimum cost is the minimum value cross all the x's.

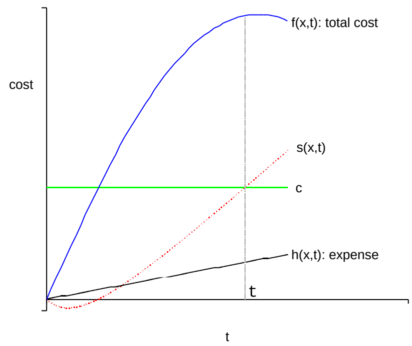

Fig. 13. Calculating the cost.

1) restoration vs. resource allocation: Fig. 14 shows the cost, expense and loss as the functions of the resource allocation rate x. Here $c = 500$ and $v = 50$. The figure shows the same cost change trends as Fig. 7 where $u(x)$ does not change over time. The similarity of the two figures is expected since $u(x, t)$ should only affect the value of the cost, but not the relation between the cost and the resource allocation rate. The minimum cost exists at x^*.

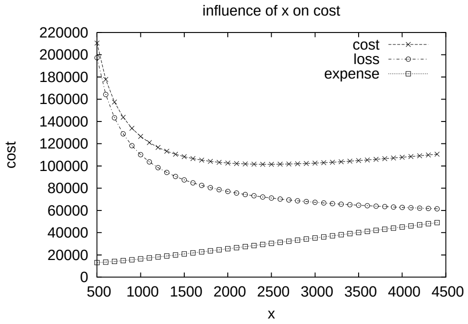

Fig. 14. The cost changes as x increases for $u(x, t)$.

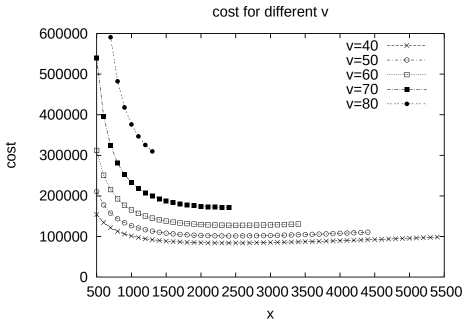

Fig. 15. Cost vs. v, for $u(x, t)$.

Fig. 15 then shows how the cost changes with increasing x at different compromise rate v. It is clear that as v increases, the cost increases. The figure also suggests that minimum costs exist. However, the figure reports early terminations of the cost curves for some large v. For example, the curve for $v = 80$ ends before x reaches 1500 while the curve for $v = 40$ extends well beyond $x = 5000$. This phenomenon indicates the failures of total recovery at these non-data points because of the exhaustion of the resources. On the other hand, the figure suggests that when the compromise rate is high and the resource is limited, a smaller rate of resource allocation, which possibly produces more efficient restoration, could yield a better result, i.e, a complete restoration. Our other results also show that corresponding curves of τ reveal the same phenomenon.

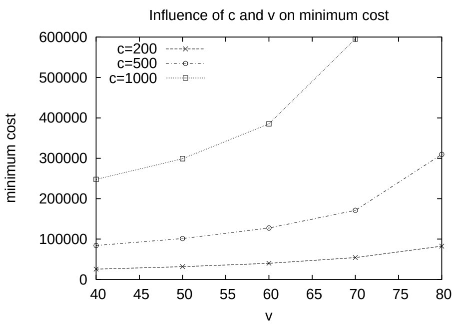

Fig. 16. Influences of c and v on minimum cost, for time-varying rate.

2) Achieving minimum cost: Figures 16 and 17 show how c and v influence the minimum cost C^* and the corresponding x values x^*. Fig. 16 shows that increasing v and c incur higher C^*, similar to the trends observed in Fig. 9. The figure also suggests that when v and/or c is too large, no total restoration is possible. Thus the minimal cost can not be achieved. Or, say, the minimal cost occurs at the largest successful allocation rate. This explains Fig. 17 where x^* decreases when v increases. Those x^* are the points that restoration is possible, but not minimum.

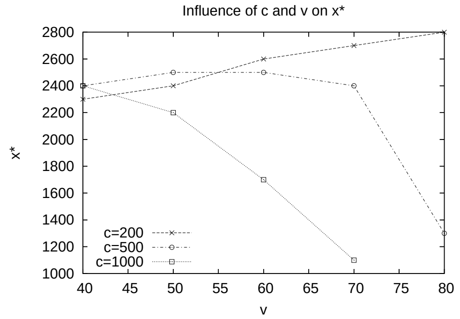

Fig. 17. Influences of c and v on x^*, for time-varying rate.

VI. Conclusions

In this paper, we provide analysis on the cost incurred when restoring compromised systems. Resource allocation is the major influential factor in our analysis. Typically the usage of the resource contributes to the restoration nonlinearly, following *the law of diminishing marginal utility*. We study both time-variable and time-invariable behaviors. Our results are presented using numerical methods. The results suggest that the minimal cost is achievable for many conditions, but tight optional regions also exist. Our future work is searching data to estimate the parameters for practical scenarios and validating the models.

References

[1] W. J. Baumol and A. S. Blinder. *Economics: Principles and Policy*. South-Western College Pub, 2004.

[2] L. A. Gordon and M. P. Loeb. The economics of information security investment. *ACM Transactions on Information and System Security*, 5(4):438–457, November 2002.

[3] J. H. Lala. Introduction. In *Proceedings of the Foundations of Intrusion Tolerant Systems (OASIS'03)*, pages x – xix. IEEE, 2003.

[4] MapleSoft. Maple. In *http://www.maplesoft.com/*.

[5] C. Meadows. A cost-based framework for analysis of denial of service in networks. *Journal of Computer Security*, 9(1-2):143–164, January 2001.

[6] D. Moore, C. Shannon, and J. Brown. Code-red: a case study on the spread and victims of an internet worm. In *Proceedings of the ACM SIGCOMM/USENIX Internet Measurement Workshop*, pages 273 – 284. ACM, November 2002.

[7] S. Ray, Q. Zheng, X. Hong, and K. Kwiat. Integrity function - a framework for server replication and placement in adversarial environment. In *submitted to IEEE Transactions on Parallel and Distributed Systems*.

[8] P. E. Verissimo, N. F. Neves, and M. P. Correia. Intrusion tolerant architectures: Concepts and design. *Architecting Dependable System, Lecture Notes in Computer Science*, 2677(44):3–36, May 2003.

[9] C. Zou, W. Gong, and D. Towsley. Code red worm propagation modeling and analysis. In *Proceedings of the 9th ACM conference on Computer and communications security*, pages 138 – 147. ACM, November 2002.