Comparative Visualization of Predicted Disaster Resilience

Christopher W. Zobel
Virginia Polytechnic Institute and State University
czobel@vt.edu

ABSTRACT

The disaster resilience triangle is a simple but effective tool for illustrating the relationship between the initial impact of a disaster event and the subsequent time to recovery. This tool can also be expanded, however, to provide an *analytic* measure of the level of resilience exhibited by a particular entity in a given disaster situation. We build upon the previous work in this area by developing a new approach for visualizing and analyzing the tradeoffs between the two primary defining characteristics of the disaster resilience triangle. This new approach supports strategic decision making in a disaster planning environment by providing a straightforward means for directly comparing the relative predicted resilience of different critical facilities within an organization, with respect to both location and type of risk.

Keywords

Resilience triangle, visualization, decision support.

INTRODUCTION

In the wake of events such as Hurricane Katrina and 9/11, there is an increasing awareness that the concept of disaster resilience can be used to help characterize an organization's ability to resume an acceptable level of functionality in the aftermath of a disaster. Although, in popular usage, the term resilience is traditionally considered to refer to "the act of rebounding or springing back" (OED Online, 2003), the resilience literature tends to be divided between this strict focus on the recovery process (Rose, 2004; Rose et al., 2007) and an expanded focus that also considers the effects of mitigation against the initial impact of the disaster event (Shinozuka et al., 2004; Tierney and Bruneau, 2007; McDaniels et al., 2008). This paper takes the viewpoint that both initial resistance to a disaster and the subsequent process of recovery are important factors that characterize an organization's ability to resume "normal" operations.

The ability to measure and analyze disaster resilience is particularly important in the context of strategically locating critical facilities such as supply chain distribution hubs or network operations centers. We propose that disaster resilience can function as an effective multi-dimensional measure by which the suitability of each potential location for such a facility can be compared, with respect to a variety of different potential hazards. In support of this, we introduce a new approach to both visualizing and analytically representing the underlying relationship between the two characteristics of resilience discussed above: robustness against initial loss, and rapidity of the recovery process. We begin the discussion by considering the background of the resilience triangle introduced by (Bruneau et al., 2003), and by discussing the issue of tradeoffs between these two primary defining factors. We then define the notion of *predicted resilience* as an analytic measure and provide a simple and straightforward approach to calculating its value. Based upon these results, we then introduce and characterize our new multi-dimensional approach to visualizing and comparatively assessing the predicted resilience of different facilities under the presence of different types of disaster risk.

THE RESILIENCE TRIANGLE

Background

As initially proposed by Bruneau, et al. (2003), disaster resilience is characterized by the extent to which the factors of robustness, rapidity, resourcefulness, and redundancy are present in a physical or social system.

Reviewing Statement: This paper represents work in progress, an issue for discussion, a case study, best practice or other matters of interest and has been reviewed for clarity, relevance and significance.

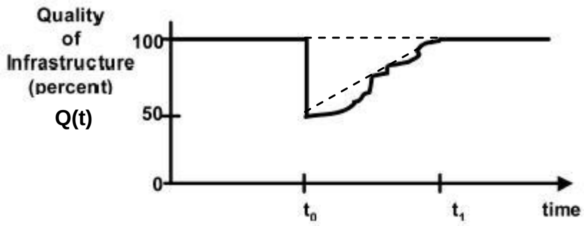

Figure 1: Original resilience triangle (adapted from (Bruneau et al., 2003))

The last two factors, resourcefulness and redundancy, are generally considered to be the "means" by which disaster resilience can be improved, and the corresponding "ends" are typically measured by the impact of such improvements on the first two factors, robustness and rapidity (Bruneau et al., 2003). With this in mind, Bruneau, et al. (2003) introduced the resilience triangle given in Figure 1. This "triangle" leads to a simple measure of the *loss of resilience* in a system, defined mathematically as follows:

$$R = \int_{t_0}^{t_1} [100 - Q(t)]dt \qquad (1)$$

where $Q(t)$ represents the quality of the system's infrastructure at a given time t. The vertical axis in Figure 1 can be thought of as representing the level of robustness of the given system and the horizontal axis as representing the associated rapidity of recovery (Bruneau and Reinhorn, 2007).

Approximating Resilience

Because the original definition of R, as given in Equation 1, represents the area between the curve $Q(t)$ and the line Q=100, we can create a simple approximation to its value by calculating the area of the actual triangle formed by (1) the initial drop in functionality and (2) the corresponding time to recovery (represented by the dotted lines in Figure 1). As we will discuss below, the area of this triangle can then be used as the basis for generating a direct measure of resilience, as suggested by the work of Cimellaro, et al. (2010).

It is important to recognize, however, that if we measure resilience as a single number that is some function of the area of the resilience triangle, then very different combinations of initial loss and time to recovery can correspond to exactly the same level of resilience. Because resilience defined in this way depends on the scaled product of these two factors, a facility which suffers little initial loss but which has a very long time to recovery may have the same measured amount of resilience as one which has significant immediate loss but a rapid recovery time (See Figure 2). These represent very different situations, however, and a given decision maker may have a relative preference of one of these two scenarios over the other. In such a case, it would be desirable to be able to clearly differentiate between the relative value of the two scenarios to that particular decision maker.

With this in mind, we propose an approach to representing decision makers' preferences that involves explicitly incorporating the multi-dimensional nature of resilience. By doing so, we are treating the actual resilience number as just one aspect of the overall resilience of a given situation, but we are also providing it with a general context within which it ultimately gains more meaning. This makes it possible to more easily compare the relative resilience of different facilities, and to support such comparisons in the presence of a number of different preference structures.

Defining Predicted Resilience

For the purposes of the following discussion, we focus on the predicted resilience of a given facility, such as a network operations center, with respect to a particular disaster event. We define such predicted resilience, R, to be a function of the predicted initial loss, X, measured as a percentage of total functionality, and the predicted time to recovery, T, measured in a relevant time unit such as weeks. Both X and T are estimates of the actual values for loss and recovery time, respectively, and they are based on the best information available before the event occurs. It may also be possible, of course, to estimate resilience immediately after an event has occurred. In this case we may replace the predicted initial loss with the actual initial loss.

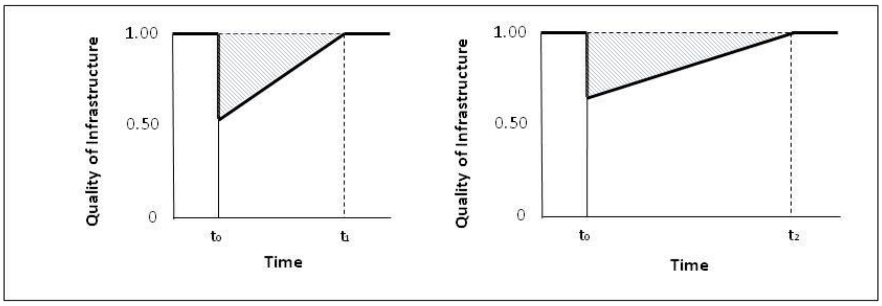

Figure 2: Two triangles representing the same calculated resilience, but with different proportions

If we further assume that the recovery rate for any given facility is linear in nature, to keep our formulation simple, then as time progresses we would expect the ratio between the current loss and the remaining time to recovery to be constant. We follow Cimellaro, et al. (2010) in defining robustness, $R_b = 1\text{-}X$, so that larger values of robustness generally correspond to fewer losses and thus more resilience. Furthermore, we define T to be a direct measure of rapidity, i.e., $R_p=T$, in that smaller values of T indicate a quicker recovery time.

Although the area of the triangle directly defined by X and T (i.e., $XT/2$) can lead to a simple measure of resilience, this simple measure must be carefully considered. In particular, because equation (1) represents *loss of resilience*, and not resilience itself, smaller triangles naturally correspond to smaller values for R. From the standpoint of defining a direct measure of resilience, we would expect the opposite to be true: that a more resilient system should be represented by a higher amount of measured resilience. In (Bruneau and Reinhorn, 2007) and (Cimellaro et al., 2010), the authors address this issue by considering the area *outside* of the resilience triangle, rather than that of the triangle itself, to represent actual resilience, so that a smaller triangle will directly correspond to a larger resilience value. They accomplish this by making use of the upper bound of 1 (100%) on the value of X and by calculating the area under the curve over a fixed time period that includes the entire recovery process.

Because we are estimating X and T *a priori* and assuming a consistent shape for our predicted resilience triangle, we can generate an appropriate value for resilience that does not require the explicit formulation of an objective function or the calculation of an integral. We do this by simply subtracting the area of the resilience triangle from a fixed larger area (see Figure 3). This larger area is determined by the upper bound on X and by a new parameter, T^*, where $T^* >= \max(T)$, for all T under consideration. This latter parameter allows us to compare multiple instances of predicted resilience on the same relative scale.

The predicted resilience measure for the ordered pair (X, T) is thus given by:

$$R(X,T) = \frac{T^* - \frac{XT}{2}}{T^*} = 1 - \frac{XT}{2T^*} \qquad X \in [0,1],\ T \in [0,T^*] \tag{2}$$

where T^* (i.e., $1 \cdot T^*$) is the larger area from which we subtract the area of the resilience triangle for X and T. Dividing the difference between the two area values by this larger value allows us to represent resilience as a percentage, as in (Cimellaro et al., 2010).

Note that the minimum value for $R(X, T)$ under this formulation (X=1 and T=T^*) is 50%, and that its maximum value (X=0 or T=0) is 100%, because we are subtracting the area of a triangle that is, by definition, completely

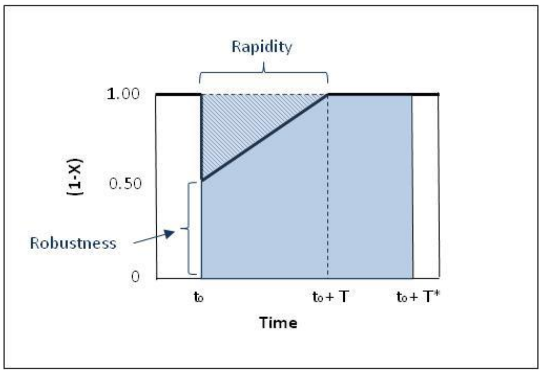

Figure 3 – Predicted resilience triangle as a proportion of *T**

contained within the rectangle of size T^*. Although the range of this function can easily be scaled up to [0,100], if desired, the minimum value of 50% is appropriate in this context since it accurately captures the assumption of a linear recovery rate.

We choose T^* to be greater than or equal to max(T) so that we can compare different instances of (X,T) on *total* predicted resilience, rather than just on the predicted amount of resilience within a particular short time frame. In general, T^* may be interpreted as the maximum acceptable amount of recovery time. For example, there may be a practical upper limit on the amount of time that a given company would be willing to wait for any of their facilities to be completely rebuilt. If the estimated time to recovery for any such facility is greater than that upper limit, then it would likely be written off as a loss.

VISUALIZING RESILIENCE

Representing the Tradeoffs

Given a formulation for resilience in terms of X and T (equation 2), we can now represent a specific instance of the predicted resilience triangle as the coordinate pair (X, T). For a fixed value of R (and also T^*), the relationship between X and T is given by

$$XT = (X\text{-}0)(T\text{-}0) = M, \quad \text{where} \quad M = (2T^*)(1 - R). \tag{3}$$

This is exactly the equation of a rectangular hyperbola that is centered at the origin and that has asymptotes lying directly on the X and T axes. Thus for any value of R, the set of possible values for X and T describe a hyperbola. By varying R, we subsequently achieve a series of equilateral hyperbolas, each of exactly the same shape, in the upper right quadrant of the plane (see Figure 4). Larger values of R correspond to smaller values of M in equation (3), and therefore curves closer to the origin have better (higher) resilience values.

Observations which lie on the lower right portion of each of these resilience curves represent scenarios for which the initial disaster-related loss is very high but the recovery time is small, whereas those which lie on the upper left portion of each curve represent a long recovery time despite a relatively small initial loss. Observations whose values of X and T place them in the middle of a curve are those which are more evenly balanced between the two factors. Because all points on a given curve share exactly the same resilience, this allows us to explicitly capture and represent the tradeoffs between the different values of X and T (and thus the different relative shapes of resilience triangles).

Support for Comparative Visualization

There are several benefits associated with representing the multi-dimensional nature of the resilience triangle in this way. First of all, by plotting the expected resilience of multiple facilities, they can all be compared simultaneously with respect to both their overall resilience and their relative levels of robustness and rapidity. For example, in Figure 4, Facility B has the same expected resilience as Facility A even though the initial loss (value of X) is much less, whereas Facility C has a slightly larger loss than Facility B and a slightly longer time to recovery than Facility A, but a better overall resilience than both of them.

Different decision makers' preferences, with respect to acceptable levels of robustness and rapidity, may also be represented by different regions within the graph. For example, we would expect that a given decision maker would be partial to observations falling on the lower portion of the graph if they believe that customers would prefer no service for a very short period of time over partial service for a longer period of time. Similarly, someone who is willing to accept a bit longer time to full recovery, in exchange for less of an initial impact, might prefer observations that cluster around the line X=T.

Because all observations which lie along any straight line passing through the origin have the same relative ratio of X to T, their associated predicted resilience triangles all share the same shape. Observations which fall below such a line have a relatively larger value of X with respect to T, and those falling above the line have a relatively smaller value of X with respect to T. Consequently, for example, a decision-maker can easily identify the range of possible resilience instances for which the relative amount of initial loss (X) is at most a given percentage of the expected time to recovery (T). As shown in Figure 5, this can also be combined with an explicit indication of the set of resilience values above a certain desired level, R^*, in order to more precisely represent a decision maker's preferences.

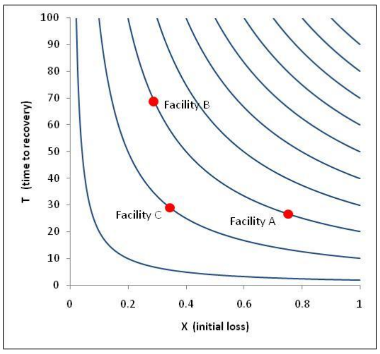

Figure 4: Basic resilience curves

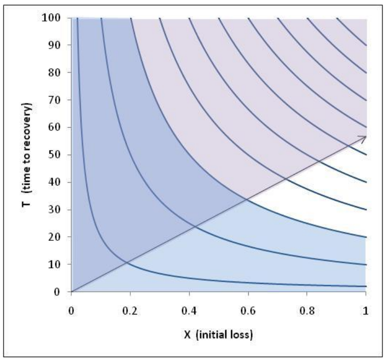

Figure 5: *X* at most a given percentage of *T*, and *R* ≥ *R**

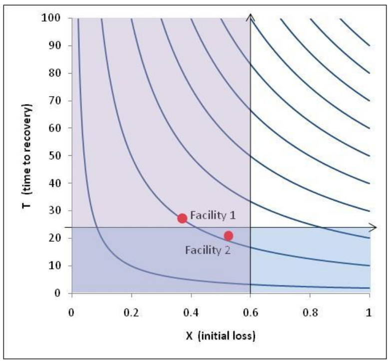

Figure 6: Acceptable regions for both recovery time and initial loss

This same approach can be taken to identify facilities whose robustness against loss (and/or expected time to recovery) falls within a particular preferred range of values. Thus, in comparing several possible alternative mitigation scenarios, those that are not expected to exceed a specified level of initial loss easily can be identified and then compared against other scenarios with respect to the relative length of their time to recovery. For example, a given decision maker may be willing to accept slightly higher initial loss (and a correspondingly larger overall resilience value) in order to avoid exceeding a given length of recovery time (see Figure 6). This helps to illustrate the importance of understanding the multi-objective nature of the measure, in order to support more comprehensive decision making.

SUMMARY AND CONCLUSIONS

This initial research effort extends the concept of resilience by providing a new approach to visualizing and comparing different combinations of the predicted measures of initial loss (X) and time to recovery (T), with respect to the predicted overall resilience that they represent. Providing such a comprehensive representation supports a more complete understanding of the contributions of these parameters and allows for improved support of the decision-making process.

The new approach concisely represents the relationship between the parameters by identifying a series of hyperbolic curves, each of which defines the set of resilience triangles that share a given area. Although all observations on a given curve have the same overall resilience value, each one has a unique combination of values for *rapidity* (T) and *robustness* ($1-X$). This ability to display the relative values of the two supporting parameters at the same time as the overall resilience provides the opportunity to simultaneously compare the relative resilience of different facilities or entities in a very rich information environment. In a similar manner, it also allows for the relative resilience of a single entity to be directly compared under several different mitigation scenarios.

REFERENCES

1. Bruneau, M., Chang, S. E., Eguchi, R. T., Lee, G. C., O'Rourke, T. D., Reinhorn, A. M., Shinozuka, M., Tierney, K., Wallace, W. A. & von Winterfeldt, D. (2003) A framework to quantitatively assess and enhance the seismic resilience of communities. Earthquake Spectra, 19, 733-752.
2. Bruneau, M. & Reinhorn, A. (2007) Exploring the concept of seismic resilience for acute care facilities. Earthquake Spectra, 23, 41.

3. Cimellaro, G., Reinhorn, A. & Bruneau, M. (2010) Seismic resilience of a hospital system. Structure and Infrastructure Engineering, 6, 127-144.
4. McDaniels, T., Chang, S. E., Cole, D., Mikawoz, J. & Longstaff, H. (2008) Fostering resilience to extreme events within infrastructure systems: Characterizing decision contexts for mitigation and adaptation. Global Environmental Change, 18, 310-318.
5. OED Online (2003), Vol. 2009, pp. Oxford University Press.
6. Rose, A. (2004) Defining and measuring economic resilience to disasters. Disaster Prevention and Management, 13, 307-314.
7. Rose, A., Adosu, G. & Liao, S. Y. (2007) Business interruption impacts on the electric power system resilience to a total blackout of a terrorist attack of Los Angeles: Customer resilience to a total blackout. Risk Analysis, 27, 513-531.
8. Shinozuka, M., Chang, S. E., Cheng, T.-C., Feng, M., O'Rourke, T. D., Saadeghvaziri, M. A., Dong, X., Jin, X., Wang, Y. & Shi, P. (2004) Resilience of Integrated Power and Water Systems. In: MCEER Research Progress and Accomplishments: 2003-2004(Ed, MCEER), pp. pp. 65-86. Buffalo, NY.
9. Tierney, K. & Bruneau, M. (2007) Conceptualizing and Measuring Resilience: A Key to Disaster Loss Reduction. In: TR News, pp. 14-17.