Innovation and Supply Chain Management, Vol. 7, No. 3, pp. 001–013, November 2013

# Towards Systems Resilience

Hiroshi MARUYAMA [1] and Kazuhiro MINAMI [1]

[1] Trans-disciplinary Research Integration Center, The Research Organization of Information and Systems, Tokyo, Japan

**Abstract** : Some systems are *resilient* – they recover from unanticipated large scale perturbations and continue to exist – while others are not. We started a new project called "Systems Resilience" that aims to study underlying principles of resilience. We study a wide variety of domains, such as biology, ecology, economics, engineering, sociology, and business management, and extract common characteristics that differentiate resilient systems from non-resilient ones. We are building a mathematical model to represent these characteristics and to quantify resilience.

**Key Words :** Resilience; extreme events; redundancy; diversity; adaptability

## 1. Introduction

After the 3.11 earthquake there has been significant increase on the discussions on events that are both rare and having significant impacts, sometimes called *X-Events* [6]. These "surprising" events occur as an outside of the anticipated envelope (e.g., Tsunami of 14m high vs the anticipated max of 5.7m), or something completely unheard of (e.g., Tokyo subway gas attack in 1995). "X-Events" are inevitable. They do happen, and our systems may fail. It is not possible, theoretically or practically, to protect our systems for all X-Events. When the system fails, we have to recover from the damage as quickly and as inexpensively as possible. This process may not be exactly a "recovery" because it may or may not restore the system into the original configuration; rather, the system can be in a completely new configuration that is also acceptable, or even desirable, to the stakeholders. The ability to make a system to withstand large perturbations and enable generalized recovery once the system fails is, in our definition, "resilience."

To shed scientific lights to the concept of resilience, we started a multi-year, multi-disciplinary project called "Systems Resilience"[*1] in April 2012, within The Trans-Disciplinary Research Integration Center in The Research Organization of Information and Systems, a national research organization under the Japanese Ministry of Education, Culture, Sports, Science and Technology. This paper overviews the project and its current status. We first give our basic assumptions and project goals. Our working hypothesis about resilience strategies are

Corresponding Author: Hiroshi Maruyama
Institute of Statistical Mathematics, 10–3 Midori-cho, Tachikawa, Tokyo 190–8562, Japan
hm2@ism.ac.jp
(Received September 10, 2013)

[*1] See http:http://systemsresilience.org/index-e.html for the details.

covered in Section 3. In Section 4, we briefly touch on the mathematical model we are building for explaining these strategies. We conclude the paper by discussions in Section 5, followed by the future directions. This paper is an extended version of our paper presented at the 1st International Workshop on Systems Resilience held in Budapest, 2013 [10].

## 2. Project Goals

Resilience is observed in many different systems; human minds exhibit resilience after suffering real hardship, ecosystems such as rain forests and the Antarctic ocean are known to recover after perturbations, biological systems such as human body can recover from injuries, the global communication systems have survived many technological faults, natural disasters, policy changes, etc., many cities demonstrated that they could withstand large historical events such as pandemics, wars, and geopolitical changes, and so on.

Our basic assumption is that there must be a common set of strategies that make these systems resilient, regardless their own domain, be it a biological, ecological, economical, organizational, or engineering system.

Provided that this assumption holds, our goal is to investigate these common strategies and organize them into an organized body of knowledge (BoK). This "Resilience BoK" will guide us when we design and operate a (social, engineering, or some other) system so that we can make better, scientifically-accountable decisions to make it more resilient. To do this, the BoK will catalogue various resilience strategies and describe when and how these strategies should be applied.

This goal is not easy to achieve. Our approach is to take the following three steps. First, we collect experiences of resilience (and also non-resilience) in many different fields as many as possible, and extract common strategies among them. Second, we build a mathematical model to explain why these strategies work and when. If the model can provide reasonable explanations to the past events, we apply the model to open problems in different fields to verify it.

## 3. Working Hypothesis

We studied various fields and came to a temporary hypothesis that we can categorize resilience strategies into three categories: *redundancy*, *diversity*, and *adaptability*.

### 3.1 Strategy 1: Redundancy

Redundancy is a frequently-used resilience strategy seen in many domains.

#### 3.1.1 Redundancy in Biological Systems

Biological systems are known to have a large redundancy. For example, E. Coli has approximately 4,300 genes, each of which has its unique function, but almost 4,000 of them are known to be redundant – that is, knocking out one of them will not hamper its ability to reproduce [25].

Three-spine stickleback is fresh-water fish that had lost their armor plates when they migrated to fresh water from sea water about 10,000 years ago. A sample caught in Lake Washington in 1957 had no armor plates as in Figure 1(a) but more recent samples have armor plates (see Figure 1(b)). One theory to explain this change is that they regained armor plates because of the predation pressure by trouts whose population had increased during this period due to the increase of the water transparency in the lake. The genotype of the armor plates


(a) Three-spine stickleback captured in 1957.


(b) Three-spine stickleback captured in 2006.

Fig. 1 Adaptation of three-spine sticklebacks.

was dormant (and thus, redundant) during the peaceful years but became active when the necessity arose [14].

3.1.2 Redundancy in Engineering Systems

In engineering systems, it is a common strategy to have back-up systems to make them more reliable. For example, mission-critical storage systems use RAID (Redundant Arrays of Inexpensive Disks) so that the system can continue to function even though one or more disks fail [22].

Before 3.11, the nuclear power had accounted for about 30% of all the electricity consumption in Japan. Within 14 months after the earthquake, every one of Japan's 50 nuclear power stations went into maintenance cycles and remained non-operational until a few of them resumed a few months later. Although Japan has lost almost a third of its electric generation capacity, Japan has never experienced major blackout during this period. This can be attributed to the centralized and monopolized system of Japanese electric industry. One of their top priorities resides in the stable supply of electricity, and for that purpose Japanese electricity systems have had a huge excessive capacity.

3.1.3 Redundancy in Management Systems

The auto industry was also affected by the earthquake because their extremely complex supply chains depend on a large number of suppliers located in the Tohoku area. Despite the unprecedented scale of damage they suffered, every major auto company in Japan survived the crisis. One of the reasons of their survival was their monetary reserve that could compensate the temporary loss of the revenue. Electricity and money can be considered to be universal resource, and having extra universal resource in reserve is a good strategy for preparing unseen threats.

When the United States was attacked by the terrorists on September 11th, 2001, the police departments, the fire departments, and the secret service had difficulty in communication and coordination due to the lack of interoperability between their communication equipments. Interoperability enables one component to function as a back-up of another component. Thus, interoperability is a form of redundancy in this context.

### 3.2 Strategy 2: Diversity

Diversity is the second category of the resilience strategies.

3.2.1 Diversity in Biological Systems

Diversity plays a central role in the survival of biological systems. The first life on the earth appeared about 4 billion years ago. Since then, there have been a number of large crisis that endangered the survival of life. For example, the PermianTriassic extinction event that occured about 251 million years ago caused up to 96% of marine species to become extinct. One of the reasons that the biological systems as a whole survived is because of their diversity – some species had better capability

to deal with changing environments.

#### 3.2.2 Redundancy in Engineering Systems

Diversity is also a common strategy in mission-critical engineering systems. The Boeing 777 is the first commercial airliner whose control systems is fly-by-wire, i.e., the flight is controlled electric signals rather than hydraulic systems. These signals are controlled by a redundant system consisting of three computers. These three computers are based on different hardware and software developed by independent vendors. If these three computers share the same design, a design flaw would make all the computers fail at the same time. By having diversity in its designs, Boeing 777 can withstand a computer failure caused by a design flaw of a single computer.

#### 3.2.3 Redundancy in Management Systems

In the domain of forest management, it is a common wisdom not to extinguish small forest fires and let the patch of the forest rejuvenate. Otherwise, every part of the forest gets older and dryer, and the risk of a large-scale forest fire would much increase. The diversity of tree ages in a forest is a key to keep the forest resilient.

In an investment management, the diversification of stocks and bonds that are to be invested is a common practice. This is not the optimum strategy if the goal is to maximize the expected returns. To invest all the money on the stock with the highest expected return is the optimal solution if that is the goal. It is also a risky strategy because the investor loses all the money if the invested company bankrupts. By diversifying the investments, the investor can significantly reduce the risk of catastrophic loss in exchange for a slightly lower expected return.

#### 3.2.4 Diminishing Return

The role of diversity in complex systems is extensively discussed by Scott Page in [21]. Diversity in an ecosystem is measured by the *Diversity Index* defined as follows. Suppose that the ecosystem has $N$ distinct species and each species $i$ has the population of $p_i$. Then the diversity index $G$ of this ecosystem is:

$$G(p_1, p_2, ..., p_N) = (\sum_{i=1}^{N} \frac{p_i^2}{N})^{-1}$$

The index takes the largest value $1/p^2$ when all the species have exactly the same size of pollution $p$ (where $p = p_1 = p_2 = ... = p_N$). The value is the smallest when one species dominates the entire ecosystem, that is, $p_1 = Np$ and $p_i = 0$ where $i > 1$. In this case, the diversity index will be $1/(p^2N)$.

Given the current environment, some species are more fit (i.e., advantageous) than others, and thus, can reproduce more efficiently, and as the result their population goes up. This is captured by the following *replicator equation*.

$$p_i^{t+1} = p_i^t \frac{\pi_i}{\pi^t}$$

where $p_i^t$ is the population of species $i$ at time $t$, $\pi_i$ is the fitness of species $i$, and $\pi^t$ is the mean fitness of all the species at time $t$. Assuming this replicator equation, it can be shown that a diverse ecosystem has better chances to survive in various conditions.

If higher diversity entails better survivability, an interesting question here is how to increase the diversity. Or, what are intrinsic mechanisms to introduce diversity into a system? We argue that *law of diminishing return* plays a significant role. With

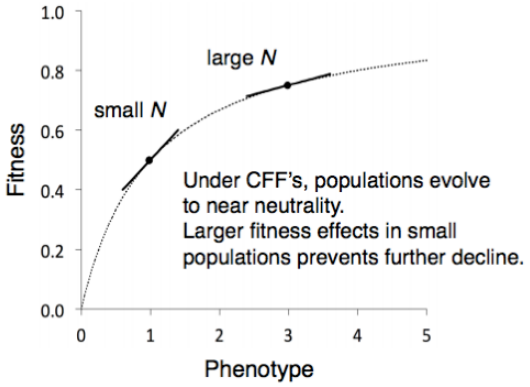


Fig. 2 Concave Fitness Function and Weak Selection

the above replicator equation, the population of a fit species will get larger by each generation, and the most fit species will ultimately dominates the entire ecosystem without a mechanism that penalizes such domination. If the fitness $\pi_i$ of species $i$ is a function of the population size, and the function $\pi_i(p_i)$ is a decreasing function, then the dominating species loses its advantage as its population increases, and this gives spaces for other species to occupy.

A gene allele is an alternative forms of a gene, often leading to no visible difference of phenotypes. Kimura [17] argued that this *neutrality* in terms of fitness is the source of gene-level diversity of biological systems. Later Ohta [20] discovered that the pure neutrality could not explain the observations of real world data, and proposed the *near*-neutral theory. Akashi et al.[1] hypothesized that a *concave fitness function*, as shown in Figure 2, could explain why we observe so much of slightly deleterious mutations in the nature. This concave function represents the law of diminishing return of cumulative advantages of alleles, because, as the species gain a larger fitness, a contribution of each advantageous mutations to the fitness declines.

Many systems, especially those appear in the nature, seem to have the law of diminishing return. For example, human sensitiveness to external stimulus is known to be logalismic. On the other hand, artificial systems are often linear, and do not follow the law of diminishing return. A prominent example is our financial system. Although the subjective value of $100 widely varies depending on whether you are a street worker or a millionaire, goods and services you can buy by your $100 stay unchanged. Thus, your money adds up linearly. This leads to polarization between the rich and the poor, and may make the society more fragile.

### 3.3 Strategy 3: Adaptability

The third resilience strategy is adaptability, which is defined as a relative speed of the systems' capability to adapt against environmental changes.

#### 3.3.1 Adaptability in Biological Systems

Biological systems are known to be very adaptive. One of the adaptability mechanisms of life is *evolution*. When a life reproduces, there are mutations on the genes. These mutations could be random, and the variations that fit the current environment most have better chances to survive. This way, species adapt themselves according to the environment change. This is a slow process whose speed is determined by the generation cycle of the species.

#### 3.3.2 Adaptability in Engineering Systems

In IT systems, IBM proposed the concept of Autonomic Computing [16] in 2003. This architecture is based on so-called the MAPE (Monitor - Analyze - Plan - Execute) cycles. It is more sophisticated than a simple feedback system, but the fundamental strategy is to make the system more adaptable – it senses the changes and react automatically to handle the situations. Thus, it is another example of the adaptability strategy.

A quicker adaptation is realized by *feedback* in both biological and engineering systems. Our body temperature is maintained by sensing the situations and secreting hormones. Air conditioning systems monitor the inside and outside temperatures and control their operations.

#### 3.3.3 Adaptability in Management Systems

A legal system is usually very rigid. Laws take a long time to be discussed at the parliament/diet and once they are passed they stay the same for many years. However, there are other regulatory approaches other than imposing legislations top-down. One approach is to self-regulation by the stakeholders, or co-regulation combining top-down guidances (sometime called "nudging") and bottom-up self-regulations. Ikegai [13] argues that co-regulation is more flexible and faster to adapt to the environment change. This approach seems particularly useful for rapidly-changing landscape of Internet-based services and privacy/security related to them.

### 3.4 Active Resilience

So far we discussed three resilience strategies, namely, *redundancy*, *diversity*, and *adaptability*. These strategies do not require human intervention and appear in any resilient systems. We call these *passive resilience*.

Some systems such as human minds, financial systems, organizations, and social systems have human intelligence in their decision loop. We call this type of resilience strategy *active resilience*. Active resilience introduces another set of dimensions to our catalogue of resilience strategies.

#### 3.4.1 Anticipation

If we can *anticipate* a large scale event, we can prepare for it. WHO defines six phrases of pandemic alert. When avian flue H5N1 pandemic was a major threat in 2009, the global society at large responded based on the phase 4-6 declarations by WHO. Japan Meteorological Agency issues warnings on large scale natural events such as typhoons, volcanic activities, and tsunami. Accurate anticipation of such events is extremely hard, and it generally requires a lot of intelligence and computation. There are three different approaches to anticipation; *prediction*, *scenario planning*, and *simulation*.

Nate Silver [24] extensively discussed why predictions are so difficult. Some predictions, such as weather forecast, can be done based on the past statistical data, but the best predictions are usually based on combinations of a large amount of high-quality data on the past phenomena and the wisdom of human experts in the domain. More generally, Scheffer et al.[8] suggested that for any dynamical systems there could be early-warning signals that indicate the system is near a tipping point.

#### 3.4.2 Modeling

*Modeling* is another possibility. When a disaster occurs, collecting information, analyzing it, building models, and making plans based on these models are other areas that human intelligence can play key roles. SPEEDI (System for Prediction of Environmental Emergency Dose Information)[*2] is designed for this purpose. It collects information on major nuclear incidents and issues predictions based on there models (unfortunately this system was not effectively used when Fukushima nuclear power plant exploded).

#### 3.4.3 Emergency Response

*Business Continuity Planning (BCP)* is now a standard practice for enterprises. ISO 22320 defines requirements on the man-

[*2] See http://www.bousai.ne.jp/vis/torikumi/index0301.html

agement processes for emergency. It stresses the importance of empowering the employees in the bottom of the hierarchy who are dealing with the situation at first hand. They need to make tough decisions. They need to improvise. The management process should be designed so that their creativity is encouraged in emergency.

#### 3.4.4 Cognitive Errors

Human intelligence generally enhances the system's resilience. However, active resilience may introduce a new source of errors unique to human intelligence – *cognitive errors* [15]. People may overestimate the threat of certain types, such as terrorism, and may overreact.

#### 3.4.5 Consensus Building

How to recover from the shock usually requires *consensus building* among stake holders. After the 2011 earthquake and tsunami, Miyagi prefecture, the largest prefecture in the Tohoku area decided to rebuild a stronger industry base in the damaged area, whereas the people in Iwate prefecture, whose main industry is agriculture and fishery, decided to focus more on wellness of the residents than its economical success.

In general, a large perturbation may present an opportunity to scrap and re-build the system from scratch. But first we have to identify the stakeholders and ask for their consensus. This is also a unique aspect for active resilience.

#### 3.4.6 Switching Modes

Nasim Taleb discussed in his book *Black Swan* [27] that common statistics based on Gaussian distribution, mean values, and standard deviations etc. do not work for extreme events because these extreme events do not follow the familiar probability distributions. Many extreme events, such as earthquakes, are known to follow a power-law distribution, and depending on the parameter, a power-law distribution may not have a finite average value or a finite standard deviation. This means that we can not rely on insurance because insurance is based on the estimated average loss of multiple incidents.

A similar discussion goes to how high the sea walls must have been to prevent the damage caused by the 2011 tsunami. The Fukushima nuclear power plant disaster could have been prevented if the sea wall were 15m high instead of 5.7m. However, it is recorded that the Meiji Sanriku Tsunami was as high as 40m in the history. It is not practical to build such a high sea wall within which people live happily without worry.

Statistician Kei Takeuchi [26] argued that, for such extreme and rare events, it would be better to ignore these risks in the normal life. If you are lucky, you will never be a victim of such a disaster. You can live a happy life without too much worrying about the worst. On the other hand, if such disaster do happen, the society has to change its mode and get ready to help each other. Under these extreme circumstances, the social norm has to inevitably change, and the people need to accept the reality and try to recover.

We call this concept *mode switching*. In the *normal* mode, the system works within the designed realm and the system follows the designed set of policy, for example, pursuing maximum economic efficiency. If an extreme event happens and the system can no longer function as designed, the system switches its operational mode to the *emergency* mode, in which the system and the people behave based on a different set of policies (e.g., helping others). Maruyama et al.[11] discussesed the mode switching concept in the context of security policy in the

face of emergency.

## 4. The Model

The second step of our research is to build a mathematical model to study why these strategies work and in what conditions. We base our model on the framework of dynamic constraint satisfaction problems (DCSPs) [9],[28] and formally define the notion of *resilience* of open dynamic systems. In this paper, we present only a simplified verson of our model. Interested readers are referred to [18] for the general model.

### 4.1 Definition of resilience

The concept of resilience appears in various disciplines such as environmental science, materials science, sociology, and so on. Holling [12] first introduced the concept of resilience in ecology. Holling defined the resilience as the capacity of an ecosystem to respond to a perturbation or disturbance by resisting damage and recovering quickly.

We choose Bruneau [5]'s definition of seismic resilience for disaster prevention and formalize it to provide quantitative metrics for the resilience of a system. Bruneau considers a situation where a system's quality degrades abruptly at time $t_0$ due to some unexpected event and fully recover to the original state at time $t_1$, as shown in Figure 3. If we denote by $Q(t)$ the quality of the system at time $t$, the resilience of the system is measured as follows:

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

As the measured triangle area gets smaller, the system becomes more resilient. That is, there are two dimensions concerning the resiliency of the system.

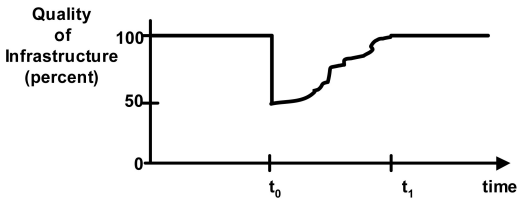


Fig. 3 Bruneau's definition of resilience.

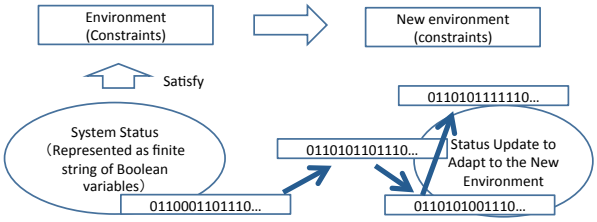


Fig. 4 Modeling as Dynamic Constraint Satisfaction Problem

- Resistance (reduced service degradation from failures at time $t_0$)
- Recoverability (reduced time to recovery (i.e., time interval between $t_0$ and $t_1$))

Although Bruneau points out that the notion of resilience is the combination of the above two properties, we only consider the recoverability aspect of the resilience to simplify our presentation in this paper.

### 4.2 Modeling as a Dynamic Constraint Satisfaction Problem

We assume that a system status can be described in a finite expression. Without loss of generality, a system status can be represented as a bit string of length $n$. At any given time, the system takes one of the $2^n$ possible configurations (see Figure 4).

A system operates in an environment. A system configuration may be *fit* against the environment. The fitness could be represented by a cost function over the set of all configurations. For simplicity, let us assume here that the cost function can be

represented as a subset $C$ of all fit configurations. A system configuration $s$ is said to be *fit* iff $s \in C$.

Suppose that there is an event (a shock) of type $D$ (say, earthquake of magnitude 7) and the environment changes from $C$ to $C'$. It is also possible for the system to change its state as a result of an event. If the current system configuration $s$ becomes unfit, that is, $s \notin C'$, the system needs to adapt to the new environment as quickly as possible by flipping some bits in $s$. One way to model this process is that the system flips one bit at a time. If the system can fix its configuration for any perturbations of type $D$ within $k$-steps, we call the system $k$-recoverable.

**Example:** We consider the hypothetical spacecraft system below. The system consists of a fixed set of $n$ components, each of which has a single binary variable $n_i$ representing the availability of the component as follows. The component $i$ is good if $n_i = 1$; otherwise, it is no good. The system's state is represented as a bit string of length $n$.

Suppose that the constraint $C = 1^n$ at every time $t$, which requires that every component of the spacecraft is good and that the spacecraft is occasionally hit by space debris causing at most $k$ component failures. Let $D$ be such a damaging event to the spacecraft.

If the spacecraft can fix one component at each time step, we consider that the spacecraft is $k$-recoverable under the presense of an event of type $D$ assuming that once the spacecraft has component failures at time $t$, it will not have another component failure until time $t + k$.

### 4.3 Reasoning uncertainty

Assuming that all possible states and events of a given system are known in advance, the notion of K-maintainability [4], which is similar to the notion of k-recoverability in Section 4.2, precisely defines the notion of resilience. We say that a system is K-maintainable if, for any non-normal state of the system, there exists a sequence of actions (i.e., events controllable by a system administrator) that move the system back to one of the normal states within $k$ steps. However, to analyze a system based on this definition requires us to know in advance all possible events, some of which could be totally unexpected. Therefore, it is not clear whether a model checking approach is applicable to evaluate the resiliency of a system with an incomplete specification. We, therefore, expect that reasoning techniques dealing with various uncertainty of a system model [7],[23] be a promising tool to explore this research space.

### 4.4 Evaluating tradeoffs

In general, there are tradeoffs among the strategies we discussed in Section 3. The available resource (e.g., budget) is limited. Should we invest our resource on redundancy, diversity, adaptability, or active resilience? Investing too much on redundancy by having n-way backup systems may delay the system update cycle and thus may hamper the adaptability for the business environment. What combination of resilience strategies is optimum under a given condition is one of the questions that we would like to answer in our project.

We plan to address that question using an evolutionary multi-agent system. Each agent in the system is a digital organism [29] that can self-replicate, mutate, or evolve, so that we can perform experiments on scale that are beyond reach with

any biological entity. We choose such a multi-agent system as our testbed since many complex systems in various domains, such as the Internet, have been implemented as a set of autonomous self-organized entities to achieve its resilience.

We quantify the three resilience properties of the system as follows. First, we consider the amount of a resource owned by an agent as the redundancy factor. An agent can remain alive until it uses up its resources even if it does not satisfy a constraint for a certain period. Second, we measure the diversity of a population consisting of multiple agents with the diversity index in Section 3.2.4. Third, we quantify the speed of an adaptation by the number of bits an agent can flip at a time.

Our focus is to identify key parameters that makes an agent population, which represents a decentralized complex system, resilient to a changing environment, by conducting various multi-agent simulations while changing the above system parameters.

### 4.5 Centralized vs. decentralized

Many complex systems consist of numerous components interacting with each other in a decentralized way, and to modularize a large system into smaller independent components seems to be a good design principle in order to contain a damage from a failure in a limited area.

However, Bak [2] shows that many decentralized systems that are modeled based on cellular automaton naturally reach a critical state with *minimum* stability without carefully choosing initial system parameters and that a small disturbance or noise at the critical state could cause *cascading* failures of the system leading to a large disaster, such as Northeast blackout of 2003.

Although subsequent research shows that there are various natural or artificial systems, such as earthquakes, DNAs, and stock exchange prices, which can be explained well with the notion of critical points, there is very little research about avoidance of critical points in complex systems. In ecological biology, to perform small destructions to an environment is known to improve the sustainability of the ecological system, and we might need to have such centrally coordinated interventions to a decentralized system in order to avoid critical points. We plan to investigate such tradeoffs between centralized and decentralized approach in the future.

## 5. Discussions

The concept of resilience is relatively new, and there are several issues that require our attention. Although there are no general consensus on these issues, certainly we should consider these issues when we design and operate resilient systems.

### 5.1 Types of Shock

First, we should consider what types of shock we should address in our resilient system design. Some types of shock, such as earthquakes, are known in the history and even their probabilistic distribution could be estimated. Other types of shocks, however, completely unexpected.

Also some shocks happen randomly and some are not. Barabasi [3] shows that network-based systems that possess the *scale-free* property are extremely robust against random failures of system components. However, when we consider a containment of a spreading virus that is deliberately designed to attack the hubs of the network, such connectivity becomes a vulnerability of the system. We need to investigate whether

there is a common property of resilience for various requirements or we need a way to dynamically switch a "resilience" mode of the system.

### 5.2 System Granularity

Conflict of resilience requirements among different levels of system granularity appears in many domains.

Take an biological system for an example. There are a few possible ways to consider the resiliency of the system. The most granular level would be the individual of a species. That individual has to survive against the environmental changes during its life span. Then there is the species level. Species can survive even if it loses some of its members during a perturbation. The most coarse level is the entire ecosystem as a whole. In this case, if at least one species survives, the system is considered to be resilient. So the definition of resilience should be relative to the granularity of the system. In general, the more coarse the system is, it is easier to make the system resilient. Although group evolution theory [19] provides a coherent theory for resolving such conflict in evolutionary biology, it may not be applicable to human societies. because this raises a moral question whether we are allowed to sacrifice individual lives to save the community. This is an open philosophical question.

### 5.3 Testing Resilience

In our resilience research we are primarily interested in unexpected and rare events. This means that, even if we build a system according to our resilience strategies, it is extremely difficult to prove that it is in fact resilient. The shock may be too large for the system to recover. Or, the shock never occurs in its entire life span.

A similar situation exists for computer security. When designing a computer system, the designer cannot predict what kind of attacks will be mounted on the system in the future. The general strategy for measuring the security of the system is two-fold. One is to evaluate the developing process to ensure that the developers take appropriate steps to take security considerations into the design. The more through the processes are, the system is considered more secure. The other is black-box testing, or testing by a so-called "tiger-team'.' In this approach, a group of highly skilled people try to attack the system.

Both ideas should be able to be applied to evaluating the resilience of a system, once we have better understanding of the types of shocks, the general strategies of resilience, and best tradeoffs among them.

## 6. Future Steps

Our catalogue of resilience strategies is by no means complete. This is an open-ended quest and we continue to look at more domains and extract insights from them.

Our first mathematical model of resilience is defined and now is ready to be applied to to the known strategies. We try to explain why these strategies work in which situations. Also we expect that the model can give some explanations to unsolved, open-questions in certain areas, such as why the ecosystem in the Antarctic Ocean is stable despite the fact that it is very simple (and less diverse).

Our *Systems Resilience* project is fundamentally transdisciplinary. We cannot achieve our goals without having collaboration with diverse fields, many of which are still beyond our radar scope. We are thus extending our network. If you are

interested in, please contact us.

**References**


[1] H. Akashi, N. Osada, and T. Ohta. Weak selection and protein evolution. *Genetics*, 192(15-31), 2012.

[2] Per Bak, Chao Tang, and Kurt Wiesenfeld. Self-organized criticality: An explanation of the $1/f$ noise. *Physical Review Letters*, 59:381–384, Jul 1987.

[3] Albert-Laszlo Barabasi and Eric Bonabeau. Scale-free networks. *Scientific American*, 288:50–59, 2003.

[4] Chitta Baral and Thomas Eiter. A polynomial-time algorithm for constructing k-maintainable policies. In *Proceedings of 14th International Conference on Automated Planning and Scheduling*, 2004.

[5] Michel Bruneau. A Framework to Quantitatively Assess and Enhance the Seismic Resilience of Communities. *Earthquake Spectra*, 19(4), 2003.

[6] John Casti. *X-Events: The Collapse of Everything*. William Morrow, 2012.

[7] Hei Chan and Adnan Darwiche. On the Revision of Probabilistic Beliefs Using Uncertain Evidence. 63:67–90, 2005.

[8] Marten Scheffer et. al. Early-warning signals for critical transitions. *Nature*, 461(3), 2009.

[9] Boi Faltings and Santiago Macho-Gonzalez. Open constraint programming. *Artificial Intelligence*, 161:181–208, 2005.

[10] H. Maruyama. Towards systems resilience. *Proc. of 1st Workshop on Systems Resilience (WSR)*, 2013.

[11] H. Maruyama, et al. Ichigan security – a security architecture that enables situatuion-based policy switching. *Proc. of 3rd International Workshop on Resilience and IT-Risk in Social Infrastructures (RISI 2013)*, 2013.

[12] Crawford Holling. Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics*, 4:1–23, 1973.

[13] N. Ikegai. Multiple structure of intermediary liability law and self-regulation: latest problems on the provider liability law. *Information Communication Policy Review*, 2, 2013.

[14] J. Kitano, et al. Reverse evolution of armor plates in the trheespine sticleback. *Current Biology*, 18, 2008.

[15] Daniel Kahneman. *Thinking, Fast and Slow*. Farra Straus & Giroux, 2013.

[16] J. Kephart and D. Chess. The vision of autonomic computing. *IEEE Computer*, 2003.

[17] M. Kimura. Evolutionary rate at molecular level. *Nature*, 217(624-6), 1968.

[18] N. Schwind, et al. Systems resilience: a challenge problem for dynamic constratint-based agent systems. *Proc. of the 12th Int. Conf. on Autonomous Agents and Multianget Systems (AAMAS)*, 2013.

[19] Martin A. Nowak. Five rules for the evolution of cooperation. 314:1560–1563, 2006.

[20] T. Ohta. The nearly neutral theory of molecular evolution. *Annual Review of Ecology and Systematics*, 23(1), 1992.

[21] Scott E. Page. *Diversity and Complexity*. Princeton University Press, 2010.

[22] Randy H. Katz Patterson, David; Garth A. Gibson. A case for redundant arrays of inexpensive disks (raid). *SIGMOD*, 1988.

[23] Chiaki Sakama and Katsumi Inoue. Abduction, unpredictability and garden of eden. In *Proceedings of Model-Based Reasoning in Science and Technology (MBR)*, 2012.

[24] Nate Silver. *The Signal and Noise: Why So Many Predictions Fail – but Some Don't*. Penguin Press HC, 2012.

[25] T. Baba, et. al. Construction of escherichia coli k-12 in-frame, single-gene knockout mutants: Keio collection. *Molecular Systems Biology*, 10, 2006.

[26] Kei Takeuchi. *What is Chance (Guuzen toha Nanika, in Japanese)*. Iwanami Publishing, 2010.

[27] Nassim Nicholas Taleb. *The Black Swan: The Impact of the*

*Highly Improbable*. Random House, 2007.

[28] Gérard Verfaillie and Narendra Jussien. Constraint solving in uncertain and dynamic environments: A survey. *Constraints*, 10(3):253–281, July 2005.

[29] Claus O. Wilke and Christoph Adami. The biology of digital organisms. 17:528–532, 2002.

**Hiroshi MARUYAMA**

Hiroshi Maruyama is a professor at Institute of Statistics in Japan. He is working on modeling engineering, cyber-physical systems, and service sciences. His team has been studying resilience properties of engineering systems by developing new modeling and simulation techniques. He worked at IBM Tokyo Research Laboratory for more than 25 years and published numerous papers on natural language processing, XML, web services, and information security. He is an author of best-selling book XML and Java: Developing Web Applications published in 2002. Before joining ISM in 2011, he served as a director of IBM Tokyo Research and a deputy group executive at Canon.

**Kazuhiro MINAMI**

Kazuhiro Minami is a research associate professor at Institute of Statistics in Japan. He has been working on security and privacy in pervasive computing particularly focusing on secure information sharing among mutually untrusted parties. He is currently working on new software engineering disciplines for developing resilient ICT systems, which could recover from disastrous situations in a flexible and graceful way. After finishing a Ph.D degree, he spent several years as a lecturer and a postdoctoral researcher at University of Illinois, and then took a position of an associate professor at National Institute of Informatics until March 2012.