

Systems Resilience: A Challenge Problem for Dynamic Constraint-Based Agent Systems*

Nicolas Schwind
National Institute of
Informatics
Tokyo, Japan
schwind@nii.ac.jp

Tenda Okimoto
Transdisciplinary Research
Integration Center
Tokyo, Japan
tenda@nii.ac.jp

Katsumi Inoue
National Institute of
Informatics
Tokyo, Japan
inoue@nii.ac.jp

Hei Chan
Transdisciplinary Research
Integration Center
Tokyo, Japan
hei@nii.ac.jp

Tony Ribeiro
The Graduate University
for Advanced Studies
Tokyo, Japan
tony_ribeiro@nii.ac.jp

Kazuhiro Minami
Transdisciplinary Research
Integration Center
Tokyo, Japan
kminami@ism.ac.jp

Hiroshi Maruyama
Institute of
Statistical Mathematics
Tokyo, Japan
hm2@ism.ac.jp

ABSTRACT

Many researchers of different fields are interested in building resilient systems that can absorb shocks and recover from damages caused by unexpected large-scale events; however, no common agreement on the definition of resilience exists. In this paper, we set out to establish a new challenging research discipline that we call “systems resilience”, which provides a set of unified design principles for building resilient systems. We define a dynamic constraint-based agent model called SR-model, whose resilience can be evaluated through a range of properties.

Categories and Subject Descriptors

G.1.6 [Optimization]: Constrained optimization; I.2.11 [Distributed Artificial Intelligence]: Intelligent agents

Keywords

Systems resilience, Dynamic system

1. INTRODUCTION

In recent years, many researchers of different fields have recognized the importance of a new research discipline concerning the *resilience* of complex agent systems, which can be informally defined as the ability “to maintain its core purpose and integrity in the face of dramatically changed circumstances” [4, 2, 7]. The reason for this is that unexpected events in terms of time and scale (such as the 3.11 earthquake in Japan, the global economic crisis, or a new strain of virus) may cause irreversible damages to the core functionality of these agent systems. The goal of resilience research is to provide a set of general principles for building resilient systems in various domains, such that the systems are resistant from large-scale perturbations caused by unexpected events and changes, and if their functionality is lost temporarily due to outside forces, the systems can recover gracefully and quickly to restore their functionality in the long run.

The concept of resilience has appeared in various disciplines such as environmental science, materials science, sociology, and

*This work is supported in part by the “Systems Resilient” project of Transdisciplinary Research Integration Center.

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, Ito, Jonker, Gini, and Shehory (eds.), May 6–10, 2013, Saint Paul, Minnesota, USA. Copyright © 2013, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

so on. Holling [4] first introduced the concept of resilience in ecology, and defined the resilience as the capacity of an ecosystem to respond to a perturbation or disturbance by resisting damage and recovering quickly. Bruneau’s [2] definition of seismic resilience for disaster prevention elaborates the concept of the resilience by introducing quantitative measures, which compute the “triangular” area of the degradation of the functionality of the system over time. In Artificial Intelligence, similar concepts to resilience, such as stabilizability [6] and maintainability [1] have also been proposed by other researchers within the structure of Discrete Event Systems.

However, while we have seen many examples of seemingly resilient systems in various fields, researchers have not agreed on a common definition on resilience among the different domains yet. The challenging research topic that we call here “systems resilience” is to provide a set of unified design principles for building resilient systems. Our first step is to define a novel system model, which we will call the *SR-model*.

The significant aspects of our SR-model are as follows. The definition of the SR-model allows the system to change dynamically over time, such that the variables, domains, constraints, and configurations of the system can evolve based on the decisions made by agents and/or outside environmental events. The flexibility of our SR-model allows the modeling of the dynamicity of systems that is required in many domains. In comparison, Constraint Satisfaction Problems (CSPs) have been applied traditionally in closed-world scenarios, where all choices and constraints are known from the beginning and fixed. A notion of Open CSP (OCSP) was investigated in [3]. In OCSP, the set of variables is closed but the domains are open, that is, extra values can be added to variable domains. However, it is more appropriate for handling dynamic real-world problems that all possible changes to a CSP can be expressed, i.e., add/remove constraints, domains and variables.

Then, given a configuration trajectory (i.e., a chain of assignments of values over the variables of a dynamic system over time), our SR-model enables us to measure four important properties that are central to the idea of resilience:

- *Resistance*: The ability to maintain some underlying costs under a certain “threshold”, such that the system satisfies certain hard constraints and does not suffer from irreversible damages.
- *Recoverability*: The ability to recover to a baseline of acceptable quality as quickly and inexpensively as possible.
- *Functionality*: The ability to provide a guaranteed average degree of quality for a period of time.

• *Stabilizability*: The ability to avoid undergoing changes that are associated with high transitional costs.

A dynamic system is resistant, recoverable, functional, or stabilizable if one can find a “strategy” and a configuration trajectory within this strategy that is resistant, recoverable, functional, or stabilizable. Following the concept of resilience in most related literature, for our SR-model we define a resilient dynamic system as one that satisfies all properties of resistance, recoverability, functionality, and stabilizability¹.

Next we will define our SR-model, and show an example of how it can be used to evaluate the robustness of a system, and afterwards discuss open questions that are important for future research.

2. SR-MODEL

In the following, we are given a finite set of variables $\mathcal{X} = \{x_1, x_2, \dots\}$ and a finite domain \mathcal{D} .

DEFINITION 1 (SYSTEM, CONFIGURATION). A system S is a tuple $\langle X, dom, c \rangle$ where $X \subseteq \mathcal{X}$ is the set of variables of S , dom is a domain function, that is a mapping associating every variable $x_i \in X$ with its domain $dom(x_i) \subseteq \mathcal{D}$, and c is a cost function, that is a mapping from $\Omega(S)$ to $\mathbb{R}^+ \cup \{+\infty\}$, where $\Omega(S) = \prod\{dom(x_i) \mid x_i \in X\}$.

An element of $\Omega(S)$ is called a configuration of S . Given a configuration α of S , for every variable $x_i \in X$ we denote $\alpha(x_i)$ the value of x_i in α . A configuration α of S is said to be optimal if it minimizes the cost function, i.e., if for every configuration α' of S , $c(\alpha) \leq c(\alpha')$. \mathcal{S} denotes the set of all possible systems. Ω denotes the set of all configurations of all possible systems, i.e., $\Omega = \bigcup\{\Omega(S) \mid S \in \mathcal{S}\}$.

The definition of a system above is similar to the one of a Distributed Constraint Optimization Problem (DCOP [5]), a fundamental problem that can formalize various applications related to multi-agent cooperation. A DCOP consists of a set of agents, each of which needs to decide the value assignment of its variables so that the sum of the resulting costs is minimized.

EXAMPLE 1. We represent the economic situation of country J which is composed of three inhabited islands. Let S^J be the system $\langle X, dom, c \rangle$ representing the country. $X = \{x_1, x_2, x_3\}$ represents the set of the three islands. We set $dom(x_1) = dom(x_2) = dom(x_3) = \{Low, Medium, High\}$ the domain of each variable, where each one of these values represents the level of expenditures the country invests for each island. For the sake of simplicity, we represent a configuration of S^J as a chain of three letters among $\{L, M, H\}$; for instance, $\alpha = MML$ represents the fact that the country spends a normal amount of money for the two first islands, while a low amount of money is invested in the third island. We define a reference cost function c_{ref} for every configuration α of Ω as $c_{ref}(\alpha) = \sum_{i=1}^3 local_{c_{ref}}(\alpha(x_i))$, where for every $x_i \in X$, $local_{c_{ref}}(\alpha(x_i)) = 0$ if $\alpha(x_i) = L$, 1 if $\alpha(x_i) = M$, 3 if $\alpha(x_i) = H$. Then, c_{ref} represents the global amount of money spent by the country for the group of the three islands. For instance, $c_{ref}(MML) = 2$. Now, assume that our system is ruled by the fact that at least “normal” expenses have to be invested in at least two islands among the group. Then, the cost function c of S^J is defined as follows, for every $x_i \in X$,

$$c(\alpha) = \begin{cases} +\infty & \text{if } |\{\alpha(x_i) = L, x_i \in X\}| \geq 2, \\ c_{ref}(\alpha) & \text{otherwise.} \end{cases}$$

¹In [7] the concept of resilience is also captured through a set of properties, though the properties we introduce here are slightly different and more adapted to our SR-model.

We now define the notions of dynamic system and strategy within a dynamic system.

DEFINITION 2 (DYNAMIC SYSTEM). A dynamic system DS is a tuple $\langle S_0, \mathcal{A}, poss, \Phi_A \rangle$, where:

- S_0 , where $S_0 \in \mathcal{S}$, is the “initial” system, which represents the specifications of the current system.
- \mathcal{A} is a non-empty set of actions (or moves),
- $poss$ is a mapping from \mathcal{S} to $2^{\mathcal{A}}$ which gives a set of possible moves for each system, and such that for every system $S \in \mathcal{S}$, $poss(S) \neq \emptyset$,
- Φ_A is a partial function from $\mathcal{S} \times \mathcal{A}$ to $2^{\mathcal{S}}$ such that for every system $S \in \mathcal{S}$, for every move $a \in poss(S)$, $\Phi_A(S, a)$ relates to some non-empty subset of systems from $2^{\mathcal{S}}$. Φ_A specifies how a given system may change in response to some moves.

A dynamic system can be represented as a graph where each vertex represents a system and each edge represents the (potential) consequence of some move that would transform the current system into another one (e.g., add/remove some variables or change the cost function.) Remark that Φ_A is a non-deterministic function, in the sense that from a specific move taken in a given system, one may fall into several possible systems at the next step. This non-determinism is due to the presence of exogenous events that may alter the result of a move².

DEFINITION 3 (STRATEGY). Given a dynamic system $DS = \langle S_0, \mathcal{A}, poss, \Phi_A \rangle$, a strategy within DS is a dynamic system $\langle S_0, \mathcal{A}, strat, \Phi_A \rangle$ such that for every system $S \in \mathcal{S}$, $strat(S) \subseteq poss(S)$ and $|strat(S)| = 1$.

Here, $strat$ describes a specific “control policy” that is adopted by agents, i.e., which specific move is planned within each system.

EXAMPLE 1 (CONTINUED). We denote DS^J as a dynamic system that represents all possible scenarios about the evolution of country J . Assume that the system S^J represents the current system, i.e., $S_0 = S^J$. Φ_A is derived from a study of the potential consequences of some possible exogenous events that could occur in the country at any time step (e.g., a tsunami in one of the islands, or some new norms set by the government), given the decision taken by the agents. These moves (captured in \mathcal{A}) depend on the given states of the dynamic system (i.e., on each system), since the corresponding moves could represent some restructuring plans that are conducted regarding some “costly” systems resulting from an exogenous event (e.g., when a disaster occurred in one of the islands and the country has to invest more money to repair the related damages.)

We now introduce the notions of system trajectory and configuration trajectory.

DEFINITION 4 (SYSTEM TRAJECTORY). Given a dynamic system $DS = \langle S_0, \mathcal{A}, poss, \Phi_A \rangle$, a system trajectory ST of DS is a (possibly infinite) sequence of systems (S_0, S_1, \dots) of \mathcal{S} , such that for every $i \in \{1, 2, \dots\}$, there is a move $a \in poss(S_{i-1})$ such that $S_i \in \Phi_A(S_{i-1}, a)$.

Accordingly, a system trajectory $ST = (S_0, \dots)$ of a dynamic system $\langle S_0, \mathcal{A}, poss, \Phi_A \rangle$ represents one possible scenario where every $S_t \in ST$ represents a “system snapshot” at time step t . Every time step could represent a day, a month, or any other frame of time depending on the application. At each time step $t \in \{1, \dots\}$, S_t is the system resulting from one of the possible consequences of a move performed in the system S_{t-1} (i.e., at time step $t - 1$.)

²In our definition the exogenous events are not explicitly represented, but their possible consequences lie in the non-determinism of Φ_A in which they take a fundamental part.

DEFINITION 5 (CONFIGURATION TRAJECTORY). Given a system trajectory $ST = (S_0, S_1, \dots)$, a configuration trajectory CT of ST is a sequence $(\alpha_0, \alpha_1, \dots)$ such that for every $i \in \{0, 1, \dots\}$, α_i is a configuration of S_i . A subtrajectory CT' of CT is a subsequence $(\alpha_a, \dots, \alpha_b)$ of CT , with $a \leq b$.

EXAMPLE 1 (CONTINUED). Suppose that in our dynamic system DS^J , the image of Φ_A is a set of systems that consider the same set of variables X and the same domain function dom (only the cost function may change), i.e., the country will never refine the possible values of expenditures $\{Low, Medium, High\}$, and no island will be added to or removed from the country. Then, the sequence $(HHH, HMM, HLM, LLM, MLM)$ is an example of a configuration trajectory of any system trajectory of DS^J .

We are now ready to introduce the properties that we believe are relevant for the characterization of resilient dynamic systems. We first introduce the properties on configuration trajectories, then extend these properties to the structures of system trajectory and dynamic system. In the following definitions, we suppose that we are given a configuration trajectory CT of a system trajectory ST .

DEFINITION 6 (RESISTANCE). Given a non-negative integer l , CT is said to be l -resistant if for each $i \in \{0, 1, \dots\}$, $c_i(\alpha_i) \leq l$. Intuitively, a configuration trajectory is l -resistant if the cost of each one of its configurations is kept under the threshold l .

EXAMPLE 1 (CONTINUED). Let $l = 7$ be the threshold above which the country does not have the capability to handle the situation due to lack of money. Figure 1 depicts a configuration trajectory that is not 7-resistant. It represents for instance the scenario where a big disaster (e.g., a tsunami) occurred in the three islands at time step 0, such that it strongly constrains the system S_1 so that no configuration of it can be found with a cost below 7.




Figure 1: A configuration trajectory that is not 7-resistant.

DEFINITION 7 (RECOVERABILITY). Given two non-negative integers p and q , CT is said to be $\langle p, q \rangle$ -recoverable if for each one of its subtrajectory $(\alpha_a, \dots, \alpha_b)$ such that for every $i \in \{a, \dots, b\}$, $c_i(\alpha_i) > p$, we have $\sum_{i=a}^b c_i(\alpha_i) \leq q$.

Intuitively, q represents the total amount of extra cost (i.e., costs above p) that is necessary for a $\langle p, q \rangle$ -recoverable configuration trajectory to get back to a “safe” state. This cumulative extra cost is similar to the “triangular” area of the degradation of the functionality of the system over time in Bruneau’s definition of seismic resilience [2].

EXAMPLE 1 (CONTINUED). Figure 2 shows an example of a $\langle 3, 6 \rangle$ -recoverable configuration trajectory. It represents for instance the scenario where a major disaster affects two islands (the second and the third one) at time step 0. At time step 1, the second island recovers from the damages and changes its state to “Medium”, and at time step 2 the third island recovers as well in the same




Figure 2: A $\langle 3, 6 \rangle$ -recoverable configuration trajectory.

way. In this scenario the configuration got back to a cost of 3 by accumulating a total extra cost which does not exceed 6, therefore this configuration trajectory is $\langle 3, 6 \rangle$ -recoverable.

Figure 3 depicts a configuration trajectory that is not $\langle 3, 6 \rangle$ -recoverable. This scenario is the same as the one depicted in Figure 2 until time step 2, when at the same time the third island of the group recovers from the damages, and there is a new disaster in the first island. The state of the first island changes to the value “High”, so that the cost of the configuration is 5. In this scenario the country needs one more time step to recover to a cost of 3. The cumulated extra cost is 8 in this scenario, i.e., it exceeds 6, therefore such configuration trajectory is not $\langle 3, 6 \rangle$ -recoverable.




Figure 3: A configuration trajectory that is not $\langle 3, 6 \rangle$ -recoverable.

Let us stress the fact that the parameters p and l play different roles in the definition of resilience. Indeed, a configuration whose cost is below p will be considered as being a “safe” state of the system; between p and l , it will fall into an “unsafe” state, an acceptable situation that still needs to be “recovered”; above l , the situation is considered to be “unacceptable” for the current system. We now introduce the notions of functionality and stabilizability.

DEFINITION 8 (FUNCTIONALITY). Given two non-negative integers f and k , assuming CT has a size of at least k , then CT is said to be f -functional over k if $\sum_{i=0}^k \frac{c_i(\alpha_i)}{k+1} \leq f$.

The notion of functionality is important in the sense that it provides a guaranteed average degree of “quality” for the configuration trajectory. Indeed, stating that a configuration trajectory is f -functional over k means that its average cost is kept under f from the initial step until k time steps.

EXAMPLE 1 (CONTINUED). The configuration trajectory depicted in Figure 2 is 4-functional over 5, while the one depicted in Figure 1 is 8-functional over 5.

We now assume that the set Ω is together with a premetric denoted d , that is, d is a mapping from $\Omega \times \Omega$ to \mathbb{R}^+ such that for every $\alpha \in \Omega$, $d(\alpha, \alpha) = 0$. d allows us to represent a cost that stands for passing from a configuration to another one. d is said to be a *transitional cost function* on Ω .

DEFINITION 9 (STABILIZABILITY). Given a non-negative integer s , CT is said to be s -stabilizable w.r.t. CT if for every $i \in \{1, 2, \dots\}$, $d(\alpha_{i-1}, \alpha_i) \leq s$.

EXAMPLE 1 (CONTINUED). Assume that some specific configuration transitions lead to some extra amount of costs (to varying degrees) for DS^J , such as when the country cuts its investment in an island (i.e., when a variable changes its value to “Low”) or when the country invests a high amount of money in another island during two consecutive time steps. Therefore, we define d as follows, for every configuration $\alpha, \alpha' \in \Omega$ as $d(\alpha, \alpha') = \sum_{i=1}^3 local_d(\alpha(x_i), \alpha'(x_i))$, where for every $x_i \in X$,

$$local_d(\alpha(x_i), \alpha'(x_i)) = \begin{cases} 6 & \text{if } (\alpha(x_i), \alpha'(x_i)) = (H, L), \\ 4 & \text{if } (\alpha(x_i), \alpha'(x_i)) \in \{(M, L), \\ & \quad (H, H)\}, \\ 2 & \text{if } (\alpha(x_i), \alpha'(x_i)) = (L, L), \\ 0 & \text{otherwise.} \end{cases}$$

We get for instance, $d(HML, HLL) = 0 + 4 + 2 = 6$.

Figure 4 represents two configuration trajectories: $\alpha = (MLL, HMH, MLH, HLL, MMM)$ and $\alpha' = (MLL, HMH, MMH, HMM, MMM)$. The transitional costs are given between each pair of successive configurations. One can see that α' is 4-stabilizable while α is not. This example shows that the concepts of stabilizability and recoverability are independent, since α is $\langle 3, 6 \rangle$ -recoverable while α' is not.




Figure 4: Two configuration trajectories that are not 3-stabilizable.

We are now ready to extend the properties introduced above to system trajectories and dynamic systems:

DEFINITION 10. Let \mathcal{P} a subset of properties from $\{l$ -resistant, $\langle p, q \rangle$ -recoverable, f -functional over k , s -stabilizable}. A system trajectory ST satisfies the set \mathcal{P} of properties if there is a configuration trajectory of ST that satisfies all properties from \mathcal{P} . A dynamic system DS satisfies the set \mathcal{P} of properties if there is a strategy of DS whose all system trajectories satisfy the set \mathcal{P} of properties.

Note that if a dynamic system satisfies two disjoint subsets $\mathcal{P}_1, \mathcal{P}_2$ of properties from \mathcal{P} , it does not necessarily satisfy their union $\mathcal{P}_1 \cup \mathcal{P}_2$. This is due to the fact that the properties from \mathcal{P}_1 on the one hand, from \mathcal{P}_2 on the other hand could be satisfied by two different strategies of the dynamic system. Similarly, a system trajectory may satisfy the sets \mathcal{P}_1 and \mathcal{P}_2 of properties independently and may not satisfy the set $\mathcal{P}_1 \cup \mathcal{P}_2$. Therefore, a dynamic system satisfying the full set of properties $\mathcal{P} = \{l$ -resistant, $\langle p, q \rangle$ -recoverable, f -functional over k , s -stabilizable} is more “robust” than a dynamic system satisfying each singleton subset of \mathcal{P} separately. We are now ready to provide our definition of a resilient dynamic system:

DEFINITION 11 (RESILIENT DYNAMIC SYSTEM). Given non-negative integers l, p, q, f, k, s , a dynamic system is said to be $\langle l, p, q, f, k, s \rangle$ -resilient if it satisfies the full set of properties $\{l$ -resistant, $\langle p, q \rangle$ -recoverable, f -functional over k , s -stabilizable}.

3. OPEN QUESTIONS

In this paper, we introduced the topic of systems resilience, and defined a new model called SR-model, which can be used to represent a dynamic constraint-based agent model. We captured the notion of resilience for dynamic systems using several factors, resistance, recoverability, functionality and stabilizability. We believe that our SR-model can provide a unified design principle for building resilient systems across different domains.

We now present open questions that are important for the future extension of our SR-model. First, we can associate a decision problem on dynamic systems for every subset \mathcal{P} of properties introduced above, though our main interest lies in the design of resilient dynamics systems: Given a dynamic system DS and non-negative integers l, p, q, f, k, s , is $DS \langle l, p, q, f, k, s \rangle$ -resilient? It is a novel and challenging problem. While solving a DCOP [5] consists of finding an optimal configuration that minimizes the cost function (an NP-hard problem in general), evaluating the resilience of a dynamic system is a more subtle problem. Indeed, some non-optimal configuration trajectories would satisfy all the properties we introduced while some optimal ones would not.

Another set of important problems for our SR-model is optimization problems. For example: (i) If our dynamic system must satisfy $\{l$ -resistant} and we are now given a cost threshold p , what is the minimum q such that we can guarantee the dynamic system to satisfy $\{l$ -resistant, $\langle p, q \rangle$ -recoverable}? (ii) If our dynamic system must satisfy $\{f$ -functional over k , s -stabilizable} how “resilient” can it be, i.e., what are the minimal values of l, p, q such that the system is $\langle l, p, q, f, k, s \rangle$ -resilient? These questions are important for the design of resilient dynamic systems when we must make tradeoffs between some properties.

As a future work, we will investigate the connections between the properties that feed our SR-model and other theories such as classical control theory and modal logics. Moreover, in the current version of our SR-model, we assume that we have a complete knowledge on all past and current configurations in the dynamic system. In reality, we may only have uncertain information on some of these configurations. Therefore, models for the probabilistic reasoning on dynamic systems, similar to those such as hidden Markov models and dynamic Bayesian networks, may need to be incorporated into our SR-model.

4. REFERENCES

- [1] C. Baral, T. Eiter, M. Bjärelund, and M. Nakamura. Maintenance goals of agents in a dynamic environment: Formulation and policy construction. *Artificial Intelligence*, 172(12-13):1429–1469, 2008.
- [2] M. Bruneau. A Framework to Quantitatively Assess and Enhance the Seismic Resilience of Communities. In *Earthquake Spectra*, volume 19, 2003.
- [3] B. Faltings and S. Macho-Gonzalez. Open constraint programming. *Artificial Intelligence*, 161:181–208, 2005.
- [4] C. Holling. Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics*, 4:1–23, 1973.
- [5] P. J. Modi, W.-M. Shen, M. Tambe, and M. Yokoo. ADOPT: asynchronous distributed constraint optimization with quality guarantees. *Artificial Intelligence*, 161(1-2):149–180, 2005.
- [6] C. M. Ozveren, A. S. Willsky, and P. J. Antsaklis. Stability and stabilizability of discrete event dynamic systems. *Journal of the ACM*, 38(3):7300–7752, 1991.
- [7] B. H. Walker, C. S. Holling, S. C. Carpenter, and A.P. Kinzig. Resilience, adaptability and transformability. *Ecology and Society*, 9(2): 5, 2004.