

Assured System-Level Resilience for Guaranteed Disaster Response

Melkior Ornik

Dept. of Aerospace Eng. and Coordinated Science Laboratory
University of Illinois Urbana-Champaign
Urbana, USA
mornik@illinois.edu

Jean-Baptiste Bouvier

Dept. of Aerospace Eng.
University of Illinois Urbana-Champaign
Urbana, USA
bouvier3@illinois.edu

Abstract—Resilience of urban infrastructure to sudden, system-wide, potentially catastrophic events is a critical need across domains. The growing connectivity of infrastructure, including its cyber-physical components which can be controlled in real time, offers an attractive path towards rapid adaptation to adverse events and adjustment of system objectives. However, existing work in the field often offers disjoint approaches that respond to particular scenarios. On the other hand, abstract work on control of complex systems focuses on attempting to adapt to the changes in the system dynamics or environment, but without understanding that the system may simply not be able to perform its original task after an adverse event. To address this challenge, this programmatic paper proposes a vision for a new paradigm of infrastructure resilience. Such a framework treats infrastructure across domains through a unified theory of controlled dynamical systems, but remains cognizant of the lack of knowledge about the system following a widespread adverse event and aims to identify the system’s fundamental limits. As a result, it will enable the infrastructure operator to assess and assure system performance following an adverse event, even if the exact nature of the event is not yet known. Building off ongoing work on assured resilience of control systems, in this paper we identify promising early results, challenges that motivate the development of resilience theory for infrastructure system, and possible paths forward for the proposed effort.

Index Terms—Smart infrastructure, resilience, uncertain systems, network systems

I. VISION AND PRIOR WORK

Resilience to catastrophic events is a crucial infrastructural challenge, recognized across communities and government levels [1]–[3]. Natural disasters, terrorist acts, and far-reaching system failures all have the potential to rapidly deteriorate the infrastructure’s capabilities to meet population needs, as well as to change those needs. Inability to adapt to such events in real time also impedes emergency services, evacuation, and supply chain operations. The need for resilience — already amply displayed during natural disasters in past decades [4]–[6] — is made stronger with growing reliance on cyber-physical systems for control of infrastructure. For instance, a team of security researchers recently demonstrated the capability to influence traffic signals over the internet in at least ten cities [7], thus causing disorder in traffic flows over a large area.

While the natural approach is to use standard cyber-physical safety and security measures to try to protect each network

component from failure [8], [9], *perfect protection is impossible* in the face of the number of connected systems in a smart city. Extreme natural events will inevitably degrade the capabilities of a part of the overall infrastructure [10], and — with increasing numbers of autonomous components in interconnected transportation, electric, and telecommunications networks — cyberattacks will focus on identifying network components vulnerable to hijacking [11]. On the other hand, autonomous control over infrastructure offers exceptional agility and adaptability, enabling rapid restoration of satisfactory performance even in situations where physical infrastructure is significantly degraded, e.g., by using traffic signals to automatically redirect traffic flow after damage to an area. A natural goal — already identified in prior work [12], [13] — is thus to develop methods of exploiting continued control authority to drive the system to perform as well as it can under the changed circumstances.

While applied work on ensuring resilience in infrastructure often deals with specific scenarios [4], [12], ensuring continued system performance in off-nominal situation is a classical problem subject to substantial research in domain-agnostic control theory and planning. Namely, the methods of robust and adaptive control [14] have been previously applied to smart city infrastructure [15], [16]. However, these methods often suffer from several drawbacks:

- they are only able to guarantee adaptation to limited disturbances — a bounded change in the system dynamics or a change in a finite number of parameters,
- they assume specific mathematical structure of the system — e.g., linearity — as well as of available knowledge about the system after an adverse event,
- they aim to control the system to achieve *the same level of performance* as prior to the adverse event, not recognizing that the system may simply no longer be capable of doing so using *any* control strategy.

These drawbacks are crucial when considering the problem of resilience of infrastructure to sudden, potentially catastrophic, adverse events. Namely, events such as natural disasters over a large area and hostile acts will *fundamentally* change the system, substantially reduce the knowledge about it immediately following the adverse event, and likely disable

the system from performing all its usual tasks, prompting a need to determine its new capabilities in real time.

Instead of classical robustness and adaptation methods, this paper proposes a less brittle framework of *resilience and guaranteed performance*, illustrated in Fig. 1. Its objective is to answer the following questions:

- 1) Can we *a priori guarantee* system resilience, i.e., prove that the system will be able to continue with its task after an adverse event, even if the exact nature of the event is not known in advance?
- 2) If so, how can the system continue with task completion after such an event, even if the consequences of the event are not entirely immediately known?
- 3) If not, what are the tasks that the system is certifiably able to achieve, given the lack of knowledge about the event's consequences?



Fig. 1. A graphical representation of the proposed framework.

In subsequent sections, we describe our initial work on this framework, investigate an illustrative scenario of shelter assignment, and identify a plethora of research challenges to be surmounted before such a paradigm can be widely applied and adapted to smart city infrastructure.

II. RESILIENCE AND GUARANTEED PERFORMANCE

Any infrastructure system is naturally a combination of static and dynamic components [17], [18]. The former consist of physical objects and are nominally — barring significant adverse events — fixed, at least over short time scales. The latter consist of time-varying functions of system users, environmental conditions, and protocols for control over the system. To take an example of a transportation network, the static components are the roads, speed limit signs, and positions of traffic signals. The dynamic features are the number of cars at every intersection at any given time, precipitation, and the traffic signal lights at a given time.

If x denotes all elements relevant to the system's mission success (i.e., *system state*), ordinary differential equation

$$\dot{x}(t) = \frac{dx(t)}{dt} = f(x(t), u(t), t), \quad (1)$$

provides a widely used way to describe the change in those features over time [13], [15]. It states that this change is described by *dynamics* f which depend on the system state

itself, the *control authority* u that can be used to influence the system, and time, and can be naturally appended by stochastic components or time-delay.

An adverse event will cause dynamics (1) to change to

$$\dot{x}(t) = \bar{f}(x(t), v(t), w(t), t), \quad (2)$$

with inputs (v, w) which model a possible partial loss of control authority. Namely, v describes inputs that are still under the controller's authority, while w describes inputs which the controller cannot influence. Function \bar{f} models the new system dynamics, which may be partially unknown to the controller.

To assess the system's capability to respond to adverse events, we are interested in pursuing two broad problems.

Problem 1 (Resilience): Given the control system described by (1), *a priori* determine whether it will be able to complete a task following an adverse event, *without knowing* the event's exact nature and behavior.

Naturally, we also want to determine conditions for a system to be resilient *regardless* of the exact nature of its task, enabling *resilient system design*. However, if the system is not resilient, we want to understand its *fundamental capabilities*.

Problem 2 (Guaranteed Performance): Determine in real time the set of tasks that a control system is *guaranteed* to be able to complete after an unexpected event, *given incomplete knowledge* of (2) at the time of the event. Consequently, generate an appropriate control signal that enables the system to complete the chosen task.

While problems of resilience and guaranteed performance can be posed mathematically rigorously — possibly also including tasks with a time-budget to complete them — they are challenging even for simple control systems. Yet, initial work [19]–[23] in domains different from infrastructure — and under often-restrictive technical assumptions — yielded preliminary success by using methods of geometric and optimal control theory. We now introduce the first foray of this theory into infrastructure planning.

III. ILLUSTRATIVE APPLICATION: SHELTER ASSIGNMENT

Consider a scenario — investigated in detail in [24] — where, due to a natural disaster or hostile action, urban population needs to evacuate to one of several community shelters. For this illustration, the shelters are in a central location and, except for a small fraction of the population that moves to each shelter directly, the majority of the population first approaches one of several checkpoints where they are directed to an appropriate shelter.

We assume that the population arrives to checkpoints and shelters at a steady pace, and that the population assigned to the shelter is immediately counted among that shelter's population. Assuming that there are N shelters and K checkpoints, and denoting the population count of shelter i at time t by $x_i(t)$, system dynamics are given by

$$\dot{x}_i(t) = \alpha_{0i} + \sum_{j=1}^K \alpha_j u_{ji}(t), \quad (3)$$

where α_{0i} is the rate of arrival directly to shelter i without previously going to a checkpoint, $u_{ji}(t)$ is the fraction of the population at checkpoint j that gets directed to shelter i at time t , and α_j is the rate of arrival at checkpoint j . Naturally, $u_{ji}(t) \in [0, 1]$ with $\sum_{i=1}^N u_{ji}(t) = 1$ for all j .

The objective is to fill each shelter to its capacity C_i ; the task fails if an individual is assigned to an overpopulated shelter without filling all shelters first. Given that x_i is increasing for all i , this goal is one of *reachability*: ensuring that there exists time T such that $x(T) = (C_1, \dots, C_N)$.

Determining whether an appropriate control policy (u_{11}, \dots, u_{KN}) for (3) exists, and computing it if it does, is trivial. However, as a consequence of the adverse event, one of the checkpoints is not operating properly: its decisions are not coordinated with other checkpoints, are likely not optimal, and might even be adversarial to the system’s objective. We wish to determine whether the task can be completed *regardless of* what happens at the “rogue” checkpoint.

To illustrate the method of resilience verification, we consider a scenario with $N = 2$ and $K = 3$, shown in Fig. 2.




Fig. 2. An instance of the shelter assignment problem. Blue checkpoints are under central authority and their shelter assignments are controlled. Orange checkpoint is not controlled and is potentially adversarial. Green arrows indicate the arrival of population to checkpoints, as well as directly to one of the shelters. The rate of these arrivals is known and cannot be controlled.

We assume that shelters have the same capacity $C_1 = C_2 = 500$ and the checkpoints have the same arrival rates $\alpha_j = 1$. Additionally, one of the shelters also receives a part of the population directly: $\alpha_{01} = 0$, $\alpha_{02} = 1$. Recalling that $u_{j2} = 1 - u_{j1}$ for all j , the system dynamics are thus

$$\dot{x}(t) = \begin{pmatrix} 0 \\ 4 \end{pmatrix} + \begin{pmatrix} 1 & 1 & 1 \\ -1 & -1 & -1 \end{pmatrix} \begin{pmatrix} u_{11} \\ u_{21} \\ u_{31} \end{pmatrix}. \quad (4)$$

We assume that the third checkpoint is malfunctioning. By introducing a “dummy state” x_3 with $\dot{x}_3 = 0$ and taking $v_1 = 2u_{11} - 1$, $v_2 = 2u_{21} - 1$ and $w = 2u_{31} - 1$, dynamics (4) take a linear form

$$\begin{aligned} \dot{x} &= Ax + B_1 v + B_2 w \\ &= \begin{pmatrix} 0 & 0 & 1.5 \\ 0 & 0 & 2.5 \\ 0 & 0 & 0 \end{pmatrix} x + \begin{pmatrix} 0.5 & 0.5 \\ -0.5 & -0.5 \\ 0 & 0 \end{pmatrix} v + \begin{pmatrix} 0.5 \\ -0.5 \\ 0 \end{pmatrix} w, \end{aligned} \quad (5)$$

with $x(0) = (0, 0, 1)$ and $v_1, v_2, w \in [-1, 1]$. Now comes the technical key to the solution, used more generally in [22]. Following [25], subspace $\mathcal{E} = \{\alpha(x_{e1}, x_{e2}, 1) \mid \alpha \in \mathbb{R}\}$ is reachable from $x(0)$ with dynamics (5) for any input signal w if and only if \mathcal{E} is reachable from $x(0)$ with dynamics $\dot{x} = Ax + z$, where control input $z \in \mathcal{Z} = B_1[-1, 1]^2 \cap (B_1[-1, 1]^2 \ominus B_2[-1, 1])$, with $B\mathcal{Y} = \{By \mid y \in \mathcal{Y}\}$ and symbol \ominus denotes the Pontryagin difference [26]. Also, because of $\dot{x}_3 = 0$ in (5), subspace $\mathcal{E} = \{\alpha(x_{e1}, x_{e2}, 1) \mid \alpha \in \mathbb{R}\}$ is reachable if and only if $(x_{e1}, x_{e2}, 1)$ is reachable.

Omitting arduous computations, we obtain $\mathcal{Z} = \{\beta(1, -1) \mid |\beta| \leq 1/2\}$. Hence, (C_1, C_2) is reachable for any input u_{31} in (4) if it is reachable by $\dot{x}_1 = 1.5 + z$, $\dot{x}_2 = 2.5 - z$ for $|z| \leq 1/2$. The reachable set of this system is simple to analyze: both shelters can be filled exactly to capacity as long as $C_1 \leq C_2 \leq 3C_1$. Since $C_1 = C_2 = 500$, methods of [25] and [22] thus show that the system is resilient to loss of authority over the third checkpoint.

Similar work for $K = 2$ shows that an analogous system is *not* resilient: *no* control law can “fix” the malfunction in one of the checkpoints and an alternative task must be chosen. This analysis thus provides the system designer with guidance on the system’s fundamental limits: if there is a chance that one of the checkpoints will not be functioning as planned and shelter capacities are fixed, there *need* to be at least three checkpoints.

IV. CHALLENGES

Notwithstanding a simple example like above, existing work [19]–[23] still faces a plethora of challenges in order to be applicable to system-wide response to disastrous events. We outline several of them.

Faithful Modeling of System Structure and Disaster Effects

The illustrative example in Sec. III only dealt with a malfunctioning — and possibly adversarial — decision-maker, with the fundamental dynamics of the system remaining the same. Yet, in a variety of disaster recovery scenarios, the adverse event causes a substantial change in the dynamics of the system, *and* also renders them not immediately known. Prior work on theory of guaranteed reachability [19], [23] considers several models of partial knowledge of the dynamics. However, these models — largely reliant on rapid learning of local dynamics combined with bounds on the growth and magnitude of the dynamics — are not amenable to the partial knowledge patterns and comparatively long time scales relevant to infrastructure problems.

One possible path forward is to decompose the system dynamics into known and unknown parts [27], then identifying, as in [22], the worst-case possible unknown dynamics, and obtaining guarantees on system performance and resilience by observing the best possible response to the worst case.

Scalability and Network Control

While the example of Sec. III is limited in the number and interaction of both states and control inputs, a large-scale infrastructure system may consist of thousands of interacting

states and inputs. Taking an example of a transportation network with traffic signals as inputs, the city of Chicago operates around three thousand signalized intersections [28].

Existing results [19]–[23] have only been applied to systems of up to a dozen states and inputs. The computational complexity of resilience and guaranteed performance verification, as well as appropriate control design, remains largely unexplored. As prior results often describe resilience through conditions on eigenvalues and singular values of matrices derived from system data, it is likely that numerical approximation methods are needed to provide such results at scale.

In addition to scalability challenges, infrastructure systems spread over a wide area naturally suffer from *delays in observations and implementation* of reactive control laws, as well as *partial observability* of the system state [12]. A possible way forward is to relate the resilience of an entire system with resilience of its tightly coupled subsystems. Doing so would enable less burdensome resilience certification, as well as simple derivation of *distributed* control inputs.

Complex and Time-Critical Tasks

Disaster response often carries a time imperative or involves a sequence of subtasks that should be completed. It is thus not only necessary to determine *whether* a system is resilient, but *how* resilient it is. To quantify such a notion for time-critical missions, [20] introduced the notion of *quantitative resilience* as the ratio of the times required to reach any target for the initial and malfunctioning systems. Such work provides a possible initial step towards more general quantification of fundamental system capabilities. However, in addition to suffering from all the drawbacks and challenges described previously, the theory of quantitative resilience has thus far only been introduced for systems with malfunctioning inputs, and does not consider tasks more complex than simple reachability. Developing metrics and algorithms for quantitative resilience applicable to a wider class of changes in the system dynamics and possible tasks is crucial for infrastructure applications.

V. CONCLUSION

This paper proposes a vision for the development of theory and algorithms that verify resilience and fundamental capabilities of large-scale infrastructure systems in response to sudden, possibly disastrous events. The envisioned effort is founded on initial work on resilience and guaranteed reachability in partially known or malfunctioning control systems. However, significant challenges remain in attempting to apply existing work: issues of scalability, partial observability, distributed control and complex system interdependencies all pose crucial theoretical, algorithmic, and computational questions. Yet, the potential reward is immense, promising *provable* resilience and *verifiably correct* system-level response to disasters, with the same fundamental theory operating across domains.

REFERENCES

[1] Council of the European Union, “Council Directive 2008/114/EC on the identification and designation of European critical infrastructures and the assessment of the need to improve their protection,” 2008, Brussels, Belgium.

[2] The White House, “Presidential Policy Directive 21: Critical infrastructure security and resilience,” 2013, Washington, USA.

[3] D. L. Alderson, G. G. Brown, and W. M. Carlyle, “Operational models of infrastructure resilience,” *Risk Analysis*, vol. 35, no. 4, pp. 562–586, 2015.

[4] S. Detzer, G. Gurczik, A. Widera, and A. Nitschke, “Assessment of logistics and traffic management tool suites for crisis management,” in *European Transport Conference*, 2016.

[5] T. Sakano, Z. M. Fadlullah, T. Ngo, H. Nishiyama, M. Nakazawa, F. Adachi, N. Kato, A. Takahara, T. Kumagai, H. Kasahara, and S. Kurihara, “Disaster-resilient networking: a new vision based on movable and deployable resource units,” *IEEE Network*, vol. 27, no. 4, pp. 40–46, 2013.

[6] L. Shen, Y. Tang, and L. C. Tang, “Understanding key factors affecting power systems resilience,” *Reliability Engineering & System Safety*, vol. 212, 2021.

[7] W. Neelen and R. van Duijn, “Hacking traffic lights,” in *DEF CON 28 Safe Mode*, 2020.

[8] C. Alcaraz and S. Zeadally, “Critical infrastructure protection: Requirements and challenges for the 21st century,” *International Journal of Critical Infrastructure Protection*, vol. 8, pp. 53–66, 2015.

[9] J. M. Yusta, G. J. Correa, and R. Lacal-Arántegui, “Methodologies and applications for critical infrastructure protection: State-of-the-art,” *Energy Policy*, vol. 39, no. 10, pp. 6100–6119, 2011.

[10] N. Yodo and P. Wang, “Engineering resilience quantification and system design implications: A literature survey,” *Journal of Mechanical Design*, vol. 138, no. 11, 2016.

[11] R. K. Pandey and M. Misra, “Cyber security threats – smart grid infrastructure,” in *National Power Systems Conference*, 2016.

[12] A. Ashok, M. Govindarasu, and J. Wang, “Cyber-physical attack-resilient wide-area monitoring, protection, and control for the power grid,” *Proceedings of the IEEE*, vol. 105, no. 7, pp. 1389–1407, 2017.

[13] N. Yodo, P. Wang, and M. Rafi, “Enabling resilience of complex engineered systems using control theory,” *IEEE Transactions on Reliability*, vol. 67, no. 1, pp. 53–65, 2017.

[14] P. A. Ioannou and J. Sun, *Robust adaptive control*. Prentice-Hall, 1996.

[15] R. Pejmanfar, M. R. Haghifam, S. Soleymani, and B. Tavassoli, “Large signal stabilization of hybrid AC/DC micro-grids using nonlinear robust controller,” *Energies*, vol. 10, no. 12, 2017.

[16] M. Zaman, M. Al Islam, A. Tantawy, C. J. Fung, and S. Abdelwahed, “Adaptive control for smart water distribution systems,” in *IEEE International Smart Cities Conference*, 2021.

[17] A. Alsubaie, K. Alutaibi, and J. Martí, “Resilience assessment of interdependent critical infrastructure,” in *10th International Conference on Critical Information Infrastructures Security*, 2015, pp. 43–55.

[18] J.-P. Rodrigue, *The geography of transport systems*. Routledge, 2020.

[19] H. El-Kebir and M. Ornik, “Online inner approximation of reachable sets of nonlinear systems with diminished control authority,” in *SIAM Conference on Control and its Applications*, 2021, pp. 9–16.

[20] J.-B. Bouvier, K. Xu, and M. Ornik, “Quantitative resilience of linear driftless systems,” in *SIAM Conference on Control and its Applications*, 2021, pp. 32–39.

[21] J.-B. Bouvier and M. Ornik, “Designing resilient linear systems,” *IEEE Transactions on Automatic Control*, 2022.

[22] —, “Quantitative resilience of linear systems,” *arXiv preprint arXiv:2201.12278 [eess.SY]*, 2022.

[23] T. Shafa and M. Ornik, “Reachability of nonlinear systems with unknown dynamics,” *IEEE Transactions on Automatic Control*, 2022.

[24] M. Ng, J. Park, and S. T. Waller, “A hybrid bilevel model for the optimal shelter assignment in emergency evacuations,” *Computer-Aided Civil and Infrastructure Engineering*, vol. 25, no. 8, pp. 547–556, 2010.

[25] O. Hájek, “Duality for differential games and optimal control,” *Mathematical Systems Theory*, vol. 8, no. 1, pp. 1–7, 1974.

[26] H.-N. Nguyen, *Constrained control of uncertain, time-varying, discrete-time systems*. Springer, 2014.

[27] A. Kullberg, I. Skog, and G. Hendeby, “Online joint state inference and learning of partially unknown state-space models,” *IEEE Transactions on Signal Processing*, vol. 69, pp. 4149–4161, 2021.

[28] Chicago Metropolitan Agency for Planning, “(Draft) regional traffic signals existing conditions,” 2019, Chicago, USA.