Computers and Chemical Engineering 123 (2019) 1–11

Contents lists available at ScienceDirect

Computers and Chemical Engineering

journal homepage: www.elsevier.com/locate/compchemeng

Control of complex sociotechnical systems: Importance of causal models and game theory

Venkat Venkatasubramanian[a,1,*], Yu Luo[b], Zhizun Zhang[a]

[a] Department of Chemical Engineering, Columbia University, United States
[b] Department of Chemical and Biomolecular Engineering, University of Delaware, United States

ARTICLE INFO

Article history:
Received 3 August 2018
Revised 5 December 2018
Accepted 5 December 2018
Available online 8 December 2018

Keywords:
Systemic risk
Complex systems
Process systems engineering

ABSTRACT

Recent systemic failures in different domains have reminded us, once again, of the fragility of complex sociotechnical systems. Although the failures occurred in very different domains, there are, however, certain common underlying mechanisms driving these disasters. Understanding these mechanisms is essential to avoid such disasters in the future. To understand them, one needs to go beyond analyzing them as independent one-off accidents, and examine them in the broader perspective of the potential fragility of sociotechnical systems. It is their scale, nonlinearity, inter-connectedness, and interactions with humans and the environment that can make these systems fragile. Here we present an overview of the challenges and opportunities in the modeling and analysis of sociotechnical systems. We highlight a control-theoretic modeling framework that unifies the social and the technical components. We discuss how certain problems can be addressed by using concepts and techniques from causal modeling, game theory, and behavioral economics.

1. Introduction

Control and risk management in complex sociotechnical systems pose unique modeling and execution challenges that go far beyond the ones faced in regulatory control problems. By sociotechnical we mean systems that comprise social elements (namely, humans) as well as technical elements (such as pumps, valves, reactors, etc.). The human elements are an integral part of the system and are often the cause of major systemic failures. The task of designing such systems, and their control mechanisms at all levels, to ensure safe operations over their life cycles is very challenging. Complex sociotechnical systems have a very large number of inter-connected components with nonlinear interactions that can lead to “emergent” behavior—that is, the behavior of the whole is more than the sum of its parts—that can be difficult to anticipate and control. Moreover, these systems are not isolated. They interact with the physical, market, and regulatory environments, resulting in complex feedback dynamics; in particular, human decision-making and the associated errors are part of the feedback processes in these systems. The cumulative effect of the nonlinearity, inter-connectedness, and interactions with humans and the environment makes these system-of-systems potentially fragile and susceptible to systemic failures.

Recent systemic failures in different domains such as the Global Financial Crisis (2007–2009), BP Deepwater Horizon Oil Spill (2010), and Indian Power Outage (2012) continue to remind us of the fragility of complex sociotechnical systems. Systemic failures occur when an entire system collapses, where the system is typically a large entity whose failure negatively impacts a large number of people and their environment, causing enormous financial losses. Examples of such systems are refineries, inter-state power grids, country-wide financial networks, large institutions, and so forth. Union Carbide’s Bhopal Gas Tragedy in 1984, in which an estimated 5000 died and about 100,000 were seriously injured by the accidental release of methyl isocynate was a systemic failure. Another example is the Piper Alpha Disaster in 1988, where an offshore oil platform operated by Occidental Petroleum in the North Sea, U.K., exploded killing 167 and resulting in about $2 billion in losses.

The Challenger (1986) and Columbia (2003) Space Shuttle Disasters, Schering Plough Inhaler Recall (1999), Northeast Power Blackout (2003), SARS Outbreak (2003), BP Texas City Refinery Explosion (2005), Johnson & Johnson Multidrug Recall (2010), and Upper Big Branch Mine Disaster (2010) are all examples of systemic failures in different domains. Examples of financial systemic failures include Enron (2001) and WorldCom (2002) Collapses, and Madoff Ponzi Scheme (2008). The Collapse of News of the World

* Corresponding author.
E-mail addresses: venkat@columbia.edu (V. Venkatasubramanian), yuluo@udel.edu (Y. Luo), zz2216@columbia.edu (Z. Zhang).

[1] To whom correspondence should be addressed.

https://doi.org/10.1016/j.compchemeng.2018.12.010
0098-1354/© 2018 Elsevier Ltd. All rights reserved.

Newspaper Organization (2011) is an example of systemic failure in the media domain. The Wells Fargo Accounts Scam (2016) and Volkswagen Emissions Scandal (2016) are recent examples.

In each case, official postmortem inquiries were conducted and reports of the accidents were produced. Chemical engineers might study, for example, the BP Texas City Refinery Explosion Report (Baker et al., 2007), and people from the financial world may browse, for example, The Financial Crisis Inquiry Report (FCIC, 2011), but rarely does one compare failures across the different domains to study their commonalities and differences. But when one undertakes such a comparative study (Venkatasubramanian and Zhang, 2016), one is struck by the commonality across different domains. There is an alarming sameness about such disasters, which can teach us important fundamental lessons. Although the failures listed above occurred in different domains, in different facilities, were triggered by different events, there are, however, common failure mechanisms that often underlie such events. Systematically identifying and understanding these mechanisms are essential to avoiding such disasters in the future. In this paper, we discuss such a modeling framework, which emphasizes the need for including causal and game-theoretic models as we begin to tackle this challenging control problem.

2. Academic vs industrial view of process control

To illustrate the need for causal models, let us start by contrasting the academic and industrial views of process control. In the typical academic view, we teach process control by introducing dynamic models of prototypical process systems (such as first-order and second-order systems), Laplace transforms and transfer functions, control block diagrams, Bode plots, instability and poles in the right half plane, etc. This is the typical format of an undergraduate process control course, the typical organization of a control textbook. The first author of this paper himself taught process control along these lines for about twenty years at Purdue and for five years at Columbia.

Let us now contrast this academic view of process control with the view as seen from a control operator's perspective on an eventful day. Consider this scenario. Pump A, which is pumping oil, has tripped and the alarm has gone off. The operator does not know the cause of this failure. He switches to Pump B, which is a backup pump installed in anticipation of such events, but that also trips within a few minutes—again, cause(s) unknown. Soon, dozens of alarms go off—the operator doesn't know whether these are all due to the same originating abnormal cause or whether there are multiple abnormal causes involved. Within a few minutes, there is an explosion and fire, which kills two people and a few more are seriously injured in the plant. To make matters more challenging, it is 10 p.m., and the operator is on an off-shore oil platform in the middle of the North Sea.

As the control board operator, how is he supposed to bring the plant under control?

As he struggles with the plant, under great time pressure, he tries to recall the concepts and techniques he was taught in his academic control course—all those Laplace transforms, transfer functions, Bode plots, and poles in right half plane—that he spent months mastering. None of them come to his help. All these are useless at this moment as far as this operator is concerned. But he has a process control problem—his process is out of control, big time! What is he supposed to do? What process control knowledge does he need, beyond what he was taught, to control this abnormal process?

Is this a plausible scenario? Can such an incident happen or did the authors of this paper make up such a scenario to drive home their point?

Unfortunately, not only such an accident can happen, but it did happen in real life. What we described is the initiating sequence of one of the worst ever chemical plant accidents in history, the Piper Alpha Disaster, which killed 167 people and resulted in about $2 billion in losses in 1988.

When the alarms go off, operators want to know, in real-time, answers to questions such as: (i) What are the abnormal causes? (ii) What are the adverse consequences? (iii) How likely are these? (iv) Why these causes and consequences? (v) What is the causal pathway? (vi) What are the control options? (vii) How much time have I got?

Can we develop a real-time operator advisory system that can provide such assistance? Answering these questions requires *cause and effect reasoning*, going back and forth between various potential abnormal causes and their adverse consequences, and analyzing how best they fit the observed alarms and failures. *Developing such causal models is missing in our traditional process control education and research.* How do we develop such models for real-time automation? This is the challenge the first author and his students have worked on for the past 30+ years. The causal models do not replace the traditional models we use in regulatory control, but they complement them. These causal models are essential as we go up on the control hierarchy in a complex sociotechnical system—as we try to model the decision-making in abnormal events management, process hazards analysis, and so on.

We view this as the next phase in the evolution of automated process control systems. The last fifty years saw great progress in the automation of low-level control, i.e., regulatory control, of chemical process plants. We believe the next fifty would see the emergence and dominance of intelligent control systems, through the application of artificial intelligence (AI) methodologies, which automate higher-level control decisions in a complex plant. We believe the promise of AI for intelligent control is finally here, thanks mostly to Moore's Law, which has resulted in tremendous advances in hardware, software, communication, and storage, coupled with the crashing of the cost of all these (Venkatasubramanian, 2018).

The following comparison puts this progress in a startling perspective. In 1985, the supercomputer Cray-2, arguably the fastest computer in the world, performed 1.9 GFLOPS using a 244 MHz processor, consumed 150 KW of power, and cost $32 million (in 2010 dollars). Contrast that with a 2015 Apple Watch, with 3 GFLOPS performance, 1 GHz processor, 1 W (!) power usage, costing about $300! It is a stunning gain of about 150,000 fold in performance/unit cost! This has serious implications in the coming decades. This is why real-time map guidance, self-driving cars, and human-like robotic automation are realities now, no longer the fantasies they used to be when one of us (VV) embarked on developing AI methods for process systems engineering applications in 1983.

3. TeCSMART: a hierarchical model of complex sociotechnical systems

Most human engineered complex systems, such as chemical plants, corporations, transportation networks, power grids, governments, societies, etc., are organized as a hierarchical network of human and non-human (e.g., machines) entities. Generally speaking, they comprise of autonomous and non-autonomous elements, which usually translate to human and non-human entities. In this paper, we are not considering non-human entities that are autonomous, such as robots, as they have not reached human-like autonomous capabilities yet, even though this is going to be an important development in a couple of decades.

We call our modeling framework as *TeCSMART (Teleo-Centric System Model for Analyzing Risks and Threats)* (Venkatasubramanian and Zhang, 2016). *Telos* means goal or purpose in Greek. *The central theme of our approach is the emphasis on recognizing and modeling goals of different agents, at different levels of abstraction, in a complex sociotechnical system.* Both individual players and groups are goal-oriented, driven to act by their goals and incentives, in a complex system. Individuals usually have different goals, or even goals with conflicts of interest with each other or with goals from other individuals. The dynamics of how goals across the system interact, transform, and disperse in the hierarchy is essential in affecting both individual and systemic performances. Modeling this results in a multi-scale modeling framework, having *seven layers* organized as a hierarchy, as shown in Fig. 1, that naturally arise and represent different perspectives of the entire system. Each layer above is a zoomed-out, aggregate, view of the immediate layer below. For example, the block representing process unit in the network of Plant View contains the individual feedback loop in Equipment View. The bottom layer of the stack is the basic building block of a system (e.g., equipment and processes). The top layer of the stack is the macroscopic view of a society.

Each layer has its own set of goals which drive the decision-making and actions taken by the agents in that level. The decisions are taken based on the inputs the layer receives from the layers immediately above and below it. Similarly, the actions are communicated to these adjacent layers as outputs. These decisions/actions are indicated, in Fig. 1, by the arrows that capture these information flows, up and down the hierarchy. These information flows are the feedback loops between the layers (i.e., *inter-layer* feedback loops). There are also feedback loops within a given layer, as depicted in the figure, which are *intra-layer* loops. Associated with each layer is a set of agents (autonomous and non-autonomous), organized in a particular configuration that is appropriate for the goals of that layer (e.g., the layout of equipment in a chemical plant, called a flowsheet). Such a multi-layered representation lends itself naturally to account for emergent phenomena that arise from one scale to another.

This is a uniform and unified input-output modeling framework that is conceptually the same across all levels. This elementary input-output model structure that serves as a building block in our framework is shown in Fig. 2. Specifying such a uniform modeling structure across all levels has the advantage of integrating and unifying the analysis of the outcomes at different levels in a consistent manner. Such a template structure allows us to systematically identify the various failure modes of the different elements at different levels of the hierarchy as we have discussed elsewhere (Venkatasubramanian and Zhang, 2016).

There are seven key elements in this control-theoretic modeling building block: (i) input, (ii) output, (iii) sensor, (iv) actuator, (v) controller, (vi) “process” unit that transforms inputs to outputs, and (vii) connection (e.g., wires and pipes).

As an organized group, these entities collect, decide, act on, report, and receive a variety of performance information and metrics. In a chemical plant, for example, in the Equipment View layer, they collect, decide, and act on individual process and equipment performance information and metrics (such as temperature, pressure, flow rate, batch times, etc.), that are vital for safe, efficient and profitable operation, and report them to the Plant View layer, and receive, in turn, local control specifications (such as temperature and pressure setpoints) from Plant View layer.

The Plant View layer agents make these decisions by considering information from all the processes and equipment under its purview as well as by considering manufacturing targets (such as what to make, how much to make, when to make, etc.). These targets, in turn, are decided by the agents in the Management View, which get translated into the associated setpoints and constraints by the agents in the Plant View, and communicated down to the Equipment View as inputs. The target metrics are decided by the agents in Management View by responding to competitive market conditions as dictated by the Market View. In a similar manner, relevant information regarding market or company stability, performance, fair competition, etc., are monitored and acted on by the agents in the Regulatory View, by enacting and enforcing appropriate regulations approved by the agents in the Government View (such as the Congress in the U.S.). In an ideal democracy, a government is elected by the citizens of that society, who have the final word in determining what kind of government and laws they would like to live by.

Consider the Management View, for example, where the agents involved are the critical decision-makers such as the CEO, Senior Vice Presidents, and Board of Directors. Their goal is to maximize profitability and create value for the shareholders by making sure the company’s business performance metrics (including safety) meet the expectations from the Market (which is the next level up). Influenced by the nature of business and accounting cycles, this layer operates in a time scale of quarter (i.e., 3-month period) to a year. As seen in the control-theoretic information model of this level in Fig. 3, this group of decision-makers (Management team) sets the overall policies that “control” (i.e., manage) the behavior and outcomes of the corporation including its autonomous and non-autonomous assets. Autonomous agents at this layer include managers and supervisors of each division, while the non-autonomous agents are corporate assets. The Market at the next level up sets and demands certain performance targets be met by the company for its survival and growth. These metrics are usually financial at this level such as ROI, ROE, market hare, sales growth, and so forth. These are the *setpoints* and *constraints* given to the Management team.

The Management team, in turn, translates these targets into actionable quantitative information such as production performance metrics, strategic deployment of resources, and so forth, at different plants (the corporation might have several plants distributed all over the world) as well as more qualitative ones that define the company culture including the safety culture. They also set the incentive policy to encourage better performance from the employees. These are communicated to the Plant View layer as their *setpoints* and *constraints*. The Management team decides on these targets by taking into account of all relevant information concerned with the survival, profitability, and growth of the company in a competitive and regulatory environment. Thus, the information flow is not only from the company’s internal sources but also from the environment, which are the two levels immediately above.

Differing from the control policies at the lower levels, which mainly focus on controlling equipment (i.e., non-autonomous agents), the policies from this layer onward, at the higher levels, focus more on achieving the desired behavior and outcomes from autonomous agents (i.e., humans). As a result, while the lower level control policies can be based on precise models of process/equipment (as captured by DAE models), the higher level policies will necessarily have to deal with imperfect models of human behavior which cannot be reduced to a set of equations.

Consider, for instance, the difficulties involved in “modeling” the culture of a corporation. At best, we might be able to identify certain key features or characteristics that define a corporation’s culture. From this level onward, we have to rely more on *graph-theoretic, game-theoretic, and agent-based modeling frameworks*. Thus, from this level onward modeling becomes trickier, and the notion of “control” of agents transitions to the “management” of agents. Moreover, the importance of TeCSMART failure modes-based examination becomes more obvious. Such a systematic risk analysis of human decision-making would help improving safety-related management activities, among other things.

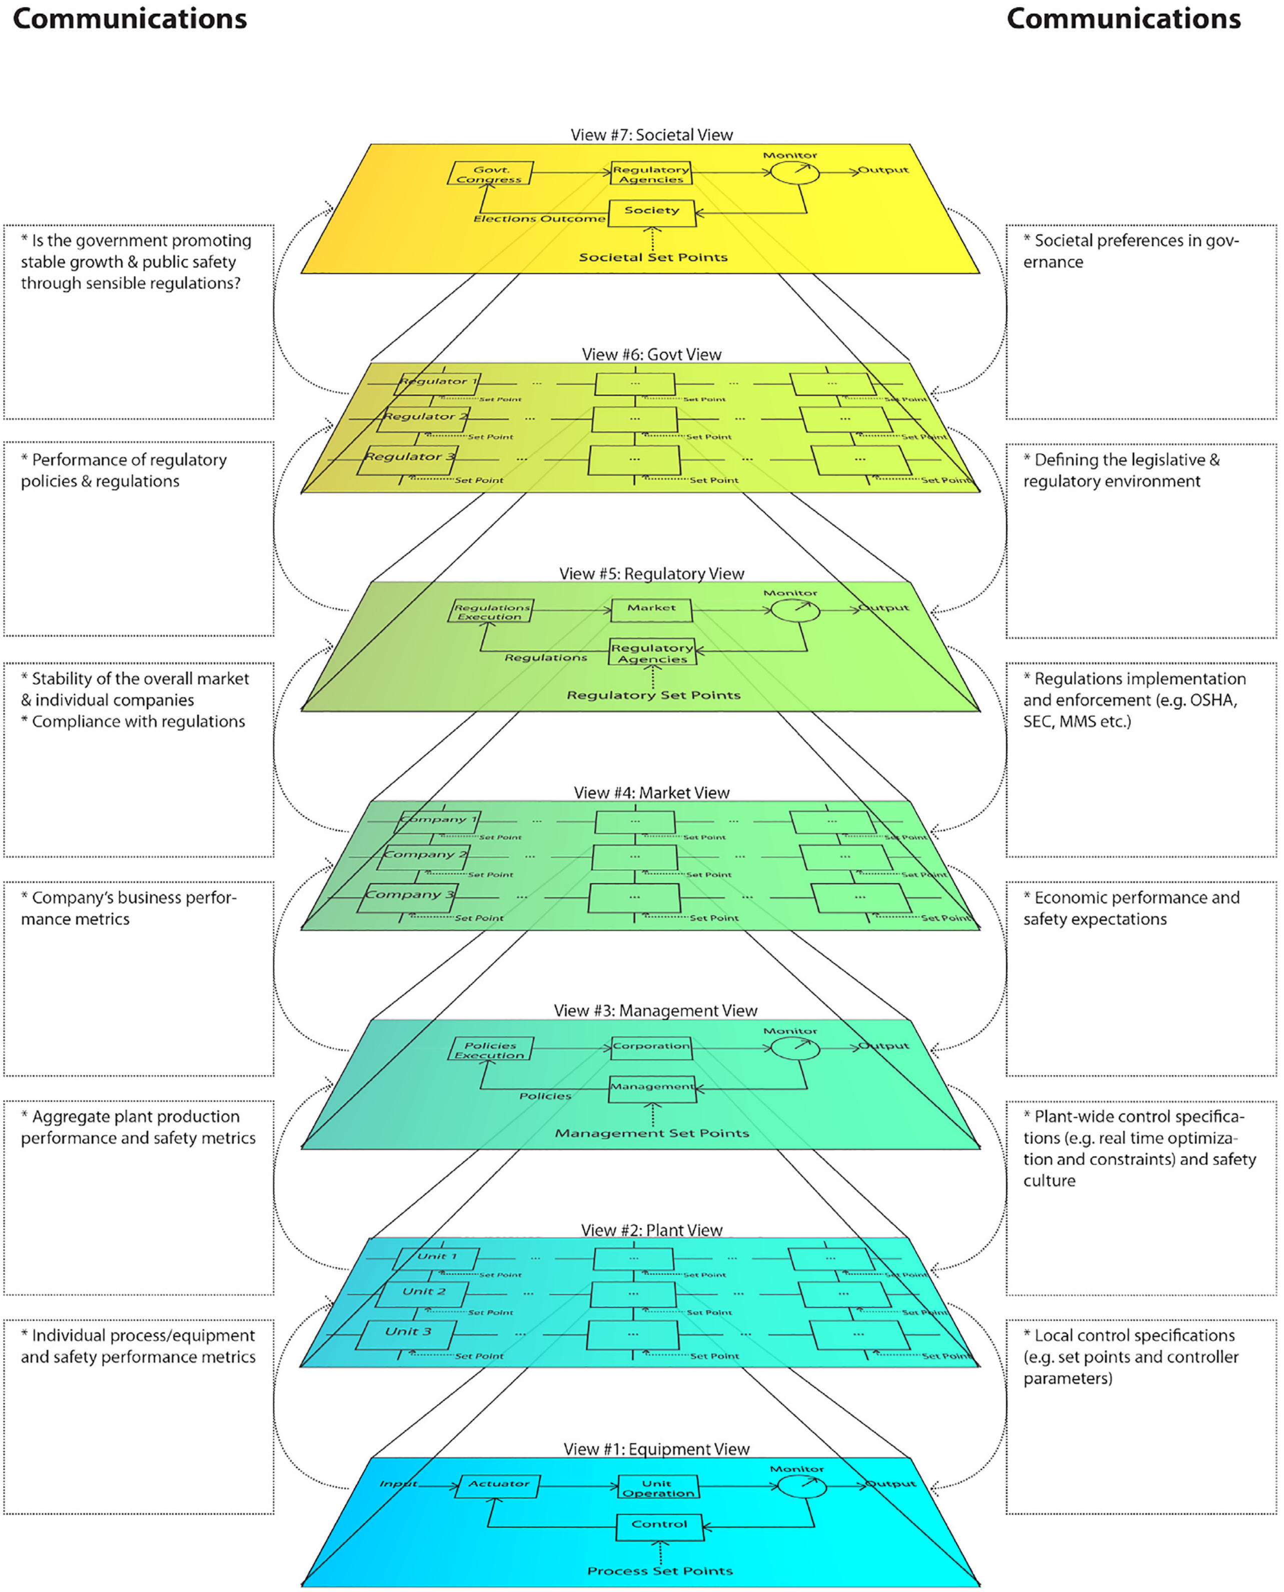
Communications
Communications
View #7: Societal View
Govt. Congress
Regulatory Agencies
Monitor
Output
Society
Elections Outcome
Societal Set Points
* Is the government promoting stable growth & public safety through sensible regulations?
* Societal preferences in governance
View #6: Govt View
Regulator 1
Regulator 2
Regulator 3
Set Point
* Performance of regulatory policies & regulations
* Defining the legislative & regulatory environment
View #5: Regulatory View
Regulations Execution
Market
Monitor
Output
Regulations
Regulatory Agencies
Regulatory Set Points
* Stability of the overall market & individual companies
* Compliance with regulations
* Regulations implementation and enforcement (e.g. OSHA, SEC, MMS etc.)
View #4: Market View
Company 1
Company 2
Company 3
Set Point
* Company's business performance metrics
* Economic performance and safety expectations
View #3: Management View
Policies Execution
Corporation
Monitor
Output
Policies
Management
Management Set Points
* Aggregate plant production performance and safety metrics
* Plant-wide control specifications (e.g. real time optimization and constraints) and safety culture
View #2: Plant View
Unit 1
Unit 2
Unit 3
Set Point
* Individual process/equipment and safety performance metrics
* Local control specifications (e.g. set points and controller parameters)
View #1: Equipment View
Input
Actuator
Unit Operation
Monitor
Output
Control
Process Set Points

Fig. 1. TeCSMART framework.

The Management team acts as a "controller" to monitor the various performance metrics (e.g., sales, expenses, revenue, profits, ROI, ROE, etc.), compare them with the setpoints, and take appropriate actions by manipulating the relevant variables (e.g., cost cutting, acquisition, etc.) in order to meet the setpoint targets. The Management level deals with the big picture and general strategy for the corporation as a whole. These get translated into more detailed prescriptions and recommendations as they are communicated from this layer to the lower layers. The failure of the elements in Fig. 3 can be modeled along the lines of Equipment View

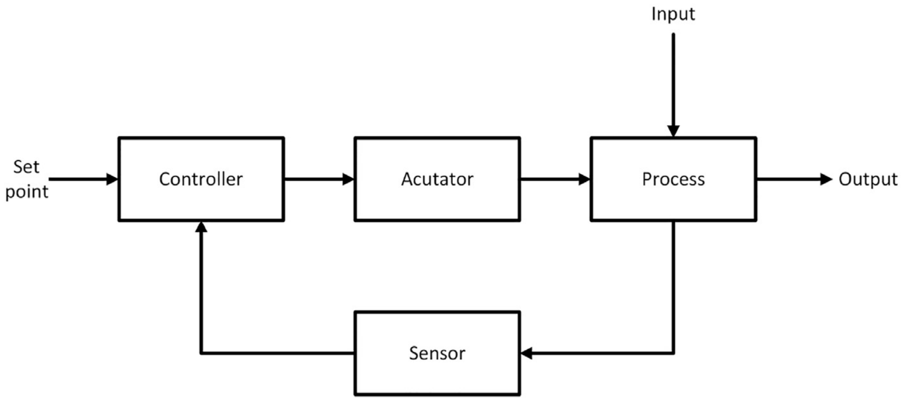

Fig. 2. Schematic of a feedback control system (Adapted from Stephanopoulos (1984), fig. 13.1b, pp. 241).

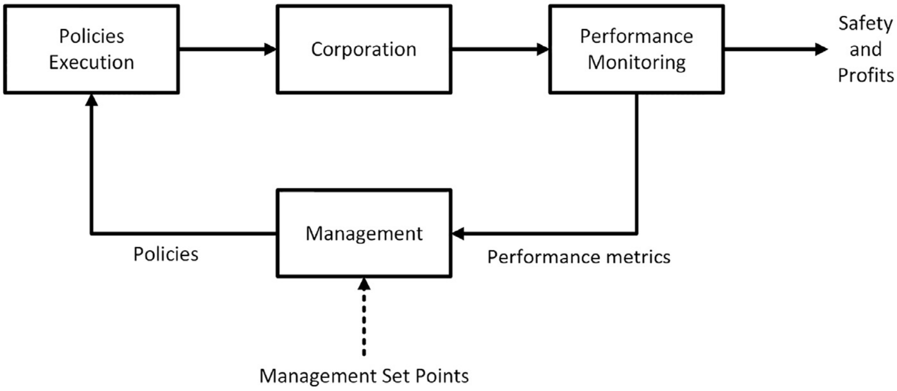

Fig. 3. Control-theoretic model of company/management layer.

and Plant View layers. For example, the Performance Monitoring task (i.e., “sensor”) may fail because of errors in the measurements or estimations (e.g., fail high, low, or zero) or they may be communicated (or not communicated at all) erroneously. One can methodically identify similar failure modes for the other elements including the connections (which are the communication channels) as we have described in Venkatasubramanian and Zhang (2016).

Thus, as we move higher up in the hierarchy, it becomes clear that we need causal models, often modeled as signed digraphs (SDG), and goal-driven, incentive-based, game-theoretic models, often addressed using agent-based models. As an example, we show, in Figs. 4 and 5, the hierarchical causal model of a sour water stripping plant that is used for automated process hazards analysis. These figures show the knowledge representation of cause-and-effect and failure modes knowledge using Petri nets and SDG. We have discussed such approaches in greater detail elsewhere (Bookstaber et al., 2015; Luo et al., 2016; Maurya et al., 2003a; 2003b; 2004; Srinivasan and Venkatasubramanian, 1996; 1998a; 1998b; 1998c; Vaidhyanathan and Venkatasubramanian, 1995; 1996; Venkatasubramanian, 2011; Venkatasubramanian and Rengaswamy, 2003a; 2003b; Venkatasubramanian et al., 2003; Venkatasubramanian and Vaidhyanathan, 1994; Venkatasubramanian and Zhang, 2016; Venkatasubramanian et al., 2000).

We applied the TeCSMART framework to analyze the following thirteen well-known systemic failures: (1) the Bhopal Disaster (1984), (2) the Space Shuttle Challenger Disaster (1986), (3) the Piper Alpha Disaster (1988), (4) the SARS Outbreak (2003), (5) the Space Shuttle Columbia Disaster (2003), (6) the Northeast Power Blackout (2003), (7) the BP Texas City Refinery Explosion (2005), (8) Global Financial Crisis (2007–2009), (9) the BP Deepwater Horizon Oil Spill (2010), (10) the Upper Big Branch Mine Disaster (2010), (11) the Chilean Mining Accident (2010), (12) the Fukushima Daiichi Nuclear Disaster (2011), and (13) the India Blackouts (2012).

We carefully reviewed the official postmortem reports of these disasters as well as other relevant sources. We analyzed and classified over 700 failures mentioned in these reports (Baker et al., 2007; Bonnefoy, 2010; Browning, 1993; CAIB, 2003; CERC, 2012; CSB, 2005; Cullen, 1993; FCIC, 2011; cha, 1986; Kurokawa et al., 2012; McAteer et al., 2011; PresidentialCommission, 2011; Task-Force, 2004; WHO, 2006). We categorized these failures into five primary classes, and nineteen subclasses, that are consistent with the typical failure modes we discussed in the previous section. The five classes are as follows: (1) Monitoring Failures; (2) Decision-Making Failures; (3) Action Failures; (4) Communication Failures; and (5) Structural Failures. Each category has sub-categories that define more detailed failures. The details of this analysis can be found at Venkatasubramanian and Zhang (2016).

4. Soft regulation: coordinating self-interested agents in sociotechnical systems

As noted, the dynamics of the autonomous agents in the higher levels of the hierarchy (e.g., regulators, policy-makers) are determined by incentive-based goal-driven behavior, whose modeling and analysis require incorporating a control-theoretic, behavioral economics approach. In this section, we present a novel framework, which we have named as *soft regulation*, for coordinating decision-makers at the industry and regulatory level. We only present an outline as the details have been discussed elsewhere (Luo et al., 2016).

In a typical regulatory environment involving conventional technologies, regulators issue mandates that have to be followed

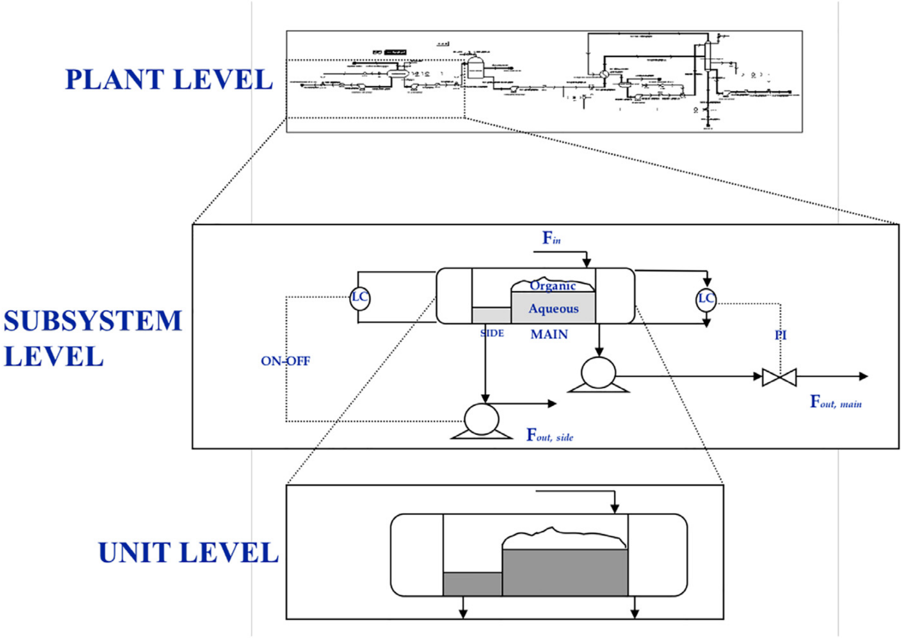

Fig. 4. Hierarchical causal model.

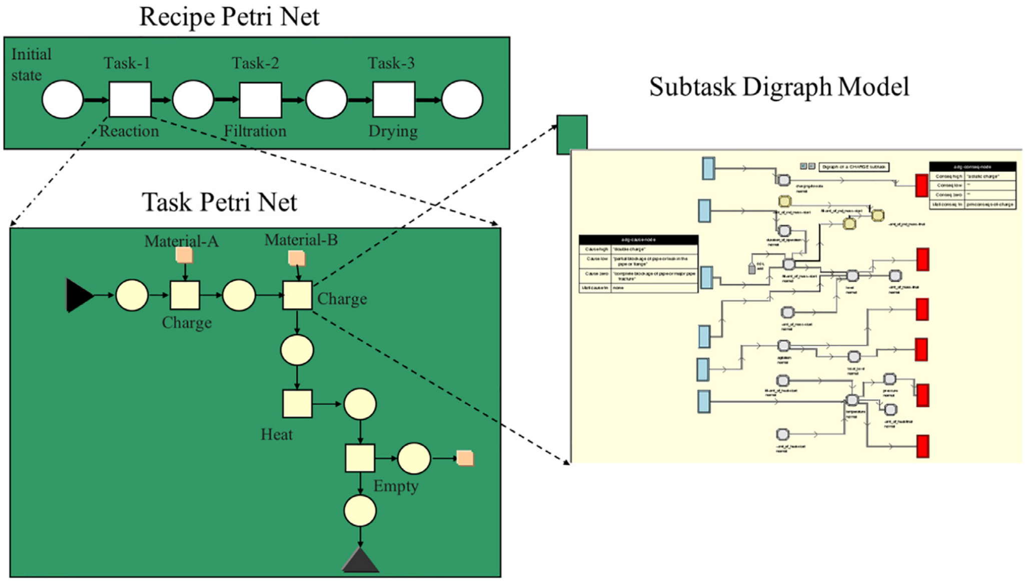

Fig. 5. SDG and petri nets representation of the causal model.

by the regulated agents. The agents face fines and other punitive consequences for non-compliance. We call this approach *hard regulation*. While hard regulation might be effective in the lower levels of the hierarchy (i.e., this is simply regulatory control), such as in the Equipment Layer or the Plant Layer, where reasonably reliable models of equipment and process are available, it is not effective in the higher levels as such models are neither available nor even possible.

One possible course of action would be to offer options to agents that are likely to be adopted because they are incentive compatible, namely *soft control* (Han et al., 2006; Zhang and Parkes, 2008). Examples of this approach include the soft paternalism approach for modifying social behavior (Thaler and Sunstein, 2003) wherein carefully designed options "nudge" people to make better decisions to be more environmentally aware or healthy (Thaler and Sunstein, 2008). However, as in the case with hard control, soft control can be used only when there is a *reliable model* and a *well-defined setpoint*. Soft paternalism and similar social mechanisms are effective because we understand saving energy and staying physically active are the right things to do.

What if we do not know what is best for the agents? Can we somehow learn this as we function? *Soft learning* is a class of learning mechanisms that appropriately incentivize agents in a social network to aggregate such important information. Examples of soft learning include social sensing and social learning (Krishnamurthy and Poor, 2014; Rendell et al., 2010; Shmueli et al., 2014) in the context of real-time traffic information and online reviews (such as Yelp).

Our new regulatory paradigm, *soft regulation*, combines features of *soft control* and *soft learning*. The regulator aggregates key system-level statistics in a "privacy preserving" (Abbe et al., 2012) manner (individuals do not need to explicitly disclose their states) and shares these statistics with all agents. The agents have the flexibility to accept, reject, or partially accept the recommendations from regulator based on their own self interests. The recommendations are simply “nudges” (Thaler and Sunstein, 2008). The mechanism does not interrupt the regulated entities who have direct access to field performance. It creates a collective learning environment for both the regulator and the agents. This partial acceptance (or “confidence level” β) of recommendation is a crucial feature of soft regulation. It critically determines the effectiveness of the mechanism. Soft regulation seeks a balance between over- and under-regulation: Agents have the freedom to rely on both individual exploration and social learning.

We expect soft regulation to be effective when the system has the following features:

1. *Imperfect information*: The action-utility payoff structure is poorly understood, i.e., the data are *noisy* and the models are absent or *incomplete*. Each individual may only possess partial information about the unknown process. Agents rely on inaccurate measurements, approximations, or subjective evaluations to optimize.

2. *Weak interaction*: The agents can optimize their own actions without taking into consideration the response of other agents, i.e., each’s utility or payoff is only a function of the agent’s own state, and the optimal setpoint is identical among agents. A good example of such a setting is the initial stages of a new technology; the resources being exploited are abundant and the profits of the agents are not limited by competition but by their ability to exploit the resource effectively. Although the reward an agent receives while operating at a setpoint may vary, the setpoint itself, however, is likely to be identical or at least restricted to a narrow range. The discovered setpoints (by soft regulation or traditional methods) will later become the industry standards when the technology matures. Another example of setting with weak or no interaction is when humans improve their own health conditions by changing habits, medications, or even environments. The interaction among agents is usually minimal. Although each has his/her own unique physiological configurations, grouped by characteristics such as age, gender, profession, etc., they are likely to exhibit common optimal setpoints within groups.

3. *Bounded rationality*: Agents are autonomous and self-interested, and they *always* move in a direction that locally improves utility, subject to available information.

Soft regulation creates a feedback system where agents have the freedom to choose to accept this feedback. Feedback has long been recognized as an essential feature of complex adaptive systems where causes and effects are intertwined. There have been several attempts over the years to understand the dynamics of social systems in terms of feedback control (see, e.g., Carver and Scheier (1982); Leveson (2011); Powers (1973); Trochim et al. (2006)). While such contributions are useful advances, much of this work, however, is conceptual and qualitative.

In contrast, soft regulation is a practical and quantitative methodology—self-interested agents use the feedback from past outcomes to determine future actions, and the regulator provides all the agents a feedback that aggregates system-level information. One can extend this to include group, organization, even societal-level feedback loops.

There are two ways to generate feedback for soft regulation: the best recommendation and the crowd recommendation. As the name suggests, best recommendation corresponds to the case where the regulator has full information and computes the feedback by solving a centralized optimal control problem. The crowd recommendation on the other hand, is simply the average of the participants’ actions. We have shown that, despite its simplicity, crowd recommendation is as good as the best recommendation for a wide range of β values (Luo et al., 2016). This is actually not surprising. The collective wisdom of groups has been discussed by scholars, for example, in the Condorcet’s jury theorem (De Condorcet, 2014) and in popular culture such as bestsellers The Wisdom of Crowds (Surowiecki, 2005) and Wiser: Getting Beyond Groupthink to Make Groups Smarter (Sunstein and Hastie, 2014). In our paper (Luo et al., 2016), we proposed a control-theoretic framework that goes beyond one-time predictions and showed the effectiveness of the wisdom of crowds for optimizing a process using continuously refined information.

We analyzed a stylized model of soft regulation that preserves the essential features discussed previously. The system consists of one regulator and n agents. Agent i wants to select an action x_i that maximizes the value of the real-valued and strongly concave utility function $f_i(x_i)$ over a convex compact set $\mathbb{X} \subseteq \mathbb{R}$. We assume that although the individual utility functions f_i are different for each agent, the optimal setpoint $\theta^* = \arg\max_{x\in\mathbb{X}} f_i(x)$ is identical across agents.

We assume that the utility function $f_i(x_i)$ is not explicitly known, nor is it deterministic; agents cannot solve the optimization problem explicitly. In theory, by averaging out the noise, one can obtain a more accurate mapping of the utility function. However, in our setting, each sample corresponds to actual utility each agent receives; therefore, they might not have the incentive to oversample at the location where the utility is low. The agents update individual actions using the following dynamics:

$$\bar{x}_i = g_i(x_i) \tag{1}$$

where g_i denotes the optimization algorithm used by the ith agent. In practice, g_i can be any function that maps an old action x_i to a new action $\bar{x}_i$. In order to converge to the optimal θ^*, the function must satisfy regularity conditions. More specifically, g_i should converge to a unique fixed point regardless of the initial value of x_i. For instance, the Kiefer–Wolfowitz stochastic gradient method (Kiefer and Wolfowitz, 1952) is one of the many techniques that satisfy such condition. We call a setting where an agent updates its action based on its own measurement the *open loop* scenario (or asocial learning as in Rendell et al. (2010)).

$$g_i(x_i) = x_i + \frac{a_t}{c_t} \cdot \big(f_i(x_i + c_t) - f_i(x_i - c_t)\big). \tag{2}$$

In the soft regulation setting the regulator computes a feedback recommendation u. The agents then combine u with $\bar{x}_i = g_i(x_i)$ to compute a new action x_i^+ and the dynamics can be described as a feedback control loop:

$$x_i^+ = h_i(x_i) \equiv (1-\beta_i)g_i(x_i) + \beta_i u = (1-\beta_i)\bar{x}_i + \beta_i u \tag{3}$$

where $\beta_i \in [0,\ 1]$ or $[0,\ 100\%]$ is a measure of the *confidence* that the ith agent puts on the feedback (or the degree of social influence if the feedback is from the peers (Luo et al., 2018)), and is therefore, called the confidence level. The confidence level β plays an important role in the resulting dynamics. Action changes are relatively independent of recommendation for agents with small β (the *explorers*), and action remains in the vicinity of u for agents with large β (the *followers*). Note that $h_i(x_i)$ can be re-written as follows:

$$h_i(x_i) = x_i + (1-\beta_i)(\bar{x}_i - x_i) + \beta_i(u - x_i). \tag{4}$$

The soft regulation feedback function resembles the feedback seen in bird flocks and swarm intelligence (Kennedy, 2010).

When the regulator is fully informed about the functions f_i, g_i, and β_i, the optimal feedback u^* can be computed explicitly by solving the following centralized optimal control problem that

Table 1
Model parameters.

n	σ_ω	θ^*	k	a_t	c_t
1000	$200/\sqrt{3}$	0	100	$1/t$	$1/(t+200)^{1/3}$

maximizes social welfare (sum of utilities) over the projected trajectory:

$$\begin{aligned} \max_{u(t)} \quad & \sum_t w(t) \sum_i f_i(x_i(t)) \\ \text{s.t.} \quad & x_i(t+1) = (1-\beta_i) g_i(x_i(t)) + \beta_i u(t) \end{aligned} \tag{5}$$

where the time varying weight $w(t)$ can favor either the present or future. We call the solution u^* to this problem the best recommendation.

Since the function f_i, g_i, and the parameter β_i are only privately known to the agents, in practice, it is unlikely that the regulator knows the functions and the parameters. Following Brotman (2014), Surowiecki (2005) and Sunstein and Hastie (2014), we assume that the regulator reports the average, i.e., $u = \frac{1}{n}\sum_i x_i$. We call this recommendation the crowd recommendation. Note that using privacy preserving computations (Abbe et al., 2012), the regulator can compute the crowd recommendation without ever learning any individual input x_i. In Luo et al. (2016), we proved mathematically that the crowd recommendation ensures the convergence to the optimal setpoint; moreover, we demonstrated using an agent-based model that it is as good as the best recommendation for a wide range of β values.

Recall that, we assumed the underlying action-utility payoff function $f_i(x_i)$ to be strongly concave. In the following simulation and closed form analytical solution, we assumed that the dynamics of a group of heterogeneous agents can be described by the same number of representative agents with an identical and quadratic utility function $f_i(x_i) = f(x_i) = -k(x_i - \theta^*)^2 + \omega$. The assumption of representative agent is helpful in identifying the effect of β value based on agent-based simulations as well as developing a closed form analytical solution. In order to study the convergence behavior, one can without loss of generality, assume that $\theta^* = 0$. This particular choice for f is motivated by the fact that any strongly concave function can be approximated by a quadratic function near its optimum. The noise is $\omega \sim \mathcal{N}(0, \sigma_\omega)$. Agents only observe the noisy function values—the underlying structure is not known to the agents.

We define the optimization *efficiency* as the percent reduction in MSE:

$$\eta_t \equiv \frac{\mathrm{MSE}_{t_0} - \mathrm{MSE}_t}{\mathrm{MSE}_{t_0}} \times 100. \tag{6}$$

The efficiency is 100 when the system reaches optimum. We simulated the agent dynamics in NetLogo. The model parameter values are listed in Table 1. The parameters do not represent practical meanings. The particular values are chosen such that the results are easily identifiable. a_t and c_t are the time-varying Kiefer–Wolfowitz step sizes used by (2).

We first ran the simulation for soft regulation with best recommendation. Given the quadratic utility, Kiefer–Wolfowitz algorithm, and system-wide β value, the regulator can easily compute best recommendation by solving (5). The best recommendation is

$$u^*(t) = -\left(\frac{(1-\beta)(1-4ka_t)}{\beta}\right) \cdot \frac{1}{n}\sum_i x_i(t). \tag{7}$$

In Fig. 6, we plot the efficiency after 200 iterations against different β values. The efficiency increases monotonically as the β value increases. This result is not surprising. As the β value increases, the regulator has a stronger influence on the agents, therefore, exerting a more efficient control. Even though for each β value, the regulator issues the best recommendation, the recommendation is only effective when the agents choose to listen.

In Figs. 7 and 8, we plot the efficiency against β value for soft regulation with crowd recommendation. The results from Fig. 6 are also included as a reference. It is remarkable that soft regulation with crowd recommendation is as good as the one with best recommendation for a wide range of β values (from 0 to 99%). The real advantage of best recommendation only appears when the β value is close to 100%. However, to achieve this best recommendation or even hard regulation, the regulator needs information about utility function, optimization algorithm, and the β value. This practice, despite being efficient under the setting of complete information, is costly, impractical, and error prone in practical settings. Especially for hard regulation, additional cost of enforcement needs to be considered.

The results in Figs. 7 and 8 show that system only reaches about 70% optimum while the system performance is more than 90% optimal when the β value is 50% (i.e., the agent takes an average between its own optimization result and the recommendation). There is a sharp decline in performance when β value is too close to 100%. Beyond this “cliff,” the agents explore very little and essentially stay where they are.

In Fig. 9, we plot the time progressions of efficiency for different β values. When β is low ($\beta = 0$ or 10%), the MSE increases (efficiency declines) before converging. This is caused by large initial step sizes. As β increases, the system begins convergence earlier. As β further increases, the system shifts from the regime dominated by exploration to the one dominated by conformity, and the recommendation does not have enough time to converge to optimum before agents start conforming. We also analytically approximated the multi-agent system. Our approximation agrees closely with the simulation (see lines in all simulation result figures). Interested readers can find detailed derivations in Luo et al. (2016).

Despite the name, soft regulation has applications beyond industrial regulation. The soft regulator module, i.e., $x_i^+ = (1-\beta_i)\bar{x}_i + \beta_i u,$ can be integrated in different control systems and problem-solving scenarios. We only analyze a specific and stylized model in this paper to illustrate the efficacy of the mechanism. In practice, soft regulation should be implemented and modified in a case by case manner. For example, when the regulator can obtain more information other than the average action, it is entirely reasonable to formulate a better recommendation based on the richer information set such as trends, histograms, etc. The agents, instead of adjusting β value via the method proposed in this paper, can also explore and compare utilities (on a much slower timescale) to adapt new β values. For a large population where centralized information collection is impractical, soft regulation might be plausible on a peer-to-peer basis. All these possibilities will be explored and analyzed in future work.

The medical domain is another applicable area of soft regulation. Powered by mobile phones and wearables, researchers can now collect timely mass medical data (via Apple's ResearchKit (Apple, 2015) for example). Soft regulation is suitable in this scenario because medical research satisfies all three features, i.e., imperfect information (unknown relationships between patient behaviors and health conditions), weak interaction (one patient's condition is not affected by another's), and bounded rationality (patients always wish to improve their own health, however, have limited information). In addition, thanks to the convenience of mobile devices, we expect good participation rate. A large population size further ensures the accuracy of recommendation. Patients can optimize their own health while contributing to medical research. Even if patients do not want to optimize themselves, medical researchers may implement the soft regulation module to do that based on data collected locally. The β value can also be explicitly controlled by the service provider. Soft regulation in this setting

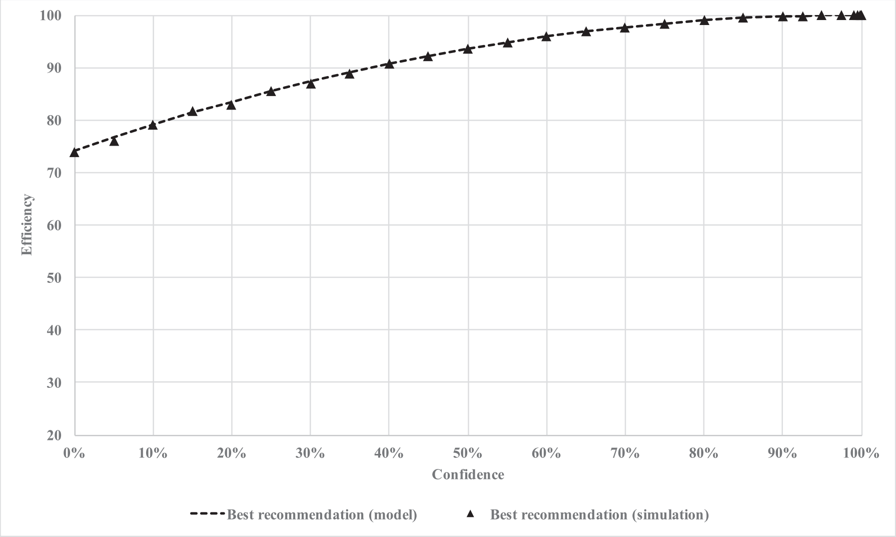

Fig. 6. Efficiency of soft regulation with best recommendation.

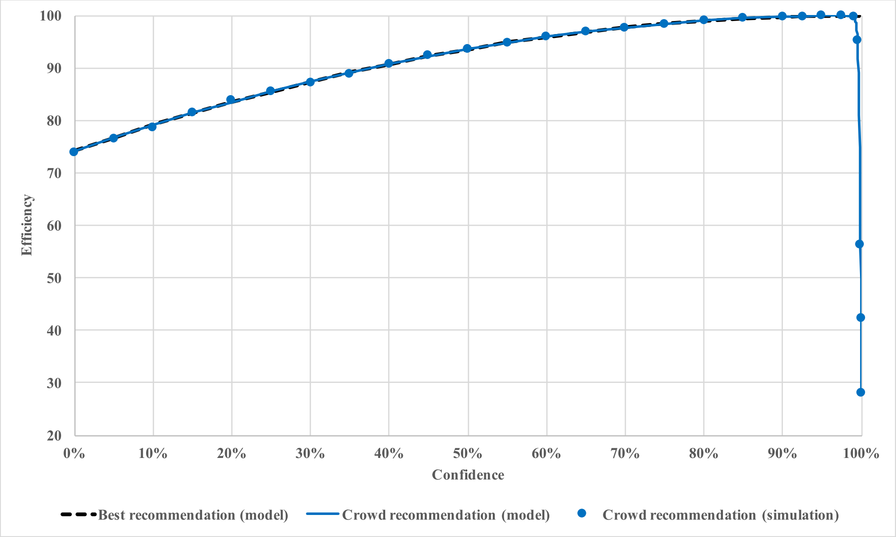

Fig. 7. Efficiency of soft regulation with crowd recommendation.

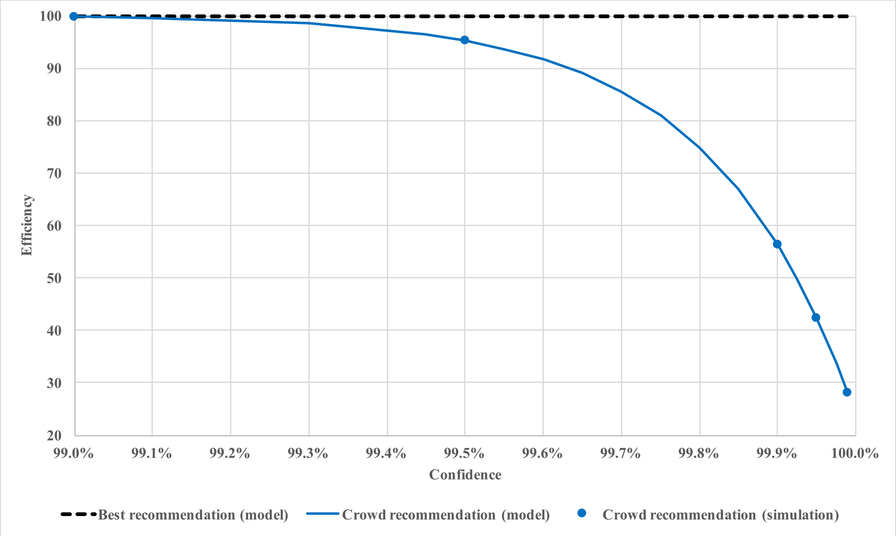

Fig. 8. Efficiency of soft regulation with crowd recommendation (large β values).

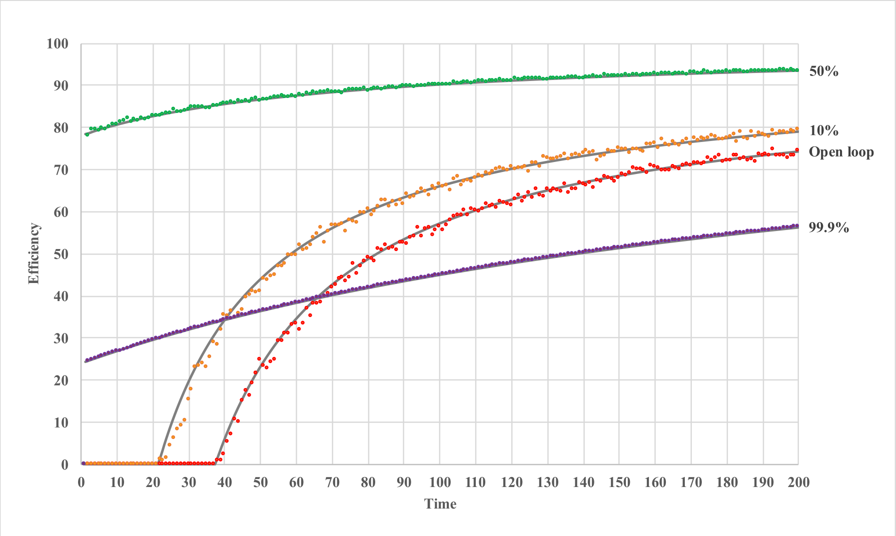

Fig. 9. Efficiency of soft regulation with crowd recommendation over time.

becomes a crowdsourcing framework. The results in this paper are expected to hold.

This work also has some implications other than our central arguments on control and regulation. It reinforces the idea that an averaged opinion can accurately predict under uncertainty, i.e., the wisdom of crowds, given the population is large, independent, and relevant. Unlike conventional takes on the wisdom of crowds, soft regulation does not stop at collecting average information but also feeds it back to the system. This dynamical mechanism suggests more flexible scenarios and applications. Recall that weak interaction is one of the features for an effective soft regulation of emerging technologies. A tightly coupled system where agents compete with each other for limited resources (e.g., regulatory setting of developed technologies) requires a more sophisticated optimal control framework with game-theoretic components such as mechanism design to guide decision-making at the regulatory level.

5. Summary and conclusions

In this paper, we have argued the need for including causal and game-theoretic models in the control of complex sociotechnical systems. Our study using a unified control-theoretic framework, namely, TeCSMART, shows how these apparently disparate failures correspond to common failure modes associated with the elements of a control system, namely, sensor, controller, actuator, process unit, and communication channels. Even though every systemic failure happens in some unique manner, and is not an exact replica of a past event, we show that the underlying failure mechanisms can be traced back to similar patterns associated with other events, thus teaching us valuable lessons for the future. We believe the use of causal models and failure models are indispensable in this regard. In a similar vein, game-theoretic behavioral economics models play an important role in the modeling of decision-making in complex sociotechnical systems. In this context, we have proposed a novel framework, called *soft regulation*, which addresses some of the challenges faced by regulators. Modeling frameworks such the ones we have proposed are the beginnings of the next era in the development of intelligent control systems for complex sociotechnical systems. This program has a long way to go, perhaps two to three decades, but it is bound to revolutionize automatic control as regulatory control and model predictive control did in the last four decades.

Acknowledgments

This paper is based on the invited presentation made at the *Foundations of Computer Aided Process Operations–Chemical Process Control Conference*, Tucson, Arizona, January 8–12, 2017. The authors thank the organizers for their kind invitation. This work is supported in part by the Center for the Management of Systemic Risk (CMSR) at Columbia University.

References

Abbe, E.A., Khandani, A.E., Lo, A.W., 2012. Privacy-preserving methods for sharing financial risk exposures. Am. Econ. Rev. 102 (3), 65–70.
Apple (2015). Researchkit for developers. [Online]. Available: https://developer.apple.com/researchkit/.
Baker, J., Leveson, N., Bowman, F., Priest, S., 2007. The report of the bp us refineries independent safety review panel. Independent Safety Review, Report.
Bonnefoy, P. (2010). Poor safety standards led to chilean mine disaster. August 29, 2010. [Online]. Available: http://www.globalpost.com/dispatch/chile/100828/mine-safety.
Bookstaber, R., Glasserman, P., Iyengar, G., Luo, Y., Venkatasubramanian, V., Zhang, Z., 2015. Process systems engineering as a modeling paradigm for analyzing systemic risk in financial networks. Off Financ Res Work Pap Ser 15 (01).
Brotman, B. (2014). With energy reports, it's game on. [Online]. Available: http://articles.chicagotribune.com/2014-03-03/features/ct-energy-comparisons-brotman-talk-0303-20140303_1_energy-hog-energy-efficiency-comed.
Browning, J.B. (1993). Union carbide: disaster at bhopal. In on managing under siege. detroit (MI) (pp. 1–15).
CAIB, 2003. Columbia accident investigation board report. report. Columbia Accident Investigation Board. [Online]. Available: http://www.slac.stanford.edu/spires/find/books?irn=317624
Carver, C.S., Scheier, M.F., 1982. Control theory: a useful conceptual framework for personality–social, clinical, and health psychology. Psychol. Bull. 92 (1), 111.
CERC, 2012. Report on the grid disturbance on 30th july 2012 and grid disturbance on 31st july 2012. Report.
CSB, 2005. Investigation report refinery explosion and fire. report. U. S. Chemical Safety and Hazard Investigation Board.
Cullen, W.D., 1993. The public inquiry into the piper alpha disaster. Report.
De Condorcet, N., 2014. Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix. Cambridge University Press.
FCIC (2011). The financial crisis inquiry report.
Han, J., Li, M., Guo, L., 2006. Soft control on collective behavior of a group of autonomous agents by a shill agent. J. Syst. Sci. Complexity 19 (1), 54–62.
Investigation of the challenger accident. (1986). Committee on Science and Technology House of Representative, Report, October 29, 1986.
Kennedy, J., 2010. Particle swarm optimization. In: Encyclopedia of Machine Learning. Springer, pp. 760–766.
Kiefer, J., Wolfowitz, J., et al., 1952. Stochastic estimation of the maximum of a regression function. Ann. Math. Stat. 23 (3), 462–466.
Krishnamurthy, V., Poor, H.V., 2014. A tutorial on interactive sensing in social networks. Computational Social Systems, IEEE Transactions on.

Kurokawa, K., Ishibashi, K., Oshima, K., Sakiyama, H., Sakurai, M., Tanaka, K., Tanaka, M., Nomura, S., Hachisuka, R., Yokoyama, Y., 2012. The official report of the fukushima nuclear accident independent investigation commission. report. The Fukushima Nuclear Accident Independent Investigation Commission.

Leveson, N., 2011. Engineering a Safer World: Systems Thinking Applied to Safety. MIT Press.

Luo, Y., Iyengar, G., Venkatasubramanian, V., 2016. Soft regulation with crowd recommendation: coordinating self-interested agents in sociotechnical systems under imperfect information. PLoS ONE 11 (3). E0150343.

Luo, Y., Iyengar, G., Venkatasubramanian, V., 2018. Social influence makes self-interested crowds smarter: an optimal control perspective. IEEE Trans. Comput. Social Syst. 5 (1), 200–209.

Maurya, M.R., Rengaswamy, R., Venkatasubramanian, V., 2003a. A systematic framework for the development and analysis of signed digraphs for chemical processes. 1. Algorithms and analysis. Ind. Eng. Chem. Res. 42 (20), 4789–4810.

Maurya, M.R., Rengaswamy, R., Venkatasubramanian, V., 2003b. A systematic framework for the development and analysis of signed digraphs for chemical processes. 2. Control loops and flowsheet analysis. Ind. Eng. Chem. Res. 42 (20), 4811–4827.

Maurya, M.R., Rengaswamy, R., Venkatasubramanian, V., 2004. Application of signed digraphs-based analysis for fault diagnosis of chemical process flowsheets. Eng. Appl. Artif. Intell. 17 (5), 501–518.

McAteer, J.D., Beall, K., Beck, J., McGinley, P., 2011. Upper big branch: the april 5, 2010, explosion: a failure of basic coal mine safety practices: Report to the governor. Report. Governor's Independent Investigation Panel.

Powers, W.T., 1973. Feedback: beyond behaviorism stimulus-response laws are wholly predictable within a control-system model of behavioral organization. Science 179 (4071), 351–356.

PresidentialCommission, 2011. Deepwater, the gulf oil disaster and the future of offshore drilling. Report. National Commission on the BP Deepwater Horizon Oil Spill and Offshore Drilling.

Rendell, L., Boyd, R., Cownden, D., Enquist, M., Eriksson, K., Feldman, M.W., Fogarty, L., Ghirlanda, S., Lillicrap, T., Laland, K.N., 2010. Why copy others? insights from the social learning strategies tournament. Science 328 (5975), 208–213.

Shmueli, E., Singh, V.K., Lepri, B., Pentland, A., 2014. Sensing, understanding, and shaping social behavior. Computational Social Systems, IEEE Transactions on.

Srinivasan, R., Venkatasubramanian, V., 1996. Petri net-digraph models for automating hazop analysis of batch process plants. Comput. Chem. Eng. 20 (96). S719–S725

Srinivasan, R., Venkatasubramanian, V., 1998a. Automating hazop analysis of batch chemical plants: part i. The knowledge representation framework. Comput. Chem. Eng. 22 (9), 1345–1355.

Srinivasan, R., Venkatasubramanian, V., 1998b. Automating hazop analysis of batch chemical plants: part ii. algorithms and application. Comput. Chem. Eng. 22 (9), 1357–1370.

Srinivasan, R., Venkatasubramanian, V., 1998c. Multi-perspective models for process hazards analysis of large scale chemical processes. Comput. Chem. Eng. 22 (98). S961–S964

Stephanopoulos, G. (1984). Chemical process control: an introduction to theory and practice.

Sunstein, C., Hastie, R., 2014. Wiser: Getting Beyond Groupthink to Make Groups Smarter. Harvard Business Review Press.

Surowiecki, J., 2005. The Wisdom of Crowds. Anchor.

TaskForce, 2004. Final report on the august 14, 2003 blackout in the United States and Canada. Report. US-Canada power system outage task force.

Thaler, R.H., Sunstein, C.R., 2003. Libertarian paternalism. Am. Econ. Rev. 175–179.

Thaler, R.H., Sunstein, C.R., 2008. Nudge: Improving Decisions About Health, Wealth, and Happiness. Yale University Press.

Trochim, W.M., Cabrera, D.A., Milstein, B., Gallagher, R.S., Leischow, S.J., 2006. Practical challenges of systems thinking and modeling in public health. Am. J. Public Health 96 (3), 538.

Vaidhyanathan, R., Venkatasubramanian, V., 1995. Digraph-based models for automated hazop analysis. Reliab. Eng. Syst. Saf. 50 (1), 33–49.

Vaidhyanathan, R., Venkatasubramanian, V., 1996. A semi-quantitative reasoning methodology for filtering and ranking hazop results in hazopexpert. Reliab. Eng. Syst. Saf. 53 (2), 185–203.

Venkatasubramanian, V., 2011. Systemic failures: challenges and opportunities in risk management in complex systems. AIChE J. 57 (1), 2–9.

Venkatasubramanian, V., 2018. The promise of artificial intelligence in chemical engineering: is it here, finally? AIChE J.

Venkatasubramanian, V., Rengaswamy, R., 2003a. A review of process fault detection and diagnosis: part i: quantitative model-based methods. Comput. Chem. Eng. 27, 293–311. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0098135402001606.

Venkatasubramanian, V., Rengaswamy, R., 2003b. A review of process fault detection and diagnosis: part ii: qualitative models and search strategies. Comput. Chem. Eng. 27, 313–326. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0098135402001618.

Venkatasubramanian, V., Rengaswamy, R., Kavuri, S.N., Yin, K., 2003. A review of process fault detection and diagnosis part iii: process history based methods. Comput. Chem. Eng. 27 (3), 327–346. 654CE Times Cited:423 Cited References Count:118. [Online]. Available: ⟨GotoISI⟩://WOS:000181475400003.

Venkatasubramanian, V., Vaidhyanathan, R., 1994. A knowledge-based framework for automating hazop analysis. AIChE J 40 (3), 496–505.

Venkatasubramanian, V., Zhang, Z., 2016. Tecsmart: a hierarchical framework for modeling and analyzing systemic risk in sociotechnical systems. AIChE J..

Venkatasubramanian, V., Zhao, J.S., Viswanathan, S., 2000. Intelligent systems for hazop analysis of complex process plants. Comput. Chem. Eng. 24 (9–10), 2291–2302.

WHO (2006). Sars: how a global epidemic was stopped, report. [Online]. Available: http://www.tandfonline.com/doi/abs/10.1080/17441690903061389.

Zhang, H., Parkes, D.C., 2008. Value-based policy teaching with active indirect elicitation. AAAI 8, 208–214.