Paying Attention to the Man behind the Curtain

Mary (Missy) L. Cummings and Kristopher M. Thornburg

In the push to develop smart energy systems, designers have increasingly focused on systems that measure and predict user behavior to effect optimal energy consumption. While such focus is an important component in these systems' success, designers have paid substantially less attention to the people on the other side of the energy system loop—the supervisors of power-generation processes. Out of sight and sadly, in terms of technological advancements, out of mind, many operators perform high-risk jobs in data-rich but information-impoverished settings.

The Three Mile Island nuclear power plant accident in 1979 was caused primarily by operators misunderstanding sensor data in an overwhelmingly complex control panel.[1] In 2003, operators in the Northeast couldn't see or understand nearby power grids' critical system states, which ultimately led to the largest blackout in North American history, contributing to at least 11 deaths and costing an estimated US$6 billion.[2] In these high-profile cases, and in countless other electric and nuclear power plant incidents, a significant problem was the lack of explicit design to support rapid data aggregation and information visualization for operators' time-pressured decision making.

Smart energy systems that leverage pervasive computing could add to these supervisory control operators' workload. They'll have to predict possible power plant load and production changes caused by environmental and plant events, as well as dynamic system adaptation in response to consumer behaviors. Contrary to many assumptions, inserting more automation, including distributed sensors and algorithms to postprocess data, won't necessarily reduce operators' workload or improve system performance.

SUPERVISORY CONTROL AND WORKLOAD

Current power-generation operations are highly automated. In normal, day-to-day operations, automation adjusts system parameters, with human operators generally acting as system supervisors, monitoring system states and typically intervening in non-monitoring operations, such as responding to an alarm, managing a plant start-up, or overseeing other off-nominal operations. Although the system is highly automated, these supervisors receive little automated decision-making support, especially in time-critical, system anomaly situations. Indeed, although digital displays are replacing analog ones in current control rooms, many plant displays, particularly in the nuclear reactor realm, simply replicate analog displays, effectively keeping 1960s-era control rooms' look and feel.

Inserting more automation, including distributed sensors and algorithms to postprocess data, won't necessarily reduce operators' workload or improve system performance.

Central to a system control paradigm with high automation levels is *human supervisory control* (HSC), which assumes that a human operator will monitor a given system, taking the role of system supervisor or manager.[3] Again, operators don't control low-level system actions, though they can intervene when the situation requires. This relationship between the system and an operator is called *human on the loop*, rather than *human in the loop*, directing focus away from constant, direct control and toward supervisory control.

Figure 1 depicts an HSC conceptual model with two human operators controlling two plants.[4] This model indicates that the automated system is responsible for controlling the physical system (labeled "Plant"). Human operators supervise and interact with the system only through automation. In addition to supervising the plant, human operators are responsible for monitoring and synthesizing different types of information coming from the smart-grid system to ensure safe and efficient plant operation.

Introducing pervasive computing

Published by the IEEE CS ■ 1536-1268/11/$26.00 © 2011 IEEE

in smart-grid settings could increase supervisory control operators' workload, especially those without any advanced decision support. A seemingly reasonable solution is to use more automation in grid management. Automation is commonly introduced into HSC systems to reduce the likelihood of operator error by reducing operator workload. However, because increased automation often just changes the nature of the work, this relationship isn't universally true.[5]

Raja Parasuraman and his colleagues reviewed several studies that suggest aircraft cockpit automation has actually increased—rather than decreased—operator workload, as intended.[6] Interestingly, little evidence supports the idea that operators will delegate tasks to automation when workload is high.[7,8] Whether we can generalize these relationships to the power-generation domain is unclear. However, several similarities exist between automated aircraft control and power-generation systems. Monitoring an aircraft on autopilot contains similar tasks as monitoring a power-generation system. Alarms or warnings activate when the system reaches a state outside predefined parameters, and operators are then expected to take over some level of control of the system, usually with increased time criticality and serious consequences if mistakes are made.

The Three Mile Island accident and the 2003 Northeast blackout required operators to move from a monitoring state to an emergency-action mode with significant time pressure. In both incidents, operators' lack of understanding of the automation, often called *mode confusion*, exacerbated the problem. Mode confusion occurs when an HSC operator attempts to take control of a highly automated system but doesn't understand the current automation mode, including its objectives. In both aviation and power-generation systems, this lack of understanding has caused catastrophic human-system failure because of confusion over who is in control, especially when an operator's desired goal state differed from that of the automated system.

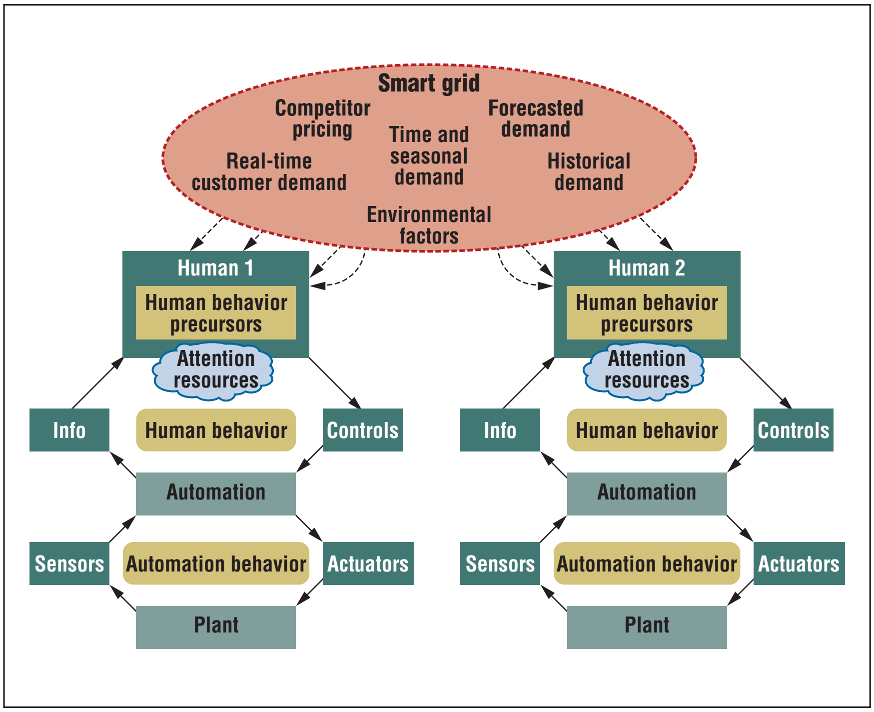

Figure 1. Human supervisory control (HSC) model for power-generation systems. Two power plants are connected through a smart grid via two HSC networks. Automation controls the actual plant processes; however, in future smart grids, a human supervisor in each plant will make changes on the basis of not only local plant information but also predictive demand models and other plants' real-time changes.

In smart energy system development, the question remains of how power plant supervisors will respond to the inevitable addition of automation. HSC operators will need to understand how a smart energy management system could affect the safety and efficiency of power-generation processes. Significant automation will be necessary to both manage these processes and represent data so the operator understands what the automation is doing. Such nested automation layers increase system opacity, and this lack of transparency, including a lack of sufficient and intelligible feedback, is a causal factor of mode confusion.[5,6]

Smart energy system designers must consider the unintended consequences of possible mode confusion and increased workload for HSC operators caused by increased automation, particularly in emergency scenarios. This requires more advanced display technology in the form of integrated software-decision-support tools that leverage more advanced information visualization and data fusion techniques for both current and predicted state representations. Furthermore, smart energy system design and operation will require sociotechnical changes; organizational and regulatory policies and procedures must be updated or changed outright. For example, it remains unclear how smart-grid technologies meant to improve efficiency—but ultimately linked to power-generation safety—will influence operating procedures and certifications.

HUMAN OVERSIGHT OF AUTOMATED PLANNING

The assumption that increased automation can reduce operator workload in smart energy systems is not only naïve

SPOTLIGHT

TABLE 1
The Fitts list for human-computer role allocation.

Humans are better at	Computers are better at
Perceiving patterns	Responding quickly to control tasks
Improvising and using flexible procedures	Repetitive and routine tasks
Recalling relevant facts at the appropriate time	Handling simultaneous complex tasks
Reasoning inductively	Reasoning deductively
Exercising judgment	Fast and accurate computation

in terms of workload management but also ignores the critical role that the human operator plays in supervisory control systems. Humans can apply reasoning in situations in which automation can't.

A critical aspect in integrated sociotechnical system design with significant embedded autonomy is role allocation—who (automation and/or human) should perform which functions and when. According to early research examining human-computer allocation in the air traffic control domain, humans and computers possess the respective strengths listed in Table 1, known as the Fitts list.[9] This early attempt at role allocation for humans and computers recognized that automation can support—not necessarily replace—human operators in large-scale computational decision-making tasks.

As Table 1 depicts, algorithms can execute repeatable, precise, and speedy computations, which is ideal for complex optimization problems such as those inherent in pervasive computing and smart-grid environments. However, automation can be inflexible and unable to adapt to changing situations. Though computer optimization algorithms are fast and can handle complex computation far better than humans, they're notoriously "brittle" in that they only take into account those quantifiable variables identified as critical during the design stages.[10,11]

In contrast, humans can improvise, learn, and reason inductively, which are precisely the skills required to adapt to unexpected circumstances. This type of problem solving is called *knowledge-based reasoning*, during which humans make decisions under novel and uncertain situations—attributes inherent in supervisory control scenarios.[12] In terms of managing the large data streams that smart-grid environments will generate, automation will be critical in handling the bulk of problem solving and system management. However, as the 2003 Northeast blackout illustrates, even highly automated systems can encounter dynamic and unexpected variables that designers don't anticipate, which can ultimately lead to catastrophe. So, whereas smart grids will be highly automated, with embedded complex algorithms to balance power input and output across a network, they won't be completely automated, primarily due to the inherent uncertainty in both the environment and the algorithms themselves.

Although much supervisory control literature has focused on keeping the human in the loop for potential interventions for low-probability events such as a blackout, we know significantly less about how human operators can provide value in assisting embedded algorithms to optimize system performance, which is the crux of smart-grid operations. For smart energy systems, every consumer represents a node that might not always behave in an expected manner, which can cause problems when expectations and forecasts don't match actual operating conditions. Given the complexity of a large problem space with layers of uncertainty, it's unclear whether algorithms will always perform optimally in all conditions across what is effectively a decentralized network. Little research has examined how operators in decentralized energy networks can aid algorithms in optimizing system performance. But, recent research in supervisory control of unmanned vehicles sheds light on the capabilities of humans working collaboratively with algorithms to achieve superior system optimization performance.

The US military envisions networks of decentralized unmanned vehicles (including air, ground, sea surface, and subsurface) that work together, with a human on the loop, to conduct resource allocation missions such as using an array of unmanned vehicles to search remote, possibly hostile areas for enemies or victims. In these networks, each unmanned vehicle computes its best plan using local negotiation with other unmanned vehicles. There's no globally optimal plan because each vehicle strives to maintain the best plan with possibly limited information. In contrast to a centralized approach, this decentralized approach protects against network vulnerabilities caused by bandwidth limitations and avoids reliance on specific vehicles for critical tasks.

These decentralized network attributes have direct mappings in smart-grid energy environments, wherein utilities could create smaller, decentralized spheres of localized smart-grid control, possibly managed by home and building owners. Similar to the decentralized unmanned vehicle network, these smaller spheres could allow for local resource optimization without requiring more complex and resource-intensive globally optimal network solutions at substations or utility-managed command centers. However, such nodes of local control still require some supervisory oversight, particularly in anomalous situations, such as major power outages and extreme weather conditions.

Because relatively high automation levels are essential to operate these decentralized networks, and humans are necessary for system oversight, questions remain of just how much human collaboration should be allowed and what the impact of human interaction could be for such a system. To partially address these issues, an experiment examined how well a decentralized vehicle network would perform without human oversight, as compared to a system with a human operator who was allowed to tweak the automation's resource allocation and scheduling plans.[13]

In this experimental setting, a network of five unmanned vehicles was tasked with searching as much of a predetermined area as possible, and tracking targets using a mixed-integer linear programming algorithm. In the automation-only condition, the system generated all the plans, which were automatically approved, and a human operator never changed the tasking or rate at which it generated plans. In the second condition with both humans and automation, humans could update the algorithm's tasking or replan more often if they thought the automation wasn't performing adequately.

Figure 2 demonstrates how much value the human provided in terms of the two primary dependent measures—percentage of area covered and number of targets found. Investigation of three different workload levels (30, 45, and 120 seconds between replanning intervals) determined how both the automated planners and the operators would respond under changing workload conditions.

As Figure 2 shows, letting a human operator evaluate and occasionally change an automated solution allowed the system to perform substantially better than if the automation was left alone. Of the six conditions in Figure 2, the 120-second interval for the area-searched metric was the only automation approach statistically comparable to the human-assisted mode, suggesting that the longer intervals between replanning benefitted the automation. For the targets-found metric, the collaboration between the human and the automation resulted in more than a 20 percent increase across all factor levels.

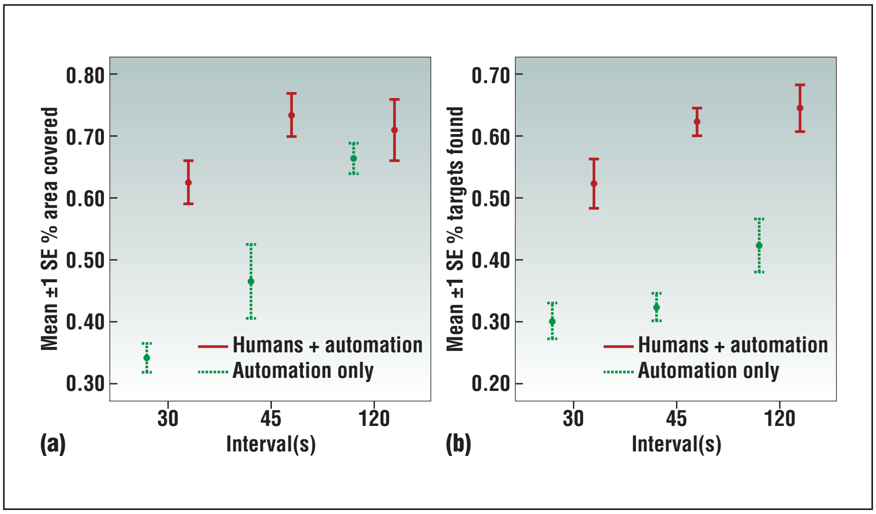

Figure 2. Value added by letting humans work with a planning and scheduling algorithm. (a) Decreasing workload levels represented by 30-, 45-, and 120-second replanning intervals are plotted against the percentage of geographic area that a group of unmanned vehicles searched. (b) The same replanning levels are plotted against the number of targets the unmanned vehicles found. Automation assisted by a human operator performed significantly better than the automation alone.

Although these results are for a decentralized unmanned vehicle planning and scheduling problem, they highlight the importance of understanding the benefit of human interaction in systems that use automation for decentralized scheduling and resource allocation, which is likely the future of smart grids. The human operator was critical in this domain because of the uncertainty inherent in the system. The autonomous planners on the unmanned vehicles operated with a priori cost functions coded by the algorithm designers, which theoretically generated an optimal solution, but in reality could be improved by occasional human judgment.

This temporal component of human interaction is important because previous related research has shown that if operators intervene too much in such distributed planning systems, overall system performance could suffer.[14] Determining a robust range of helpful human interaction is key. In addition, whereas military command and control settings possibly contain more uncertainty than power-generation settings, many sources of uncertainty in smart-grid system management—such as weather, customer behaviors, algorithm design, and system failures—can lead to similar problems.

In the envisioned future of smart energy systems, pervasive computing systems will measure and infer user behavior to mitigate and optimize energy use. Such systems will require significant embedded algorithms as well as some level of human supervisory control—the proverbial men and women behind the curtain. In domains in which uncertainty exists, including probabilistic algorithms and behavioral inference, considering the human role not only as a monitor of anomalous system states but also a collaborator is critical.

It's generally recognized that in power-generation environments, automation is necessary in safety-critical

SPOTLIGHT

monitoring tasks such as fault detection, situation assessment and diagnosis, and response planning.[15] Unfortunately, no organized effort that we're aware of focuses on developing algorithms and associated decision tools to support supervisors managing dynamic and adaptive smart grids. Such research is necessary to determine the required degree of interactivity between supervisory control operators and a pervasive computing system's automation, how to manage the voluminous data streams that smart energy systems will generate, and how to balance the competing objectives of safety and optimal energy production.

Mary L. (Missy) Cummings is an associate professor at the Massachusetts Institute of Technology's Aeronautics and Astronautics Department and director of the Humans and Automation Laboratory. Contact her at missyc@mit.edu.

Kristopher M. Thornburg is a postdoctoral associate at the Massachusetts Institute of Technology's Humans and Automation Laboratory. Contact him at k_thorn@mit.edu.

REFERENCES

1. J.S. Walker, *Three Mile Island: A Nuclear Crisis in Historical Perspective*, University of California Press, 2004.
2. J. Minkel, "The 2003 Northeast Blackout—Five Years Later," *Scientific American*, Aug. 2008; www.scientificamerican.com/article.cfm?id=2003-blackout-five-years-later.
3. T.B. Sheridan, *Telerobotics, Automation and Human Supervisory Control*, MIT Press, 1992.
4. P.E. Pina et al., "Identifying Generalizable Metric Classes to Evaluate Human-Robot Teams," *Proc. 3rd Ann. Conf. Human-Robot Interaction*, ACM/IEEE Press, 2008, pp. 13–20.
5. C.E. Billings, *Aviation Automation: The Search for a Human-Centered Approach*, Lawrence Erlbaum Associates, 1997.
6. R. Parasuraman, T.B. Sheridan, and C.D. Wickens, "A Model for Types and Levels of Human Interaction with Automation," *IEEE Trans. Systems, Man, and Cybernetics Part A: Systems and Humans*, vol. 30, no. 3, 2000, pp. 286–297.
7. V. Riley, "A Theory of Operator Reliance of Automation," *Human Performance in Automated Systems: Current Research and Trends*, M. Mouloua and R. Parasuraman, eds., Lawrence Erlbaum Associates, 1994, pp. 8–14.
8. R. Parasuraman and V. Riley, "Humans and Automation: Use, Misuse, Disuse, Abuse," *Human Factors*, vol. 39, no. 2, 1997, pp. 230–253.
9. P.M. Fitts, *Human Engineering for an Effective Air Navigation and Traffic Control System*, Nat'l Research Council, Division of Anthropology and Psychology, Committee on Aviation Psychology, 1951.
10. B.G. Silverman, "Building a Better Critic: Recent Empirical Results," *IEEE Expert*, vol. 7, no. 2, 1992, pp. 18–25.
11. P. Smith, E. McCoy, and C. Layton, "Brittleness in the Design of Cooperative Problem-Solving Systems: The Effects on User Performance," *IEEE Trans. Systems, Man, and Cybernetics Part A: Systems and Humans*, vol. 27, no. 3, 1997, pp. 360–371.
12. J. Rasmussen, "Skills, Rules, and Knowledge; Signals, Signs, and Symbols, and Other Distractions in Human Performance Models," *System Design for Human Interaction*, A.P. Sage, ed., IEEE Press, 1987, pp. 291–300.
13. M.L. Cummings et al., "The Impact of Human-Automation Collaboration in Decentralized Multiple Unmanned Vehicle Control," to be published in *Proc. IEEE*, 2011.
14. M.L. Cummings, A. Clare, and C. Hart, "The Role of Human-Automation Consensus in Multiple Unmanned Vehicle Scheduling," *Human Factors*, vol. 52, no. 1, 2010, pp. 17–27.
15. J.M. O'Hara et al., *Human Factors Considerations with Respect to Emerging Technology in Nuclear Power Plants* (NUREG/CR-6947), Nuclear Regulatory Commission, 2008.