



HUMAN-CENTERED COMPUTING

Editors: **Robert R. Hoffman, Jeffrey M. Bradshaw, and Kenneth M. Ford,**
Institute for Human and Machine Cognition; rhoffman@ihmc.us

Command and Control Requirements for Moving- Target Defense

**Marco Carvalho, Jeffrey M. Bradshaw, Larry Bunch, Tom Eskridge, Paul J. Feltovich,
and Robert R. Hoffman, Institute for Human and Machine Cognition
Daniel Kidwell, Department of Defense**

The Moving-Target Rule of macrocognitive work systems, as introduced by Sidney Dekker and his colleagues in 2003,¹ states that:

The sociotechnical workplace is constantly changing, and constant change in environmental constraints may require constant adaptation in cognitive work, even if domain constraints remain constant.

This rule is a consequence of the bounded ecology of macrocognitive work systems.² A work system can never match its environment completely; there are always gaps in fitness, because what it means to be fit is itself a moving-target. There is always a struggle to adapt, which can ease or intensify as events unfold. The adaptive capacity of a macrocognitive work system is therefore a tradeoff between optimality and resilience: increasing the scope of what counts as “routine” merely serves to increase the opportunities for surprise. Resilience requires a capacity to adapt to surprising events, but the ability to anticipate surprise requires additional resources whose contribution might be missed and at any given moment might be mistaken for inefficiency.

In this article, we look at a domain where the workplace is a moving-target in three ways:

- new technology and work methods are continually being introduced,
- domain constraints are not constant—the work itself is changing in terms of its new goals and requirements, and
- next to nothing is routine and anything can be surprising.

This is the domain of cyberdefense. Cyberdefense is a constant game of “catch-up” or “staying

one step ahead,” but events can transpire at a rate exceeding the capabilities of humans to comprehend and make decisions. And because of the new demands posed by clever adversaries, new technologies are being developed and injected into the work. So, the three ecological constraints amplify one another.

New tools and approaches are being developed to help cyberworkers cope with the problems of data volume, work tempo, and problem complexity. Many tools embody a designer-centered design perspective, where algorithms and visualizations are used to query databases or display ongoing data statistics. The development of such systems should hinge on a deep understanding of cyberworkers’ needs in the context of their real work. *Moving-target network defense* provides a case study of some of these problems and how they can be approached from a human-centered perspective.

What Is Moving-Target Defense?

Current computer network systems typically operate with a relatively static layout and configuration of applications and services. Once such systems are deployed, attackers can observe, probe, and study them over long periods of time seeking potential vulnerabilities or entry points, without worrying about the acquired system information becoming stale.

Moving-target defense has been proposed as a game-changing capability for the protection of important network systems and critical computing infrastructure.³ Moving-target defense lets networked computers change their structure and configuration dynamically while maintaining their functionality and availability to legitimate users.

Based on the analogy of homogeneity and diversity in ecological systems, moving-target defense avoids unnecessary consistency. The key idea is to exploit diversity as a way to hinder an aggressor from using easily replicated attacks that rely on the static nature of computer networks to maintain target viability. The goal of diversity defenses is to present attackers with an uncertain and unpredictable target. If the target changes quickly enough, it will be too difficult for attackers to succeed in their malicious intent.

Moving-target defenses rely on two broad sets of capabilities:

- low-level tools and mechanisms to support mobility of processes and communication, application and operating system diversity, and monitoring; and
- high-level command and control (C2) mechanisms that implement the logic for system mobility and adaptive response to failures and attacks.

While the concept of moving-target defense for cybersecurity is relatively new, defense mechanisms implementing similar principles were introduced more than 10 years ago.⁴ Rather than focusing on low-level moving-target defense capabilities that have been covered extensively elsewhere, here we discuss the requirements and capabilities that will be needed for a next-generation high-level C2 mechanism for moving-target defense—and, in the process, explore some lessons about human-centered computing that might be applied more generally.

Moving-Target Defense for C2

Moving-target defense for C2 must focus on two competing aspects of human-computer systems. First, the

C2 must present a clear picture of the status of the network and attacks to operators so that awareness is maintained. Second, the C2 must act and react quickly enough to ensure that the network defenses will be effective. In a proactive mode, moving-target defense could make scheduling changes at periodic or aperiodic time intervals, thereby making the system more unpredictable and difficult for attackers to map. Proactive moving-target defense operates independently from perceived attacks or other stimuli. For example, the port and address hopper in BBN Technologies' Applications That Participate in Their Own Defense project randomly changes the IP address that clients use to communicate with a given service every few seconds, thereby making it difficult for malicious scanners to pinpoint service locations.⁵

In a reactive mode, a moving-target defense responds to detecting an intrusion, observing harmful behavior, or sensing a major network equipment state change. It may reorganize its services and diversify its configuration to improve resilience while maintaining mission continuity.⁶

Previous Moving-Target Defense Work

Some of the early adaptive defense capabilities that implemented the concept of moving-target were designed to operate on a predefined mobility pattern or in a closed control loop. The Intrusion Tolerance by Unpredictable Adaptation (ITUA) project, for example, proposed middleware for intrusion tolerance based on network and service adaptation, and used a set of predefined control loops to respond to specific security events (such as specific firewall or intrusion detection events).⁵ Upon detection of a security event, the ITUA middleware would distribute the fault information to all

nodes involved in the defense control process, leading to a potential termination and reinstantiation of the faulty process.

Figure 1a shows a generalization of the control process, where the box labeled "Mobility space" represents the actual network and service resources as well as the specific network and service adaptation mechanisms. In the case of ITUA, for example, the mobility space would represent the services and communication resources of the system, and the communications middleware used for event sharing, coordination and service control. The C2 component represents the coordination algorithm used to control the mobility space. It receives feedback from the mobility space through a set of network and service defense monitoring tools. The external inputs to the system under control (that is, the service and communication resources in the mobility space) are the interactions with external clients, both legitimate and malicious.

This general formulation describes nearly all current moving-target defense implementations, which tend to rely solely on automated control mechanisms. Most of these tools are designed merely to collect information for an analyst or to automatically respond to attacks and system disruptions. The role of the analyst is one of controlling the tools that control the network, rather than using the tools to understand and control the network.

Resilient Human-Automation Teamwork

Most prior and current approaches to moving-target defense—and cyberdefense in general—tend to focus on specific mechanisms, with the analyst being relegated to the role of compensating for various shortcomings in

the opaquely constructed automated control loop rather than being part of the perception, decisions, and actions taking place within the loop itself.¹ We assert that better, more resilient performance can be obtained by leveraging the capabilities that humans and automation can bring jointly to the proactive anticipation of, and response to, cyberevents—so long as the system is designed to support such teamwork.

Capabilities that take advantage of the power of human-automation teamwork are especially important in moving-target defense, where the mobility of the computational and communications infrastructure can be adjusted in response to changes in context and the dynamics of the situation. As illustrated in Figure 1b, the C2 component of the moving-target defense can be designed from a human-centric perspective, to support the capability to create uncertainty on the part of the attacker and to obfuscate the protected computing and network infrastructure. Building on previous research,⁷ we believe that moving-target defense tools and mechanisms may work best when humans can keep the technology aligned to context, monitor their progress, verify their ongoing effectiveness, and contribute the human powers of perception and decision-making to the work.

Moving-target defense systems must not only allow for emergent phenomena, they must be designed to allow emergence to occur. They must not only allow for resilience, they must be designed to insure resilience by incorporating semantically rich models of human-machine teamwork.

A Sensemaking Strategy for Moving-Target Defense

One of the major activities of the analyst is *sensemaking*, a continuous

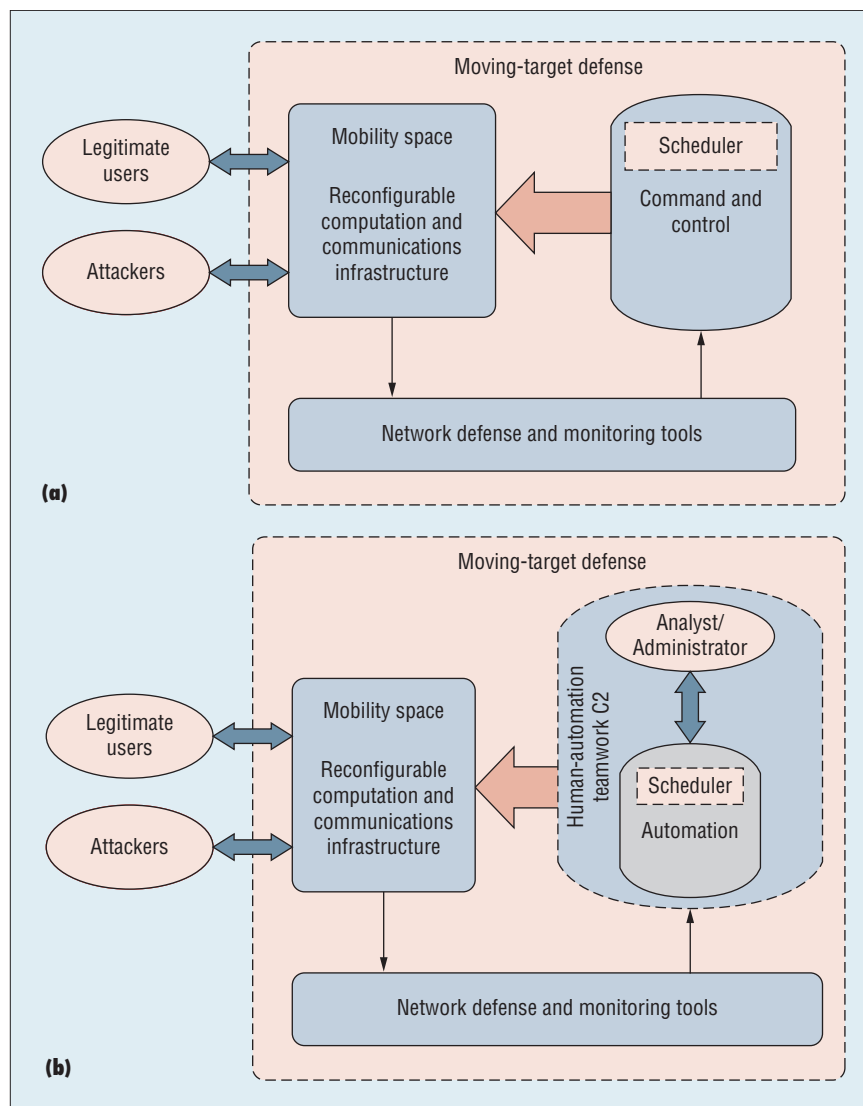


Figure 1. The command and control (C2) component controls the resources in the mobility space on the basis of feedback from monitoring tools. (a) Most current implementations rely on automated mechanisms; (b) a human-agent teamwork approach to the C2 component can make the system more resilient.

effort to understand how the relevant components of the world connect and interact so that their future behavior may be anticipated and acted upon. Current applications of sensemaking theory to intelligence analysis have focused on ways to shape the sensemakers' investigative steps to help them validate lines of reasoning and counteract misconceptions.^{8,9} Building on this foundation, the next step has to be toward implementation. In particular, we need an understanding of the potential

impact of new forms of visualization and automation on the sensemaking process and of how such tools ought to be designed in light of what we already know. The emphasis of our own work on sensemaking is to put questions about the role and benefits of computer interaction with people front and center.¹⁰

In light of the current emphasis on validation using multiple methods within the sensemaking literature, the question for the system designer becomes not only, "How can we help

analysts know whether their hypotheses are correct?” but also, “How can we, to the greatest possible degree, use visualization, automation, and collaboration tools to help them expose their hypotheses to the light of experience and inquiry, in order to evaluate and refine them as thoroughly as possible?” In complex and high-tempo work, we can’t afford anything less than full engagement of the perceptual strengths, experience, and know-how manifested in both humans and automation as we grapple with the increasing number and severity of cyberattacks.

In systems theory, *emergence* denotes the phenomenon whereby unexpected phenomena or behaviors arise from interactions among the system’s functional components. The classic example from philosophy is the impossibility of predicting the properties of water on the basis of knowledge of the properties of hydrogen and oxygen. Higher up the scale of cumulative science is the classic example of consciousness, often said to be unpredictable on the basis of knowledge about neurons. Complex systems that show emergence can nonetheless be rule-governed. A classic example is the complex patterns that emerge from simple equations, such as fractals, where the initial conditions can lead to unexpected outcomes. Systems can be bounded by structural and environmental constraints, described as rule governance, but the interactions of the components can result in things that are new.

At a higher level of emergence, new rules or patterns of interaction arise as a result of the operation of the rules that initially governed the system’s functional components. “The system is able to detect, amplify, and build upon emergent behavior. The latter can only happen by operating

on the behavior programs that causally influence behavior, similar to the way genetic evolution operates on the genes.”¹¹ This can be illustrated by analogy to human culture. The normative pressure of culture serves to reduce the number of alternatives available for acceptable behavior in a given setting, thus greatly simplifying the problem of human choice and interpredictability of players in routine situations. However, the governing constraints of culture are themselves subject to change.

In the context of macrocognitive work, *coactive emergence* is the phenomenon by which the work system not only changes what it does but changes how it changes, as an adaptation to circumstances that had not been previously encountered. However, in contrast to systems that effect second-order changes in response to environmental influences that are indifferent to the objectives of the system, in this context the humans and software members of a team deliberately seek to influence the direction of adaptations in ways that converge on shared objectives. The term “coactive” is meant to emphasize this explicit aim of synergy in joint development of a common set of relevant hypothesis for a given situation by humans and machines working interdependently. In other words, coactive emergence occurs when the emergence is a result of the interactions of functional subsystems, including the human agents whose interactions influence one another as a part of what it means to work collaboratively to achieve shared goals.

Building on this concept, we take interdependence as the central organizing principle among people and machines (that is, software agents) working together.¹² However, while coactive design has primarily emphasized the role of interdependence in

joint activity to support *doing* things in the world, moving-target defense is meant to support cognitive work that involves *making sense* of a given situation. Our emphasis on interdependent activity of humans and machines proceeds from the premise that a proper interface and regulatory framework can better support moving-target cyberdefense, increasing the range, richness, and utility of models beyond those that could be explored by humans or agents alone.

What Would a Moving-Target System Look Like?

Normally, cyberworkers use what they know in order to discover new patterns of attacks and then to define software components to detect and monitor the new patterns. Once created, the automated software components simply notify analysts of matches, flagging them for further investigation.

In a coactive environment, software agents can still be created by cyberworkers to perform specific monitoring and control tasks. However, software agents can also create new hypotheses to support or modify the understanding of the cyberworker and to provide evidence that the human can accept or reject. Such interdependent exploration could be expected to improve the sensemaking process and the collaborative monitoring and control of the computer network infrastructure and its defenses.

A collaborative human-machine C2 system for cyber-sensemaking and computer network defense would enable the following capabilities:

- new (collective) human and machine capabilities through an ongoing iterative cycle, combining agent learning of threat patterns and

positive or negative reinforcement by the analyst;

- new lines of reasoning and hypotheses based on what analysts and agents are learning together about anticipated threat patterns;
- new policies and shared expectations as human analysts direct and redirect agent activities, based on appraisals of progress on assigned tasks reported by agents; and
- new patterns in shared representations (such as visualizations) as analysts and agents progressively converge on useful interpretations of the situation, identifying and characterizing new kinds of threats as their outlines gradually emerge from seemingly chaotic background chatter.

In interdependent activity, both top-down policy constraints and bottom-up individual behavior are simultaneously shaped, enabling continuous adaptive refinement—that is, coactive emergence. In addition to shaping human-machine collaboration, top-down policies may regulate, for example, constraints on software agents working with specific kinds of data in specific contexts.¹³ In addition to shaping human-machine collaboration, top-down policies may regulate, for example, constraints on how software agents work with specific kinds of data in specific contexts. They may also be used to support service orchestration, define resource utilization constraints and system performance bounds, and regulate agent lifecycle and mobility activities.

Bottom-up emergence, on the other hand, would take the form of novel actor strategies for achieving individual tasks to refine and optimize. By allowing individual software agents in the system to adapt within policy bounds, and with feedback from humans, the moving-target defense

infrastructure would enable the bottom-up emergence of new policies.

With regard to the moving-target nature of cyberwork, the goal is to rapidly propagate lessons learned about productive and unproductive actions and to avoid undesirable states and events. In their complementary role, machines have the potential to help people cope, for example, with the volume, tempo, computational complexity, and highly distributed nature of joint tasks. In addition to supporting appropriate aspects of taskwork, agents can be used to help support coordination and other aspects of team process.

In mixed human-agent teams, people occupy a privileged position because, among other things, they generally know more about the way joint tasks interact with broader ongoing activities and with the situation at large. For these reasons, humans have an important role in keeping machine taskwork aligned with its wider contexts.⁸ That said, the machines themselves must be able to model the relationships between taskwork, teamwork, and the relevant situation.

Visualizations will likely be one of the primary means of communicating the structure and performance of such jointly constructed models.¹⁰ This communication must be bidirectional: the structure of the visualization communicates to the analyst the structure and processing performed by the machines (software agents in the system), while at the same time the user interactions with the visualization are used to direct and modify the activity of the software agents. For example, software agents that classify incoming NetFlow data into classes such as “normal,” “communications to whitelisted sources,” and “communications from blacklisted sources” may be visualized as a flow between a source and three sinks.

Adding a fourth sink representing the intersection between the second and third classes can both be informative to the operator and instruct the agent system to spawn an agent to look for and tag such an intersection.

This proposed framework relies not only on the collaboration between analysts and software agents but also on the collaborative work of multiple software agents, organized as teams to address specific tasks. For example, a special parser may be combined with a data source provider and an intrusion detection component (all implemented as software agents) to create an aggregate capability for intrusion detection. The organization of software agents may be defined through some combination of explicit policies and self-organization. Self-organizing strategies for software agents can be created to provide higher levels of resilience to the composed capability.

Resilience and Semantically Rich Policy Governance

Our approach relies on software agents that (ideally) support graceful, robust, and adaptive performance in the face of stressors and surprise. Policies define the operational ranges of different services and the tradeoff strategies between different compositions. For example, self-organized capabilities might have a specific response time or control delay, depending on the capability. Policies regulate the behavior of software agents or composed capacities—for example, by enforcing their start-stop, suspend-resume, and move functions. Policies could also express regulatory constraints on the self-organization process that are important and necessary to avoid self-inflicted denial under abnormal or unfeasible configuration requirements.

Resilience of the defense infrastructure would derive from the capacity to

quickly organize and assemble computation capabilities from multiple agents with specific functions. A self-reorganizing system would provide a set of redundant execution paths and fallbacks in response to localized failures and attacks. It would also have to be able to recover from partial disruptions that could occur in localized failures, and to restore its capacity to absorb failures. To implement these capabilities, the collection of software agents with basic functionalities must exchange information and organize in different ways, ensuring a desired outcome for the composed system.

For example, consider an intrusion detection capability that reports enriched details of security events. In addition to specific parsers and detectors, it must have access to lookup services (for reverse address lookup) and other supporting directories such as registration services and blacklists. To enable this joint capability, we could predefine the necessary combination of services (or software agents providing the services) and their interactions. But alternatively, we could let individual components signal, locate, and negotiate connections among themselves on the basis of their intrinsic capabilities, availability, and locality.

Effectively, such a designed-in emergence capability might allow cyberworkers to define higher-level descriptions of the capabilities desired—not just a specific structure of predefined components and interconnections—that other software agents or subsystems could provide.

The signaling between components would rely on the exchange of promote/inhibit messages announced for different capabilities, specifications of capability requirements, and bindings to aggregate capabilities. Software components might also differentiate when their internal software

design allows for the instantiation of specific functionalities. This concept is conceptually similar to the notion of multipotentialiation in biological systems, in this case generated by feedback among software components (promote/inhibit messages) and redundant associations between components.

Policy constraints are typically understood as governance on individual software agents. But the software agents in a moving-target defense system must be collectively obligated to perform certain tasks, such as the creation of a composed capability for intrusion detection and data enrichment. Furthermore, the specific responsibilities assigned to individual software agents aren't completely sorted out in advance, even though groups of agents within the work system can be collectively responsible for jointly executing various tasks. The goal is to allow the agents to self-organize within the constraints of their individual capabilities and current availability, while building redundant and potentially diverse execution paths.

Collective obligations explicitly represent a given agent's responsibilities within an agent group to which it belongs, without specifying in advance which agent must do what.¹⁴ In other words, the agent group as a whole becomes responsible, with individual members of the group sharing the obligation at an abstract level.

The execution and enforcement of collective obligations requires different mechanisms for different contexts. For some applications, a specialized planning system spanning a group of agents may be the best approach. However, our approach requires that the agents themselves, rather than some centralized capability, organize the work. The work system must be self-organizing so that the agents can revisit responsibilities

and resource allocations as needed, on an ongoing basis.

Combined, these capabilities may enable a new approach to a moving-target defense C2, letting humans closely collaborate with software components that are themselves adaptive and self-organizing, to enable the capabilities required by the cyberworker. A human-automation teamwork approach, relying on the principles of interdependent activity and resilience based on design for emergence via semantically rich policy governance, might be a path to coping with the moving-target problems of cyber-defense that mandate the creation of a work system that is itself a moving target. ■

References

1. S.W.A. Dekker, J.M. Nyce, and R.R. Hoffman, "From Contextual Inquiry to Designable Futures: What Do We Need to Get There?" *IEEE Intelligent Systems*, vol. 18, no. 2, pp. 74–77.
2. R.R. Hoffman and D.D. Woods, "Beyond Simon's Slice: Five Fundamental Tradeoffs That Bound the Performance of Macrocognitive Work Systems," *IEEE Intelligent Systems*, vol. 26, no. 6, pp. 67–71.
3. S. Jajodia et al., *Moving-Target Defense: Creating Asymmetric Uncertainty for Cyber Threats*, Springer, 2011.
4. D. Kewley et al., "Dynamic Approaches to Thwart Adversary Intelligence Gathering," *Proc. DARPA Information Survivability Conf. & Exposition II (DISCEX 01)*, vol. 1, IEEE, 2001, pp. 176–185.
5. M. Atighetchi et al., "Adaptive Use of Network-Centric Mechanisms in Cyber-defense," *Proc. 6th IEEE Int'l Symp. Object-Oriented Real-Time Distributed Computing (ISORC 03)*, IEEE, 2003, pp. 179–188.

6. J.D. Touch et al., "DynaBone: Dynamic Defense Using Multi-layer Internet Overlays," *Proc. 3rd DARPA Information Survivability Conf. and Exposition (DISCEX 03)*, IEEE, 2003, pp. 271–276.
 7. D.D. Woods and E. Hollnagel, "Mapping Cognitive Demands in Complex Problem-Solving Worlds," *Int'l J. Man-Machine Studies*, vol. 26, 1987, pp. 257–275.
 8. D.T. Moore, *Sensemaking: A Structure for an Intelligence Revolution*, Nat'l. Defense Intelligence College, 2011.
 9. D.T. Moore and R.R. Hoffman, "Sense-making: A Transformative Paradigm," *Am. Intelligence J.*, vol. 29, 2011, pp. 26–36.
 10. J.M. Bradshaw et al., "Sol: An Agent-Based Framework for Cyber Situation Awareness," to be published in *Künstliche Intelligenz*, 2012.
 11. C.G. Langton, ed., *Artificial Life: Proceedings of an Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, Addison-Wesley, 1989.
 12. M. Johnson et al., "Beyond Cooperative Robotics: The Central Role of Interdependence in Coactive Design," *IEEE Intelligent Systems*, vol. 26, no. 3, 2011, pp. 81–88.
 13. A. Uszok et al., "Toward a Flexible Ontology-Based Policy Approach for Network Operations Using the KAoS Framework," *Proc. 2011 Military Comm. Conf. (MILCOM 11)*, IEEE, 2011, pp. 1108–1114.
 14. J. van Diggelen et al., "Implementing Collective Obligations in Human-Agent Teams Using KAoS Policies," *Coordination, Organizations, Institutions and Norms in Agent Systems V*, LNCS 6069, Springer, 2010, pp. 36–52.
- Marco Carvalho** is an associate professor at the Florida Institute of Technology and a research scientist at the Institute for Human and Machine Cognition. Contact him at mcarvalho@fit.edu.
- Jeffrey M. Bradshaw** is a senior research scientist at the Institute for Human and

Machine Cognition. Contact him at jbradshaw@ihmc.us.

Larry Bunch is a research associate at the Institute for Human and Machine Cognition. Contact him at lbunch@ihmc.us.


Tom Eskridge is a research scientist at the Institute for Human and Machine Cognition. Contact him at teskridge@ihmc.us.

Paul J. Feltovich is a research scientist at the Institute for Human and Machine Cognition. Contact him at pfeltovich@ihmc.us.

Robert R. Hoffman is a senior research scientist at the Institute for Human and Machine Cognition. Contact him at rhoffman@ihmc.us.

Daniel Kidwell is a senior computer systems researcher with the US Department of Defense. Contact him at dlkidw2@tycho.ncsc.mil.

cn Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.



handles the details *so you don't have to!*

- Professional management and production of your publication
- Inclusion into the IEEE Xplore and CSDL Digital Libraries
- Access to CPS Online: Our Online Collaborative Publishing System
- Choose the product media type that works for your conference:
Books, CDs/DVDs, USB Flash Drives, SD Cards, and Web-only delivery!

Contact CPS for a Quote Today!

www.computer.org/cps or cps@computer.org

