

Managing the Hidden Costs of Coordination

LAURA M.D. MAGUIRE

**CONTROLLING
COORDINATION
COSTS WHEN
MULTIPLE,
DISTRIBUTED
PERSPECTIVES
ARE ESSENTIAL**

IT STARTED WITH 502 ERRORS. *Almost immediately a flood of user reports swamped the service's community Slack channel.*

A user posted "Getting 502s?" at 9:22 a.m., and within minutes 40 other users responded with the Yes and MeToo emojis.

Also at 9:22 a.m., in an ops channel, an incident had been opened by an on-call engineer, and the site reliability engineers responsible for the service had been paged out. By 9:23 a.m. five responders were checking logs and dashboards.

At 9:25 a.m.—less than two minutes after an initial tentative question indicated there may be an issue—the first notification was pushed out to users. This was aimed at slowing the influx of user reports from the 77,000-plus user community.

In less than seven minutes, eight hypotheses about

the nature of the problems had been proposed by the responders. In that same period, five of those had been investigated and discarded.

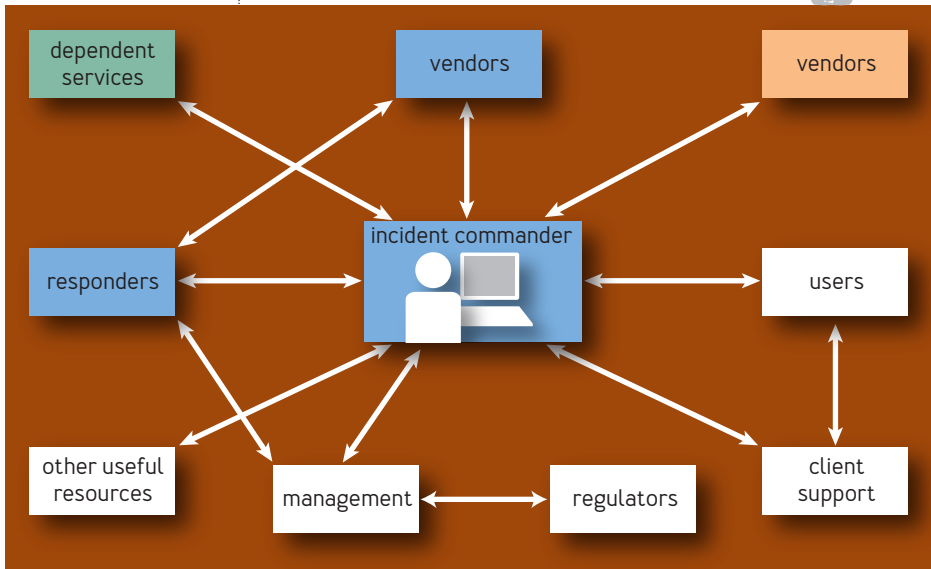
Within the first 10 minutes of the incident, the responders had been directly in touch with the 4,700 users in their community channel, opened tickets with three dependent services' support teams, and coordinated among a response squad of 10.

Diverse players are engaged when IT systems run at speed and scale. This becomes immediately apparent when the service is disrupted. Those whose work depends on the system functioning, both directly and indirectly, are compelled to get involved either to help with resolution or to seek more information so they can adjust their goals and priorities to account for the degraded (or absent) service.

Often, because of the business-critical nature of the service or four nines service-level agreements, a service outage triggers an all-hands-on-deck page for multiple responders. This core group represents a small fraction of the roles involved, however. Even with a brief look at an incident response, it becomes apparent that performance in resolving service outages in these systems is about rapid, smooth coordination of these multiple, diverse players, as expressed in figure 1.

Joint activity distributed among this collective takes place across scripted and unscripted efforts such as recognizing the disruption, taking actions that safeguard the system from further decline, diagnosing the source(s) of the problems, determining potential solutions, cross-

1

FIGURE 1: **RESOLVING SERVICE OUTAGES**

checking a fix before the code gets pushed, as well as a whole suite of after-action activities.

Even in relatively smaller scale systems, the incident response can become less about diagnosis and repair of service outages and more about managing the needed capabilities of multiple responders, the potential benefits that could be realized by having more participants available to assist, and the needs of the stakeholder groups. This coordination incurs additional demands. For example, for their skills and experience to be useful to the current flow of events, incoming responders need to be briefed and understand the tasks they've been delegated relative to the sequencing of activity. Doing this requires a

substantial amount of effort—particularly as the severity of the outage or number of responders increases or the uncertainty grows.

In the high-consequence world of managing service delivery for critical digital infrastructure, the time pressure to diagnose and repair an outage is enormous.¹ While resources may be readily available, it can be extraordinarily challenging to use them as the tempo of the incident escalates and the efforts to stop a cascade of failures occupy all the attention of the response team.

Herein lies the crux of the issue: The collaborative interplay and synchronization of roles is critical,^{12,13,15} but prior research has shown poor coordination design incurs cognitive costs for practitioners, specifically, the additional mental effort and load required to participate in joint activities.^{5,6} This can be particularly exacerbated in the digital services domain where it plays out across geographically distributed groups. Using examples from critical digital services, this article explores the nature of coordination costs and how software engineers experience them during a service outage. These findings provide new directions for design to control costs of coordination in incident response.

HIDDEN COSTS OF COORDINATION

The choreography needed for smooth operation is effortful,⁷ particularly when the system is under stress. But these efforts are difficult to discern and typically not separated from expected “professional practice” within a field. This choreography arises as “an escalating anomaly can outstrip the resources of a single responder quickly.

There is much to do and significant pressure to act quickly and decisively. *“To marshal resources and deploy them effectively requires a collection of skills that are related to but different from those associated with direct problem solving. But to be effective, these resources must be directed, tracked, and redirected. These activities are themselves demanding.”*¹⁸

That this collection of skills goes largely unnoticed is not surprising. The fluency with which expert practitioners manage these coordination demands minimizes the visibility of the efforts involved.¹⁹ It is only when the coordination breaks down that it comes to the forefront. Difficulties in synchronizing activities, disruptions to the smooth flow of task sequences, or conversation explicitly aimed at trying to organize multiple parties are examples of evidence that coordination breakdowns have occurred.

It is worth separating out the choreography needed for coordination from the costs that those activities incur. An example of this occurs when recruiting new resources to an incident response – just one function in joint activity. The associated overhead costs include:

- Monitoring current capacity relative to changing demands
- Identifying the skills required
- Identifying who is available
- Determining how to contact them
- Contacting them
- Waiting for a response
- Adapting current work to accommodate new engagement (waiting, slowly completing tasks to aid coordination)

- Preparing for engagement
 - Anticipating needs
 - Developing a ‘critsit’ or status update
 - Giving access/permissions to tools and coordination channels
 - Generating shared artifacts (dashboards, screenshots)
- Dealing with access issues (inability to join web conferences or trouble establishing audio)

These overheads seem relatively benign—they are implicit features of any joint activity. And that is precisely the point: They can be a minimal burden in normal operations and therefore disregarded as worthy of support in explicit design. In high-tempo, time-pressured, and cognitively demanding scenarios, however, these burdens increase to the point of overloading already burdened responders. Think of a loss of engine power during the first few minutes of flight or an unexpected event during a spacewalk—seconds count here and any additional friction in cognitive work matters. Now think of the speed at which critical digital services operate—*microseconds* count and the hidden coordination costs can matter in previously unconsidered ways. The cognitive costs of coordination matter in incident-response processes. Now let’s consider how poor coordination design impacts engineering teams responsible for system reliability.

THE NEED FOR COORDINATION DESIGN

Highly technical system operation is increasingly non-collocated. Demands for near-perfect reliability and the burnout this can generate for on-call engineers has given

rise to different models of 24/7 systems management to distribute calls across time zones. Even when a team may be geographically collocated, outages happen in off-hours or when members of the team may be traveling, in meetings, or otherwise unavailable for face-to-face interactions. This means incident response should be designed to accommodate entirely remote joint activity.

The need for good coordination design transcends the software community: Increasingly, other industries that were not typically geographically dispersed in the past are taking advantage of technological capabilities to distribute their workforces to optimize cost or available talent (providing just-in-time expertise).

Current coordination design focuses on the structure of handling support, including triage methods whereby runbooks or troubleshooting algorithms are used by less experienced support engineers before escalating to experts or through geographically dispersed support networks that “follow the sun.” These formats can decrease the need to wake up expert resources when the system goes down, but these configurations do not eliminate the need for coordination design. The requirements are shifted in ways that can escalate situations, compounding the coordination demands of the event as other stakeholders get engaged.

Let’s follow this through with an example. When anomalies generate the need to page the on-call staff, these direct responders begin gathering. Simultaneously, other stakeholders with an interest in the problem are also drawn in. Users may begin flooding support channels and ticketing systems trying to determine if the service is

degraded or if their system is wonky, or dependent services may experience problems and begin asking for information. This coordination “noise” makes it challenging to determine if these are all the same problems, related, or unrelated.

With diagnostic and safeguarding activities commanding substantial attention, additional resources are then needed to triage this influx of reports and sort through the incoming data to minimize data overload.¹⁶ As the incident progresses and the concern over impact grows, escalations to management bring in even more participants as senior executives begin pressing for more details or demanding the service be restored. Customer support roles facing urgent requests from clients will seek information to pass along.

Despite the substantial number of parties involved, systems are rarely designed with explicit attention to the coordination requirements. When they are, typically it is to: (1) centralize response coordination through an incident commander; (2) design an overly prescriptive process management perspective that fails to account for the hidden cognitive work of coordination; or (3) depend on tooling that fails to fully support the dynamic, nonlinear manner in which incident response happens. These methods do not necessarily support the cognitive work of coordination the way they are intended.

ATTEMPTS AT SUPPORTING COORDINATION

Some would argue that coordination design is fundamental for developing and deploying technology in distributed systems such as CDI (critical digital infrastructure). But process-driven coordination design—emphasizing

distributed tasks instead of joint activity—will not address the needs described earlier. One example of process-driven industry best practice surrounding coordination during service outages—borrowed from disaster and emergency response domains—is an ICS (incident command system). Central to this model is assigning an IC (incident commander) and ensuring disciplined adherence to the shared ICS across the roles and groups involved. Let's look at how these two tenets can actually limit resilient incident-response practices.

Attempt 1: Assigning an Incident Commander

The intent of the IC role is to manage the coordination requirements of the involved parties by directing the activities of others and holding the responsibility for taking timely decisions. Under certain conditions (in low-tempo scenarios with few involved parties or reasonably known and predictable event outcomes, for example) this may be an appropriate configuration. In these contexts (or these phases of an incident), the cognitive and coordinative demands are manageable without design for coordination.^{7,12,13} Routine events can be handled without undue stress.

Escalations that move a situation to nonroutine or exceptional, however, dramatically increase the cognitive activities needed to cope and generally do not follow a predictable course. As demands grow, an incident-command structure tends to become a workload and activity bottleneck that slows response relative to the tempo of cascading problems.²⁰ Working both *in* and *on* the incident forces attention to be divided across the

“inherent” roles of the position. For example, the IC needs to be tracking the details of the incident to be prepared to anticipate and adapt to rapidly changing conditions, but too much effort spent on forming an accurate assessment of the situation takes away from managing the coordination across roles. In reverse, trying to centrally manage who does what when tends to fall behind the pace of events and challenges, making the trouble harder to resolve and the joint activity harder to synchronize.

This is not an inconsequential point. Being an effective choreographer of the joint activity demands current, accurate knowledge *and* the ability to redirect attention to the orchestration of the players coming in and out of the event alongside their changing needs. In addition, what is seen as the IC maintaining organizational discipline during a response can actually be undermining the sources of resilient practice that help incident responders cope with poorly matched coordination strategies and the cognitive demands of the incident.

Attempt 2: Enforcing operational discipline to follow the ICS

Previous studies in software have shown different strategies for coping with workload demands such as dropping tasks (known as *shedding load*), deferring work to later, or reducing the quality of the work performed.² Other attempts to balance the workload sink with the value of the coordination call for adding more resources, but this comes with costs as well. In poorly designed systems, resources needed to help handle the demands are unable to be brought into play smoothly without disrupting the

work under way to control the adverse effects of the event.

Herein lies a paradox: You have resources available but are unable to make them useful. Concurrently, their attempts to become useful are counterproductive—new responders coming into an audio bridge or ChatOps channel need to ask for a briefing, and the updating disrupts the flow of activity. This can drive the formation of side channels among select responders where diagnostic work can take place uninterrupted. Creating this peripheral space is necessary to accomplish cognitively demanding work but leaves the other participants disconnected from the progress going on in the side channel.

Unless you have been “on the fireline” of an event of this sort, it can be easy to minimize the tension inherent in these situations. It’s worth restating: the systems studied in coordination research are often life-critical or otherwise high-consequence. Despite the importance of coordination, timely actions must be taken to cope with anomalies as they threaten to produce failures. When high costs of coordination could undermine the ability to keep pace with the evolving demands of the anomalous situation, people responsible for the outcomes will, of necessity, adapt. Incident response in critical digital infrastructure systems is not exempt. In fact, the speed and scale at which CDI operates, coupled with the challenges of a distributed team connected through technology, make the domain particularly susceptible to interference from excessive costs of coordination.

In observations of critical events and post-mortems, adaptations to create subgroups in different channels that

are separate from the “official” incident response occur repeatedly.⁹ Often, postmortems misinterpret these forms of adaptation to high costs of coordination. Retrospective discussions portray these adaptations as contrary to the ICS protocols and therefore lead to efforts to block people from forming these channels. The behavior is actually an adaptive strategy to cope as coordination becomes too expensive. Rather than forcing responders to bear significant attentional and workload costs, it is advisable to facilitate shifting various lines of work to subgroups while supporting connecting the progress or difficulties into the larger flow of the response.

The emergency services community has begun to recognize the limitations of the ICS,⁴ as have other domains where command and control or hierarchical methods are giving way to more flexible teaming structures.^{10,11} When practices such as ICS are adopted across domains, it is important to pay close attention to the critiques and findings from other large-scale, multi-agent coordination contexts. In doing so, it is possible to limit the unintended adverse impacts when real-world demands of one setting challenge the practices imported from another.

These findings about how people in an incident response adapt when high costs of coordination threaten the critical cognitive work are an important source of design seeds to guide innovations.

Attempt 3: Using technology to facilitate coordination

The term *computer-supported cooperative work* (CSCW) was coined by Irene Greif and Paul Cashman in the early 1980s to describe the emerging field of computers

mediating the coordination of activity across people and roles.³ Since then, advances in technological capabilities, the omnipresence of the computer in the workplace, and the proliferation of automated processes have solidified the importance of CSCW, while rendering it redundant since almost all forms of joint activity have become computer-mediated.

Still, this field has three main themes that are of particular interest in CDI: the use of collaboration software platforms; the coordination of joint activity between humans and bots; and the nature of reciprocity in human-automation teaming.

Collaboration software platforms

Not surprisingly, because of the changing needs of the work environment and the technical capabilities of the workforce, software engineering has driven innovation and the development of tooling and practices for collaborative work. Online software platforms take traditional offline activities such as project management planning, issue tracking, group discussion, and negotiation of shared work and enable real-time collaboration of participants across a distributed network.

The platforms have shifted from expensive, proprietary forms of file sharing to broadly accessible, cloud-based tools that can be quickly adopted across both formal and ad hoc groupings. Lowering the barrier to collaboration in this way eases the coordination costs of transient, single-issue demands and of early exploratory efforts. This means collaborative work can be facilitated more rapidly with less overhead. Flexible coordination structures also provide

the ability to adapt their use to the problem demands.

The resilience demonstrated in the earlier example of forming side channels to manage high costs of coordination was facilitated by the ease with which direct messages could be sent or new channels could be spun up. Supporting rapid reconfiguration into smaller, ad hoc teams enables smooth transitions as activity is distributed across continuously changing groups of participants. This collection of attributes—adapting to changing problem demands, dynamic reconfiguration of resources, and smooth coordination—is critically important in high-consequence work and a prominent feature of groups that are skilled at distributed joint activity in many domains.

Designing technology that can aid these capabilities is a means to control the costs of coordination. While many of these platforms optimize coordination costs on one criterion (rapid reconfiguration), ChatOps platforms exact penalties in coordinating *with* the tools themselves. For example, while the practice of ChatOps allows traceability that *could* support bringing new responders up to speed, the packed message-list format of the tooling is poorly designed to do so.¹⁴ Responders coming into an event that is under way must scroll through the list of text, searching for the relevant lines of inquiry still in consideration, key decision points, and other important contextual information to gain a current understanding of the situation.

These seemingly trivial aspects of design matter greatly. Think back to the tension inherent in high-tempo operations when seconds matter and expert resources are in high

demand. Those who are likely to be drawn in to join in the response efforts on a service outage frequently possess specialized skills that are often scarce. As such, they may not be brought into the event until later stages, at which time the tempo or propagation of failure drives a need for taking urgent action. Poor design renders ChatOps nearly useless as a tool for sensemaking as people come into an evolving and increasingly pressured situation.

*Coordinating joint activity
across humans and machines*

The last subsection shifted the framing of controlling the costs of coordination. Initially, cost of coordination referred to the additional efforts to accommodate the tasks and interactions inherent in joint activity. In human-human coordination the costs of the interaction are borne by both parties, and “investments” may be made by relaxing individual or short-term goals in the service of accommodating shared or longer-term goals. Working jointly distributes the costs across the participants. The preceding subsection introduced an important distinction: Interacting with tools and automation skews the costs. There are many coordination costs in human-machine teaming that go unnoticed or are exacerbated by tool design.

For example, the initial expenditure of effort to set up tooling designed to aid in various functions of anomaly response such as monitoring or alerting can be substantial. Engineers responsible for assembling their own stacks spend considerable effort in: assessing the appropriateness of a tool for a given purpose; evaluating

it relative to their team's needs; considering the technical capabilities needed to understand how it functions; learning how it works; maintaining an accurate mental model as new features are added; determining appropriate configurations; performing maintenance to ensure that old configurations are removed or updated as demands change; tolerating the lack of context sensitivity that can result in unnecessary alerting; providing access and permissions to the users on the team; constructing security measures to prevent inadvertent changes; and making changes and adjustments as new tools are integrated. (The list could continue.) These are all examples of how coordinating with machines have costs for their human counterparts. If the tool were a human colleague, the amount of effort you would need to expend to ensure it remained a relevant team member might give you pause; however, this fundamental asymmetry that unduly burdens the human team members with additional costs to compensate for the limitations of automation is characteristic of current-day human-machine teams.^{6,7}

A key (and often overlooked) aspect of the dynamics of teamwork across human-human and human-machine networks is the degree to which the participants in the joint activity *consider* the goals, workload, and needs of others and adapt their actions accordingly.

Recognizing the dynamics of reciprocity

Choreographing technologically mediated joint activity can enable greater opportunities for reciprocity when the technology is designed to combat excessive costs of coordination.¹⁷ For example, studies of NASA's space-

shuttle mission control during critical events reveals many patterns of effective joint activity. Of particular interest, many people join in beyond those who are titularly responsible. The technology that mediates communication in the control room and backrooms facilitates bringing people up to speed as they join in from being off duty, with low burdens on the people currently handling the anomalies.¹³ The additional personnel provide diverse perspectives, especially as each flight controller increasingly focuses on his or her scope of responsibility as the anomalous situation unfolds. The ability to “look in and listen in” has been widely documented as a benefit to smooth coordination.^{8,12}

It’s not difficult to see the parallel between mission control and CDI in the rapid escalation in the number of stakeholders (other responders, users, customer support, management) during a service outage. Technologies that enable this and other abilities for joint activity in a fully distributed network without adding extra burdens provide a means for people whose skill, experience, and knowledge *could* be useful to the event but who have not been explicitly drawn in can ready themselves to assist should the need arise. Being current on the event progression, yet untethered to specific responsibilities, offers an opportunity for reframing through fresh perspectives (Grayson, this issue).

In outlining these three attempts at supporting coordination, it’s clear that technology both affords lower-cost coordination by supporting adaptive capacity and exacerbates high-cost coordination through asymmetrical burdens on the human side. In CDI environments, where

technology can be rapidly developed and deployed, designs can easily add unintended costs for joint activity unless the tools are explicitly designed to support coordination.

CONCLUSION

Coordination remains an integral part of large-scale, distributed work systems, but the lack of coordination design for joint activity continues to add hidden cognitive costs for practitioners. These efforts and load are related to the additional work of enabling smooth synchronization across multiparty groupings as the cognitive work of anomaly response is completed in high-tempo, evolving incident scenarios. Recall the opening case, in which the escalating incident brought in multiple, diverse, and distributed perspectives, each with a vested interest in the event progression.

Each participant was necessary to managing the outage both directly and indirectly, and the ChatOps forum enabled their participation. Closer examination across a number of cases, however, reveals a paradox: The platforms themselves both facilitate and hinder coordination. The easy formation of side channels enables engineers to adapt through flexible reconfiguration outside of the main response efforts, but bringing new responders up to speed is made difficult by the structure of a packed message-list design.

Some of the common tactics thought to control the costs of coordination include adopting incident command structures, specifically the IC role. Using collaborative software platforms and adopting technologies to aid in coordination have been shown in actual cases to

Related articles

- ➔ The Calculus of Service Availability
You're only as available as the sum of your dependencies.
Ben Treynor, Mike Dahlin, Vivek Rau, Betsy Beyer
<https://queue.acm.org/detail.cfm?id=3096459>
- ➔ Collaboration in System Administration
For sysadmins, solving problems usually involves collaborating with others. How can we make it more effective?
Eben M. Haber, Eser Kandogan, Paul Maglio
<https://queue.acm.org/detail.cfm?id=1898149>
- ➔ Distributed Development
Lessons Learned
Michael Turnlund
Why repeat the mistakes of the past if you don't have to?
<https://queue.acm.org/detail.cfm?id=966801>

reveal limits and unrecognized implications for cognitive work. Nevertheless, all of these areas provide opportunities to choreograph smoothly in high-tempo, multi-agent events, especially by supporting the ability to adapt when the costs of coordination climb too high.

Some initial considerations to control cognitive costs for incident responders include: (1) assessing coordination strategies relative to the cognitive demands of the incident; (2) recognizing when adaptations represent a tension between multiple competing demands (coordination and cognitive work) and seeking to understand them better rather

than unilaterally eliminating them;

(3) widening the lens to study the joint cognition system (integration of human-machine capabilities) as the unit of analysis; and (4) viewing joint activity as an opportunity for enabling reciprocity across inter- and intra-organizational boundaries.

Controlling the costs of coordination will continue to be an important issue as systems scale, speeds increase, and the complexity rises in the problems faced during anomalies that disrupt reliable service delivery.

References

1. Allspaw, J. 2015. Trade-offs under pressure: heuristics and observations of teams resolving Internet service outages. M.S. thesis, Lund University; <https://lup.lub.lu.se/student-papers/search/publication/8084520>.
2. Grayson, M. R. 2018. Approaching overload: diagnosis and response to anomalies in complex and automated production software systems. M.S. thesis, The Ohio State University; https://etd.ohiolink.edu/pg_10?::NO:10:P10_ETD_SUBID:174511.
3. Grudin, J. 1994. Computer-supported cooperative work: history and focus. *Computer* 27(5), 19-26; <https://www.microsoft.com/en-us/research/wp-content/uploads/2017/01/IEEEComputer1994.pdf>.
4. Jensen, J., Waugh Jr., W. L. 2014. The United States' experience with the Incident Command System: what we think we know and what we need to know more about. *Journal of Contingencies and Crisis Management* 22(1), 5-17; 22(1), 5-17; https://onlinelibrary.wiley.com/doi/pdf/10.1111/1468-5973.12034?casa_token=s3009HeH72kAAAA:GF6Yz4LwClKY3tEhbF5s32VW2by189SiefXqJc5xST2dMq1xwZ_fZbqhCaemkEDNvSlgDilky1p6fc.
5. Klein, G. 2006. The strengths and limitations of teams for detecting problems. *Cognition, Technology & Work* 8(4), 227-236; <https://link.springer.com/content/pdf/10.1007%2Fs10111-005-0024-6.pdf>.
6. Klein, G., Feltovich, P. J., Bradshaw, J. M., Woods, D. D. 2004. Common ground and coordination in joint activity. In *Organizational Simulation*, ed. W. Rouse and K. Boff, 139-184. New York: Wiley; http://jeffreymbradshaw.net/publications/Common_Ground_Single.pdf.

7. Klein, G., Woods, D.D., Bradshaw, J., Hoffman, R.R., Feltovich, P.J. 2004. Ten challenges for making automation a “team player” in joint human-agent activity. *IEEE Intelligent Systems* 19(6), 91-95; <https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=1363742>.
8. Luff, P., Heath, C., Greatbatch, D. 1992. Tasks-in-interaction: paper and screen-based documentation in collaborative activity. In *Proceedings of the ACM Conference on Computer-supported Cooperative Work*, 163-170; <https://dl.acm.org/citation.cfm?id=143475>.
9. Maguire, L.M. 2020, forthcoming. Controlling the cognitive costs of coordination in large-scale, distributed systems: in search of a model of choreography to support joint activity. Dissertation, The Ohio State University; https://www.researchgate.net/profile/Laura_Maguire5.
10. Nemeth, C. P. 2007. Groups at work: lessons from research into large-scale coordination. *Cognition, Technology & Work* 9(1), 1–4; <https://link.springer.com/article/10.1007/s10111-006-0049-5>.
11. O’Leary, R., Bingham, L. B. (eds.). 2009. *The Collaborative Public Manager: New Ideas for the Twenty-first Century*. Georgetown University Press; <https://muse.jhu.edu/book/13036>.
12. Patterson, E. S., Watts-Perotti, J., Woods, D. D. 1999. Voice loops as coordination aids in space shuttle mission control. *Computer Supported Cooperative Work (CSCW)* 8(4), 353-371; <https://www.semanticscholar.org/paper/Voice-Loops-as-Coordination-Aids-in-Space-Shuttle-Patterson-Watts-Perotti/9cbe3fe4e0067be8b1>

- acc7816152cef172b410ae.
13. Patterson, E. S., Woods, D. D. 2001. Shift changes, updates, and the on-call architecture in space shuttle mission control. *Computer Supported Cooperative Work (CSCW)* 10[3-4], 317-346; <https://link.springer.com/article/10.1023/A:1012705926828>.
 14. Potter, S. S., Woods, D.D. 1991. Event-driven timeline displays: beyond message lists in human-intelligent system interaction. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics*; <https://ieeexplore.ieee.org/abstract/document/169864>.
 15. Watts-Perotti, J., Woods, D. D. 2009. Cooperative advocacy: an approach for integrating diverse perspectives in anomaly response. *Computer Supported Cooperative Work (CSCW)* 18[2-3], 175-198; <https://link.springer.com/article/10.1007/s10606-008-9085-4>.
 16. Woods, D.D. 1994. Cognitive demands and activities in dynamic fault management: abduction and disturbance management. In *Human Factors of Alarm Design*, ed. N. Stanton. London: Taylor & Francis, 63-92.
 17. Woods, D.D. 2019. Essentials of resilience, revisited. In *Handbook on Resilience of Socio-Technical Systems*, eds. M. Ruth and S. G. Reisemann. Edward Elgar Publishing, 52-65; https://www.researchgate.net/profile/David_Woods11/publication/330116587_4_Essentials_of_resilience_revisited/links/5c2e448ba6fdccd6b58f871e/4-Essentials-of-resilience-revisited.pdf.
 18. Woods, D. D., ed. 2017. STELLA Report from the SNAFUcatchers Workshop on Coping with Complexity.

- SNAFU Catchers Consortium; <http://stella.report/>.
19. Woods, D.D., Hollnagel, E. 2006. *Joint Cognitive Systems: Patterns in Cognitive Systems Engineering*. Boca Raton, FL: CRC Press (Taylor & Francis).
 20. Woods, D. D., Patterson, E. S. 2001. How unexpected events produce an escalation of cognitive and coordinative demands. In *Stress, Workload, and Fatigue*, eds P.A. Hancock and P.A. Desmond. Mahwah, NJ: L. Erlbaum; http://csel.eng.ohio-state.edu/productions/laws/laws_mediapaper/2_4_escalation.pdf.

Laura M.D. Maguire studies human performance in high risk/ high consequence work as a researcher at the Cognitive Systems Engineering Lab at The Ohio State University. Her research interests lie in resilience engineering, coordination, and enabling adaptive capacity across distributed work teams and forms of systems regulation and control. She has been a researcher with the SNAFU Catchers Consortium since 2017 and works closely with large- and medium-sized digital service companies on incident response practices, tool development, design, and contextual research. Laura has a master's degree in Human Factors & Systems Safety and is currently completing her Ph.D. in cognitive systems engineering at OSU. She draws from experience working in forestry, oil & gas, investment banking, and industry associations in her research. As a backcountry skier and alpine climber, she also is interested in and has written on cognitive work & resilient performance in mountain environments.

Copyright © 2019 held by owner/author. Publication rights licensed to ACM.