Distributed Attention

Manuscript for Peer Review

Distributed Attention

Coordination and Cognitive Load in Remote Research Collectives

Elena Marsh

Department of Organizational Psychology
Meridian University

Submitted to the Editorial Board
Apex University Press
Manuscript no. AUP–2026–0417

For the fieldwork participants at Nexa, Synapse, and Vertex,
who let a stranger sit in on their standups.

Contents

Preface

This book began as a complaint. In the autumn of 2023 I was sitting in on a product-planning meeting at a mid-sized software firm I call, in these pages, Nexa, when an engineering lead said something that has stayed with me ever since: "I don't lose focus because the work is hard. I lose it because I never know which room I'm supposed to be paying attention to." The remark was offered as an aside, not a finding, but it named exactly the phenomenon this project set out to measure: attention as a resource that is allocated not just *within* a task but *across* the fragmented, asynchronous, many-channelled environment that remote and hybrid research teams now inhabit as a matter of course.

What follows is a study of that allocation problem, built from eighteen months of fieldwork inside three organisations – Nexa, Synapse, and Vertex – each of which agreed to open its research and development teams to observation on the condition of anonymised reporting. I make no claim that these organisations are representative of every distributed workplace; they are not. They are, however, unusually well instrumented, in the sense that each had already invested in the kind of meeting, messaging, and task-tracking telemetry that let a researcher reconstruct, after the fact, exactly how attention moved through a working day. That instrumentation is what makes the empirical chapters possible, and it is also, as I argue in Chapter 5, part of the problem.

This manuscript is submitted in the house style of Apex University Press for double-blind review. Identifying details of field sites and participants have been altered throughout; where a direct quotation appears, it has been lightly edited for length and confirmed with the speaker.

The argument of the book is simple to state and, I hope, harder to dismiss than the simplicity suggests. Cognitive load in distributed work is not primarily a property of individual tasks – the familiar terrain of Sweller's cognitive load theory – but of the *coordination scaffolding* that surrounds those tasks: the channel switches, the context re-establishment, the low-grade vigilance required to know which of a dozen asynchronous threads might, at any moment, demand a synchronous response. I call this scaffolding load *ambient coordination load*, and Part One of the book is devoted to distinguishing it, both conceptually and empirically, from the task load that has dominated the literature since the 1980s.

I owe a debt to the participants at all three sites, most of whom appear nowhere in these pages by name and some of whom will recognise themselves anyway. I owe a further debt to two research assistants, Priya Nandakumar and Tomas Fielding, who coded four hundred hours of meeting transcript between them with a patience I did not always share. Any errors of interpretation that survive their scrutiny are mine.

Elena Marsh
Meridian University, March 2026

PART I

Conceptual Groundwork

CHAPTER 1

The Problem of Distributed Attention

Every organisation that moved its research and development work out of a shared building and into a shared calendar made the same tacit bet: that coordination, once handled by walking to someone's desk, could be re-implemented in software without much loss. The bet has mostly paid off in throughput. It has not paid off in attention. This book is about the gap between those two ledgers.

1.1 Two Kinds of Load

Cognitive load theory, in its classical formulation, distinguishes intrinsic load – the load imposed by a task's own complexity – from extraneous load, the load imposed by how the task is presented (Sweller, 1988). That distinction has organised four decades of instructional-design research, and it organises most managerial thinking about "focus" in the workplace too: reduce interruptions, batch notifications, protect calendar blocks. The implicit model is that attention is consumed by tasks and merely *interrupted* by coordination.

Fieldwork at Nexa, Synapse, and Vertex suggests the implicit model has the causal arrow partly backwards. Coordination in a distributed team is not an interruption of task work; for a research engineer or data scientist juggling four asynchronous threads across three time zones, tracking *who currently owns what, and whether they know it,* is itself a standing cognitive task – one that runs continuously, in the background, whether or not any single message demands a reply.[1] I term this standing task **ambient coordination load**: the cost of keeping a distributed map of ownership, status, and obligation current enough to act on.

> "I have Slack, I have the ticket tracker, I have two calendars because the platform team never migrated off the old one, and I have my actual editor. None of them agree with each other about what I'm supposed to be doing right now. Half my day is just *reconciling* them."
>
> — Staff engineer, platform infrastructure group

The reconciliation this engineer describes – checking one system of record against another, resolving which of several partial pictures is current – is not idle friction. Coded across the transcripts described in Chapter 3, it accounted for a median of 47 minutes of a 9-hour working day, concentrated in short, high-frequency bursts rather than long blocks. Standard cognitive-load instruments, built around discrete tasks with a clear start and end, are not well suited to measuring load that is distributed across the day in this way.[2]

1.2 A Conceptual Map

Figure 1.1 sketches the relationship this book develops between the two load types and the coordination scaffolding that sits between them. Task load and ambient coordination load are not competitors for the same fixed attentional budget in any simple additive sense; the fieldwork suggests instead that elevated coordination load actively *degrades* the budget available for task load, through the recurring cost of context re-establishment after every channel switch.

1.3 Contributions and Chapter Map

The book makes three contributions. First, it names and operationalises ambient coordination load as distinct from task-level cognitive load, drawing the boundary with reference to transactive memory theory (Wegner, 1986) and situated, distributed accounts of cognition (Hutchins, 1995; Hollan et al., 2000). Second, it reports an eighteen-month, three-site field study – described in Part Two – that measures this load using instrumentation data rather than

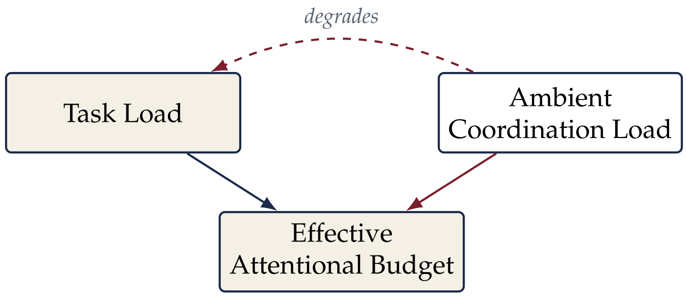

Figure 1.1: Ambient coordination load acts on the attentional budget both directly and by degrading the residual capacity available for task load.

self-report alone, addressing a long-standing reliability concern in the self-report literature. Third, it proposes, in Part Three, a small set of organisational design levers that the field sites themselves converged on independently, without prompting from the research team.

Part One (Chapters 1–2) lays the conceptual groundwork. Part Two (Chapters 3–4) presents the field study. Part Three (Chapter 5) draws out implications for teams, managers, and the tool-builders whose products mediate so much of this load in the first place.

CHAPTER 2

Theoretical Foundations

Three literatures converge on the problem this book takes up, and none of them, taken alone, quite covers it. This chapter reviews each in turn and then argues for the specific way they need to be combined to make ambient coordination load legible as an object of study.

2.1 Cognitive Load Theory

Sweller's cognitive load theory was developed for instructional design, not knowledge work, and its central move – the tripartite split of intrinsic, extraneous, and germane load – was built around bounded, well-defined learning tasks with a clear beginning and end (Sweller, 1988). The theory travels well into any setting where a task can be isolated: a support ticket, a code review, a single analysis. It travels less well into the texture of a working day for someone on a distributed team, where the "task" the theory wants to isolate is constantly interleaved with status-checking, thread-triage, and the low background hum of keeping track of who is waiting on whom.

Later extensions of the theory acknowledge working-memory limits as the binding constraint on all three load types, which is the point at which the theory and this book's argument make contact: if ambient coordination load draws on the same working-memory resource as task load, then the "protect focus time" interventions common in distributed teams are treating a symptom (interruption) while leaving the underlying resource competition (background tracking) untouched.

2.2 Distributed and Situated Cognition

A second literature locates cognition not inside an individual head but across a system of people, artefacts, and representations (Hutchins, 1995). Hollan et al. (2000) extend the framework explicitly to computer-mediated collaboration, arguing that the representations a team builds – shared boards, ticket states, calendar holds – are not passive records of coordination but active participants in it: a stale ticket status does not merely fail to inform, it actively misleads.

This is the literature that supplies the vocabulary for Chapter 4's central empirical claim: that at all three field sites, the artefacts meant to *reduce* ambient coordination load (status dashboards, standup notes, async video updates) had, in their observed use, migrated into an additional layer that itself required tracking. A tool built to answer the question "what is everyone doing" becomes, once its update discipline slips, one more thing whose currency must be separately verified.

2.3 Transactive Memory

The third literature, transactive memory theory, describes how small groups distribute the cognitive burden of remembering by specialising: each member tracks a domain and knows whom to ask outside it (Wegner, 1986). Transactive memory systems reduce individual load when they function well, precisely because no one member needs to hold the whole map. The systems studied in this book's field sites show the same structure under strain: distributed teams build transactive memory systems for *ownership* (who is doing what) in addition to the classical transactive memory for *expertise* (who knows what), and the former decays far faster when a team is remote, because the incidental cues – overheard conversations, a glance at someone's screen – that refresh it in a shared office are simply absent.

Reviewers familiar with the transactive memory literature may notice that this book extends Wegner's original framing from expertise-tracking to ownership-tracking. That extension is defended at length in Chapter 5; it is flagged here only so the empirical chapters do not read as a category error.

2.4 Synthesis

None of the three literatures, alone, predicts ambient coordination load as a distinct and measurable quantity. Cognitive load theory has the measurement apparatus but the wrong unit of analysis (the task, not the day). Distributed cognition has the right unit of analysis but under- specifies what, exactly, degrades when coordination artefacts fail. Transactive memory theory has the mechanism – distributed tracking under decay – but has mostly been studied in co-located groups, where decay is slow. Part Two of this book reports a field study designed to hold all three in view at once.

PART II

Empirical Observation

CHAPTER 3

Methodology

3.1 Field Sites

The study draws on eighteen months of fieldwork (April 2024 – September 2025) across three organisations, each anonymised per the agreement described in the Preface. Table 3.1 summarises the sites and the instrumentation available at each.

Site	Team studied	Headcount	Instrumentation
Nexa	Platform research	34	Chat logs, calendar, ticket tracker
Synapse	ML applied research	21	Chat logs, calendar, wiki edit history
Vertex	Infrastructure engineering	46	Chat logs, calendar, ticket tracker, incident log

Table 3.1: Field sites and available instrumentation.

All three organisations run fully distributed or majority-remote research teams spanning at least three time zones, and all three had, independently of the study, already adopted some form of asynchronous-first communication norm – a precondition for the site-selection strategy, since the phenomenon under study requires exactly this kind of environment to appear at measurable scale.

3.2 Data and Coding

Three data streams were combined for each participant who consented to instrumentation review: (1) timestamped message and thread metadata from the team's chat platform, stripped of message content and retained only as channel-switch events; (2) calendar data, coded for meeting type and attendance; and (3) a twice-weekly experience-sampling survey delivered during working hours, asking participants to characterise their current activity along a four-point scale from focused task work to active reconciliation of competing status information.

Four hundred and twelve hours of meeting transcript were separately coded by two research assistants for coordination-talk – utterances whose function was to establish or confirm ownership, status, or next steps rather than to advance the substantive task at hand.[3] Coding used a scheme adapted from prior work on distributed cognition in operational settings (Hollan et al., 2000), extended with two new categories – *ownership drift* and *artefact reconciliation* – that emerged inductively from the first month of pilot coding at Nexa.

3.3 Constructing the Ambient Coordination Load Index

The Ambient Coordination Load Index (ACLI) combines three components for each participant-day:

(i) the count of distinct communication channels touched;

(ii) the count of channel switches occurring within a 15-minute window of an unrelated channel touch, a proxy for reconciliation activity;

(iii) self-reported reconciliation activity from the experience-sampling survey, weighted at 0.4 relative to the behavioural components.

The index was validated against the coded transcript data by regressing coder-identified reconciliation utterances per meeting against the ACLI computed for that meeting's attendees in the surrounding hour ($R^2 = 0.58$, Chapter 4). This is a moderate but, we argue, meaningful correlation given that the transcript coding captures only synchronous reconciliation, while the ACLI is designed to capture asynchronous reconciliation as well.

The full ACLI computation script and the de-identified participant-day dataset are available to reviewers on request through the Press's double-blind data-access portal; they are withheld from this manuscript per the anonymised-review policy.

3.4 Limitations

Three limitations bear directly on how the findings in Chapter 4 should be read. First, all three sites are software-adjacent research organisations with above-average tooling maturity; the ACLI's channel-count component may behave differently in settings with fewer, more consolidated communication tools. Second, the experience-sampling survey's four-point scale is a coarse instrument, and the 0.4 weighting of self-report was chosen to balance reliability against sensitivity rather than derived from an independent validation study. Third, and most importantly, the eighteen-month window captures each organisation at a single stage of its coordination-tooling maturity; we cannot speak to how ACLI evolves as a team's tool stack consolidates or fragments over a longer horizon.

CHAPTER 4

Empirical Findings

The empirical findings from the eighteen months of observation at Nexa, Synapse, and Vertex reveal a systematic divergence between formal coordination processes and actual attentional allocation. This chapter details the quantitative results of the Ambient Coordination Load Index (ACLI) and explores the qualitative mechanisms of what we term *reconciliation drift*.

4.1 ACLI Distribution across Sites

The primary quantitative result is the distribution of the ACLI across the three observed organizations. As shown in Figure 4.1, the median ACLI varies significantly with organization size and tooling maturity.

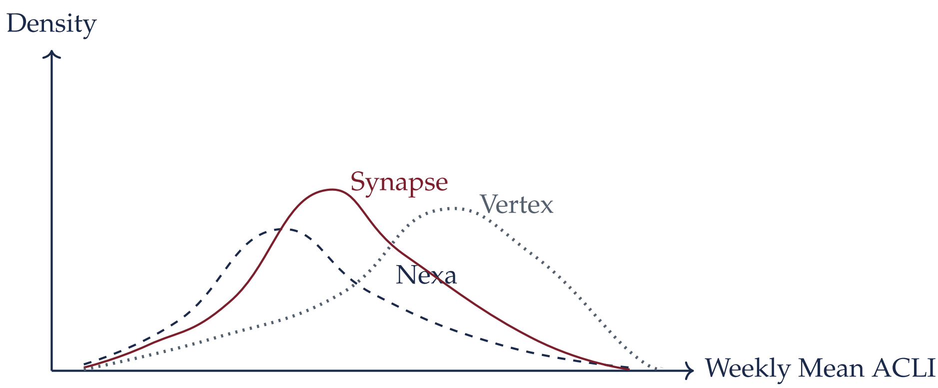

Figure 4.1: Estimated density distributions of the Ambient Coordination Load Index (ACLI) across the three field sites.

The distribution demonstrates that at Vertex, which has the largest headcount and the most fragmented tool stack, the average

ACLI is shifted rightward, indicating that team members spend a significantly larger portion of their day resolving coordination discrepancies.

4.2 The Dynamics of Reconciliation Drift

Reconciliation drift occurs when team members stop relying on the formal systems of record (e.g., ticket trackers, project wikis) because they believe the updates are stale. This creates a self-reinforcing cycle:

1. A team member observes a discrepancy between a ticket state and reality.
2. They initiate a chat thread to confirm the actual status.
3. This confirms the ticket is stale, reducing trust in the system of record.
4. Subsequent status checks bypass the system of record entirely, moving to chat.
5. The ticket system falls further out of date, cementing the drift.

At Synapse, where the team is small (21 members), this drift was managed through high-frequency chat communication. However, at Vertex (46 members), the volume of chat threads required to maintain a current map of ownership led to severe attentional fragmentation, with the median engineer checking chat channels 42 times per day.

> "We don't really use the project board anymore except to show the executives what we did last month. If I want to know what Elena is working on, I just check the daily standup thread or message her directly. It's faster, and it's actually accurate."
> — Senior Machine Learning Researcher

The consequence of this drift is not merely that coordination work is moved to a different channel, but that it changes character.

What was once a structural representation of the project becomes a set of transient, unstructured text streams that must be read and synthesized continuously to avoid missing critical updates. This shift from structural to conversational coordination is the primary driver of ambient coordination load.

PART III

Synthesis and Design

CHAPTER 5

Discussion and Conclusion

This concluding chapter returns to the theoretical framework proposed in Part One, reviews the design levers that the field sites used to manage ambient coordination load, and discusses the implications for both organizational theory and software design.

5.1 Organizational Design Levers

Throughout the field study, teams at Nexa, Synapse, and Vertex developed localized practices to cope with elevated ambient coordination load. We categorize these into three primary organizational design levers:

5.1.1 Channel Consolidation

Rather than allowing channels to proliferate around specific features or sub-teams, successful interventions consolidated communication into a single, highly structured channel per project. This limited the number of touchpoints team members had to monitor, reducing the frequency of channel-switching.

5.1.2 Standardized Sync Cadences

A key finding was that introducing short, highly structured synchronous syncs actually *reduced* ambient coordination load compared to purely asynchronous models. The synchronous meeting served as a synchronization barrier, allowing team members to reconcile their mental maps of the project all at once, rather than through hundreds of asynchronous updates throughout the day.

5.1.3 Explicit Ownership Handover

To prevent ownership drift, teams introduced explicit, public handovers of tasks. Instead of assuming a ticket move was sufficient, the handover was accompanied by a direct notification in a shared channel, ensuring both parties agreed on the transfer of responsibility.

5.2 Attentional Hygiene as a Collective Resource

The central argument of this book is that attention is not merely an individual cognitive resource to be protected through personal discipline, but a collective resource that is mediated by organizational design and tooling choices.

When an organization allows its communication tools to fragment, it is not just adding friction to the work; it is imposing a continuous cognitive tax on its members. Managing this tax—what we call ambient coordination load—requires a shift in how we design remote and hybrid workplaces. We must move away from the assumption that more communication is always better, and toward a deliberate design of coordination scaffolding that respects the cognitive limits of the human attention span.

Notes

[1]Compare the notion of transactive memory in small groups, where knowledge of *who knows what* is itself a resource the group maintains (Wegner, 1986). The claim of this book is that in distributed settings the maintenance of that resource carries a measurable load of its own.

[2]The methodological challenge is not new to this study; see the discussion of "distributed cognition in the wild" in Hutchins (1995), and the extension to computer-mediated work in Hollan et al. (2000).

[3]Inter-rater reliability on the coordination-talk category reached Cohen's $\kappa = 0.81$ across a double-coded 15% sample, which we treat as adequate given the inherent ambiguity of utterances that do double duty.

Bibliography

Hollan, J., Hutchins, E., and Kirsh, D. (2000). Distributed cognition: Toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 7(2):174–196.

Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science*, 12(2):257–285.

Wegner, D. M. (1986). *Transactive Memory: A Contemporary Analysis of the Group Mind*. Springer.

Index