

W H I T E P A P E R

# The Paradox of AI-Accelerated Software Engineering:

---

Why More AI Output Creates More Human Work  
*Introducing the AI-driven Software Engineer Resourcing Formula*

**Eric Larsen**

March 2026



**Aurix AI Solutions**

<https://aurixai.solutions>

*Enterprise AI Governance, Compliance & Marketing Intelligence*

*Based on empirical observations from AI-driven development of enterprise-scale software  
across parallel AI coding sessions with real-time architectural coordination.*

## Abstract

The rapid adoption of large language model (LLM) coding assistants has created a widespread assumption that AI will reduce the need for human software engineers. This paper presents empirical evidence from the deployment of AI-driven development practices within enterprise software engineering, demonstrating a counterintuitive finding: **at scale, AI coding assistants create more work for human engineers, not less**, even as they multiply output by an order of magnitude.

Through direct observation of parallel AI coding sessions building a production enterprise platform, we document the emergence of a coordination bottleneck that scales nonlinearly with the number of concurrent AI sessions. We introduce the AI-driven Engineer Resourcing Formula (AAERF), a mathematical model for predicting the human coordination overhead required to manage parallel AI engineering sessions and propose guidelines for organizations planning AI-augmented development teams.

**Keywords:** *AI-assisted software development, LLM coding, parallel development sessions, engineering management, resource planning, coordination overhead, software engineering productivity*

## 1. Introduction

The narrative surrounding artificial intelligence in software development has been dominated by a singular theme: replacement. Industry analysts, technology journalists, and venture capitalists have collectively projected a future in which AI systems progressively absorb the responsibilities of human programmers, eventually rendering the profession obsolete or dramatically reduced in headcount.

This paper challenges that narrative with empirical evidence. Drawing on direct observation from organizations that have adopted AI coding assistants as primary development resources—including cases where AI sessions serve as the equivalent of full engineering teams—we document a phenomenon we term the **AI Coordination Paradox**: as the number of parallel AI coding sessions increases, the human effort required to manage them grows superlinearly, even as raw code output scales linearly.

We find that at three or more concurrent AI sessions, the human engineer managing them is working at or above 100% capacity—not writing code, but performing architectural coordination, context synchronization, deployment sequencing, and quality validation. The work has not decreased; it has transformed from production labor into management labor.

This has profound implications for workforce planning, organizational design, and the economics of AI-assisted software development. Organizations that plan for AI to reduce headcount may find instead that they need to restructure roles, with fewer traditional developers and more AI orchestration specialists.

## 2. Background and Prior Work

## 2.1 Historical Software Engineering Productivity

Empirical studies of software engineering productivity have consistently found that individual developer output varies by an order of magnitude, a finding first documented by Sackman, Erikson, and Grant (1968) and confirmed by DeMarco and Lister (1985) in their seminal work on the subject. The Standish Group’s annual CHAOS reports indicate that the average developer produces between 325 and 750 lines of debugged code per month in enterprise contexts, with significant variation by domain, language, and organizational maturity.

Brooks’ observation in *The Mythical Man-Month* (1975) that adding manpower to a late software project makes it later remains one of the most frequently cited principles in software engineering. The underlying mechanism—communication overhead growing as  $O(n^2)$  with team size—is directly relevant to our findings, though in the AI context, the “team members” are AI sessions that cannot communicate with each other.

## 2.2 AI Coding Assistant Capabilities (2024–2026)

The current generation of LLM-based coding assistants, including systems from Anthropic, OpenAI, Google, and others, can generate production-quality code across multiple languages, frameworks, and architectural patterns. In controlled benchmarks, these systems demonstrate proficiency that would place them in the upper quartile of human developers for isolated task completion.

However, benchmarks measure *task completion* in isolation. Enterprise software development involves sustained context across sessions, awareness of concurrent work by other contributors, adherence to evolving architectural decisions, and knowledge of runtime environment specifics (database schemas, API contracts, deployment configurations). These capabilities remain limited in current AI systems, creating the coordination gap we document in this paper.

## 2.3 The Missing Literature

While substantial literature exists on AI pair-programming (single-session assistance) and on theoretical projections of AI impact on software engineering labor markets, we find a notable gap: there is virtually no published empirical research on the operational dynamics of multi-session parallel AI development. This paper aims to begin filling that gap.

# 3. Methodology and Observations

## 3.1 Observation Context

Our observations are drawn from consulting engagements where organizations have adopted AI coding assistants as primary development resources. In the most intensive configuration observed, a single technical lead manages three to four concurrent AI coding sessions, each operating on different subsystems of the same enterprise application. The application under development in our primary case study comprises:

- A Vue.js 3 frontend with 30+ views and 50+ components
- A PHP backend with 40+ API endpoints
- A PostgreSQL database with 200+ tables including version tracking
- Integration with external GPU computing services, AI model APIs, and cloud storage
- A workflow automation engine with visual designer
- Cost management, job scheduling, and model lifecycle subsystems

### 3.2 Observed Work Patterns

The following table summarizes the distribution of human effort as the number of concurrent AI sessions increases, based on aggregated observations:

Sessions	Human Capacity Used	Code Output Multiplier	Context Switches / Hour	File Conflict Risk	Primary Human Activity
1	40–50%	3–5×	2–4	None	Own work + review
2	75–85%	6–10×	8–15	Low	Coordination + review
3	100–120%	10–18×	20–30	Moderate	Full-time orchestration
4+	>120%	15–25×	30+	High	Triage + damage control

Table 1. Human effort distribution by concurrent AI session count

A critical observation is that at  $N = 3$  sessions, the human operator’s role shifts entirely from *producer* (writing code, making designs) to *orchestrator* (coordinating outputs, resolving conflicts, sequencing deployments, maintaining architectural coherence). The operator is no longer a developer who uses AI; they are a manager of AI developers.

### 3.3 The Coordination Taxonomy

We identify six categories of coordination work that emerge when managing parallel AI sessions:

1. **Context Synchronization:** Passing summaries of what each session has done to other sessions so they don’t produce conflicting code. This includes handoff documents, schema updates, and “don’t touch this file” instructions.
2. **Deploy Sequencing:** Determining which session’s output should be deployed first when multiple sessions modify the same system. Database migrations must precede API changes, which must precede frontend updates.
3. **Architectural Arbitration:** Making real-time design decisions that multiple sessions are blocking on. Each session proposes solutions; the human must evaluate, choose, and propagate the decision.
4. **Schema Validation:** Verifying that column names, table structures, API contracts, and data types match across sessions. AI sessions frequently reference schemas from memory that have since been modified by another session.

5. **Quality Triage:** Identifying and resolving issues in AI-generated code, including incorrect database column references, CSS style conflicts, missing route registrations, and mismatched function signatures.
6. **Runtime Debugging:** Diagnosing issues that only appear at deployment—500 errors from wrong file paths, Cloudflare blocks on API calls, PostgreSQL constraint mismatches, and Vite build failures from invalid Vue SFC structure.

## 4. The AI-driven Software Engineer Resourcing Formula

Based on our observations, we propose a mathematical model for estimating the human coordination effort required to manage parallel AI coding sessions. We term this the AI-driven Software Engineer Resourcing Formula (AAERF).

### 4.1 Core Formula

$$H(N) = \alpha N + \beta \cdot N(N-1) / 2 + \gamma \cdot \delta(N)$$

Where:

$H(N)$	Human coordination effort required, expressed as a fraction of one full-time engineer (1.0 = 100% capacity)
$N$	Number of concurrent AI coding sessions
$\alpha$	Per-session base load coefficient. The minimum human attention each session demands for review and direction. Empirically observed: $\alpha \approx 0.25$
$\beta$	Pairwise coordination overhead coefficient. The additional effort created by each pair of sessions that might produce conflicting output. Empirically observed: $\beta \approx 0.12$
$\gamma$	Runtime debugging multiplier. The additional effort caused by deployment issues that only manifest when outputs from multiple sessions are integrated. Empirically observed: $\gamma \approx 0.08$
$\delta(N)$	File conflict probability function, modeled as $N^2/C$ where $C$ is the total number of independent files in the codebase. For a typical enterprise application with $C \approx 200$ files and sessions touching 10–20 files each, this produces meaningful values at $N \geq 3$ .

### 4.2 Model Predictions vs. Observations

Applying the empirically fitted coefficients to the formula yields the following predictions:

Sessions (N)	Base Load ( $\alpha N$ )	Pairwise Overhead	Debug Overhead	Total H(N)	Observed Capacity
1	0.25	0.00	0.00	<b>0.25</b>	~40–50%
2	0.50	0.12	0.02	<b>0.64</b>	~75–85%
3	0.75	0.36	0.07	<b>1.18</b>	~100–120%
4	1.00	0.72	0.13	<b>1.85</b>	>120%
5	1.25	1.20	0.20	<b>2.65</b>	Requires 2nd engineer

Table 2. AAERF predictions vs. empirical observations

The model correctly predicts the critical threshold observed at  $N = 3$ , where  $H(N)$  exceeds 1.0 and the human operator is saturated. At  $N = 5$ , the formula predicts  $H(N) = 2.65$ , indicating that approximately 2.5 full-time-equivalent human engineers are needed for coordination alone—producing zero lines of code themselves.

### 4.3 The Output Multiplier

While coordination cost grows superlinearly, raw output grows approximately linearly:

$$O(N) = N \cdot \mu \cdot (1 - \lambda \cdot \ln(N))$$

Where  $\mu$  is the single-session output multiplier (empirically 5–7× vs. a human developer) and  $\lambda$  is the quality degradation factor due to context fragmentation ( $\lambda \approx 0.08$ ). The logarithmic degradation term reflects the observation that each additional session slightly reduces per-session quality as the human operator has less time for review.

## 5. Implications for Workforce Planning

### 5.1 The Role Transformation

Our findings suggest that AI coding assistants do not eliminate engineering roles; they transform them. The traditional software engineer role bifurcates into two distinct functions:

7. **AI Orchestration Engineers** — senior technical staff who manage parallel AI sessions, make architectural decisions, sequence deployments, and maintain system coherence. These roles require deep system knowledge, architectural judgment, and the ability to context-switch rapidly. They write little to no code directly.
8. **AI Session Specialists** — the AI systems themselves, which produce code at high velocity but require continuous human direction. Each session is, in effect, a junior developer with photographic memory but no awareness of what the developer at the next desk is doing.

### 5.2 Optimal Team Sizing

The AAERF provides a framework for team sizing. For a team managing  $N$  total concurrent AI sessions:

$$\textit{Engineers Required} = \lceil H(N) / W \rceil$$

Where  $W$  is the sustainable work capacity factor per engineer. For standard 8-hour shifts,  $W = 0.85$  (accounting for meetings, breaks, and non-coordination tasks). For intensive development sprints,  $W$  may approach 1.0 but is not sustainable long-term.

Note that the formula explicitly does not account for the unsustainability of extended working hours. In observed cases where a single engineer managed three sessions over multi-hour overnight sessions, output was extraordinary but the pace was not replicable across a standard workforce.

### 5.3 Economic Analysis

Consider an organization currently employing 10 mid-level developers at \$120,000/year each (\$1.2M annually). Under an AI-augmented model:

- 3 AI Orchestration Engineers at \$180,000/year: \$540,000
- Each managing 3 AI sessions: 9 total sessions
- AI service costs (API, compute): ~\$50,000/year
- Total cost: \$590,000 (51% reduction)
- Output: 9 sessions  $\times$  5 $\times$  multiplier = ~45 $\times$  single-developer output (vs. 10 $\times$  previously)

The organization achieves **4.5 $\times$  the output at half the cost**—but with three highly skilled engineers working at full capacity, not ten developers working at moderate capacity. The skill premium for orchestration engineers will likely drive salaries above the \$180,000 estimate as demand grows.

## 6. Limitations and Future Work

This paper presents initial empirical observations from a limited number of organizations and development contexts. Several limitations should be noted:

- Observations are drawn primarily from full-stack web application development; applicability to other domains (embedded systems, data science, mobile development) has not been established.
- The AAERF coefficients are empirically fitted to a small sample and should be refined through broader data collection.
- AI capabilities are evolving rapidly; the coordination overhead may decrease as AI systems develop better awareness of concurrent work.
- The formula does not model the quality-adjusted output, which may differ significantly from raw output metrics.

Future work should focus on validating the AAERF across diverse development contexts, refining coefficients through controlled experiments, and investigating tools and workflows that may reduce the coordination overhead—particularly automated conflict detection, shared context protocols between AI sessions, and AI-to-AI coordination systems.

## 7. Conclusion

The prevailing narrative that AI will replace software engineers oversimplifies a complex transformation. Our observations demonstrate that AI coding assistants are extraordinarily powerful production tools that multiply output by an order of magnitude—but they introduce a coordination overhead that grows superlinearly with parallel usage.

At three or more concurrent sessions, the human managing them is no longer a developer. They are an **orchestrator**—a role that requires deeper technical judgment, faster decision-making, and broader system awareness than traditional development. This role is harder to fill, not easier.

Organizations that approach AI adoption with a headcount-reduction mindset may be disappointed. Those that approach it as a role-transformation opportunity—investing in orchestration capabilities and treating AI sessions as team members that need management—will unlock the extraordinary output potential while maintaining system quality and architectural coherence.

The AI-driven Software Engineer Resourcing Formula provides a starting framework for this planning. We encourage the software engineering community to contribute additional empirical data to refine the model and to develop the tooling that will eventually reduce the coordination overhead that currently constrains AI-augmented development at scale.

---

For inquiries about AI-driven engineering consulting and the AAERF framework:  
[Aurix AI Solutions](#) | [aurixai.solutions](#) | [agenta.red](#)