

# From $O(n^2)$ to $O(\log n)$ : The Complexity Theory of Software Teams

A Whitepaper on Why AI Agents Break Brooks's Law

---

## Software Teams Have a Scaling Problem

---

In 1975, Fred Brooks published *The Mythical Man-Month* and gave us one of computing's most durable laws: **adding people to a late software project makes it later**. Fifty years later, every engineering leader has lived this truth. The math behind it is straightforward — and damning.

A team of  $n$  people has  $n(n-1)/2$  communication channels. Two engineers have one channel. Five have ten. Ten have forty-five. Twenty have one hundred and ninety. The work scales linearly, but the coordination overhead scales quadratically. This is  $O(n^2)$  complexity, and it is the fundamental constraint on every engineering organization ever built.

Every process innovation in software development has been an attempt to manage this quadratic explosion. Agile breaks teams into small squads to keep  $n$  low. Microservices reduce inter-team dependencies. Conway's Law tells us our systems mirror our communication structures — because communication is the bottleneck.

We've spent fifty years optimizing around a constraint we couldn't eliminate. AI agents eliminate it.

---

# The Complexity Classes of Software Development

---

To understand why agents change the equation, we need to be precise about the complexity of different development models. Let's define  $n$  as the number of tasks or features to be delivered.

## Traditional Teams: $O(n^2)$

In a human team, each new task doesn't just add work — it adds coordination. A new feature needs design review, code review, QA handoff, deployment coordination, and ongoing communication about dependencies and scope. Every task interacts with every other task through the humans working on them.

Real-world scaling confirms this. Doubling a team's size rarely doubles its output. Studies consistently show that individual productivity *decreases* as team size grows. A ten-person team doesn't produce five times the output of a two-person team — it produces maybe two to three times the output, because the rest is consumed by coordination.

Brooks's Law is  $O(n^2)$  in disguise. The total cost of delivering  $n$  features with a human team is:

$$T(n) = \textit{work}(n) + \textit{coordination}(n) = O(n) + O(n^2) = O(n^2)$$

The coordination term dominates. Always.

## Parallelized Human Teams: $O(n \log n)$

Smart organizations fight  $O(n^2)$  by decomposing into independent squads. This is the Spotify model, the Amazon two-pizza team, the microservices boundary. By reducing cross-team dependencies, you bring coordination down from quadratic to something closer to  $O(n \log n)$  — each team operates independently, with only logarithmic overhead for cross-team alignment.

This is the best-case scenario for human organizations. It requires excellent architecture, clear domain boundaries, and strong technical leadership. Most organizations never get here. And even when they do, they hit a ceiling: you still need humans in every team, and each human still has a finite throughput.

## AI Agents: $O(\log n)$

AI agents don't have communication channels with each other. They don't need standups, design reviews, or sprint planning. They don't misunderstand requirements because they had a bad morning. They don't take two days to context-switch into a new codebase.

When you assign  $n$  tasks to AI agents, they execute in parallel. The total time is bounded not by coordination but by the **depth of dependency** — the longest chain of tasks where each depends on the output of the previous one. In a well-architected system, this depth is logarithmic relative to the total number of tasks.

$$| \quad T(n) = O(\log n)$$

Ten tasks? A few steps deep. A hundred tasks? Maybe twice that. A thousand tasks? Three times. The coordination overhead doesn't just shrink — it functionally disappears. What remains is the irreducible dependency chain, and that grows logarithmically.

This isn't theoretical. This is what happens when you give Claude Code a codebase and a list of twenty independent changes. It doesn't do them sequentially. It parallelizes, and the wall-clock time barely increases with the number of tasks.

---

## The Math Behind the Speedup

---

Let's make this concrete. Consider a project with 100 features to deliver.

Model	Complexity	Relative Cost (n=100)	Relative Cost (n=1000)
Traditional team	$O(n^2)$	10,000	1,000,000
Optimized squads	$O(n \log n)$	664	9,966
AI agents	$O(\log n)$	6.6	10
Theoretical limit	$O(1)$	1	1

The numbers are illustrative, not literal — you can't reduce 1,000 features to a cost of 10 in absolute terms. But the *ratios* are real. Moving from  $O(n^2)$  to  $O(\log n)$  on a 100-feature project is a **1,500x improvement** in coordination efficiency. At 1,000 features, it's **100,000x**.

This is why “30x faster” undersells it. The speedup isn't constant — it *increases with scale*. The larger your project, the more dramatic the advantage. A solo founder building a ten-feature MVP might see a 3-5x improvement. An enterprise delivering a thousand-feature platform sees orders of magnitude.

---

## Why This Isn't Just Parallelism

---

A reasonable objection: “This is just parallel computing applied to project management. We've known about Amdahl's Law since 1967.”

Yes — and that's exactly the point. Software engineering has been weirdly resistant to applying its own principles to itself. We build distributed systems that scale horizontally, then manage the teams building them with a process that scales quadratically.

Amdahl's Law tells us the speedup is limited by the *serial fraction* — the portion of work that can't be parallelized. In software development, the serial fraction is:

- **Architectural decisions** that constrain everything downstream
- **Strategic direction** that determines what to build
- **Trust boundaries** where human judgment is irreplaceable
- **Sequential dependencies** where feature B requires feature A

Everything else — implementation, testing, code review, documentation, deployment, monitoring — is parallelizable. And agents can parallelize it without the coordination overhead that humans impose.

The serial fraction of most software projects is surprisingly small. Maybe 5-15% of the total work requires sequential human decision-making. The rest is implementation, and implementation is embarrassingly parallel when you remove human coordination from the equation.

# Brooks's Law, Revised

---

Brooks's original law: *Adding manpower to a late software project makes it later.*

The reason: each new person adds communication channels quadratically, and the time spent coordinating eventually exceeds the time saved by the additional hands.

## Brooks's Law doesn't apply to agents.

Adding a new AI agent adds zero communication channels. It adds zero meetings. It adds zero context-switching overhead for existing team members. It adds capacity with near-zero marginal coordination cost.

The revised law for the agentic era:

*Adding agents to a software project makes it faster — with diminishing but never negative returns.*

The returns diminish because of Amdahl's Law: eventually you hit the serial fraction and more agents can't help. But they never go *negative* the way adding humans does. You never make a project *later* by adding agents.

This is the fundamental break from fifty years of software engineering orthodoxy. The constraint that shaped every methodology, every org chart, every process — that human coordination is expensive and scales poorly — no longer applies to the majority of software development work.

---

## What $O(\log n)$ Leadership Looks Like

---

If the execution layer runs at  $O(\log n)$ , what does the leadership layer look like?

### The CTO as Architect, Not Manager

Traditional CTOs spend 60-70% of their time on people management: hiring, one-on-ones, performance reviews, team structure, conflict resolution, career development. This is necessary because the humans are the execution layer, and humans need management.

When agents are the execution layer, the CTO role collapses to its essential core:

- **Define the architecture** that determines how agents decompose work
- **Set the quality gates** that ensure output meets standards
- **Make the judgment calls** that require experience and context no agent has
- **Own the trust boundary** between automated and human decision-making

This is a CTO who operates at  $O(1)$  relative to project size. The number of strategic decisions doesn't scale with the number of features. Whether you're shipping ten features or a thousand, the architectural decisions, quality standards, and trust boundaries are roughly the same.

## The Engineer as Orchestrator

Individual engineers shift from writing code to orchestrating agents. The skill set changes:

$O(n^2)$ Engineer	$O(\log n)$ Engineer
Writes code line by line	Defines intent and reviews output
Attends standups and sprint planning	Configures agent pipelines and quality gates
Estimates story points	Defines success criteria and constraints
Does manual code review	Builds automated review systems
Context-switches between tasks	Runs tasks in parallel through agents
Productivity limited by typing speed	Productivity limited by clarity of thought

The best engineers in the  $O(\log n)$  world won't be the fastest coders. They'll be the clearest thinkers — the ones who can decompose a problem into parallelizable intents with precise success criteria.

# The $O(1)$ Horizon

---

$O(\log n)$  isn't the end state. The theoretical limit is  $O(1)$  — constant time regardless of project size. This seems impossible, but consider what it actually requires:

1. **Fully autonomous decomposition:** agents break down any intent into parallelizable sub-tasks without human guidance
2. **Self-healing architecture:** agents detect and resolve conflicts between parallel workstreams automatically
3. **Unlimited horizontal scaling:** you can spin up as many agents as the problem requires
4. **Zero serial fraction:** no human decision is required between intent and delivery

We're not there yet. The serial fraction — strategic decisions, trust boundaries, regulatory judgment — will likely always require humans. But the serial fraction is shrinking every quarter as agents get more capable.

The practical target for 2026-2027 is  $O(\log n)$  with a small constant. That alone represents a paradigm shift in what's possible.

---

# Implications for Your Organization

---

## If You're a Startup

Stop hiring to scale output. Every engineer you add brings  $O(n^2)$  communication overhead. Instead, invest in agent infrastructure that lets a small team (2-5 people) operate at the output level of a 50-person organization. Keep  $n$  small for humans. Let agents handle the parallelizable work.

## If You're an Enterprise

Your existing  $O(n^2)$  structure — hundreds of engineers across dozens of teams — is your biggest liability. Each cross-team dependency is a coordination tax. Start by identifying which teams are doing  $O(\log n)$ -compatible work (implementation, testing, deployment) and replacing those coordination channels with agent pipelines.

## If You're a Technical Leader

Your value just shifted from managing  $O(n^2)$  complexity to eliminating it. The CTOs and engineering leaders who thrive in the  $O(\log n)$  world will be the ones who understand both the computer science and the organizational implications — who can redesign their orgs around agents the way the last generation redesigned them around microservices.

---

# Conclusion

---

For fifty years, Brooks's Law has been the immovable constraint at the center of software engineering. Every methodology, every org structure, every process has been designed to manage the  $O(n^2)$  reality of human coordination.

AI agents don't just improve the constant factor. They change the complexity class. From  $O(n^2)$  to  $O(\log n)$ . This is not an incremental improvement — it is a *phase transition* in how software gets built.

The organizations that recognize this will restructure accordingly. The ones that don't will keep adding humans to late projects and wondering why they keep getting later.

The math doesn't lie. It never has.

---

# References

---

- Brooks, Fred. *The Mythical Man-Month: Essays on Software Engineering*. Addison-Wesley, 1975.
  - Amdahl, Gene. "Validity of the Single Processor Approach to Achieving Large Scale Computing Capabilities." AFIPS Conference Proceedings, 1967.
  - Jansen, Remo. "Agent Driven Development (ADD): The Next Paradigm Shift in Software Engineering." DEV Community, July 2025.
  - Putnam, Lawrence & Myers, Ware. "Team Size Can Be the Key to a Successful Project." QSM Associates, 2013.
  - Scholtes, Peter R. "An Elaboration on Amdahl's Law." Parallel Computing, 2022.
- 

*27 delivers AI-powered technical leadership as a monthly subscription.*

*Strategy, architecture, and execution — not hourly billing.*

*Learn more at [27consulting.com/pricing](https://27consulting.com/pricing).*