ANDYGROVE.SYSTEMS CONTINUING PROFESSIONAL UPSKILLING BRIEF // PP-03 (arXiv:2604.02460)
Open-Access Executive & Engineering Study Guide on Fresh AI Operations Preprints
Full Preprints Library: https://andygrove.systems/preprints
42-Hour CPD Curriculum: https://andygrove.systems/education
Canonical Brief URL: https://andygrove.systems/preprint/pp-03
Primary Source URL: https://arxiv.org/abs/2604.02460

TITLE: Single-Agent LLMs Outperform Multi-Agent Systems on Multi-Hop Reasoning
  Under Equal Thinking Token Budgets
AUTHORS: Dat Tran, Douwe Kiela (Stanford University) (Apr 2026 * Stanford)
PILLAR: TOKEN BUDGET & TOPOLOGY
---
1. EMPIRICAL & MATHEMATICAL FINDING FOR CONTINUING PROFESSIONAL UPSKILLING:
  Uses the Data Processing Inequality (DPI) to prove that under equal thinking-token
  budgets, multi-agent message passing loses mutual information and trails
  single-agent reasoning unless context rot on noisy long inputs degrades
  single-agent focus.
---
2. ACTIONABLE COO & INDUSTRIAL ENGINEERING PRODUCTION RULE:
  Audit every multi-agent architecture at equal token budget; eliminate decorative
  inter-agent chatter unless noisy-context compression or privilege isolation
  requires it.
---
3. SELF-STUDY & TEAM UPSKILLING INTEGRATION (GOOGLE CLASSROOM & COLAB):
  * Interactive Calculators (L0-L3 Wilson TRM & L = lambda*W): https://andygrove.systems/calculators
  * Runnable Python Queueing & Routing Colab Lab: https://andygrove.systems/colab.ipynb
  * 1-Page Agent Line Audit Specimen: https://andygrove.systems/audit
---
NODE 01 (Part-Time Remote IE/OR/BizOps Pods, $20-45/hr): https://andygrove.systems/#apply
NODE 02 (COO Office Audit Desk): mailto:output@andygrove.systems
Independent platform; not affiliated with Intel Corporation or the Grove family.