

I'm SO Glad I'm Uncoordinated!

Coordination Is Increasingly Painful... What Can Be Done?

Pat Helland

Salesforce

phelland@salesforce.com

ABSTRACT

In my 44 years building software, technology trends have dramatically changed what's difficult and what's hard. In 1978, CPU, storage, and memory were precious and expensive but coordinating across work was effectively free. Now, we see an abundance of CPU, storage, memory, and bandwidth but we lose a LOT of opportunities for work when we coordinate across parallel activities.

This presentation examines the nature of coordination and concludes that it is really aligning disparate arcs of the computations within the work's directed acyclic graph. Parallel systems employ a partial order or DAG (Directed Acyclic Graph) of their work. By aligning these arcs of computation, coordination is "the partial collapse of partial order". There are many tricks we use to reduce how often we coordinate and we attempt to taxonomize some of these and explain how they are frequently used to enhance the performance of our systems. We consider the separation of work in space, time, layering, equivalence, and also confluence. Looking at these properties helps us understand the nature of our systems and the implicit or explicit parallelism we use within them. By accepting that coordination is challenging, we can design systems to reduce the pain of these operations.

Personally, I am a clumsy and uncoordinated person. I've adapted to this in my personal life by reducing the need to coordinate. Similarly, our computer systems designs work best when designed to be uncoordinated.

ACM Reference Format:

Pat Helland. 2022. I'm SO Glad I'm Uncoordinated!: Coordination Is Increasingly Painful... What Can Be Done?. In *The 16th ACM International Conference on Distributed and Event-based Systems (DEBS '22)*, June 27–30, 2022, Copenhagen, Denmark. ACM, New York, NY, USA, 1 page. <https://doi.org/10.1145/3524860.3544407>

1 BIOGRAPHY

Pat Helland has been building distributed systems, database systems, high-performance messaging systems, and multiprocessors since 1978, shortly after dropping out of UC Irvine without a bachelor's degree. That hasn't stopped him from having a passion for academics and publication. From 1982 to 1990, Pat was the chief architect for TMF (Transaction Monitoring Facility), the transaction logging and recovery systems for NonStop SQL, a message-based fault-tolerant system providing high-availability solutions for business critical solutions. In 1991, he moved to HaL Computers where he was chief architect for the Mercury Interconnect Architecture, a cache-coherent non-uniform memory architecture multiprocessor. In 1994, Pat moved to Microsoft to help the company develop a business providing enterprise software solutions. He was chief architect for MTS (Microsoft Transaction Server) and DTC (Distributed Transaction Coordinator). Starting in 2000, Pat began the SQL Service Broker project, a high-performance transactional exactly-once in-order message processing and app execution engine built deeply into Microsoft SQL Server 2005. From 2005-2007, he worked at Amazon on scalable enterprise solutions, scale-out user facing services, integrating product catalog feeds from millions of sellers, and highly-available eventually consistent storage. From 2007 to 2011, Pat was back at Microsoft working on a number of projects including Structured Streams in Cosmos. Structured streams kept metadata within the "big data" streams that were typically 10s of terabytes in size. This metadata allowed affinity placement within the cluster as well as efficient joins across multiple streams. On launch, this doubled the work performed within the 250PB store. Pat also did the initial design for Baja, the distributed transaction support for a distributed event-processing engine implemented as an LSM atop structured streams providing transactional updates targeting the ingestion of "the entire web in one table" with changes visible in seconds. Starting in 2012, Pat has worked at Salesforce on database technology running within cloud environments. His current interests include latency bounding of online enterprise-grade transaction systems in the face of jitter, the management of metastability in complex environments, and zero-downtime upgrades to databases and stateful applications. In his spare time, Pat regularly writes for ACM Queue, Communications of the ACM, and various conferences. He has been deeply involved in the organization of the HPTS (High Performance Transactions Systems - www.hpts.ws) workshop since 1985. His blog is at pathelland.substack.com and he parsimoniously tweets with the handle @pathelland.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

DEBS '22, June 27–30, 2022, Copenhagen, Denmark

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9308-9/22/06.

<https://doi.org/10.1145/3524860.3544407>