KALEIDO RESEARCH ABSTRACT

Graph Databases for Project Controls

Author: Jake Williams

Published: 2026-04-09

Source: AACE Technical Paper DSAA-4619

Abstract

This paper presents a framework for developing knowledge graphs for project controls, connecting cost, schedule, risk, and procurement data so teams can reason across dependencies that siloed systems routinely hide.

Why this matters for Kaleido

For Kaleido, the paper is a research basis for treating programme context as a decision asset rather than a document archive. It underpins the product direction around connected knowledge, project context intelligence, and graph-backed reasoning.

Source excerpt

This paper presents a framework for developing knowledge graphs (KG) for project controls in large-scale capital projects. Current project control systems often suffer from fragmented data sources, with cost, schedule, risk, and procurement information typically residing in separate platforms. As McKinsey research shows, 98% of megaprojects suffer cost overruns exceeding 30%, with 77% running at least 40% late [1]. Part of this problem can be attributed to information silos that prevent holistic project visibility and optimal decision-making. The paper outlines a structured methodology for implementing knowledge graphs within existing project control frameworks, including ontology design, query capabilities, and practical applications. The proposed approach leverages graph database technology to overcome the limitations of static risk registers and rigid relational databases by creating a connected network of project entities (tasks, costs, risks, constraints) and their relationships. Through a hypothetical case study, the paper demonstrates how this framework enables complex queries like "If our assumption is invalid, which interface agreements and milestone dates are at risk, and what is the total potential cost impact?" that would be nearly impossible to trace with siloed, legacy systems. By adopting a KG-driven approach with potential AI integration, project teams can transition from static, siloed data management to a dynamic, interconnected system that improves decision-making, risk detection, and execution efficiency. This seeks to address the "iron law of megaprojects: over budget, over time, over and over again" [2] by providing real-time impact analysis and continuous risk updates across traditional project control boundaries.

Publication note: This PDF is a launch-managed public abstract generated from Kaleido's local vault source. It is not the final publisher-approved technical paper PDF.

Vault source: wiki/research/dsaa-4619-graph-databases.md