

John Quevedo

(475) 777-8330 | john.quevedo@yale.edu | <https://linkedin.com/in/john-quevedo> | <https://github.com/johnquevedo>

EDUCATION

Yale University

New Haven, CT

Bachelor of Science in Computer Science and Mathematics, GPA: 3.84

Expected May 2028

Awards/Honors: Jane Street Unboxed Scholar, Hispanic Heritage Foundation Fellow

Relevant Coursework: Data Structures, Algorithms, Randomized Algorithms, Machine Learning, Statistical Learning Theory, Advanced Probability, Theory of Statistics, Linear Algebra, Combinatorics, Abstract Algebra

TECHNICAL SKILLS

Languages: Java, Go, C++, Python, TypeScript, SQL

Frameworks/Data: Spring Boot, Node.js, React, Kafka, PostgreSQL, Redis, Spark

Infrastructure: AWS, Docker, Kubernetes, GitHub Actions, Git, Linux

WORK EXPERIENCE

Yale University

New Haven, CT

Algorithms Research Intern

Jun 2026 – Present

- Designed and implemented a novel parallel graph algorithm in C++ that updates clique counts after batches of edge insertions and deletions, avoiding a full recount after each batch.
- Optimized clique-counting routines through parallel work partitioning, achieving up to a 50.8× speedup on 64 cores over single-threaded execution.
- Debugged a concurrency race causing nondeterministic undercounts; passed 26 audits against full recomputation and produced consistent counts across 105 configurations.

Management Leadership for Tomorrow

Remote

Generative AI Software Engineer

Jun 2026 – Present

- Built a full-stack financial document analysis platform using Python, AWS Lambda, Bedrock, CDK, LangChain, RAG, MCP, and SEC EDGAR APIs.
- Developed CIK lookup and EDGAR ingestion modules to map companies to filings, retrieve 10-Q/10-K data, and prepare documents for LLM question answering.
- Deployed serverless inference infrastructure with AWS SAM/CDK and integrated it with a web interface for end-to-end SEC filing analysis.

PROJECTS

[CI Failure Capture and Replay](#) (Go, Python, PostgreSQL, Docker, Kubernetes, GitHub Actions)

- Built a CI failure capture-and-replay pipeline that preserved code and runtime state, reproducing 50/50 controlled failures versus 0/50 clean reruns.
- Engineered multi-tenant orchestration with idempotent jobs, lease-based queues, and checkpointed workers; preserved job state through duplicate webhooks, worker crashes, and database restarts.
- Hardened replay containers with non-root, read-only, network-disabled execution and artifact verification; completed 172/172 controlled replays, including two 50-job concurrent runs.

[Online Payment Repair](#) (Java, Kafka, PostgreSQL, Spark, Kubernetes)

- Built a data-repair workflow that traces corrupted payments to responsible code or configuration versions and generates hashed before-and-after repair plans.
- Prevented 92 stale repairs from overwriting concurrent updates with optimistic locking while 8,787 writes completed without errors.
- Repaired 100,000 records at 865 records per second with zero duplicates, incorrect fees, or ledger imbalances in controlled tests; completed all 5,000 repairs after injected failures.

LEADERSHIP AND PROFESSIONAL DEVELOPMENT

Management Leadership for Tomorrow

Washington, DC

Career Preparation Fellow

Jan 2026 – Present

- Accepted into a selective 18-month professional development program that accelerates the career growth of emerging leaders through structured coaching, mentorship, and targeted skill-building.
- Engage with leading partner organizations (e.g., LinkedIn, Bloomberg, and Deloitte) to gain industry exposure and insight into tech-talent pathways.

ADDITIONAL SKILLS AND INTERESTS

Fluent in Spanish; open-source contributor; backend system design