Robust and Consistent Distributed Storage as a Service

Andria Trigeorgi

PhD Supervisor: Chryssis Georgiou, Industrial Supervisor: Nicolas Nicolaou

University of Cyprus & Algolysis Ltd.

OVERVIEW

Challenge: Keeping Everything in Sync

Maintaining consistent data across multiple devices can be difficult due to network delays and potential message loss. Atomic Distributed Shared Memory (ADSM) tackles these issues and is a cornerstone of our research.

ADSM: The Key to Reliable Storage

ADSM creates the illusion of a sequential memory space over asynchronous, fail-prone, message-passing nodes. Existing ADSMs face challenges such as large data support, speeds, liveness, and scalability.

Our focus: Distributed Storage Solutions

We develop algorithms for a Distributed Storage Service (DSS), ensuring strong consistency, large data handling, and high concurrency, even in dynamic environments with changing nodes.

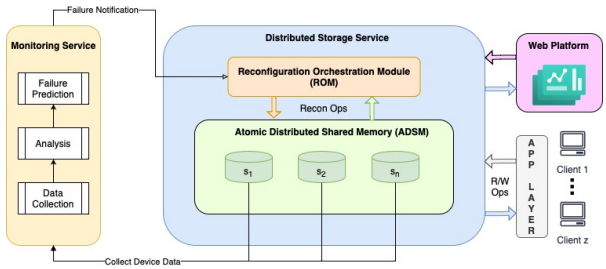

KEY CONTRIBUTIONS

- **CoBFS Framework:** Fragmentation strategy for handling large data, maintaining strong consistency, and allowing concurrent modifications.
- **Integrate the Dynamic ADSM ARES with CoBFS:** Provides dynamic reconfiguration and fault tolerance.
- **Performance Improvement:** Distributed Tracing identifies bottlenecks. Optimizations minimize communication overhead.
- **Evaluation and Comparison:** Tested on emulation and overlay testbeds. Competitive performance against commercial solutions.

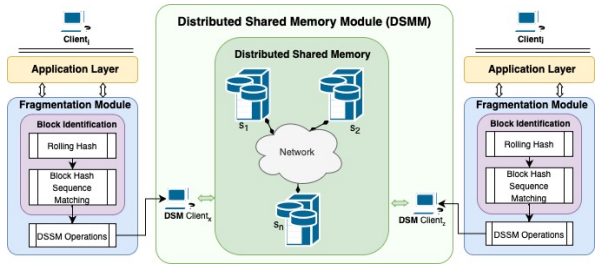

OBJECTIVES

- **Study and Formally Define Principles:** Design a DSS capable of handling large shared atomic data.
- **Reconfiguration Orchestration Module (ROM):** Enable seamless node additions and removals without service interruption.
- **Algorithms' Implementation and Optimization:** Implement, integrate, evaluate, and optimize the proposed algorithms.
- **Fully Functional DSS:** Develop, deploy, and evaluate a DSS with ROM and security guarantees, accessible through a web-based platform.

TESTBEDS: Emulab, AWS EC2, Fed4FIRE

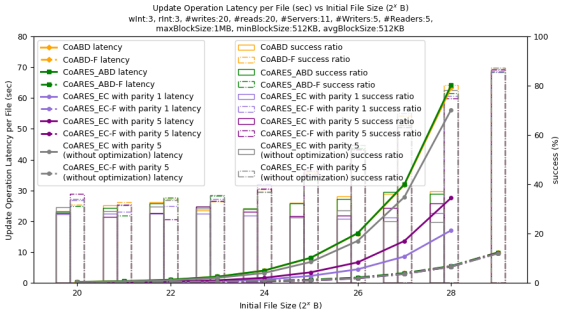

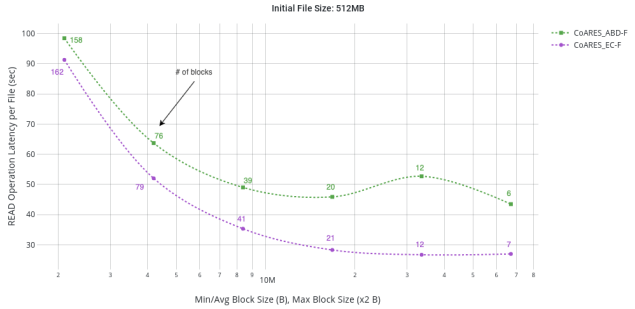

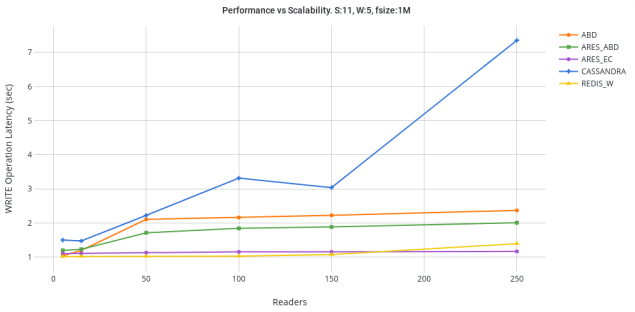

OPTIMIZING ADSM: Insights from Distributed Tracing Analysis

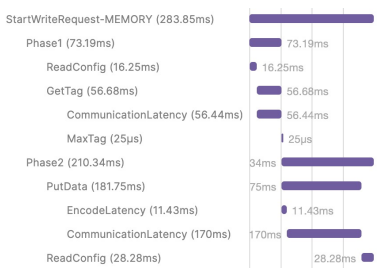

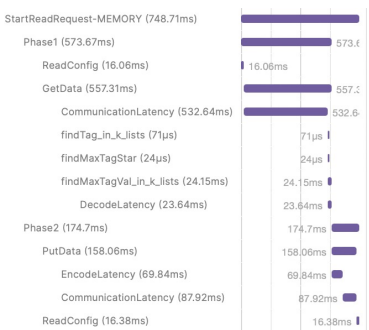

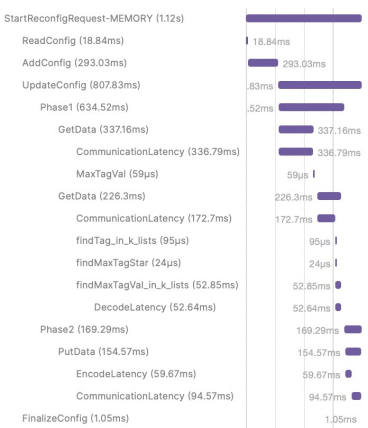

FUNDED BY MaRS-PHD IN INDUSTRY/1222/0121, CHARISMA-DUAL USE/0922/0048

Funded by the European Union NextGenerationEU

Cyprus tomorrow RECOVERY AND RESILIENCE PLAN

Republic of Cyprus

RESEARCH & INNOVATION FOUNDATION

gec24