



Raytheon
BBN

Using Distributed Ledger Technology for Command and Control and Decentralized Operations


Dr. David Last, Raytheon BBN

International Conference on Networks 2022
27 April 2022

- Motivating scenario
- Requirements
- Solution – Distributed Ledger Technology
- CARDIAC – Exploratory study
 - Solution alternatives
 - Prototype systems
 - Experimental results
- Future work


Motivating Scenario

- Air Force base is under threat of immediate attack
- Joint Forces Commander (JFC) splits forces, deploys them Forward Operating Bases
 - Divide forces into Dispersed Units (DU) – four fighters, one tanker, associated ground crews
 - Organize DUs into several levels of Parent Dispersed Units (PDU)
 - Top-level PDU includes JFC, air planning components in Air Operations Center (AOC)




Motivating Scenario

- Prior to dispersion, AOC develops air campaign plan (Commander's Objective, target lists, target assignments) and pushes it to DU aircraft
- During the campaign, DU-PDU communication cut off by jamming, cyber attack, terrain
 - DUs need authority to modify plans in response to pop-up targets, changing battlefield conditions
 - Authority resides with AOC – need to delegate to DUs
 - When communications restored, plan changes must be communicated back to AOC



Dispersed C2 Requirements

- Air Force needs a next-generation Multi-Domain Command and Control (MDC2) system that provides...
 - Distributed information management
 - Conditional Authorities
 - Merging parallel data evolution histories
 - Data evolution provenance




Dispersed C2 Solution

- Build next-gen C2 system that supports
 - Disconnected operations - Conditional authorities calculus
 - Data integrity/provenance
- Our solution – Leverage Distributed Ledger Technology (DLT e.g., blockchain) for C2
 - Leaderless distributed database - Permissioned Access control
 - Smart Contracts (conditional authorities) - Data provenance/integrity
- CARDIAC (Study)
 - Establish feasibility/calibrate C2 data handling on DLT
 - Identify C2 sweet spots/chart path for a full DLT-based C2 program

Distributed Ledger Technology

- What is DLT?
 - Distributed Ledger Technology is a distributed database
 - Multiple nodes maintain a copy of the database
 - Database transactions are bundled into “blocks”, which are appended into a “chain”
 - Blocks are cryptographically signed by their own contents and those of previous block – modifying the block will make these signatures invalid
 - Guarantees **immutability** – any tampering with blocks is easily detected
 - Transaction history can be traced – **provenance traceability**
 - Any network node can be the leader for purposes of creating a block – **decentralized** operation




CARDIAC Study Goals

- Develop Use Case
- Investigate DLT options
- Build Prototypes
 - DLT-based prototype
 - Traditional database baseline
- Experimentation
- Metrics
- Recommendations


DLT Options

- Investigated multiple DLT options
- Selected blockchain – Hyperledger Fabric
 - Open-source blockchain sponsored by Linux Foundation, contribution from IBM
- Why Hyperledger Fabric?
 - Permissioned blockchain
 - Access is tightly controlled to authorized participants
 - Energy efficient
 - Avoids energy-heavy blockchain implementations like Bitcoin




Prototype MDC2 systems

- Goal – Compare DLT to traditional database options
- Approach
 - Develop Hyperledger Fabric-based prototype
 - Develop baseline prototype (PostgreSQL)
 - Perform head-to-head experiments



- ## DLT prototype message flow



Experimentation Plan

Performed Experiments

- **Experiment 0** – Determine Hyperledger Fabric optimal configuration for CARDIAC use case
 - Metrics: Latency, Throughput
- **Experiment 1** – Continuous write Link 16 J2.2 messages, fully-connected network, vary bandwidth
 - Metrics: Latency, Throughput
- **Experiment 3** – Continuous write Link 16 J2.2 messages, fully-connected network; disconnect one database node for a period of time, then re-connect (no writes to disconnected node); vary bandwidth
 - Metrics: DB reconciliation time, DB reconciliation network overhead

Future Experiments

- **Experiment 2** – Same as Experiment 1
 - Metrics: Disk storage, processing, network overhead
- **Experiment 4** – Same as Experiment 3, except (non-conflicting) Link 16 J2.2 messages also written to disconnected database node
 - Metrics: DB reconciliation time, DB reconciliation network overhead
- **Experiment 5** – Continuous write JIPTL messages, fully connected network; disconnect one database node, continue writing (conflicting) JIPTL messages to disconnected node; re-connect node, system must reconcile/de-conflict, database; vary bandwidth
 - Metrics: DB reconciliation/de-confliction time, DB reconciliation/de-confliction network overhead



Results – Experiment 0

- Batch Timeout – wait for more transactions before bundling into blocks
 - High batch timeout: Fewer blocks → Less network overhead → Higher throughput, higher latency
 - Low batch timeout: More blocks → More network overhead → Lower throughput, lower latency
 - BUT Very low batch timeout: many blocks → much network overhead → higher latency
- **Takeaway – DLT is well-suited for low volume, limited bandwidth/compute power networks (battlefield C2), but not low-latency applications (Real time control systems)**

Results – Experiment 0, DLT Batch Timeout

Experimental Conditions

- 2 Clients writing to DB every 500 ms
- Experiment length: 300s





Results – Experiment 0

- Traffic density
 - consensus algorithm generates network overhead
 - Traffic density causes higher latency as message queues grow
- **Takeaway: DLT is more suitable for low volume, high value data (like air battle plans and target lists), not high volume data (like video streams)**

Results – Experiment 0, Traffic Density

Experimental Conditions

- Clients writing to DB every 500 ms
- Experiment length: 300s
- Batch Timeout: 0.5s





Results – Experiment 1

- DLT vs. PostgreSQL
 - Throughput – PostgreSQL outperforms DLT
 - Latency – PostgreSQL outperforms DLT by 1 order of magnitude
 - Key difference – PostgreSQL is single-master – single point of failure (not resilient to contested comms)
 - **Takeaway – DLT provides high integrity/assured propagation in resource-constrained/contested comms environments, as opposed to traditional database solutions**

Results – Experiment 1

Experimental Conditions

- 2 Clients writing to DB every 300 ms
- Experiment length: 300s
- Batch Timeout: 0.5s




Results – Experiment 3

- DLT vs. PostgreSQL
 - Reconciliation time for disconnected node to catch up on transactions it is missing
 - **Takeaway – DLT introduces complexity over traditional database solutions, but that complexity is an acceptable tradeoff in high-disconnection scenarios**

Results – Experiment 3

Experimental Conditions

- Client writes 500 messages
- One node disconnected
- Client writes another 500 messages
- Disconnected node reconnects, downloads new messages



- CARDIAC prototype
 - Self-contained CARDIAC node
 - Configurable network size/topology
 - Deliverable to third parties
- CARDIAC functionality
 - Data evolution
 - Conditional authorities
 - Blockchain forking/merging
- Additional experiments

- Air Force needs a decentralized MDC2 solution
 - Resilient to disrupted communications
 - Empower frontline units with dynamic authorities
- Distributed Ledger Technology is a promising option for these requirements
- CARDIAC study investigated feasibility
- Results – DLT is slower, but more resilient to disrupted environments


Q&A



Test harness stats for experiments

- Single Ubuntu Linux VM running on Virtual Box
 - 4096MB RAM
 - 4 CPUs, 1.9 GHz, 100% execution cap
 - 21GB Hard Drive



Hyperledger Network



Results – Experiment 0, Traffic Density

Experimental Conditions

- Clients writing to DB every 500 ms
- Experiment length: 300s
- Batch Timeout: 0.5s



Results – Experiment 0, Traffic Density

Experimental Conditions

- Clients writing to DB every 500 ms
- Experiment length: 300s
- Batch Timeout: 0.5s

