

9th Workshop on Principles and Practice of Consistency for Distributed Data
Geo-located data for better dynamic replication

Luís Silva, Frederico Aleixo, Albert van der Linde, João Leitão, Nuno Preguiça
NOVA LINCS & Universidade NOVA de Lisboa
05 April 2022




Outline

- Motivation
- Problem to solve
- Protocols (Overlay and Bully)
- Evaluation
- Future Work



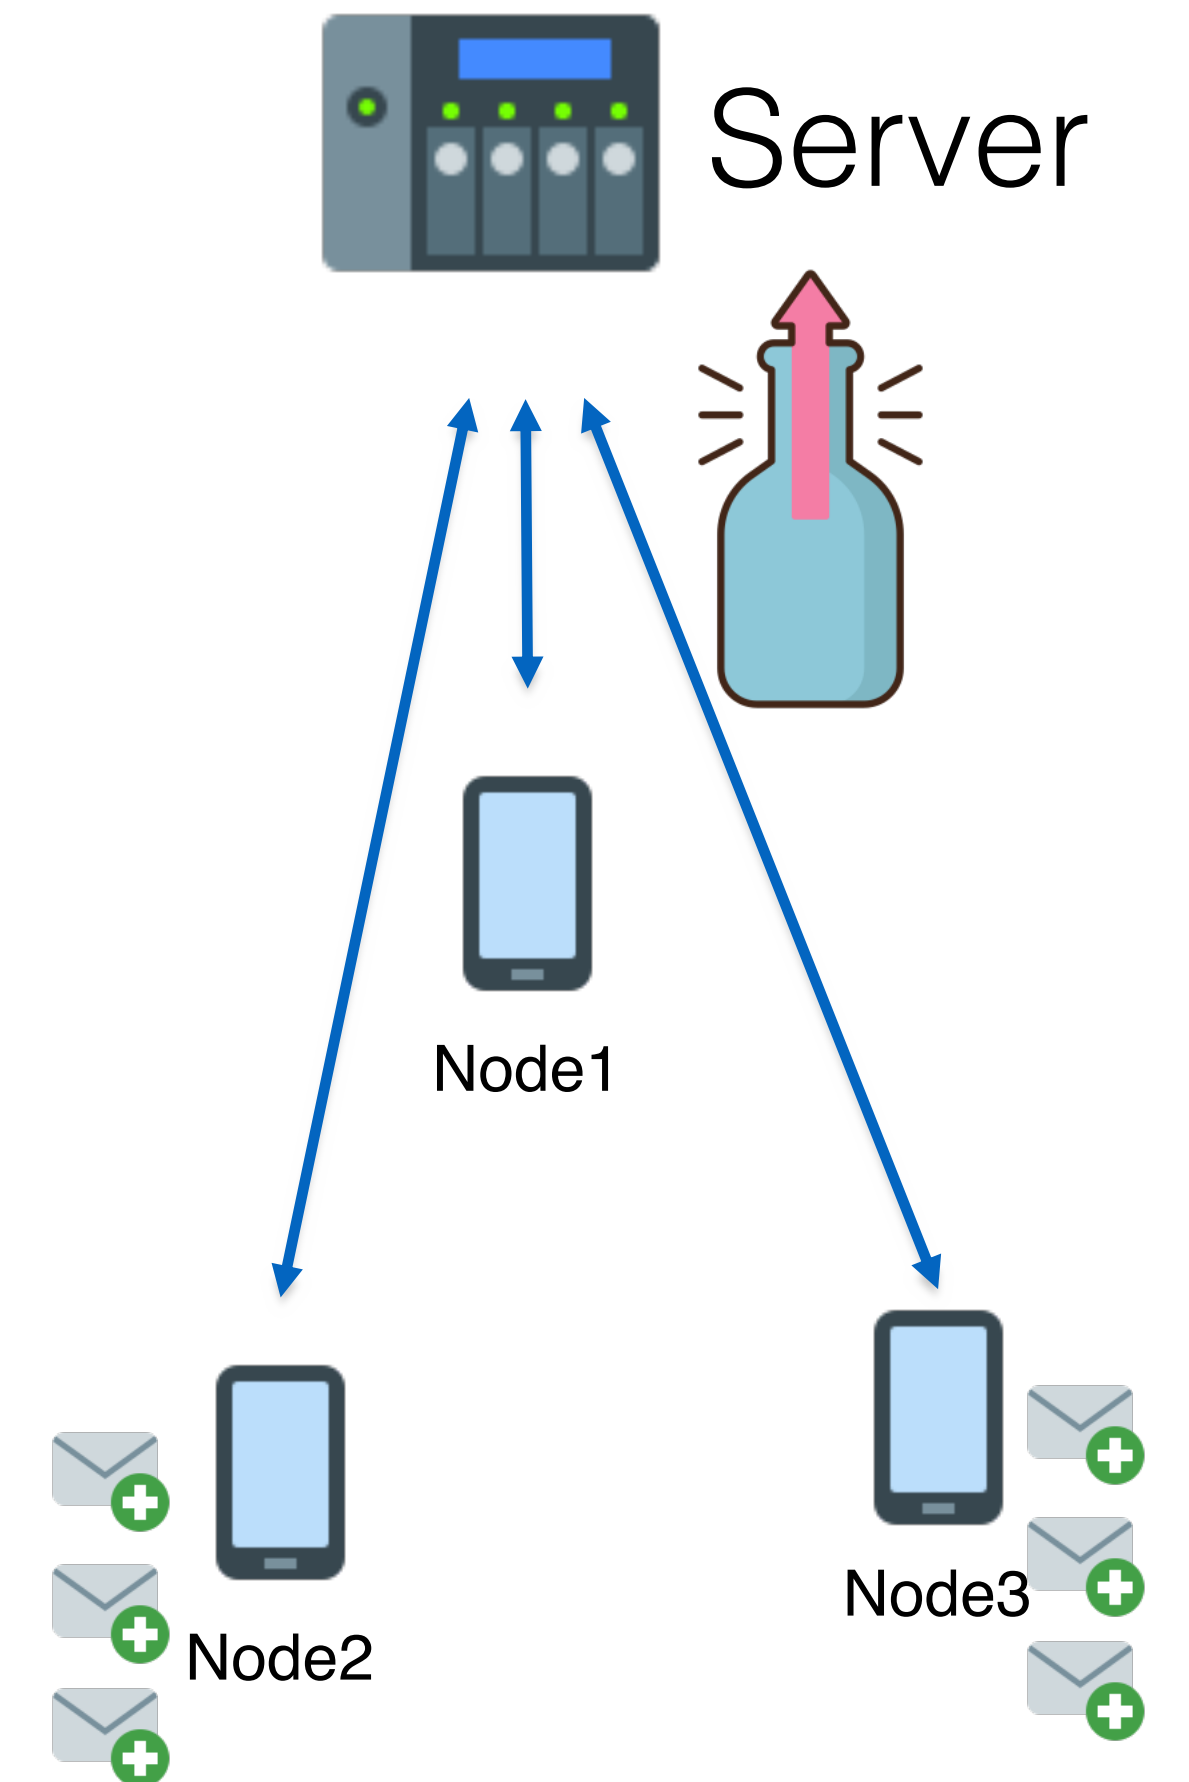
Introduction

- Current applications rely on centralised infrastructures for data storage and coordination.



Introduction


- Current applications rely on centralised infrastructures for data storage and coordination.
- A fast-expanding number of clients and workloads on the server-side swiftly grow as a scalability bottleneck.





Introduction


- A possibility is to rely on peer-to-peer architectures.





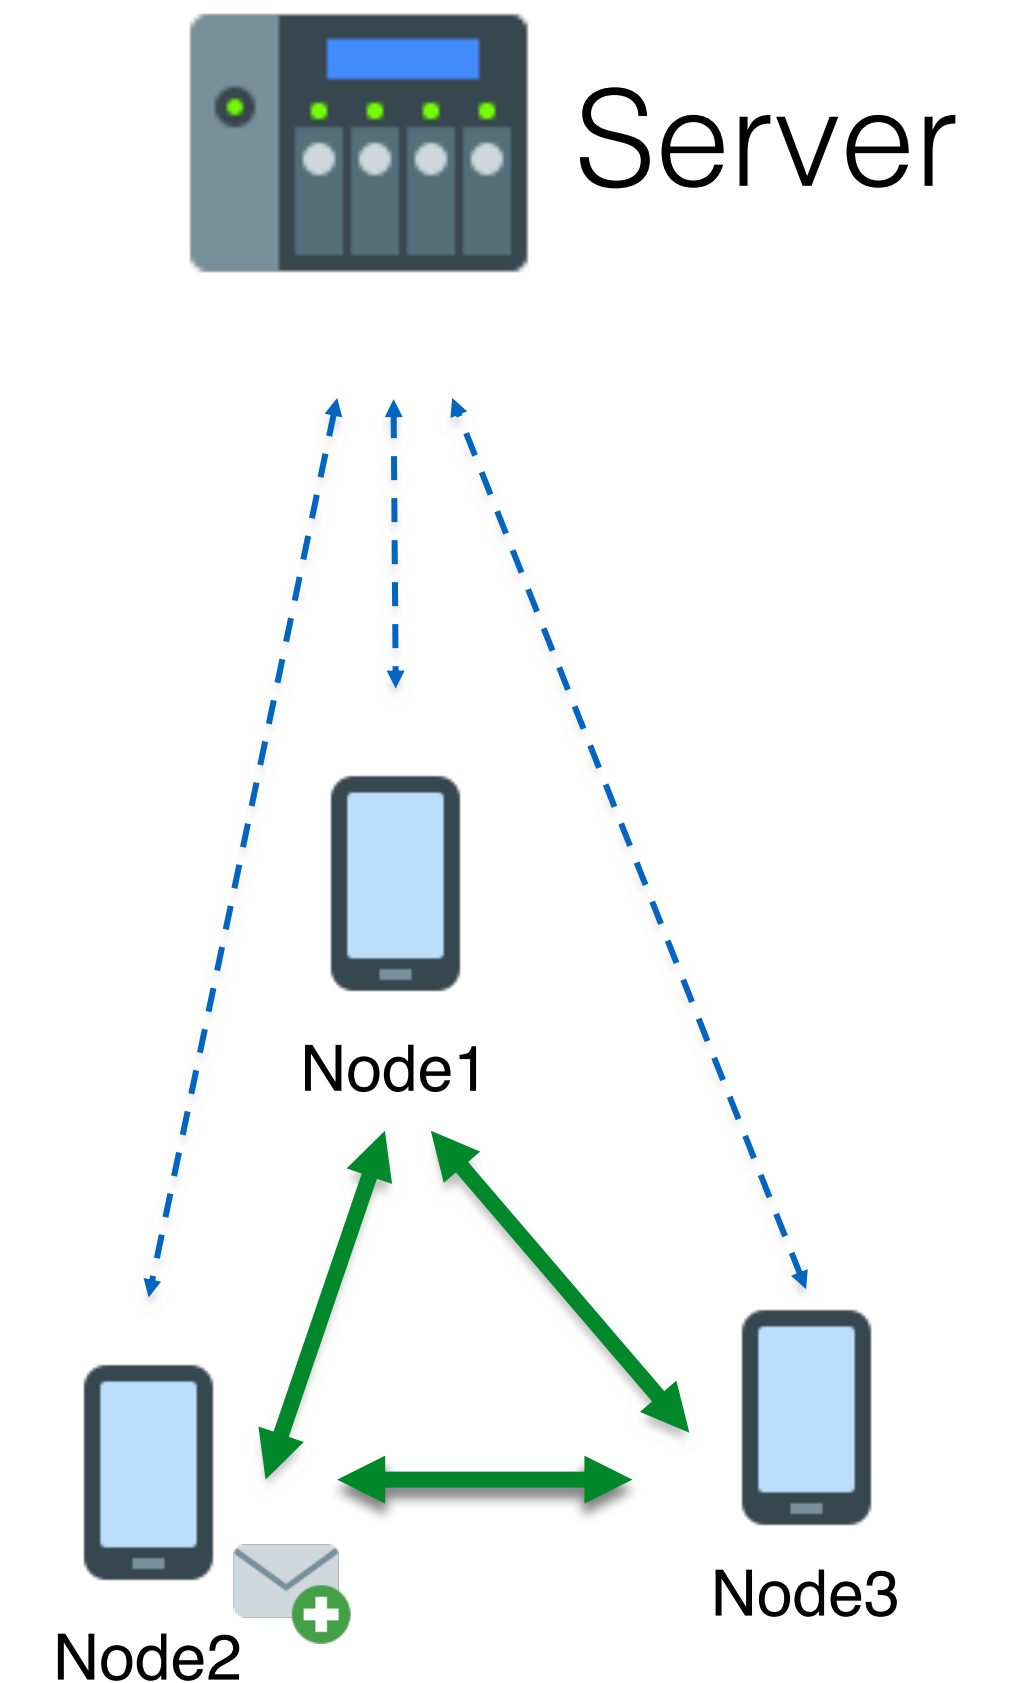
Introduction

- A possibility is to rely on peer-to-peer architectures.
- Where clients communicate and share information directly.



Introduction

- A possibility is to rely on peer-to-peer architectures.
- Where clients communicate and share information directly.
- Systems exist that provide P2P communication among peers and full replication. [Legion]



[Legion] Albert van der Linde, Pedro Fouto, João Leitão, Nuno Preguiça, Santiago Castiñeira, and Annette Bieniusa. 2017. Legion: Enriching Internet Services with Peer-to-Peer Interactions. In Proceedings of the 26th International Conference on World Wide Web (Perth, Australia) (WWW '17). 283–292.



Problem

- These systems don't consider node location.



Server



Node1




Node2





Node3



Problem





- These systems don't consider node location.
- The client's interest set on objects may change with location.






Problem



Server



- These systems don't consider node location.
- The client's interest set on objects may change with location.
- Don't account for possible Edge locations.



Node2



Node3

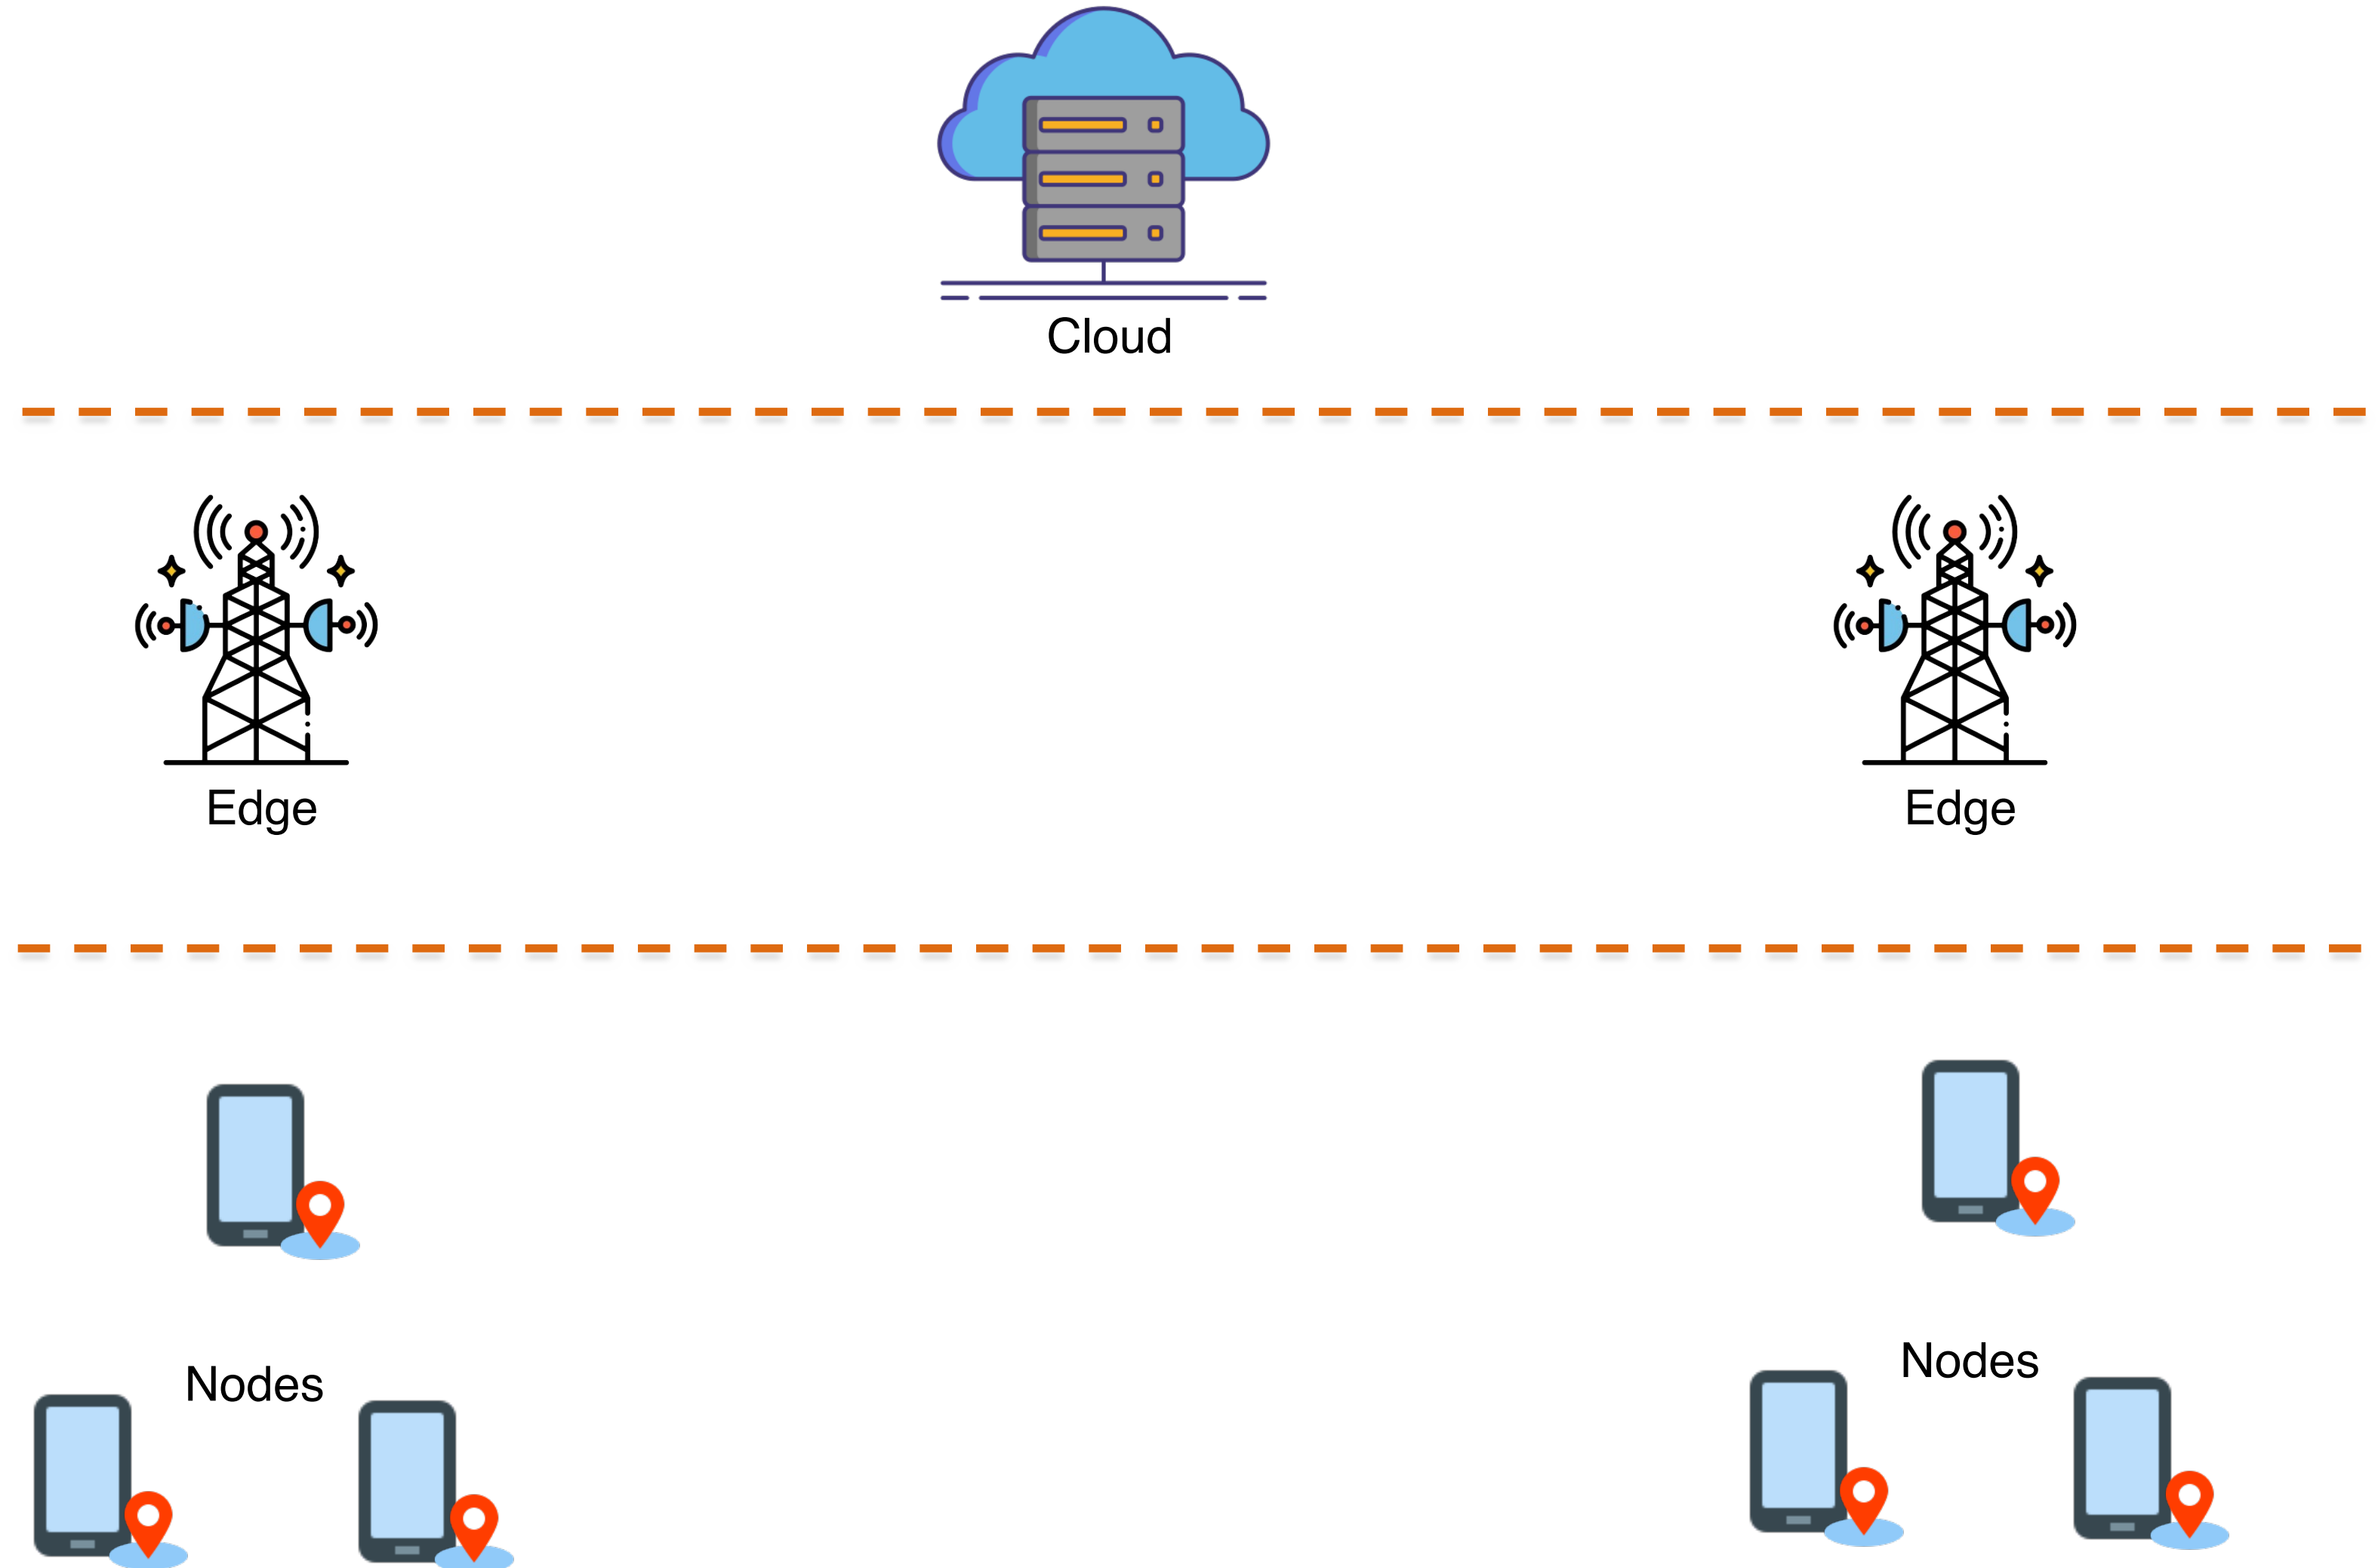


Applications

- Several popular mobile applications use geo-located data as a core data attribute.
 - Augmented reality games
 - Scooter and bike rental
 - Parking Spot finder

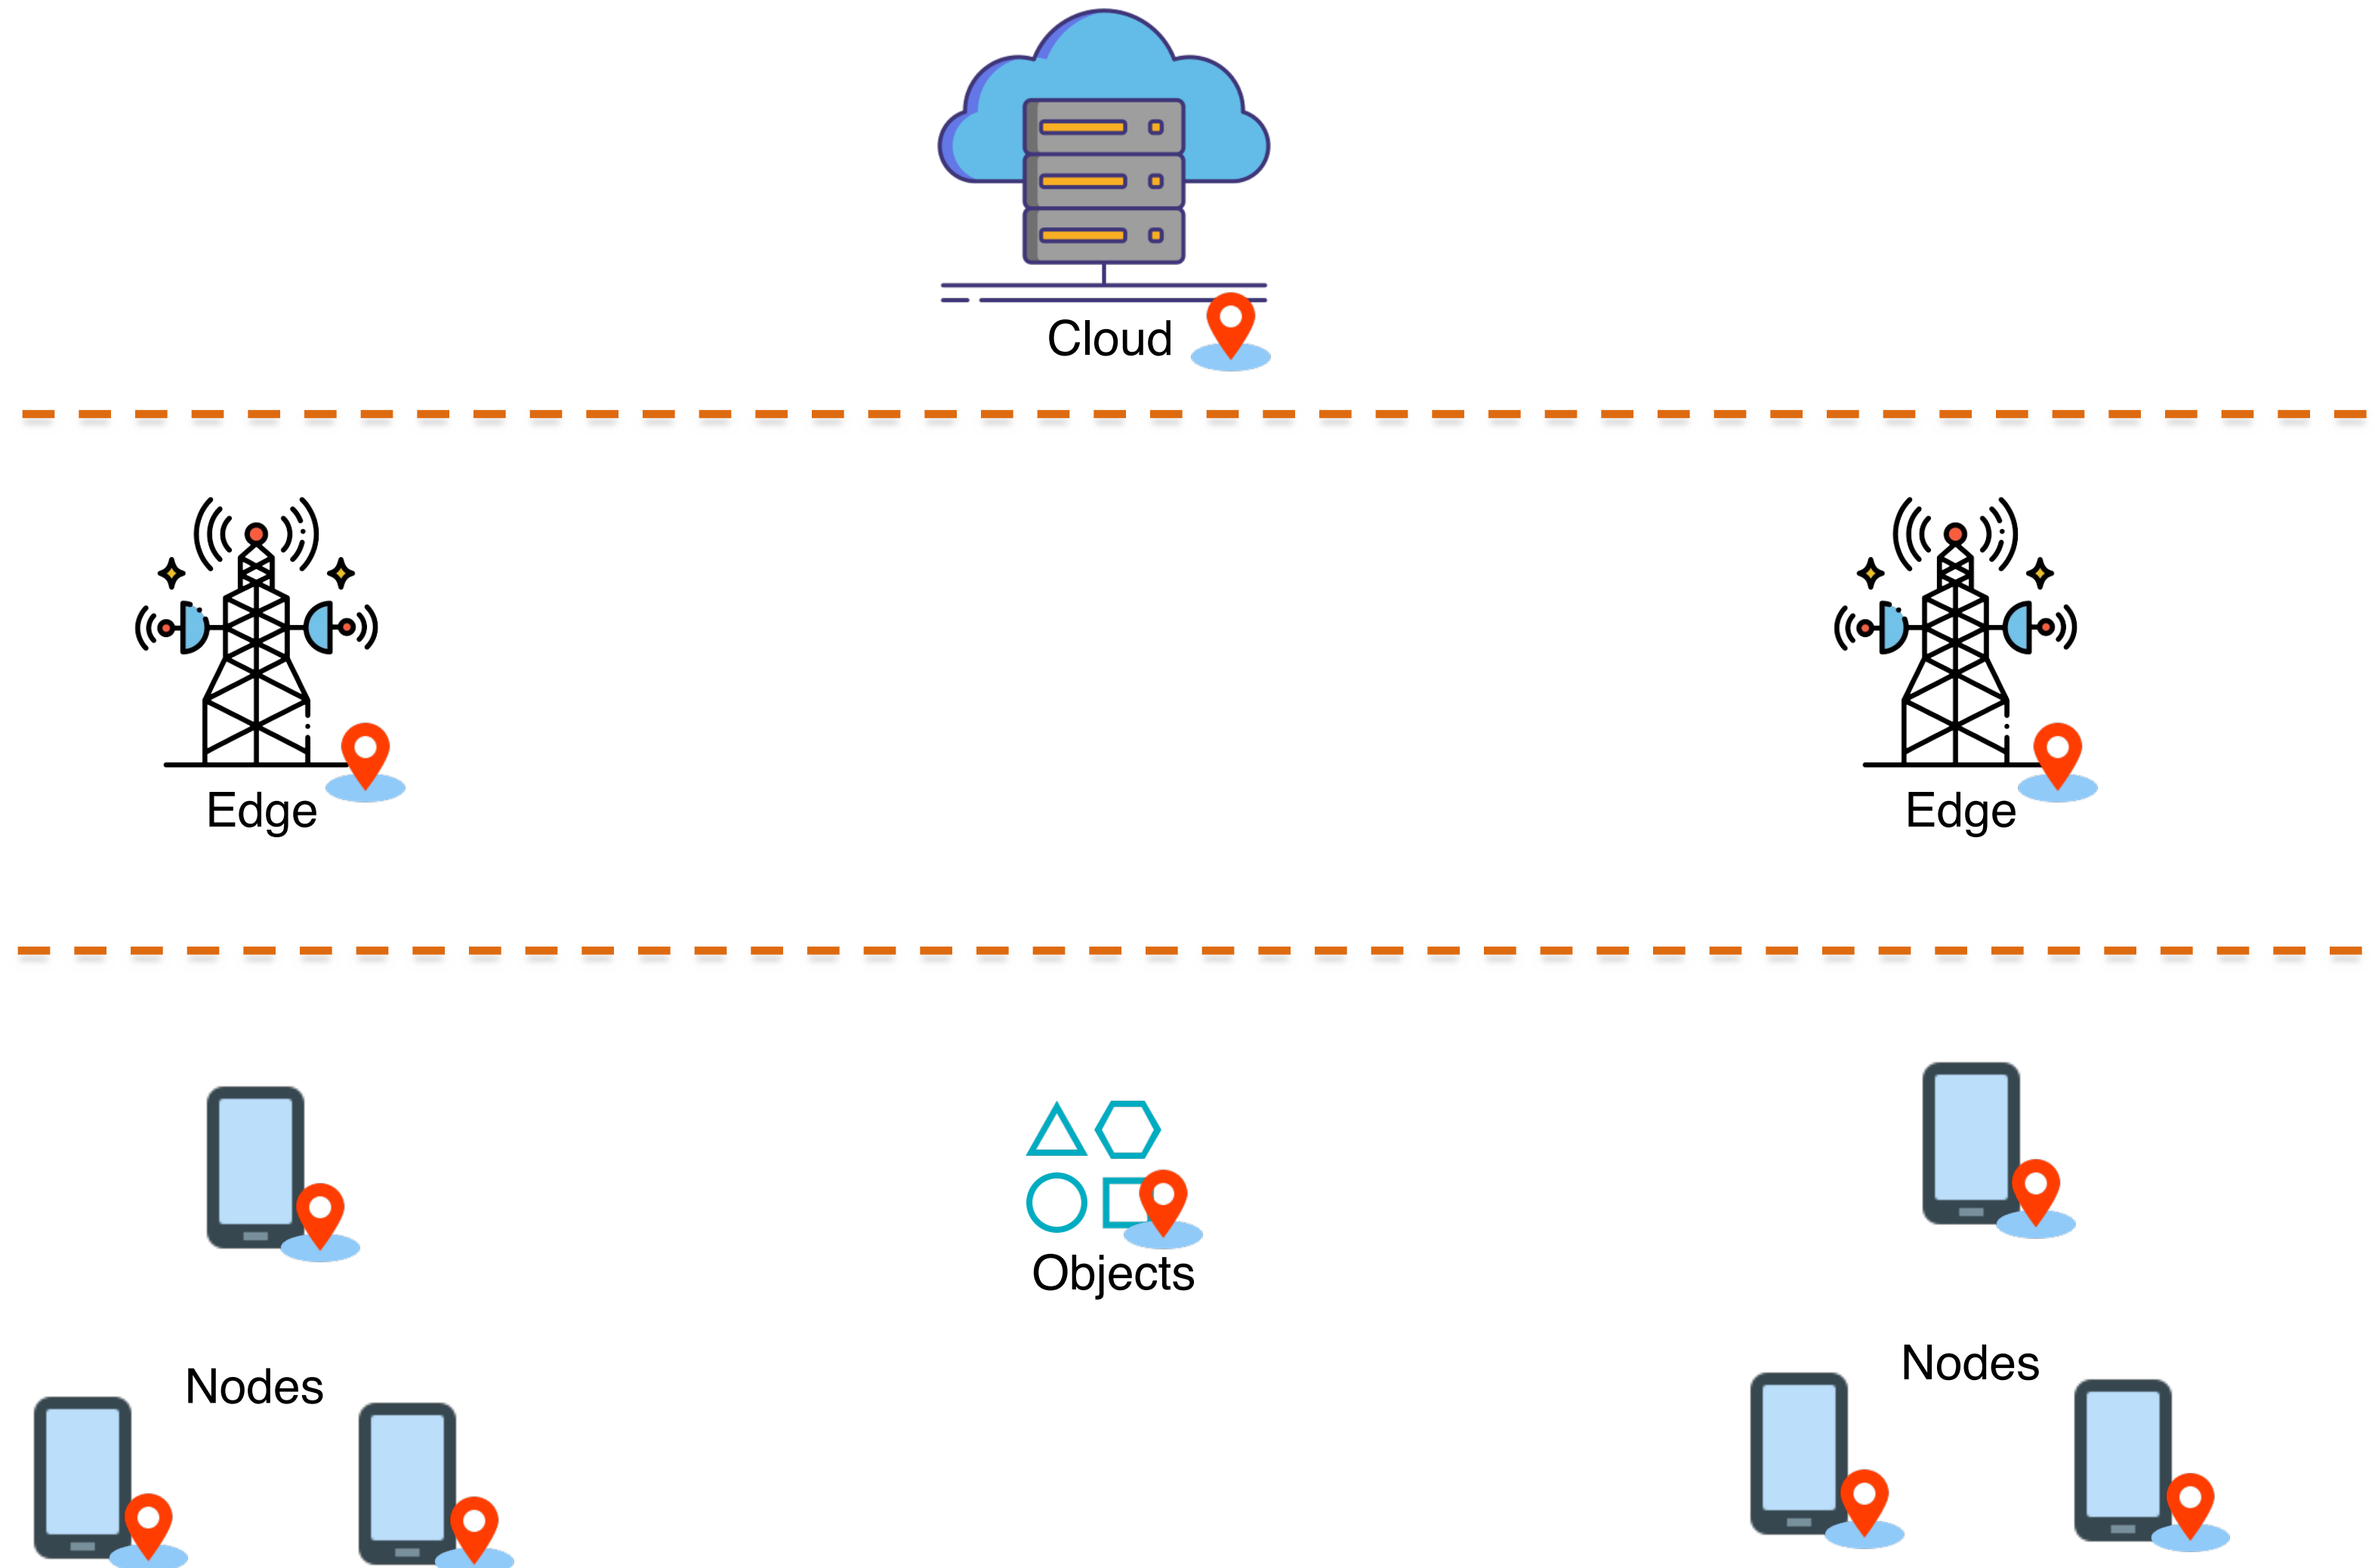
GeoLoc Framework

- Hierarchical Structure

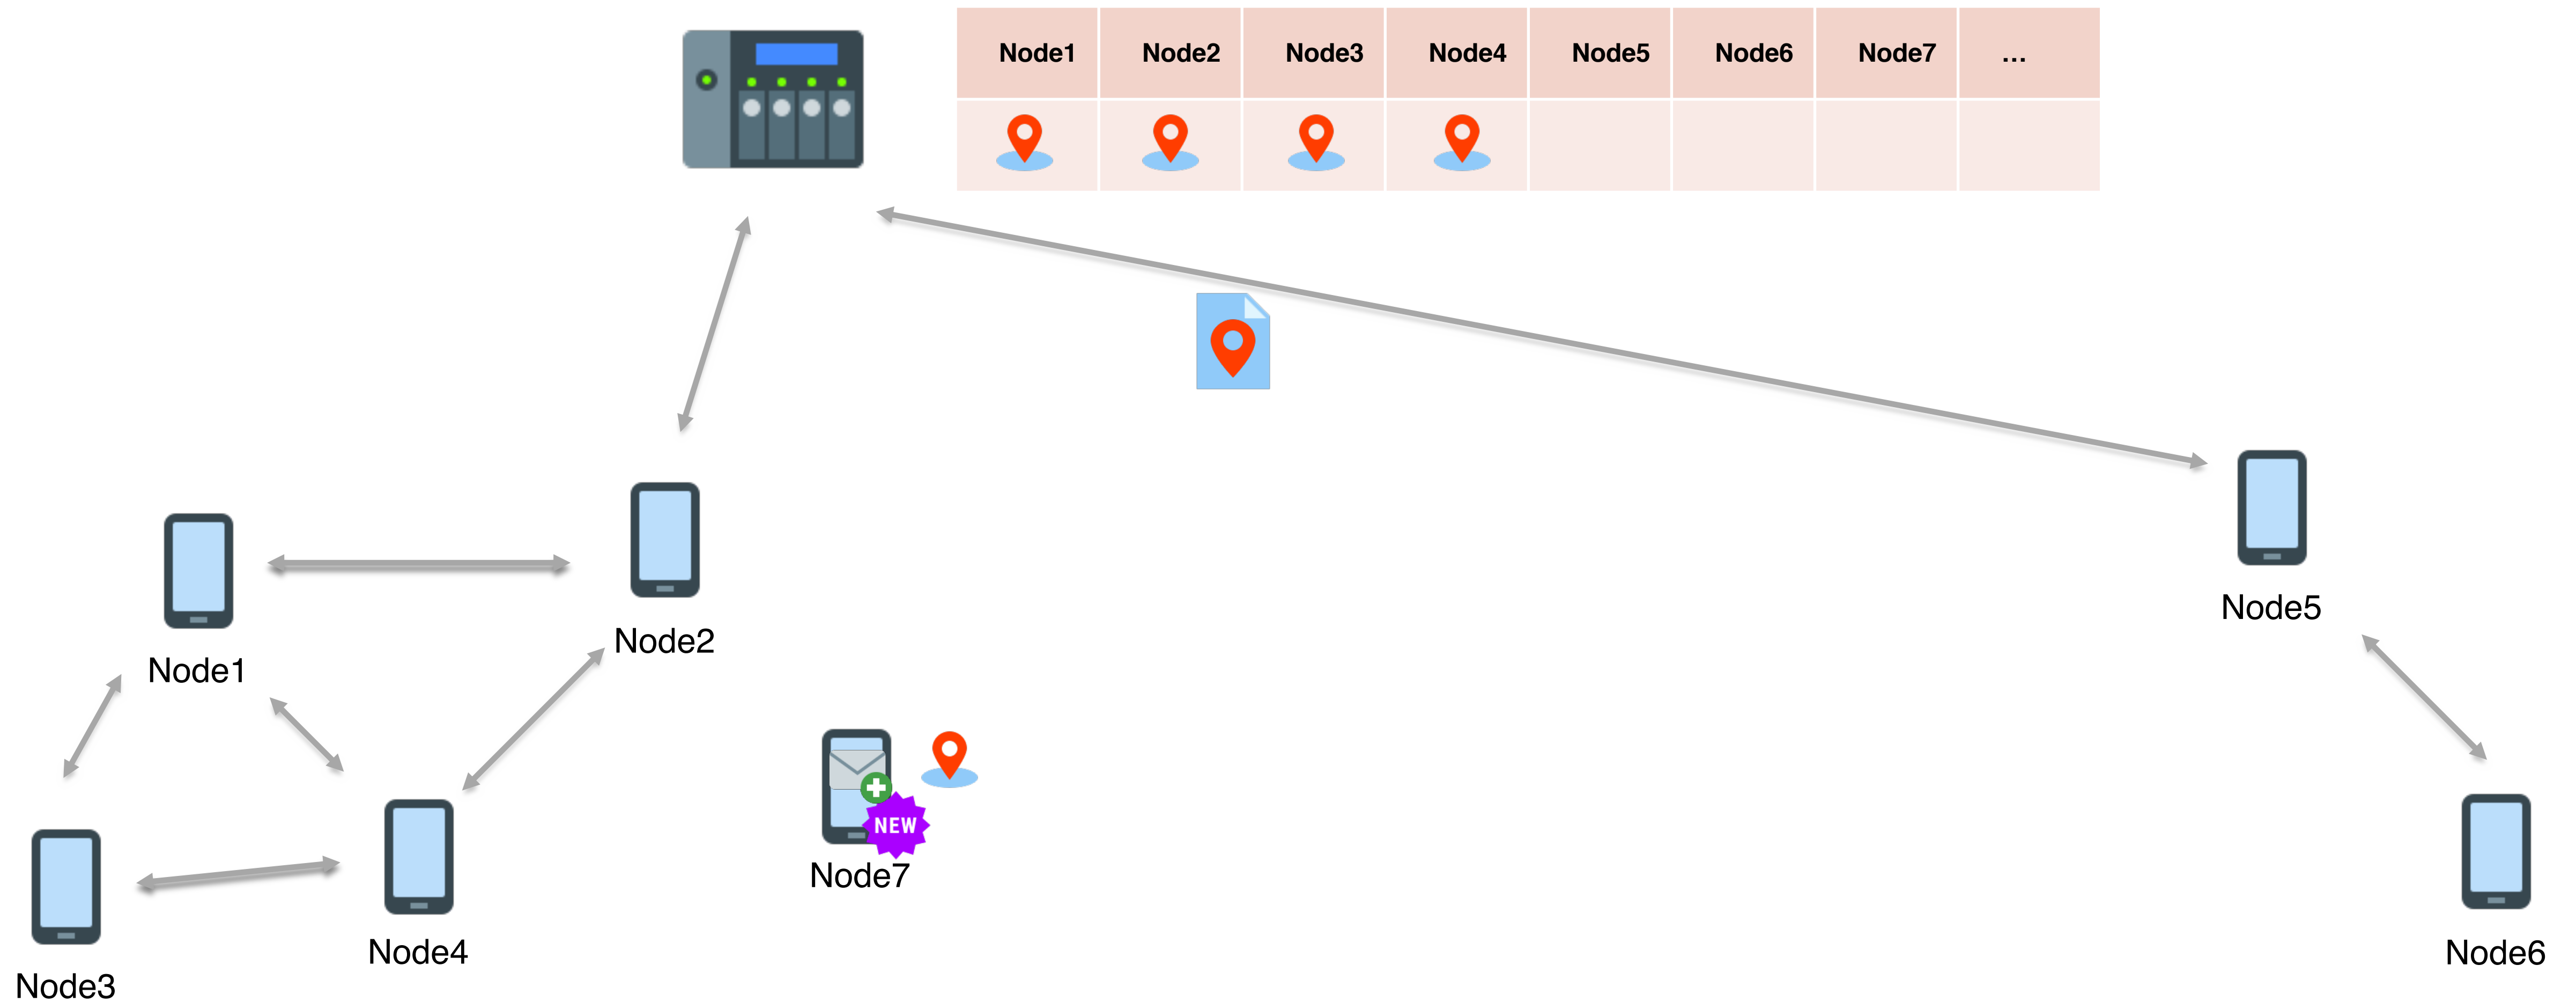


GeoLoc Framework

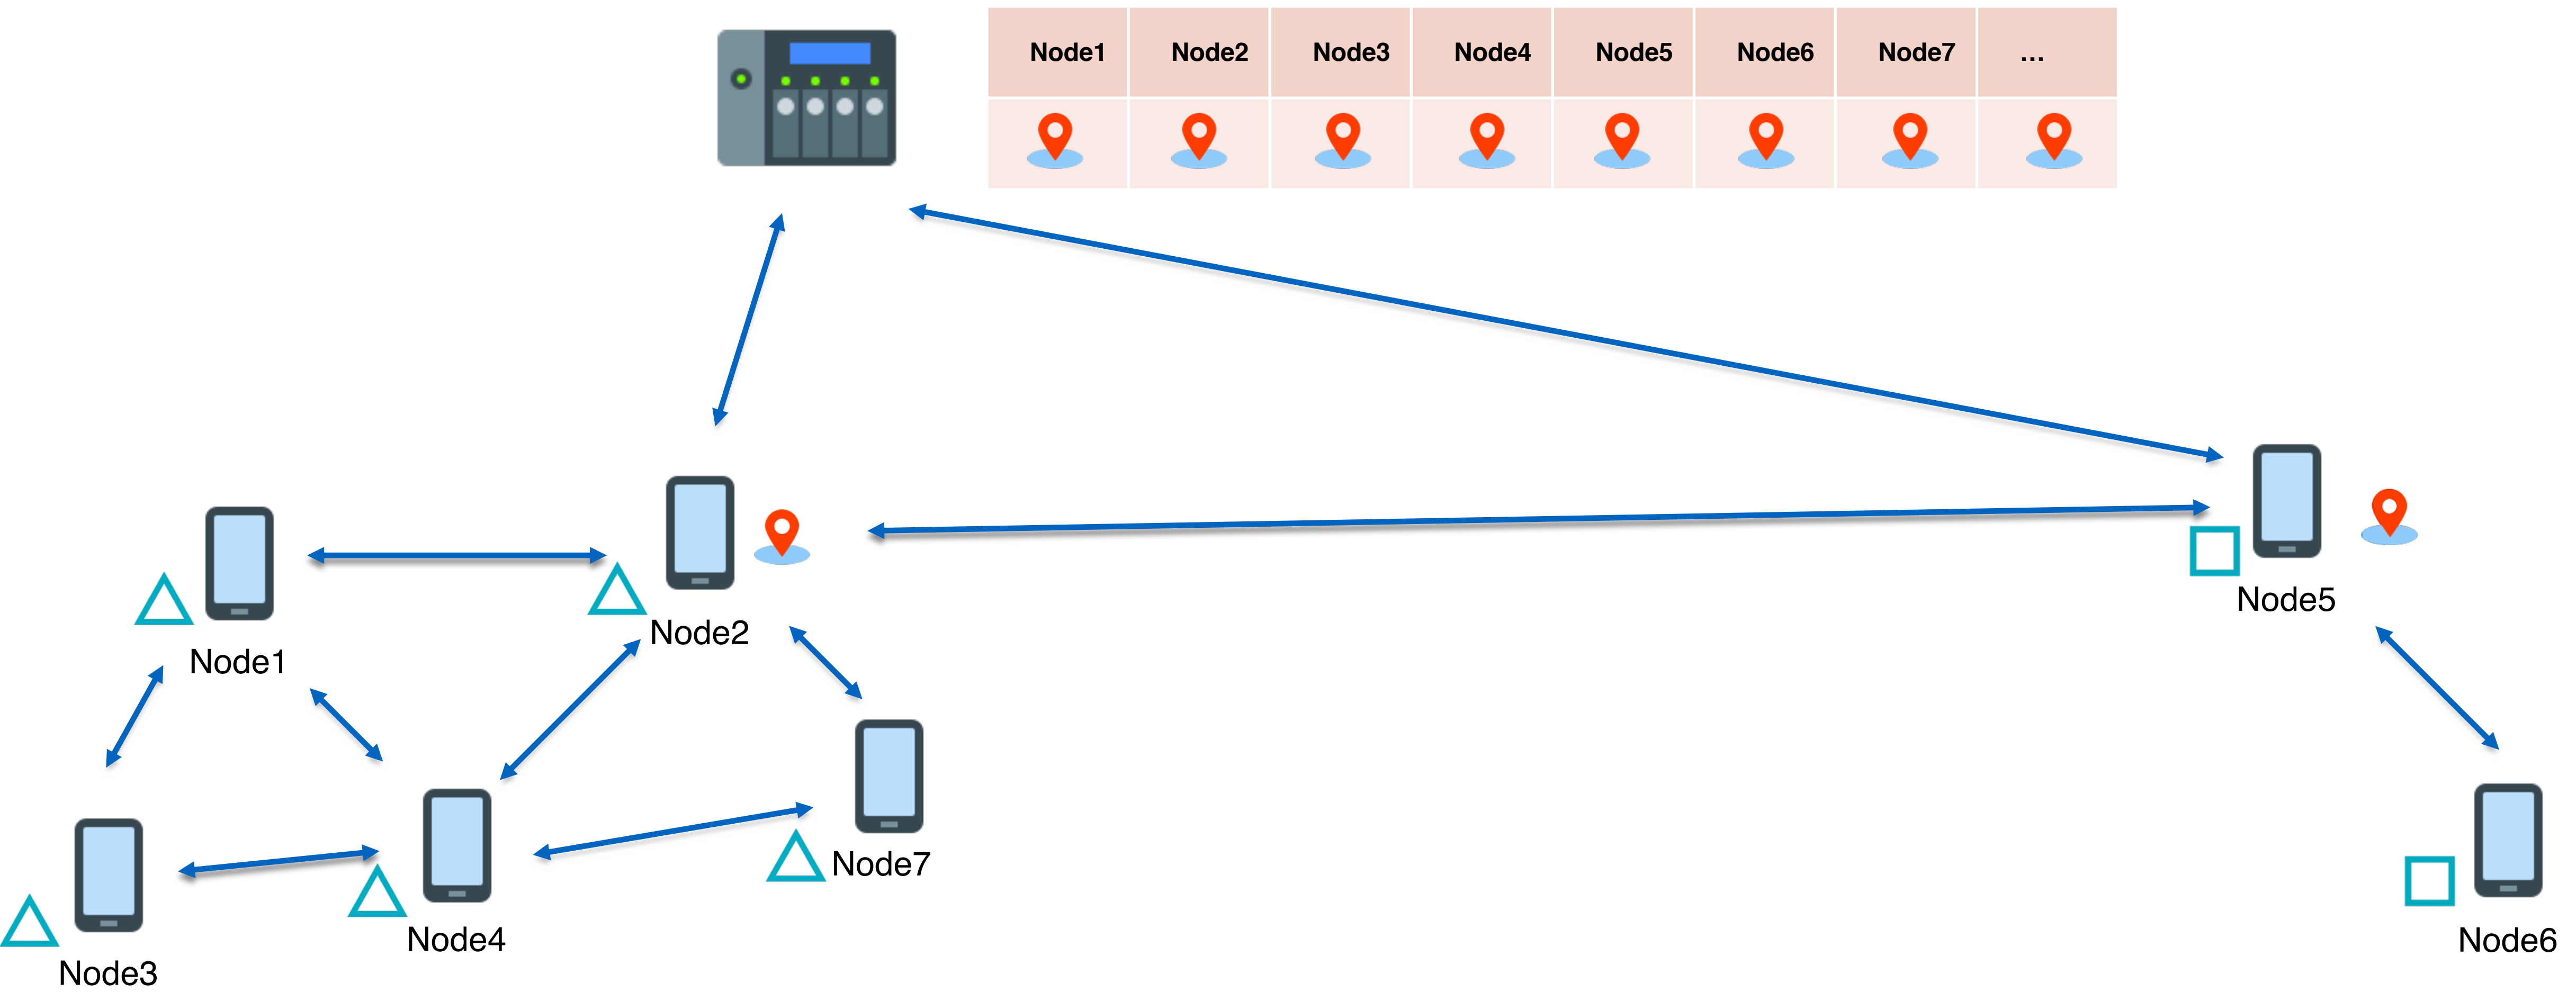
- Hierarchical Structure
- GeoLoc Overlay
- GeoLoc Bully




Geo Location Overlay







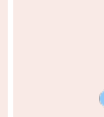
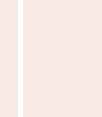


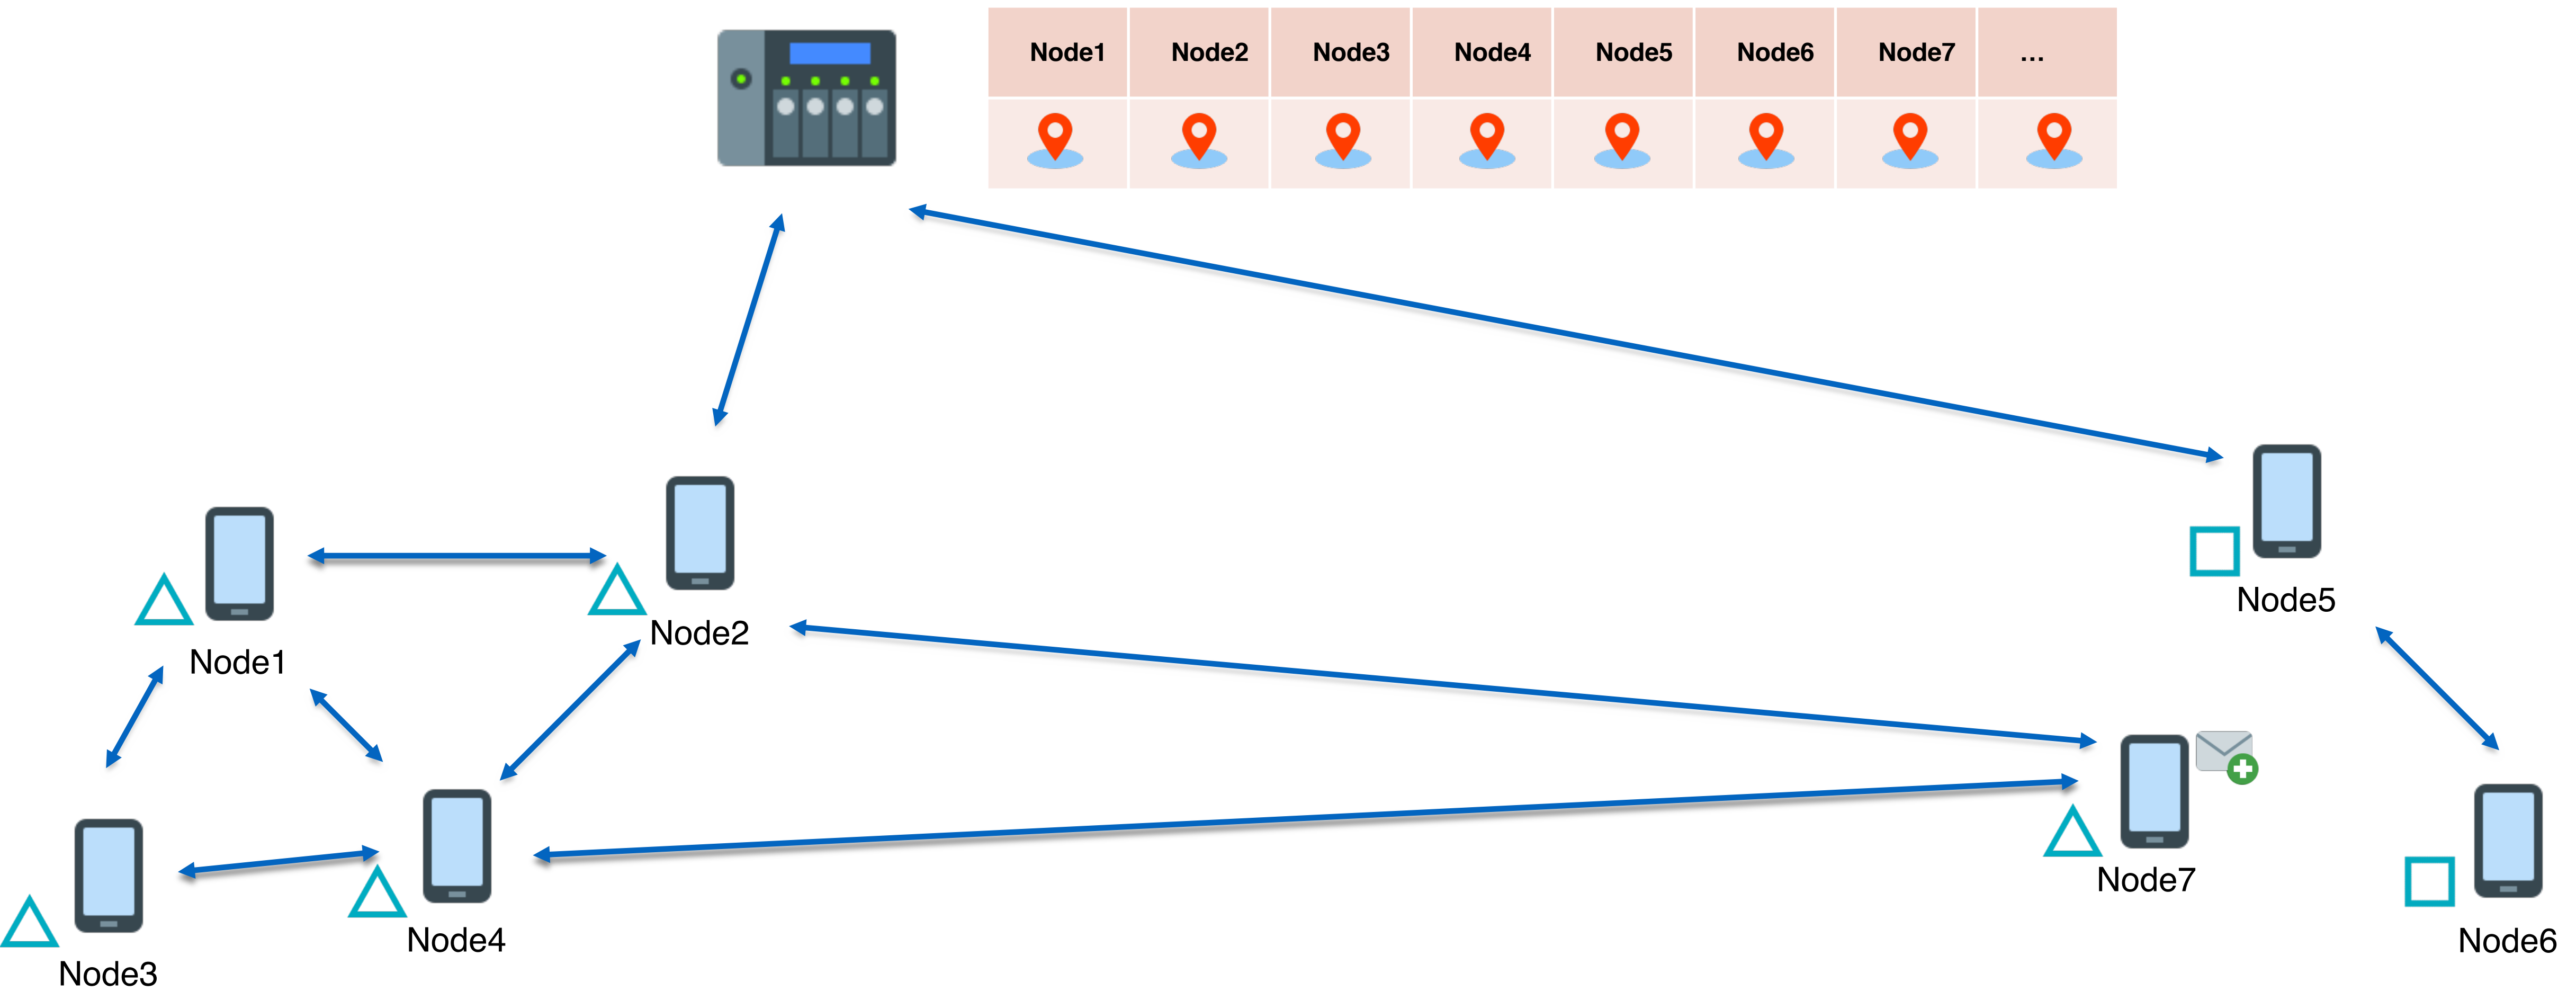
Geo Location Overlay




Geo Location Overlay







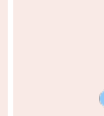
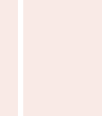



Node1	Node2	Node3	Node4	Node5	Node6	Node7	...
							



Geo Location Overlay



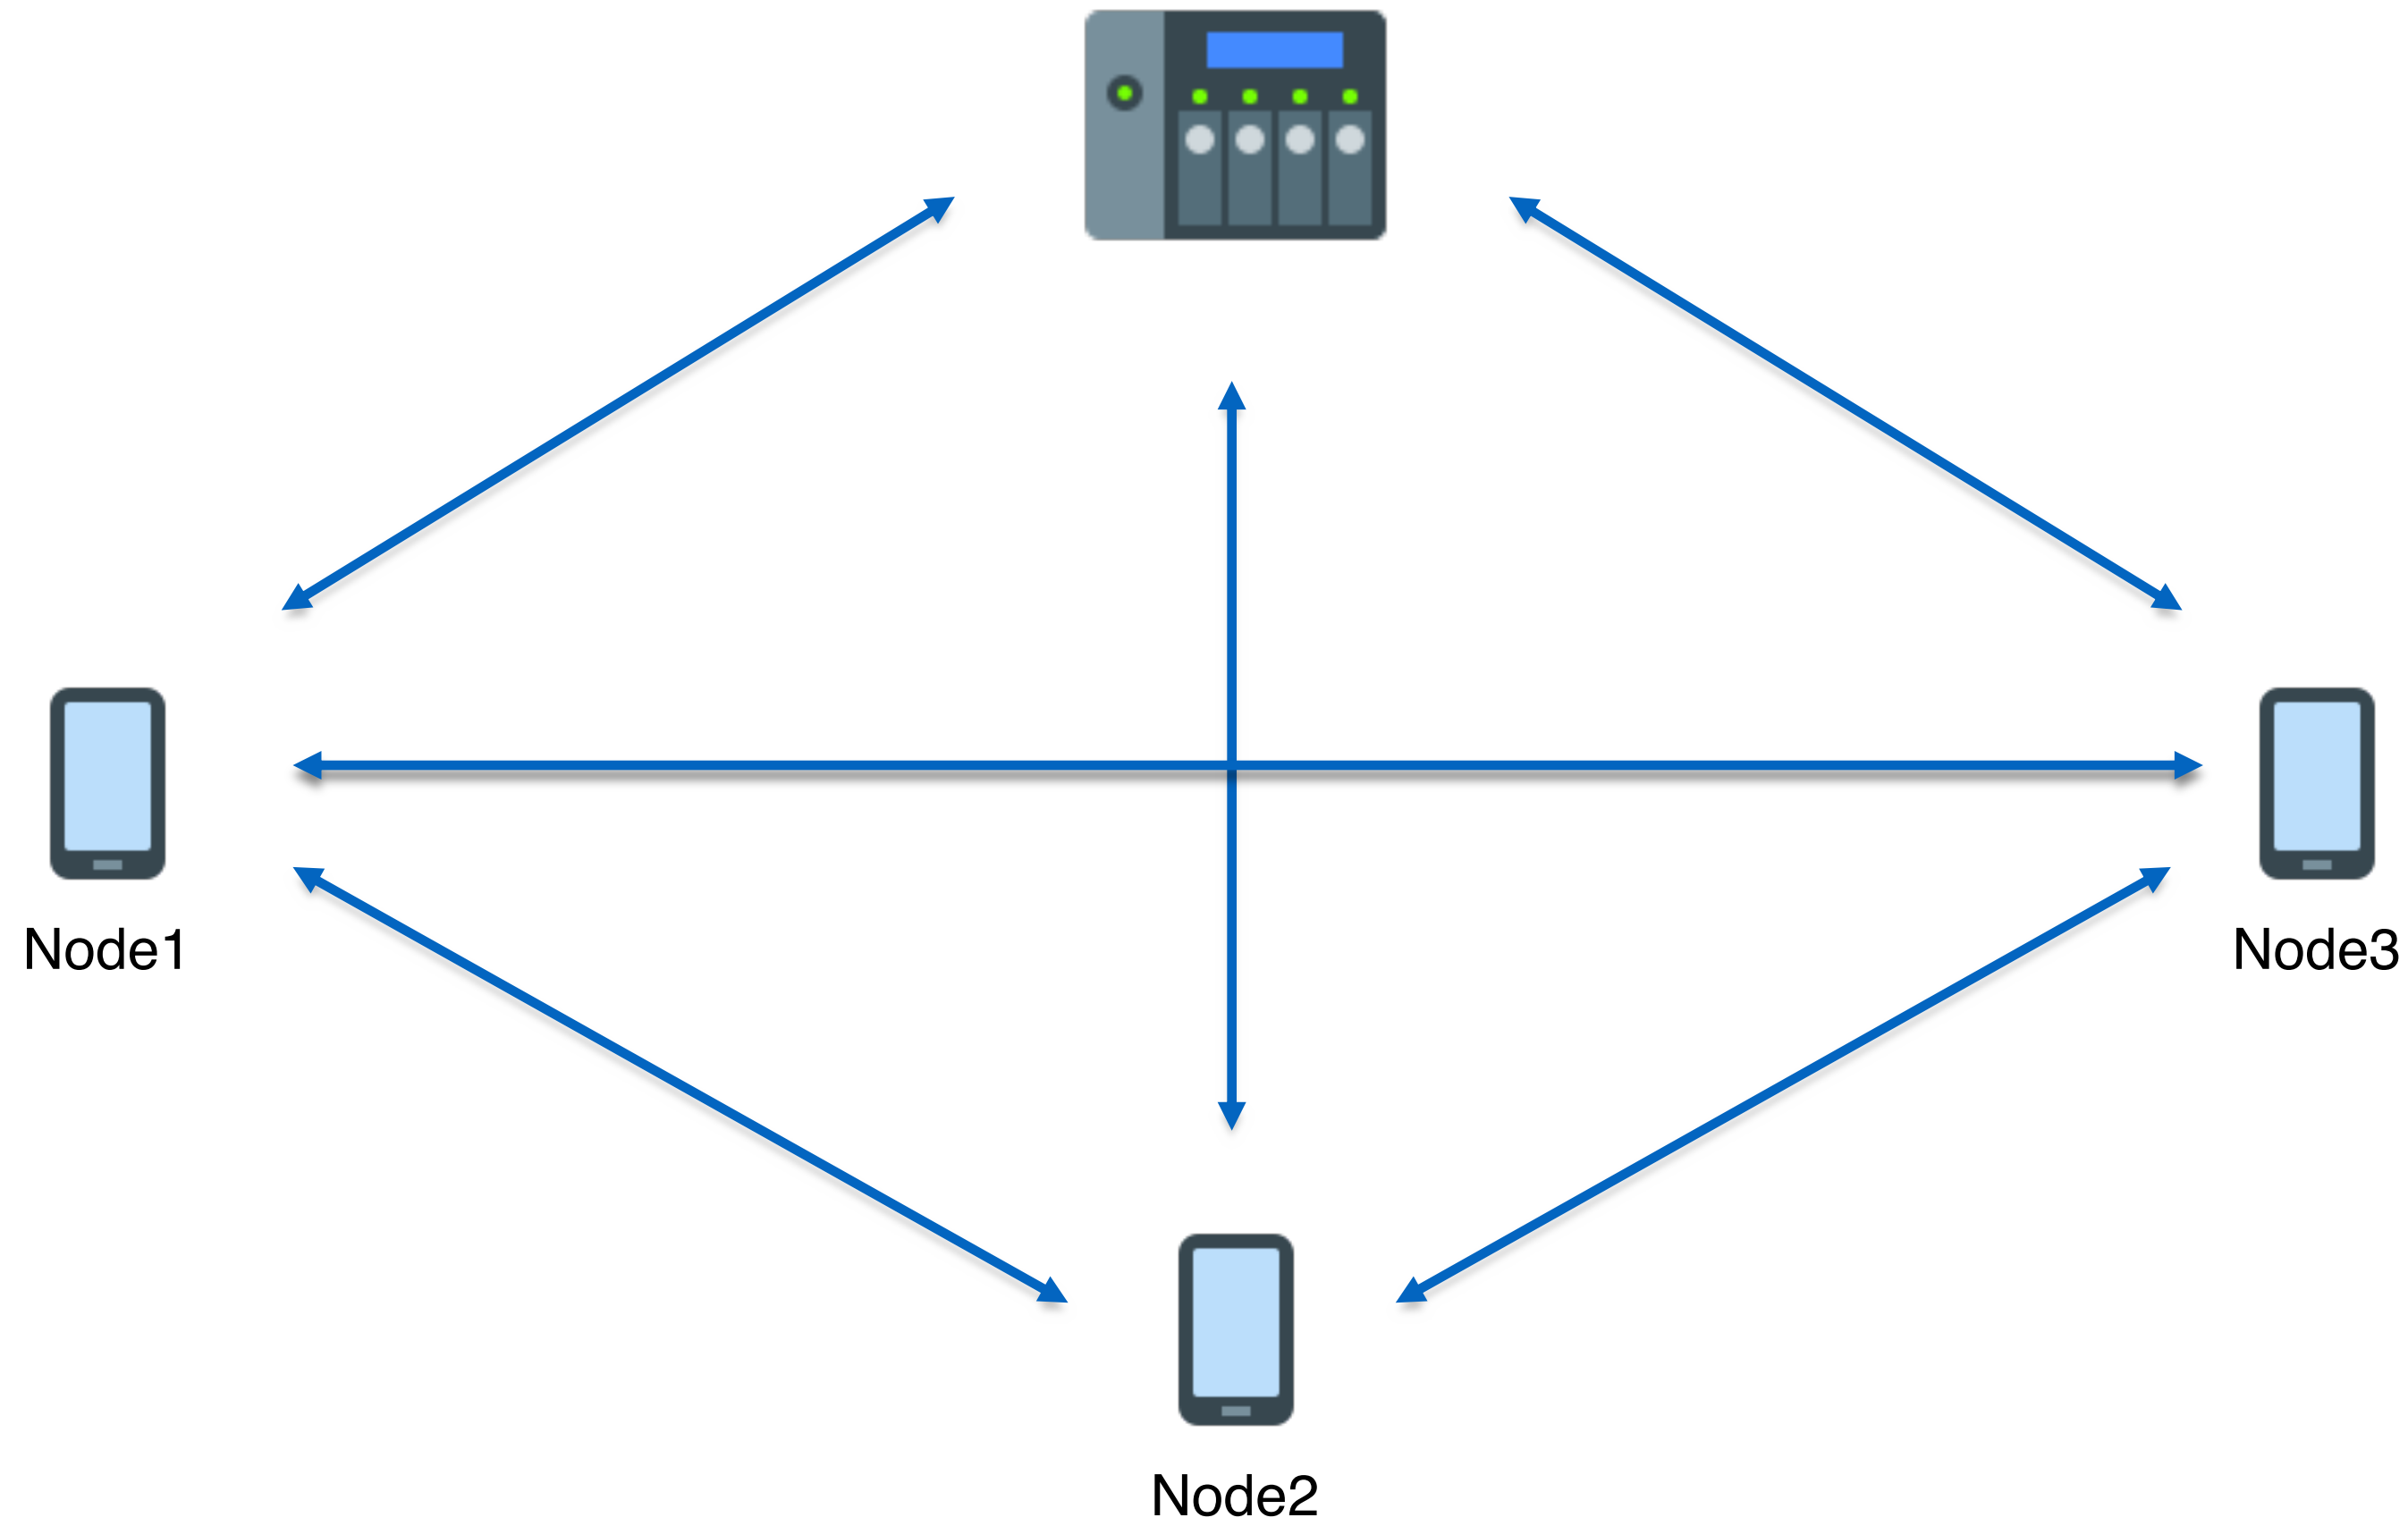
Node1	Node2	Node3	Node4	Node5	Node6	Node7	...
							



Geo Location Overlay




Geo Location Bully

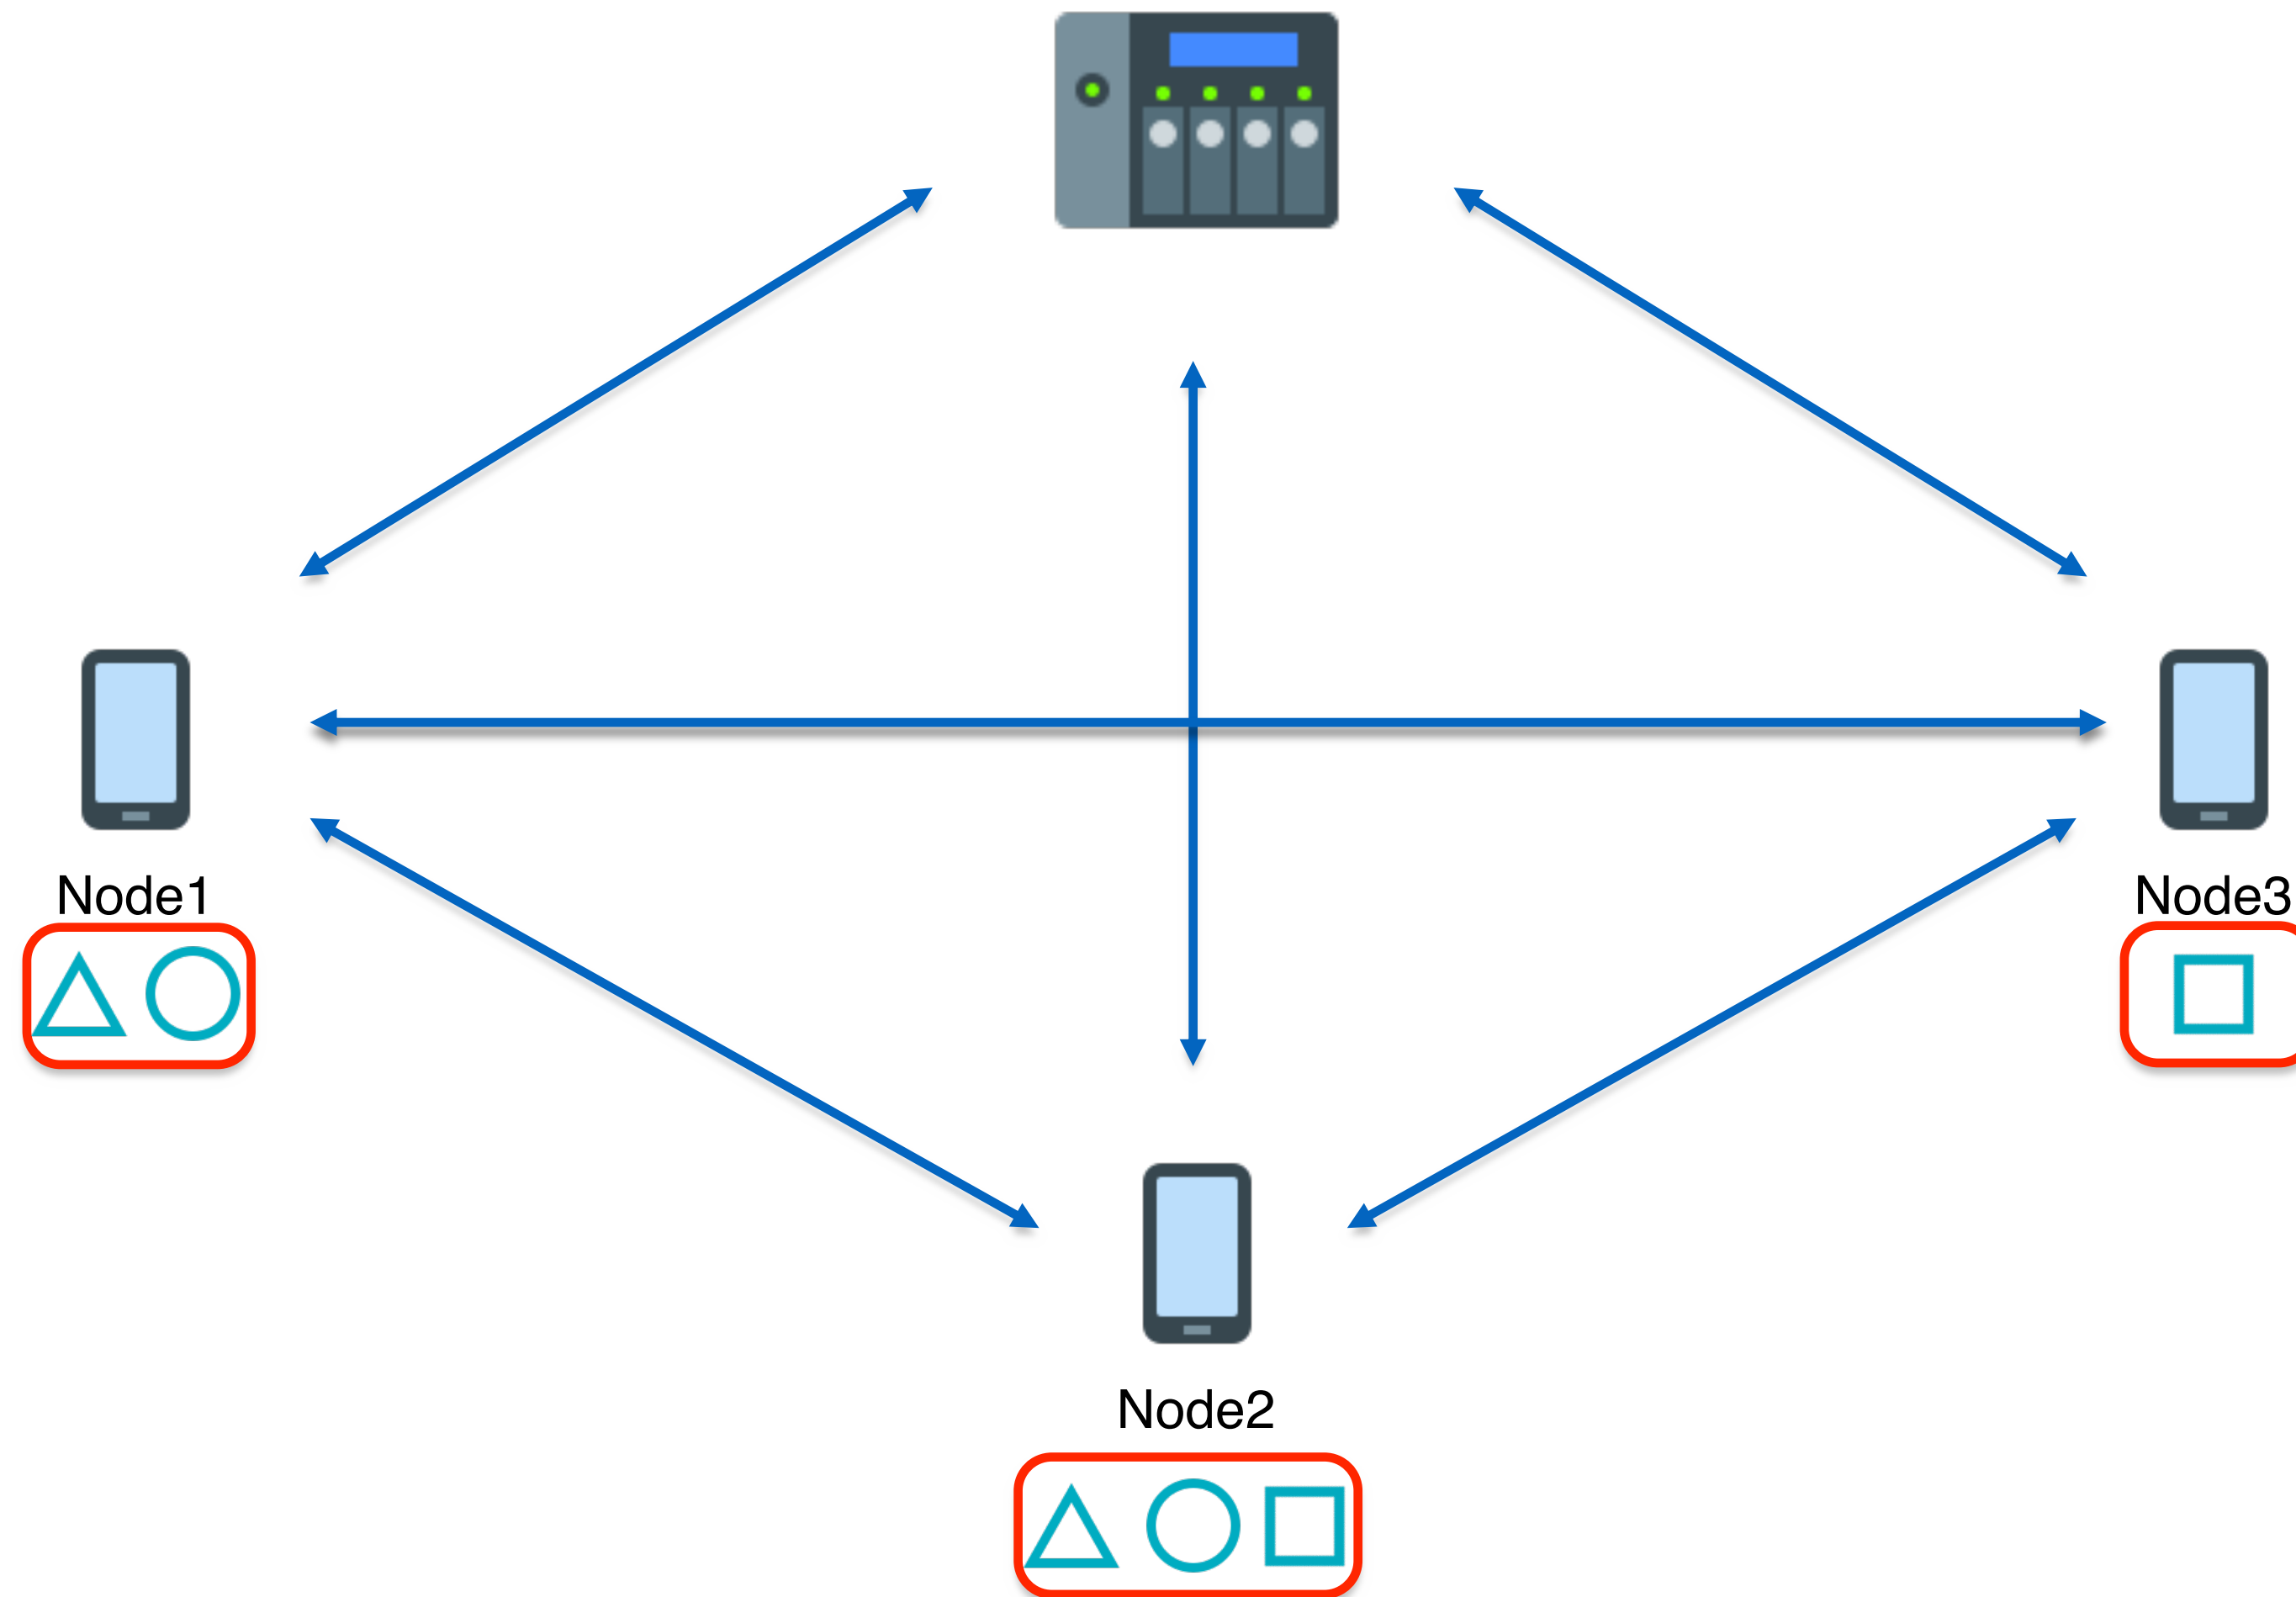





Geo Location Bully



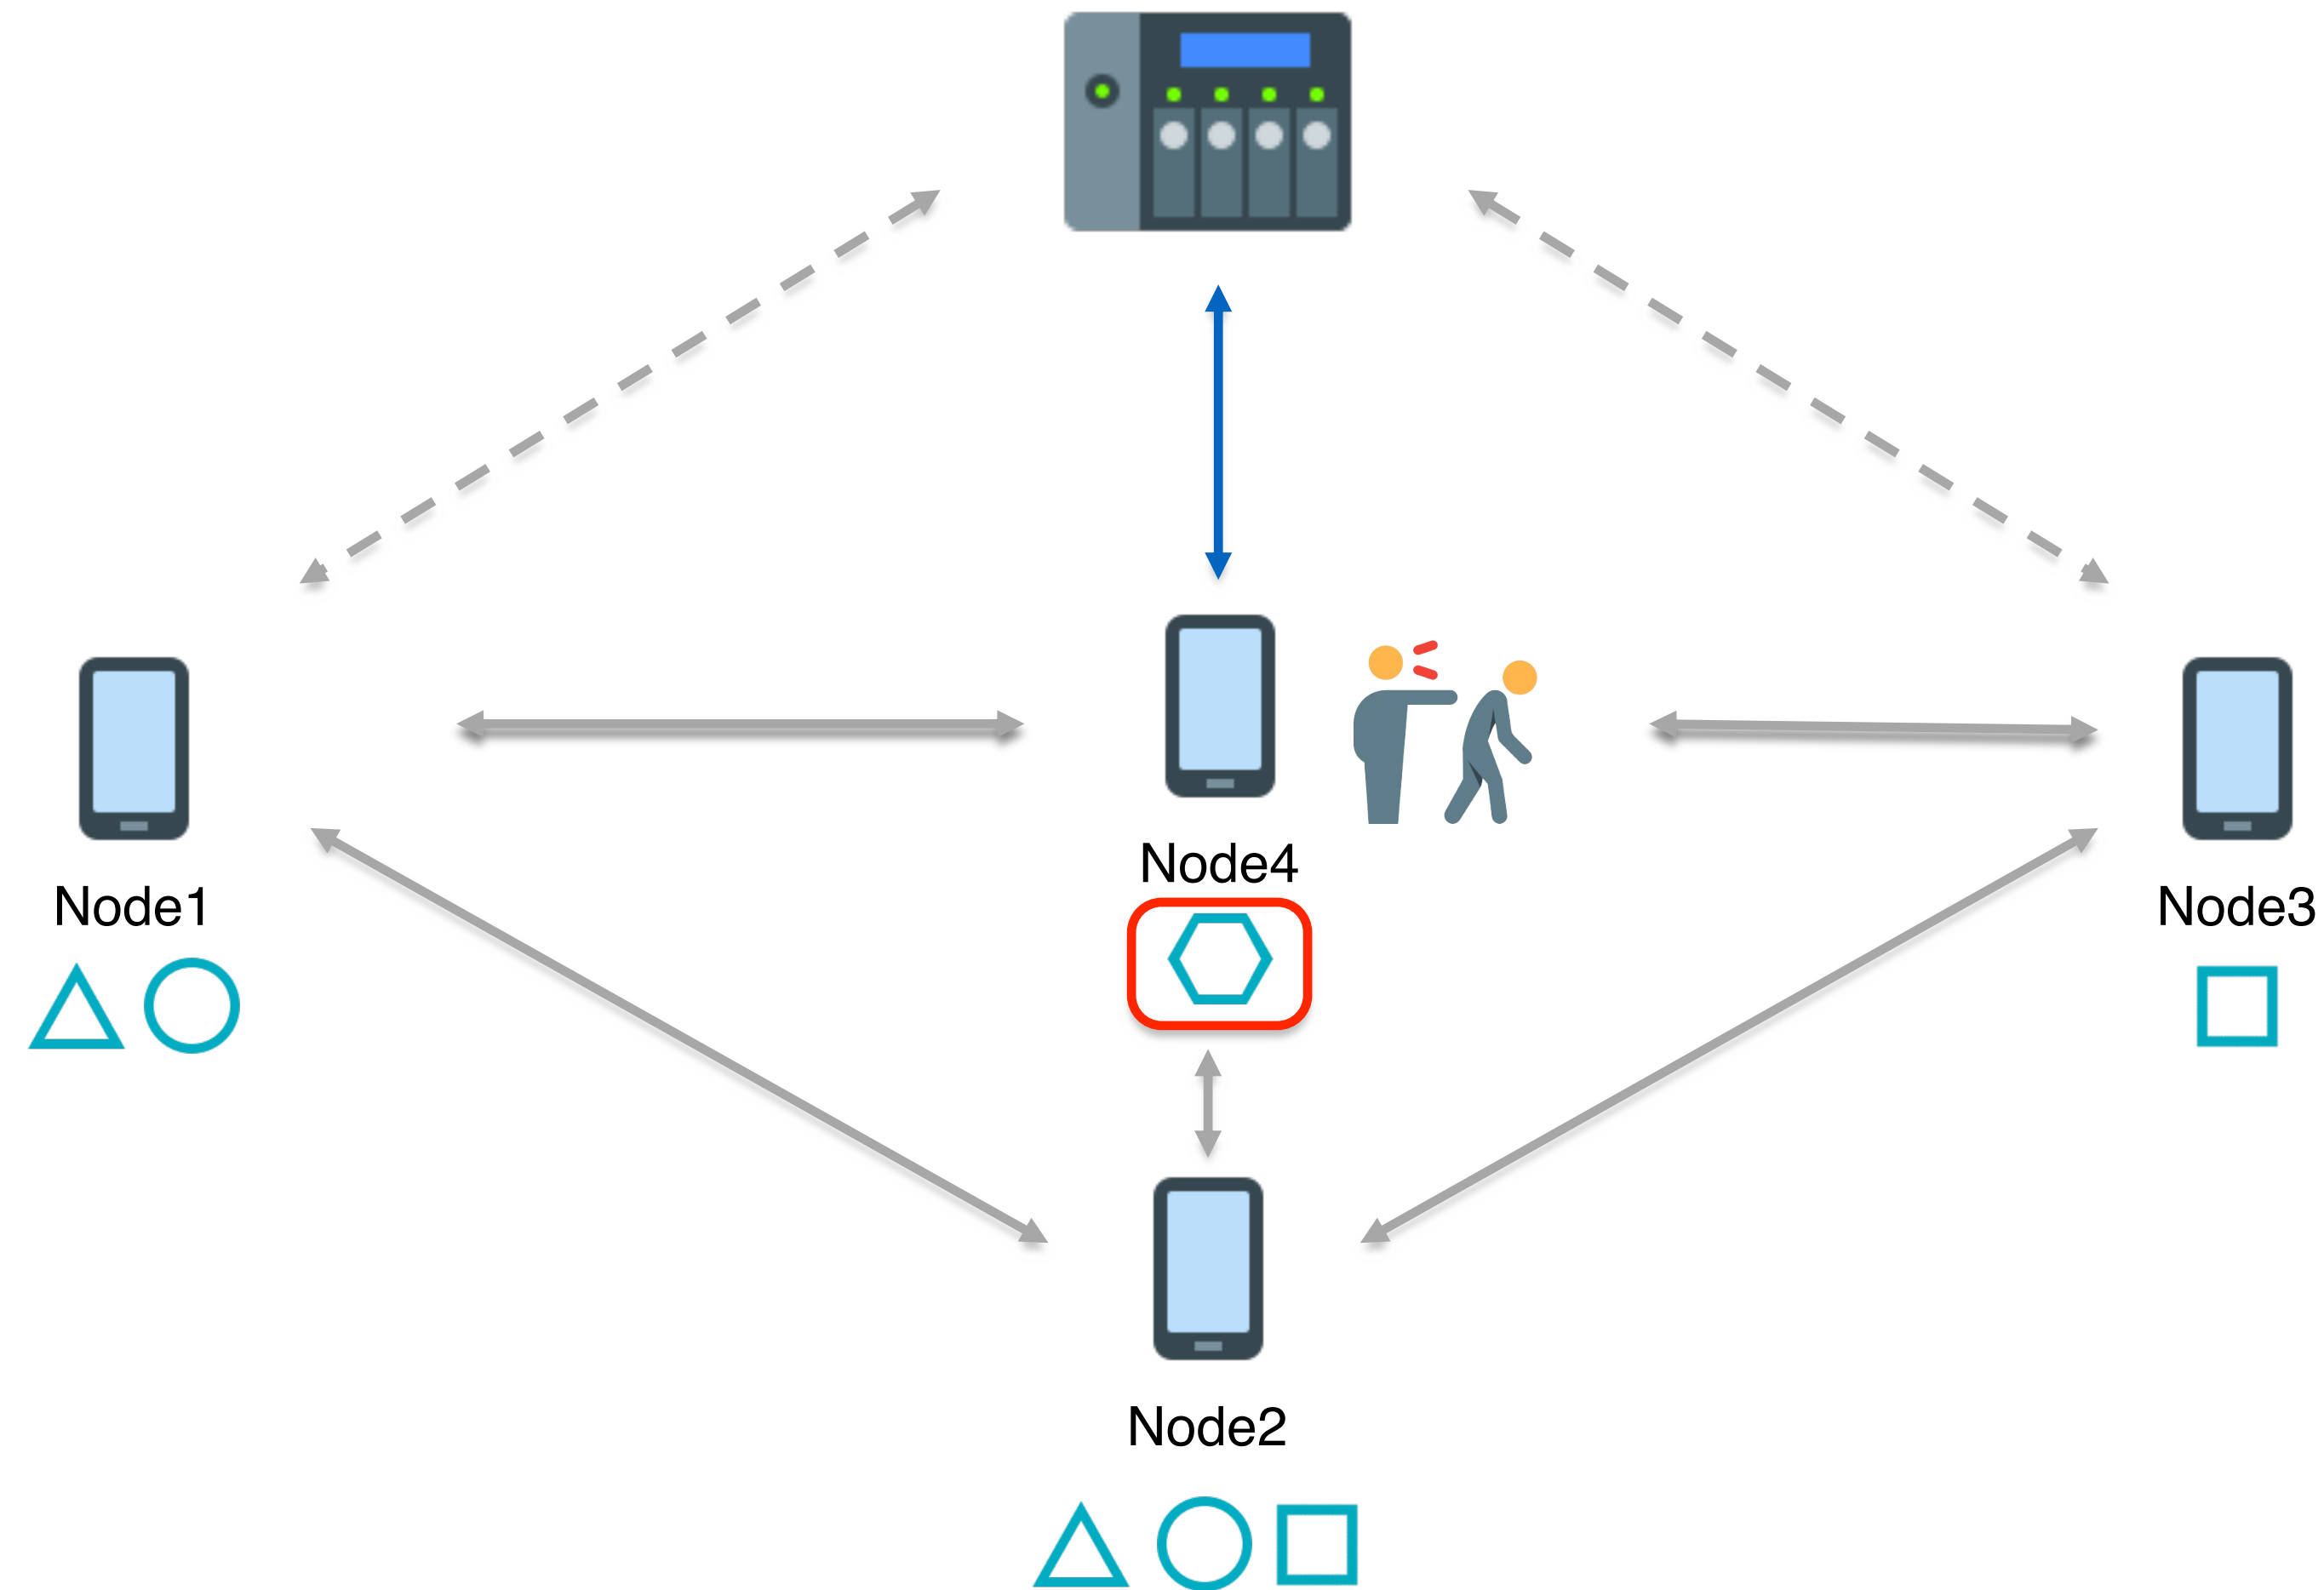
Geo Location Bully



Geo Location Bully



Geo Location Bully





Evaluation

- Naive Foursquare City Guide Application
- Built on top of the Legion framework
- Introduced changes to Legion codebase for implementing our overlay and bully algorithms.
- Evaluated load on server and peers and operation latency



Evaluation

- Two datasets – one for client movement and the other for object placement.
- Clients follow a path with Points of Interest in that route.

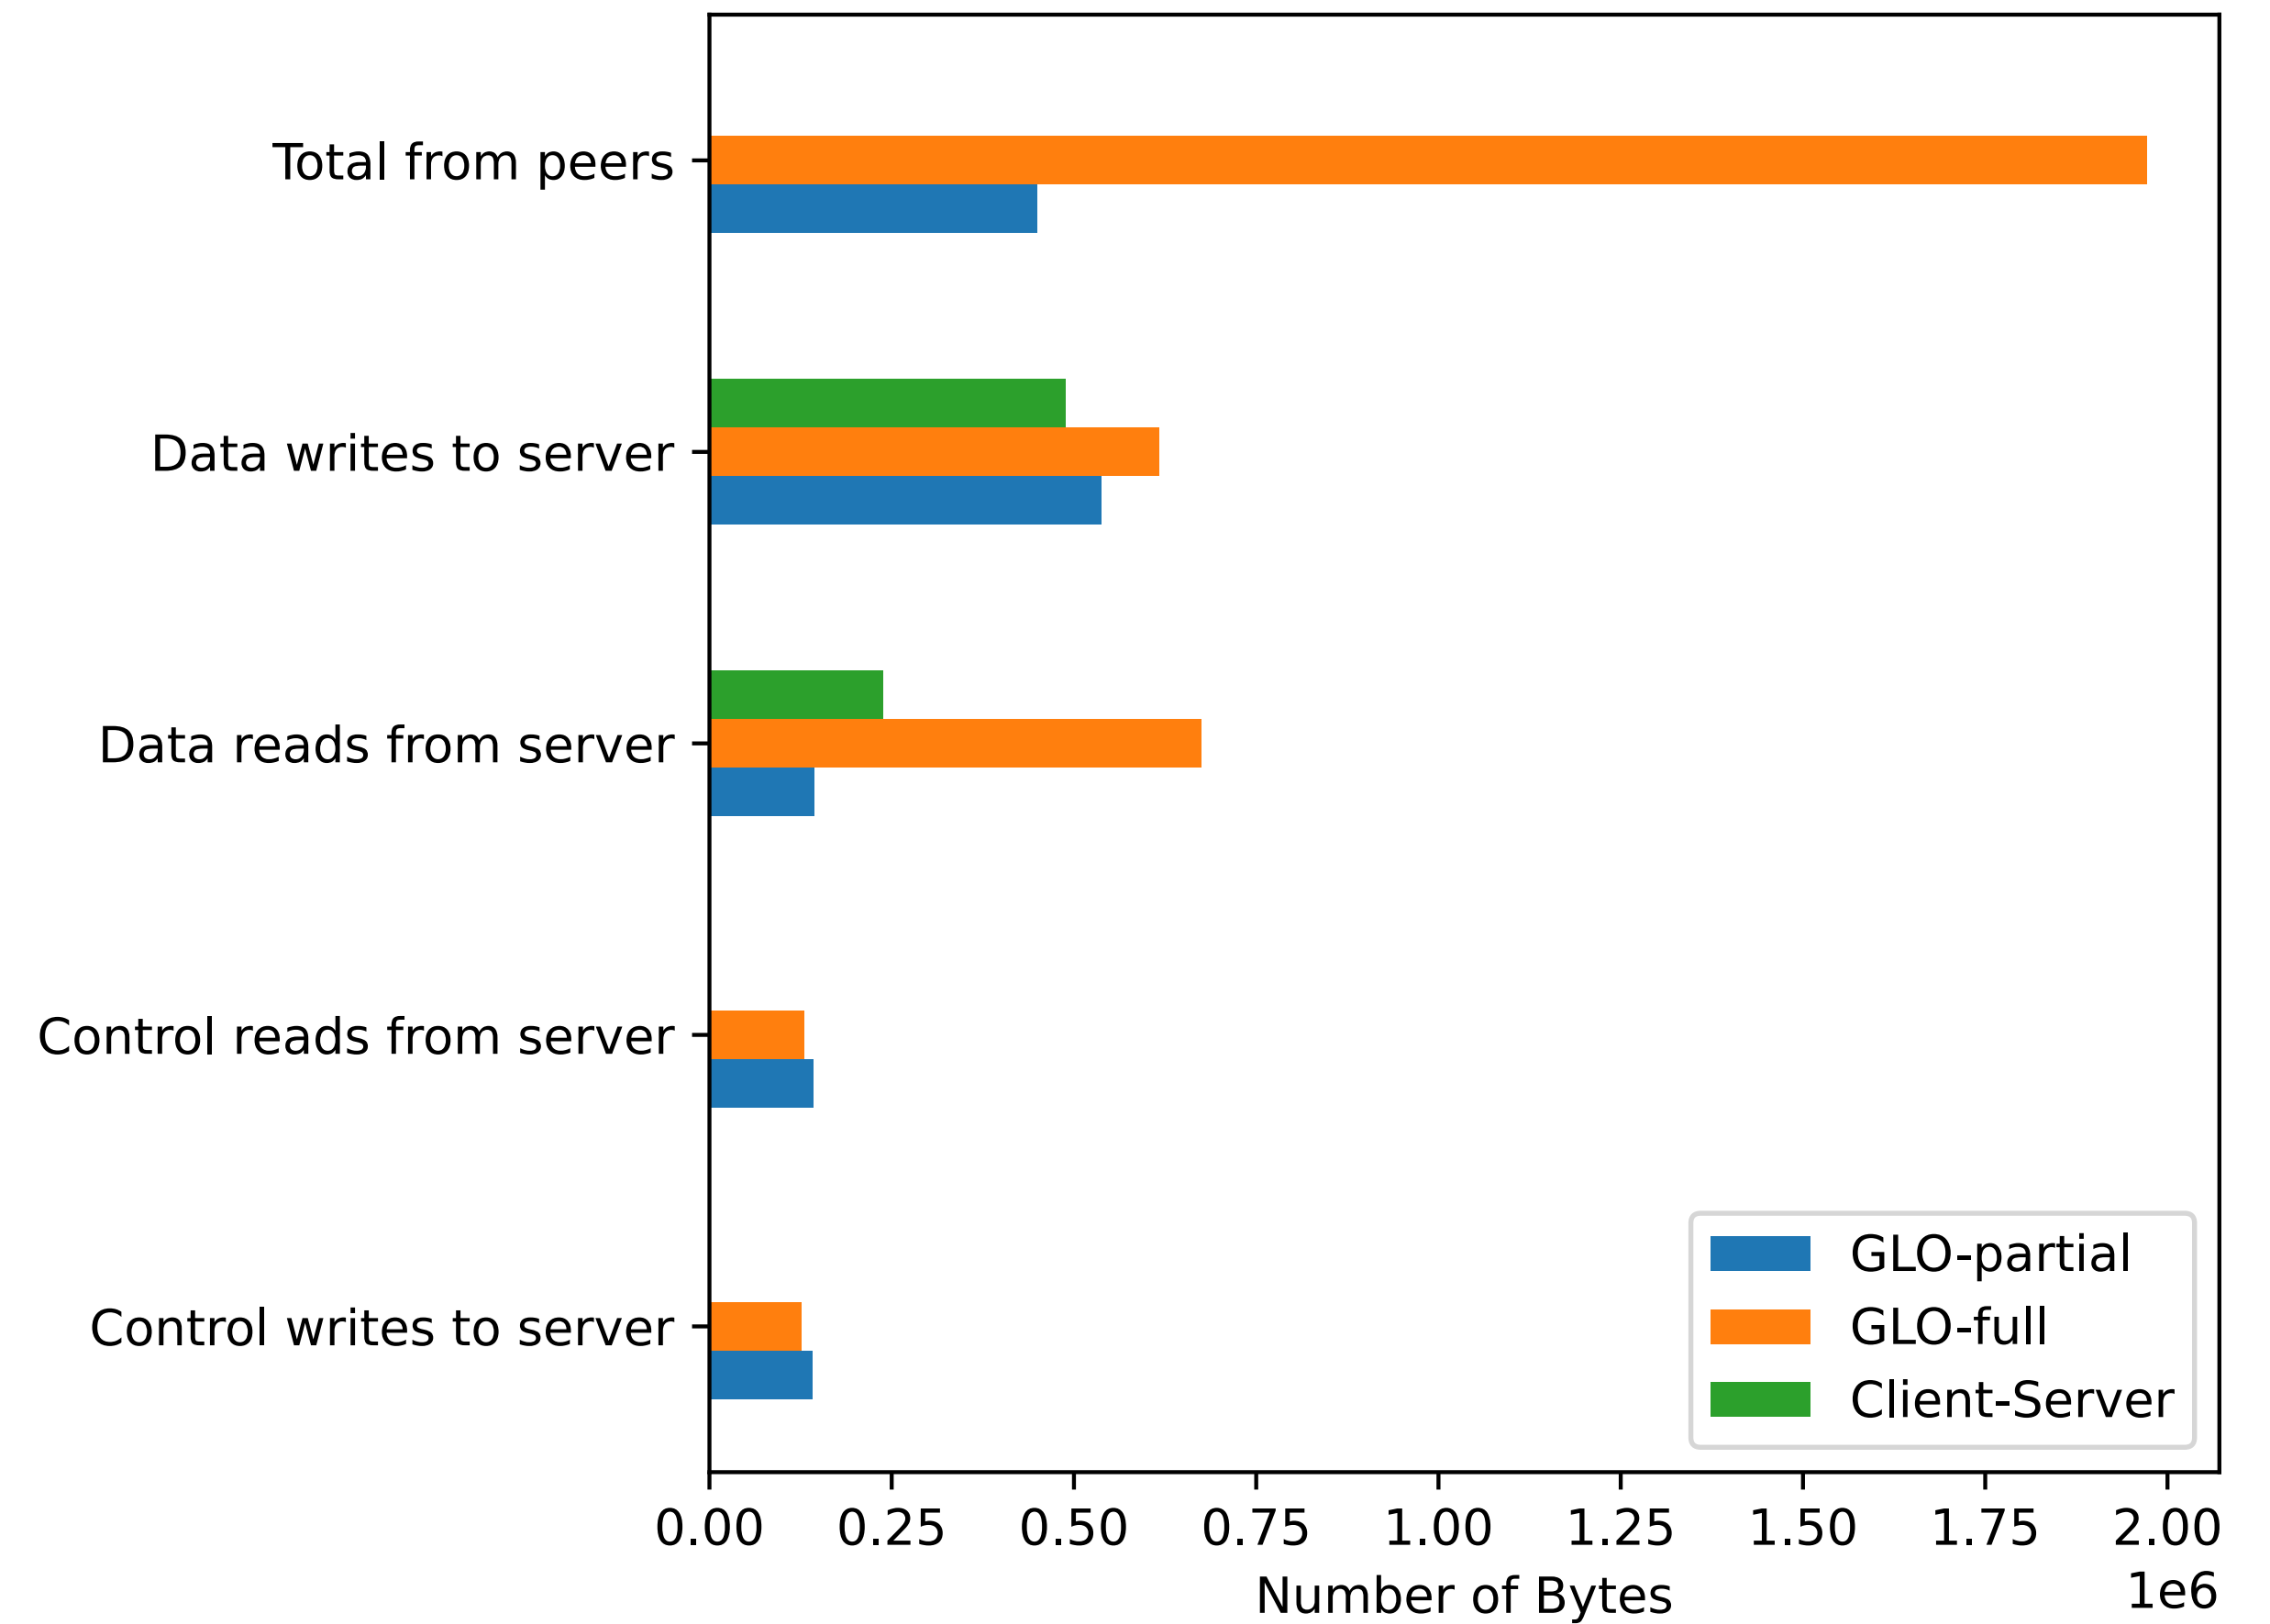



Evaluation

- Foursquare allows users to discover information on nearby locations, permitting them to perform a "Check-In" or "Write a review" on that POI.
- Three different overlay models
 - GeoLoc Overlay in partial replication mode (GLO-partial)
 - GeoLoc Overlay in full replication mode (GLO-full)
 - Client-server in partial replication mode (C-S)


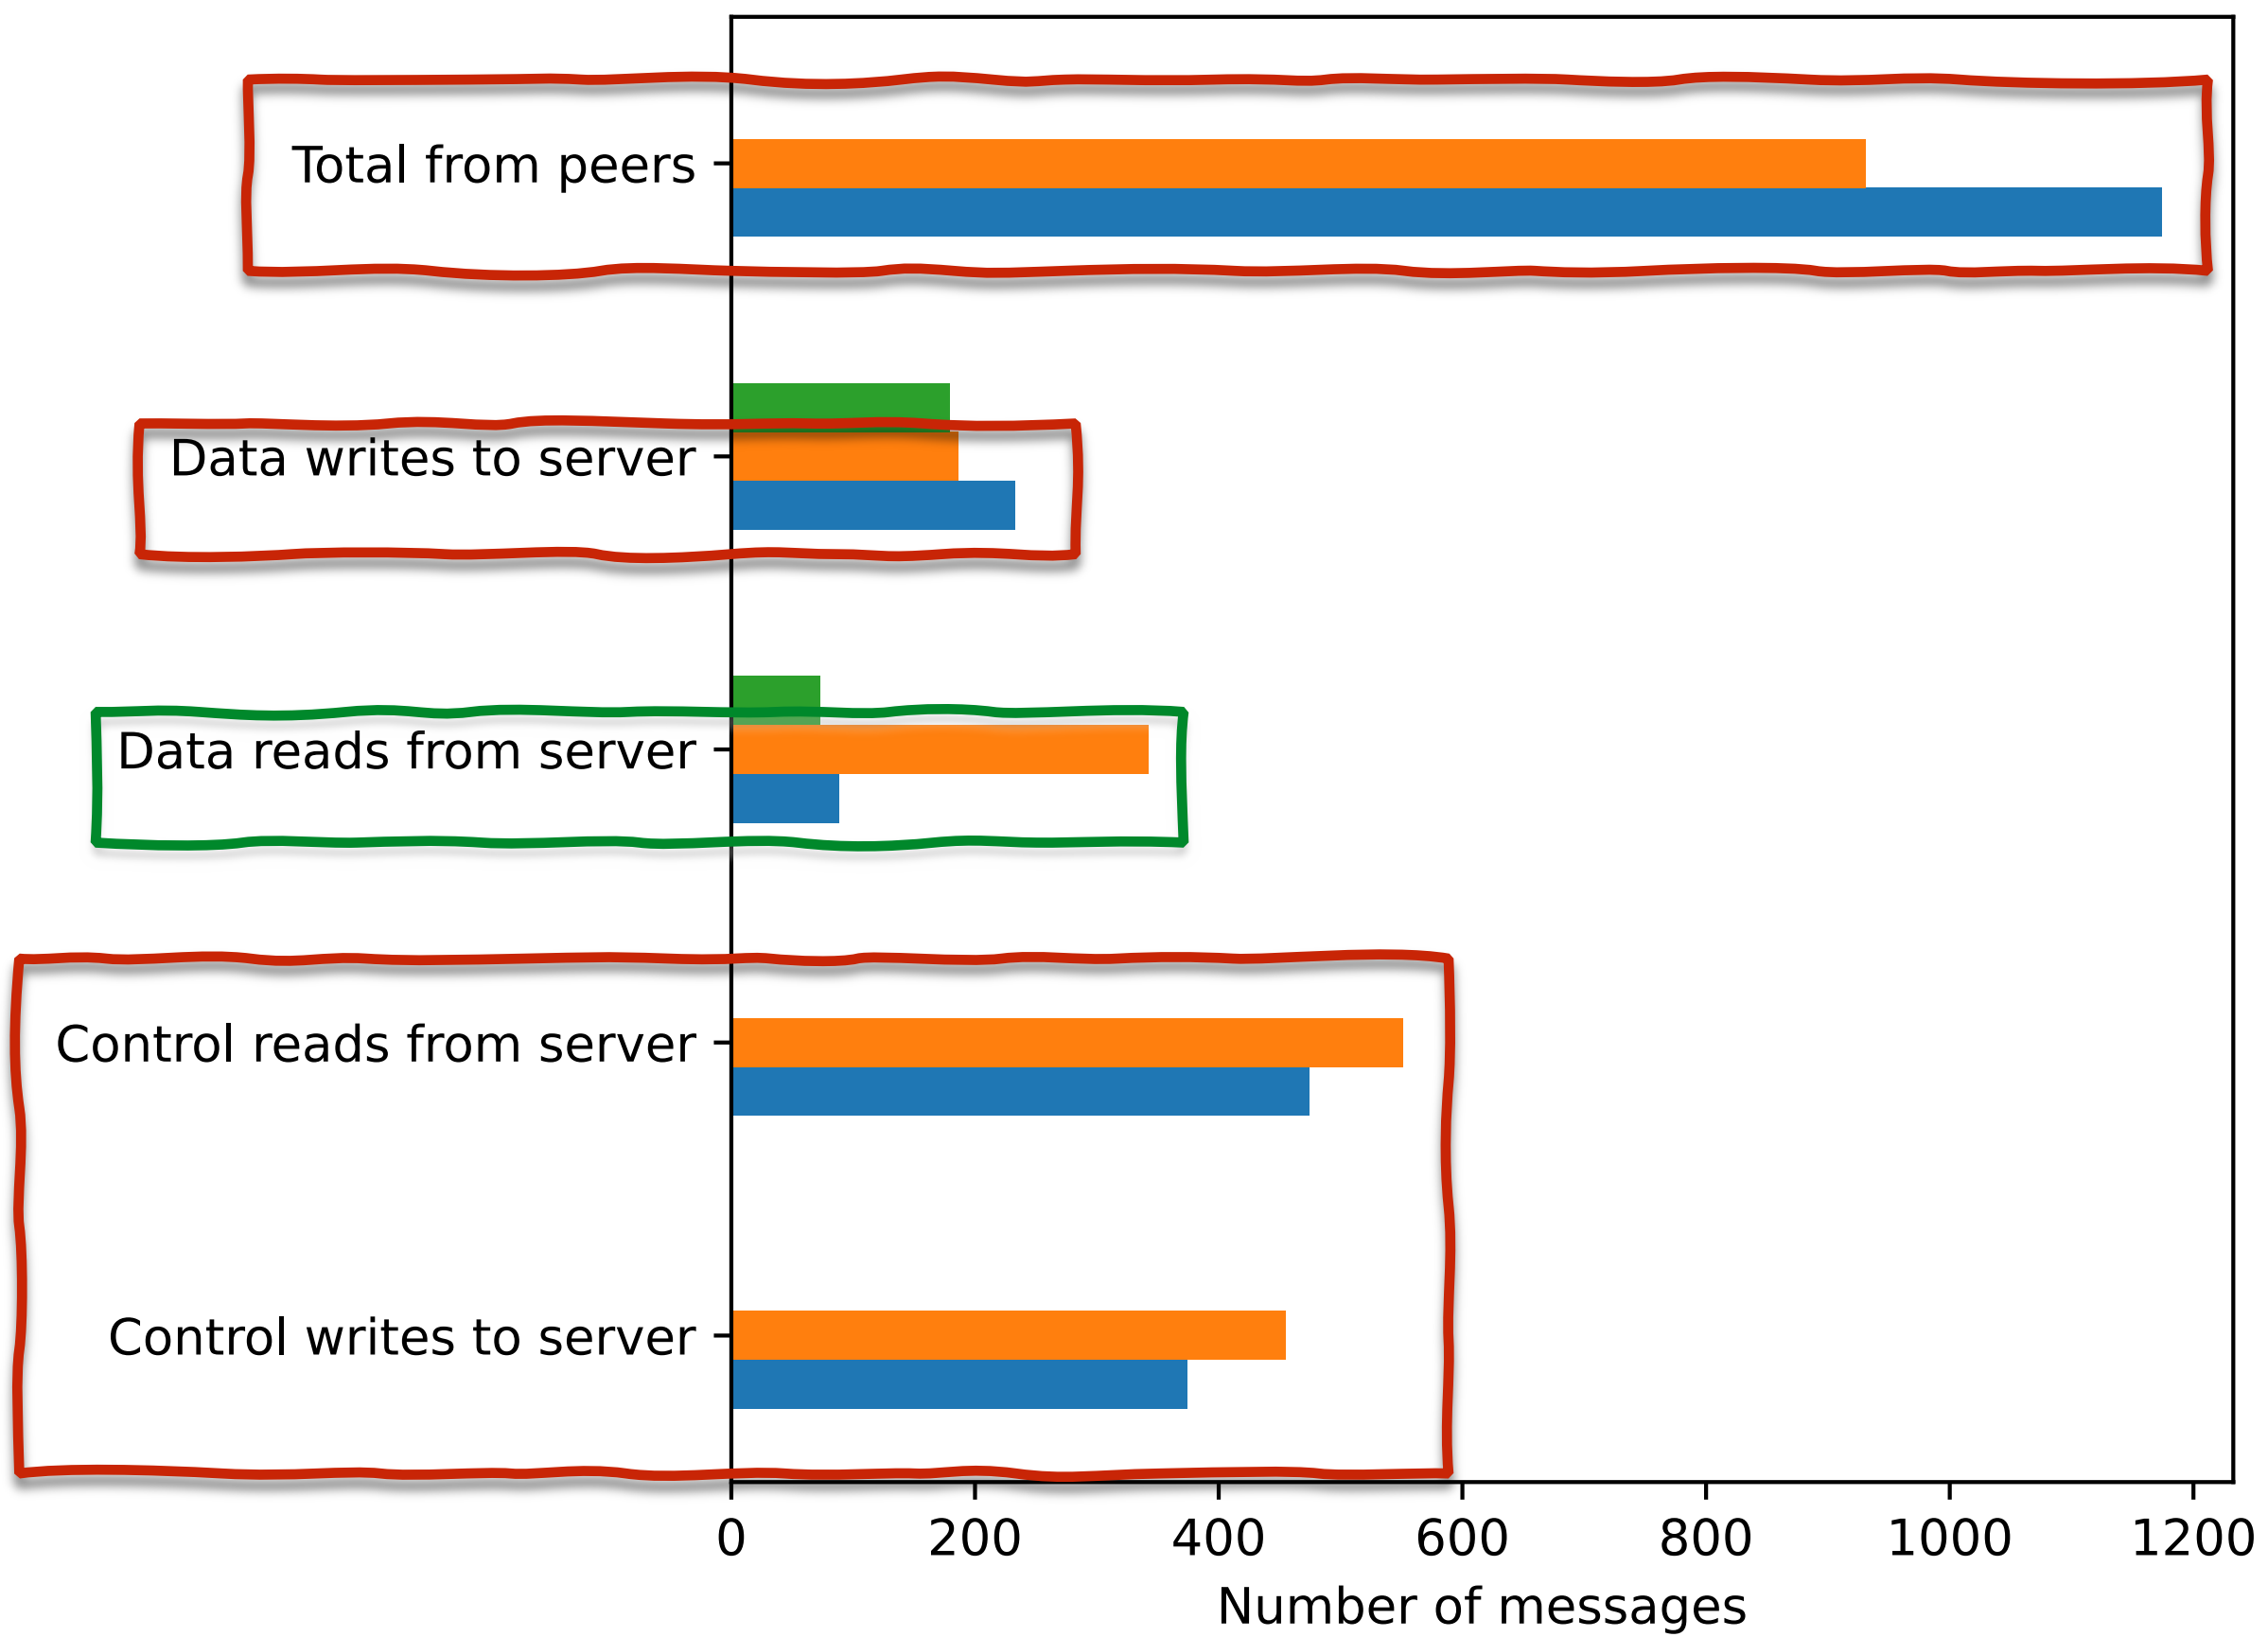
Evaluation

Check-in scenario




Evaluation

Post a review scenario



Evaluation

Latency





Future Work

- Fine tune our protocol to decrease the number of messages in the overlay
- Extend the evaluation
 - Broader the scale with more clients
 - Simulate the protocol stack



Summary

- Introduced the GeoLoc Framework with a overlay and bully algorithm
- Introduced a data model that accounts for objects and clients' locations.
- Enabling clients only to replicate data they hold an interest in, e.g. partial replication.
- Provide lower latency when compared to the client-server model.
- The drawback comes in the form of more metadata travelling the network.

Thank You!