Towards a Replication Service for Data-Intensive Fog Applications

Jonathan Hasenburg, Martin Grambow, David Bermbach

Mobile Cloud Computing Research Group

TU Berlin & ECDF, Germany

PROBLEM

Fog computing promises significant benefits for applications but is hard to use:

- Existing storage systems are poorly suited for Fog environments
- Data distribution across systems needs to be implemented inside the application
- Data-intensive Fog applications often need data streams with distributed persistence

APPROACH

- Application-controlled data placement and movement using a **declarative programming** style
- A replication service handles actual data distribution across Fog nodes
- Groups of data that should be handled in the same way are part of the same **keygroup**
- Machines at the same location are exposed as a single fog **node** to other nodes

Application Scenarios
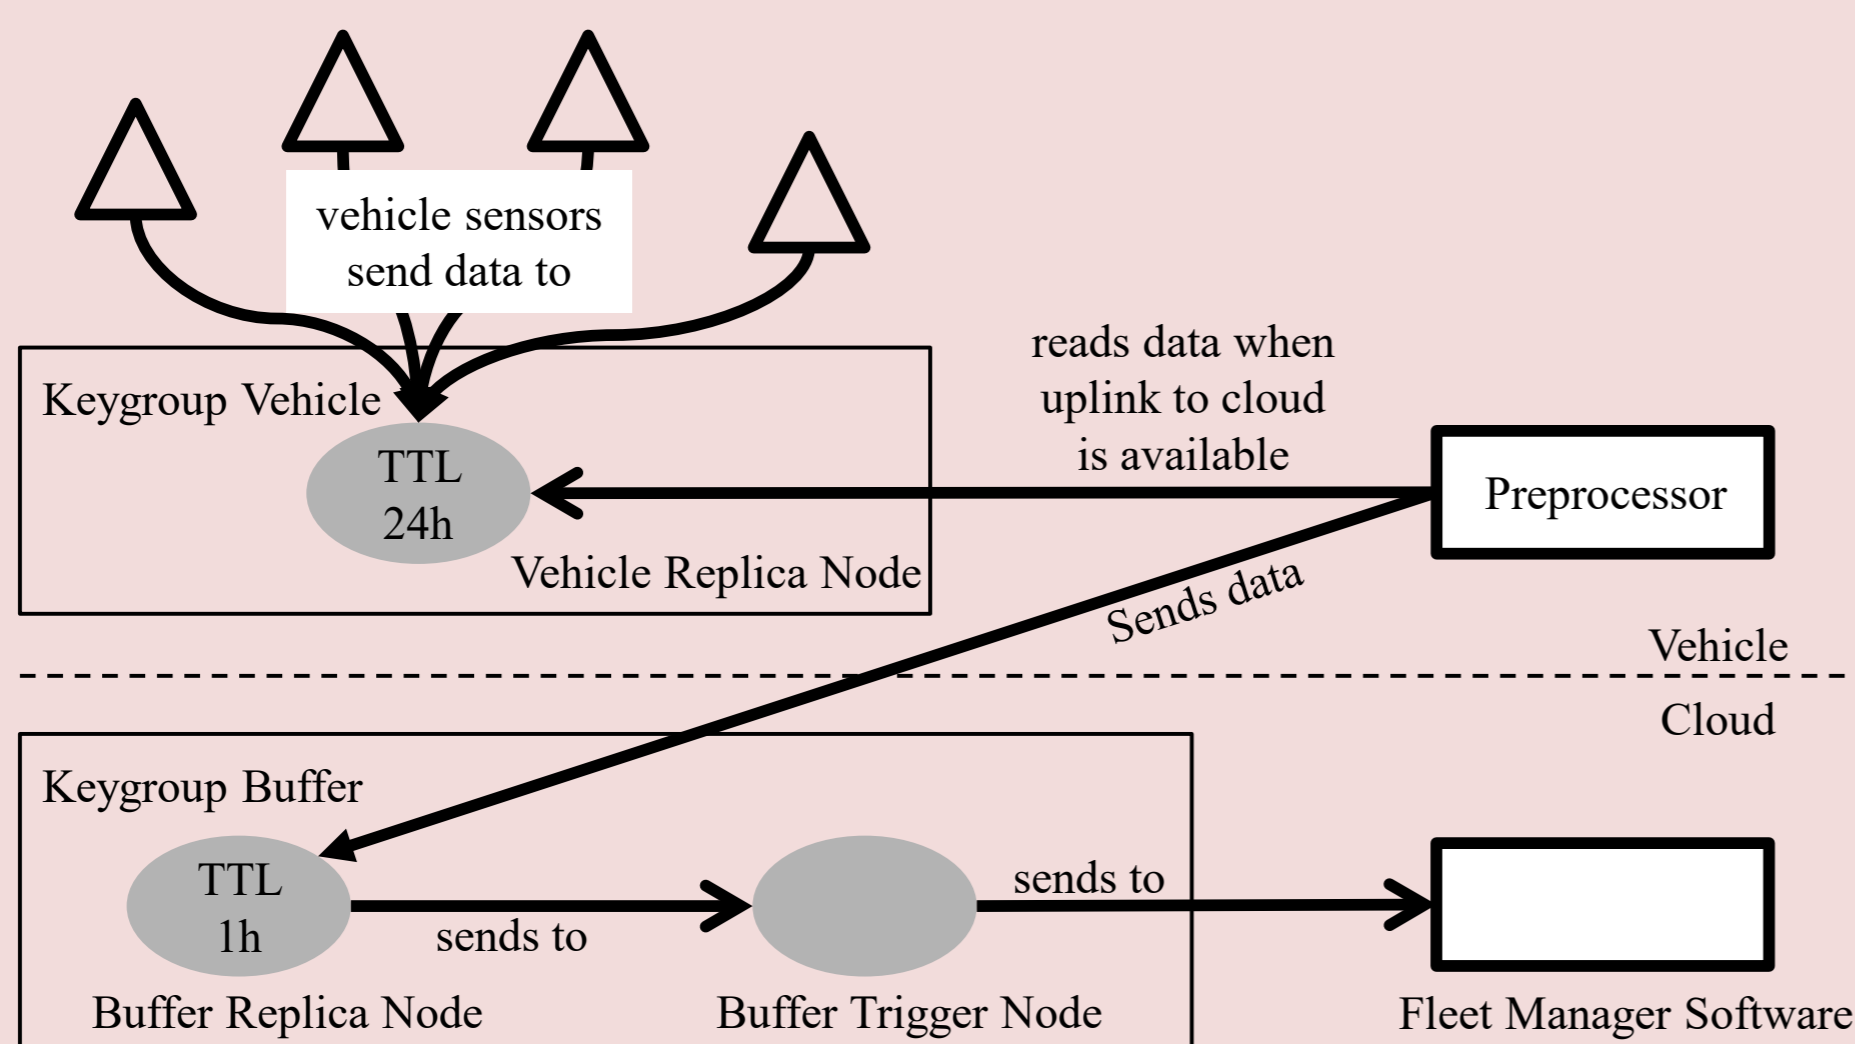
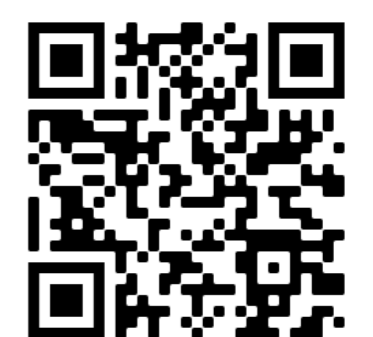
Carsharing Fleet Management

- Vehicle sensors collect information for the carsharing provider
- Keygroup "Vehicle" defines that this data is stored for 24 hours inside the vehicle
- A preprocessor aggregates the data and transmits it to the Cloud when possible
- Keygroup "Buffer" defines that data is transmitted to a trigger node which here acts as an event-based connector to the fleet manager software
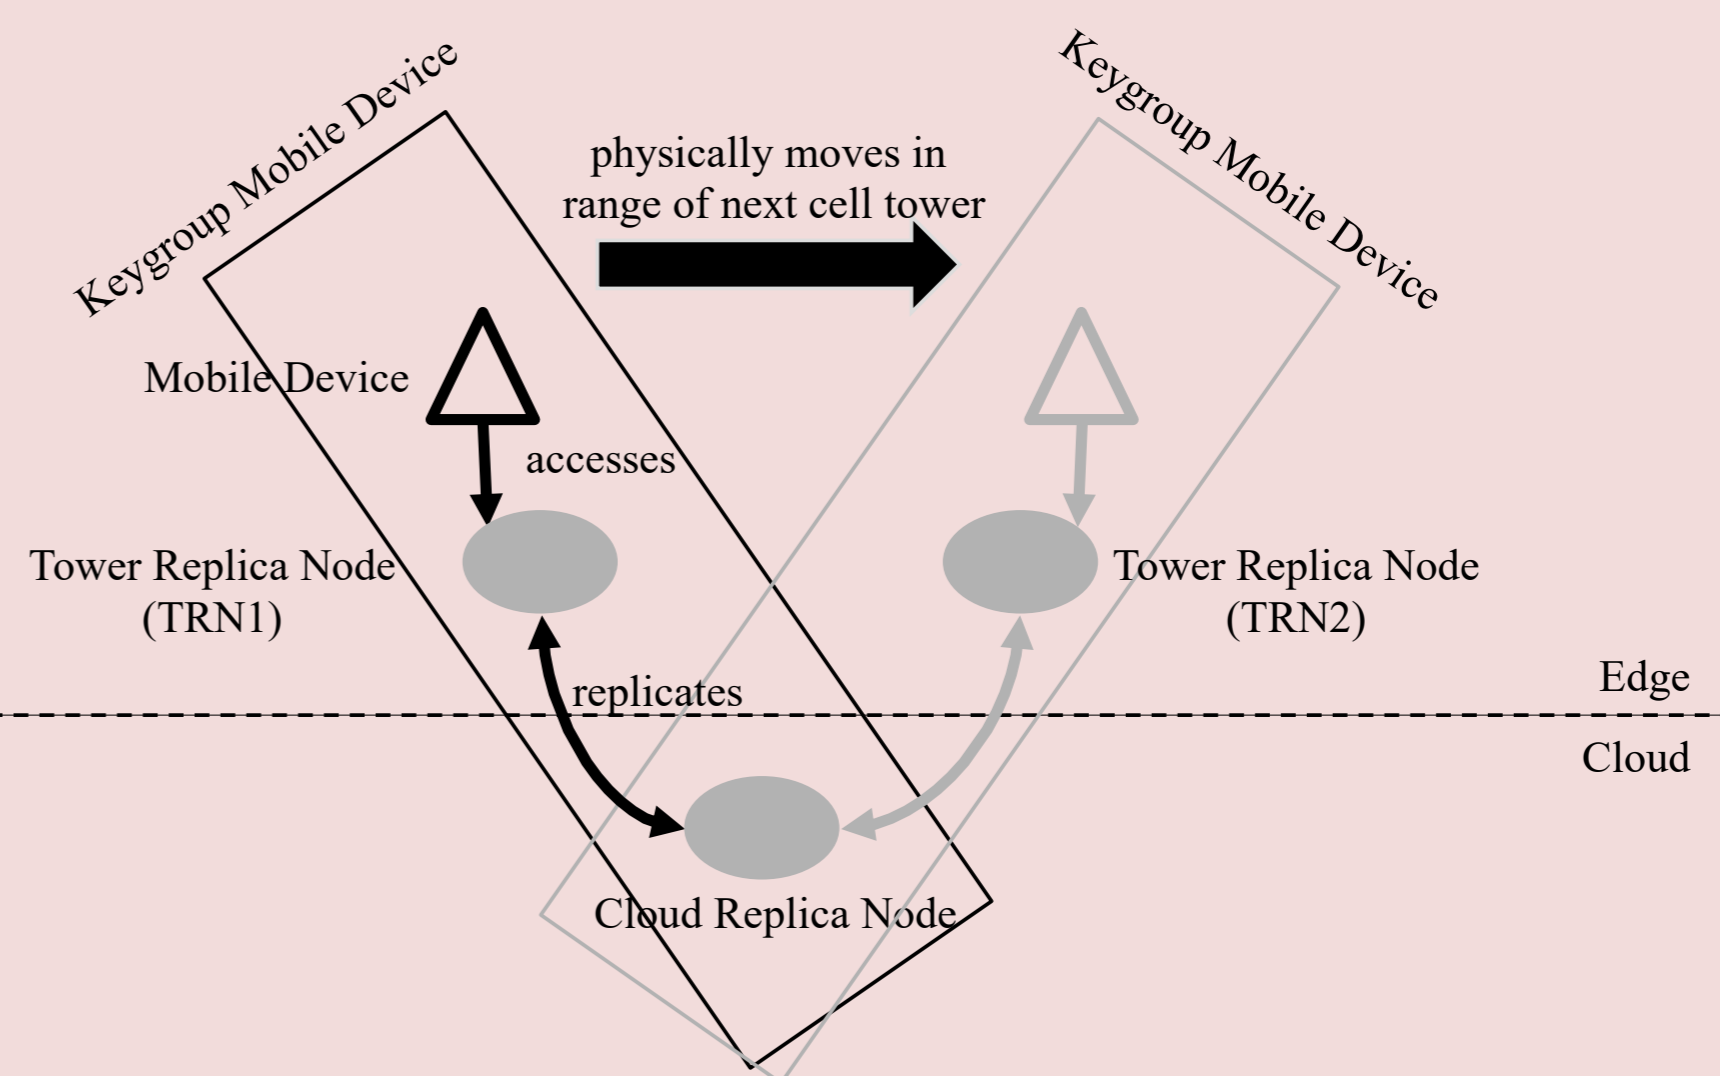
Mobile App Scenario

- A mobile device runs an app that uses the cloud for persistence
- Data is also available at the edge to speed up data retrieval
- Upon user movement the keygroup configuration is updated and the service migrates the application data.

Software Prototype
Technical report