



Technische
Universität
Braunschweig



ZUGCHAIN: Blockchain-Based Juridical Data Recording in Railway Systems

Conference on Dependable Systems and Networks 2022

Signe Rüsch¹, Kai Bleeke¹, Ines Messadi¹, Stefan Schmid¹,
Andreas Krampf², Katharina Olze², Susanne Stahnke³, Robert Schmid⁴, Lukas Pirl⁴,
Roland Kittel⁵, Andreas Polze⁴, Marquart Franz³, Matthias Müller²,
Leander Jehl^{1,6}, and Rüdiger Kapitza¹


¹TU Braunschweig, Germany {ruesch, bleeke}@ibr.cs.tu-bs.de

²Siemens Mobility ³Siemens ⁴HPI, University of Potsdam

⁵DB Systel ⁶University of Stavanger


Current Juridical Data Logging

- Juridical Recording Unit (JRU):
 - Train's “black box”, records all juridically relevant data
 - E.g. speed, brake activity, door activity, timestamps, ...
 - Data to be logged received via bus communication




Current Juridical Data Logging

- Juridical Recording Unit (JRU):
 - Train's “black box”, records all juridically relevant data
 - E.g. speed, brake activity, door activity, timestamps, ...
 - Data to be logged received via bus communication
- Detect malfunctions, accident root cause analysis
- Built to withstand **physical damage**




Juridical Data Logging: Problems

- **Single point of failure** [Hartong et al., 2008]
 - Potential data loss
 - Destruction in crash
 - Data manipulation during extraction




Juridical Data Logging: Problems

- **Single point of failure** [Hartong et al., 2008]
 - Potential data loss
 - Destruction in crash
 - Data manipulation during extraction
- **Complex data extraction**
 - Access via physical key
 - No remote extraction available




Juridical Data Logging: Problems

- **Single point of failure** [Hartong et al., 2008]
 - Potential data loss
 - Destruction in crash
 - Data manipulation during extraction
- **Complex data extraction**
 - Access via physical key
 - No remote extraction available



Juridical Data Logging: Problems

- **Single point of failure** [Hartong et al., 2008]
 - Potential data loss
 - Destruction in crash
 - Data manipulation during extraction
- **Complex data extraction**
 - Access via physical key
 - No remote extraction available
- **Expensive and proprietary device**
 - One company has logging authority
 - No distribution of trust



Juridical Data Logging: Goals

- Distributed:
 - Replace single device with **distributed, replicated system**
 - ↳ High reliability and availability
 - ↳ Ensure data integrity

Juridical Data Logging: Goals

- Distributed:
 - Replace single device with **distributed, replicated system**
 - ↪ High reliability and availability
 - ↪ Ensure data integrity
- Cost factor:
 - JRU: expensive, dedicated device
 - ↪ Use **on-train commodity hardware**
 - ↪ No changes to safety-critical infrastructure

Juridical Data Logging: Goals

- Distributed:
 - Replace single device with **distributed, replicated system**
 - High reliability and availability
 - Ensure data integrity
- Cost factor:
 - JRU: expensive, dedicated device
 - Use **on-train commodity hardware**
 - No changes to safety-critical infrastructure
- Allow **multiple companies** to contribute to juridical recording

Juridical Data Logging: Goals

- Distributed:
 - Replace single device with **distributed, replicated system**
 - High reliability and availability
 - Ensure data integrity
- Cost factor:
 - JRU: expensive, dedicated device
 - Use **on-train commodity hardware**
 - No changes to safety-critical infrastructure
- Allow **multiple companies** to contribute to juridical recording
- Facilitate **easy data export** for predictive maintenance

Distributed Juridical Data Logging

- **Mathematical analysis** of distributed JRUs [Braband et al., 2020]
 - Approach: distributed JRU functionality on commodity hardware

Distributed Juridical Data Logging

- **Mathematical analysis** of distributed JRUs [Braband et al., 2020]
 - Approach: distributed JRU functionality on commodity hardware
 - Examined train crash data provided by the UIC and ERA
 - Analysed probabilities for different scenarios (e.g. fire) and their effects on the recording hardware

Distributed Juridical Data Logging

- **Mathematical analysis** of distributed JRUs [Braband et al., 2020]
 - Approach: distributed JRU functionality on commodity hardware
 - Examined train crash data provided by the UIC and ERA
 - Analysed probabilities for different scenarios (e.g. fire) and their effects on the recording hardware
- Findings:
 - Commodity hardware easier to destroy
 - Crashes rarely affect whole train
 - Distributed JRU reaches **comparable reliability and availability**

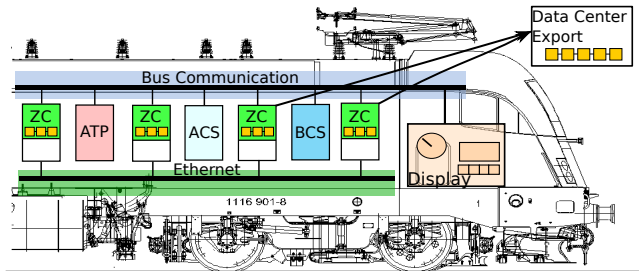
Distributed Juridical Data Logging

- **Mathematical analysis** of distributed JRUs [Braband et al., 2020]
 - Approach: distributed JRU functionality on commodity hardware
 - Examined train crash data provided by the UIC and ERA
 - Analysed probabilities for different scenarios (e.g. fire) and their effects on the recording hardware
- Findings:
 - Commodity hardware easier to destroy
 - Crashes rarely affect whole train
 - Distributed JRU reaches **comparable reliability and availability**

➔ How can we ensure that all nodes **consistently** log the same data?

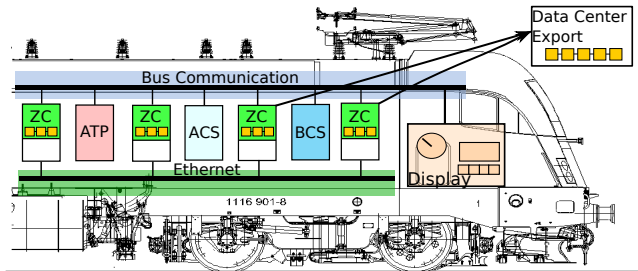
ZUGCHAIN Design

- ZUGCHAIN nodes distributed across the train
- Read safety-critical signals from **bus**




ZUGCHAIN Design

- ZUGCHAIN nodes distributed across the train
- Read safety-critical signals from **bus**
- Nodes **agree** on order and log received signals in **blockchain**



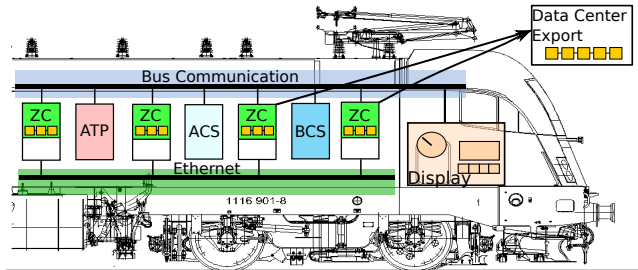
ZUGCHAIN Design

- ZUGCHAIN nodes distributed across the train
- Read safety-critical signals from **bus**
- Nodes **agree** on order and log received signals in **blockchain**
- Use **existing, non-critical links** for consensus



ZUGCHAIN Design

- ZUGCHAIN nodes distributed across the train
- Read safety-critical signals from **bus**
- Nodes **agree** on order and log received signals in **blockchain**
- Use **existing, non-critical links** for consensus
- Blockchain structure ensures **data integrity** and enables **export**



Alternative Designs

- Multiple independent nodes
 - Logs are not synchronized
 - Can lead to **data loss** in case of a crash

Alternative Designs

- Multiple independent nodes
 - Logs are not synchronized
 - Can lead to **data loss** in case of a crash
- Simple leader/follower architecture
 - Only one node needs to be read in case of a crash
 - However, data can still be **compromised**
 - Bus system can also lead to **corrupted transmissions**

Alternative Designs



- Multiple independent nodes
 - Logs are not synchronized
 - Can lead to **data loss** in case of a crash
- Simple leader/follower architecture
 - Only one node needs to be read in case of a crash
 - However, data can still be **compromised**
 - Bus system can also lead to **corrupted transmissions**
- Agreement
 - ZUGCHAIN nodes may be **co-located** on observed machines
 - Arbitrary faults on machine can influence recording functionality

Alternative Designs

- Multiple independent nodes
 - Logs are not synchronized
 - Can lead to **data loss** in case of a crash
- Simple leader/follower architecture
 - Only one node needs to be read in case of a crash
 - However, data can still be **compromised**
 - Bus system can also lead to **corrupted transmissions**
- Agreement
 - ZUGCHAIN nodes may be **co-located** on observed machines
 - Arbitrary faults on machine can influence recording functionality

↪ Byzantine agreement is necessary for ZUGCHAIN

Bus Communication Instead of Clients





Traditional clients

- Authenticated requests
- Replicas can forward requests
- Client can re-transmit requests

Bus communication

- Unauthenticated broadcast
- Synchronous communication
- Not interested in replies

Bus Communication Instead of Clients




- Traditional clients
 - Authenticated requests
 - Replicas can forward requests
 - Client can re-transmit requests
- Bus communication
 - Unauthenticated broadcast
 - Synchronous communication
 - Not interested in replies

Bus clients have to be treated differently


ZUGCHAIN Design: Communication Layer

- Underlying Byzantine agreement protocol, e.g. **PBFT**
- Nodes can read **identical** or **diverging** data from bus
- Communication layer ensures **no payload duplication**
 - Log all juridically relevant data, but **filter** duplicates
 - Prevent **omission** of data




ZUGCHAIN Design: Communication Layer

- Nodes assume others read identical data
- Primary** proposes request, **backups** don't propose
- But they set a **“soft” timer** for each request:
 - If it expires before matching pre-prepare, **broadcasting** to all
 - Start **“hard” timer**, before view change




Export Protocol: Goals

- Storage on the node is limited
 - Recent data is of highest priority
 - Block headers necessary for blockchain
- ➔ We regularly export data to avoid deletion





Export Protocol: Goals

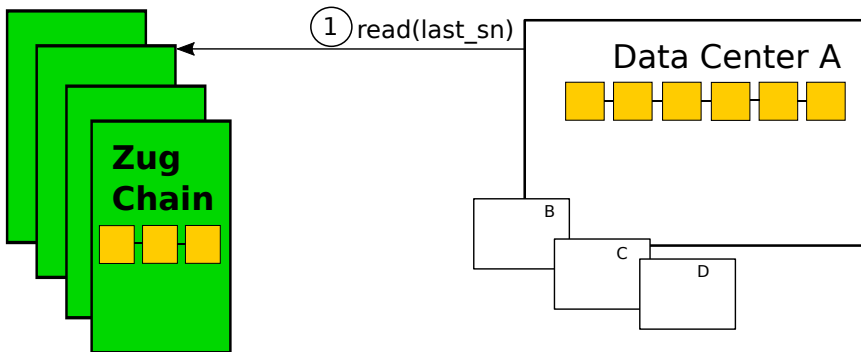
- Storage on the node is limited
 - Recent data is of highest priority
 - Block headers necessary for blockchain
 - ➔ We regularly export data to avoid deletion
- Export to data centers of railway companies
 - Multiple parties sign off on deletion
- Blockchain facilitates export
 - Verify integrity from a single response



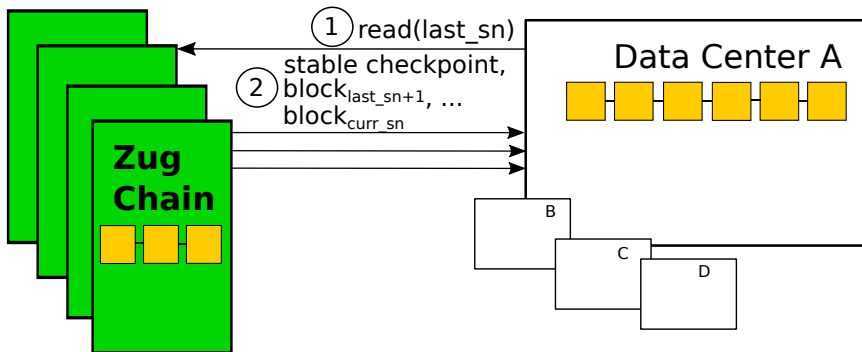
ZUGCHAIN Export Protocol: Steps



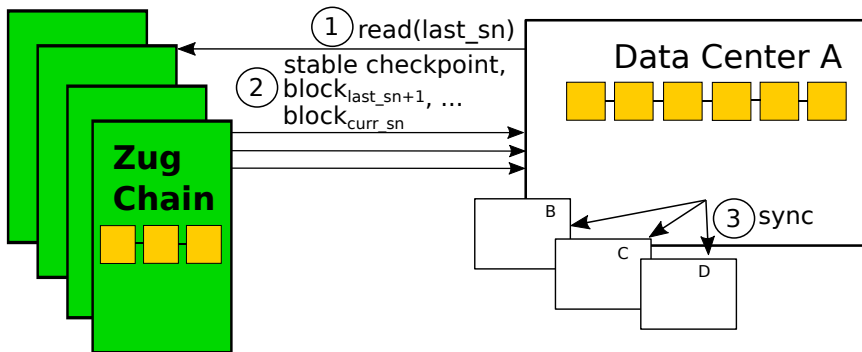
ZUGCHAIN Export Protocol: Steps



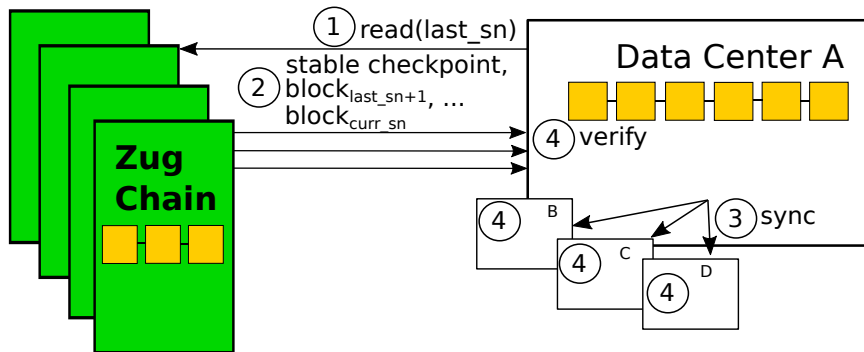
ZUGCHAIN Export Protocol: Steps



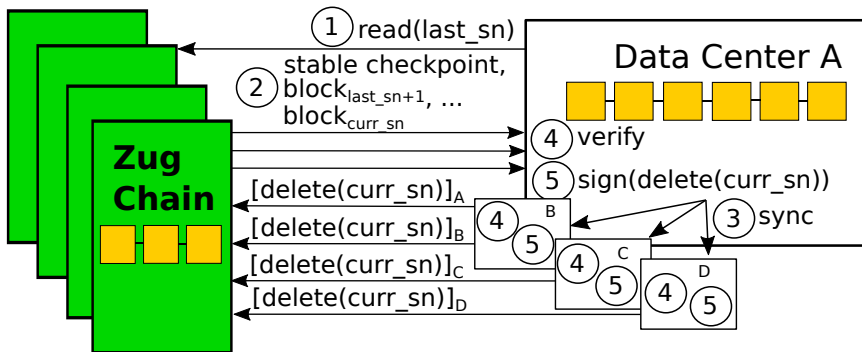
ZUGCHAIN Export Protocol: Steps



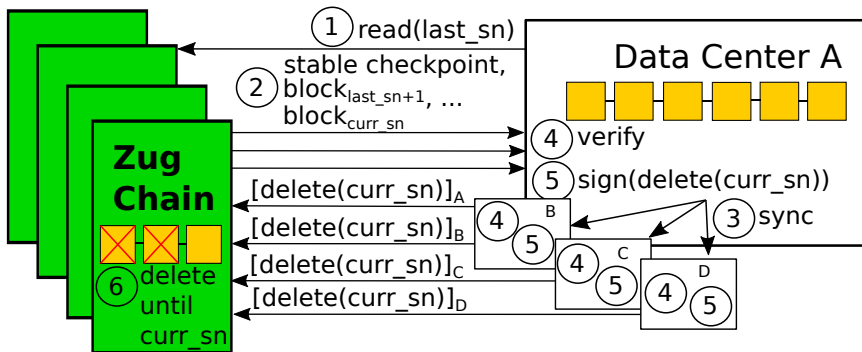
ZUGCHAIN Export Protocol: Steps



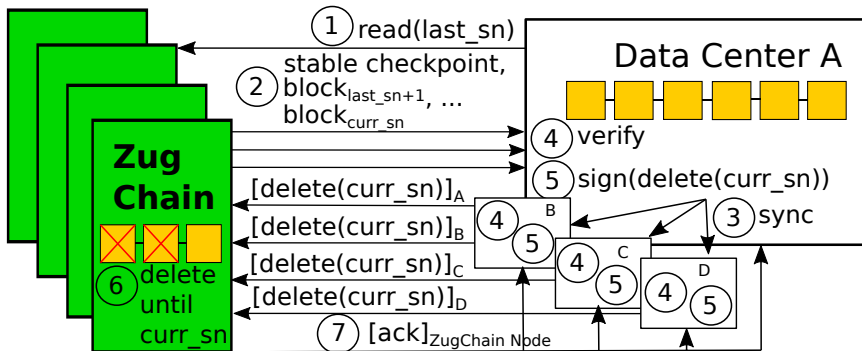
ZUGCHAIN Export Protocol: Steps



ZUGCHAIN Export Protocol: Steps




ZUGCHAIN Export Protocol: Steps



ZUGCHAIN Evaluation Setting

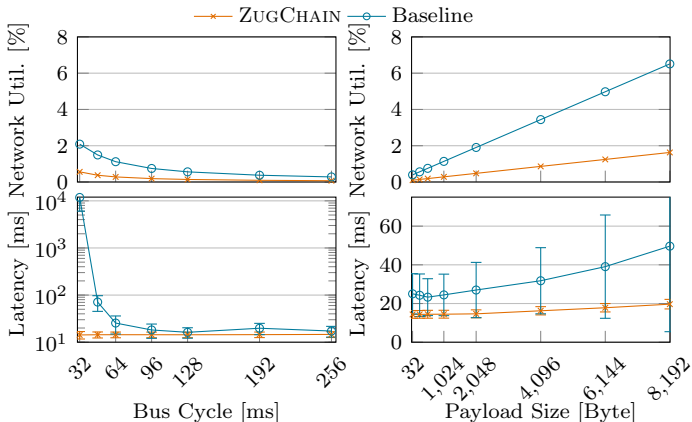
- 4 M-COMs:
 - Quad-Core CPU
 - 2GB RAM
 - 100Mbit/s Ethernet and MVB connection
 - Yocto Linux, kernel v3.10.17
- Communication layer vs. PBFT:
 - Naïve baseline
 - Order **all** data, i.e. requests logged 4×
- Vary bus cycle time and payload size



Background: Multifunction Vehicle Bus (MVB)


- Fieldbus supported by Siemens and ABB
- Synchronized leader/follower communication
 - Bus master sets the **cycle time** with each follower
- Three types of data:
 - Process Data
 - Message Data
 - Supervisory Data
- Focus on **process data** for juridical recording
- Communication errors can still occur

Evaluation – Network Utilization and Latency



➔ ZUGCHAIN: less **bandwidth** and lower, more stable **latencies!**


Evaluation – Resource Usage



➔ ZUGCHAIN better suited for **shared, resource-constrained devices!**

ZUGCHAIN Export Evaluation – LTE


- 5 minutes (500 blocks) to 3 hours (16k blocks) of data
 - Exported to AWS VM t2.xlarge



➔ Export feasible during train stops or continuously!

ZUGCHAIN Export Evaluation – LTE



- 5 minutes (500 blocks) to 3 hours (16k blocks) of data
 - Exported to AWS VM t2.xlarge



- Export feasible during train stops or continuously!
- View change and Byzantine fault measurements in paper!



Conclusion

- Further **digitalization** in railway operations
- Maintain JRU's **reliability and availability**





Conclusion

- Further **digitalization** in railway operations
- Maintain JRU's **reliability and availability**
- Allow data **integrity** verification after crashes by using blockchains
 - Even with only one remaining copy





Conclusion

- Further **digitalization** in railway operations
- Maintain JRU's **reliability and availability**
- Allow data **integrity** verification after crashes by using blockchains
 - Even with only one remaining copy
- Easy **data export** to railway companies





Conclusion

- Further **digitalization** in railway operations
- Maintain JRU's **reliability and availability**
- Allow data **integrity** verification after crashes by using blockchains
 - Even with only one remaining copy
- Easy **data export** to railway companies
- Efficient filtering in communication layer for good **performance** on real train hardware



Conclusion

- Further **digitalization** in railway operations
- Maintain JRU's **reliability and availability**
- Allow data **integrity** verification after crashes by using blockchains
 - Even with only one remaining copy
- Easy **data export** to railway companies
- Efficient filtering in communication layer for good **performance** on real train hardware

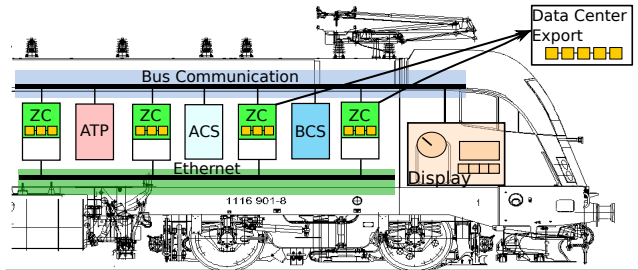


Thank you for your attention! Questions?
 {ruesch,bleeke}@ibr.cs.tu-bs.de



ZUGCHAIN Design: Crash Scenario

- After a crash, any tampering of data should be **detected**
- Probability of multiple nodes surviving crash are high [Braband, 2020]
- If only one copy survives ...
 - **Signatures** on blocks and consensus messages in blockchain
 - Verify that this is an authentic log
 - Tampered data cannot contain **$2f+1$ valid replica signatures!**



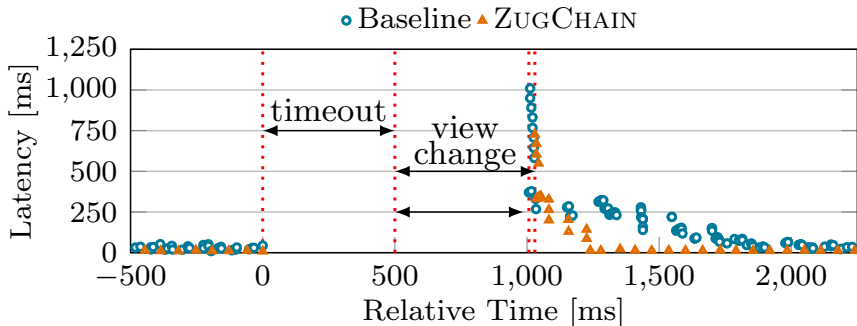
ZUGCHAIN Interface

Module	Call	Explanation
① down	PROPOSE(<i>r</i>)	proposes request to consensus group
① down	SUSPECT(<i>id</i>)	suspect node to be faulty, init. view change
① up	DECIDE(<i>r</i> , <i>sn</i>)	totally ordered request and seq.no.
① up	NEWPRIMARY	returns new primary after view change
② down	RECEIVE(<i>req</i>)	read parsed request from bus
② up	LOG(<i>req</i> , <i>id</i> , <i>sn</i>)	append request to totally ordered log

Table 1: Interfaces of BFT (①) and ZUGCHAIN (②).

ZUGCHAIN View Change Evaluation

- Complex operation: after timeout, exchange all open requests
- ↪ ZUGCHAIN has to handle fewer messages
- Duration: ZUGCHAIN 530ms, PBFT 507ms
- Stabilization time: ZUGCHAIN 210ms, PBFT 824ms



ZUGCHAIN Byzantine Behaviour Evaluation

