Mark Raasveldt, Hannes Mühleisen

Don’t Hold My Data Hostage

A Case For Client Protocol Redesign

What is a Client Protocol anyway?

- Protocol is how a client communicates with a server
  - ODBC, JDBC, psql
- Every database that supports remote clients has a client protocol

- Using this protocol, clients can:
  - Connect to the database
  - Query it
  - Receive the query results

What is a Client Protocol anyway?

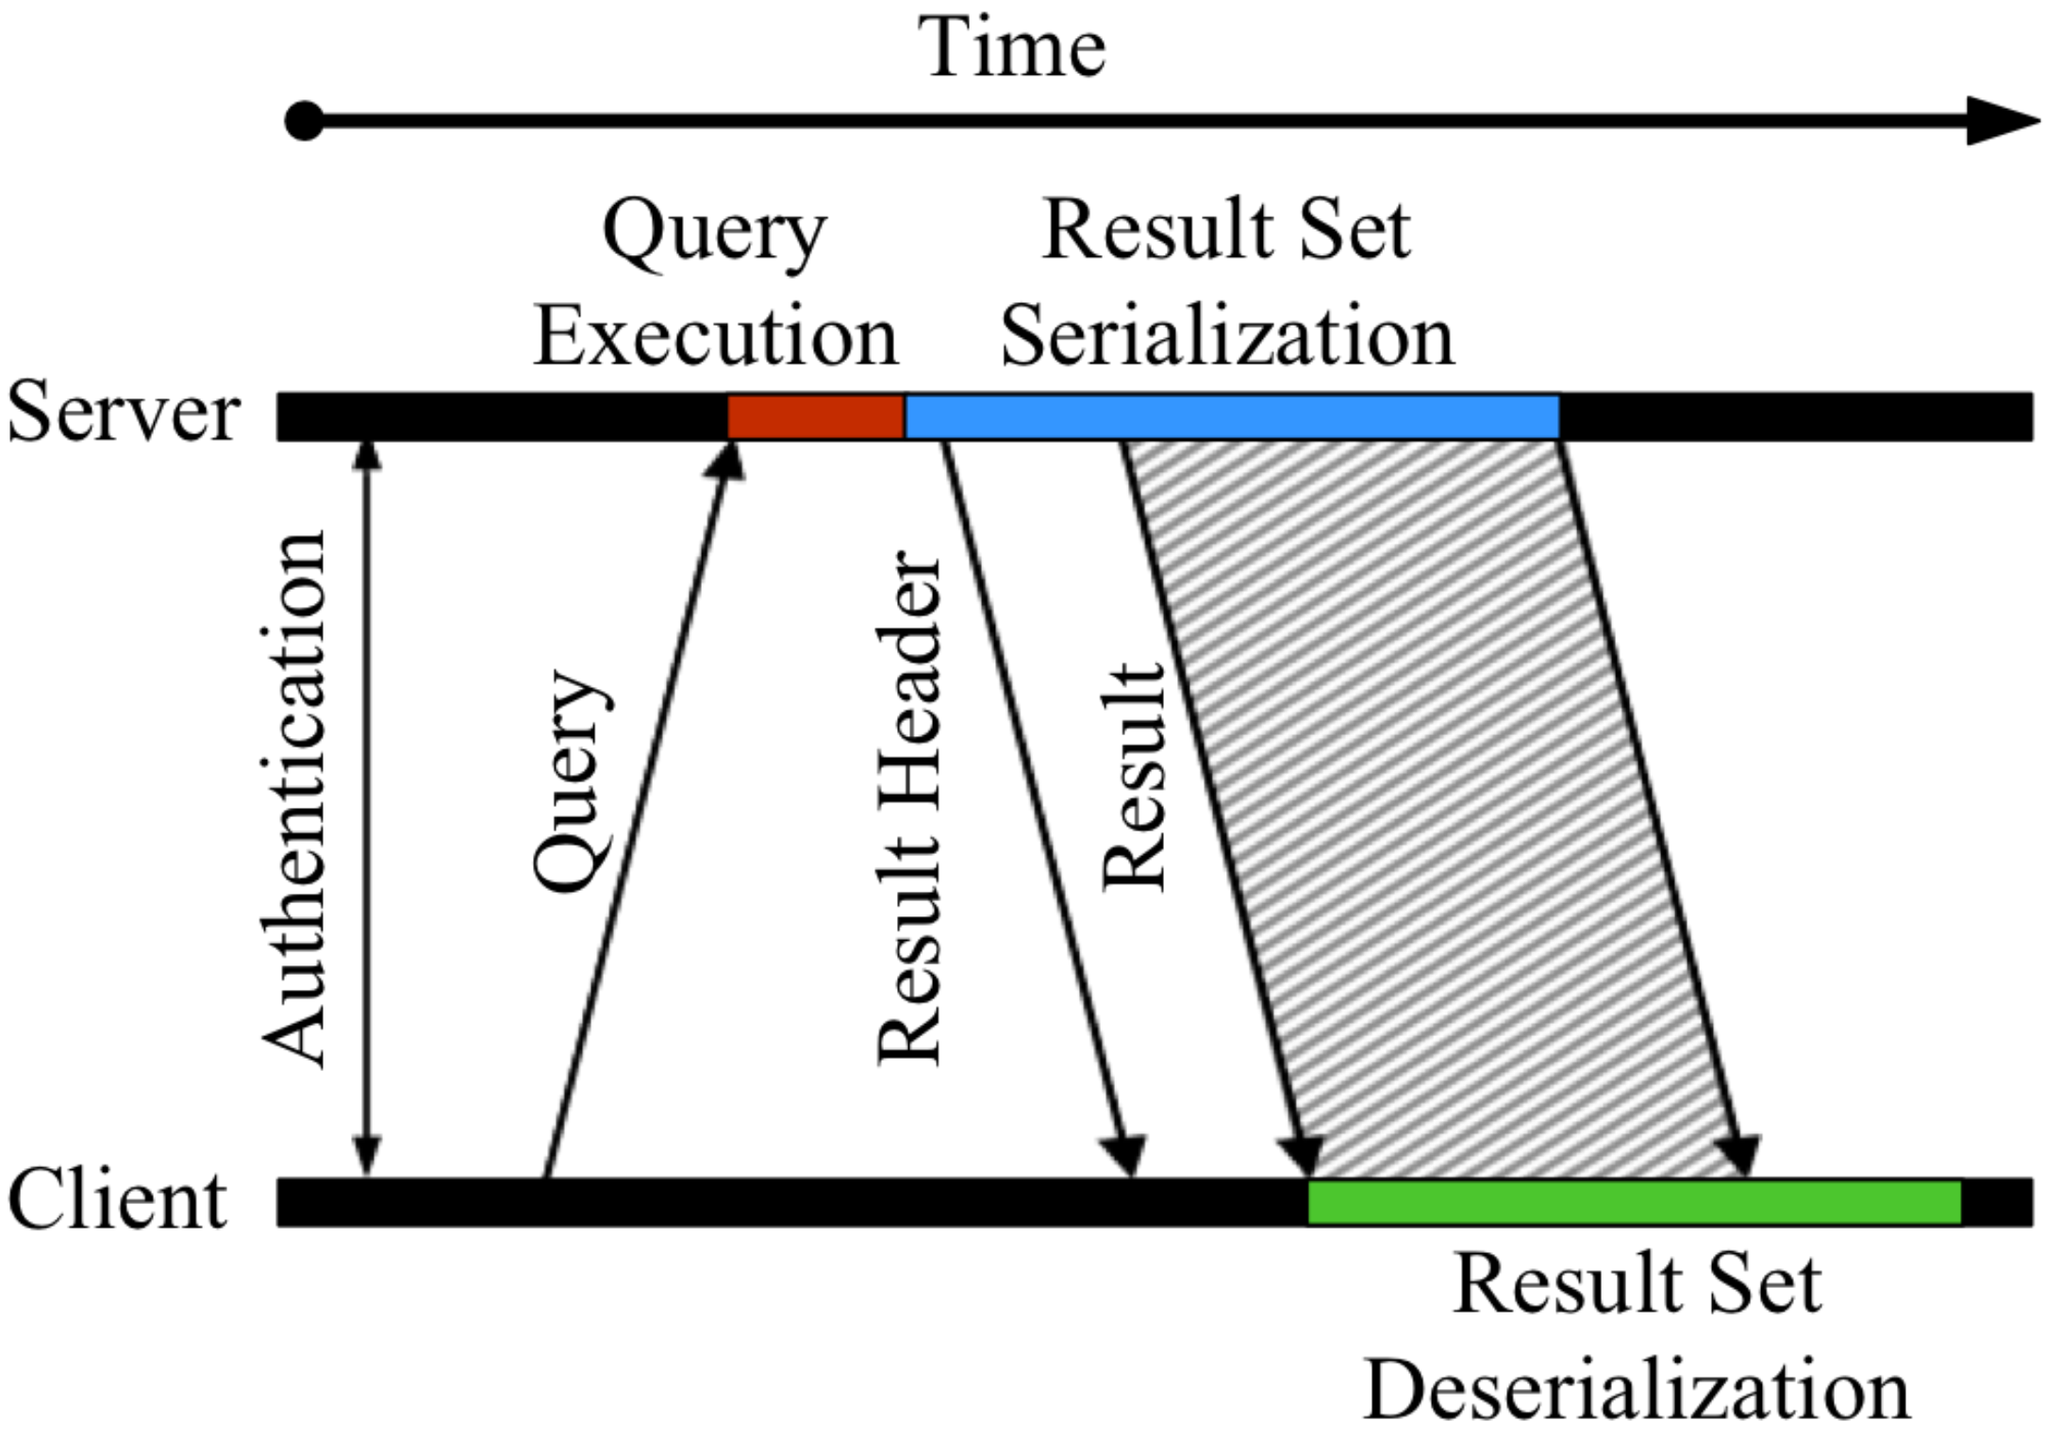

Motivation

- Problem: Current protocols were designed for exporting small amount of rows
  - Displaying results on screen
  - OLTP use cases
  - Exporting aggregates
- Exporting large amounts of data using these protocols is slow
  - Analytical tools (e.g. R/Python)
- Data export is a bottleneck when result sets are large!

Motivation

```
SELECT * FROM LINEITEM;
```

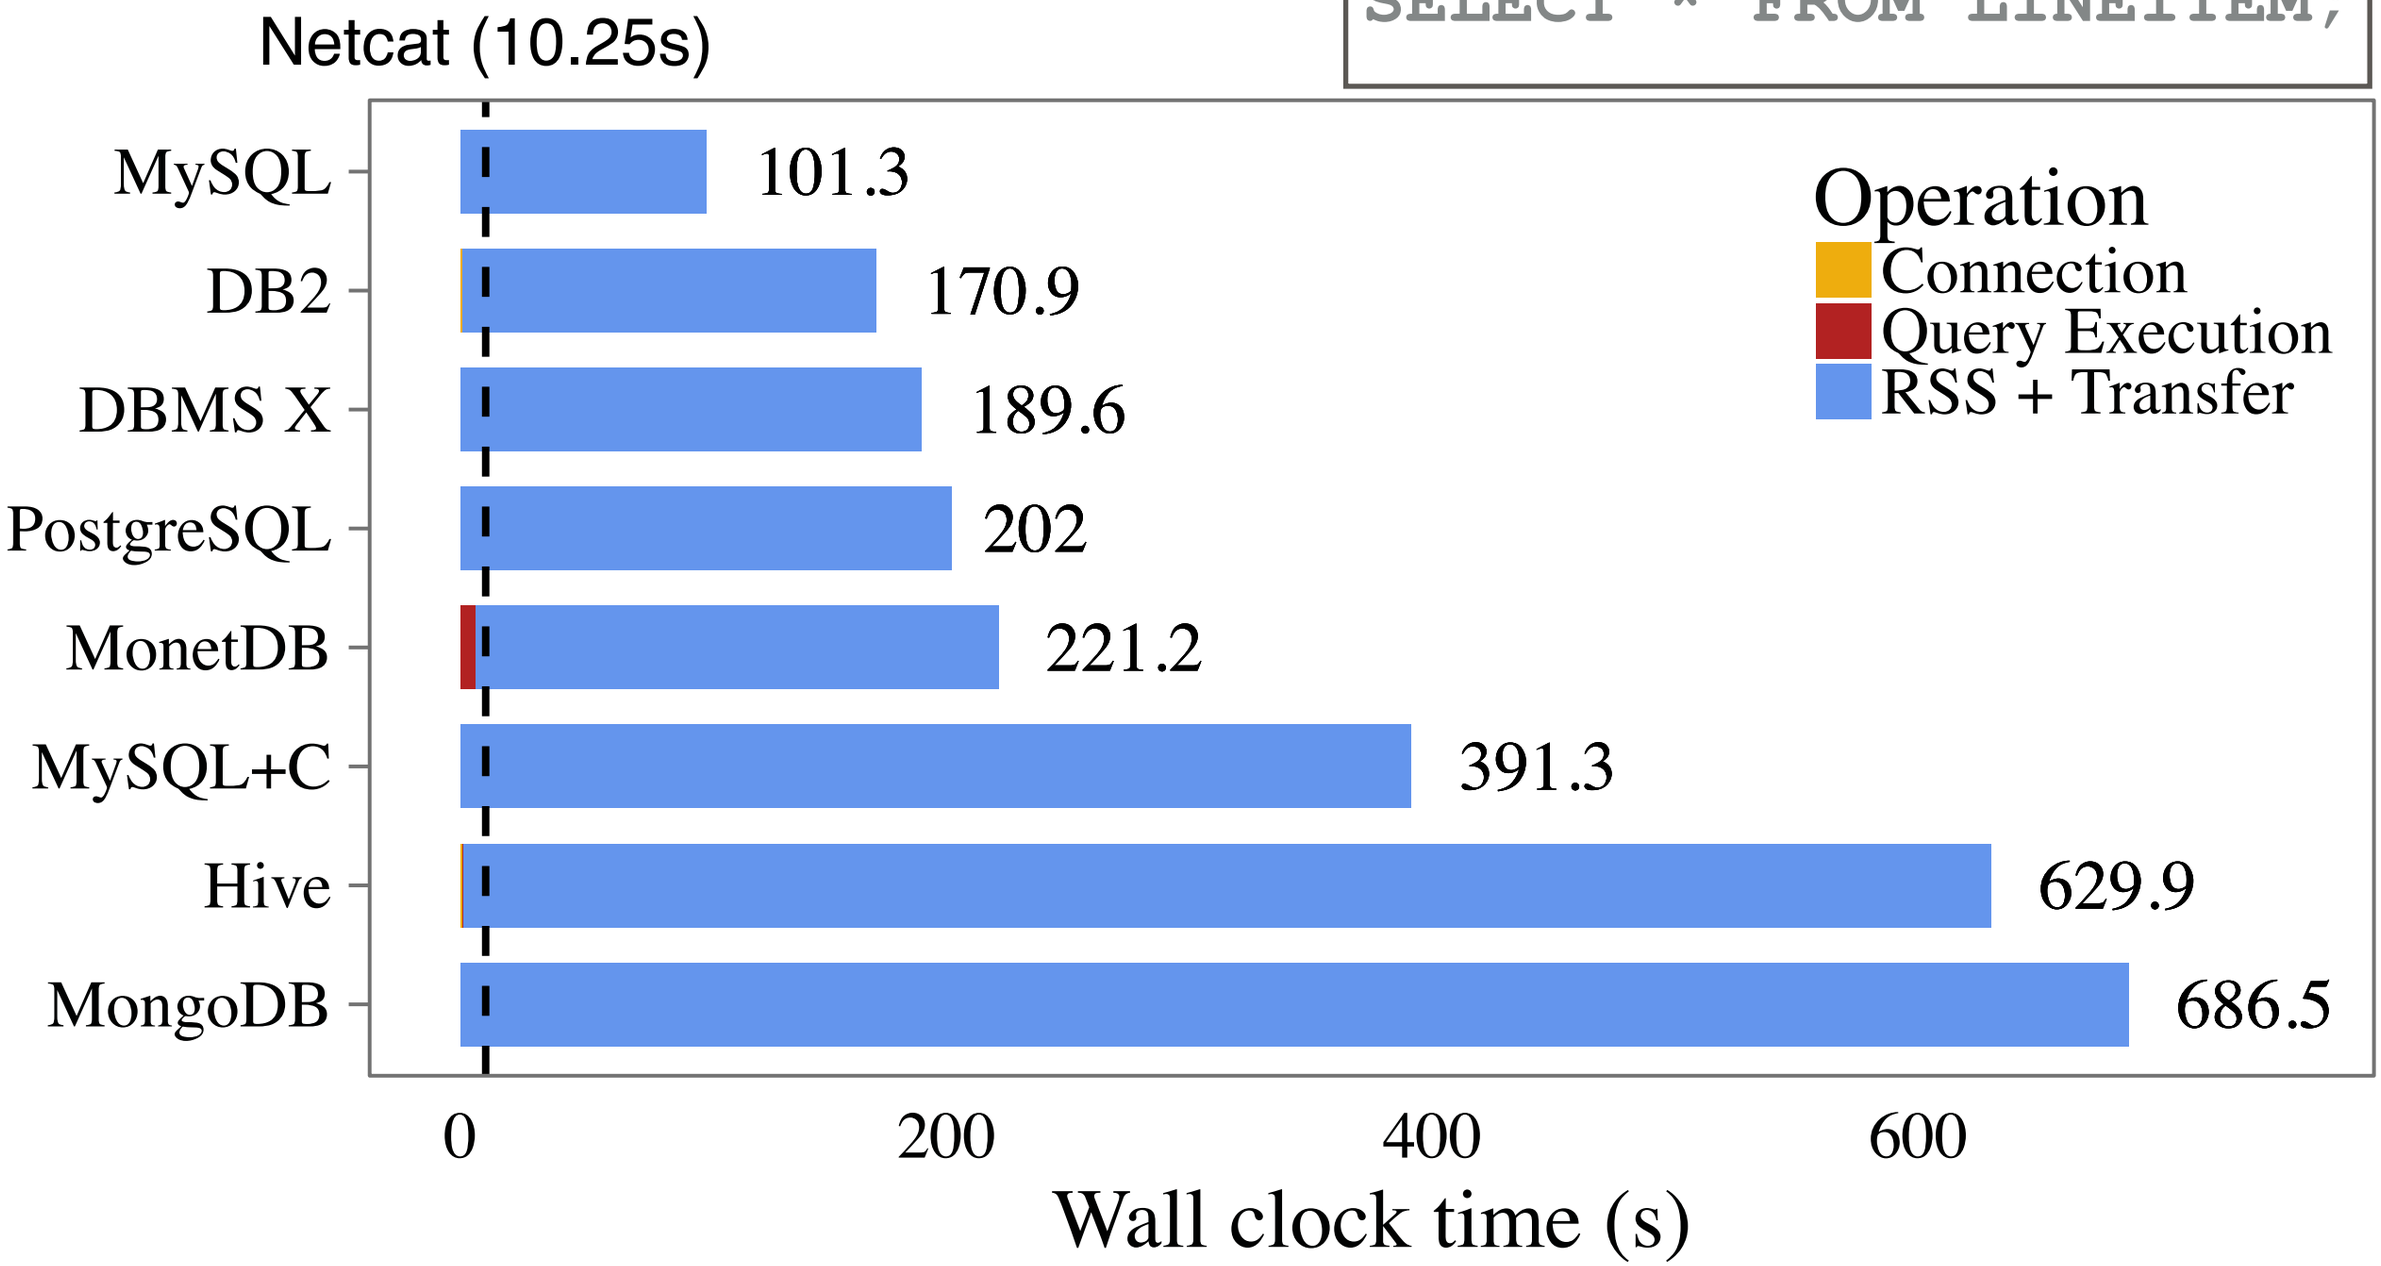

- Cost of exporting the SF10 lineitem table from TPC-H (7.2GB in CSV format) on localhost

Motivation

- We are not the first ones to notice this problem
- A lot of work on in-database processing, UDFs, etc.
- However, that work is database-specific, requires adapting of existing work flows and introduces safety issues

- Why is exporting large amounts of data from a database so inefficient?
- Can we make it more efficient?

State of the Art Protocols

- We reverse engineered how different databases transfer the following table “on the wire”
  - Source code/documentation
  - Decompilation of JDBC Drivers
  - Wireshark

INT32	VARCHAR10
100,000,000	OK
NULL	DPFKG

State of the Art Protocols

PostgreSQL

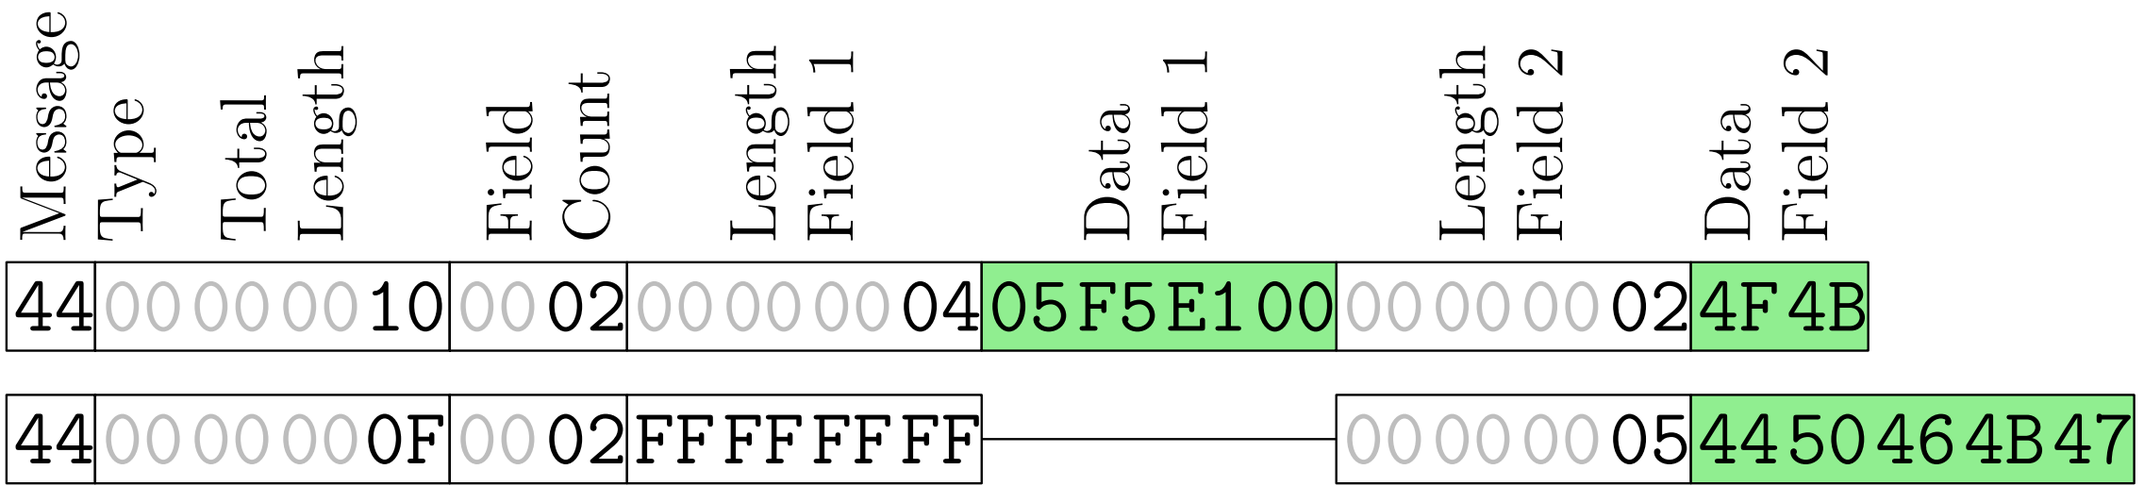

- Significant per-row overhead
- Used by many other systems:
  - Redshift, HyPer, Greenplum and Vertica

INT32	VARCHAR10
100,000,000	OK
NULL	DPFKG

State of the Art Protocols

MySQL

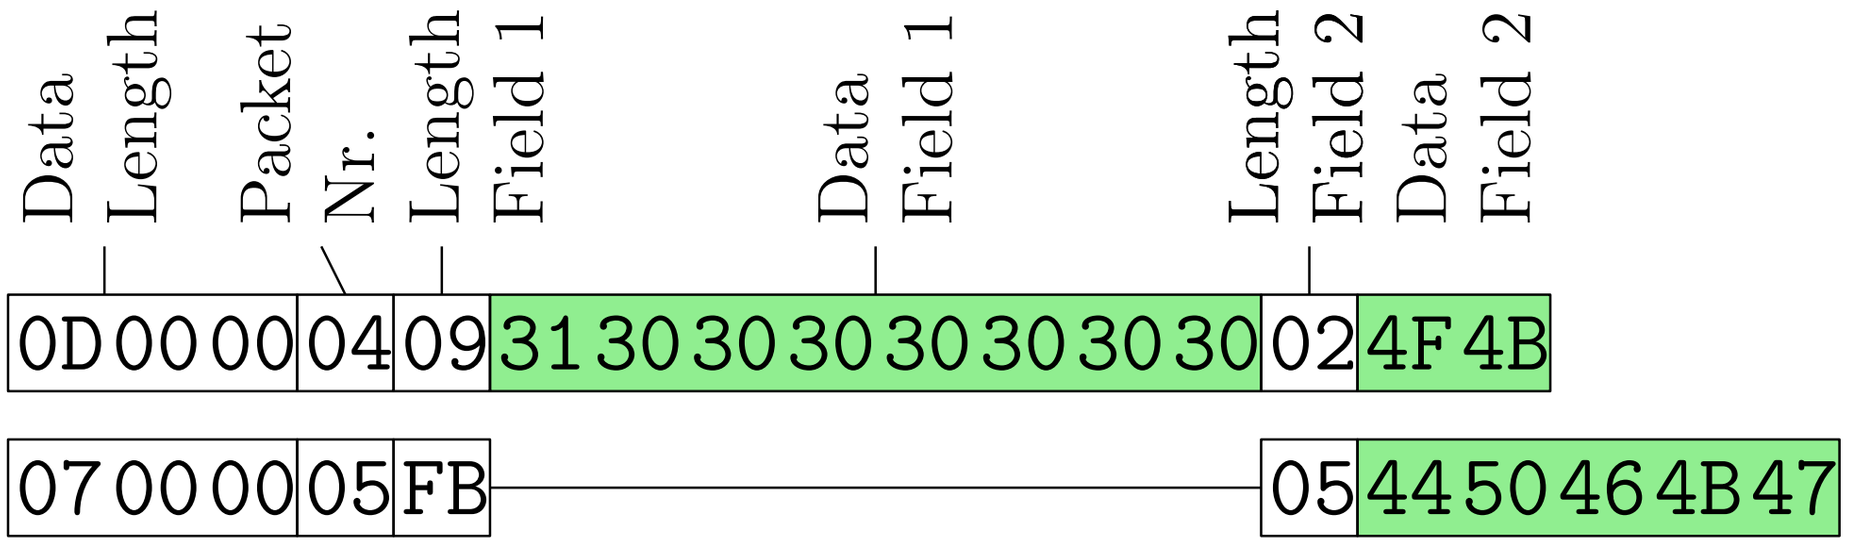

- ASCII protocol
- Every value has a length field
- Supports compression with GZIP

INT32	VARCHAR10
100,000,000	OK
NULL	DPFKG

State of the Art Protocols

Hive

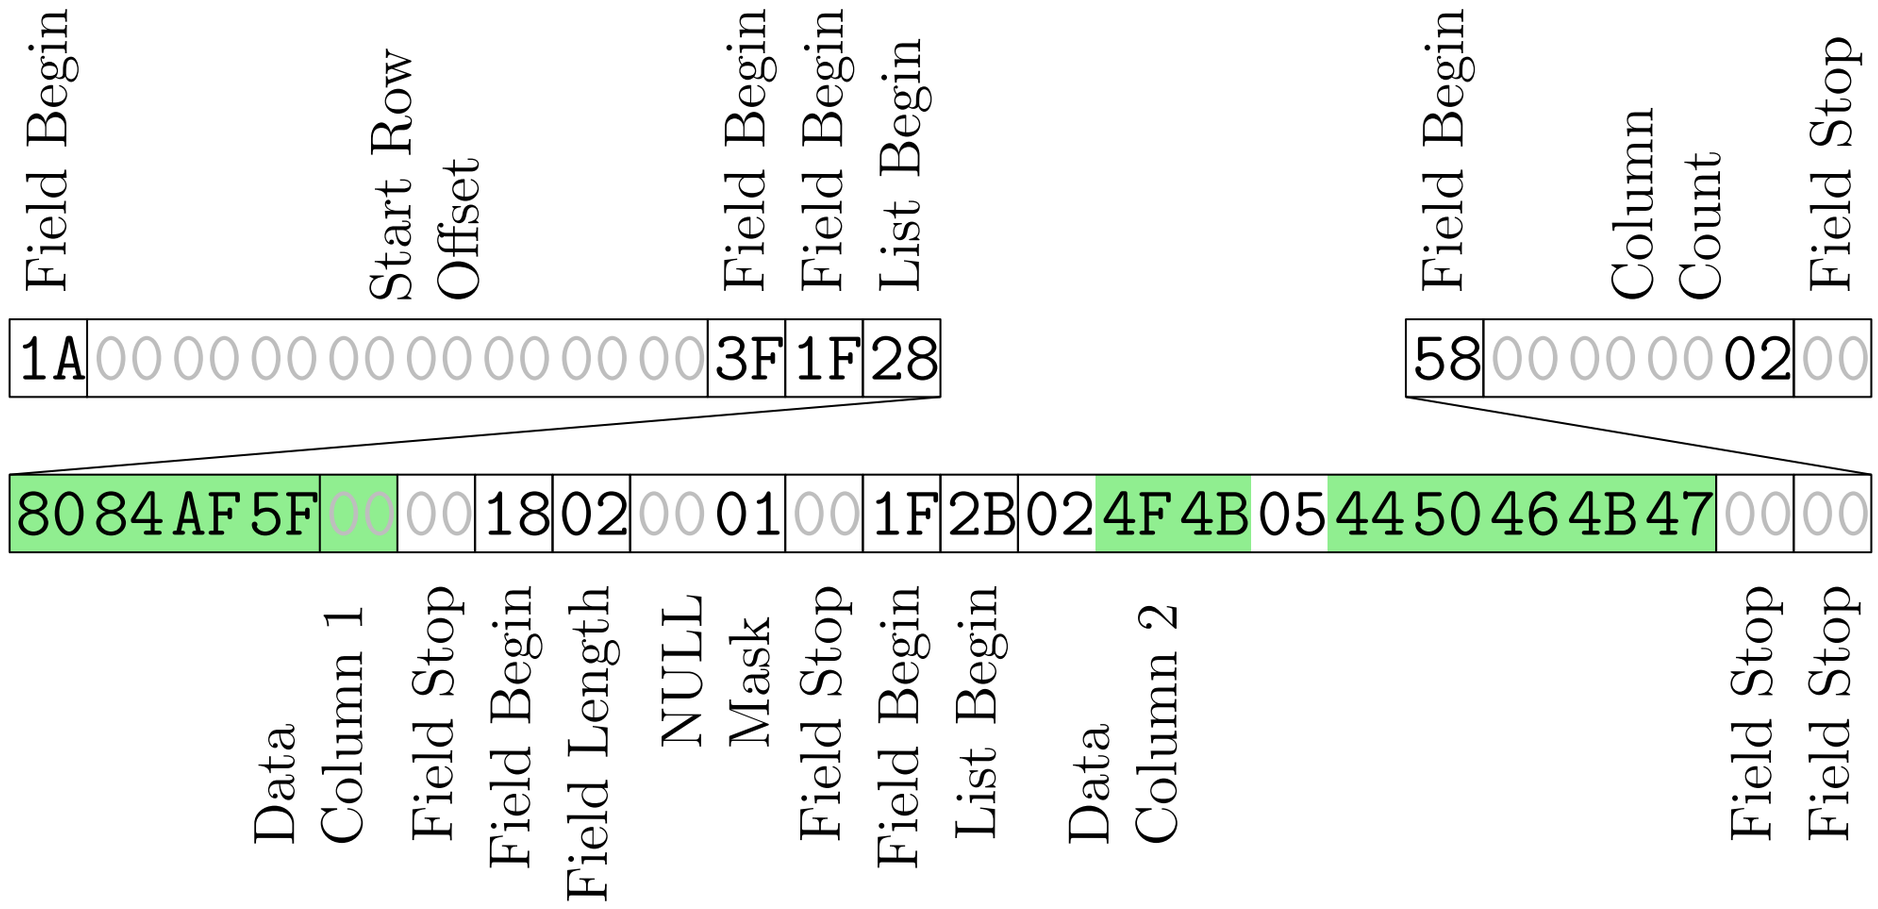

- Columnar protocol
- Uses generic Thrift serialisation library
- One byte per value in NULL mask
- Also used by SparkSQL

INT32	VARCHAR10
100,000,000	OK
NULL	DPFKG

Protocol Design Space

- Implemented prototype in PostgreSQL and MonetDB

- Serialisation Format (ASCII, Custom Binary, Generic)
  - **Custom Binary**
- Row Major or Column Major
  - **Column-Major (but chunked)**
- Data Compression Methods
  - **No compression local, stream compression on remote**
- Null Handling
  - **Close to native formats**

Proposed Protocol in MonetDB

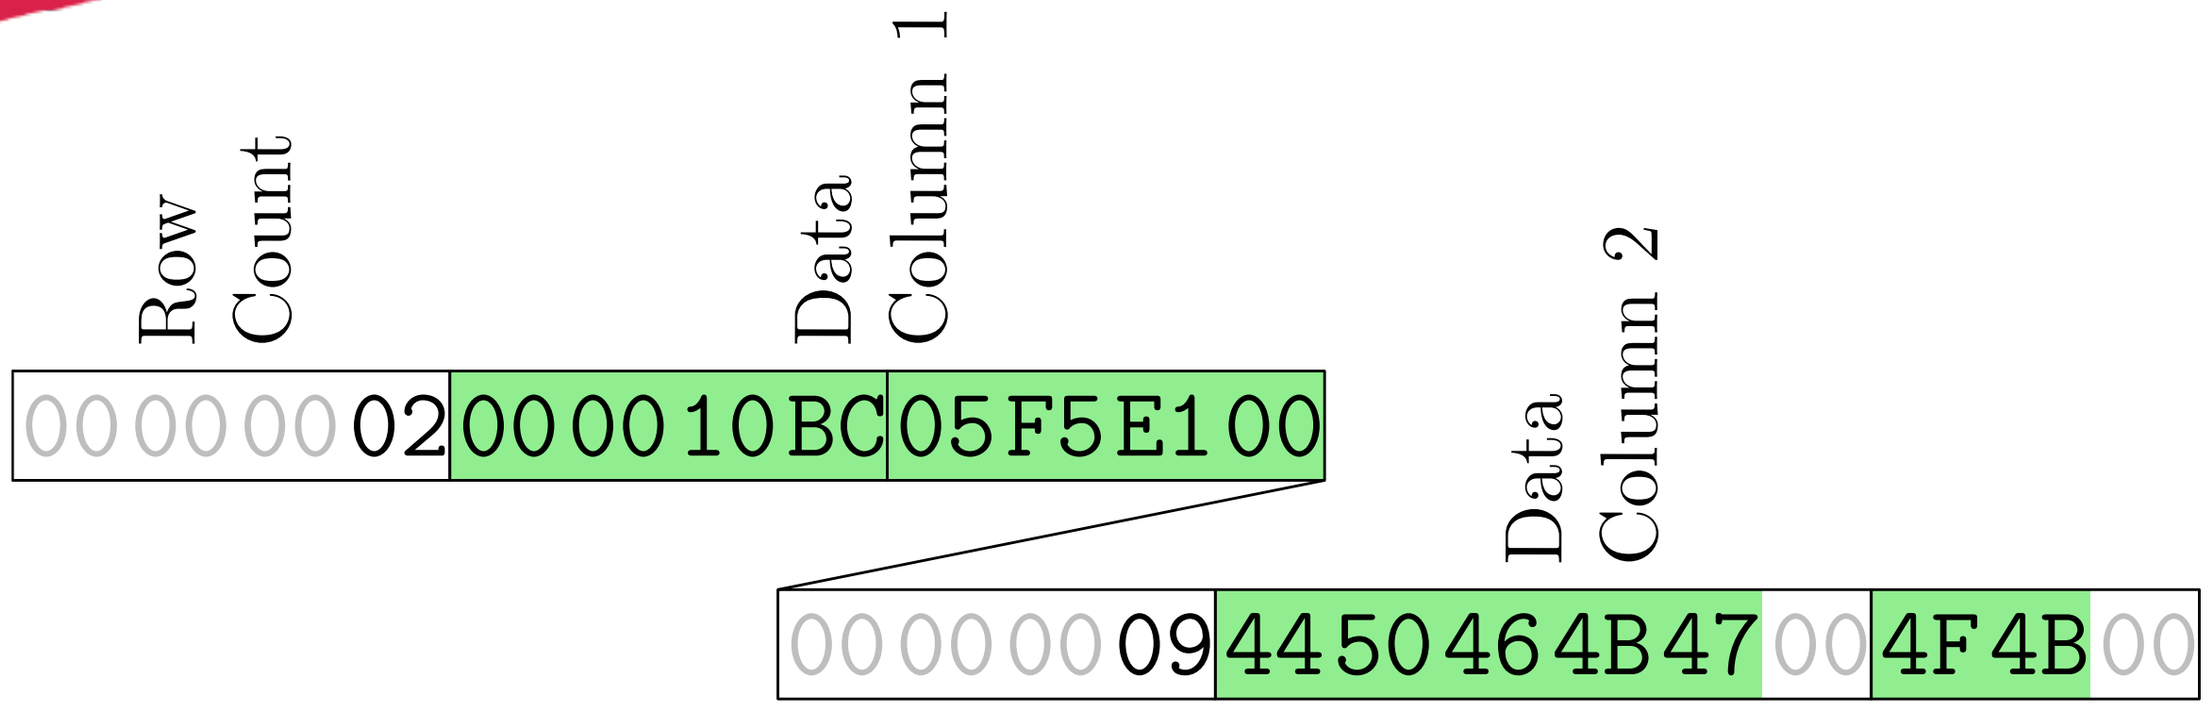

- Columnar, 1MB chunks prefixed with row count
- Missing values stored as special values in domain
- Variable-length columns prefixed with their length

INT32	VARCHAR10
100,000,000	OK
NULL	DPFKG

Proposed Protocol in PostgreSQL

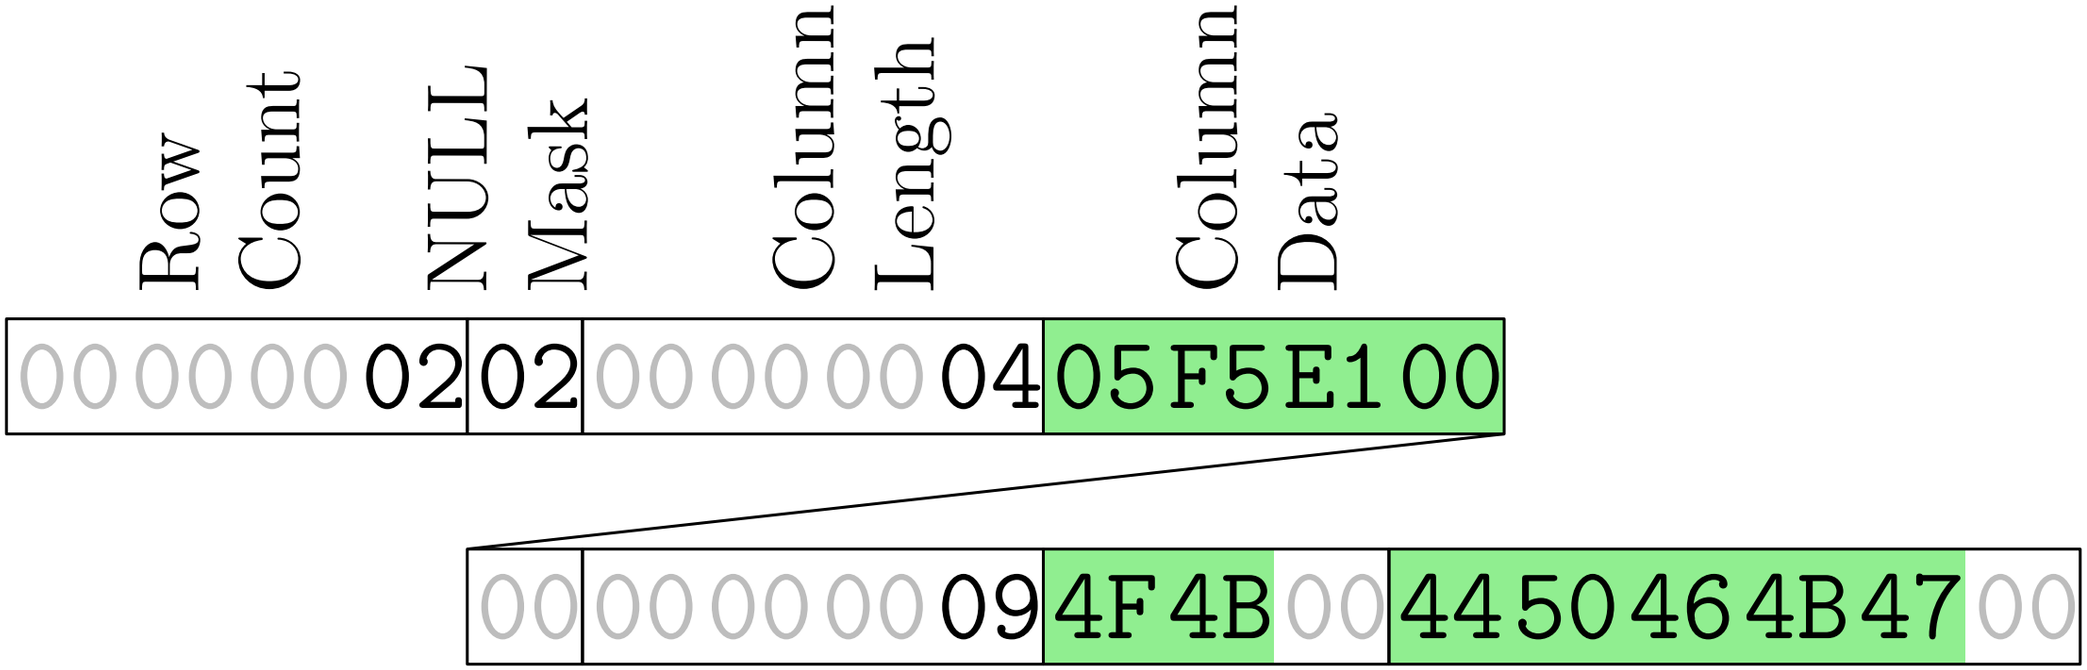

- Missing values (PostgreSQL)
- NULL bitmask for each column, 1 bit per value
- Only add mask if column has missing values
- Columns with mask have a column-length as well

INT32	VARCHAR10
100,000,000	OK
NULL	DPFKG

Benchmarks

- Three different network configurations
  - **Localhost:** No network restrictions
  - **LAN:** 1000 Mb/s throughput, 0.3ms latency
  - **WAN:** 100 Mb/s throughput, 25ms latency
- **Lineitem:** SF10, 60 million rows, 16 columns, 7.2GB in CSV format
- 1 hour timeout

Benchmarks

Localhost (No network restrictions)

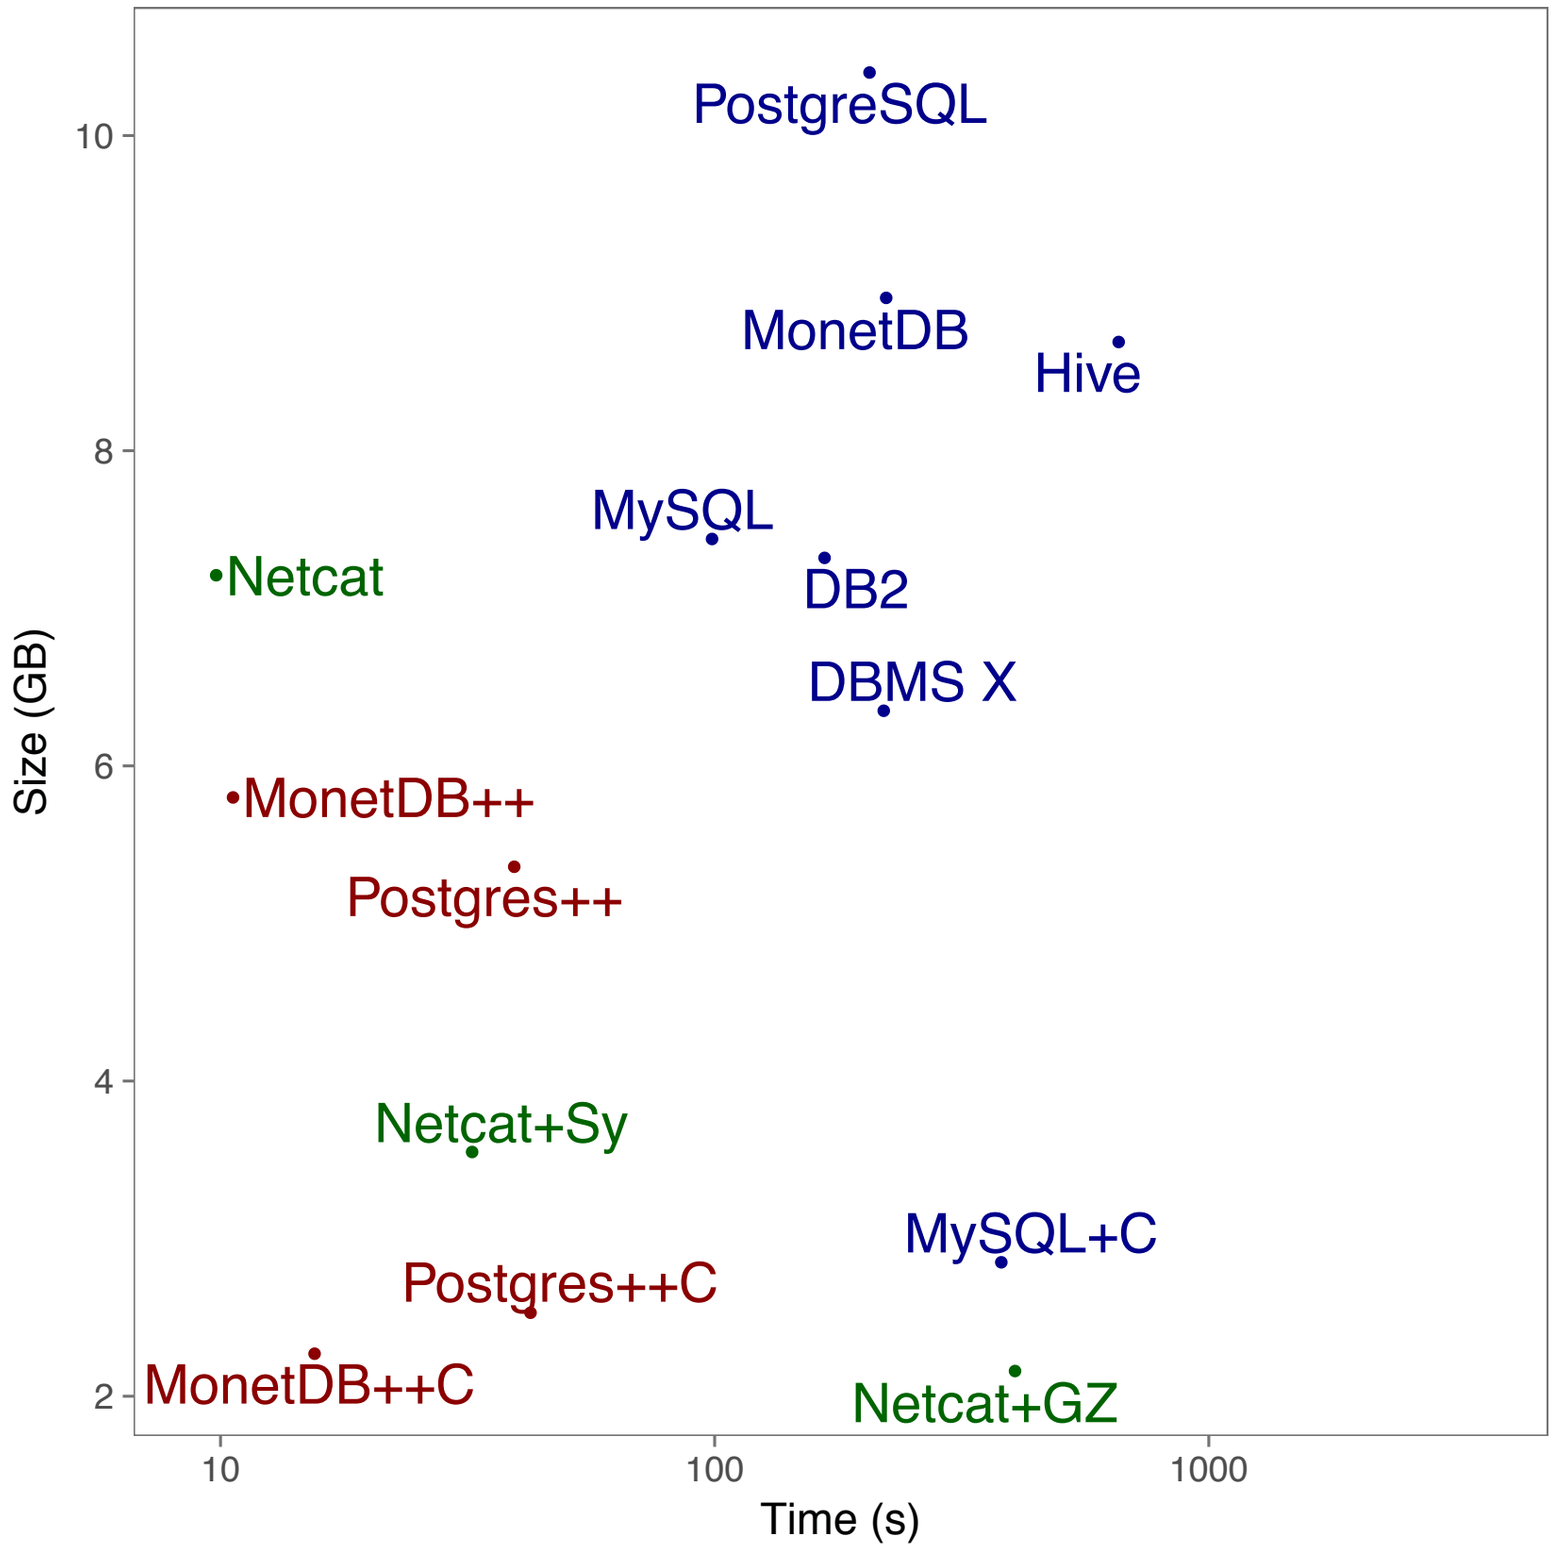

Benchmarks

LAN (1000 Mb/s throughput, 0.3ms latency)

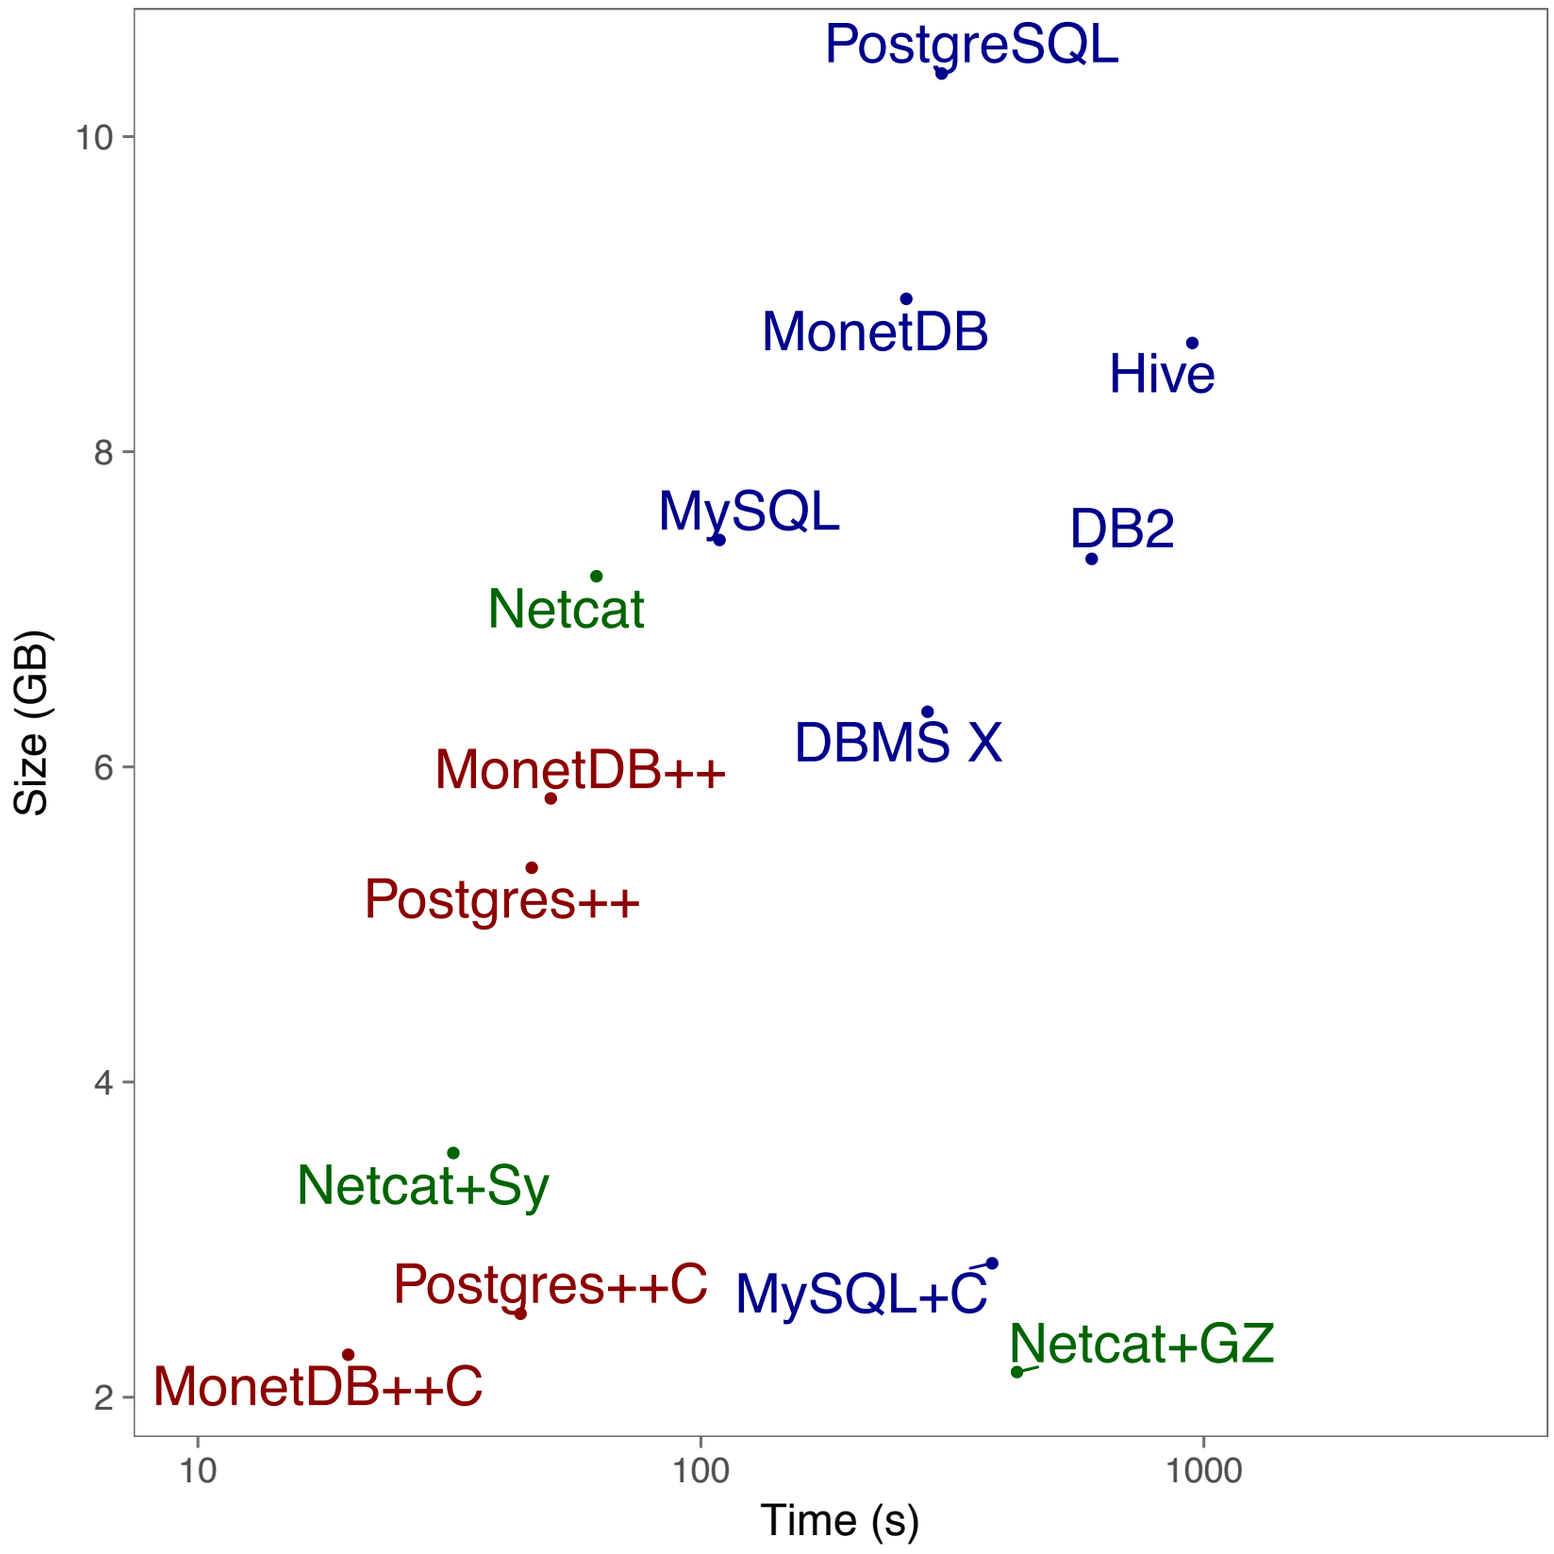

Benchmarks

WAN (100 Mb/s throughput, 25ms latency)

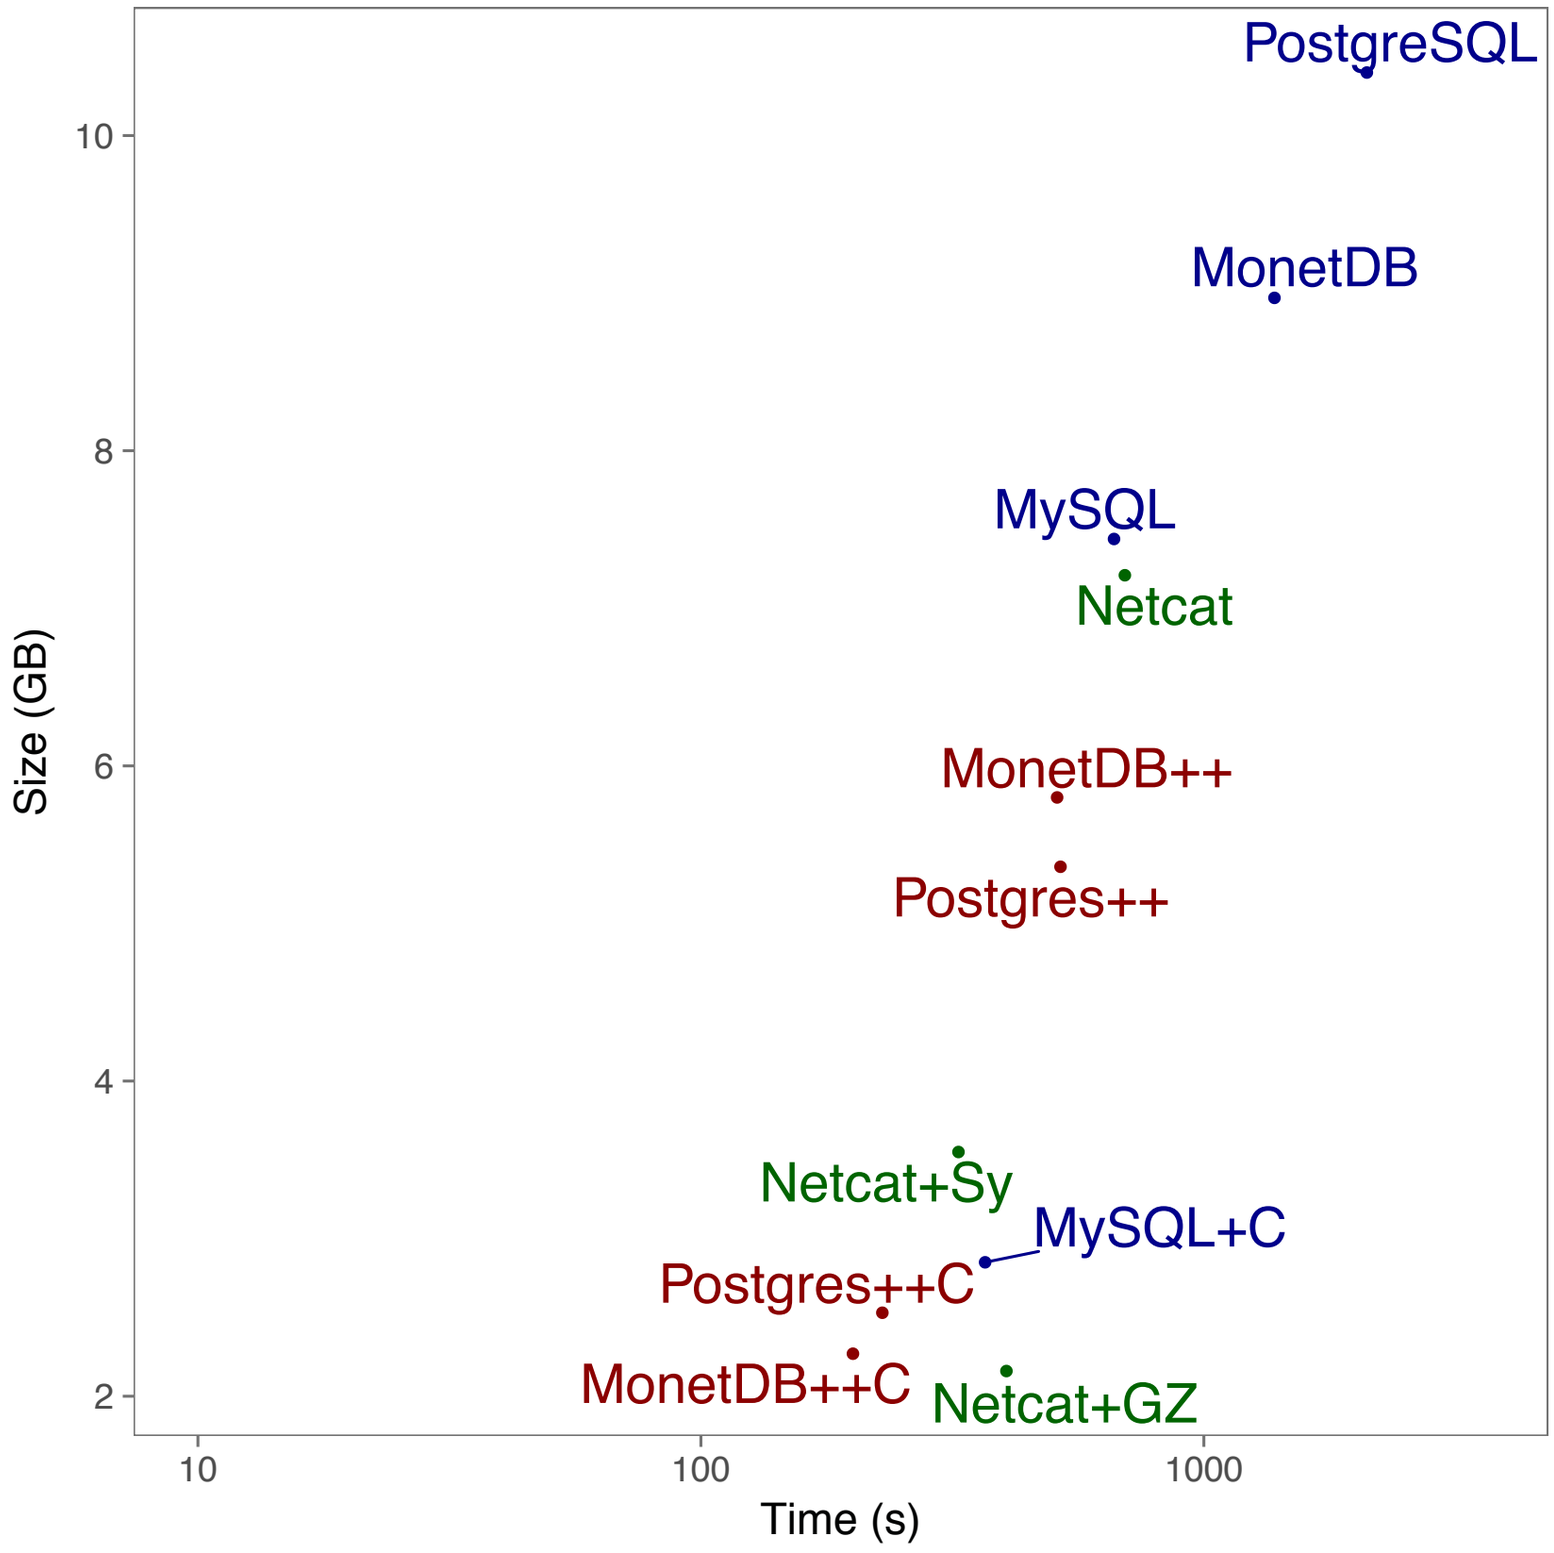

Conclusions

- Exporting data from a database does not have to be so inefficient
- State of the art database protocols can be improved for this use case
- We show this by implementing prototypes in two databases (MonetDB and PostgreSQL)
  - Avoid per-row overhead, bulk transfer
  - Stay close to database native formats
  - Avoid unnecessary copying and conversion
  - Lightweight compression on remote

- MonetDB implementation is already released.
- Benchmark information: https://goo.gl/usjfyJ