

SYRCoDIS 2011


Development of a Parallel DBMS on the Basis of PostgreSQL

Constantin Pan

South Ural State University

Moscow, 2011

This work is supported by the Russian Foundation for Basic Research (grant No. 09-07-00241-a).

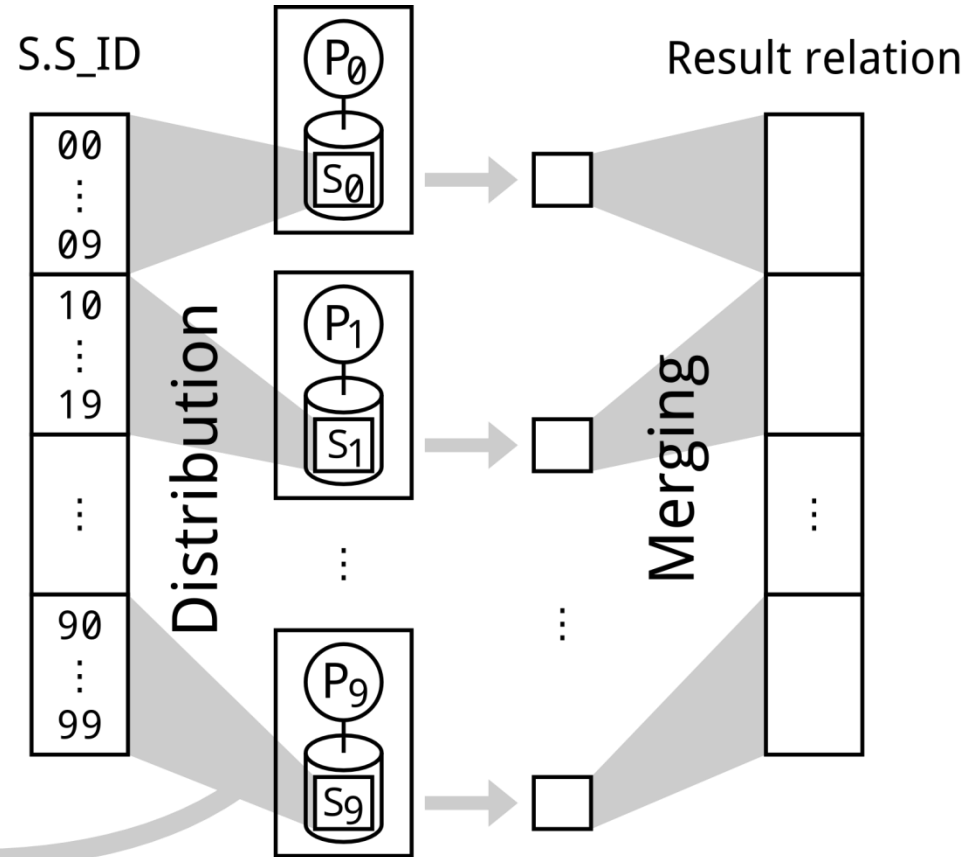


PargreSQL


Partitioned Parallelism

$$S_i = \{t | t \in S, \phi(t) = i\}$$
$$i = 0, \dots, 9$$

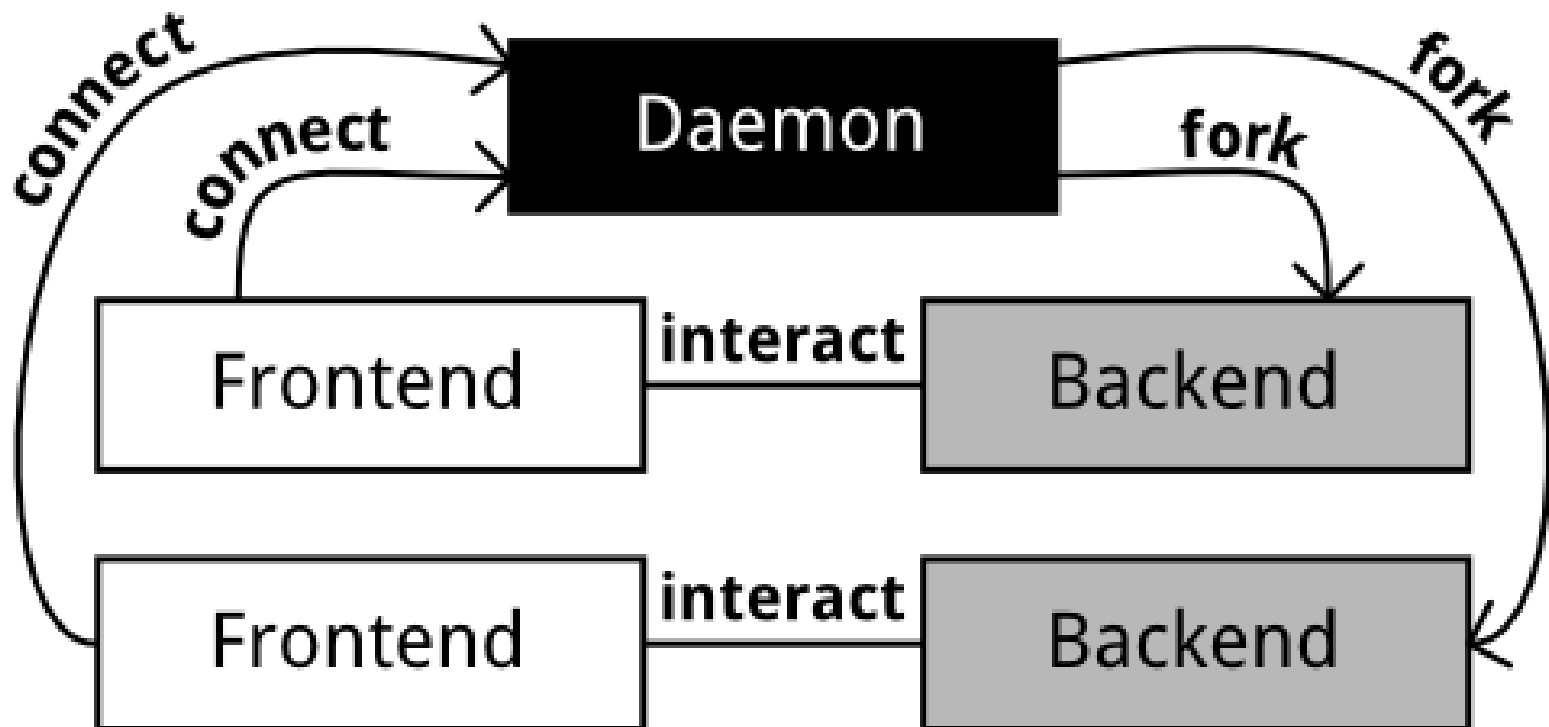
Partitioning function
 $\phi(t) = (t.S_ID \text{ div } 10) \text{ mod } 10$



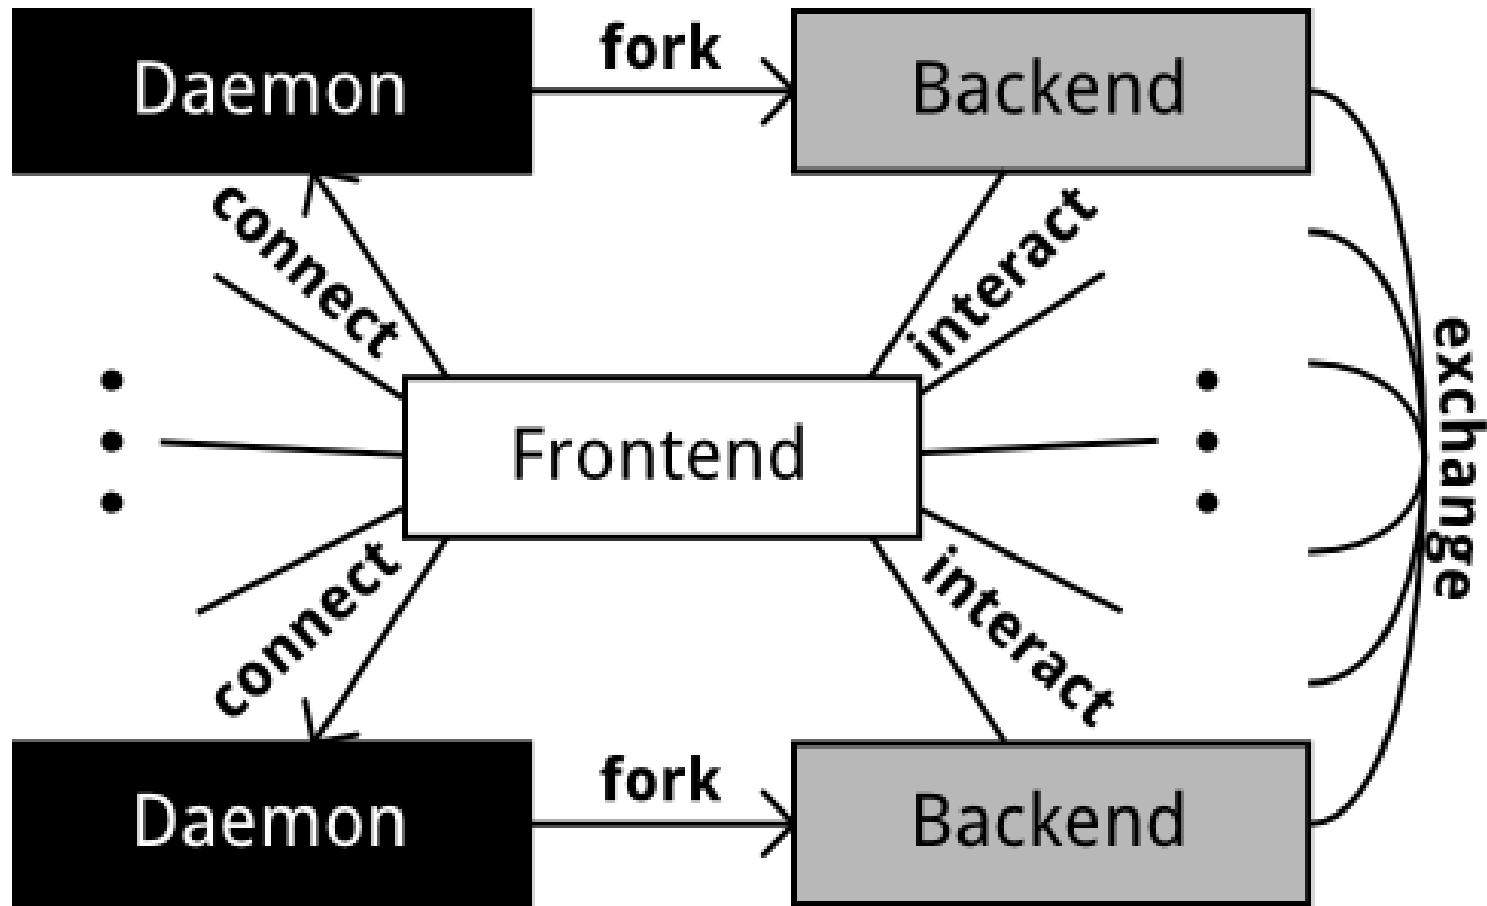
Architecture



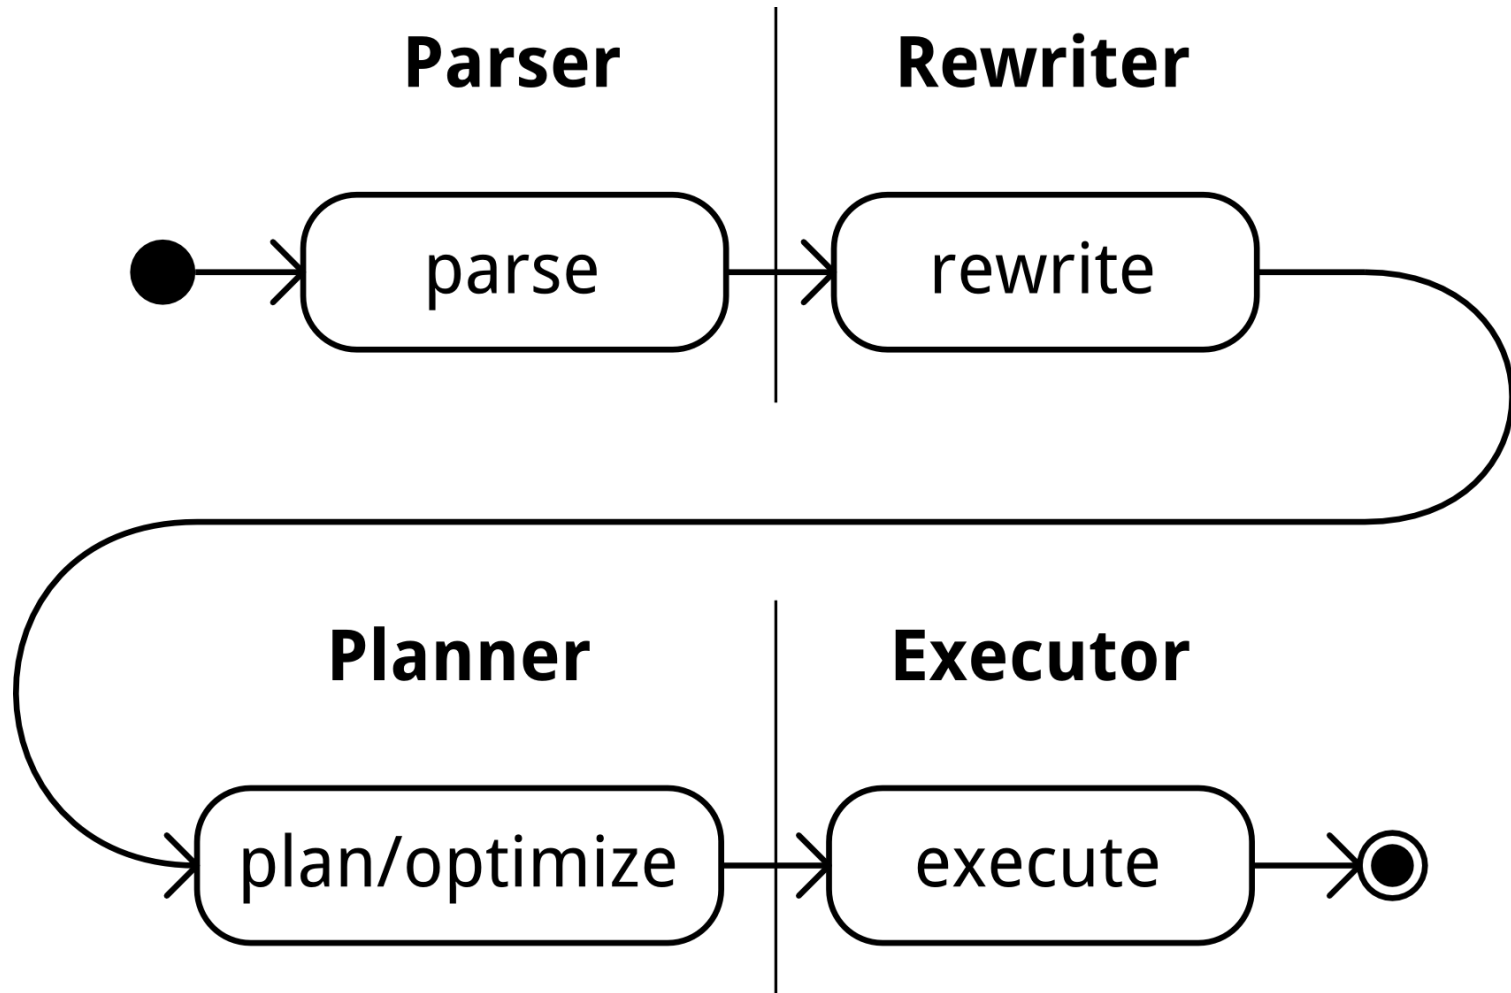
Client-server model



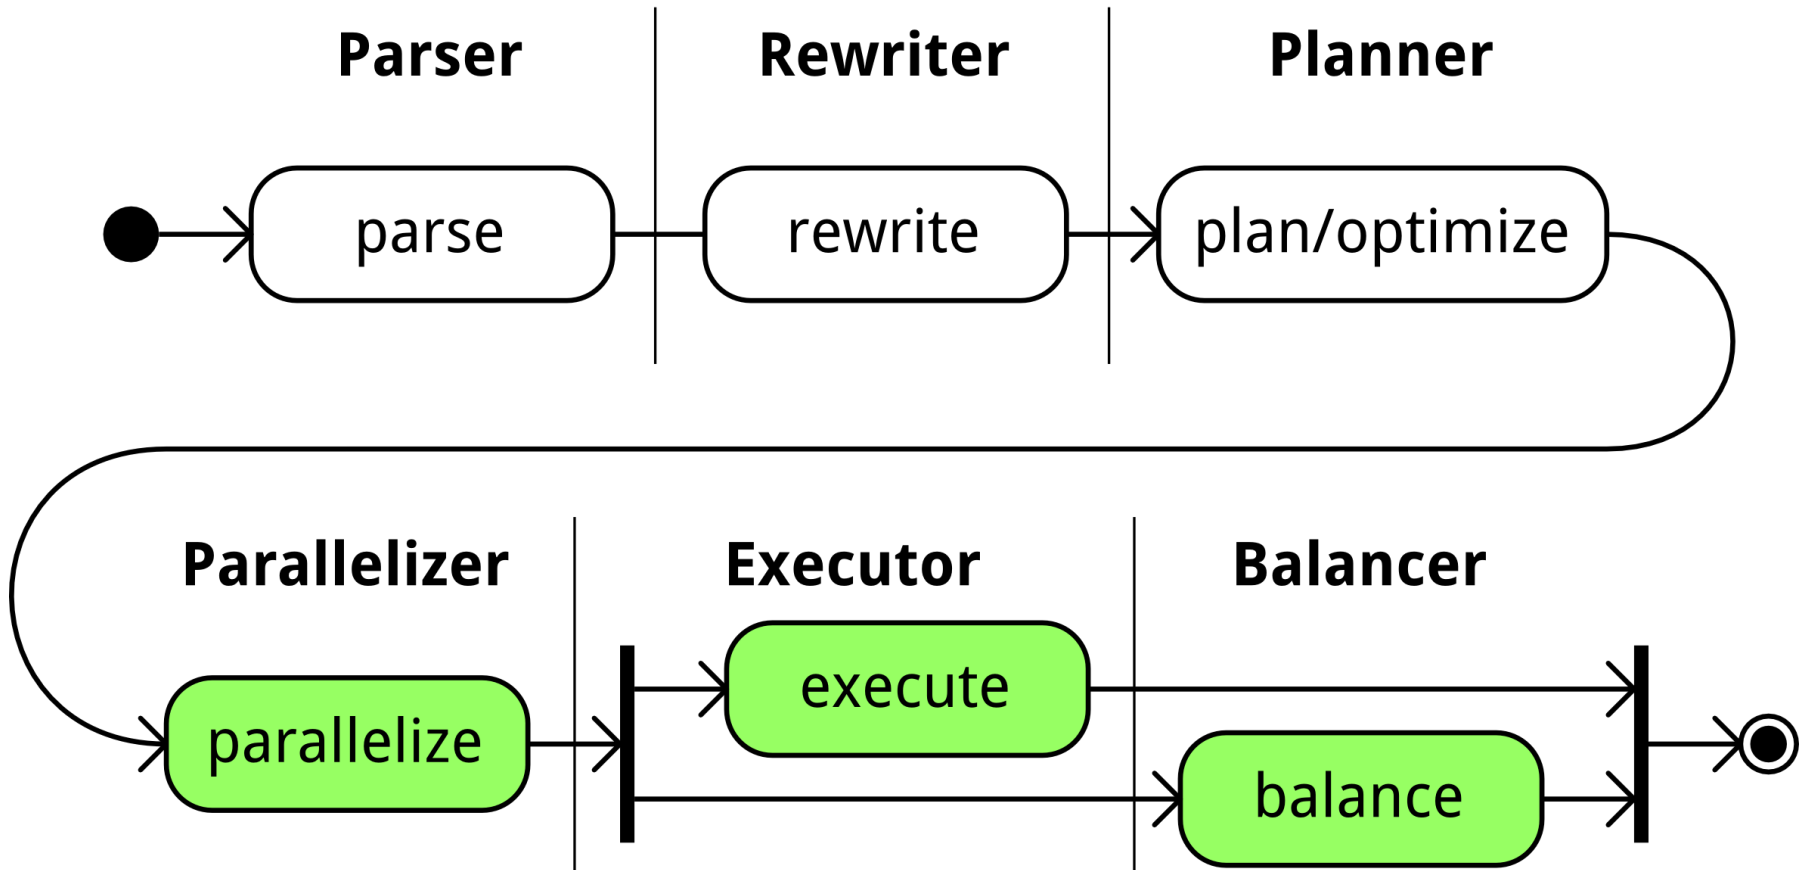
Client-server model (PargreSQL)



Query Execution




Query Execution (PargreSQL)




Exchange Operator

Node 0


Node 1



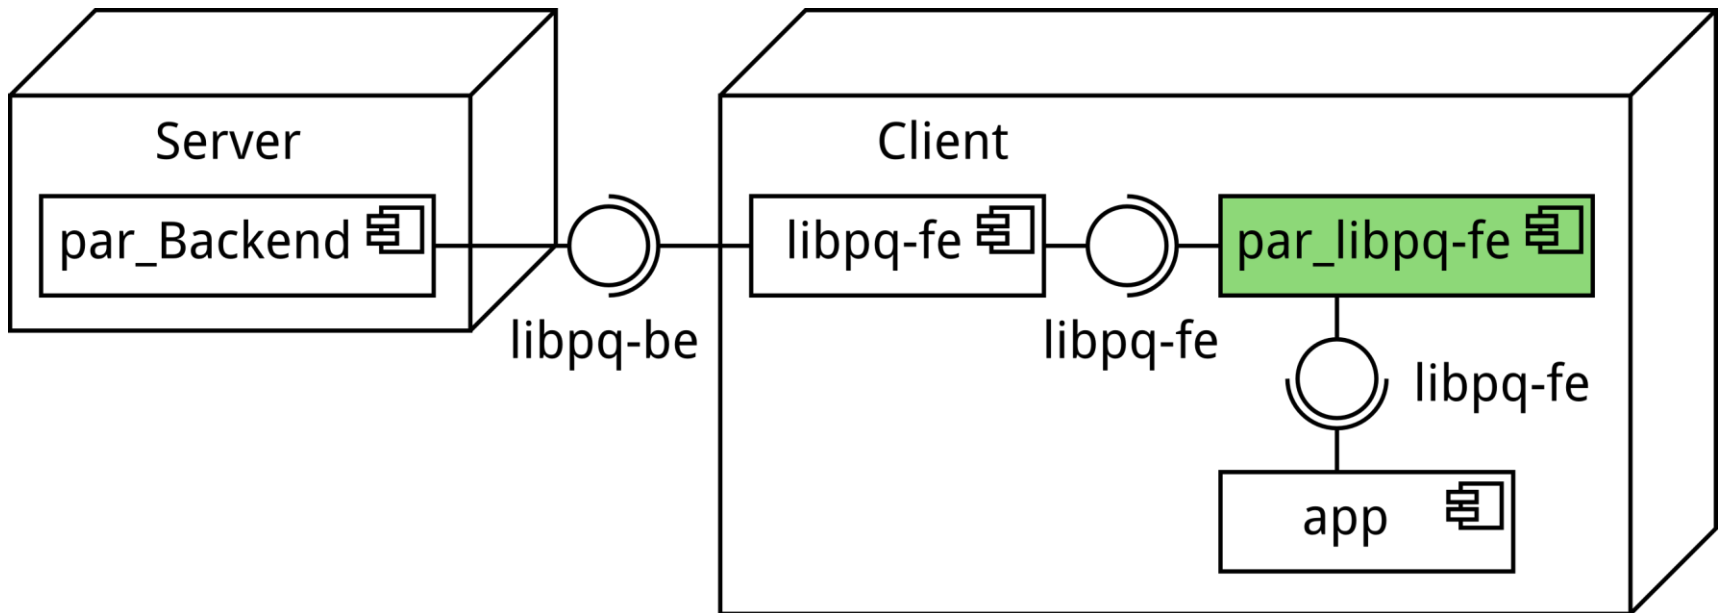
Exchange Operator Architecture



Parallelizer



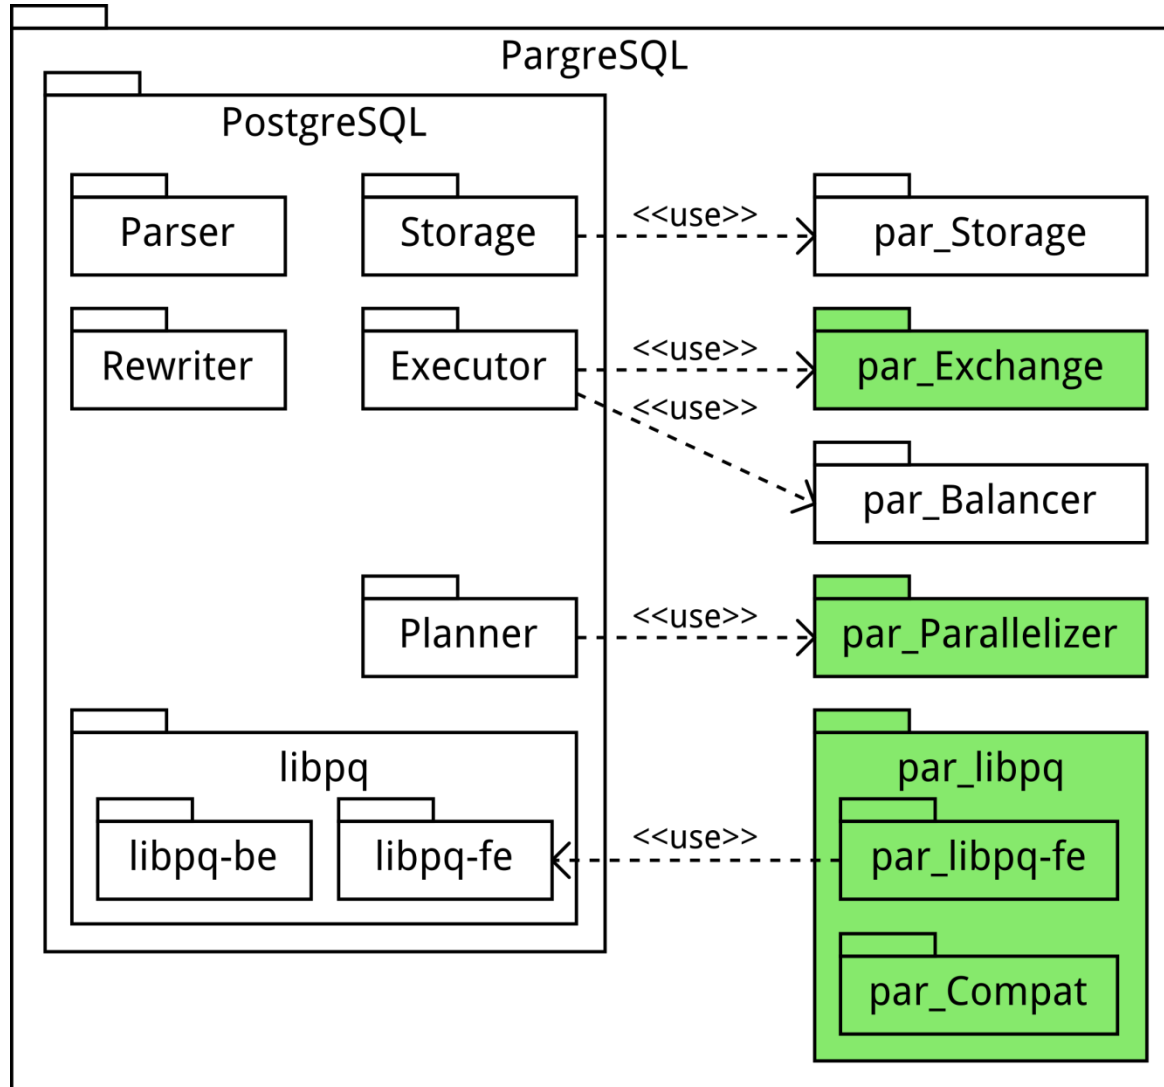
Client Side Changes



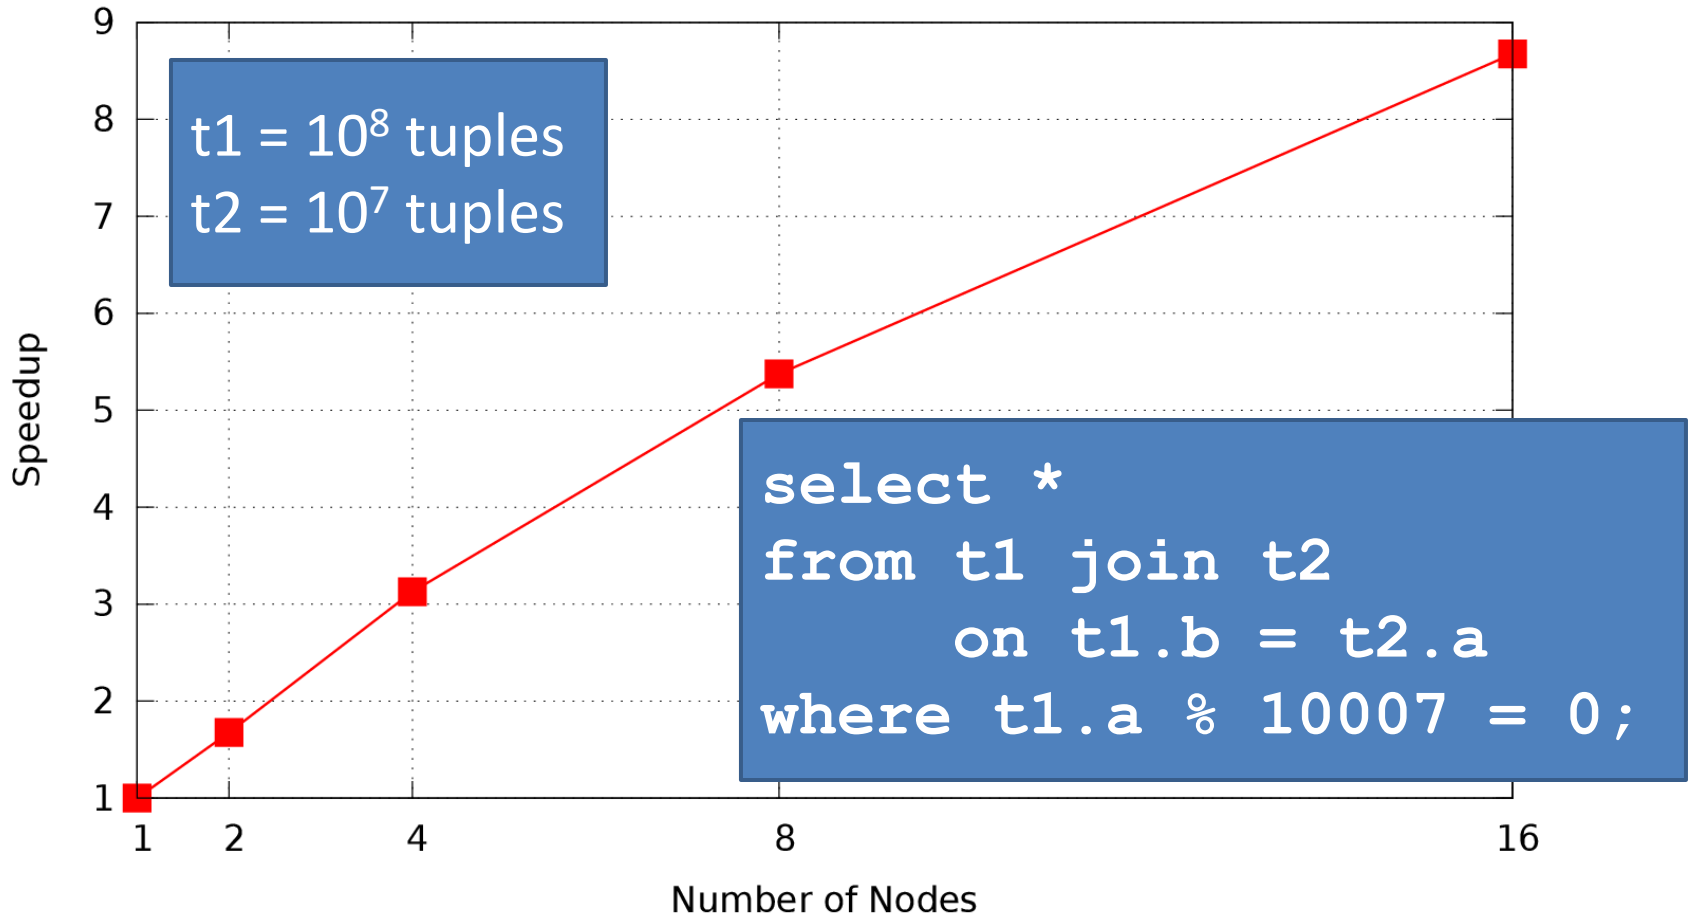
Transparency for the Application

```
#define PGconn par_PGconn
#define PQconnectdb(X) par_PQconnectdb()
#define PQfinish(X) par_PQfinish(X)
#define PQstatus(X) par_PQstatus(X)
#define PQexec(X,Y) par_PQexec(X,Y)
```

Current Results



Experiments



Future Research Directions

- load balancing
- data updates
- transactions
- fault tolerance

Thank you

- Constantin S. Pan
 - South Ural State University
 - kvapen@gmail.com

Related Work

- Extending PostgreSQL
 - Native XML type support [10]
 - Adding data types to support HL7 medical standard [4]
 - Image-handling extensions [3]
 - Integration with the Semantic Web [8]
- Adoption for parallel processing
 - Middleware for parallel processing with PostgreSQL (ParGRES, GParGRES) [9, 5]
 - MapReduce + PostgreSQL (HadoopDB) [1]

Contribution

- Embedding partitioned parallelism into open source relational DBMS PostgreSQL