


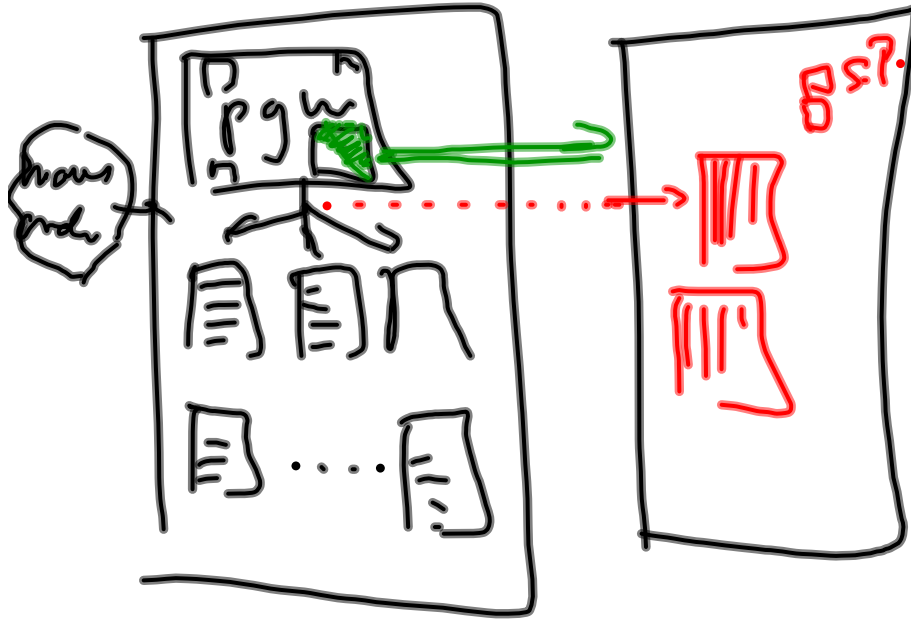
1. real data + real programs
2. data organization (→ c. store)
3. no data conversion
4. as much as possible do the calculation while tuples are in cache
5. indices help
bin. tree → prim. ind.
group by → grouping
6. SQL-execution code



7. parallel execution
8. "good mix" between memory access
and calculation
9. lock free data structures
and processors
10. in sed only

onyssyshi

IMDB



geht das wirklich nicht?

Schwa

