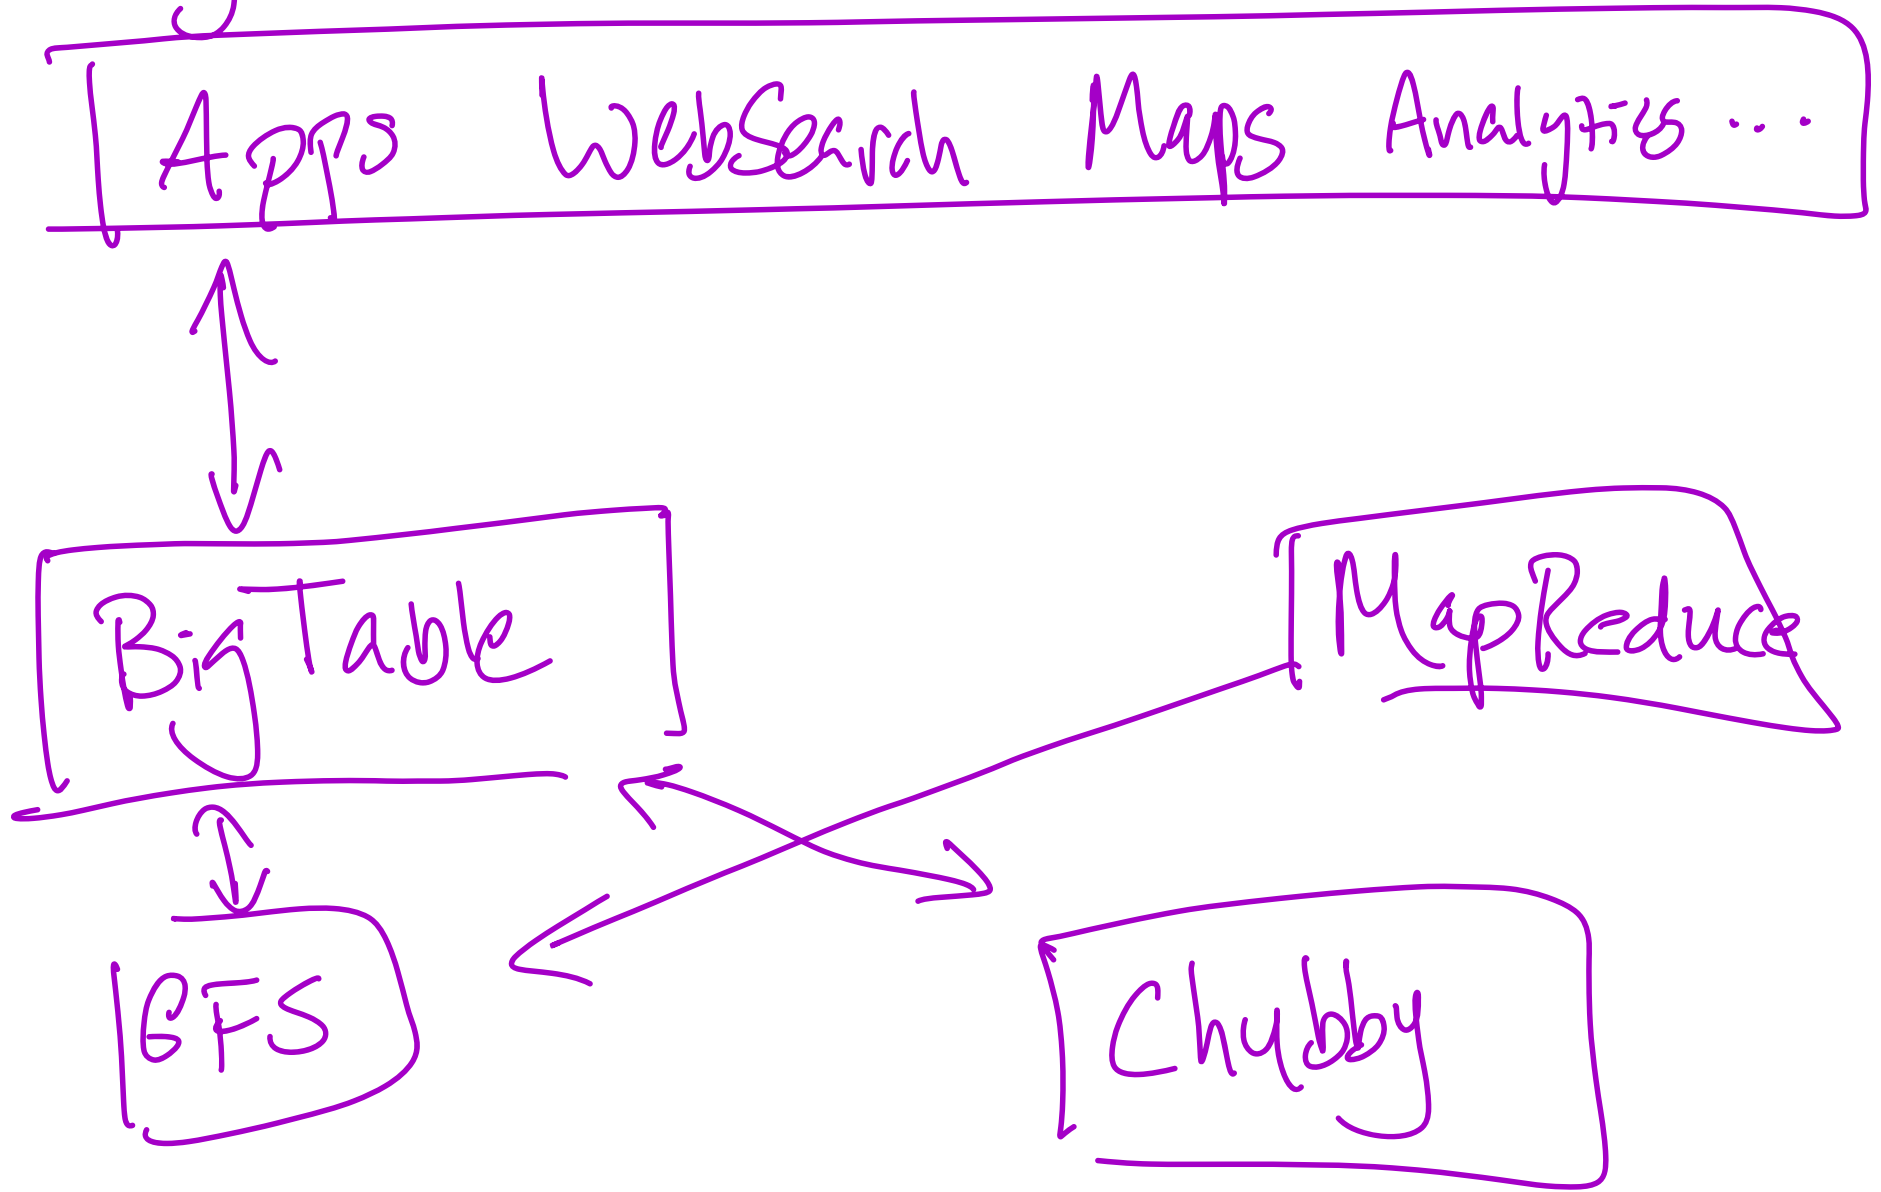


CSE 452

Distributed Systems

Big Table

Google Stack



Goals

storage system structured data

- Control over data layout
- control over locality
- non relational
- performance at scale

Data Model

(row: String, col: String, time: int(64)) → String

	contents	anchor: cnssi.com	anchor: my-book.ca
row1	t1: v1 t2: v2		
row2			
row3			

API

- Read + write rows
- Unit of atomicity is the row
- Scans of ranges of rows

not in API

~ cross row txns

Design

- SSTables

Sorted String table

immutable

- all data is stored in tablets

tablet = unit of sharding

- tablet server in charge of ops to that tablet

- only one per tablet

