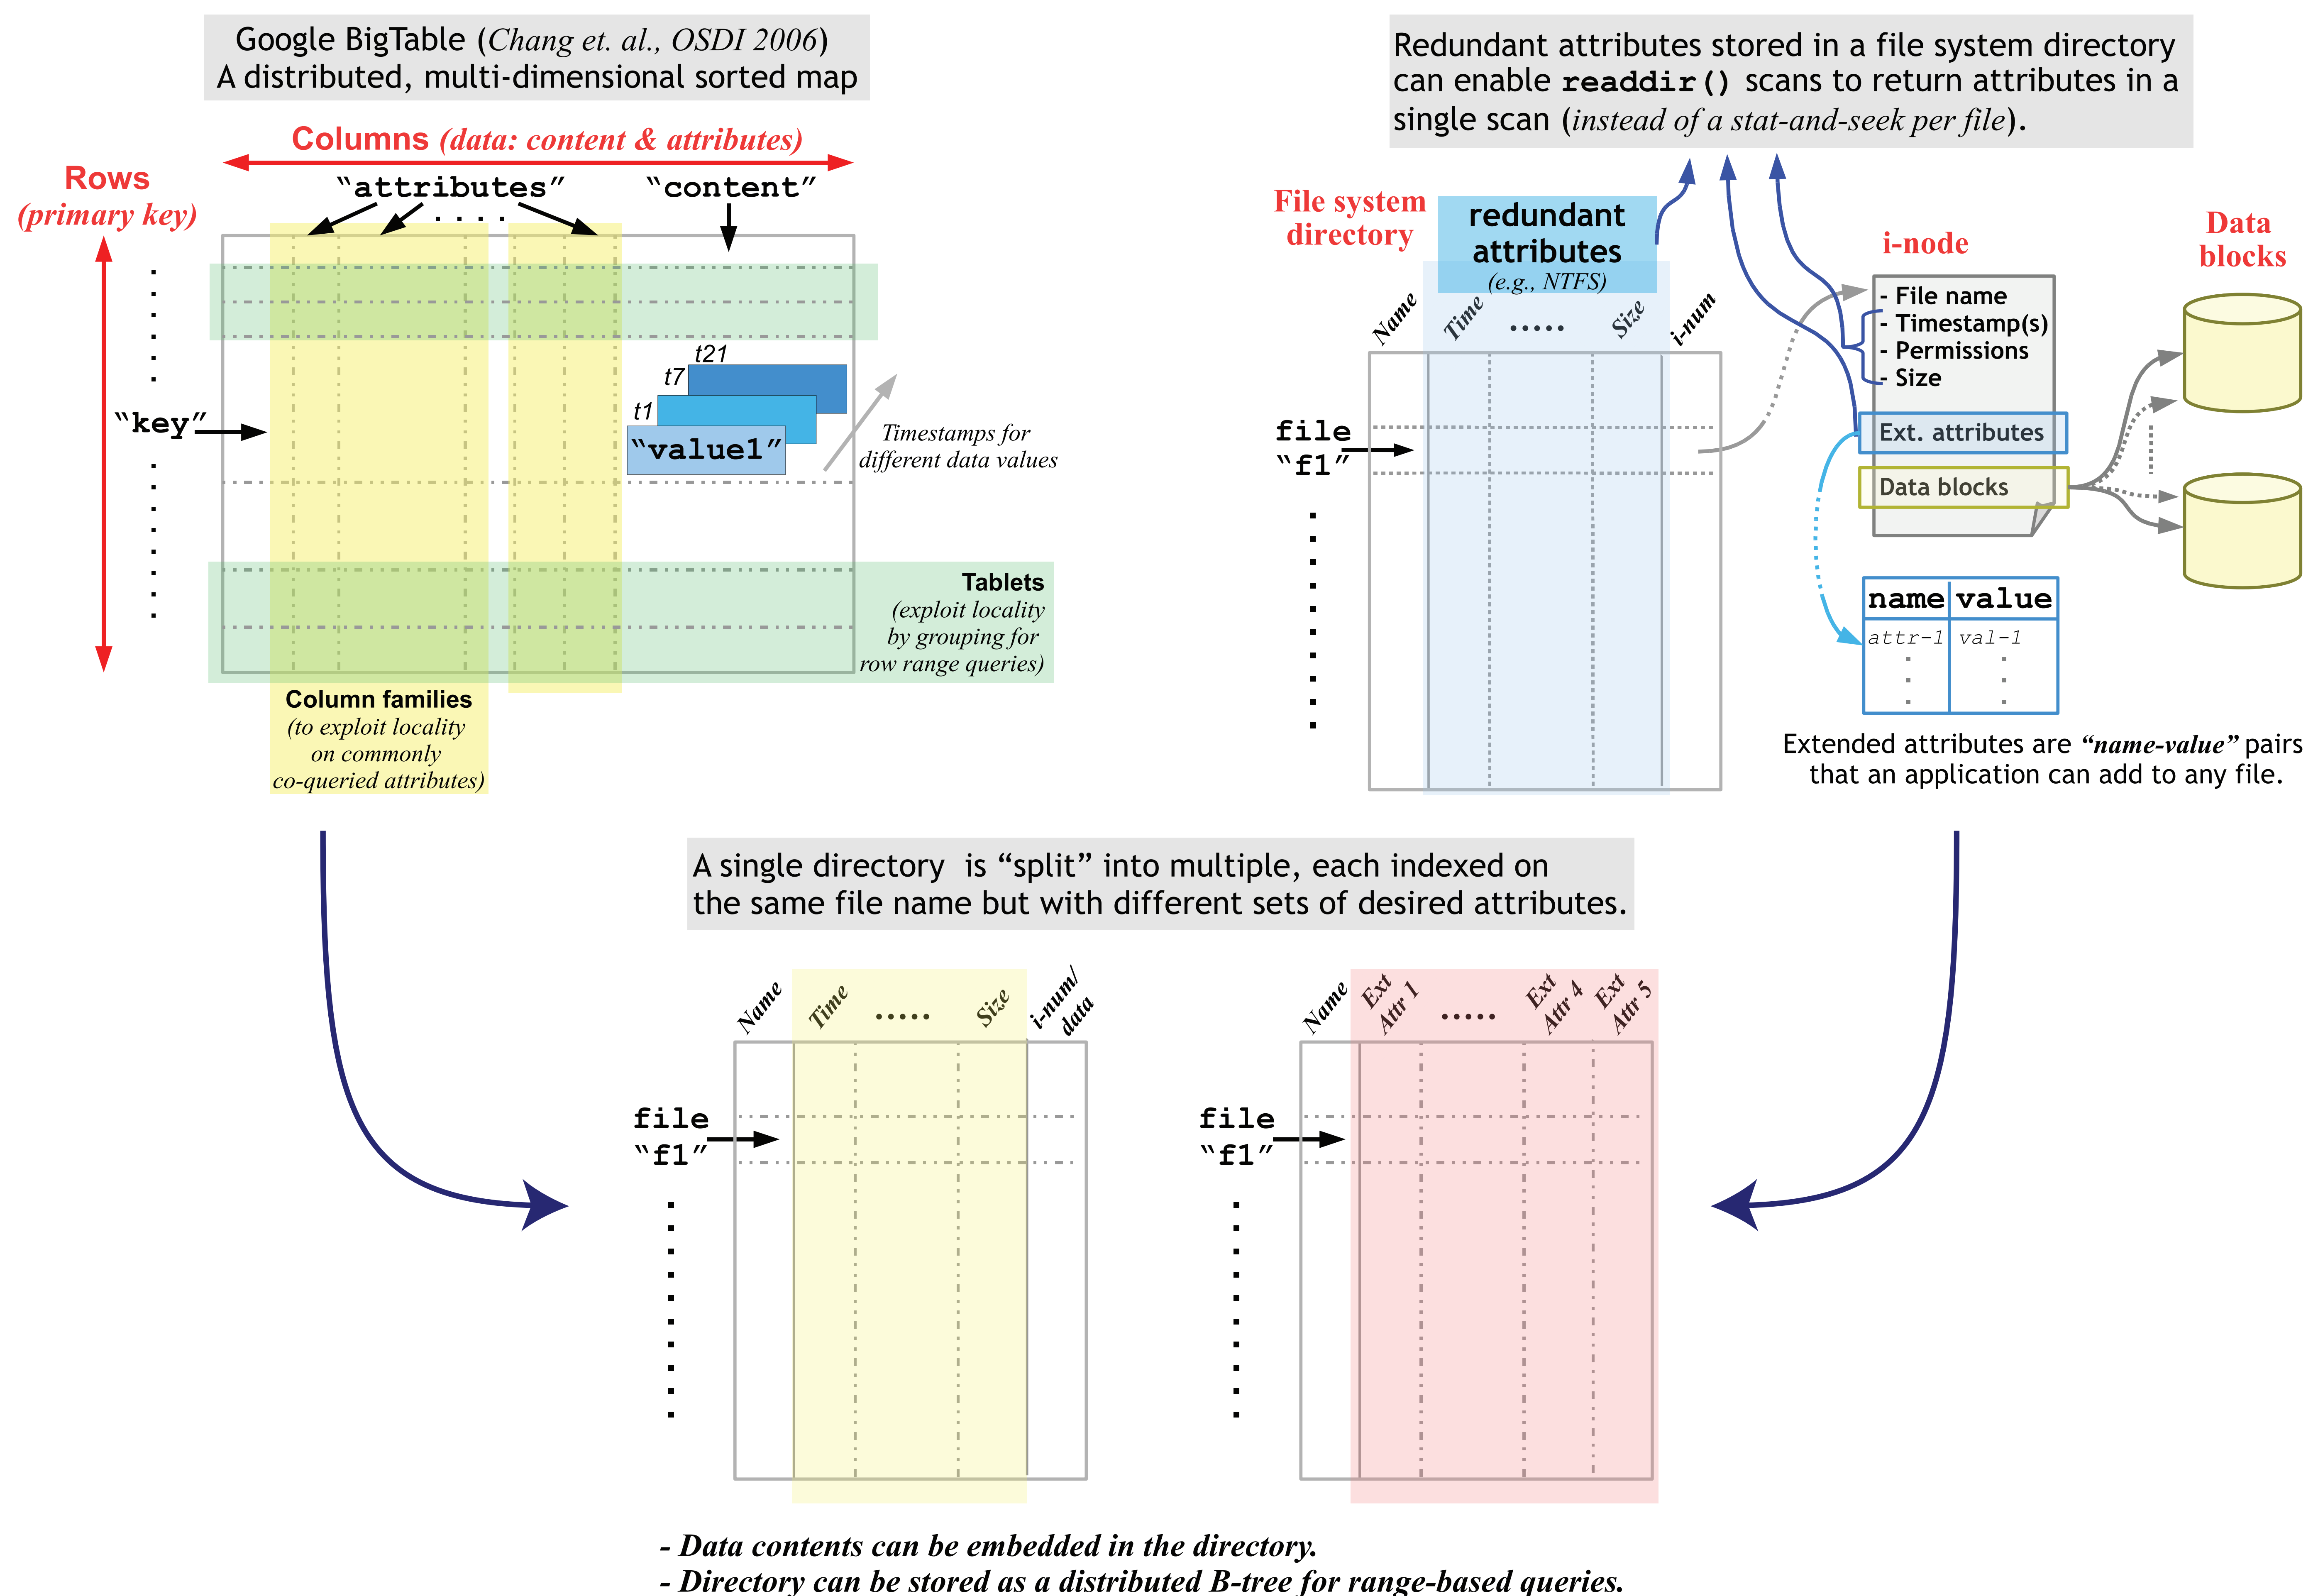


Tables and File Systems: Moving Into the Cloud

Swapnil Patil, Garth Gibson, Greg Ganger, Julio Lopez, Milo Polte, Wittawat Tantisiroj, Lin Xiao

Central Idea

- With BigTable and friends (HBase, Hypertable, ...) traditional databases took a shot across the bow
- Scalable tables are simpler than traditional databases
- Simple enough to be closer to file systems than databases?
- How much Table support can be easily added to scalable file systems?
- Or, GIGA+ version 2 embeds BigTable



Status and Directions

- A new project, part of OpenCloud exploration of table-based scalable software
- Use extended attributes for defining columns
- Extend `readdir()` with range queries
- Use a B-tree on file name
- Use directory-embedded redundant attributes for efficient scans
- Split directory into multiple directories with different embedded attributes for locality groups
- Embed small files in directory or create data-free files
- Implementation probably in FUSE, similar to GIGA+ and PLFS