Aries issues:

Crash during recovery – do it again.

CLRs are used for escrow xacts – cannot be undone multiple times. Do an example

Group commit – why required

Application errors: roll forward to a specific point in time, then undo backward – just not to the present.

HA: standard wisdom; active-passive. Roll log forward at passive site. Failover, by recovering. In flight transactions get aborted; not exactly HA. I.e. failover in seconds.

Active-active: 2 active sites, each does all xacts. No log. One is primary – other is secondary. If primary crashes, then keep going from secondary. Best with a stored-procedure interface.

To go fast: H-store data pie.

Buffer pool
Locking
Threading
WAL

See times-10

See NoSQL.

Solution: main memory, one-xact at a time, single thread, no log – failover to a backup – active-active.

Draw H-store picture.

Yabut; multicore

Yabut: multi-shard xacts – spec X

What about network partitions:

Primary can't talk to secondary. Both up. Either:

Primary continues, secondary blocks (less availability)

Or

Both continue – no consisitency

Give up one or the other. Brewer has a CAP theorem – says you can't have all 3.

Application errors, human errors, resource issues, [run out of mem, run out of disk, run out of …] install new software, reprovision – these dwarf network partitions.

Byzantine failures.

Have to have 3 replicas and voting. Nobody worries about this – except theoreticians

---

6.830 / 6.814 Database Systems

Fall 2010