THE CLOUD GOES BOOM!

Berkeley Orders Of Magnitude

- OOM bigger systems
  - OOM less code and development time
- We did it for networking protocols
  - Now we are trying to generalize
  - Make more attractive to programmers
  - How are we going to do this?

The Baby BOOMers

Peter Alvaro, Tyson Condie, Neil Conway, and Russell Sears
Advisor: Prof. Joe Hellerstein

- High level language for programming in the Cloud
  - Lincoln project led by Peter Alvaro
- Lean runtime for this language
  - C4 project led by Neil Conway

- Build a Big Data cloud stack in Overlog
  - A pedagogical exercise
  - Why is distributed programming so hard?
    - Start with the fundamentals

Baby Steps

- Consensus protocols
  - Pocket sized Paxos and 2PC
- What did we learn?
  - How distributed protocols and declarative languages go together
- Overlog gives higher-level primitives than Java or C
  - Build even higher-level “idioms”

Idioms in Paxos

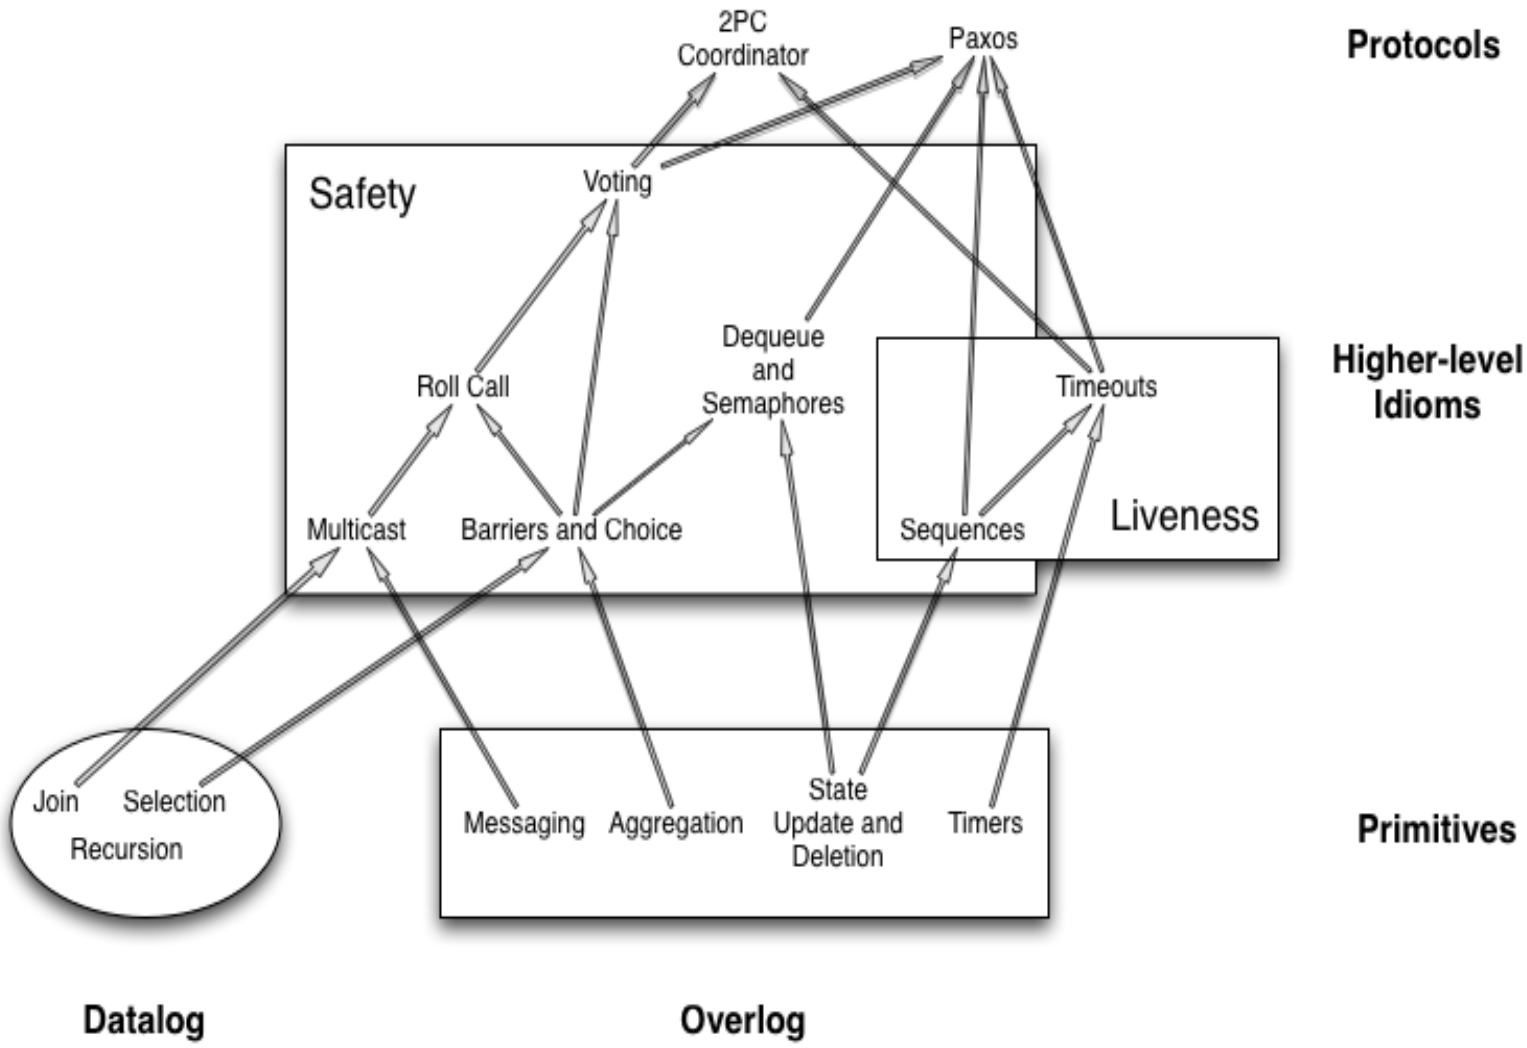

Growing Pains

- Prototype: basic Hadoop functionality
- Subsequent revisions (Hadoop fast-forward)
  - availability rev: hot-standby masters
  - scalability rev: scale out master state
  - monitoring rev: invariant checking, logging

- 9 months, 4 grad student developers most work in a 3-month span

Prototype

- JOL
  - Java-based OverLog interpreter
- BOOM-MR
  - Hadoop MapReduce “brain transplant” with Overlog
- BOOM-FS
  - Hadoop Filesystem (HDFS) rewrite in Overlog
  - API-compliant Java skin

What did we learn?

- Everything is data
  - persistent stuff (e.g. FS metadata)
  - runtime state (e.g. Hadoop bookkeeping)
  - summary stats (e.g. LATE metrics)
  - in-flight msgs and system events
- Because everything is data...
  - easy to design scale-out via partitioning
  - high availability via replication
  - *interposition* (classic OS goal) easy via dataflow
  - *invariant checking* maps cleanly to continuous queries
    - *SPJ over event streams and system state*
  - *simpler concurrency?*
    - *data derivation vs. locks on object updates*

Want more?

- I Do Declare: Consensus in a Logic Language
  - Declarative implementation of MultiPaxos and 2PC

  http://db.cs.berkeley.edu/papers/netdb09-idodeclare.pdf

- BOOM: Data-Centric Programming in the Datacenter
  - Declarative implementation of Hadoop and HDFS

  http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-113.html

- Our advisor's blog

  http://databeta.wordpress.com/