



Distributed data management

- What we say (distributed)
 - Information is ubiquitous, everywhere on the Web
 - Applications are choreographed with autonomous servers
- What we do (centralized)
 - *Centralized* relational & XML databases
 - Orchestration with *centralized* workflows
 - Cloud computing with *centralized* services Google, Facebook, Amazon
 - And this is where most of the research is focused
 - Notable exceptions: P2P storage/indexing



Alternative: move databases around

- **Active XML artifact** [Time 2009]
 - Objects moved around the internet
- With their logic and processing
 - Not just static data
 - Security issue for their hosts: code from someone else running on the host
- With their access control
 - Use cryptography to control access in travelling objects
- With possibly other database features
 - Transaction and locking, recovery management, etc.



Not clear why you would do all that, but this is clearly novel and challenging