

A Relational Model of File Organization

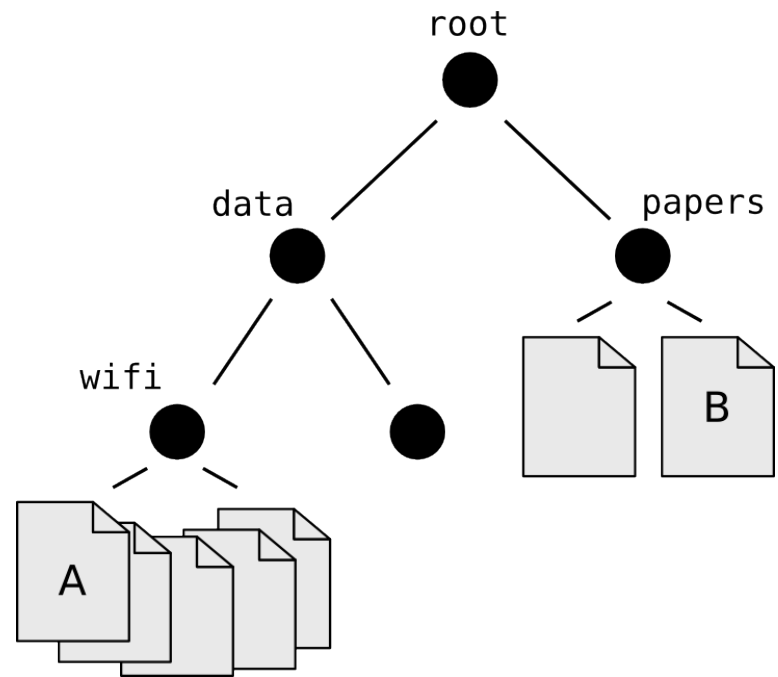
December 13, 2006
Jeff Smith

Agenda

- The *status quo*: a hierarchical model
- Introducing a relational model
- An implementation of the new model

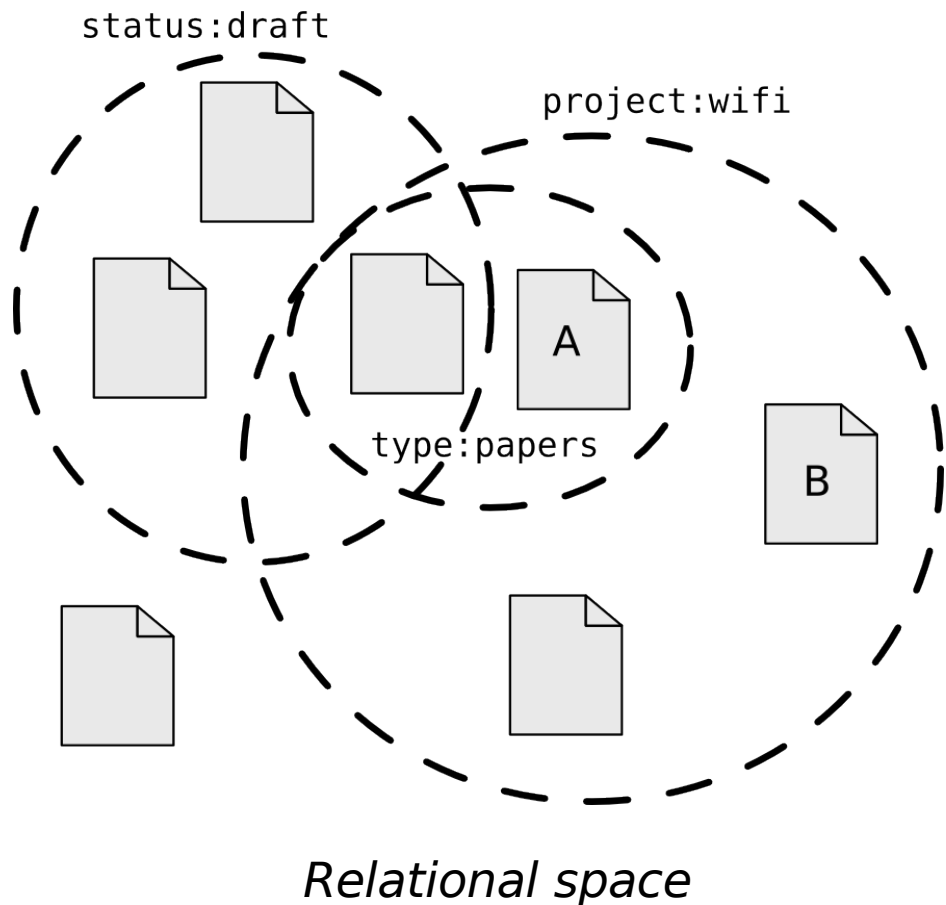
The Hierarchical Model

- Every object has a parent (except root)
- Files live in a tree
- Files are associated with parent and siblings



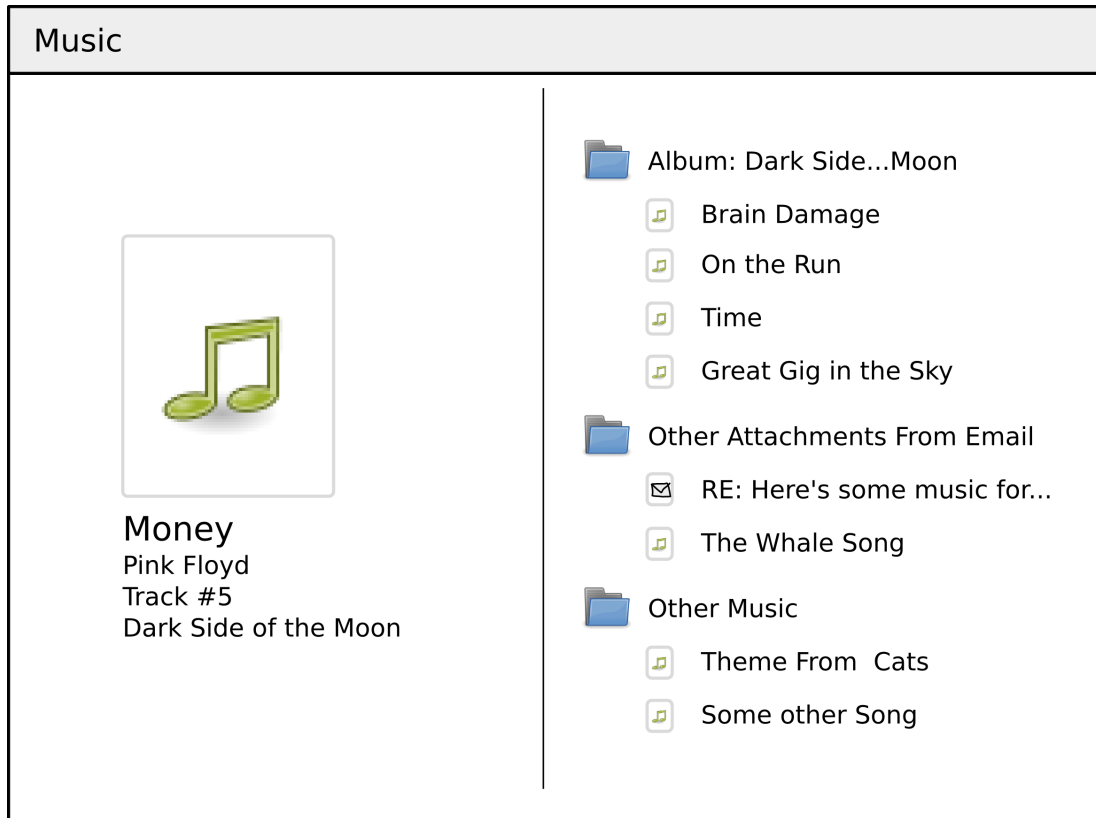
A filesystem's hierarchy

The Relational Model



- Files are tagged with (Tag, Value) tuples
- Files do not have a “location”
- Views associate files by their tags

Browsing a Relational System



Mockup of browse interface

- Tags are structured, so computers can automatically find relations between files
- Relationships can be arbitrary

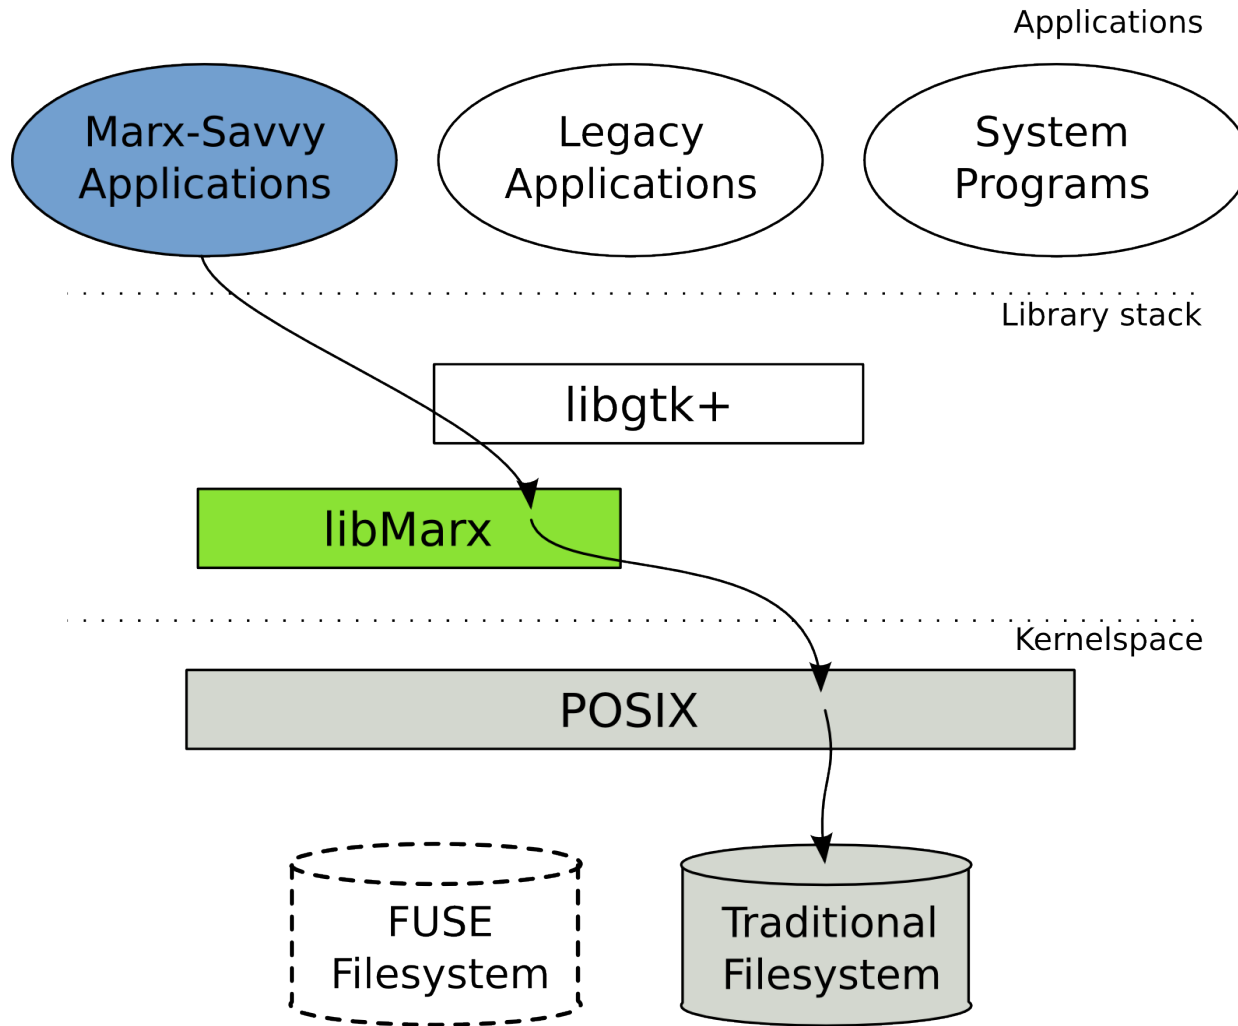
Finding in a Relational System



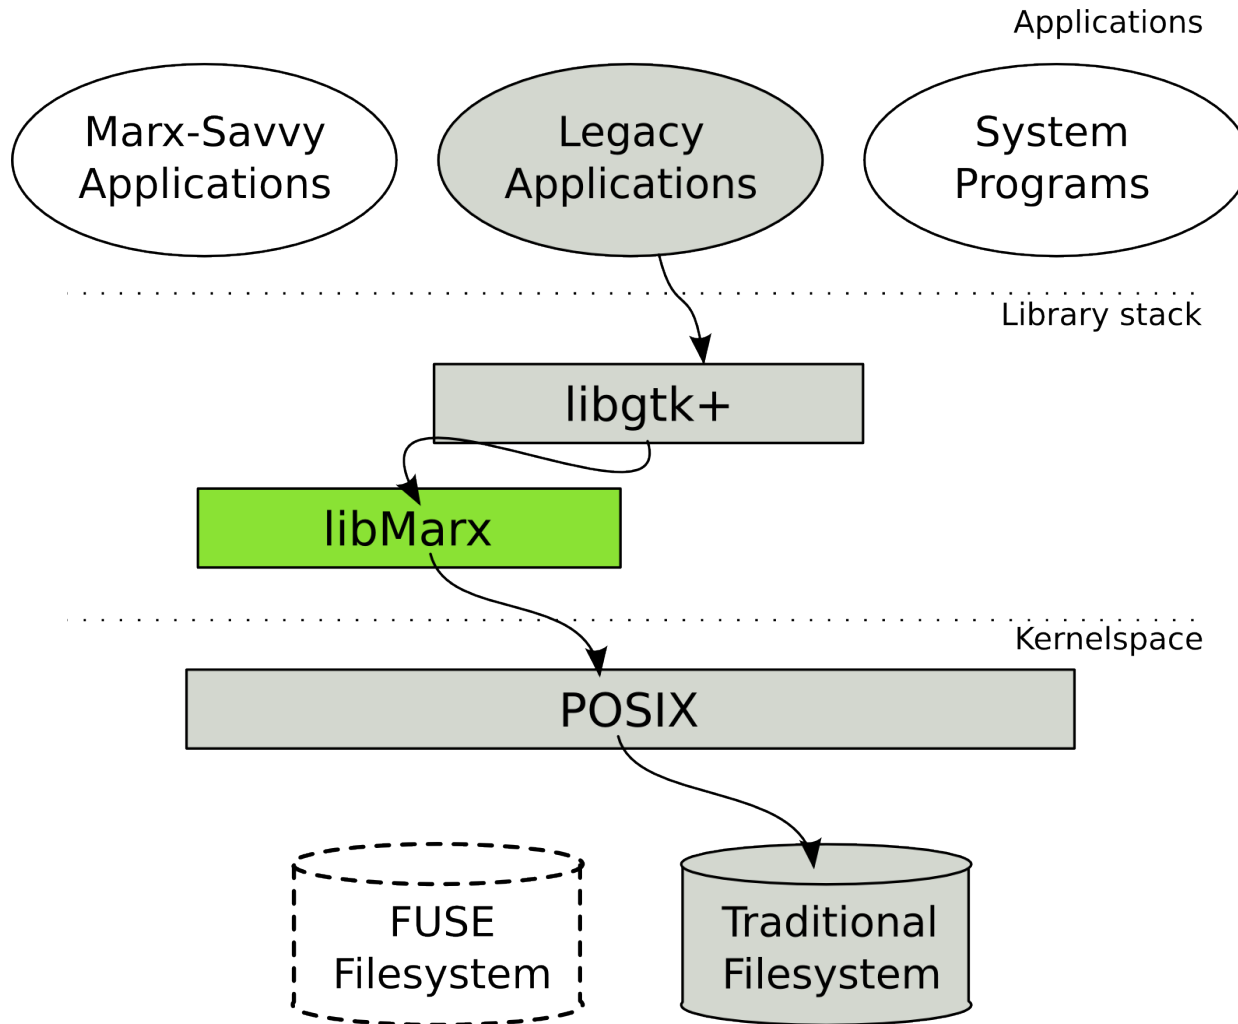
Mockup of find interface

- Thanks to tags, queries can be powerful
- Search is speedy, because metadata is managed by a relational database

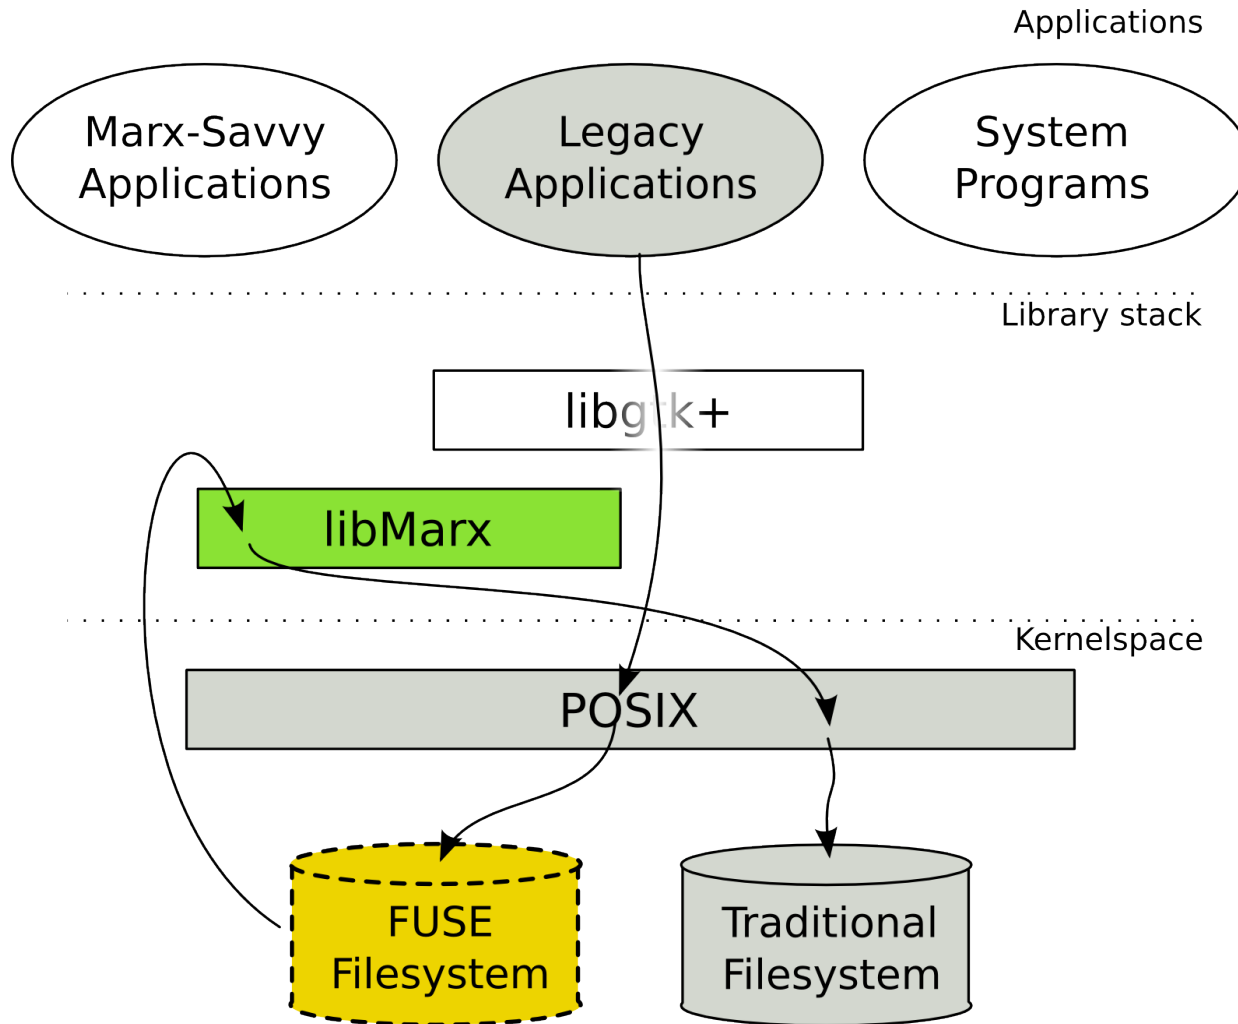
Model Architecture



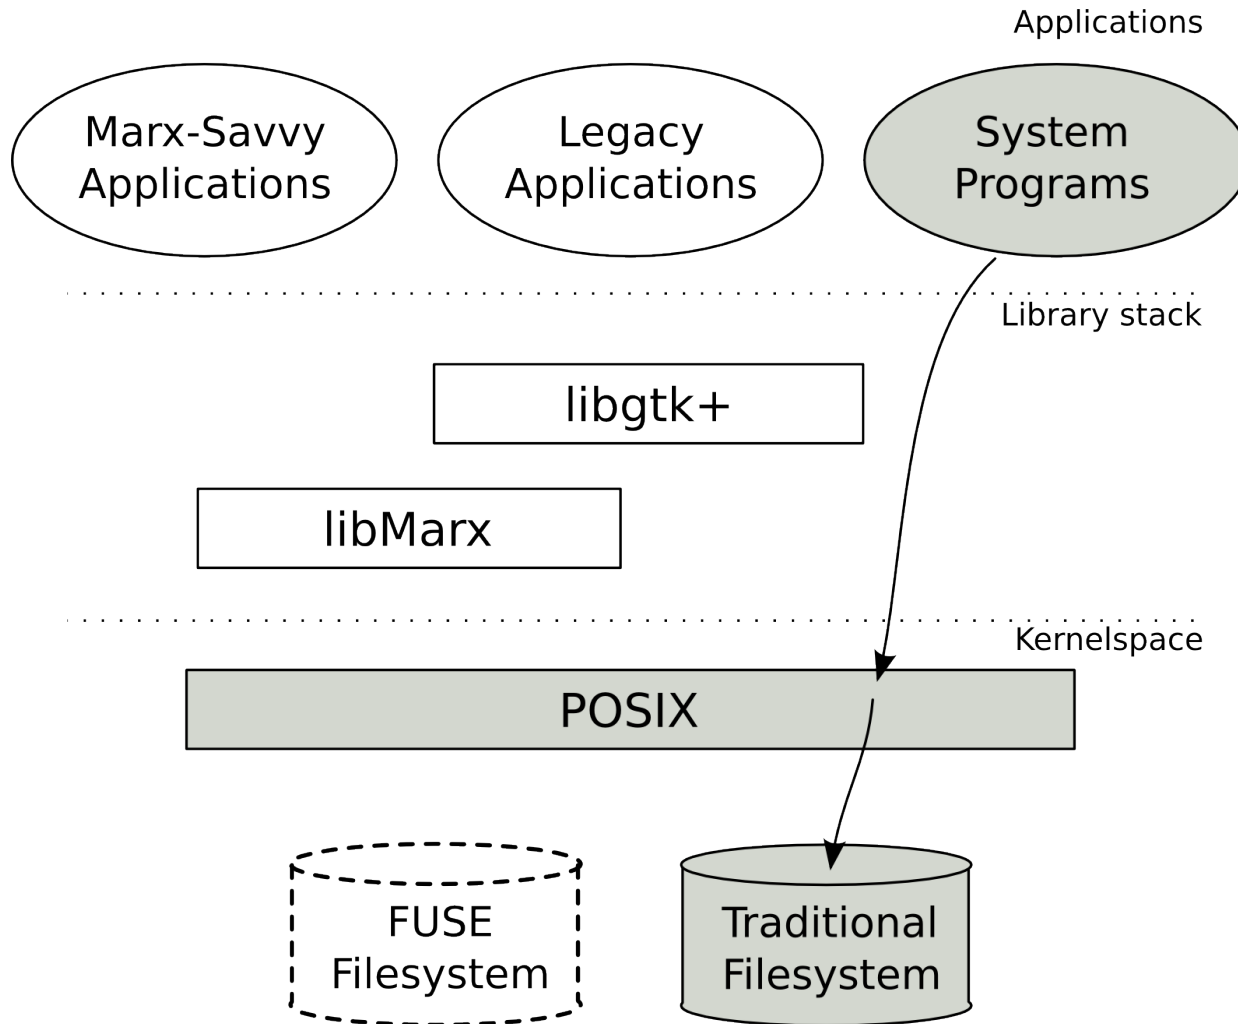
Model Architecture



Model Architecture



Model Architecture



Prototype

- Can search through 200,000 files (with 2,300,000 tags) in one second
- Strong tag semantics for predefined tags
- C interface is flexible and asynchronous