MFS: a Network-Aware Mobile File System

Benjamin Atkin and Kenneth P. Birman

Reliable Distributed Systems Group
Cornell University
Ithaca, NY

Bandwidth adaptation

- Mobile hosts see wide bandwidth variations
- Adaptation is required for good performance
- Traditionally adapt to bandwidth via modes
  - Coarse-grained
  - Can underutilise
- Alternative: adapt modelessly based on the current active RPCs

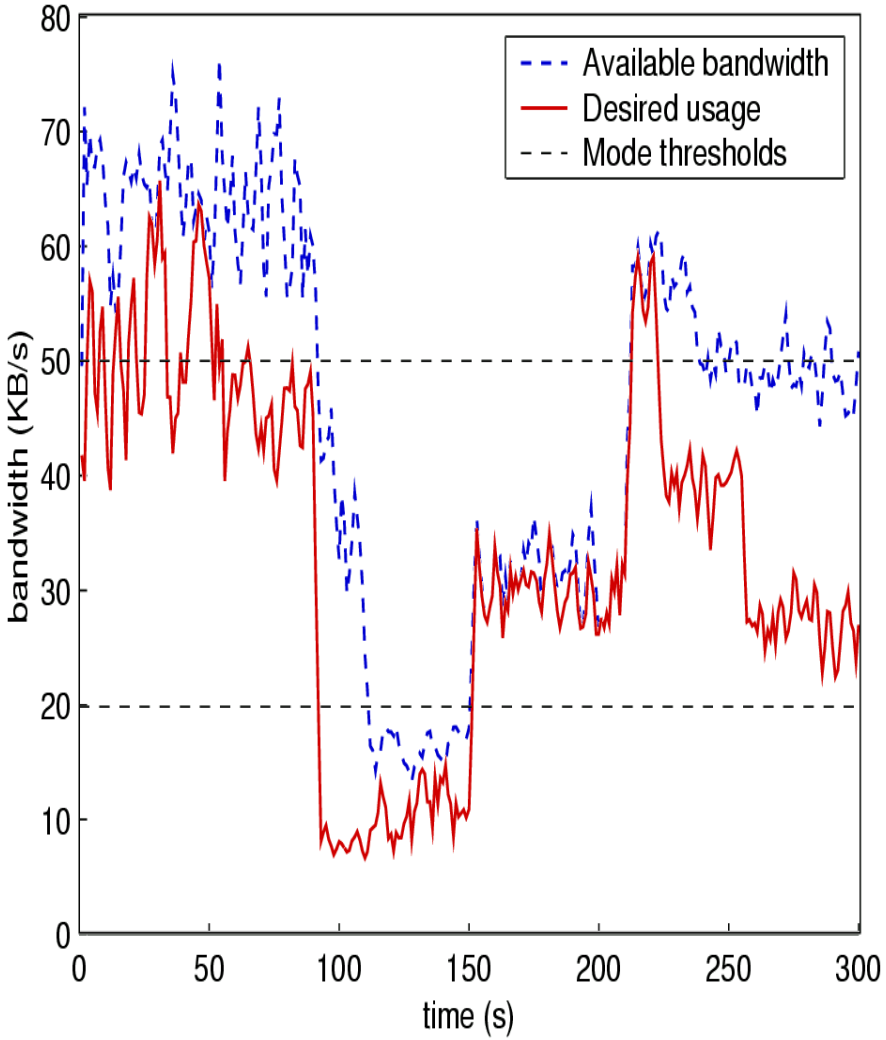

MFS Mobile File System

- Distributed file system for mobile hosts
- Bandwidth allocation via RPC priorities
  - Small or "foreground" RPC => high priority
  - Implemented using ATP [Infocom 2003]
- New cache consistency algorithm
  - Focuses on consistency for environments with high sharing, e.g. collaborative design
  - Adaptive invalidation policy
  - Differentiates shared/unshared files

Effect of priorities

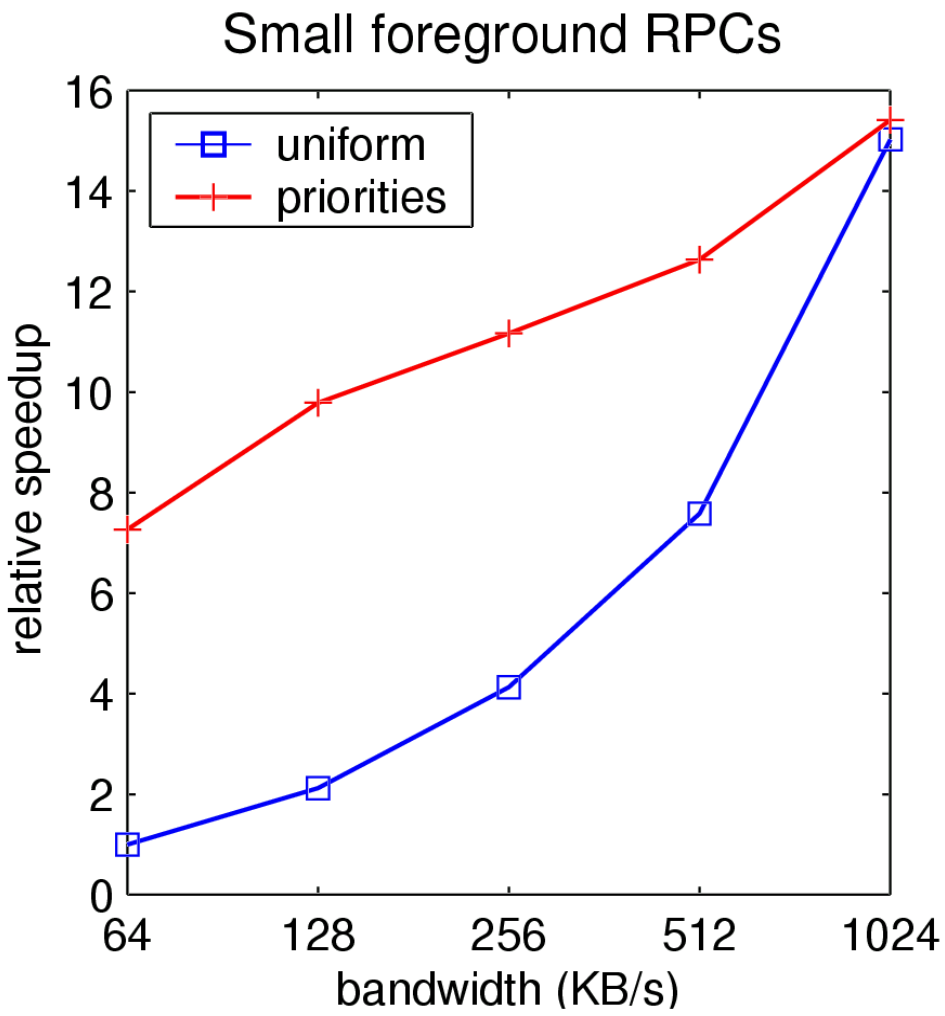

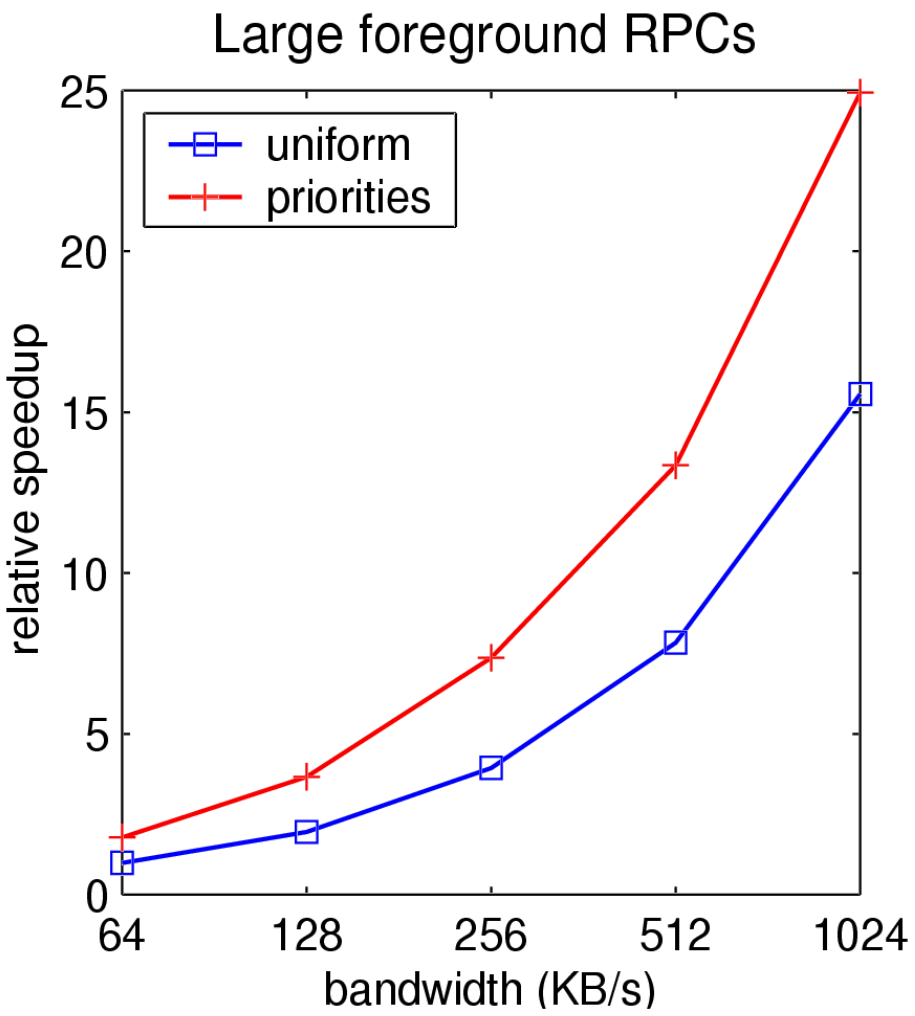

Cache consistency algorithm

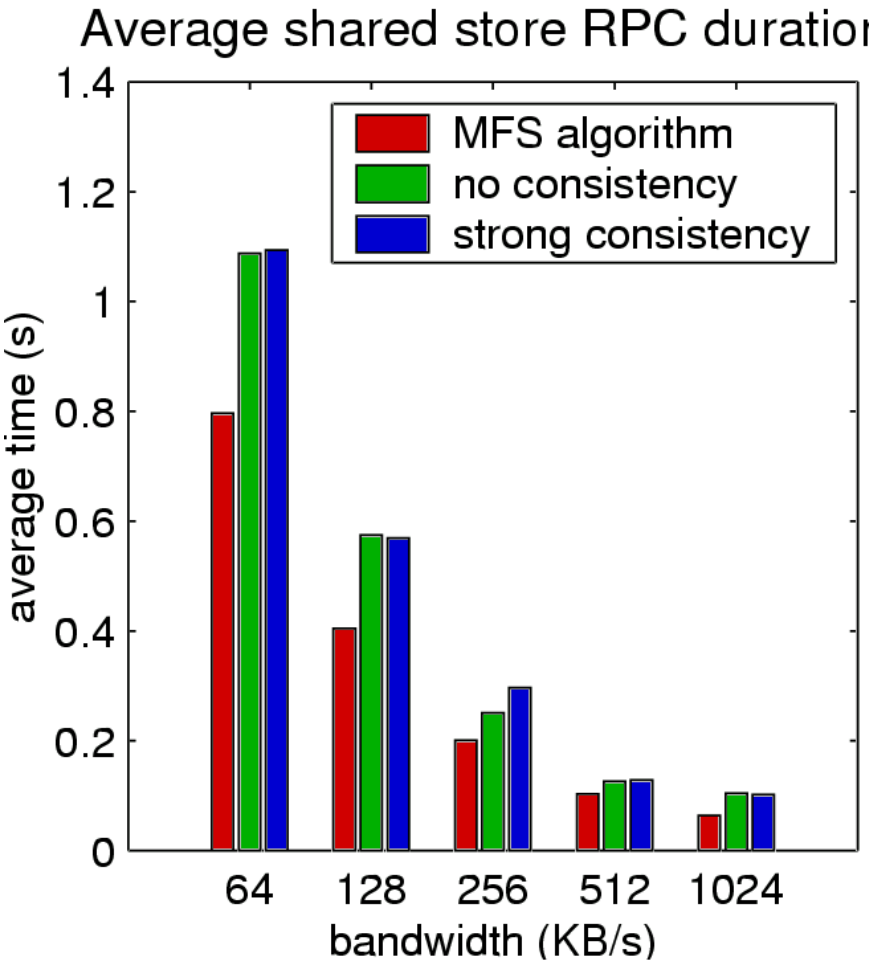

1. Effect on write delay

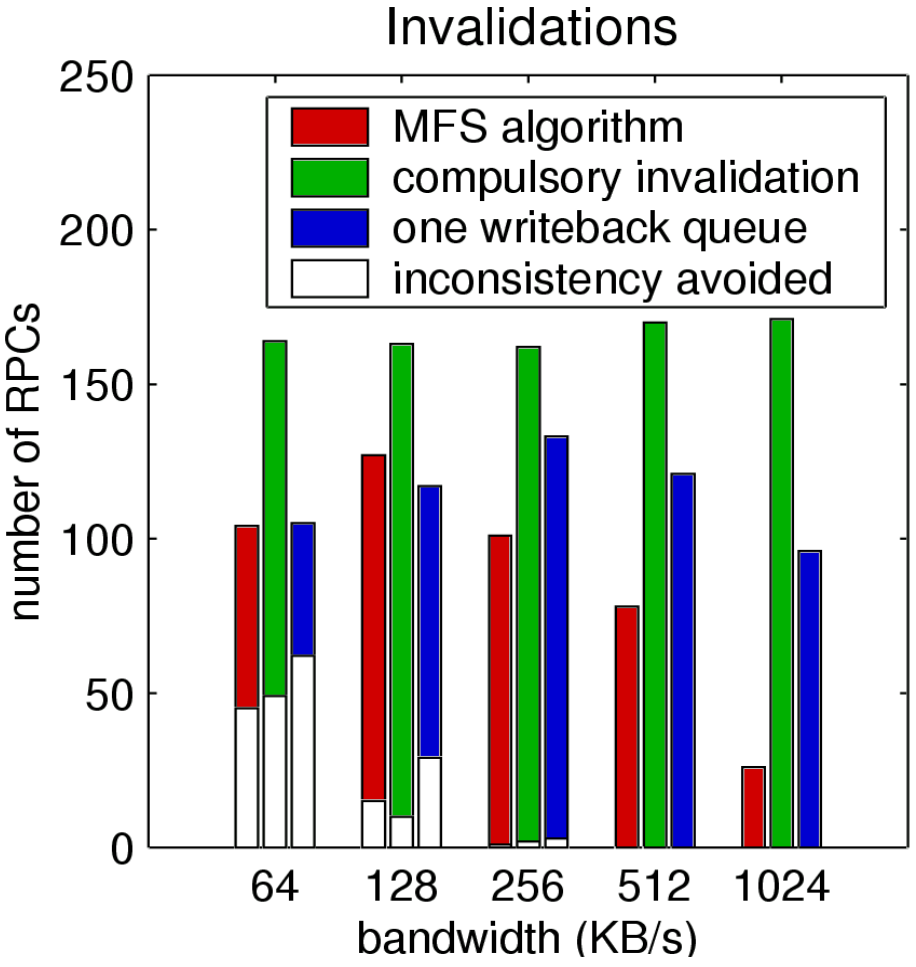

2. Effect on consistency

Summary

- Modeless adaptation in MFS looks very promising ...
  - Investigating other application domains and more general environments
- More information available on the web

`http://www.cs.cornell.edu/batkin/modeless`