**A. Arpaci-Dusseau**
**CS739: Distributed Systems**

**Department of Computer Science**
**University of Wisconsin, Madison**

# Taming aggressive replication in the Pangaea wide-area file system - OSDI'02

## 1 Introduction

- What environment is Pangaea targeting? What workload? Is this well motivated?
- What consistency model do they provide? Why do they think this model is okay? Why don't they want to provide strong data consistency?

## 2 Pangaea: A structural overview

- Where is your data stored in Pangaea? Do you think this is a good idea?
- What are the two purposes of replication? How do Gold and Bronze replicas in Pangaea fulfill these different purposes?
- What meta-data must be tracked for each type of replica? Why don't you want to make all copies Gold?

## 3 Replica set management

- How does a user create a file? (Example: User on server S creates file F in directory D.)
- How does a user on node S access an existing file F?
- After S has a replica, what must it do?

## 4 Propagating updates

- Pangaea chooses to update replicas through *flooding*; i.e., pushing the new contents to every neighbor. Since the graph of replicas is strongly connected, all replicas will eventually receive the updates. What optimizations do they implement beyond flooding?
- How does Pangaea determine whether or not it is correct to update a replica?
- Updating directory entries is particularly interesting. For example, if simultaneously execute `mv /foo /alice/foo` and `mv /foo /bob/foo`, hard to detect conflict since `alice` and `bob` might be stored on different nodes. How does Pangaea solve this problem?

- If a user is worried about eventual consistency, what can they do?

# 5 Evaluation

- How much of the system is implemented?
- What do they show in Table 2 and Figure 7?
- What is the point of Figure 8?
- What is the point of Figure 9?

# 6 Class Discussion Questions

- What do you think is this major weakness of this work?
- Given its goals and assumptions, do you think Pangaea is designed well? Why?
- There is no “central authority” for determining who owns a file or who should host gold replicas. What complications does this raise in Pangaea?
- What do you think are the pros and cons of the allocation policy used in Pangaea?
- Could Pangaea be modified to perform replication at the block level instead of the whole file level? What complications would this raise?
- Are there any disadvantages to hosting a bronze replica? Why do you sometimes want to remove a bronze replica?