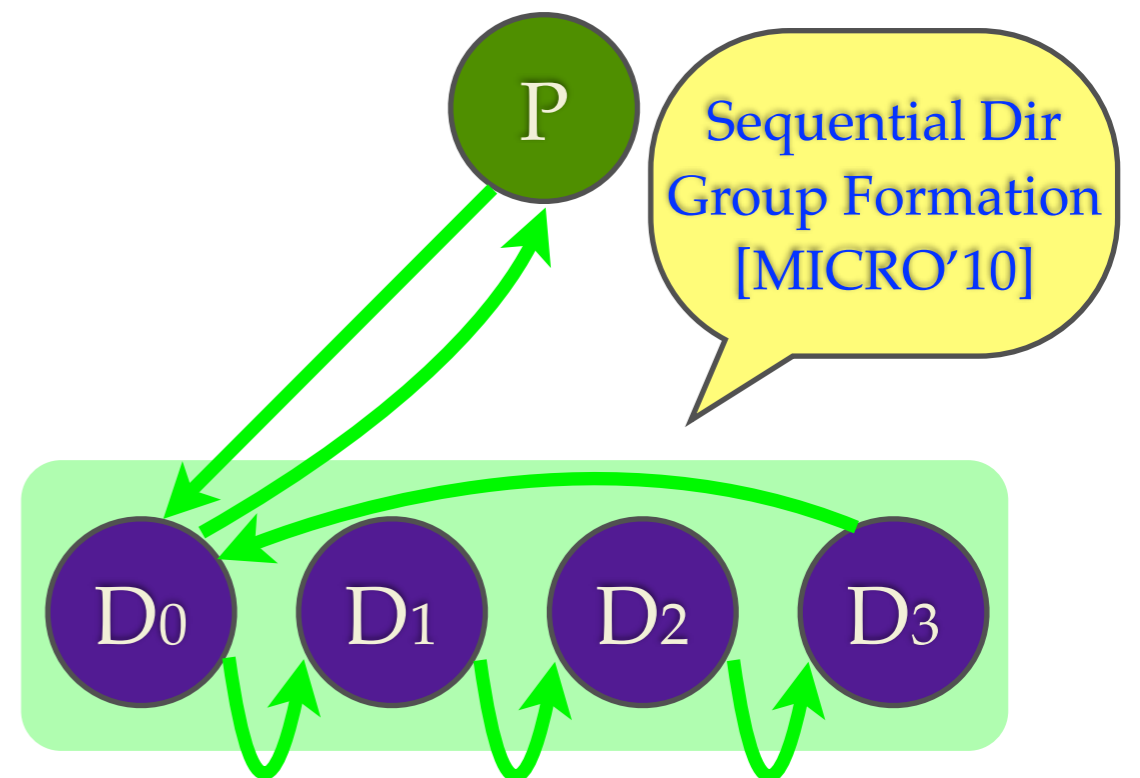



BulkCommit: Scalable and Fast Commit of Atomic Blocks in a Lazy Multiprocessor Environment

Xuehai Qian, Josep Torrellas (UIUC)


Benjamin Sahelices (Univ of Valladolid) and Depei Qian (Beihang Univ.)

- Problem:
 - Current atomic block (chunk) execution incurs unnecessary squashes
 - Atomic block commit operation in a lazy environment has sequential bottlenecks
- Our solution:
 - IntelliSquash: no squash on WAW-only conflict
 - IntelliCommit: parallel directory group formation




IntelliSquash: No Squash on WAW-only Conflict

- Insight: WAW is a name dependence. It does not break semantic atomicity
- Similarity with two conflict stores from two processors
- If two chunks only have WAW conflicts, IntelliSquash serializes them without squash



IntelliCommit: Parallel Directory Group Formation



- On chunk commit:
 - Processor sends commit requests to all the relevant directory modules
 - Directory module receives commit request:
 - Locks the memory lines
 - Responds with commit_ack
 - Processor counts the number of commit_acks received
 - Processor sends commit_confirm when it receives the expected number of commit_acks
- Challenge: resolving conflicts from two processors
- ChunkSort: ordering all the conflicting chunks in the same order in all relevant directories by **preemption**