

InterWeave: A Middleware System for Distributed Shared State

Deqing Chen, Sandhya Dwarkadas, Srinivasan Parthasarathy,

Eduardo Pinheiro, and Michael L. Scott

Computer Science Department

University of Rochester

{lukechen, sandhya, srini, edpin, scott}@cs.rochester.edu

May 26, 2000


Outline

- Motivation
- Application Programming Interface
- Coherence Models
- Implementation
- Prototype Application
- Conclusion and Current Status

Motivation

- Remote visualization and steering
- Heterogeneous distributed data sharing
- InterWeave is the merger of the InterAct and Cashmere-2L systems
 - We want to provide a system that combines tightly coupled system with remote access

A Three Level Structure



Application Programming Interface

- Segment and data creation
 - IW_create_segment, IW_malloc, IW_free
- Segment coordination
 - IW_wl_acquire and IW_wl_release
 - IW_rl_acquire and IW_rl_release
- Segment access
 - IW_get_segment, IW_get_handle
 - IW_mip_to_ptr, IW_ptr_to_mip
- Others

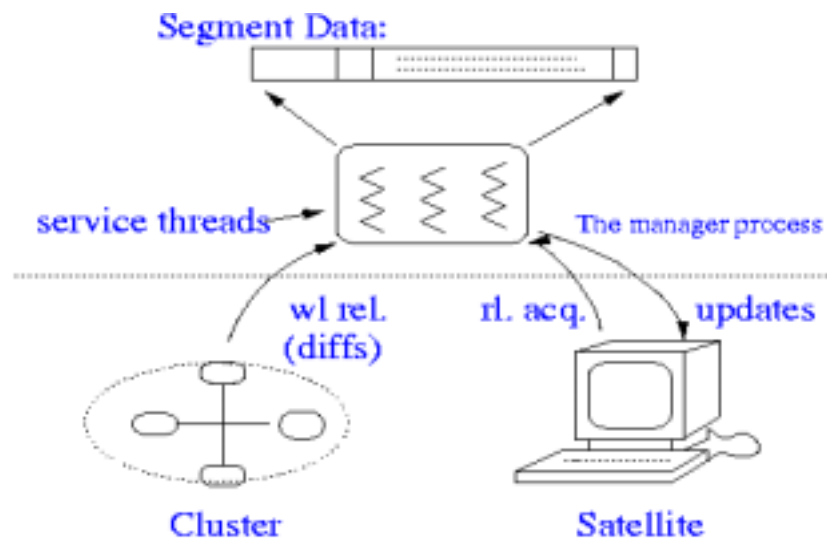
Coherence Models

- The Coherence models used in InterWeave are adapted from the InterAct system
- Strict writer coherence
- Strict reader coherence
- Relaxed reader coherence
 - Default
 - Diff
 - Temporal
 - Delta

Coherence Examples


- Barnes simulation and visualization: temporal coherence
- Interactive association data-mining: diff coherence
- Stock quotes messenger: default relaxed reader coherence

Implementation



Segment server and satellites

Implementation(cont.)



Satellite on a two level system


Prototype Application

- Remote visualization for Barnes simulation
 - Simulation space
 - simulated results are shared in an InterWeave segment
 - updates are protected by writer locks
 - Visualization satellites
 - using relaxed temporal coherence to keep its data up-to-date

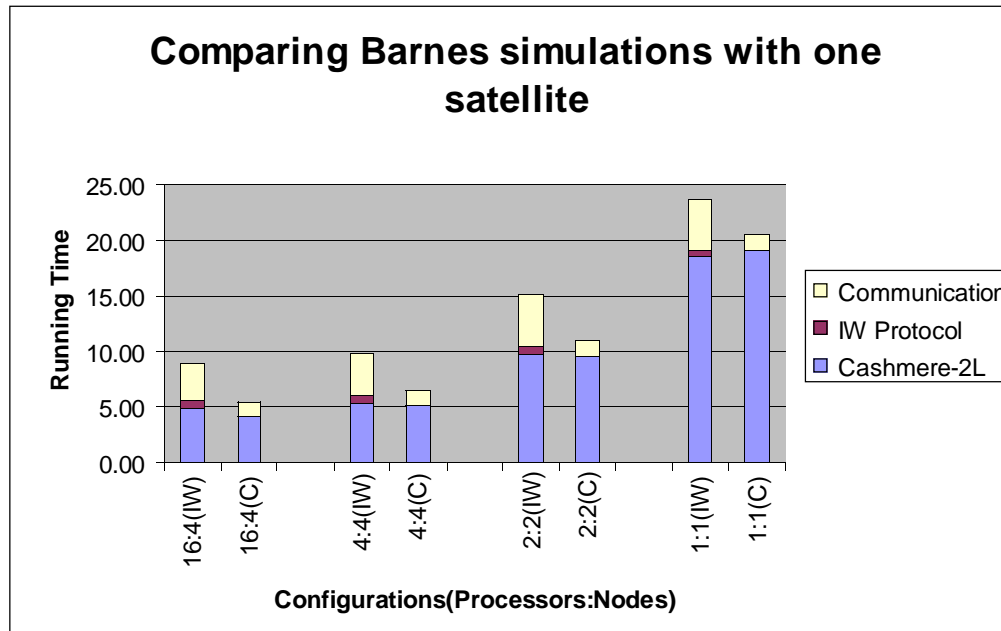
Performance Experiments

- Setup
 - the barnes simulation and visualization application
 - the simulator is running on a 4-node Cashmere cluster
 - the satellite runs on another Alpha node
- Measurement
 - we measure the performance of the InterWeave barnes simulator on the Cashmere cluster

Results(1)



Results(2)



Conclusion and Current Status

- InterWeave can be a useful system for applications that need to share a large data set over a heterogeneous environment
- Prototype implementation and application are running
- Topics that haven't been covered in this talk:
 - machine and language heterogeneity
 - cross-segment consistency
 - fault tolerance, security, etc

Heterogeneity

- **Wire format**
 - data objects are defined by an IDL language
 - a wire format is defined for communication between the server and satellite
- **Machine independent pointer**
 - URL path names
 - pointer swizzling between the MIP and local pointers

Cross-segments causality

- Causality caused by cross segments dependence
- Solution: attached a timestamp vector for each segment
 - the timestamp records the version numbers of every segment this segment dependent on
 - vector is fixed size, segment get a slot by a predefined hash function
 - version number roll-up