

MACEDON

April 20, 2004

I. Background

Distributed systems:

- o old model: RPC and/or distributed objects
- o new model: cooperating finite-state machines
- o claim: overlay networks, DHTs, etc. can be built using an FSM framework

II. Finite-State Machines

State machine for each node

- o big states: joining, leaving, init, etc.
- o state variables, e.g. neighbor list, timers
- o atomic transitions: read/write locks
 - control packets = write
 - data packets = read only

Three kind of transitions:

- o API calls: before_state(s) API <api_name> [locking read;] { <actions> }
- o message reception: before_state(s) rcv|forward <message_name> { <actions> }
- o timer events: before_state(s) timer <timer_name> [locking read;] { <actions> }

Actions: work done on a transition

- o actions occur on transitions (not in states)
- o therefore locking on transitions is typically sufficient
- o examples:
 - schedule timer (e.g. timeout)
 - transmit message
 - change state
 - modify state variables

Note that FSM is hard to check statically, uses dynamically linked handlers

III. Auxiliary Support

Basic API:

- o init
- o forward: route message to next hop
- o deliver: arrival at last hop (?)
- o notify: upcall about network changes (or lower layers in general)
- o groups: create/join/leave
- o route/multicast/anycast

Libraries:

- o SHA and crypto functions, e.g. MD5, public keys
- o basic data structures: hash tables, Bloom filters
- o locks
- o tracing/logging support
- o neighbor lists (e.g. pick a random neighbor)

IV. Code generation

Generates C++ for PlanetLab or ModelNet or NS

input code size is very small (big win): 100-600 lines of FSM code

V. Validation

very important: need to know that generated code is comparable with the real version

Validation done on ModelNet (rather than PlanetLab)

Decent validation for NICE and Chord (app-level multicast and DHT respectively)

Pastry results: Macedon has better performance than FreePastry (probably due to Java)

- o implies there is no real performance penalty for the higher-level of abstraction

Also used to explore parameter settings for SplitStream