Supporting Rapid Mobility via Locality in an Overlay Network

Ben Y. Zhao
Anthony D. Joseph
John D. Kubiatowicz

Sahara / OceanStore Joint Session
June 10, 2002

Ubiquitous Computing to the Edge

- Trends
  - Network infrastructure growing fast (802.11b, 3G)
  - Edge devices getting smaller, more powerful
  - Edge devices increasing in storage capacity
- Legacy mobility support not enough
  - Mobile IP focused on single roaming nodes
  - # of mobile devices reaching critical mass
    - E.g. ~400 million Cellphones, PDAs
    - Sheer volume can overwhelm mobility infrastructure
  - Devices need object location support for P2P applications (IM, file sharing, directory services)

Pushing the Limits of Mobility

- What is the next step?
- Scenario I
  - UCB Professor roaming to MIT, directing streaming media to his laptop, participating in video-conference
  - Need: persistent connection with low latency delivery, and fast handover between proxies
- Scenario II
  - Morning ride on BART, 100-1000 networked commuters switching from proxy to proxy in tandem
  - Need: scalable solution to proxy handover for large mobile groups (tourist groups, airplane passengers, etc…)

Legacy Mobility Support

- Mobile IP review
  - Mobile Host (MH), Correspondent Host (CH) Home Agent (HA), Foreign Agent (FA)
  - Path: CH → HA → FA → MH
- Issues
  - Triangle routing (High RDP)
  - More fault-prone
  - Higher BW

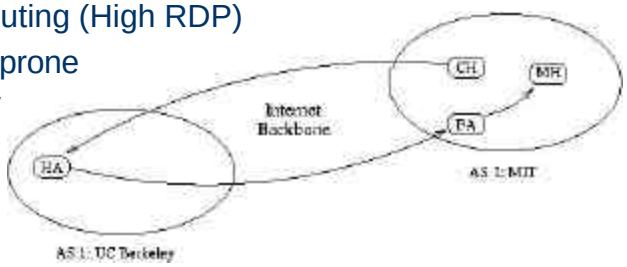

What Do We Need and Why?

- Mobile IP weakness: triangle routing
  - High latency for messages and location updates
  - Messages / updates incur high costs in wide-area
  - Messages / updates susceptible to wide-area routing failures
  - Home agent single point of load congestion and failure
- Mobile IP / TCP migrate / ROAM: fixed indirection point
- Key insight: scalability via locality-awareness
  - Client requests result in load on network infrastructure
  - Dampen / confine the impact of network operations: reduce wide-area storage / communication costs, faults
  - In practice: minimize latency and faults by localizing communication (registrations, msg routing, location updates)
- Strategy:
  - Leverage locality-aware overlay infrastructure: Tapestry

Tapestry Review

- Decentralized object location and routing overlay
  - 2000: Bayeux, Brocade, Mobile Tap., Fault-tolerant Tap.
  - Related projects: PRR97, Pastry, Chord, CAN, Kademlia
- Routing infrastructure based on prefix routing
  - Route to nearest node increasing prefix match by 1
    AA93 → 4A6D → 4361 → 437A → 4378
- Locality-aware design:
  - Proximity routing: local traffic confined to local network
  - Object location: replicate Log(N) pointers to object, place disproportionate more replicates near object
  - Result: RDP of routing to node is small
    Finds closest copy of object with low RDP

Object Location and Routing

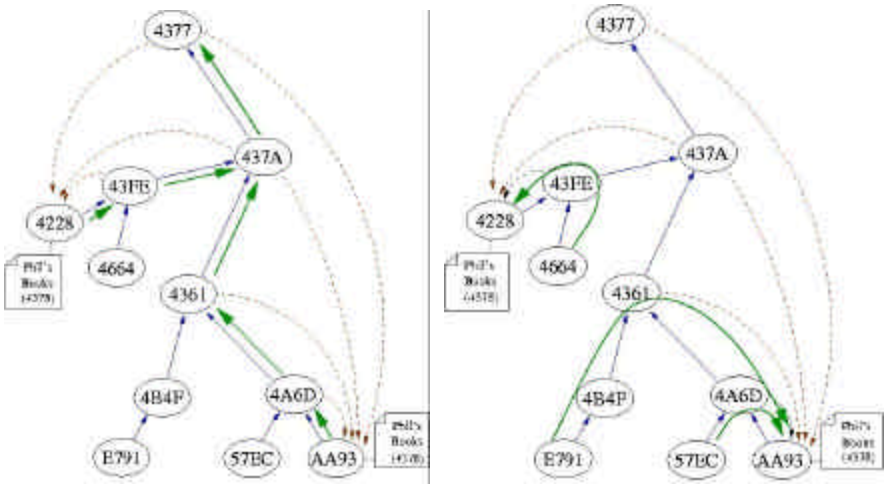

Mobile Tapestry

- Reduce mobility problem to object location
- Register (*mn*) = Proxy "publishes" object *mn* with reverse forwarding path
- Routing to *mn* = Find object *mn* + follow pointers to proxy and *mn*
- Multiple levels of transparent and efficient indirection in the infrastructure
- No home agent, no foreign agent, no triangle routing
  The Internet is your home network!

Registration and Routing

- Mobile node *mn* registers with nearby proxy P
- RegisterMsg *R* publishes object named *mn* at node P
- Nodes store reverse pointer trail from proxy P to root (mn)
- Message for *mn* routes towards R(mn), then P, then *mn*
- Object location locality → distance traveled by message proportional to actual distance between CH and P

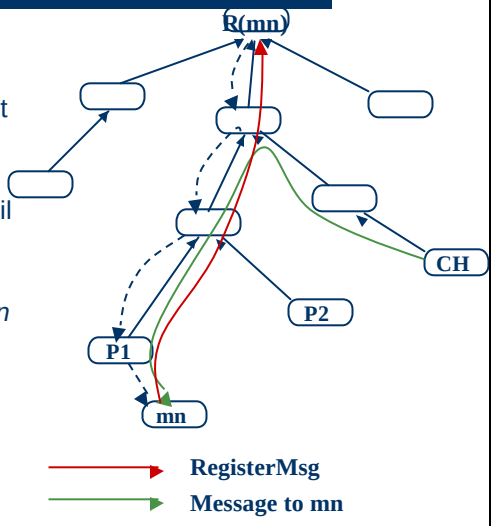

Location Updates

- *mn* moves from *P1* to *P2* sends ProxyHandoverMsg *U*
- *mn* sets up forward link from *P1* to *P2*
- Update message *U* routes up until intersects old route at B
- *U* backtracks to *P1* to delete old references
- Message M forwarded as RouteObjMsg to proxy *P2*, then to *mn*
- Routing locality → distance traveled by update *"proportional"* to distance between proxies

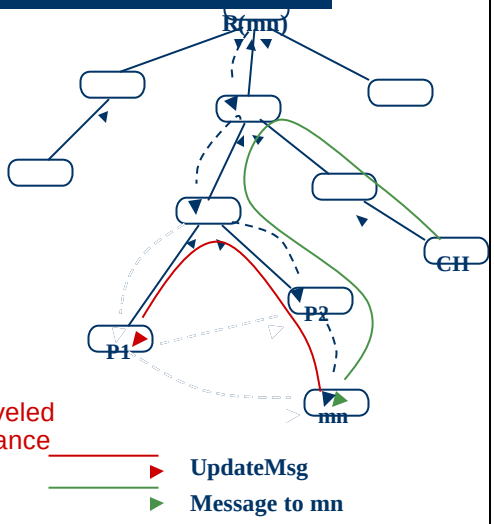

Mobile Objects

- Extending object location to mobile nodes
- Mobile object location
  - Extra step of indirection
  - Publish <Mobile Object, Mobile Node> mapping
- CH routes message to mobile object *MO*
  - Route msg to *MO*, find <MO, MN>
  - Route msg to mobile address *MN*, find <MN, P>
  - Route msg to proxy P, then forward

Mobile Objects

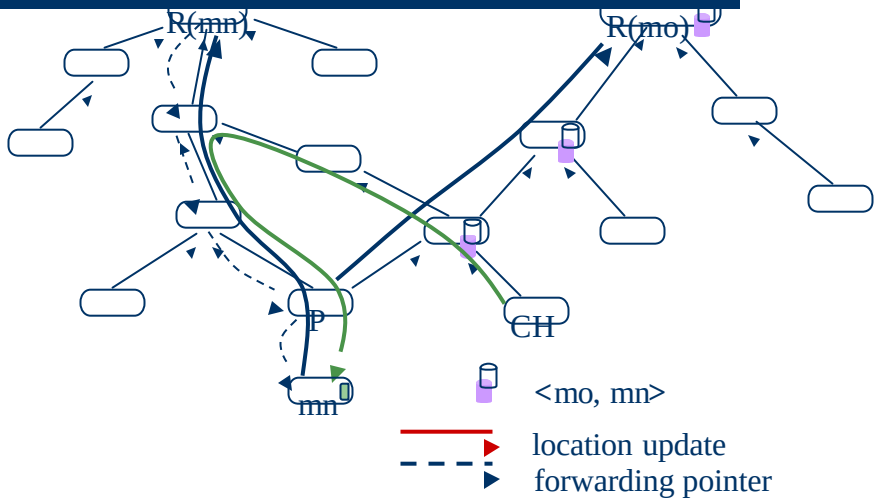

Hierarchical Mobility

- Scenario:
  - Networked bullet train in motion
  - Carrying passengers with 1000 networked devices
- Train server is node in mobile Tapestry
  - Server registers with nearby proxy *P* as node *S*
  - Nearby proxies update location of *S* as necessary
- Each mobile device registers with train as MD_i
  - Server publishes MD_i as local object $<MD_i, S>$ for all *i*
- CH sends message to a device
  - Message routes to MD_i, finds mapping for *S*
  - Message routes towards *S*, then to *P*, then MD_i.

Intro Algorithms Evaluation Conclusion

Hierarchical Mobility

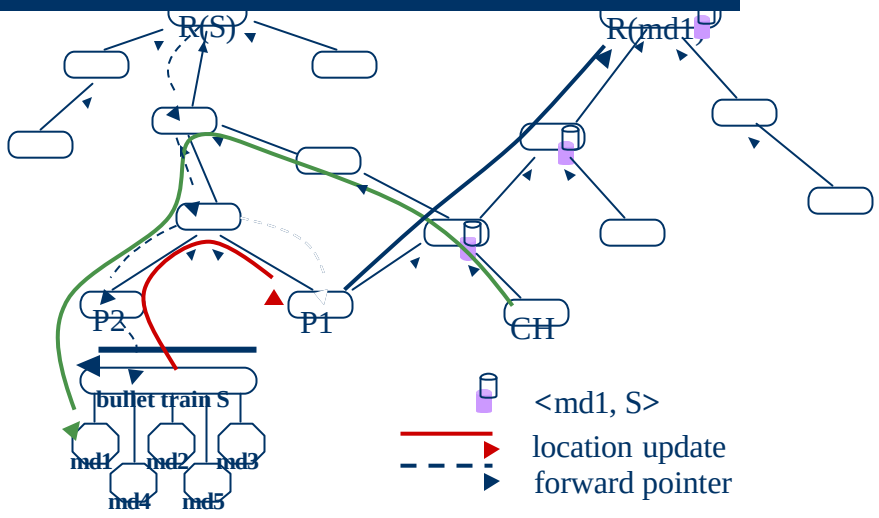

ravenben@eecs.berkeley.edu June 10, 2002

Levels of Type Indirection

- Functionality enabled by multiple levels of indirection
- Mobile node = object on static node (proxy)
- Mobile object = object on mobile node
- Mobile child = object on mobile node (proxy)
- Potentially limitless levels of indirection for hierarchy of mobile devices

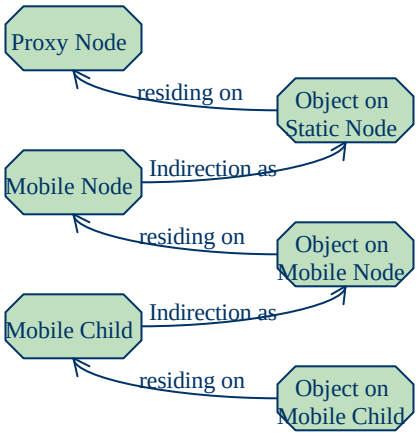

GUID Aliasing

- Utilize redundancy for performance and stability
- Each node gets several aliases
  - Generated by SHA-1(nodeID, int i) for (0 < i < 5)
- Usage
  1. CH sends message to ALL aliases
     MN receives first, drops duplicates
  2. On initializing connection, CH sends message to ALL aliases. MN responds and indicates preferred alias. All subsequent communication uses chosen alias.

Simulation Framework

- Packet level simulation
  - Simulation of network topologies (5000 nodes, 4096 Tapestry nodes, 6 digit base 4 IDs for nodes / objects)
  - No simulation of congestion or network-level traffic flow
- Experiments on transit stub topology
  - Aggregated over 9 topologies
  - Each run over 3 random overlay placements
- GUID aliasing assumptions
  - Initiate connection by using 3 parallel aliases
  - Quickly measure and pick alias for good performance

GUID Aliasing vs. Mobile IP

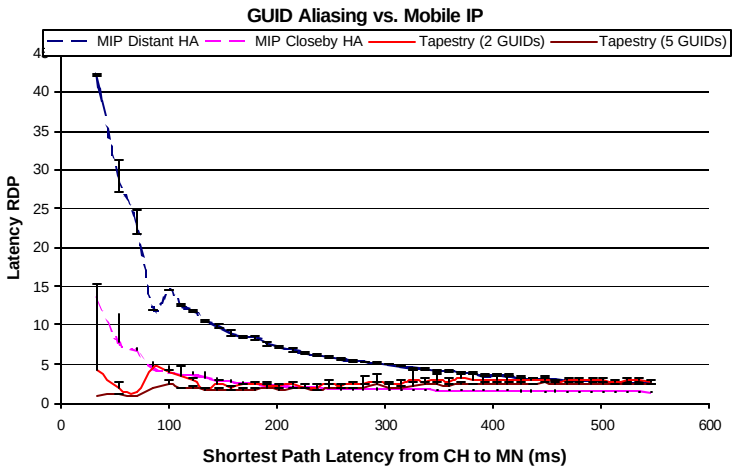

Update Latency w/ Rapid Mobility

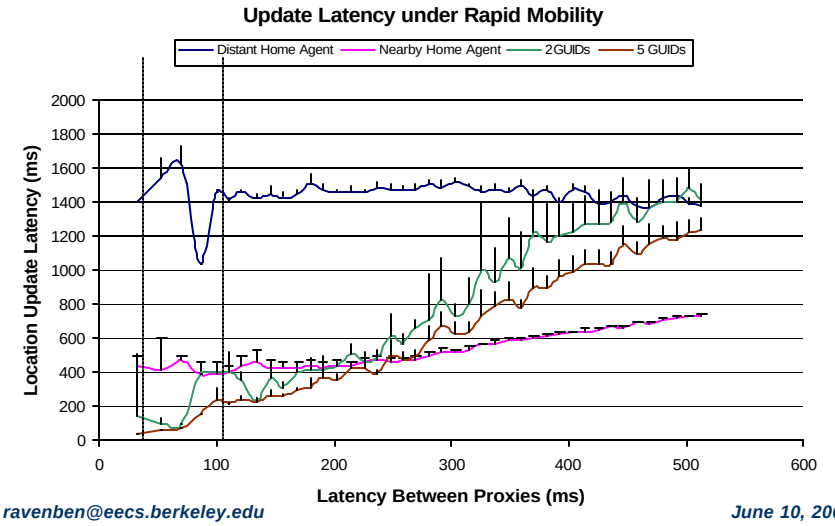

ravenben@eecs.berkeley.edu

June 10, 2002

Router Load vs. Mass Mobility

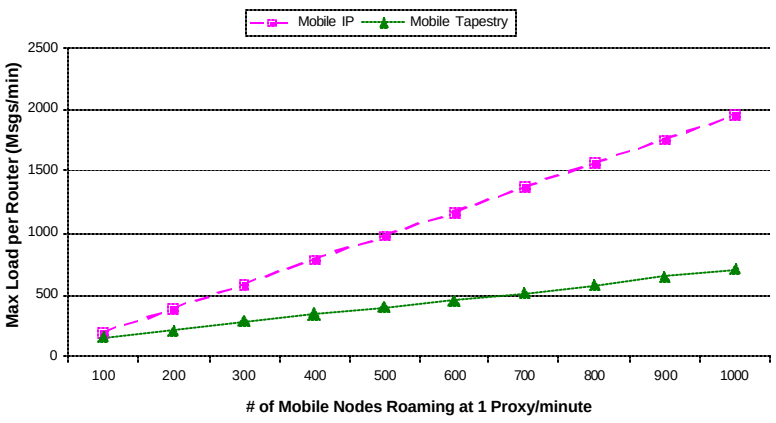

ravenben@eecs.berkeley.edu

June 10, 2002

Summing It Up…

- Mobility via indirection
  - Mobile nodes register as "objects" in infrastructure
  - Messages to mobile node redirect through location pointer
- Distinctions from ROAM:
  - Efficiency *and* flexibility using locality
- Performance
  - Tapestry routes with locality → low RDP for routing (even with overlay indirection), fast location updates
- Extreme scalability
  - Localize traffic around shortest path, stabilize wide-area
  - Use multiple levels of indirection for hierarchical mobility
- Fault-tolerance
  - Reduce susceptibility to faults by localizing wide-area traffic

For More Information…

- Tapestry / Bayeux / Brocade
  - http://www.cs.berkeley.edu/~ravenben/tapestry
  - PDFs of papers available at the retreat and online.
- OceanStore
  - http://oceanstore.cs.berkeley.edu
- Other relevant material:
  - http://www.cs.berkeley.edu/~ravenben