Telematics
Computer Systems

Freie Universität Berlin

Towards Self-Organizing, Integrated Service Placement in Ad Hoc Networks

Georg Wittenburg and Jochen Schiller

Service Placement in Ad Hoc Networks

- Process of selecting a set of nodes in a network that is best suited to host instances of a service
  - Adapt to changing service demand and network topology
- Motivation / Benefits:
  - Service configuration, i.e., set of nodes to host a service, can be optimized automatically at run-time
  - Reduction of overall network traffic and latency
  - Support for optimization according to service-specific metrics and requirements, e.g., minimal number of instances

Subproblems

- *Where*, i.e., on which nodes, should service instances be placed?
- *How many* service instances are best suited for cost optimal operation?
- *When* should the current configuration of a service be changed?
- *How* can service instances be transferred efficiently between nodes?

Pitfalls

- Excessive control overhead
  - Information required on network topology and regional service demand
  - Signaling between service instances
- Cost of changing service configuration
  - Transfer of implementation and state
  - Service and route discovery induced by changing service hosts

Synchronization Between Service Instances

- Service configuration depends on service-specific requirements on synchronizing shared data between servers:

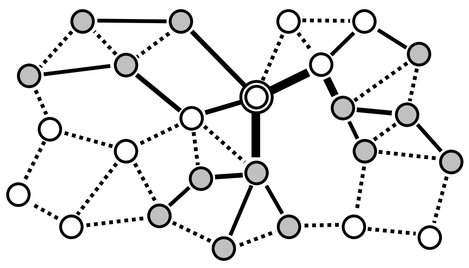

- Example #1: Transactional database
  - ➢ One single service instance

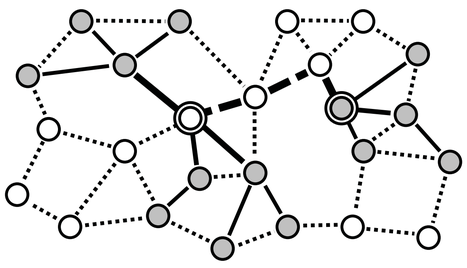

- Example #2: Directory service
  - ➢ Several service instances

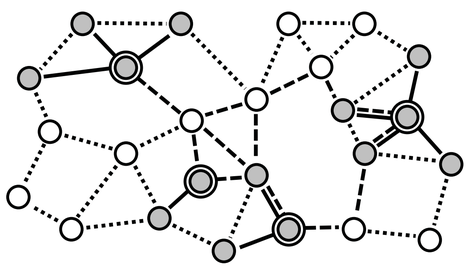

- Example #3: Web server
  - ➢ Multiple, per-cluster service instances

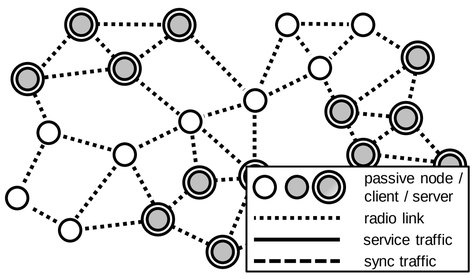

- Example #4: Spell checker
  - ➢ One service instance per node

The SP*i* Service Placement Architecture

- Service placement middleware (running on each server)
  - Collect and aggregate statistics on local/global service demand
  - Coordinate service placement process among current servers
- Enhanced routing protocol
  - Collect network topology information for middleware
- Enhanced service discovery protocol
  - Proactively inform clients about current service configuration

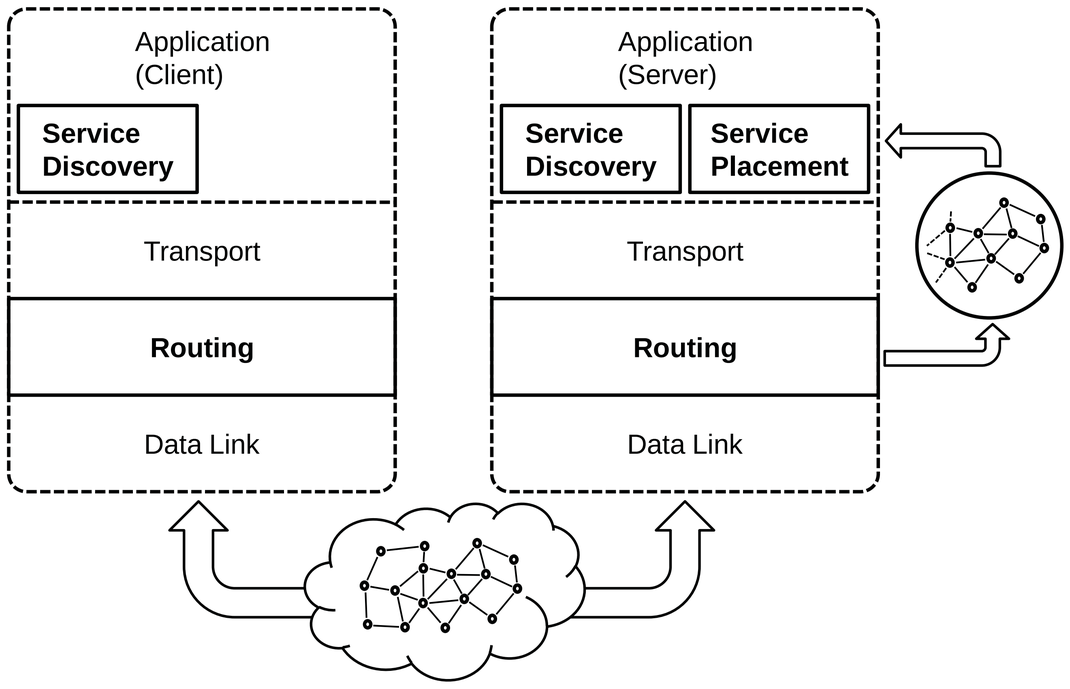

Service Placement Process in SP*i*

1) Aggregate information on network topology and regional service demand on dynamically assigned, service-specific coordinator node
   - Partial network topology provided by routing protocol
   - Statistics on past service demand provided by middleware on each host
2) Calculate optimal configuration by finding minimum of cost function
   - Cost function comprises network for service provisioning and synchronization
3) Decide whether to change service configuration (including coordinator node)
   a) Compare cost of current and optimal service configuration
   b) Compare savings of optimal configuration to estimated cost of adaptation
4) If adaptation is beneficial, issue commands to current service hosts
   - Commands: Replicate, migrate, or shutdown a service instance

Preliminary Results

- Results of IEEE 802.11 network with 100 nodes and median node degree of 8.0 simulated with ns 2.33
- Global minimum of service cost against number of service instances depends on synchronization requirements

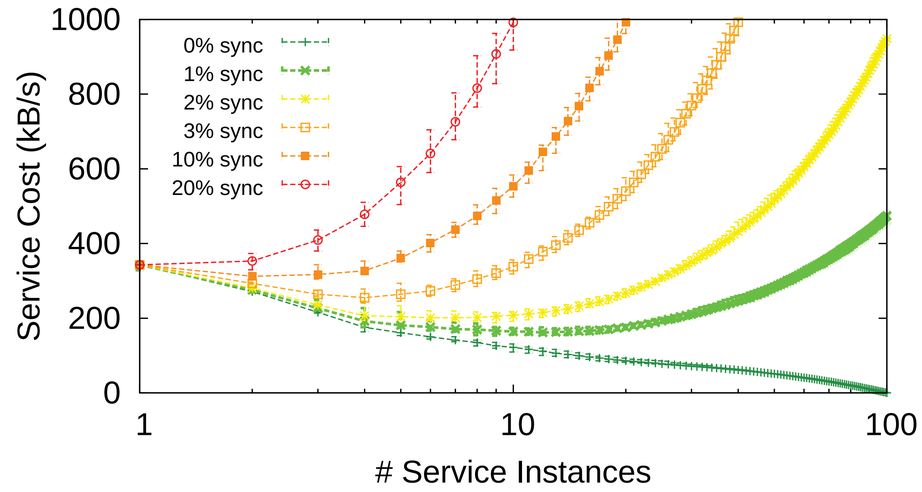

- Optimal number of service instances increases with decreasing synchronization requirements

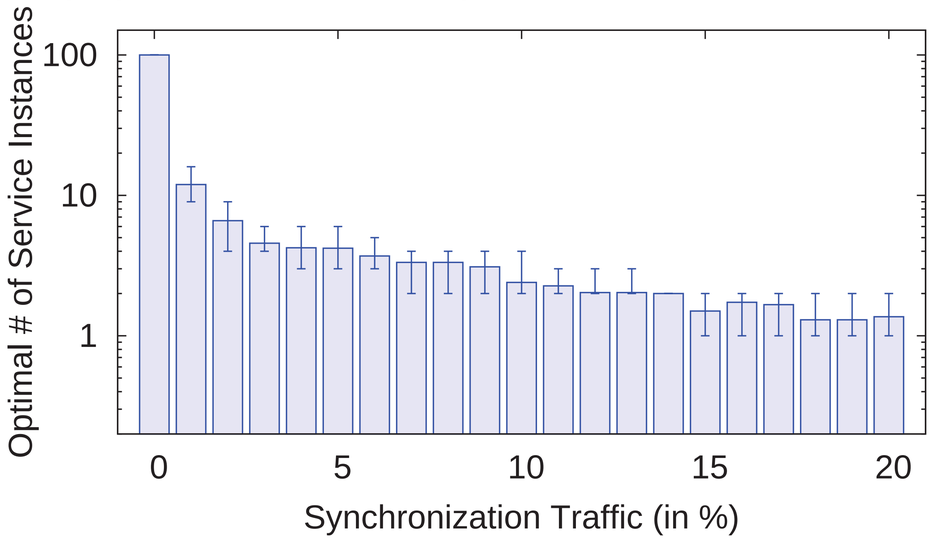

- Service placement with SP*i* is beneficial for services with synchronization traffic below 15%
  - Applicable to wide-spread services such as DNS and WWW