

Mobile IP and Multi-hop Wireless Networks

CMU CS 15-744: Computer Networks



Brad Karp

bkarp@cs.cmu.edu

22nd October, 2002

Overview: Large-Scale Wireless Systems

Small-Scale: How to build single-hop wireless LAN; how to make TCP perform well over it

Large-Scale: How to build multi-hop wireless systems (MANs? WANs?); how to support mobile nodes and users

- **Mobile IP:** How can a mobile user keep connections open and be reachable at a fixed address when roaming around the Internet?
- **Multi-hop ("ad hoc") wireless routing:** How do we find routes when the topology is highly dynamic, and when the network diameter is great?
- **Multi-hop wireless capacity:** How much user traffic can we carry on a large-scale, multi-hop wireless network?

Problem: Node Mobility

Model: single-hop wireless LANs at edges of the Internet; users with laptops, from some **home network**

Fundamentally, Internet architecture combines **end-system identifier** and **attachment-point identifier**


Address aggregation lets Internet routing scale: backbone routers store one table entry per destination **leaf subnet**

What restriction does aggregation place on IP address assignment? Can a node use a Berkeley IP address on CMU's LAN?

Result: machine cannot keep a fixed IP address when it moves!

Breaks: reachability of mobile nodes when roaming from home; long-lived connections from/to mobile nodes that span moves

Mobile IP: Entities



MH: Mobile Host, host that moves to different Internet leaf networks

HA: Home Agent, server on MH's home network

FA: (optional) Foreign Agent, server on network where MH roams

CH: Correspondent Host, any 3rd-party host that wishes to communicate with MH

Goal: Berkeley (128.32.112/24) MH retains Berkeley IP address while roaming at CMU (128.2/16)

Mobile IP: Triangular Routing

1. MH arrives at CMU; obtains **care-of address** in 128.2/16 (e.g., by DHCP)
2. MH sends **registration** to its HA at Berkeley, informing Berkeley HA of care-of-address
3. CH sends packet to MH's Berkeley IP address; HA intercepts, **encapsulates** to MH's care-of address
4. MH **decapsulates** CH's packet, receives packet from CH to MH's home IP address

Note **inefficiency**, asymmetry in CH-MH path

MH now has stable IP address for long-lived connections and servers

How does introduction of triangular routing weaken Internet security model?

Security and Mobile IP

Assumption: **Internet backbone routers deliver packets to their destinations** (if we don't have this, we don't have much . . .)

Introducing redirection of traffic at edges of Internet architecture lets attackers pervert redirection service, even if routers behave

Fix: use **shared-secret authentication** between MH and HA to authenticate MH-HA registrations, with **nonce replay protection**

IP address spoofing: a DDoS attacker forges an arbitrary IP address on a packet sent to a victim

Ingress filtering: to combat spoofing, backbone border routers look up *source* address S on packet from leaf network on interface I; refuse to forward it if forwarding table doesn't forward S on I

Impact of ingress filtering on packets originated by roaming MH?

"Fix:" **reverse tunnel** MH packets to HA from care-of address

Security and Route Optimization

If CH runs updated IP stack, we can avoid triangular route:

- CH sends initial packets to MH, via HA, as before
- HA sends CH **binding update** packet, informing CH of MH's care-of address
- CH encapsulates packets for MH via its care-of address
- MH can reverse-tunnel directly to CH, if necessary

Now how can an attacker redirect traffic without subverting backbone routers?

Why is shared-secret authentication between CH and HA not a good proposal?

Fix: Use **nonce** in packets from CH to HA; HA must repeat nonce in periodic binding cache updates to CH

Context and Evaluation: Mobile IP

Hot area starting from early 90s through 2000 (still holding on)

Proposed for IPv4, proposed for IPv6 (main differences: mandate optional features from v4 in v6; eliminate FA)

Multiple free implementations available (Monarch, NUS, SUNY Binghamton, DEC, . . .)

DHCP wins; Mobile IP deployment virtually nil! Why?


What's the driving application?

- Do you run servers on your laptop?
- HTTP uses short connections
- Do you do telephony on your laptop?
- What long-lived connections are inconvenient to restart?

Opinion: the world doesn't need another Mobile IP system

Multi-Hop Motivation: Rooftop Networks

Metropolitan-area network comprised of customer-owned and -operated radios: *Rooftop Networks*




An alternative architecture to single-hop cellular systems:

Self-organizing, rapidly deployable, potentially lower cost

Great demand already! Hardware ubiquitous; scalable algorithms for routing sorely needed

The Routing Problem



Packet-switched networks


End-to-end path: **route**

Each router chooses neighbor to which to forward received packet onward toward destination, D

Topology may be dynamic: **routes change**

Another Motivating Example

Vast wireless network of mobile temperature sensors, floating on the ocean's surface: *Sensor Networks*



Motivation (cont'd)

Enable three new classes of networks:

- **Ad-hoc networks**: mobile, infrastructureless, small-scale [Broch *et al.*, '98]
- **Sensor networks**: mobile, large-scale
- **"Rooftop" networks**: fixed, large-scale, no common administrative authority [Shepard, '96]

A mix of these characteristics:

- **Mobility**
- **Scale (number of nodes)**
- **Lack of static hierarchical structure**

Scalability Goals for Mobile, Wireless Routing

As number of nodes increases, and mobility rate increases:

- **Routing protocol message cost:** MINIMIZE
- **Application packet delivery success rate:** MAXIMIZE
- **Route length:** MINIMIZE
- **Per-node state:** MINIMIZE

Prior Work

Wired, Intra-Domain Internet Routing:

- Link-State (Dijkstra) and Distance-Vector (Bellman-Ford) routing on flat addresses to find **shortest** (in hops) **paths**
- Describe *entire* topology to *all* routers (LS) or push distances across network diameter (DV), for **$O(N)$ state** per router
- **Each link change** must be communicated to all routers to avoid loops and disconnection [Zaumen, Garcia-Luna Aceves, '91]

Ad Hoc Routing:

- Algorithms target low-bandwidth, high-mobility networks
- Many proposals (DSDV, DSR, TORA, AODV, GPSR, ZRP, ...)
- Diverse approaches: DV, source routing, geographic, proactive, on-demand ...

Ad Hoc Routing: DSDV

Destination-Sequenced Distance-Vector Routing:

- **Send increasing sequence number** with route advertisements
- **Greater seqno** takes precedence over lesser metric
- On detecting disconnection to D, router advertises route with **infinity metric and incremented seqno**
- D increments seqno on hearing advertisement with infinity metric
- Use **triggered updates** to propagate seqno increases rapidly and eliminate potentially looping routes

Ad Hoc Routing: DSR

Dynamic Source Routing:

- **On-demand routing:** only generate routing protocol traffic when forwarding requires it
- **Flood queries** to learn source routes
- **Cache replies**
- **Source routes break more frequently** as mobility and network diameter increase; **caching steadily less effective**
- **Which metrics are Broch *et al.* interested in?**
Which do they omit?
Exploration of limits of DSR?

Prior Work: Scaling

Dominant factors in scaling of DV, LS, DSR algorithms:

- Rate of change of topology
- Number of routers in the routing domain

Scaling strategies:

- Hierarchy: at AS boundaries (BGP) or on a finer scale (OSPF)

Goal: Reduce number of nodes in a routing domain

Assumptions: Level boundaries relatively fixed; administrative authority can choose level boundaries

- Caching: Store source routes overhead (DSR)

Goal: Limit propagation of future source route queries

Assumption: Source route remains fixed while cached

Assumptions invalid for highly mobile or unstructured networks!

Geography

Central Idea: Machines can know their geographic locations.
Route using GEOGRAPHY.

Established positioning methods:

- GPS outdoors (single chip, low-cost)
- Surveying (stationary routers)
- Inertial sensors (vehicles)
- Acoustic and radio range-finding (indoors, [AT&T Cambridge, 1997], [Priyantha *et al.*, 2000])

Efficient node location lookup/registration system [Li *et al.*, 2000]

All nodes know own position; packet source marks packet with destination's location

Assumptions

Bi-directional radio links (e.g., IEEE 802.11 with link-level acknowledgements)

Network nodes placed roughly in a plane

Radio propagation in free space; distance from transmitter determines signal strength at receiver (*two-ray ground reflection model*)


Fixed, uniform radio transmitter power

Greedy Forwarding

Nodes learn immediate neighbors' positions through beacons/piggybacking on data packets

Locally optimal, **greedy** forwarding choice at a node:

Forward to the neighbor geographically closest to the destination



In Praise of Geography

Self-describing

As node density increases, shortest paths through wireless networks correspond increasingly to Euclidean straight line between source and destination


Each node's state concerns only immediate neighbors:

- Tiny per-node state
- Routing protocol pushes state only one hop—tiny routing protocol overhead
- Local forwarding decisions—robust to topology changes

Compare with lookup in $O(N)$ table under DV, LS

Greedy Forwarding Failure


Greedy forwarding not always possible! Consider:



Voids


When the *intersection* of a node's circular radio range and the circle about the destination on which the node sits is empty of nodes, greedy forwarding is impossible

Such a region is a **void**:



Node Density and Voids

Existing and Found Paths, 1340 m x 1340 m Region




The probability that a void region is empty of nodes increases as nodes become more sparse

Void Traversal: The Right-Hand Rule

Well-known graph traversal: **right-hand rule**:


When arriving at x from y , move to the next vertex counterclockwise about x from y



Traverses interior faces in clockwise edge order; exterior faces in counterclockwise edge order

Planar vs. Non-Planar Graphs

The right-hand rule may not tour enclosed faces on graphs with edges that cross (*non-planar graphs*)




Seek a distributed algorithm that removes crossing edges without partitioning the network, using only neighbors' positions as input

Planarized Graphs

Relative Neighborhood Graph (RNG) [Toussaint, '80] and **Gabriel Graph (GG)** [Gabriel, '69] are long-known planar graphs

Assume an edge exists between any pair of nodes separated by less than a threshold distance (*i.e.*, the nominal radio range)

RNG and GG can be constructed using only neighbors' positions, and can be shown not to partition the network!




RNG

GG

Planarized Graphs: Example

200 nodes, placed uniformly at random on a 2000-by-2000-meter region; radio range 250 meters



Full Network

GG Subset

RNG Subset

Full Greedy Perimeter Stateless Routing

All packets begin in greedy mode


Upon greedy failure, node marks its current location in packet, and marks packet in perimeter mode

Perimeter mode packets follow simple planar graph traversal:

Forward along successively closer faces by right-hand rule, until reaching destination, or node closer to it than perimeter mode entry point

Return packets to greedy mode when they reach a node closer than their perimeter mode entry point

Perimeter Mode Forwarding Example



Traverse face closer to D along \overline{xD} by right-hand rule, until reaching the edge that crosses \overline{xD}

Repeat with the next closer face along \overline{xD} , &c.

Record first edge on face to **detect disconnection**

GPSR: Protocol Techniques for Dynamic Networks

Use of MAC-layer failure feedback: As in DSR [Broch, Johnson, '98], interpret retransmit failure reports from the 802.11 MAC as indication a neighbor has gone out-of-range

Interface queue traversal and packet purging: Upon MAC retransmit failure for a neighbor, walk the interface queue and remove packets to that neighbor to avoid head-of-line blocking of 802.11 transmitter during retries on those packets

Promiscuous network interface: Reduce beacon load and keep positions stored in neighbor tables current by tagging all packets with the forwarding node's position

Planarization triggers: Re-planarize upon acquisition of a new neighbor and every loss of a former neighbor, to keep planarization up-to-date as topology changes

Simulation Environment

ns-2 with wireless extensions [Broch *et al.*, 1998]: full 802.11 MAC, physical propagation; allows comparison of results

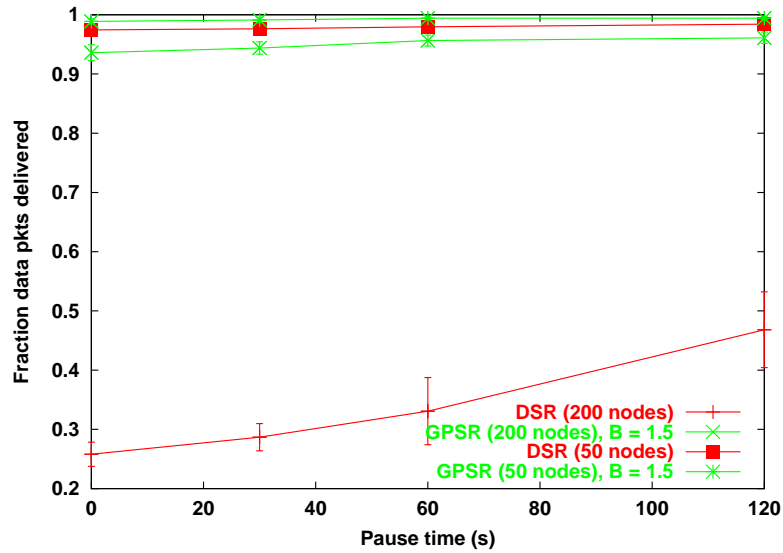
Topologies and Workloads:

Nodes	Region	Density	CBR Flows
50	1500 m × 300 m	1 node / 9000 m ²	30
200	3000 m × 600 m	1 node / 9000 m ²	30
50	1340 m × 1340 m	1 node / 35912 m ²	30

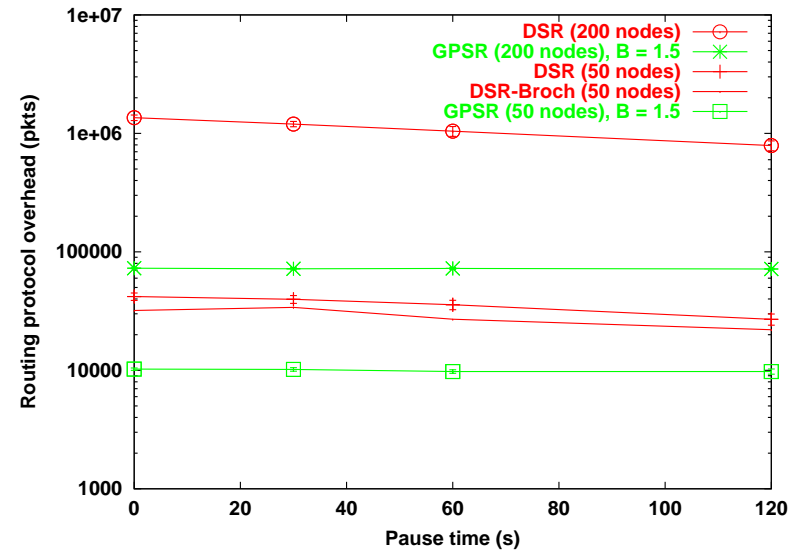
Simulation Parameters:

Pause Time: 0, 30, 60, 120 s	Motion Rate: [1, 20] m/s
GPSR Beacon Interval: 1.5 s	Data Packet Size: 64 bytes
CBR Flow Rate: 2 Kbps	Simulation Length: 900 s


Packet Delivery Success Rate (50, 200; Dense)




Routing Protocol Overhead (50, 200; Dense)



Path Length (200; Dense)

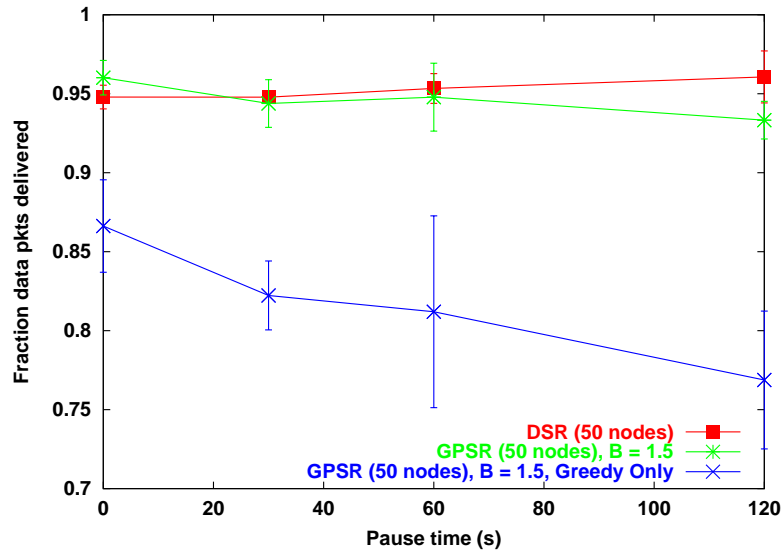


State Size (200; Dense)

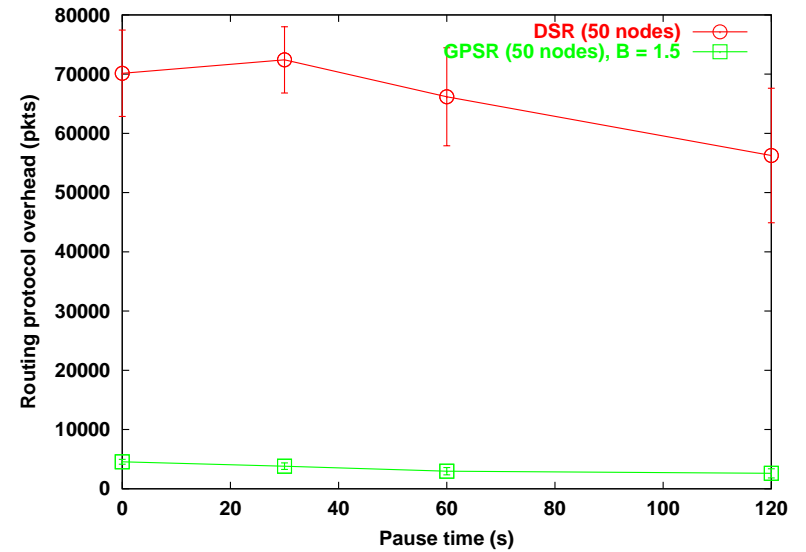


GPSR requires state proportional to node density; DSR stores state at each router proportional to the sums of the lengths of source routes


Packet Delivery Success Rate (50; Sparse)



Routing Protocol Overhead (50; Sparse)



Path Length (50 nodes, Sparse)



Shepard: Rooftop Wireless MAN Scaling

MACAW contention model: propagation to a fixed distance only; focus on floor acquisition among mutually near stations

Shepard's contention model: propagation to "radio horizon", far greater than successful communication distance; focus on S/N ratio, effects of distant transmitters

Fundamental observation: there are many more distant stations than near ones; interference from them is greater concern than nearby collisions

Feasibility of minimum-energy routing?

Feasibility of hundreds of hops, or more?

Very useful concept: bisection bandwidth

Shepard's conclusion: scaling to millions of nodes possible where nodes can still communicate with nearby neighbors at a high rate

Capacity of Ad Hoc Wireless Networks

Context: Mobicom 2001, on the heels of years of ad hoc routing research, nearly exclusively in simulation

Goals:

- Explain details of 802.11 MAC when used for forwarding, as regards network capacity
- Provide simple model for capacity of ad hoc networks, as related to traffic matrix

Fundamental phenomenon: nodes use their own one-hop transmission capacity not only for data they originate, but also for data they forward

Ad Hoc Capacity: Intuition

Some depressing intuition:

- Spatial reuse lets distant radios transmit simultaneously, as they don't interfere
- For constant node density, one-hop capacity, sum of all single-hop transfer rates possible in the network, grows as $O(n)$
- As network diameter grows, for random source/destination pairs, average path length grows as $O(\sqrt{n})$
- Total end-to-end capacity: $O(n/\sqrt{n})$, and so per-node capacity is $O(1/\sqrt{n})$.

Throughput per node approaches zero as number of nodes increases!

Forwarding and 802.11

The energy required to garble another's transmission is far less than that required to be received properly

Interfering range is 550 m, while transmission range is 250 m

What's the best throughput we can expect in a chain


$A \rightarrow B \rightarrow C \rightarrow D$, if ranges were equal?

- A and B can't transmit simultaneously; nor can B and C ; nor can A and C
- Best throughput: 1/3 link rate

With 550 m interference range, best throughput drops to 1/4; now D interferes with A 's transmissions to B

In simulations of greedy 802.11 senders arranged in a chain, throughput is closer to 1/7 than 1/4; boundary effect

Forwarding and 802.11 Backoff



$D \rightarrow E$ will clobber $A \rightarrow B$

Yet A doesn't know of D 's transmission

Result: repeated exponential backoff by A

Note similarity with MACAW four-hop, left-to-right example

Traffic Matrix and Multi-Hop Wireless Capacity

Capacity available to each node inversely related to expected flow physical path length

Traffic matrix typically studied in ad hoc routing: **uniformly randomly selected flow endpoints**

Expected path length for a uniform random traffic pattern on a network of area A : $2\sqrt{A}/3$

For n nodes and fixed node density, $A \propto n$

So the capacity available to each node is $O(1/\sqrt{n})$

Perhaps **this** is why published ad hoc routing studies use ca. 60 Kbps total application traffic workloads!

Power-Law Traffic Patterns and Capacity

Power-law traffic patterns, where probability of communication with node x distance away is given by $x^{-\alpha}$, offer **constant per-node capacity**

For $\alpha = 2$, expected communication distance scales as $O(\log_2 A)$

A useful design rule for systems for multi-hop wireless networks, e.g., GLS location database [Li *et al.*, '00]

Power-law construct makes analysis tractable; meaning is intuitively useful

Evidence of power-law communication patterns in the wild?