


Mobilize!

David Tse
Department of EECS, U.C. Berkeley

September 23, 2003

University of Toronto

Mobility Induces Time Variations



- **small** spatial-scale fading due to constructive and destructive interference between multipaths;
- **large** time-scale variation due to shadowing effects and varying path loss: change in network topology.

Role of Mobility

Mobility has traditionally been viewed as adding **complexity** to the design of wireless networks:

Examples:

- fading counter-measures
- cellular handoffs
- frequent location and route updates

Role of Mobility

Mobility has traditionally been viewed as adding **complexity** to the design of wireless networks:

Examples:

- fading counter-measures
- cellular handoffs
- frequent location and route updates

In this talk, I would like to convince you that mobility can in fact be **taken advantage of**.

Talk Outline

To support my claim, I will survey results in three topics:

- scheduling in fading wireless links
- capacity of large-scale mobile ad hoc networks
- last encounter routing


Talk Outline

To support my claim, I will survey results in three topics:

- scheduling in time-varying wireless links
- capacity of large-scale mobile ad hoc networks
- last encounter routing


Motivation: HDR Downlink

HDR (1xEV-DO): a 3G wireless data standard operating on IS-95 band (1.25 MHz)




- HDR downlink operates on a time-division basis.
- Scheduler decides which user to serve in each time-slot.

Opportunistic Communication



Channel variations can be exploited by scheduling transmissions to the user with the good channel.


Multiuser Diversity



- In a large system with users fading independently, there is likely to be a user with a very good channel at any time.
- Long term total throughput can be maximized by always serving the user with the **strongest** channel.


(Knopp and Humblet 95)

Fairness and Delay




Challenge is to exploit multiuser diversity while sharing the benefits **fairly** and **timely** to users with **asymmetric** channel statistics.

Hitting the Peaks



- Want to serve each user when it is near its **peak** within a latency time-scale t_c .

Hitting the Peaks



- Want to serve each user when it is near its **peak** within a latency time-scale t_c .
- In a **large** system, at any time there is likely to be a user whose channel is near its peak.

Proportional Fair Scheduler

(Tse 99)

At time slot t , given

1) users' average throughputs $T_1(t), T_2(t), \dots, T_K(t)$ in a past window

2) current requested rates $R_1(t), R_2(t), \dots, R_K(t)$

transmit to the user k^* with the largest

$$\frac{R_k(t)}{T_k}.$$

The past window can be made equal to t_c to match the latency requirement.

Theoretical Property

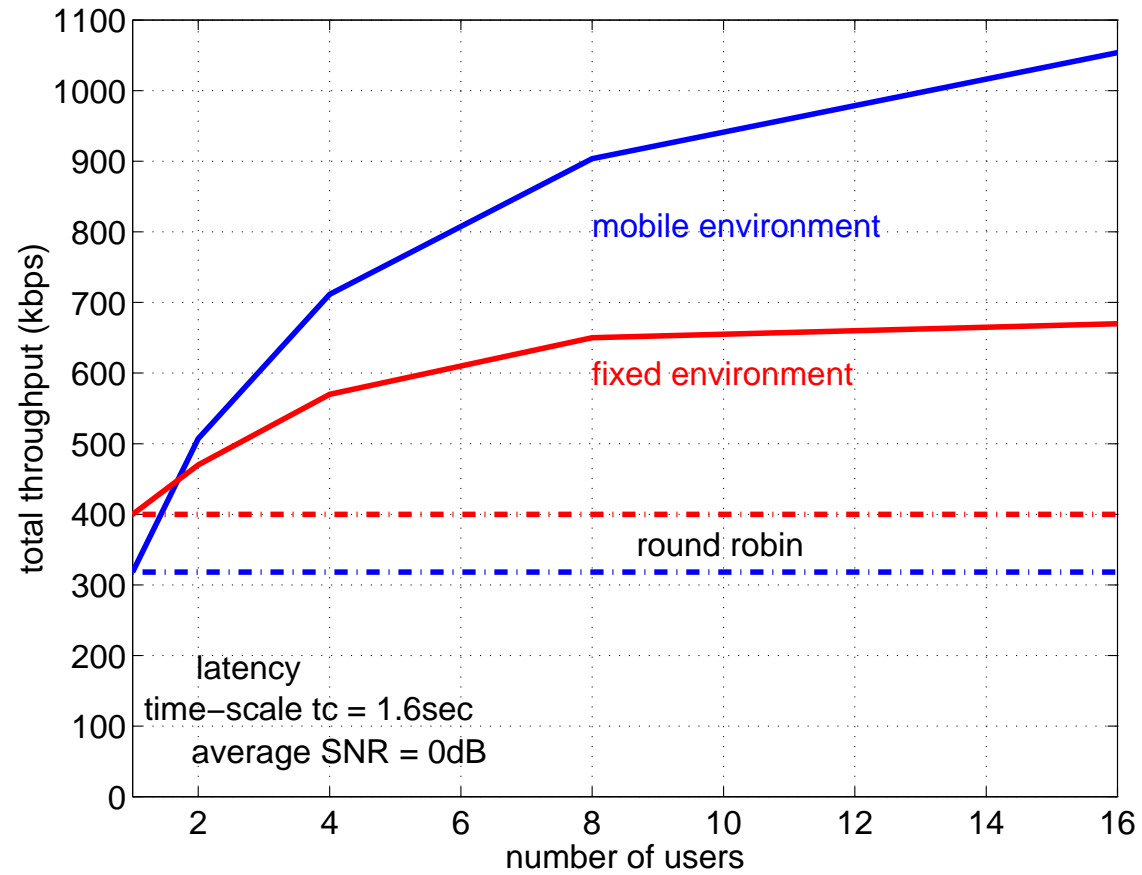
Under stationary assumptions and $t_c = \infty$, long-term average throughputs T_1^*, \dots, T_K^* of the scheduler maximizes

$$\sum_k \log T_k$$

among all schedulers, i.e. **proportional fair**.

The scheduler can be viewed as a stochastic gradient ascent algorithm to solve the optimization problem.

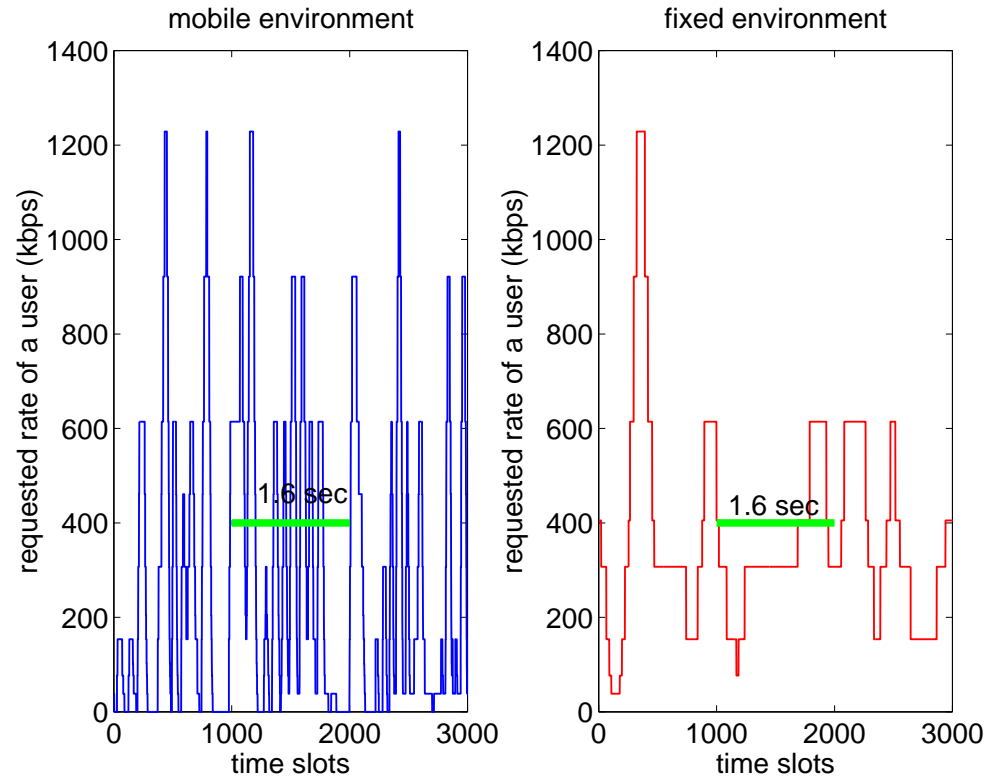
Throughput of HDR Scheduler



Mobile environment: 3 km/hr, Rayleigh fading

Fixed environment: 2Hz Rician fading with $E_{\text{fixed}}/E_{\text{scattered}} = 5$.

Channel Dynamics



Channel varies faster and has more dynamic range in mobile environments.

In typical HDR operating environments, throughput gains of 50% to 100% are common.


Scheduling for Inelastic Traffic

Proportional fair scheduling is good for **elastic** traffic: it allocates bandwidth to users only as a function of channel conditions and not users' demand.

Multiuser diversity scheduling algorithms have been proposed for **inelastic** traffic. (Shakkotai and Stoytar 01)

Scheduling is based on a metric as a function of both the **queue length** as well as **channel state**.

Throughput optimality is shown.




So far we have exploited small spatial scale mobility

How about the large scale mobility?

Talk Outline

- scheduling in time-varying wireless links
- capacity of large-scale mobile ad hoc networks
- last encounter routing

Scalability of Ad Hoc Networks



Point-to-point traffic: Suppose each node has a stream of traffic for a particular destination node.

Gupta and Kumar: Throughput per source-destination pair goes to zero like $1/\sqrt{n}$ with the number of nodes n per unit area.

Result assumes nodes stay fixed for the duration of communication.

What about if the nodes are mobiles?

Mobility Can Help!

Main result: (Grossglauser and Tse 01)

Suppose nodes move randomly and independently.

A long-term throughput of $O(1)$ per S-D pair can be achieved.....

Mobility Can Help!

Main result: (Grossglauser and Tse 01)

Suppose nodes move randomly and independently.

A long-term throughput of $O(1)$ per S-D pair can be achieved.....


.....if one is willing to wait.

Throughput is averaged over the time-scale of mobility.

Idea #1


Communicate only when the source and destination are nearest neighbors to each other.

Direct Communication Does Not Work




- The source and destination are nearest neighbors only $O(1/n)$ of the time
- In fact, can show S-D throughput is at most $O\left(n^{-\frac{1}{1+\alpha/2}}\right)$ for **any** policy that does not use relays.

Phase I: Source to Relays




- At each time slot, source relays a packet to nearest neighbor
- Different packets are distributed to different relay nodes.

Phase II: Relays to Destination



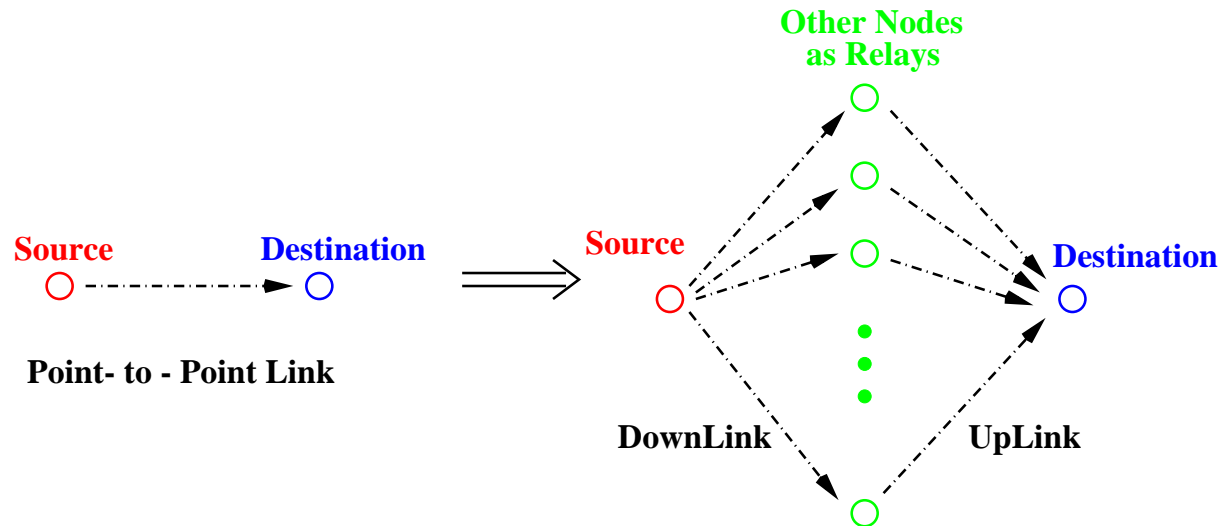
- Steady state: all nodes have packets destined for D
- Each relay node forwards packets to D only when it gets close.

Phase I and II Staggered



- $O(1)$ throughput from S to D
- Communication is confined to nearest neighbors, but each packet goes through at most two hops.
- In contrast, when nodes are fixed, $O(\sqrt{n})$ hops are required. (Gupta and Kumar)

Multiuser Diversity via Relaying



- Multiuser diversity created artificially using all other nodes as relays.
- Channel variation comes from **large** rather than **small** spatial-scale mobility.

Network Capacity


- The above discussion pertains to a **single** source-destination pair.
- It turns out that every S-D pair can follow the same strategy **simultaneously**.
- Key fact to show: Under a model for power law decay in interference, $O(n)$ simultaneous nearest neighbor communication **is** possible. (full spatial reuse)

Restricted Mobility

In the mobility model, every node wanders all over the domain.

What happens when each node's mobility is restricted?

One-Dimensional Mobility



Suppose the domain is now the surface of a sphere.

Each node moves randomly on a fixed great circle on the sphere, a circle for each node.

Traffic Bottlenecks

High throughput **cannot** be attained for all configurations of great circles.

Example: $n/2$ nodes on the same great circle, and $n/2$ nodes on another great circle.

Nodes on the same great circle are nearest neighbors with probability $O(1/n)$, but nodes on different great circles with probability only $O(1/n^2)$.

The intersections of the great circles become bottlenecks in conveying traffic between nodes in the two circles.

In contrast each node spends the same order of time as the nearest neighbor to every other node in the original model.

Random Configurations

Nevertheless.....

Theorem: (Diggavi, Grossglauser and Tse 02)

Suppose the great circle of each node is independently and uniformly chosen on the sphere. Then there exist a constant $c > 0$ such that the throughput per pair c is feasible for almost all configurations as $n \rightarrow \infty$.

Basically, we get $O(1)$ throughput on “typical” configurations.

“Mobilizing” Other Architectures

- Infostations (WINLAB) : a network of fixed “gas stations” providing pockets of high speed short-range coverage.
- Suppose they are used as a **content distribution** network to deliver information to mobiles.
- **Classic** Infostations: mobiles can only download information when they are close to the fixed Infostations.
- **Mobile** Infostations: mobiles can also serve as Infostations, exchanging information when they are close to each other.

Mobile Infostations

(Yuen, Yates and Mau 03)

- They showed that if all users want the same content (multicast), the throughput per user is increased by a factor of K , the number of mobiles per Infostation.
- Again, a **multiuser diversity** effect: users can now get data from the other mobiles in addition to the Infostations.
- In general, there is significant gain as long as there is sufficient “common interest” among users.

Mobile Infostations

(Yuen, Yates and Mau 03)

- They showed that if all users want the same content (multicast), the throughput per user is increased by a factor of K , the number of mobiles per Infostation.
- Again, a **multiuser diversity** effect: users can now get data from the other mobiles in addition to the Infostations.
- In general, there is significant gain as long as there is sufficient “common interest” among users.
- **Non-cooperative setting:**
Similar gain still holds if mobiles exchange information only when each is getting something new.

Recap

General principle:

Mobility provides a mechanism to propagate information without costly “over the air” communication.

This principle can be applied to other problems.

Talk Outline


- scheduling in time-varying wireless links
- capacity of large-scale mobile ad hoc networks
- last encounter routing

Geographical Routing

- In geographical routing, packets can be routed directly to destination node based on its location.
- Each node knows its own current position, but the information has to be conveyed to the source nodes.
- This requires a **location service**.
- In an ad hoc network, this is typically implemented by flooding and continuously updating location information across the network.
- But in fact mobility of nodes provides an alternative (and cheaper) means to diffuse the information.


Last Encounter Routing

(Grossglauber and Vetterli 03)




Location information is diffused by nodes remembering time and location of last encounters.

Mobility Diffuses Location Information



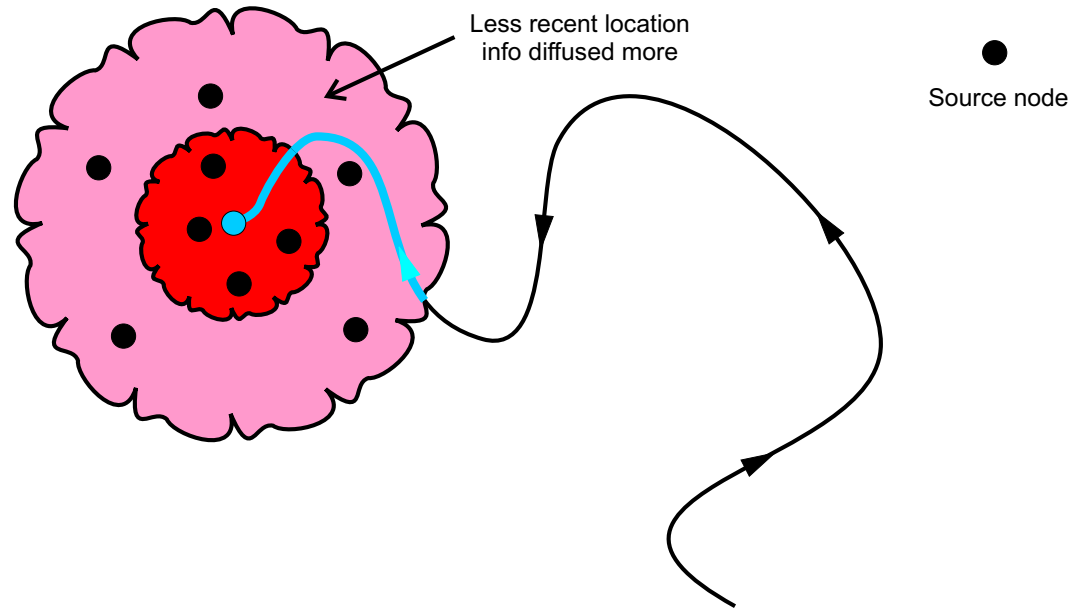
Source asks: where is the destination node?

Mobility Diffuses Location Information




Precise location information has only diffused over a limited area.

Mobility Diffuses Location Information



Precise location information has only diffused over a limited area.
Crude information available over a larger area.

Mobility Diffuses Location Information




Precise location information has only diffused over a limited area.

Crude information available over a larger area.

Distance effect: Crude information is sufficient when the packet is far away.

Mobility Diffuses Location Information



Precise location information has only diffused over a limited area.

Crude information available over a larger area.

Distance effect: Crude information is sufficient when the packet is far away.

This calls for a **successive refinement** strategy.

Routing Protocol (EASE)

Current time is 0.


Let T_0 be the time of last encounter of source node with destination node D .

Flood neighboring nodes until finding a node which last encountered D at time later than $T_0/2$.

Route packet directly to location of that last encounter. Set T_1 to be the time of that last encounter.

Repeat until finding the destination node.

EASE in Action



Scaling Property

Suppose the domain is a \sqrt{n} by \sqrt{n} grid, and there are $O(n)$ mobile nodes.

Each node moves according to a 2-D random walk on the grid.

For a random source-destination pair, let

$$C_n = \text{cost of flooding} + \text{routing}.$$

Scaling Property

Suppose the domain is a \sqrt{n} by \sqrt{n} grid, and there are $O(n)$ mobile nodes.

Each node moves according to a 2-D random walk on the grid.

For a random source-destination pair, let

$$C_n = \text{cost of flooding} + \text{routing}.$$

Theorem: (Grossglauser and Vetterli 03)

$$\mathcal{E}[C_n] = O(\sqrt{n}).$$

Since the direct route takes $O(\sqrt{n})$ hops on the average, last encounter routing is no more than a constant factor worse.

How much Searching Per Iteration?

Currently: distance $O(\sqrt{|T_k|})$ from the destination


Objective: find a node which has seen the destination during $[T_k/2, 0]$.

Facts:

- Each such “messenger” node is about $O(\sqrt{|T_k|})$ from the destination.
- There are $O(|T_k|)$ of them.

$\Rightarrow O(1)$ of them per unit area.

Simulation Results



Conclusion

We discussed how several wireless problems can be “mobilized”:

- scheduling
- routing
- network architectures

Hopefully this point of view will inspire other ways of using mobility.