Mobile P2P Databases

Ouri Wolfson
University of Illinois, Chicago

Environment

Pda's, cell-phones, sensors, hotspots, vehicles, with short-range wireless capabilities

A central server does not necessarily exist

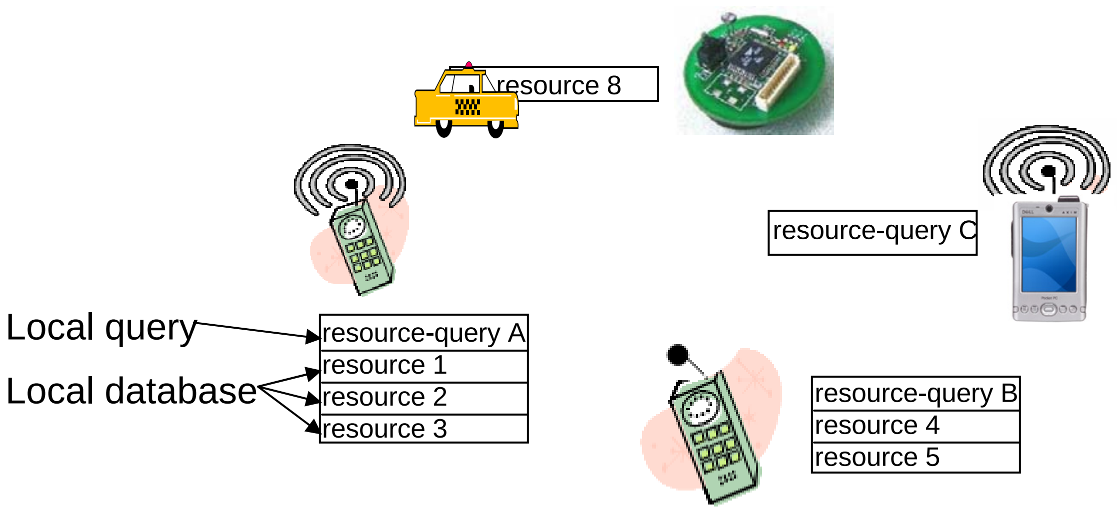

Short-range wireless networks
- wi-fi (100-200 meters)
- bluetooth (2-10, popular)
- zigbee

Unlicensed spectrum (free)

High bandwidth

3 communication modes
- listening
- transmitting (amount)
- receiving (amount)

Power/search tradeoff

"Floating database"

Resources of interest
- in a limited geographic area
- possibly for short time duration

Applications coexist

Mobile Local Search: applications

- social networking (wearable website)
  - Personal profile of interest at a convention
  - Singles matchmaking
  - Games
  - Reminder
- mobile electronic commerce
  - Sale on an item of interest at mall
  - Music-file exchange
- transportation
- emergency response
  - Search for victims in a rubble
- military
  - Sighting of insurgent in downtown Mosul in last hour
- asset management and tracking
  - Sensors on containers exchange security information => remote checkpoints
- mobile collaborative work
- tourist and location-based-services
  - Closest ATM

Transportation

- Safety
  - Vehicle in front has a malfunctioning brake light
  - Vehicle is about to run a red light
  - Patch of ice at milepost 305
  - Vehicle 100 meters ahead has suddenly stopped

Transportation (cont.)

- Improve efficiency/convenience/mobility:
  - What is the average speed a mile ahead of me?
  - Are there any accidents ahead?
  - What parking slots are available around me?
  - Taxi cab: what customers around me need service?
  - Customer: What Taxi cabs are available around me?
  - Transfer protection: transfer bus requested to wait for passengers
  - Cab/ride sharing opportunities

Applications – Common features

- Mobile/stationary peers
- Resources of interest
  - in a limited geographic area
  - Short time duration
- Can be solved by fixed servers, but
  - Unlikely solution
  - Proposed mp2p paradigm can enhance fixed solution (reliability, performance, coverage)

Mobile P2P Network

- Not really sensor networks
  - Resource constraints +
  - Reliability, mobility, data collection –
- Not really P2P
  - Incentives +
  - Mobility matters a lot (DHT probably not applicable) –
- Not really Mobile Ad Hoc Networks
  - Dynamic topology +
  - Network maybe sparsely connected (DTN) -
  - Routing to known net-id's –

Mp2p vs. client-server

- Mp2p advantages
  - Zero cost
    - Unregulated communication
    - No central database to maintain
  - Independent of infrastructure
  - Higher reliability
  - Privacy preservation
- Mp2p disadvantages
  - Weaker answer-completeness guarantees
  - Density requirements

How to enable Mobile P2P applications?

- Develop a platform for building them

Platform components

- Communication system: Intra-vehicle, vehicle-to-vehicle, and vehicle-to-infrastructure
  - Prototypes: Cartalk, UC Irvine
- Data Management: collect, organize, integrate, model, disseminate, query data
- Software tools:
  - Data mining
  - Travel-time prediction
  - Trip planning
  - Regional planning
  - Competitive resource discovery strategies
  - ……

Problems in data management

- Query processing
- Dissemination analysis
- Value of information, comparison with Client/Server
- Participation incentives
- Remote Querying
- Data Integration of sensor and higher level information (maps, trip plans, ride-sharing profiles)
- Related approaches

Query Processing outline

- Basics of MP2P search
- Data push and pull
- Power, memory, bandwidth management
- Ranking of information
- Experimental/commercial projects, Performance metric, and comparison
- Backchannel exploitation

The players

- Moving/stationary peers
  - Personal digital assistants (pda's)
  - Computers in vehicles
  - Processors embedded in the infrastructure
  - Hotspots

Collect, Organize, Disseminate, information about resources

- Spatio-temporal resource types -- examples
  - Gas stations
  - Parking slots
  - Cab-customers
  - Ride-share partners
  - Malfunctioning brake-light
  - Accidents

  - Tickets availability at events
  - Disaster victims
  - Matching profile

Resource-reports: description of (mostly physical) resources

- Each resource-report may include:
  - *create-time*
  - *home-location*

- 100 to 50,000 bytes (image)

- some resources alternate: available/unavailable

  → reports have limited useful life

Peers

- Each peer may serve as:
  - Producer of reports: sensor (rfid)
  - Consumer of reports: has queries
  - Broker of reports: has local broker database

- or combination of roles

Queries

- A query Q maps each report R to a match degree:

$$0 \le match(R, Q) \le 1$$

- Examples:
  - Top parking slots given my current location
  - Profile with expertise "children-periodontics"
  - Similarity between two images

Query/report Dissemination

- two peers within transmission range exchange queries and reports

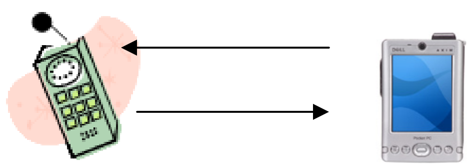

- Least relevant reports that do not fit in local broker database are purged
- Exchange not necessarily synchronous (periodic broadcast)

Avoid duplication

- Send only reports unknown to peer:
  - Saves transmission power/bandwidth
  - Saves receiving power
- Two types of duplication
  - Key
  - Content (same patch of ice reported 2 vehicles)
    - Bloom filters?

Synchronous vs. asynchronous

- When to communicate?

- when encounter
  - neighborhood awareness, heartbeats
  - what if no new encounters, i.e. peers static
- When new reports received
  - Broadcast new id's
  - Neighbors respond with id's of interest (collisions?)
- Blindly, Periodically
  - What if nobody around?

Saving listening energy

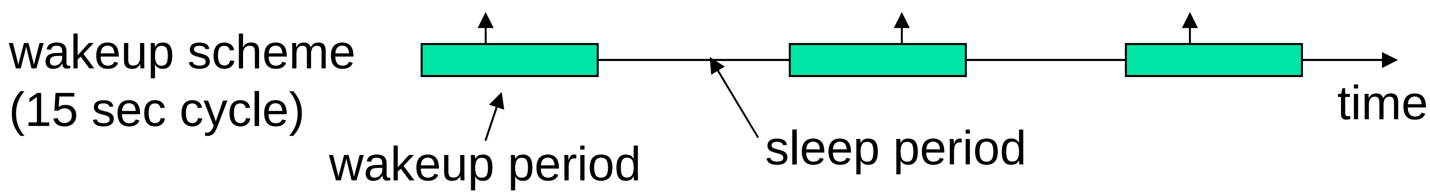

- Synchronization by rendezvous (same wakeup time)

Broadcast vs. unicast

- Broadcast of reports saves bandwidth since all neighbors hear
- Broadcast may waste receiving power
- Unicast → neighborhood awareness

Query Processing outline

- Basics of MP2P search
- Data push and pull
- Power, memory, bandwidth management
- Ranking of information
- Experimental/commercial projects, Performance metric, and comparison
- Backchannel exploitation

Push/Pull taxonomy

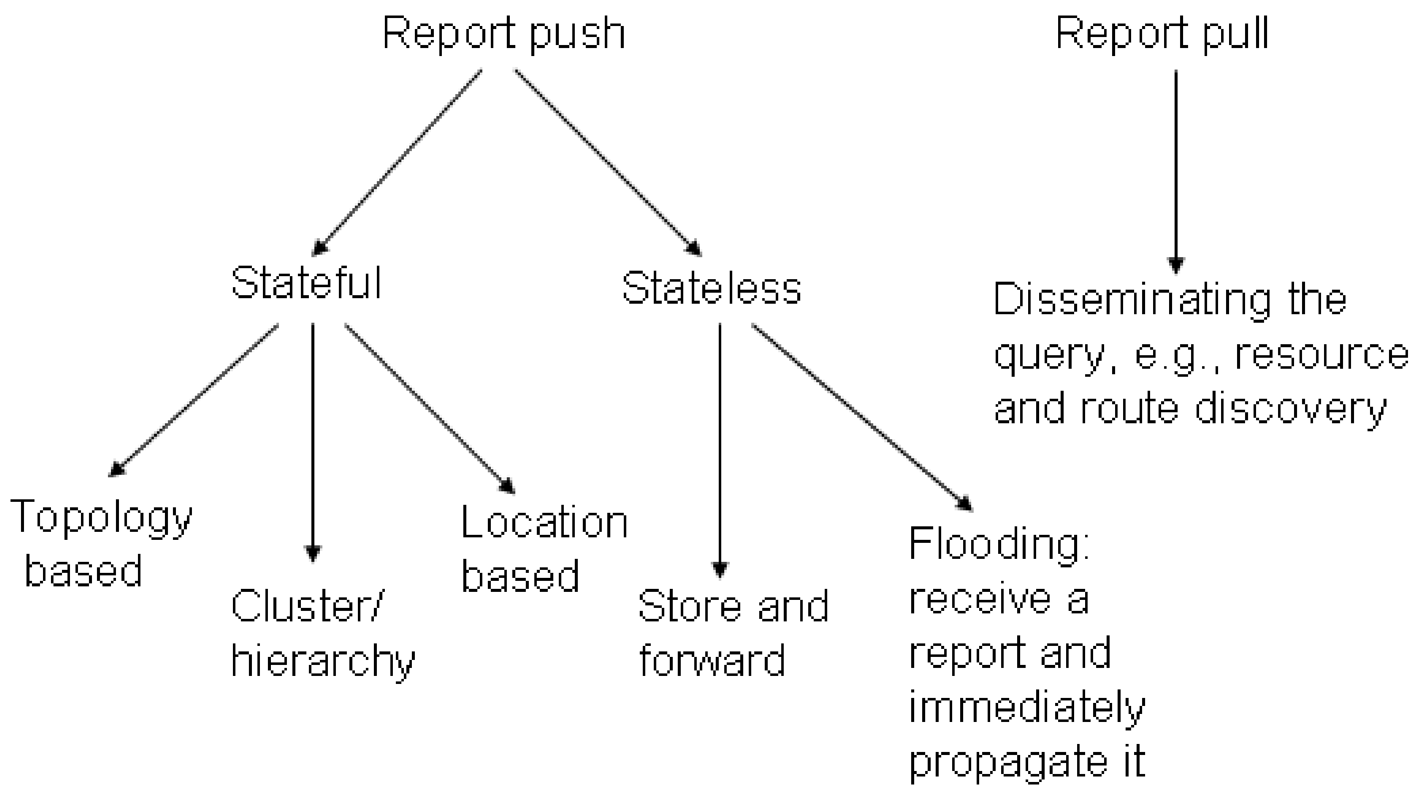

Report Pulling

- peer queries reports
- the whole network is flooded with query
- reports of interest will be pulled from peers that have them

- Widely used in resource discovery
  - route discovery in mobile ad hoc networks

Pull Example: Resource Discovery in MANET

A MANET routing protocol is augmented to enable addressing based on resource type or resource key rather than network ID

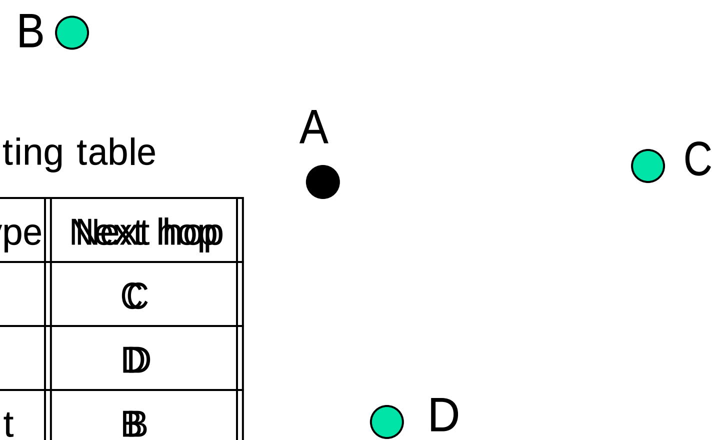

A's routing table

Ip address	Next hop
1234	C
2345	D
3456	B

Resource type	Next hop
printer	C
music	D
restaurant	B

Pull Example: Construction and Maintenance of Routing Table

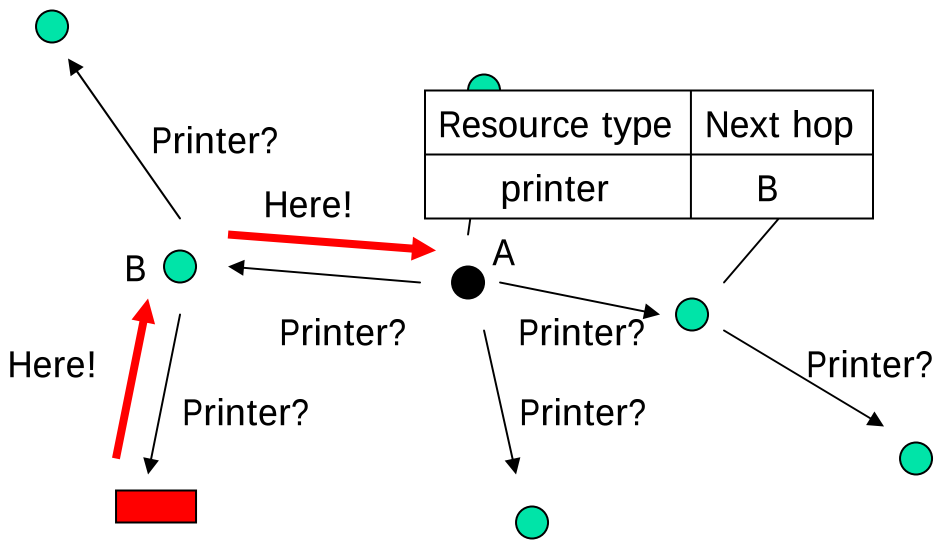

Problems when applied to our context:

- Does not work when consumer and resource are disconnected.
- Resources are transient. Consumer has to constantly poll.
- Constructed routing structure easily becomes obsolete.
- May take awhile to construct in Bluetooth networks

Report Pushing

- Report pushing
  - the dual problem of report pulling
  - reports are flooded
  - consumed by peers; query is answered by received reports
- Different types of pushing
  - broadcast: to the complete network
  - geocast: to a specific geographic area
  - unicast: to any one specific mobile node

push/pull in distributed systems

- Depends on ratio of queries/update
- Continuous queries (subscriptions)
- Extension to Mobile p2p ?

Pushing: Stateless or Stateful?

- Stateful methods: maintain global information
  - topology-based approaches
  - cluster- or hierarchy-based approaches
  - location-based approaches
- Stateless methods
  - flooding-based approaches
  - gossiping-based approaches
  - negotiation-based approaches
  - <u>store-and-forward approaches</u>

Push, stateful, hierarchy: Publish/Subscribe tree (Huang, Garcia-Molina 2004)

- Resource-reports are propagated along the arcs of a rooted tree
  - Problem: high mobility → tree obsolete.

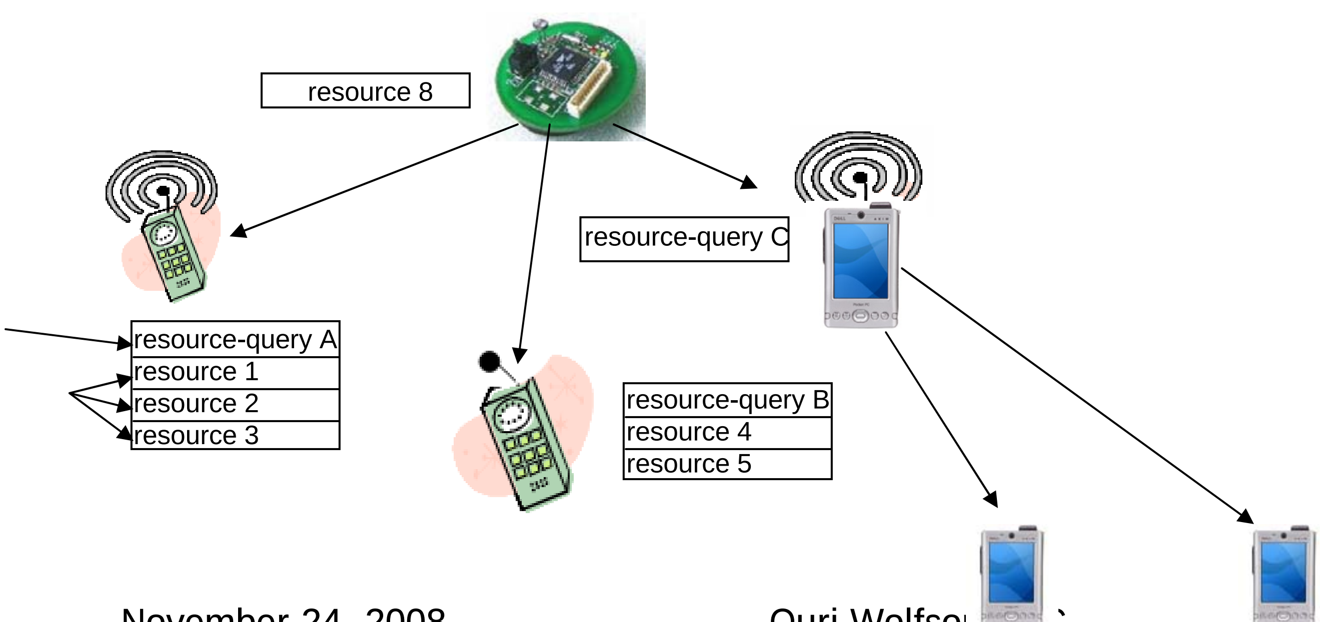

Push, stateful, Location-based: Greedy Forwarding

Location Awareness

- Assumptions
  - Each node knows
    - its location
    - the location of its neighbors
    - The location of destination node(s)
- Advantages
  - No route discovery necessary
  - No maintenance of routes necessary
  - Facilitates geo-casting: delivery of packets to all nodes in a region

Localization Techniques

- GPS (Global Positioning System)
  - Self-localization within a few meters
  - But does not work
    - indoors,
    - under foliage,
    - next to high-rise buildings
- Atomic Multilateration
  - Landmark: a node that knows its own location
  - Compute a node’s location from 3 or more landmarks using distance measurements
  - Distance by
    - Signal strength
    - ultrasound

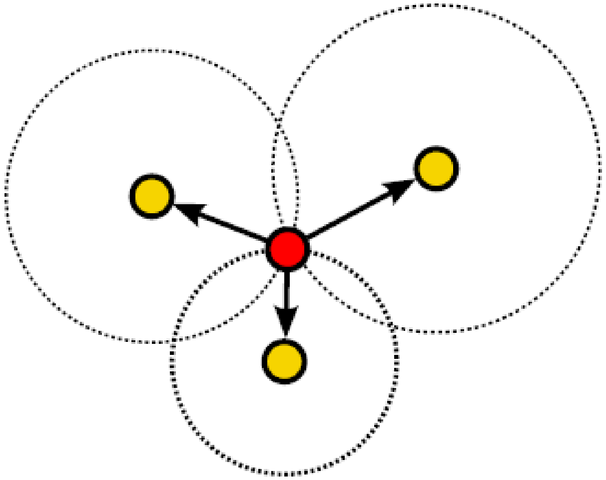

Greedy Forwarding

- Local strategy
  - A node forwards a message to a neighbor node that is "closer" to the destination than itself

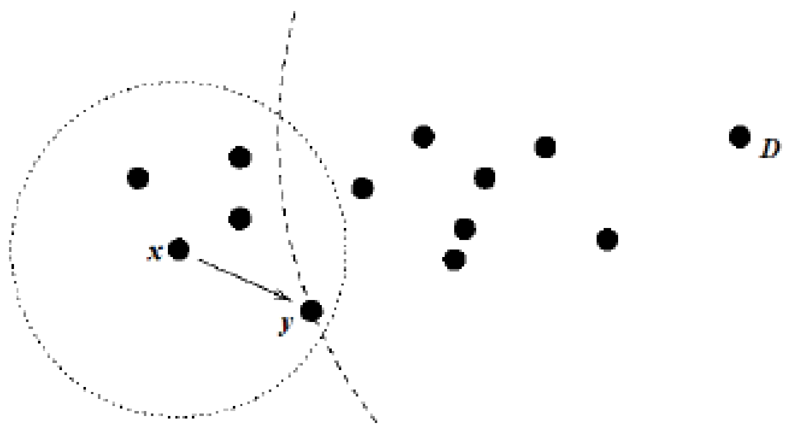

y is *x*'s closest neighbor to *D*

  - Repeat until the destination is reached
  - Instead of distance use *positive progress* or *direction*

Greedy Forwarding

- Loop-free
  - If nodes have consistent location information
- Require recovery strategy
  - Messages can get stuck in local minima
  - Greedy forwarding can fail (get stuck at *x*)
  - Recovery strategy: GPSR (Karp, Kung, 2000)

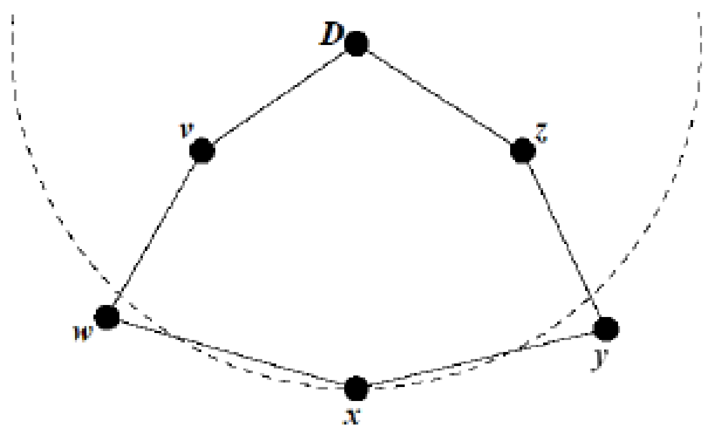

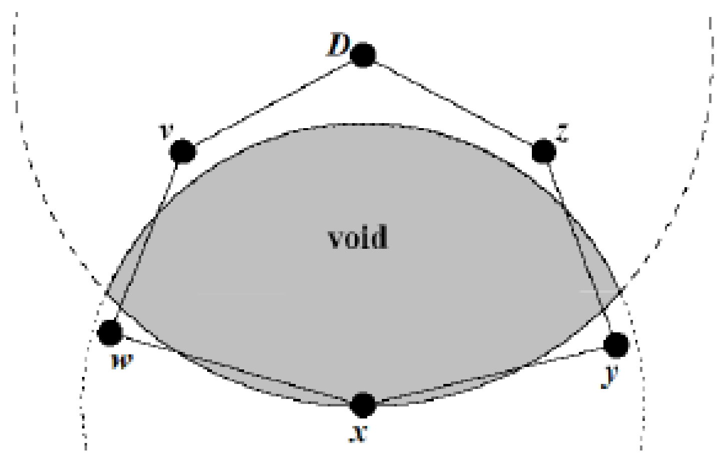

Report push, stateless: flooding, or data to query

- Each moving sensor broadcasts each report received exactly once. Problems:
  - S is disconnected; even S connects later -> miss R
  - Receipt by nodes that don't need
  - Duplicate receipts; broadcast storm
  - Resource blindness

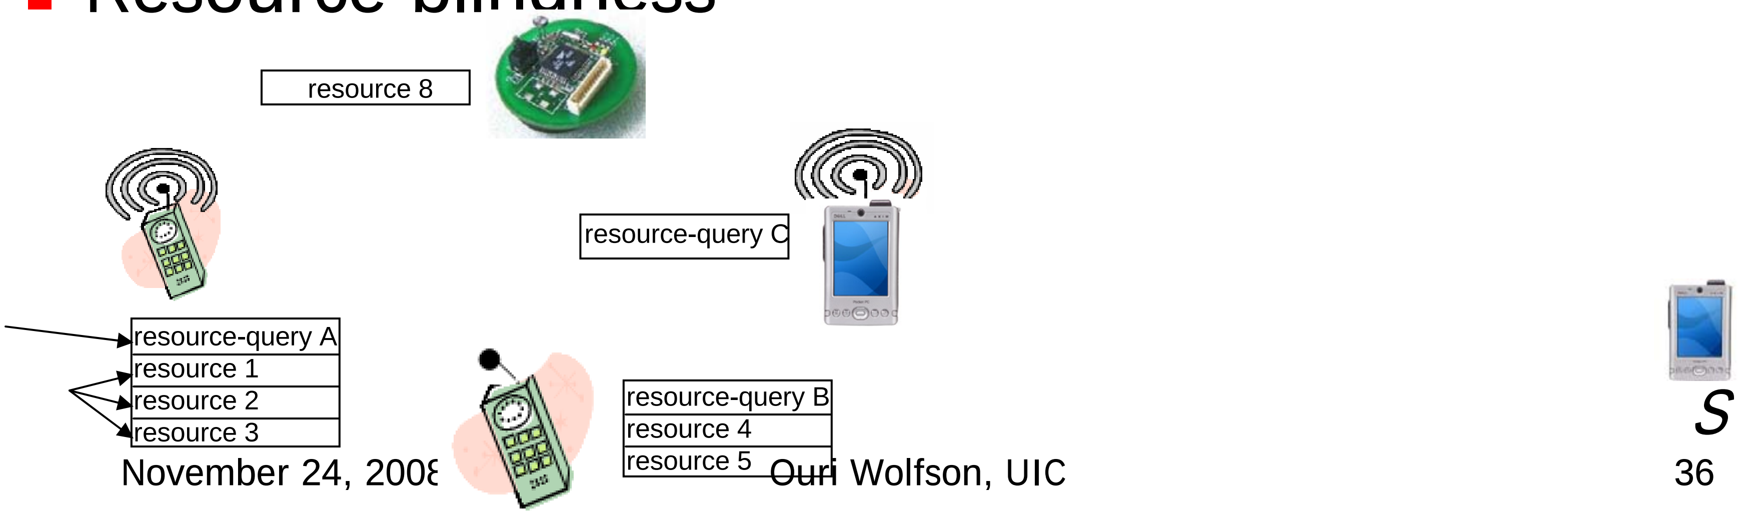

Disadvantage of Flooding

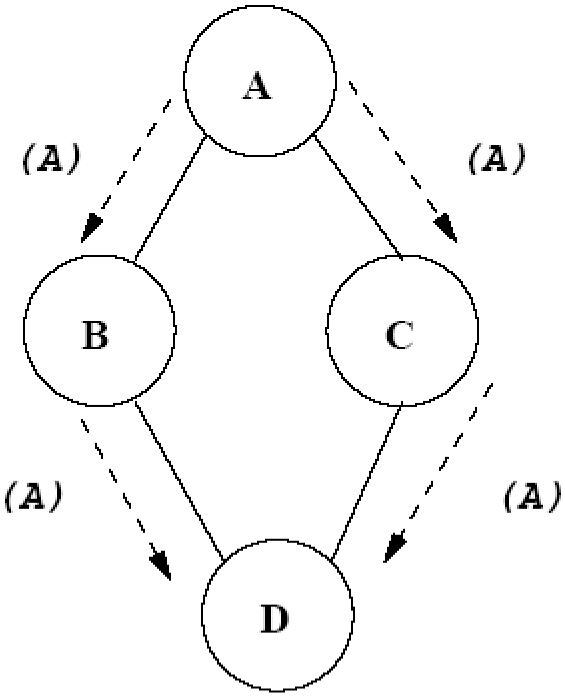

duplication

Remedy 1: Negotiation-based Approach

- SPIN (Sensor Protocols for Information via Negotiation)
  Kulik, Heinzelman, and Balakrishnan (2002)
  - Communicating raw data is expensive, but meta-data is not
  - meta-data descriptors
  - negotiating data transmissions using meta-data

- Solve duplication problems in classic flooding

- SPIN nodes can adapt their communication to
  - Application-specific knowledge of the data
  - Resources that are available to them, such as energy

- SPIN messages
  - ADV: a node A advertises data
  - REQ: an interested node B requests this data
  - DATA: the node A sends the actual data to B

SPIN

- Node A advertises data to node B (a)
- Node B responds by requesting data (b)
- Receiving the requested data (c)
- Node B then sends out advertisements to its neighbors (d), who in turn send requests back to B (e,f).

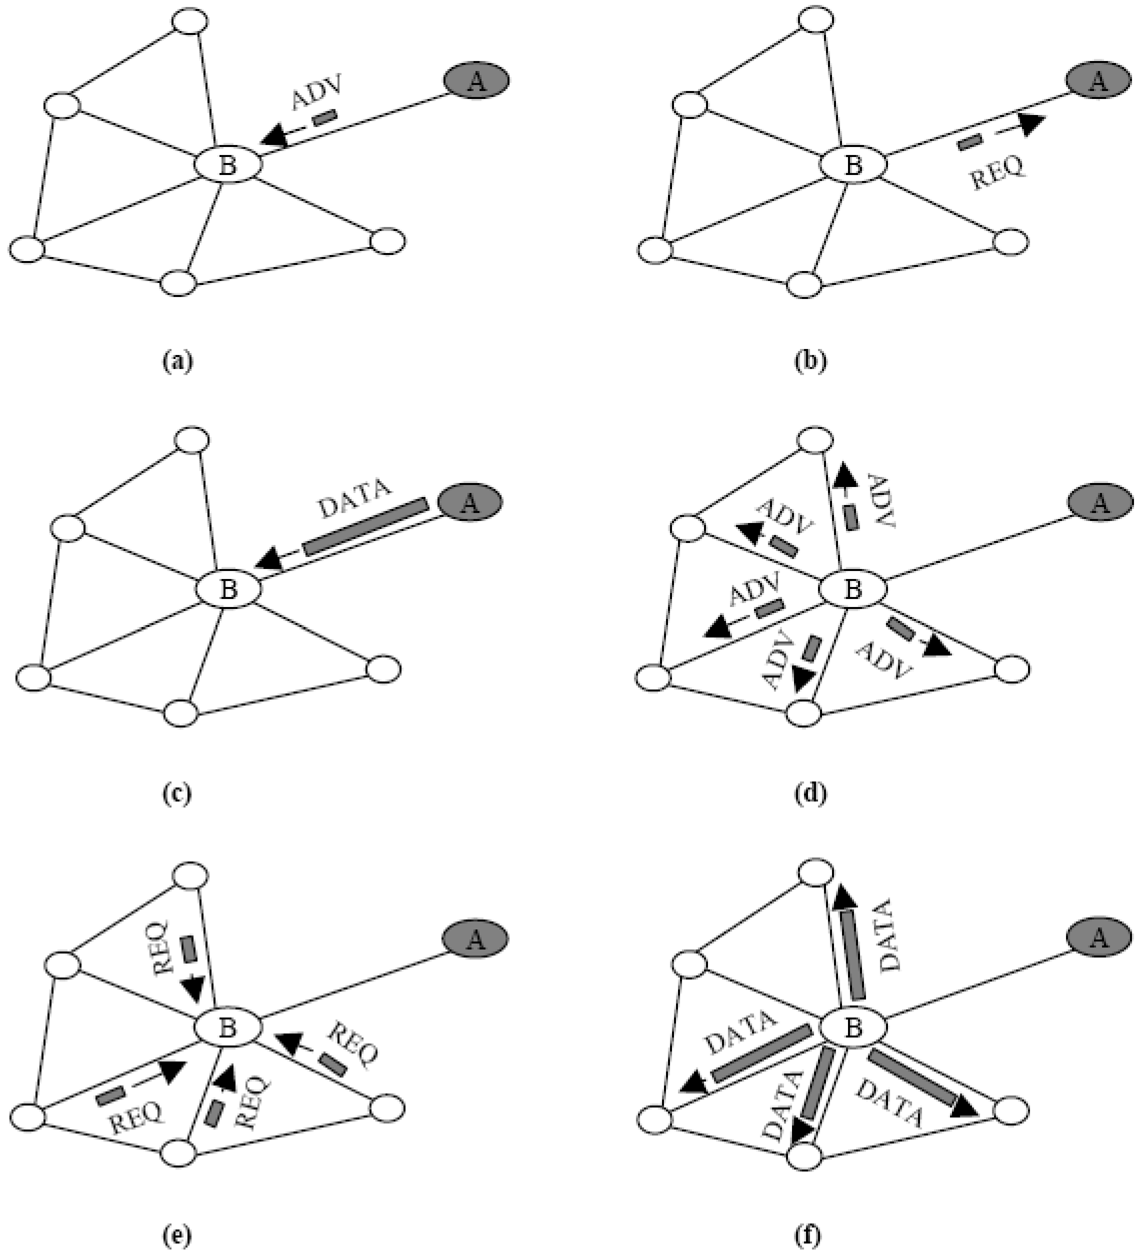

Problems:

- Handles one report at a time
- Mobility: need to keep advertising old reports to new neighbors
- Collision of requests

Remedy 2: store and forward

- Which reports in each transmission?
- How manage limited memory, power and bandwidth?
- How often to transmit?
- What is the size of each transmission?
- What is the range of each broadcast?

Remedy 2: store and forward

- Which reports in each transmit?
- How manage limited memory, power and bandwidth?

RANKING

- How often to transmit?
- What is the size of each transmission?
- What is the range of each broadcast?

Optimizing Power and Bandwidth Consumption

Query Processing outline

- Basics of MP2P search
- Data push and pull
- Power, memory, bandwidth management
- Ranking of information
- Experimental/commercial projects, Performance metric, and comparison
- Backchannel exploitation

802.11 Energy efficiency

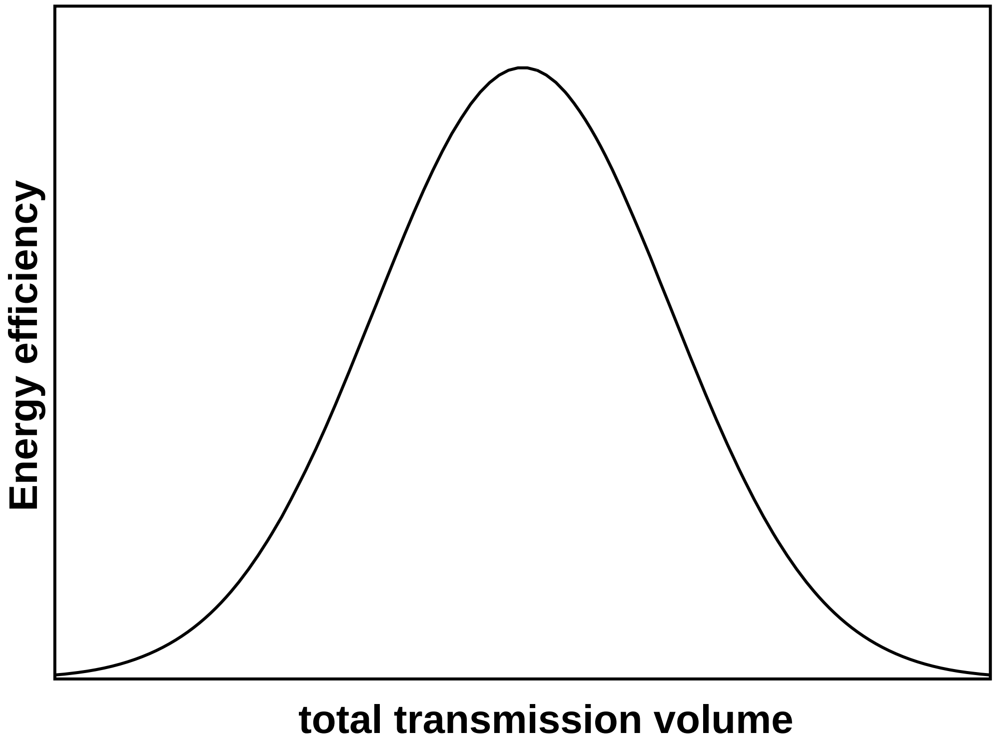

Energy efficiency = (amount of data correctly received) / (unit of transmission energy)

Tradeoff

- Frequency of transmission <=>
  transmission size

- How to quantify the tradeoff?

802.11 Basics

- 3 modes: transmitting, receiving, listening (order of power consumption)
- When listening: if detecting a message destined to host → receive-mode
- Time divided into slots, 20microsecs each
- Transmission:
  - Listen for 1 time slot
  - If channel free start broadcast (observe collision possible)
  - Broadcast may last for many time slots

Energy Efficiency of a Broadcast

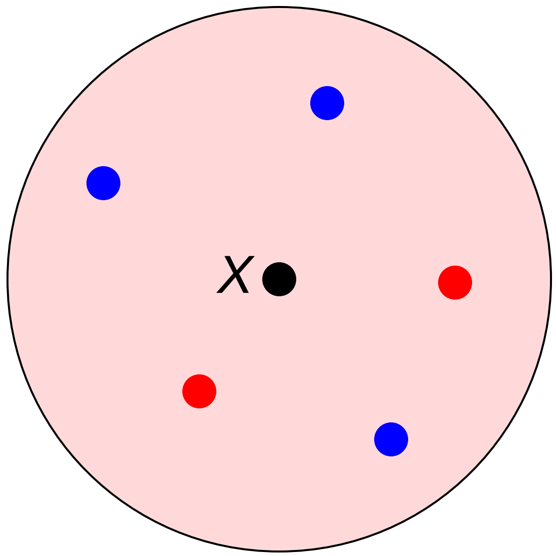

- successfully receive the broadcast from *x*
- Collisions occur at neighbor

Throughput (*Th*) =
(expected number of neighbors that successfully receive broadcast) × (broadcast size)

Power efficiency (*PE*) = $\frac{Th}{En}$

Computation of Throughput

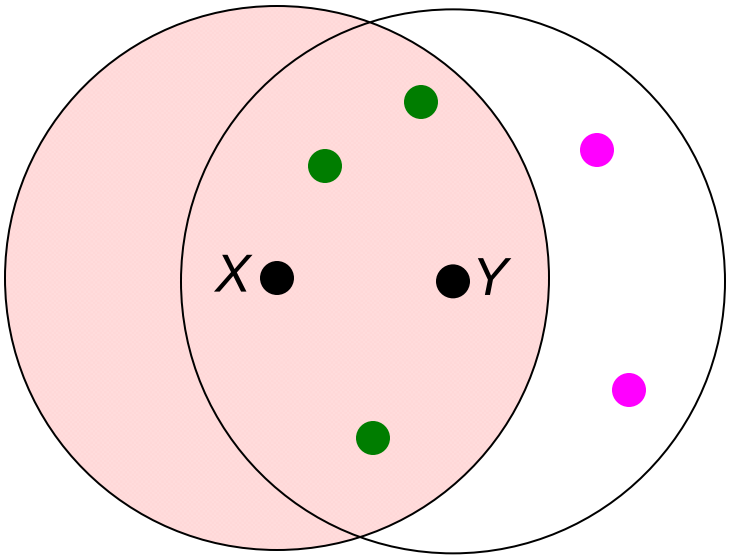

Conditions for successful reception at an arbitrary node *Y*

1. No green node inside starts to broadcast at the same time slot with *X*
2. No transmission from any purple node overlaps with that from *X*

Computation of Throughput (cont'd)

- $g(\delta)$: the probability that a neighbor at distance δ from X successfully receives the broadcast.

- $$g(\delta) = (1-p')^{\lambda S(\delta)(2T+1)} \cdot e^{\int_0^{r-\delta} 2\pi\lambda\varepsilon \ln(1-q(\varepsilon))d\varepsilon + \int_{r-\delta}^{r} 2\lambda\varepsilon \arccos(\frac{-r^2+\varepsilon^2+\delta^2}{2\varepsilon\delta}))\ln(1-q(\varepsilon))d\varepsilon}$$

where
$$S(a) = \pi r^2 - 2r^2 \cdot (\arccos(\frac{a}{2r}) - \frac{a}{2r}\sqrt{1-(\frac{a}{2r})^2}), 0 \le a \le r;$$
$$q(\varepsilon) = p \cdot (1-p')\lambda S(\varepsilon)\tau_{data}, 0 \le \varepsilon \le r; T = \frac{(M+h)\cdot 8}{b\tau}$$

λ Density of mobiles
r Transmission range in meters.
b Data transmission speed in bits per second
M Size of the broadcast (assume unique)
p The probability that a moving object attempts (by sensing the channel) to start a broadcast
p' The probability that it succeeds in starting a broadcast (unique) (broadcast frequency)
τ Length of the medium access time slot
h Size of the MAC header

Optimal Broadcast Size

$$E(Th) = M \cdot \int_0^r \lambda \cdot 2\pi\delta \cdot g(\delta)d\delta \qquad E(PE) = \frac{M \cdot \int_0^r \lambda \cdot 2\pi\delta \cdot g(\delta)d\delta}{M \cdot f + g}$$

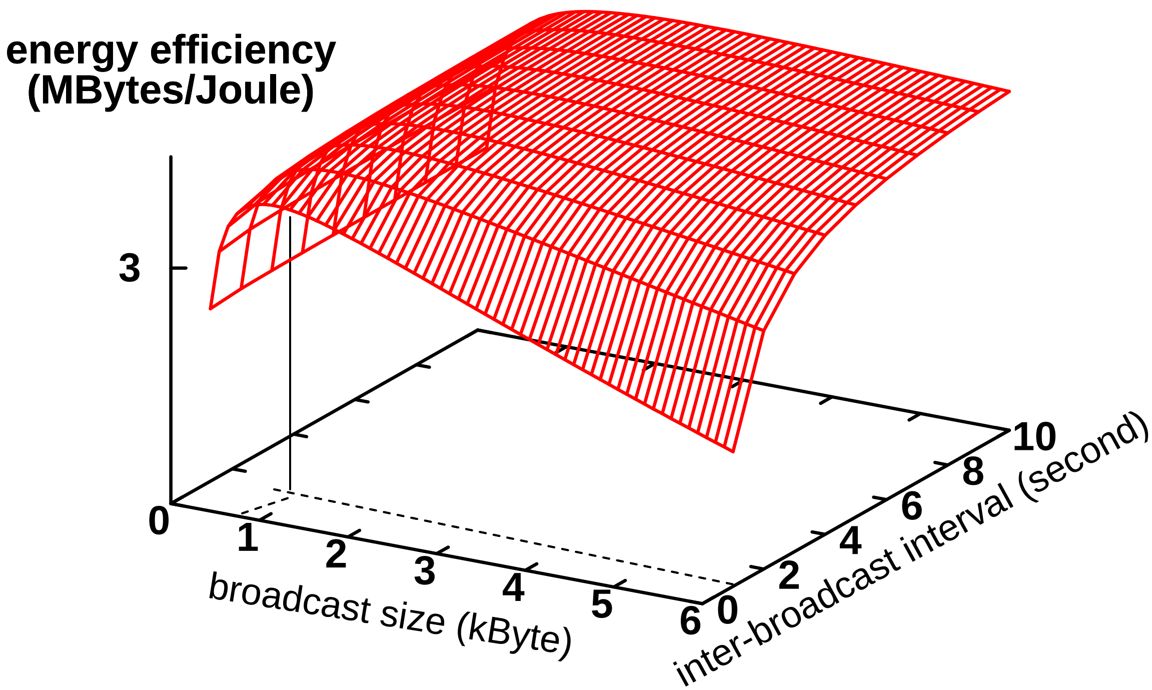

r=100 meters
λ=0.000625 per m^2
(inter-object distance=40m)

M = broadcast size

E.g., when inter-broadcast interval is 1 second, the optimal broadcast size is 800 Kbytes

Energy Constraints

- "From now until 8 o'clock the MP2P system is allowed to use 10% of the remaining energy" (The rest is used for voice communication, internet access, etc.) allowance/time-unit

- Energy consumed by a 802.11 network interface for transmitting a message of size *M* bytes

$$En = f \cdot M + g$$

For 802.11 broadcast, g=266$\times 10^{-6}$ Joule, f=5.27$\times 10^{-6}$Joule/byte

consumption

Number of reports transmitted

- Adaptively determined by
  - optimal transmission size
  - maximum transmission size
- Optimal transmission size determined
  - based on the E(*PE*) formula
  - given inter-broadcast interval (doesn't mean all broadcasts same size)
  - to optimize energy efficiency
- Maximum transmission size: determined based on the energy accumulated since last broadcast

- Actual transmission size =

  Min(optimal transmission size, maximum transmission size)

Query Processing outline

- Basics of MP2P search
- Data push and pull
- Power, memory, bandwidth management
- Ranking of information
- Experimental/commercial projects, Performance metric, and comparison
- Backchannel exploitation

Report Ranking: sample demand

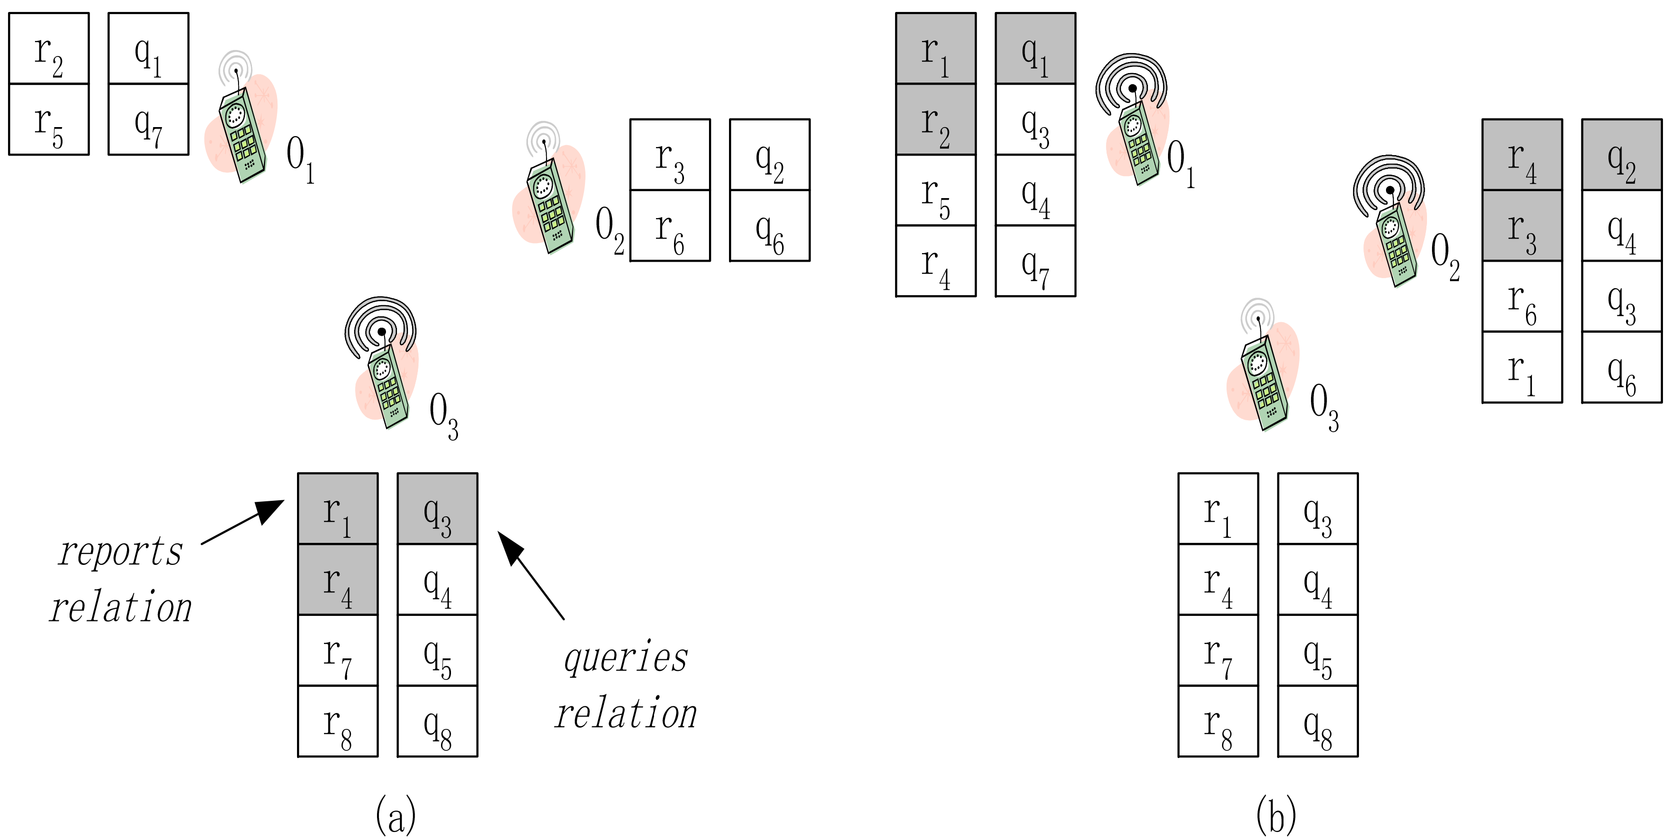

Queries relation is FIFO maintained

Rank of Reports

- Rank of a report R is determined by
  - Demand for R $\sum_{i=1}^{n} match(R, Q_i)$
    - Q_i's are the members of the queries relation
  - Supply (global parameter)

$$rank(R) = demand(R) / \text{supply}(R)$$

Query Processing outline

- Basics of MP2P search
- Data push and pull
- Power, memory, bandwidth management
- Ranking of information
- Experimental/commercial projects, Performance metric, and comparison
- Backchannel exploitation

Experimental MP2P projects (Pedestrians)

- 7DS -- Columbia University (web pages)
- iClouds – Darmstadt Univ. (incentives)
- MoGATU – UMBC (specialized query processing, e.g., collaborative joins)
- PeopleNet -- NUS, IIS-Bangalore (Mobile commerce, information type → location baazar)
- MoB – Wisconsin, Cambridge (incentives, information resources e.g. bandwidth)
- Mobi-Dik – Univ. of Illinois, Chicago (brokering, physical resources, bandwidth/memory/power management)

Vehicular projects

- **Inter-vehicle Communication and Intelligent Transportation:**
  - **CarTALK 2000** is a European project
  - **VICS** (The Vehicle Information and Control System) is a government-sponsored system in Japan with an 11-year track record
  - **FleetNet**, an inter-vehicle communications system, is being developed by a consortium of private companies and universities in Germany
  - **IVI (Intelligent Vehicle Initiative) and VII (Vehicle Infrastructure Integration)**, the US DOT
- MP2P provides data management capabilities on top of these communication systems
- Grassroots – Rutgers, p2p dissemination of traffic info to reduce travel times

Existing companies

- Messaging and Social Networks
  - Networking events: talk only to the right people
  - Social group: alert when buddy in near-by

- Mobiluck (Bluetooth, $14.95/year, Europe)
- Bedd (Bluetooth, free)
- Jambo (Wifi)
- Sixsense (Bluetooth, free)
- Inventop

RANk-based DIssemination (RANDI)

- Ranking of reports
- Bandwidth/energy aware
- Exchange enhances
  - Consumer functionality
  - Broker functionality
- Consumer: Answer local query (pull)
- Broker: Transmit reports most likely requested by future-encountered peers (push)
- Transmission trigger:
  - Encounter
  - New reports

7DS

P2P mode: each node periodically broadcasts its query and receives reports from neighboring peers. No strategy to determine query frequency and transmission size. Cache management based on web-page expiration time.

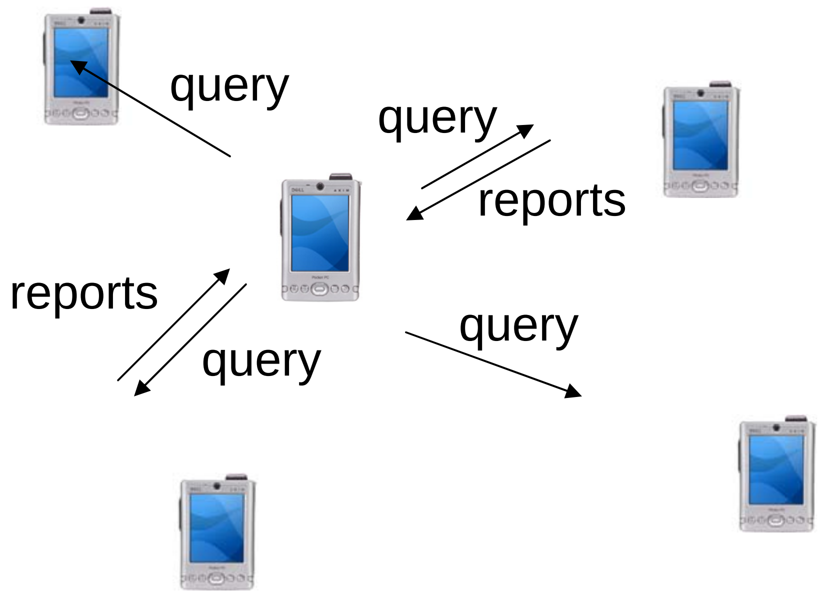

PeopleNet

Reports are randomly selected for exchanging and saving upon encountering.

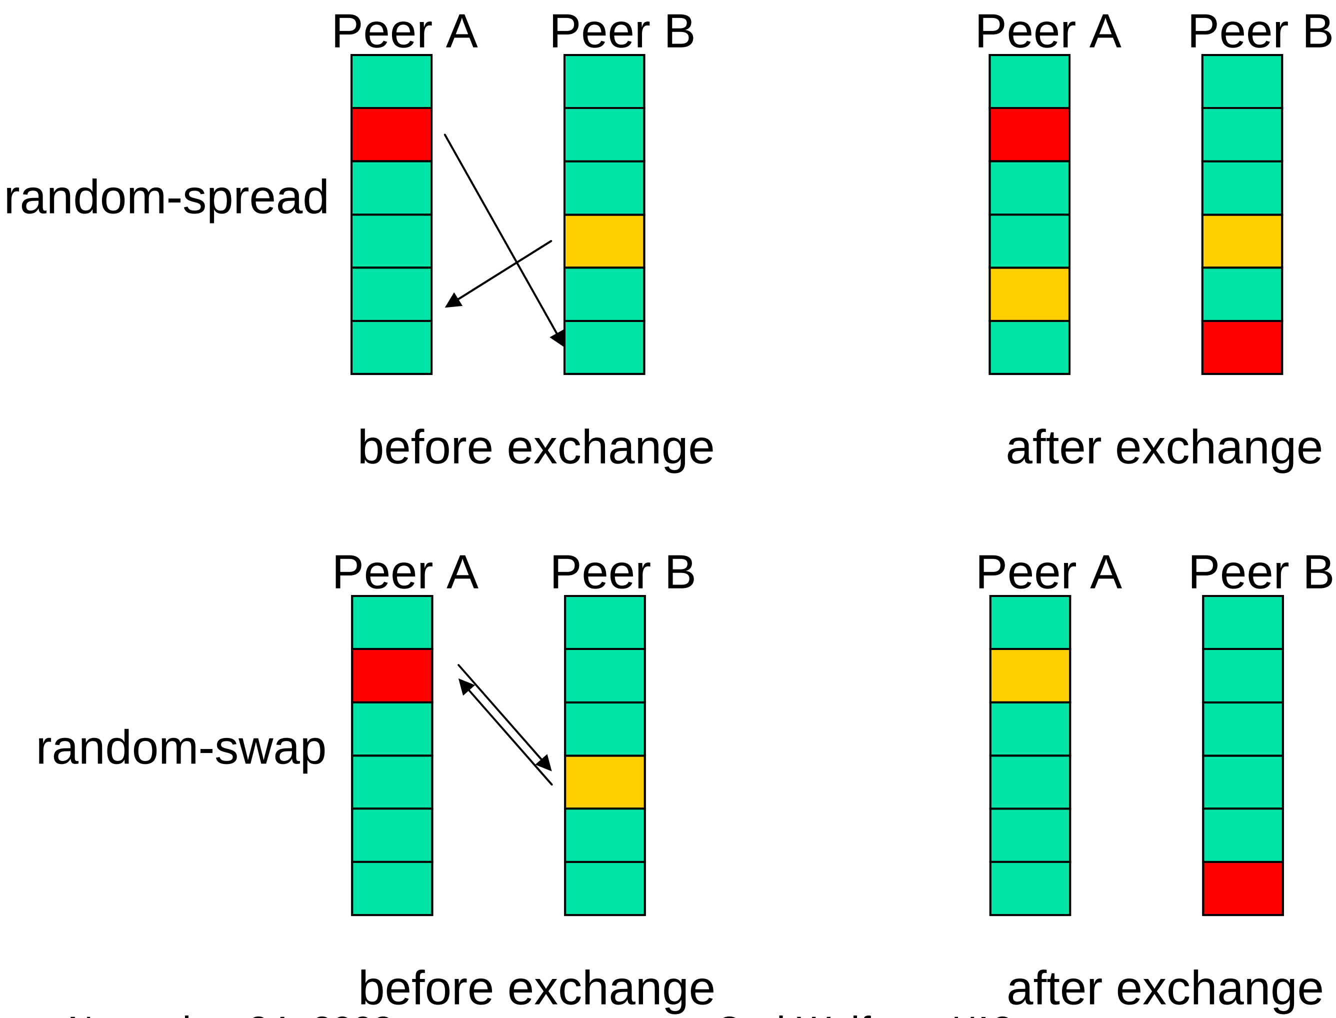

Comparison with 7DS and PeopleNet

- 7DS: Periodically querying neighboring nodes
- PeopleNet: Exchanging upon encounters
- Differences from RANDI:
  - There is no energy management for determining the transmission size;
  - The broker function is much more simplistic (no ranking);
  - 7DS does not have a good strategy to determine when to communicate.
- Also compared with PStree and flooding

transmission range=100 meters, energy allocation=0.01, battery life-time = 8 hours
mean of reports database size=100Kbytes
report size uniformly distributed between 100 and 2000 bytes
inter-device distance=40 meters, 0.1 report produced per second
mean of life span = 900 seconds, standard deviation of life span = 300 seconds

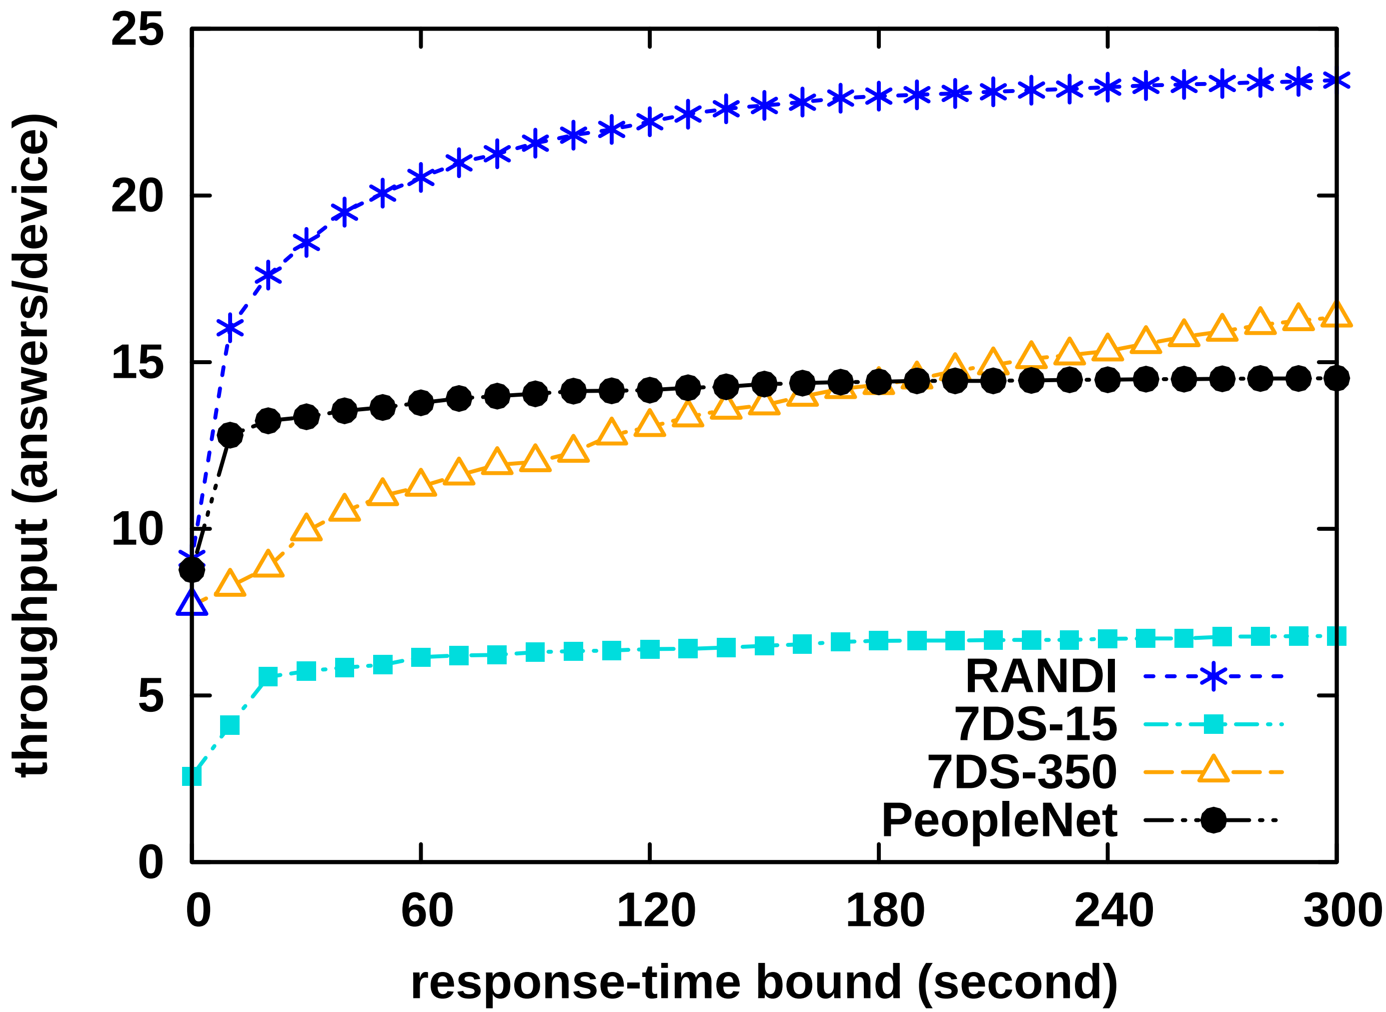

Optimum Broadcast Size is Important

energy allocation during 8 hours	percentage of transmissions that determine their size based on maximum transmission size
0.001	99.7%
0.01	23%
0.1	4.4%
1	4%

Query Processing outline

- Basics of MP2P search
- Data push and pull
- Power, memory, bandwidth management
- Ranking of information
- Experimental/commercial projects, Performance metric, and comparison
- Backchannel exploitation

Backchannel availability

- If so, why mp2p?
- Answers:
  - Backchannel does not imply server; may be
    - Costly
    - Inefficient due to frequent updates
    - Require some hierarchical architecture
  - Solution that does not necessitate server

Benefit of Backchannel

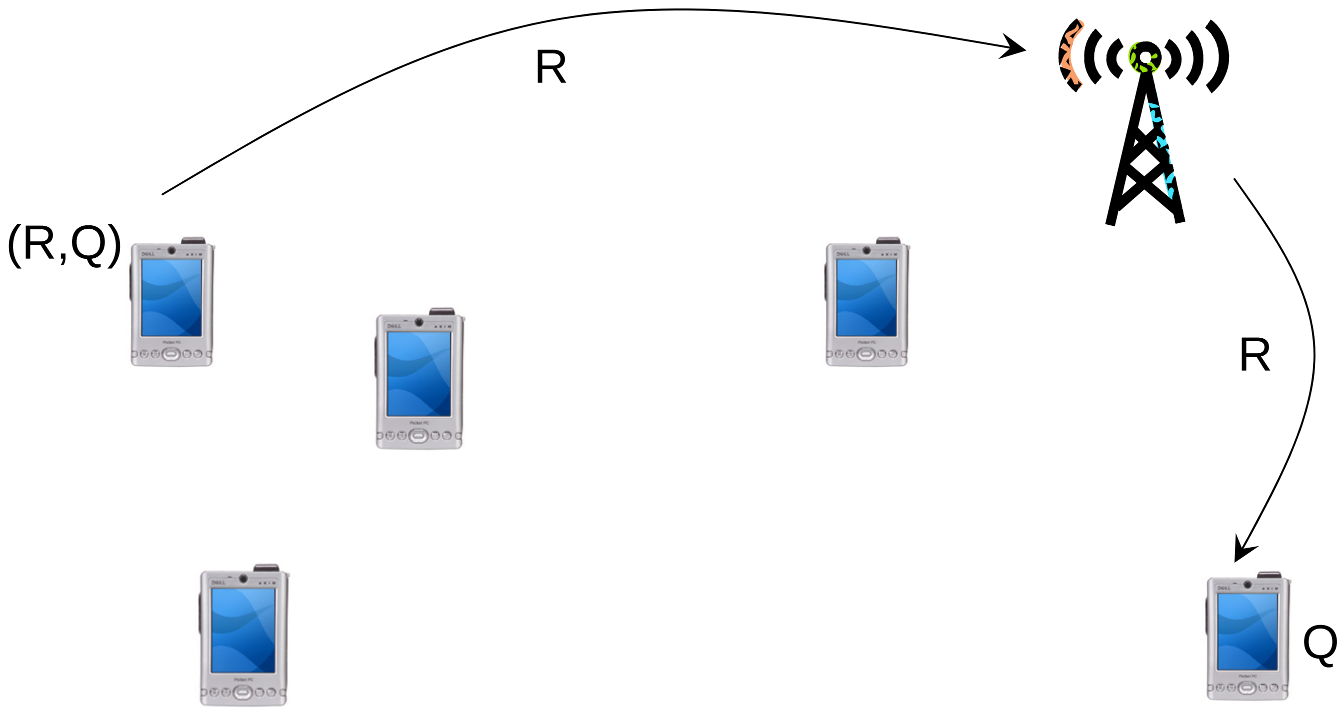

- When a match occurs at a moving object, the object sends the answer report to the query originator via the fixed infrastructure
- Backchannel communication:
  - Subject to energy constraints
  - Consumes more energy

transmission range=100 meters, energy allocation=0.01
reports database size=1Kbytes, backchannel message header=20 bytes
report size uniformly distributed between 10 and 100 bytes
inter-device distance=40 meters, 1 report produced per second
mean of life span = 3600 seconds, standard deviation of life span = 0 seconds

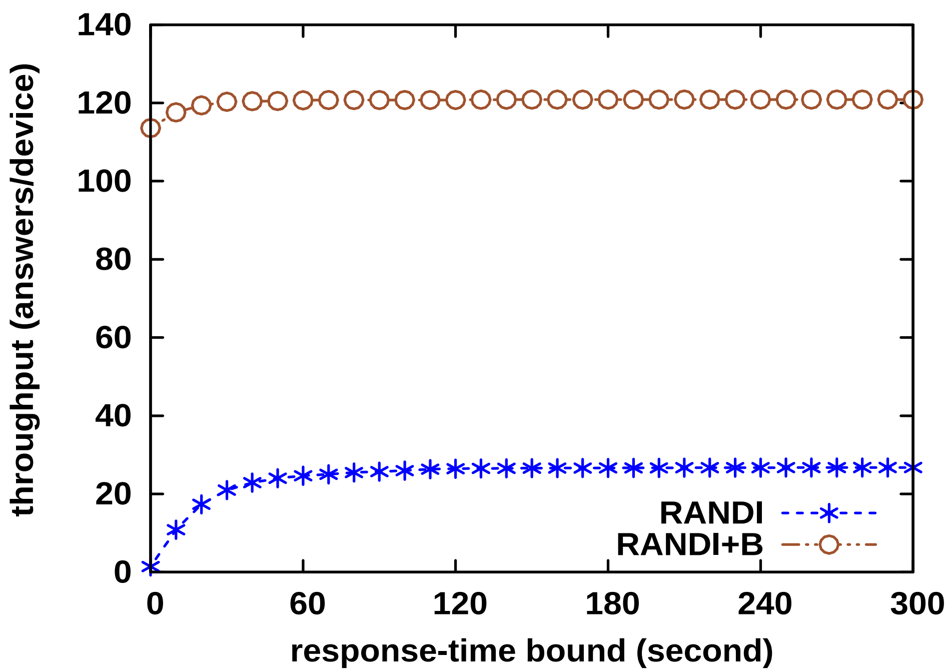

Research issues in data management

- Query Processing
- Dissemination analysis
- Value of information and comparison with C/S
- Participation incentives
- Remote Querying
- Data Integration of sensor and higher level information (maps, trip plans, ride-sharing profiles)
- Other relevant work

Dissemination analysis issues

- Dissemination pattern: how a report spreads
- Dissemination coverage: % of moving objects receiving the report

Dissemination pattern/density

Smart flooding: Locality of diffusion is ensured by demand/supply ranking and limited memory and bandwidth

How fast/far a resource is disseminated?

In a Mobile Opportunistic p2p system, the answer depends on:

- Memory allocation to the resource type
- Transmission range
- Traffic speed
- Vehicle density
- Resource density
- Average resource-validity time

Dissemination coverage

- Gas stations, restaurants, ATM's, etc., report continuously
- An report is acquired by the peers within the wireless transmission range, and disseminated transitively
- Alternative location-based-services paradigm to
  - Cellular-service provider database (privacy)

Advertising Request

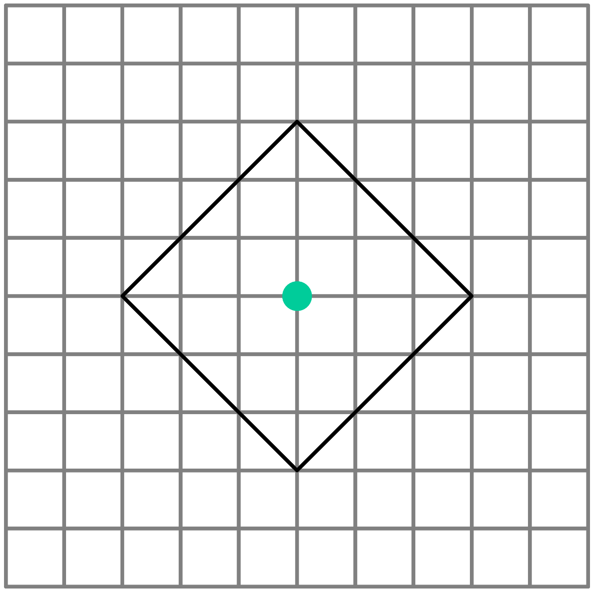

- Advertise to peers within the **target area**
- The route-distance from the *producer* is smaller than 3 blocks
- Announcement strategy: to each passer-by

Coverage: % of peers reached within target area

Transmission range = 50 meters, motion speed =40 miles/hour, diameter of target area = 1 miles

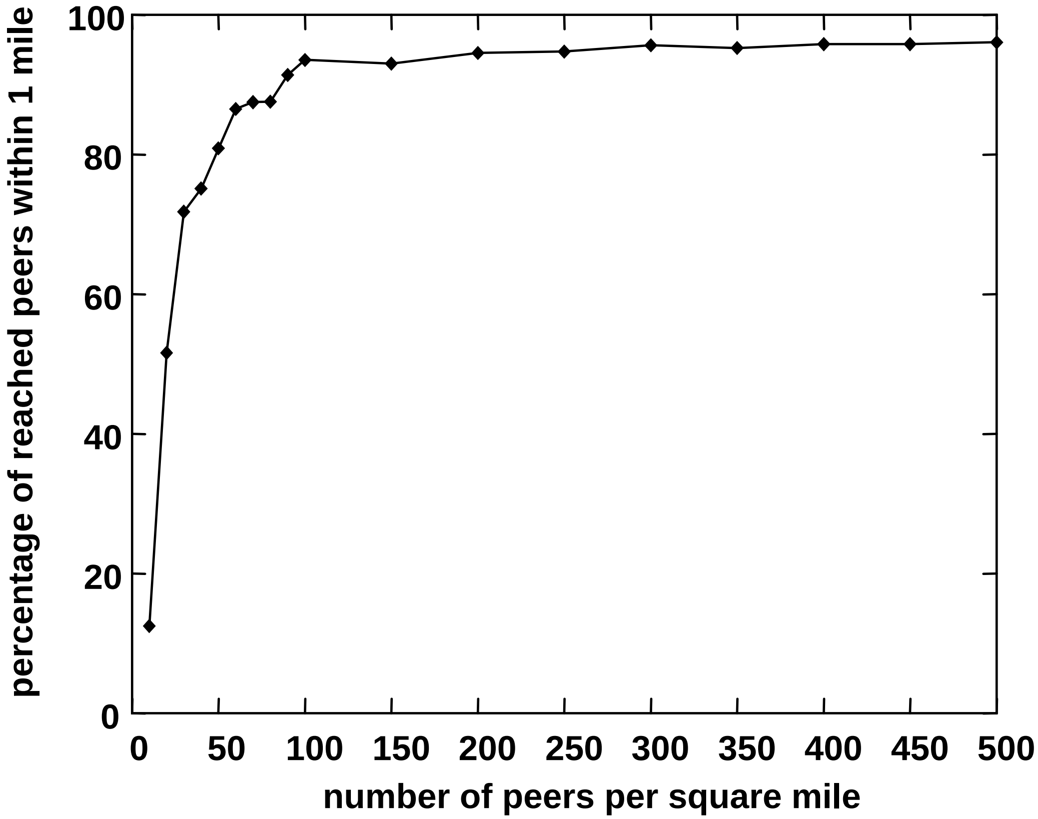

Further research in dissemination analysis – mathematical model

- Spread resembles epidemiological models of (Bailey 75) but there are important differences
  - Spatio-temporal relevance function
  - Interaction of multiple infectious-diseases (resources)
  - Random graphs relationship
- For a given
  - mobility model:
    - topology (grid, free-space, synthetic graph, map),
    - interference (independent),
    - constraints (random walk, random waypoint)
  - Resource/report generation model
  - Communication network model

Should answer questions:

- what is the probability that peer at (x,y,t) receives report generated at (0,0,0)?

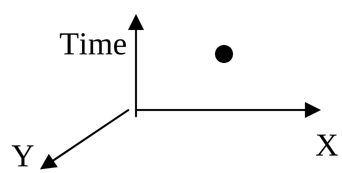

- What is the coverage for a given announcement strategy?

Research issues in data management

- Query Processing
- Dissemination analysis
- Value of information and comparison with C/S
- Participation incentives
- Remote Querying
- Data Integration of sensor and higher level information
- Other relevant work

A resource classification

- Competitive – acquired by one consumer at a time (parking slots, cab-customers)
- Semi-competitive (ride-sharing partners)
- Noncompetitive (malfunctioning brake lights, speed of a vehicle at (x,y,t))

Strategies for Capturing (semi-) Competitive Resources

- Blind Search: move in some pattern, and when encounter available resource – capture it

- Information Guided Search– as BS, except
  - Break pattern to capture resource if report received
  - Not successful ➔ resume pattern-motion
  - If while moving to r1 a higher-relevance report is received for r2, move to r2.

Blind Search

Consumer follows a pattern (e.g. clockwise along a rectangle) to capture a resource

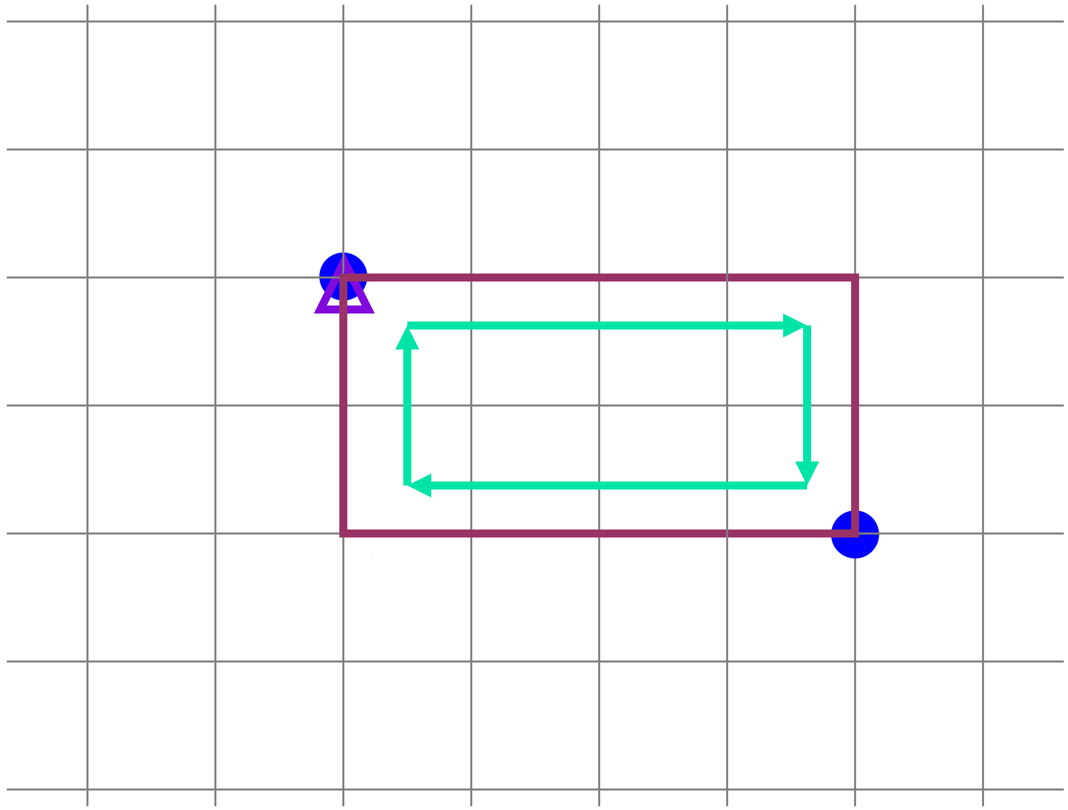

Information Guided Search

Consumer changes direction to a reported resource

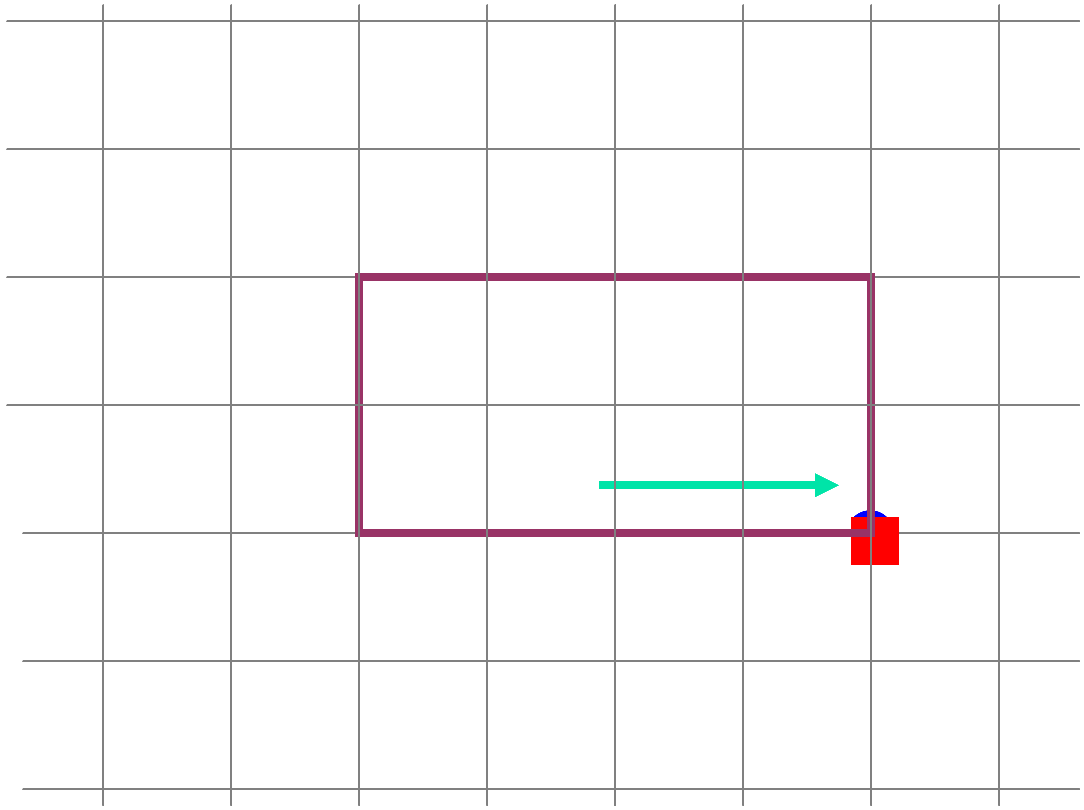

Information Guided Search

Consumer resumes the pattern if the target resource is not captured

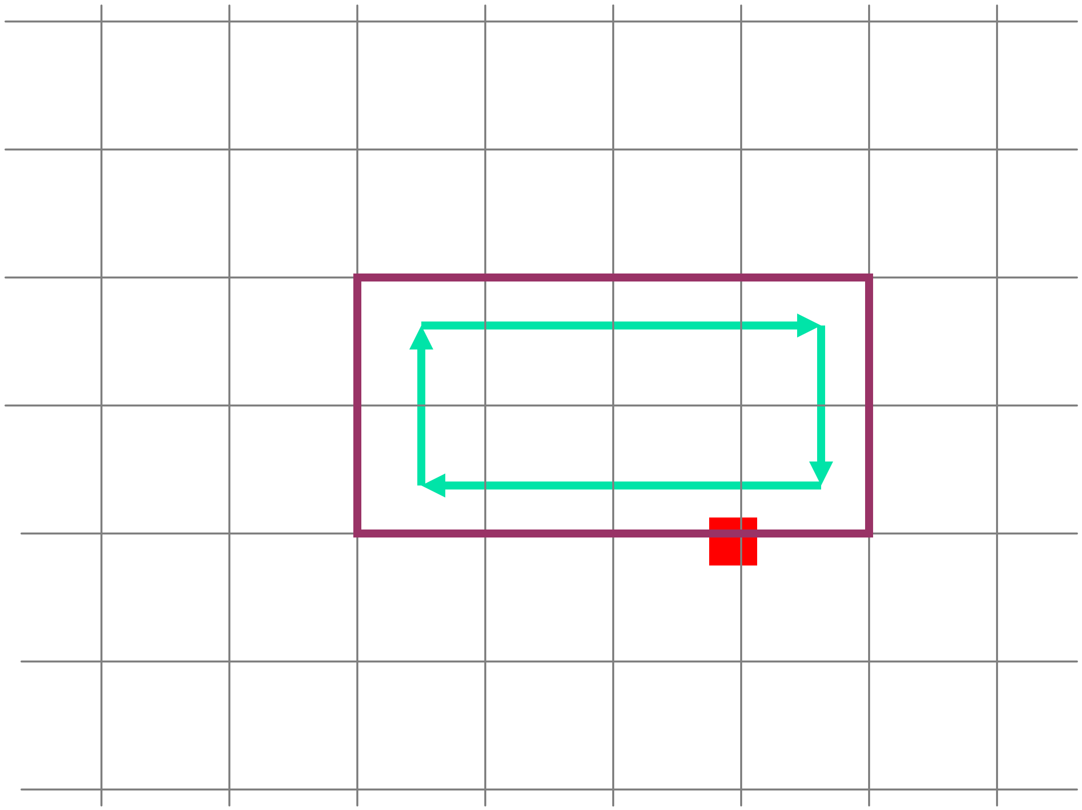

Information Guided Search

Consumer chooses the most relevant resource to go

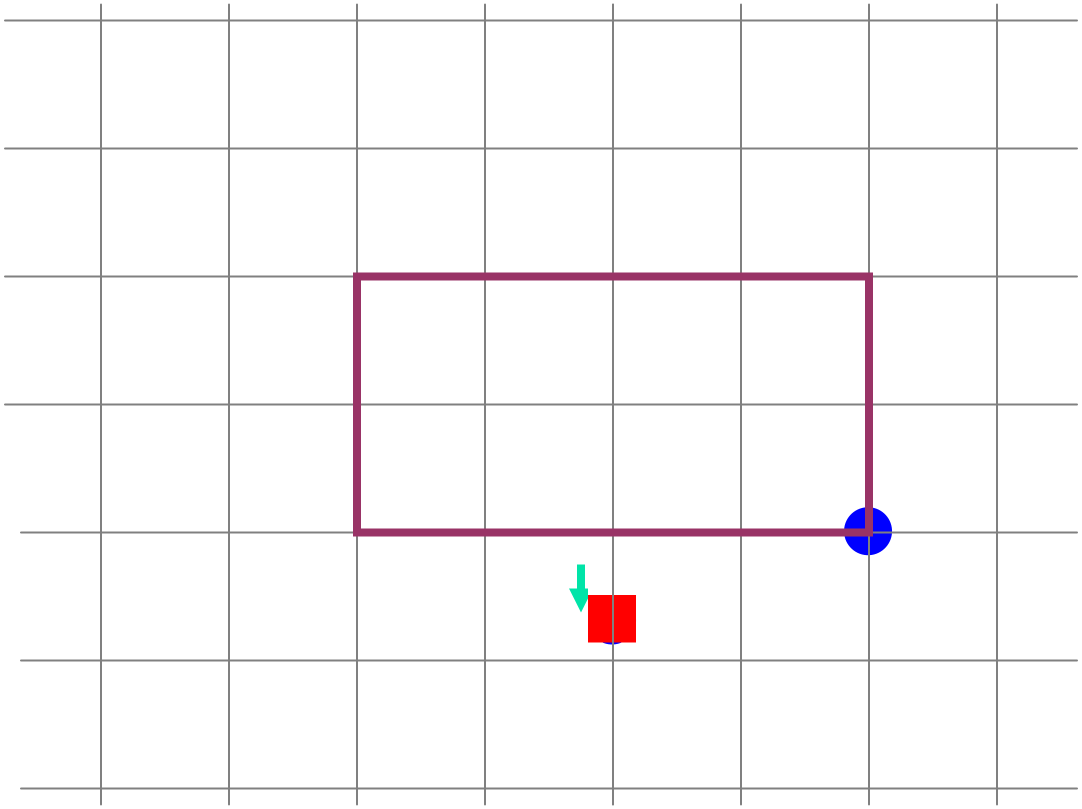

Value of information: IGS versus BS in Client/Server Mode

motion speed = 10 miles/hour, 100 consumers/mile², mean of invalid duration = 20 minutes, delay = 0

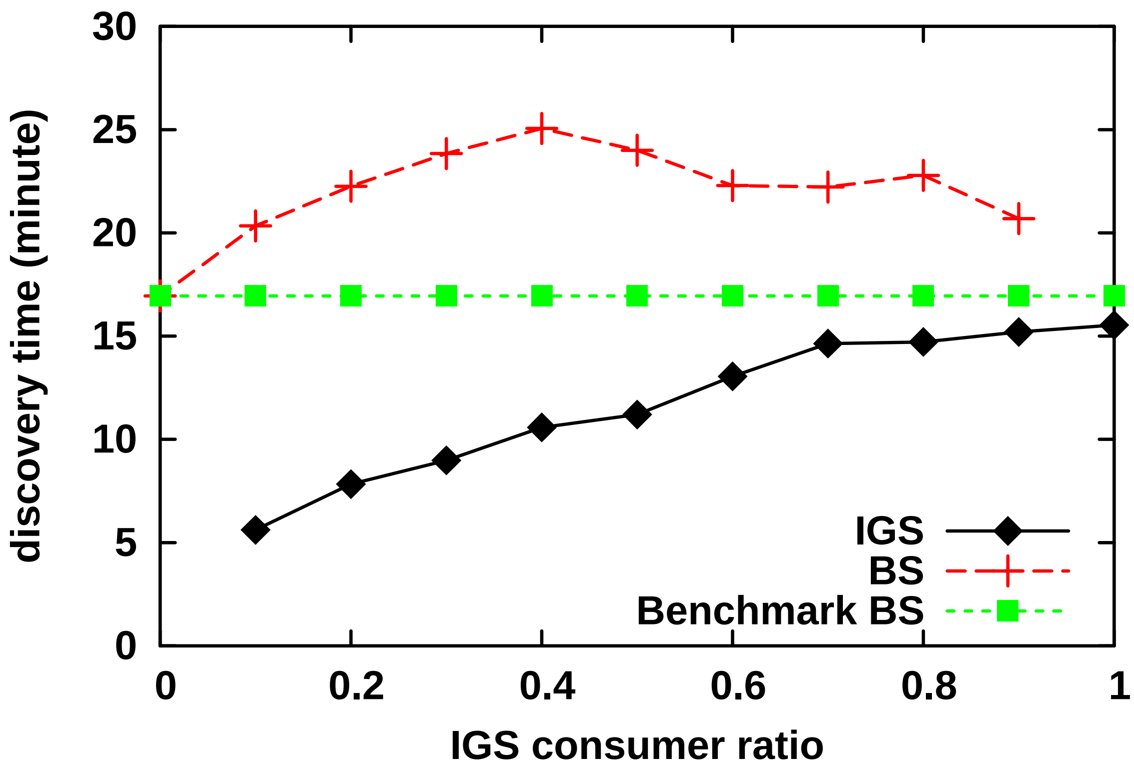

- IGS outperforms benchmark BS
- Benchmark BS outperforms partial-BS

IGS versus BS in Client/Server Mode

motion speed = 10 miles/hour, 100 consumers/mile^2,
10% IGS consumers, delay = 0

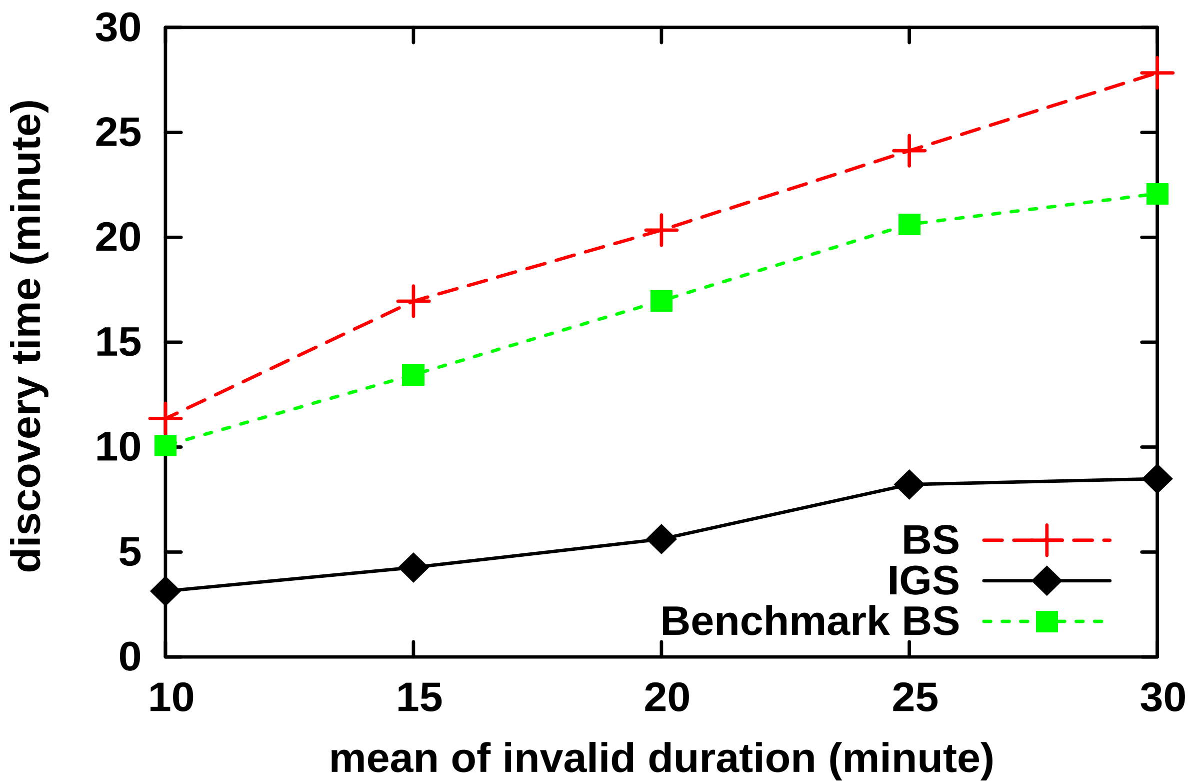

- IGS outperforms benchmark BS
- Benchmark BS outperforms partial-BS
- IGS especially suitable for long-occupancy environment

Client/server versus OP2P

transmission range = 150 meters, motion speed = 10 miles/hour, mean of invalid duration = 20 minutes, 100 consumers/mile^2, 10% IGS consumers, delay = 20 seconds

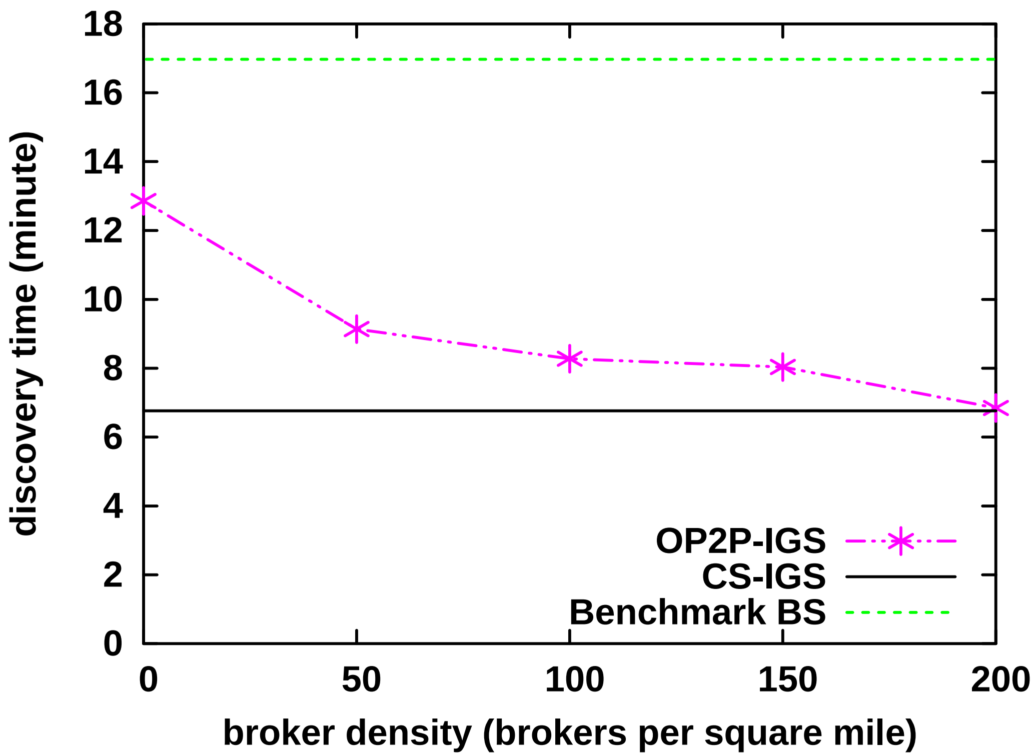

- IGS outperforms BS
- IGS approaches C/S as broker density increases

Research issues in data management

- Query Processing
- Dissemination analysis
- Value of information, comparison with Client/Server
- Participation incentives
- Remote Querying
- Data Integration of sensor and higher level information
- Related approaches

Incentive Mechanisms

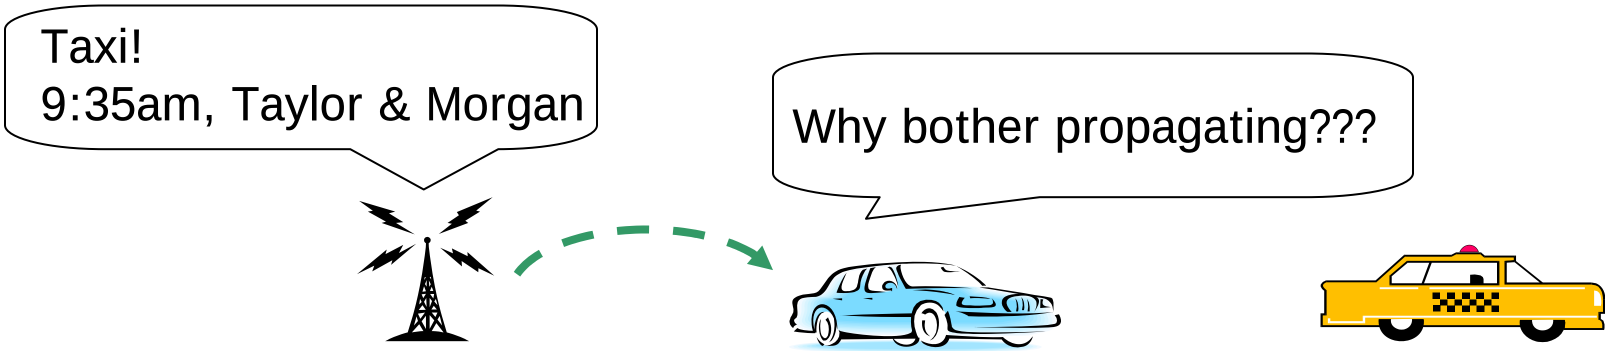

- Peers owned by different authorities, each with its own goal
- Owner may decide not to participate in information dissemination, i.e. turn off OP2P module
- Incentive mechanisms are needed

Who pays?

- Consumer: economic benefits of storing reports
- Producer

Possible solution: Virtual Currency and Security Module

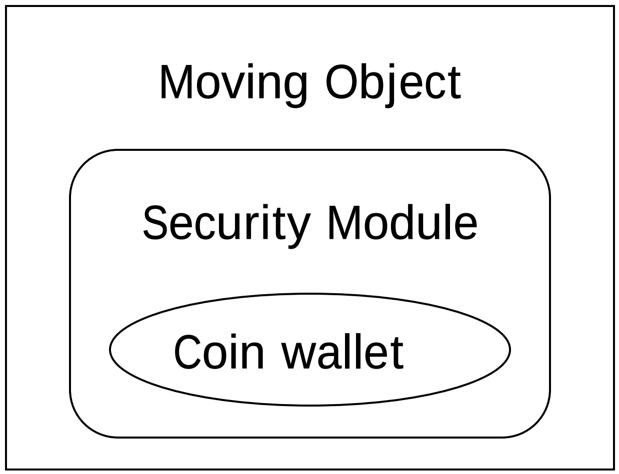

e.g.
1 coin = 1 cent

- System circulates a virtual currency, *coins for* payment
- Coin counter stored in a trusted and tamper resistant *security module* (hardware + software)
- Security module exposes a fixed set of operations.
- Example of security module implementation: microprocessor smart card
- Report dissemination is paid by producer

Advertising Report

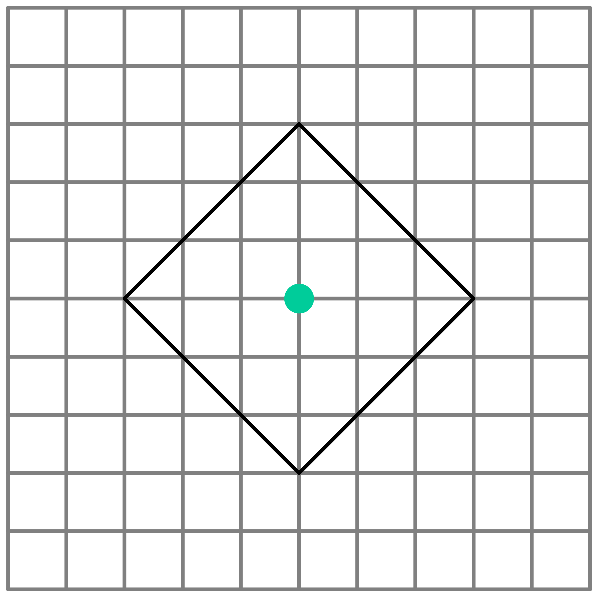

- Disseminate a data item to moving objects within the **target area**
- The route-distance from the *producer* is smaller than 3 blocks.

Offline Producer-paid (OPP) strategy

1. Announce to each passer-by
2. Producer attaches its ID (credit card) to announcement
3. Each broker records number of transmissions for producer
4. Record is submitted offline to clearance center for redemption

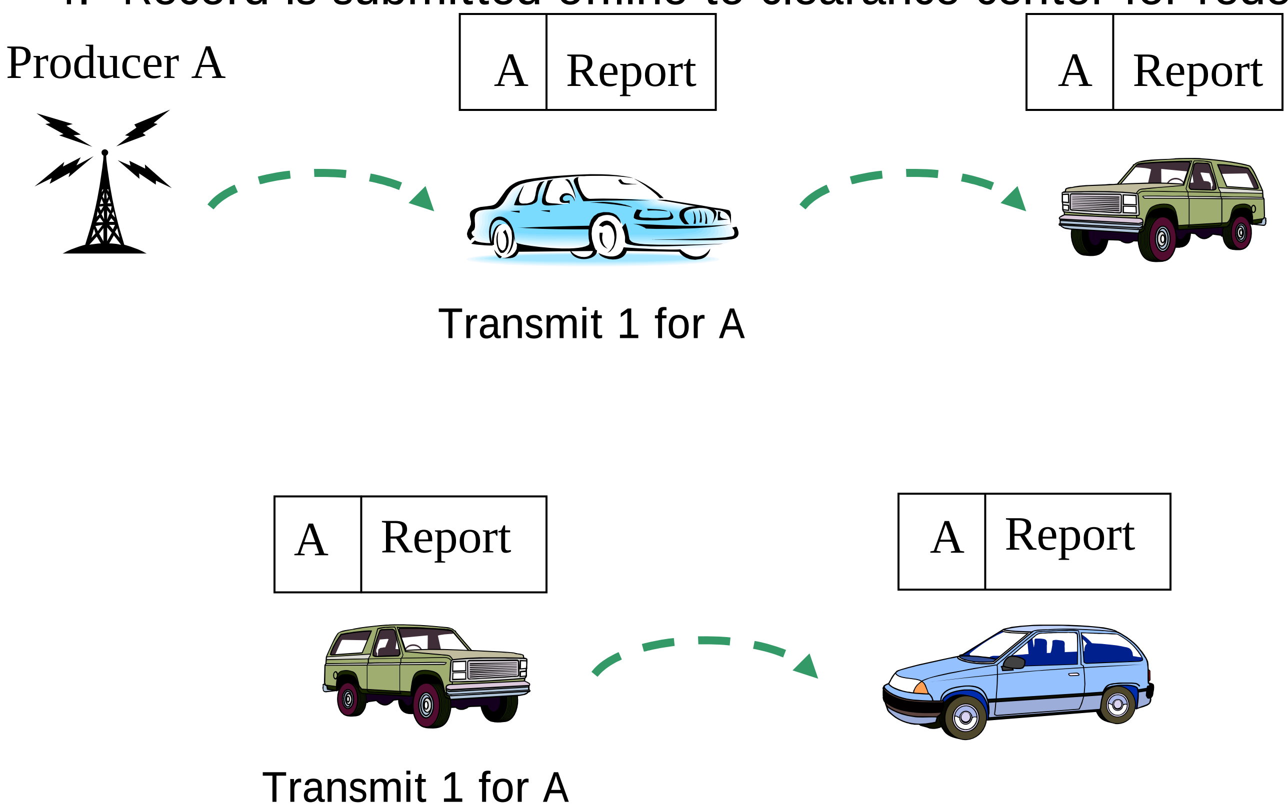

Opp drawback – clearance center

Producer-paid – Logarithmic Dissemination Strategy (LDS)

1. Announce to each passer-by
2. Producer includes an initial budget in announcement
3. Each broker withdraws a flat commission fee f
4. The rest of the budget is split between sender and receiver

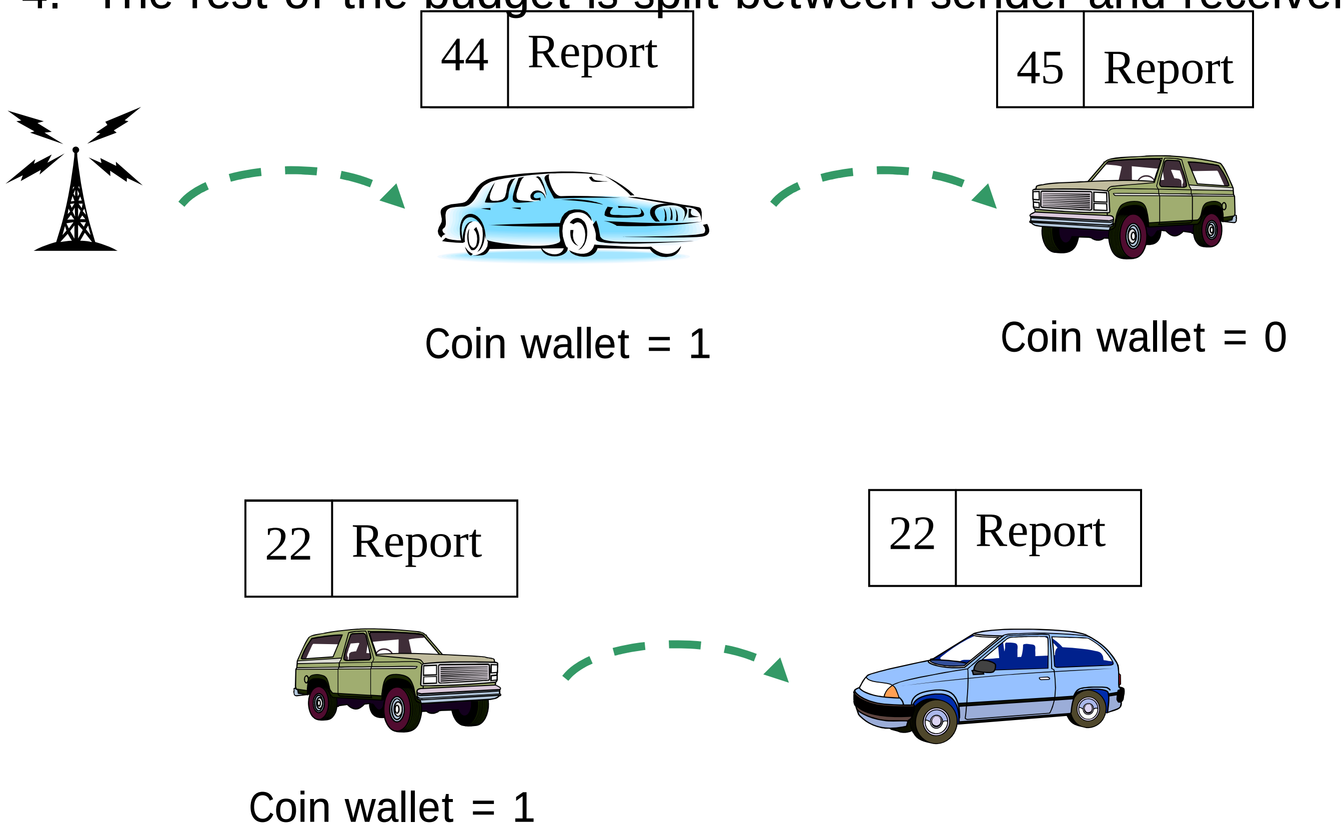

Initial credit of LDS

Tradeoff: coverage vs. initial credit

Transmission range = 50 meters, motion speed =40 miles/hour, diameter of coverage area = 2 miles

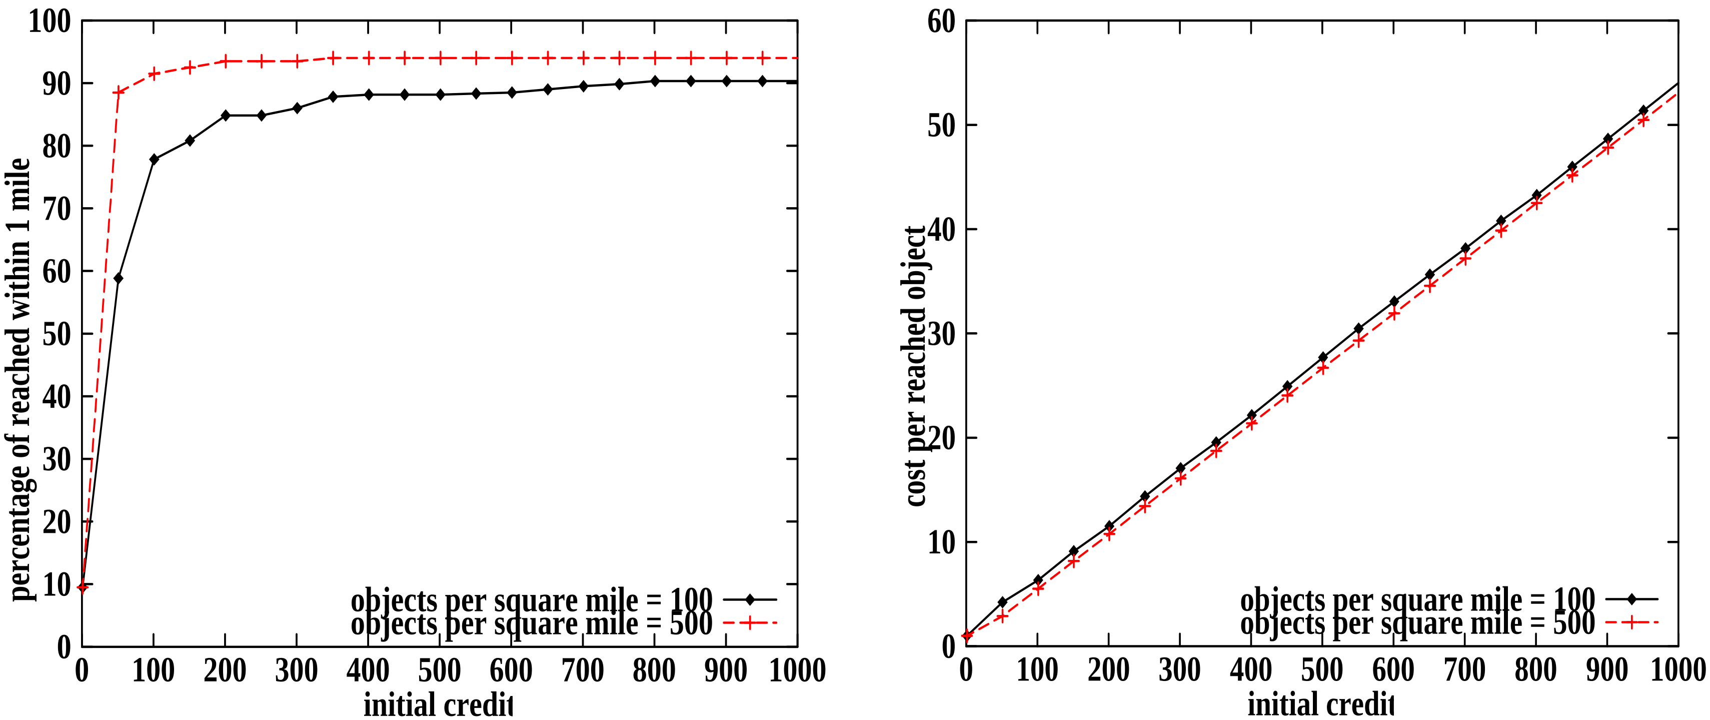

Inflection point (50) indicates optimal initial credit

At inflection point 2 coins wasted for each peer reached.

LDS versus OPP

- OPP maximizes coverage with minimum cost per reached object (exactly f)
- LDS enables complete decentralization

Atomicity Issue

- Commit protocol may not complete, leaving one participant not knowing the final status at the other participant.

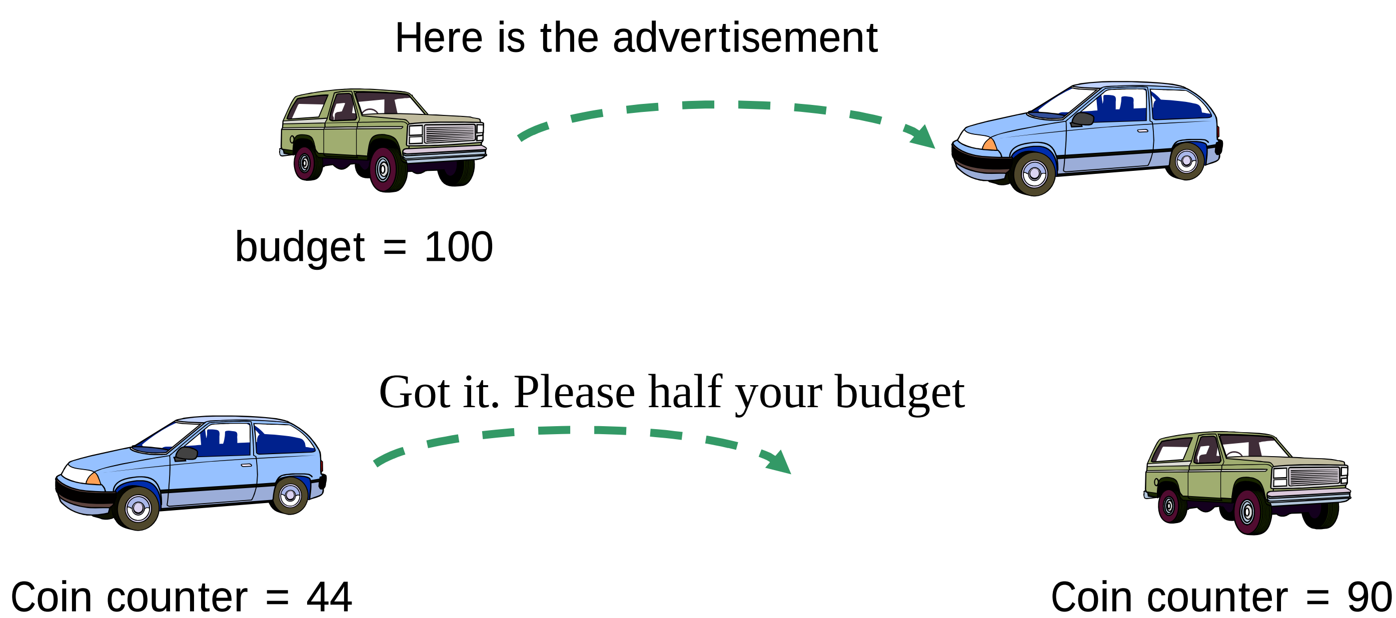

Possible solution

- Clearance center
- Peers remember unsuccessful transactions.

Research in incentive models

- Other incentive models (reputation based??)
- Pricing, negotiation, auctions (report value may depend on number of recipients)
- Cost optimizations in such models
  - Example: minimize advertisement cost per potential customer
- Transactions and atomicity issues
- Security
  - fake resources
  - tampering to gain unfair advantage, create havoc
- Incorporate P2P concerns (reward both producers and brokers), and MANET concerns (energy management)

Research issues in data management

- Query Processing
- Dissemination analysis
- Value of information, comparison with Client/Server
- Participation incentives
- Remote Querying
- Data Integration of sensor and higher level information
- Other relevant work

Spatio-temporal resource query modes

- Local: peer queries local database
- **Remote: peer queries a region R, i.e. all the peers in R**

Query language

- Spatio-temporal SQL, or OWL
- + operators:
  - Remote region of dissemination
    - May be implicit, e.g route of bus #5
  - Budget

Remote query-- research issues

- Query dissemination to R
- Query processing in R
- Answer delivery to query originator, o

Query dissemination to R

Relationship of o and R-- 2 cases:

- Originator o is in (the center of R), "find the taxi cabs within 1km of my location"
  - Query flooded within R
  - Query dropped by a peer outside R
- Originator outside R, "what is the average speed one mile ahead on the highway"
  - Flooding may be too resource-expensive

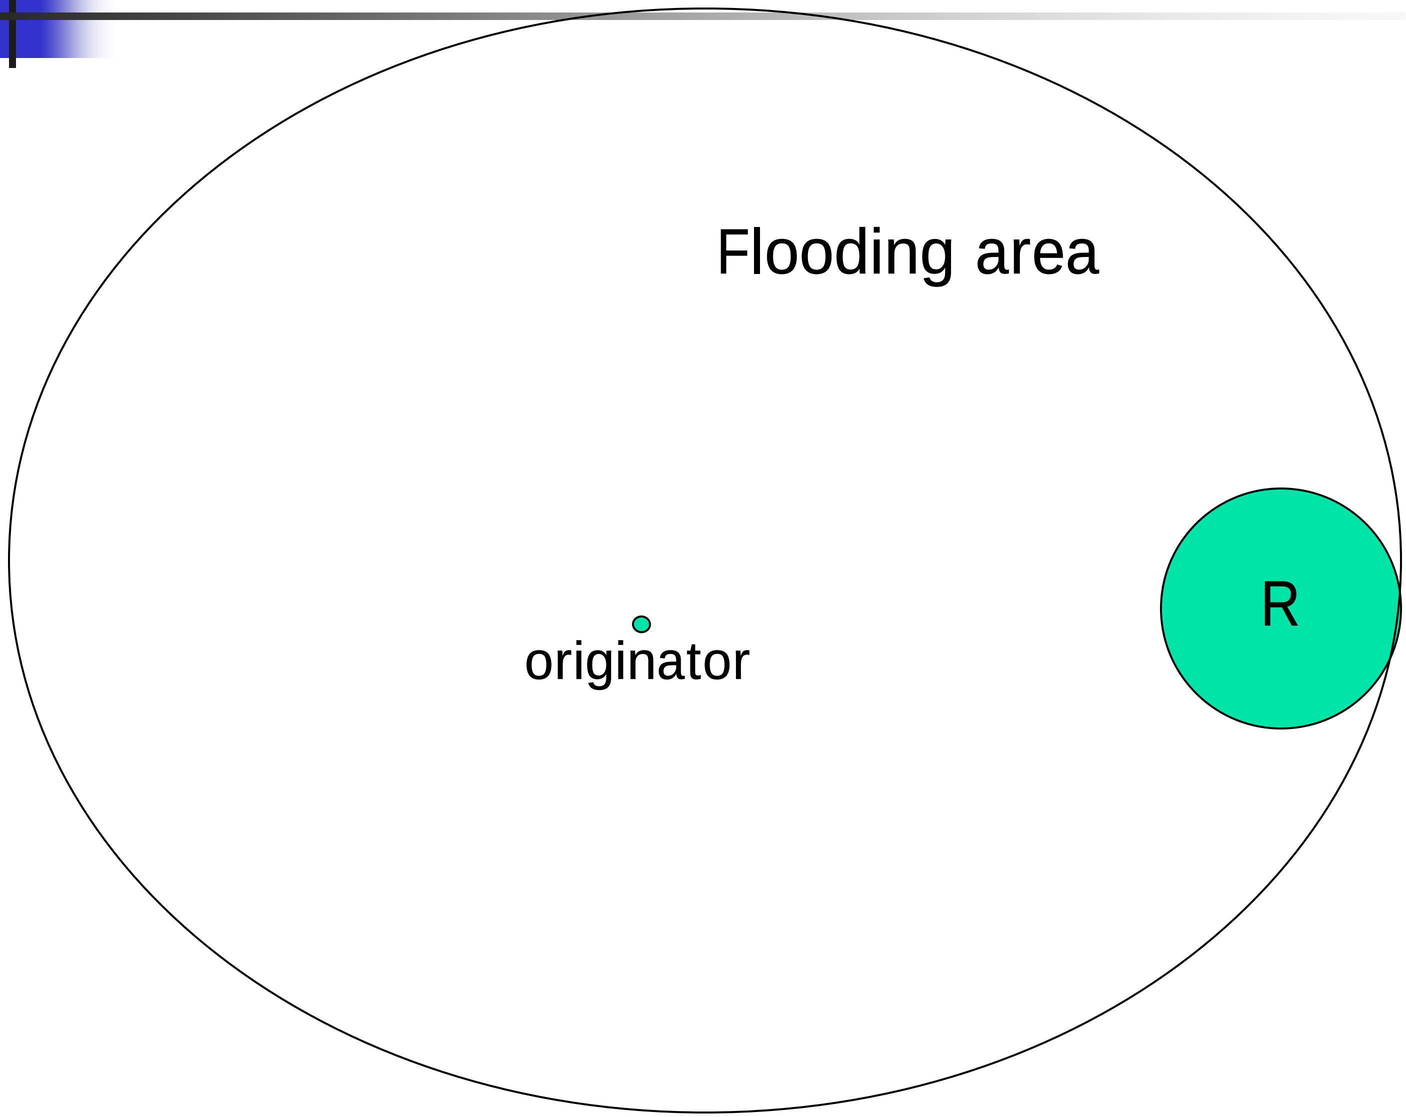
Flooding area
R
originator

Two Cases

- Each peer knows the trajectories of each other peer
  - Trajectories exchanged as resources

- Each peer does not know the trajectories of other peers except that of the originator

Known Trajectories

- Encounter graph: each edge represents the time interval during which two peers can communicate

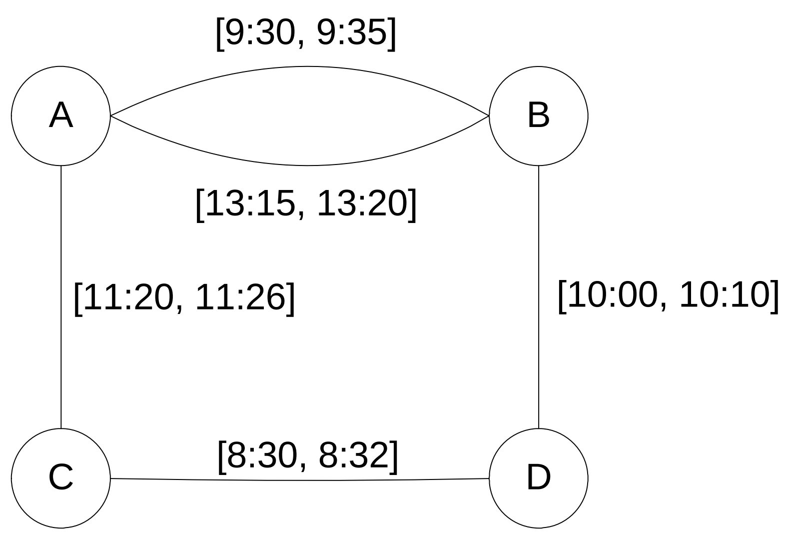

Known Trajectories

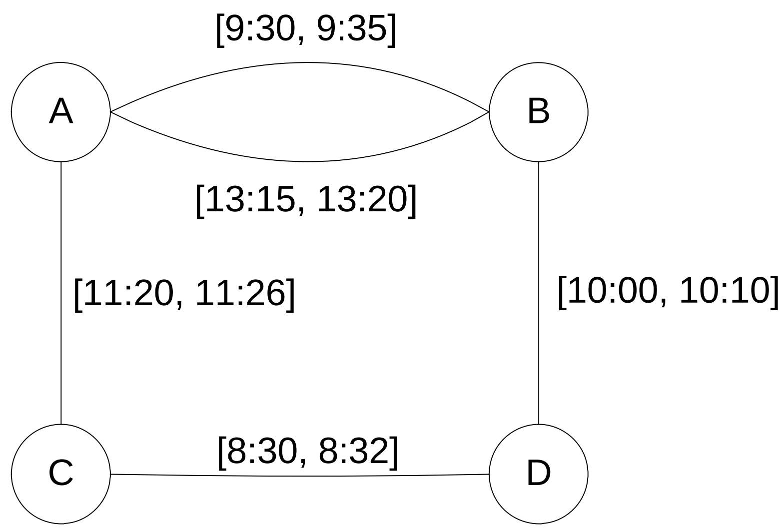

- A revised Djikstra algorithm is used to find
  - the shortest path between the originator and peers in the query destination area (for query dissemination)

Unknown trajectories: Transmitting to a peer p depends on

- Moving direction of p relative to R
- Budget of query
- Density of peers

Several related works in Mobile ad hoc networks, eg Location-based-multicast (LBM), DREAM, but

  - Low mobility
  - Connected communication graph

Remote query-- research issues

- Query dissemination to R
- Query processing in R
- Answer delivery to query originator, o

Query Processing Modes (1)

- Response to originator by each queried vehicle

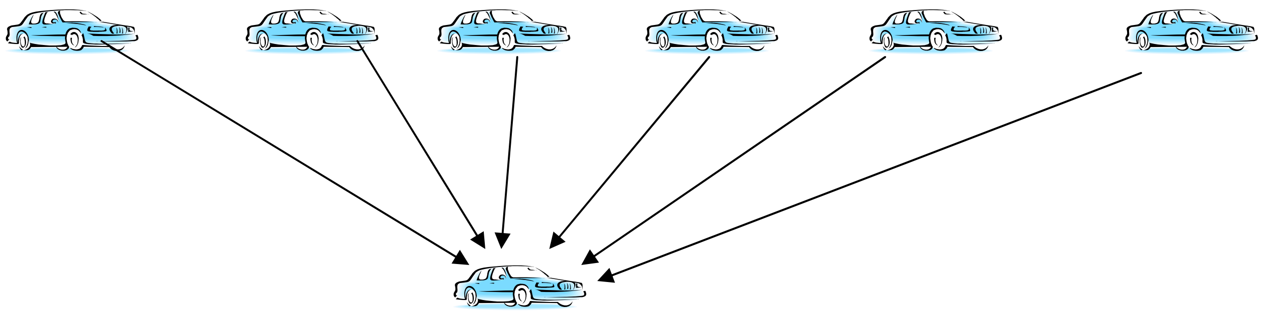

Query Processing Modes (2)

- Response to *leader* by each queried vehicle; *leader* consolidates and responds to originator

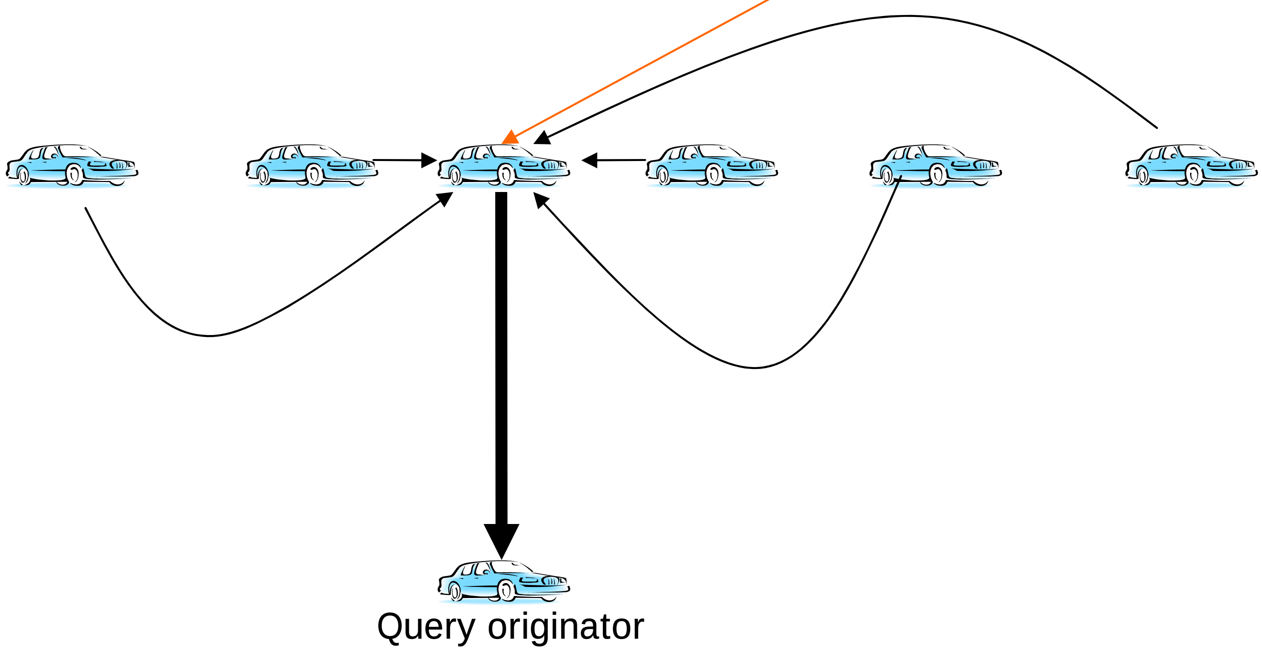

Query Processing Modes (3)

Hierarchical solution

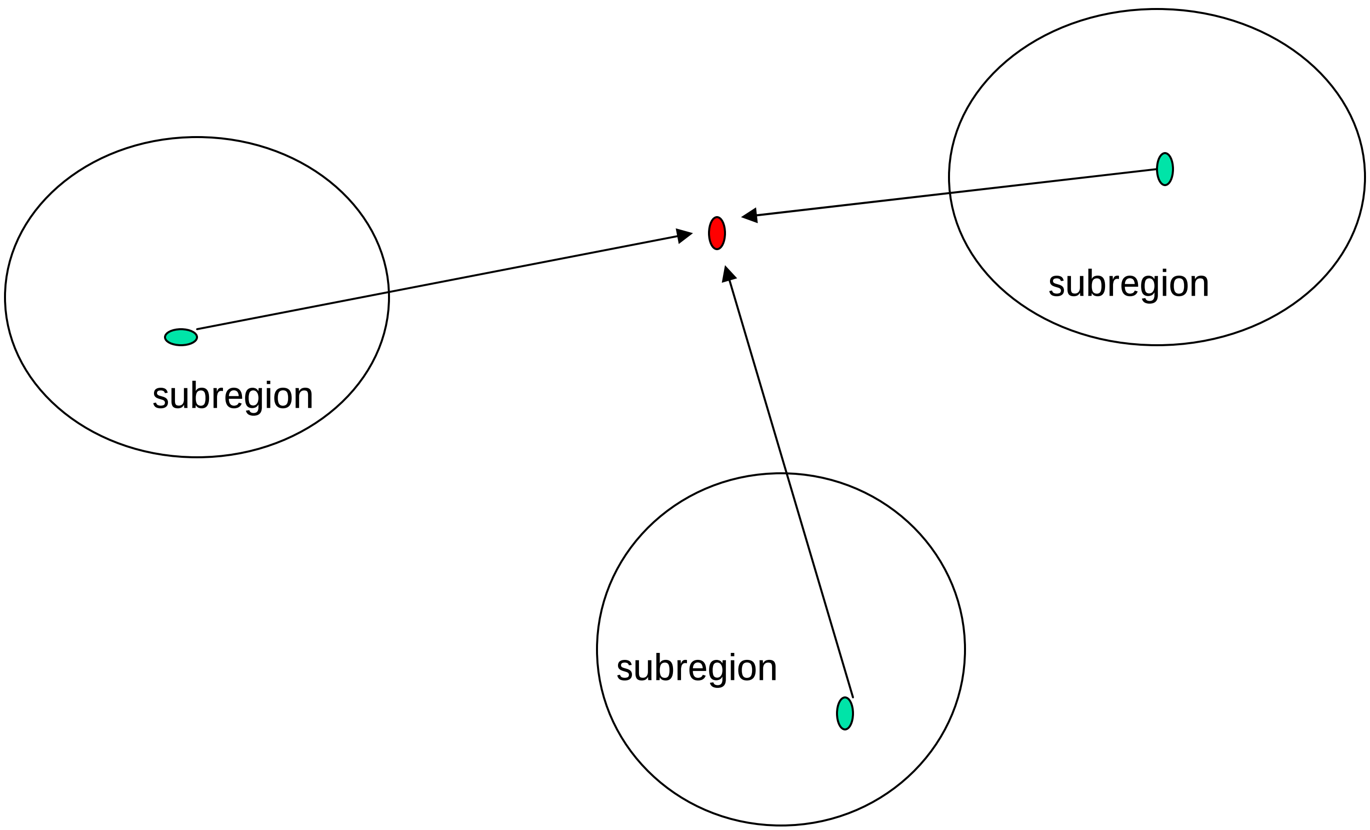

Another research issue

- <u>In network consolidation</u> of answers – delivery of query and partial answer
- Example: Q= “average speed in region R”
- When transmitting the query to a peer, the partial answer computed so far is also transmitted (static sensor network technique)

Remote query-- research issues

- Query dissemination to R
- Query processing in R
- Answer delivery to query originator, o

How is query originator v found?

- Via the infrastructure using node-id
  - May be costly
- In p2p mode
  - v sends future trajectory in query
  - Issues similar (more difficult) to query-region delivery

Other research issues in Remote Querying

- Graceful degradation of precision, depending on: peers density, budget, etc.

- Dynamic/adaptive use of infrastructure.

Research issues in data management

- Query Processing
- Dissemination analysis
- Participation incentives
- Remote Querying
- Data Integration of sensor and higher level information (maps, trip plans, ride-sharing profiles)
- Other relevant work

Moving Objects Database Technology

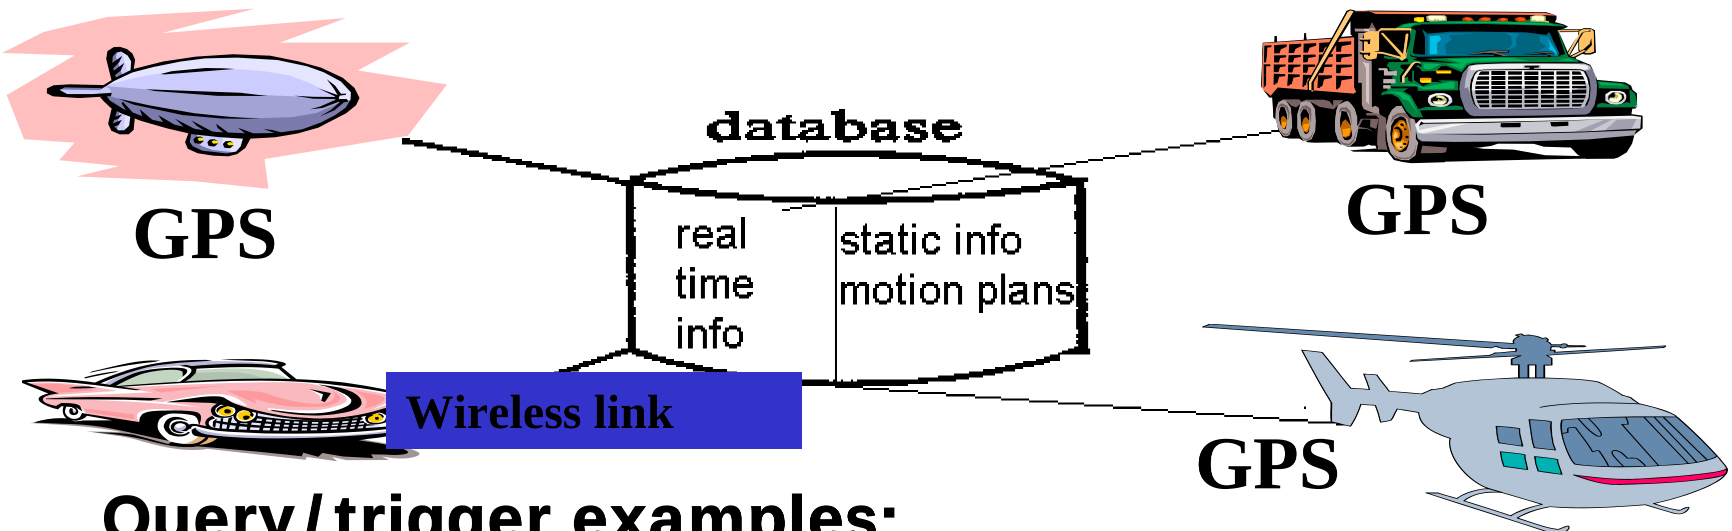

Query/trigger examples:

- During the past year, how many times was bus#5 late by more than 10 minutes at station 20, or at some station (past query)
- Send me message when helicopter in a given geographic area (trigger)
- Trucks that will reach destination within 20 minutes (future query)
- Tracking for "context awareness"

Data placement/replication

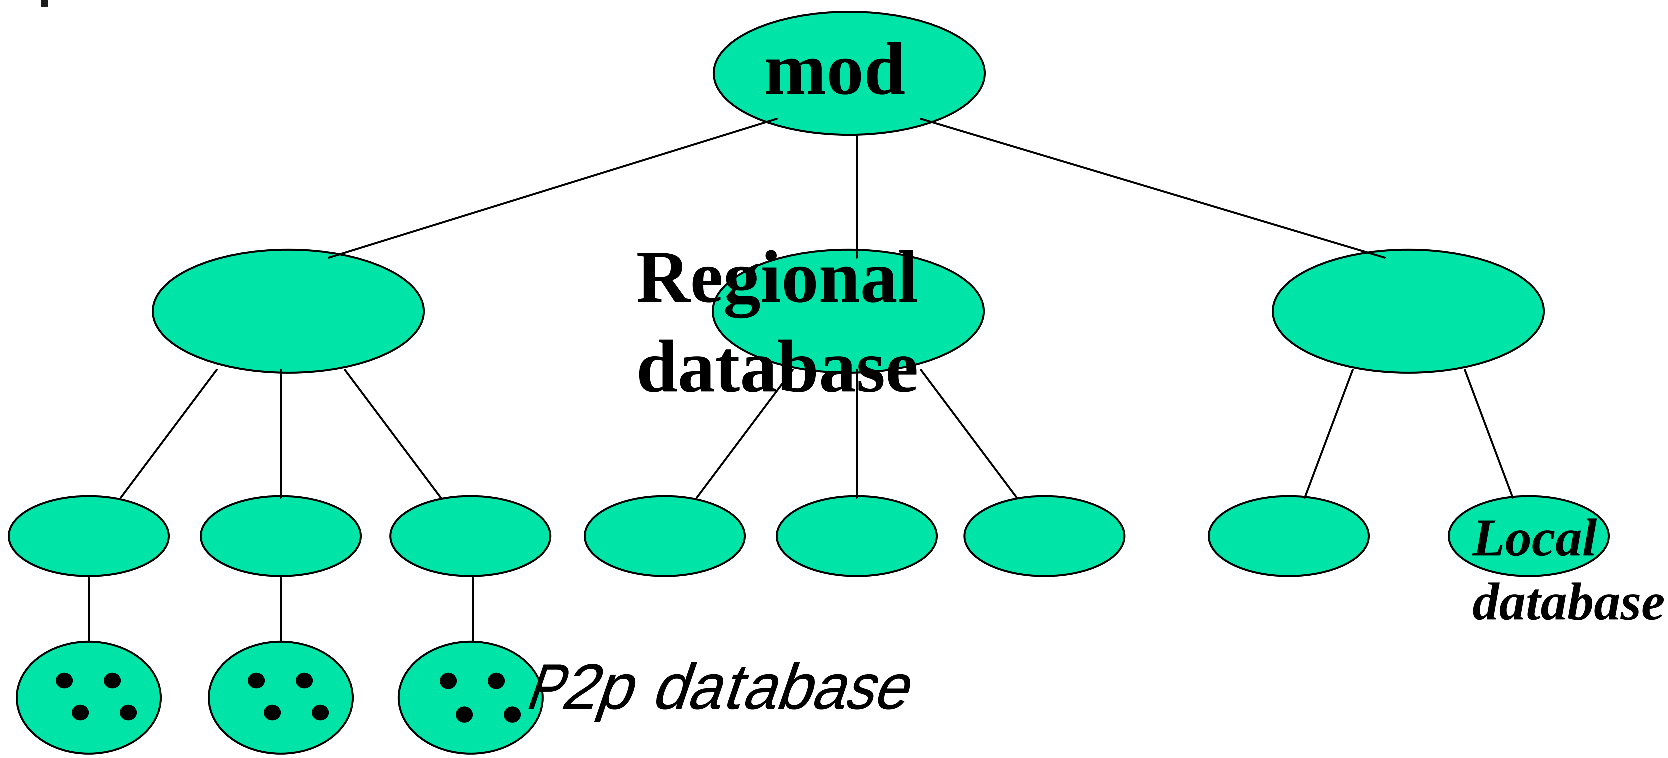

- Dynamic data replication (Wolfson, Jajodia Huang (1997)

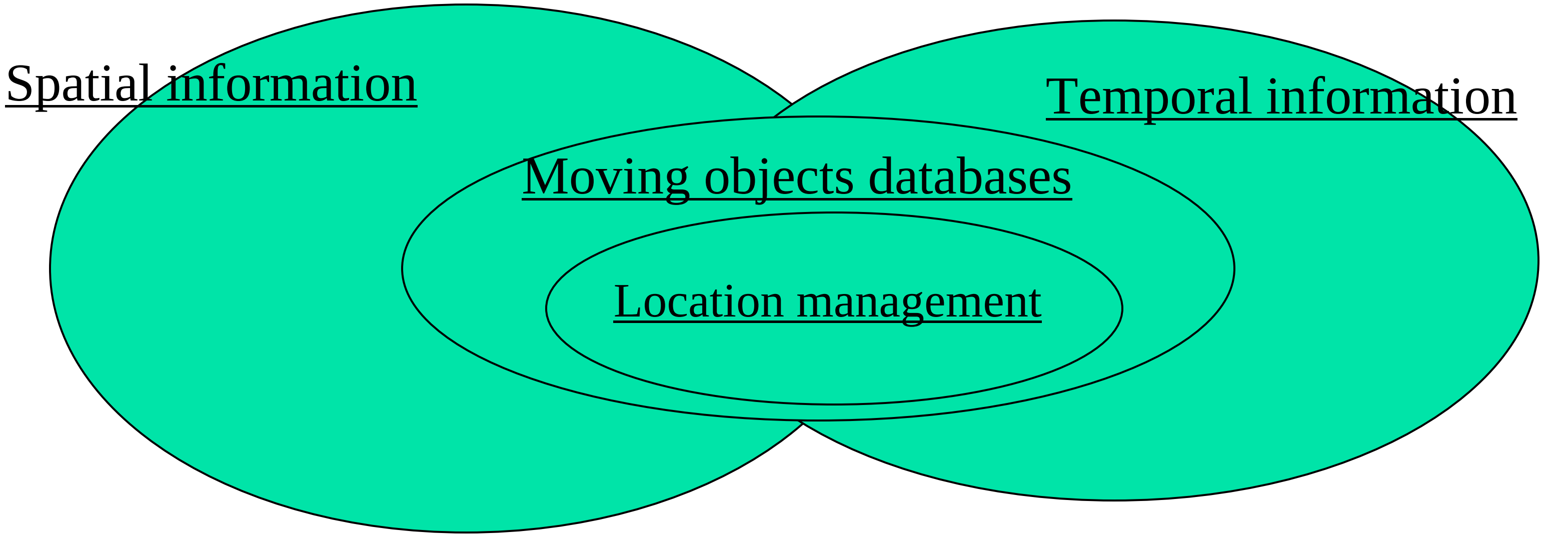
Spatial information
Temporal information
Moving objects databases
Location management

Spatial databases

- Geospatial
- Layout of a vlsi design
- 3d model of human body
- Protein structure

Main concepts: points, lines, regions

Common Instances: plane partitions (eg US into states), networks (roads, telephone)

Operations:

Intersection (line,line) -> points
(region,region) -> region
Inside (point,line) -> bool
Adjacent (region,region) -> bool

Implementation: indexing structure e.g. R-tree, extensible DBMS

Temporal databases

- Main concepts
  - Time instant
  - Period, 1/2/05 – 1/15/05
  - Interval, eg 2 months
  - Date
  - Transaction time vs. valid time (eg salary updated 1/1/05, effective 12/1/04)

Research issues in moving objects databases

- Location modeling/management
- Linguistic issues
- Uncertainty/Imprecision
- Indexing ←
- Synthetic datasets
- Compression/data-reduction
- Joins and data mining

Indexing

Used for fast processing of range queries

• Range query: Given a spatial region R and a time interval [t1, t2], retrieve all objects that will be in R at some time during the interval.

• Restricted range query: Range query when the time interval is a single point t1.

• In case of one dimension, region R is a line-segment; in higher dimensions it is a hyper-rectangle.

Performance measure: number of I/Os

Indexing Methods

•**Primal Space Method**:

• Consider the space together with time as an additional axis. Object movements form straight lines in this space (when moving with constant velocity).

•Consider the hyper rectangle X formed by the region R of the query and the given time interval. The answer is the set of objects whose lines intersect X.

Primal Space

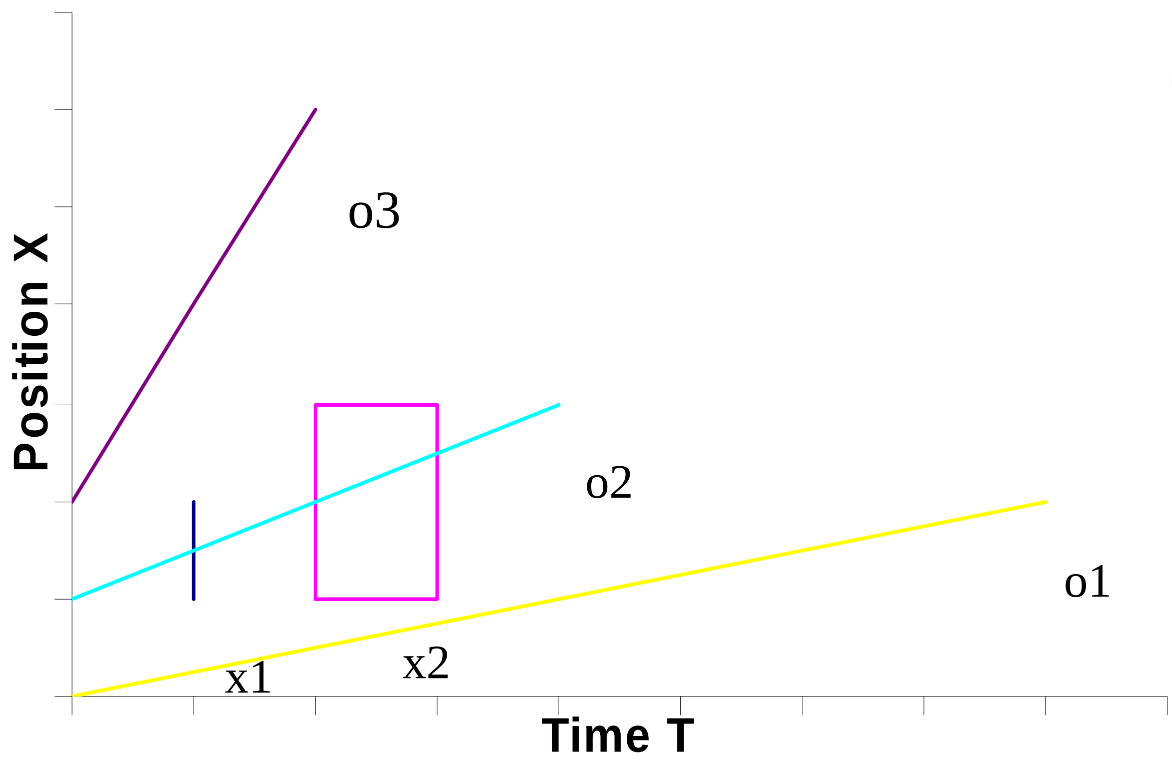

Query Q1: region R is the x-interval [1,2], time interval is the single point 1.

Q2: R is the x-interval [1,3], time interval is [2,3]

Dual Space Method: Consider the axes to be the coefficients in the equations of object motions.

Ex: In 1-dimension, the equation of motion of an object is $x=at+b$.

- The dual space has two axes a and b.
- Each object is represented by a single point (a,b) in the dual space.
- The answer to a range query is the set of points in the dual space enclosed in a region satisfying some linear constraints.

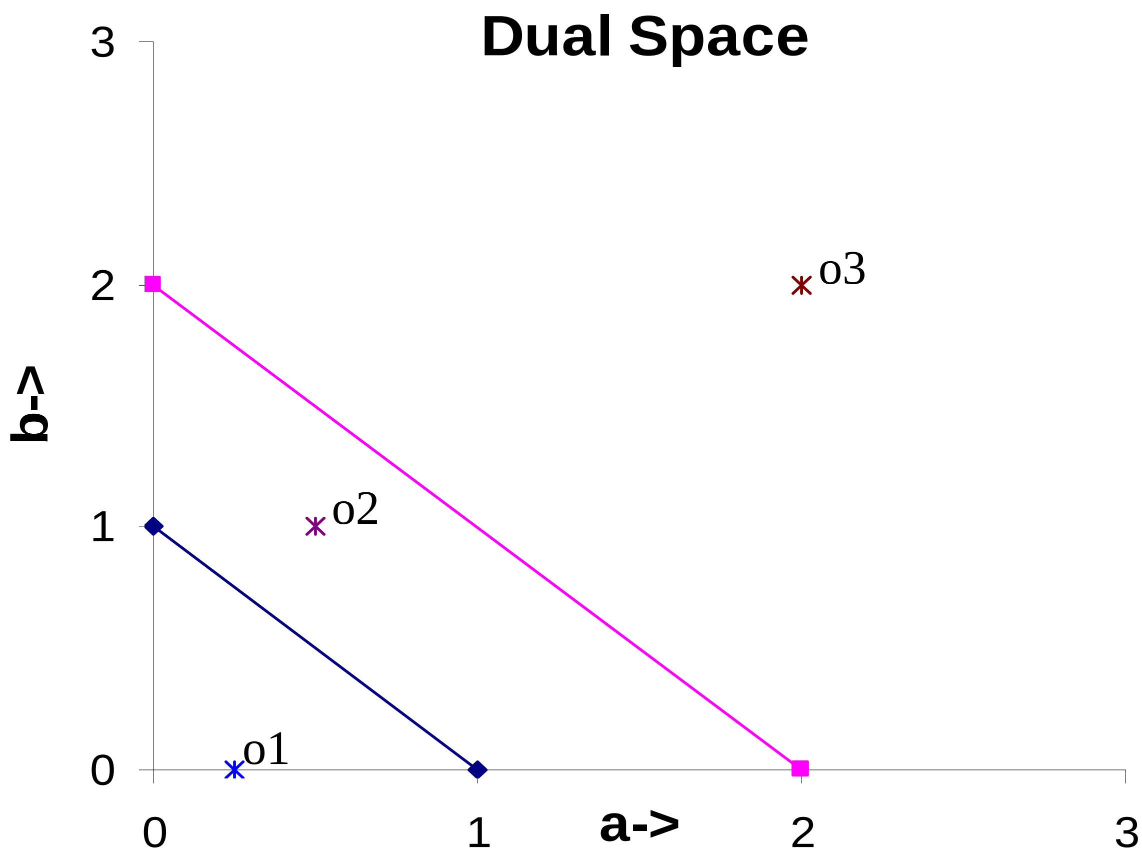

Answer to query Q1: all the points in the strip.

Q: Retrieve objects for which dynamic attribute has value v a≤v≤b at time t.

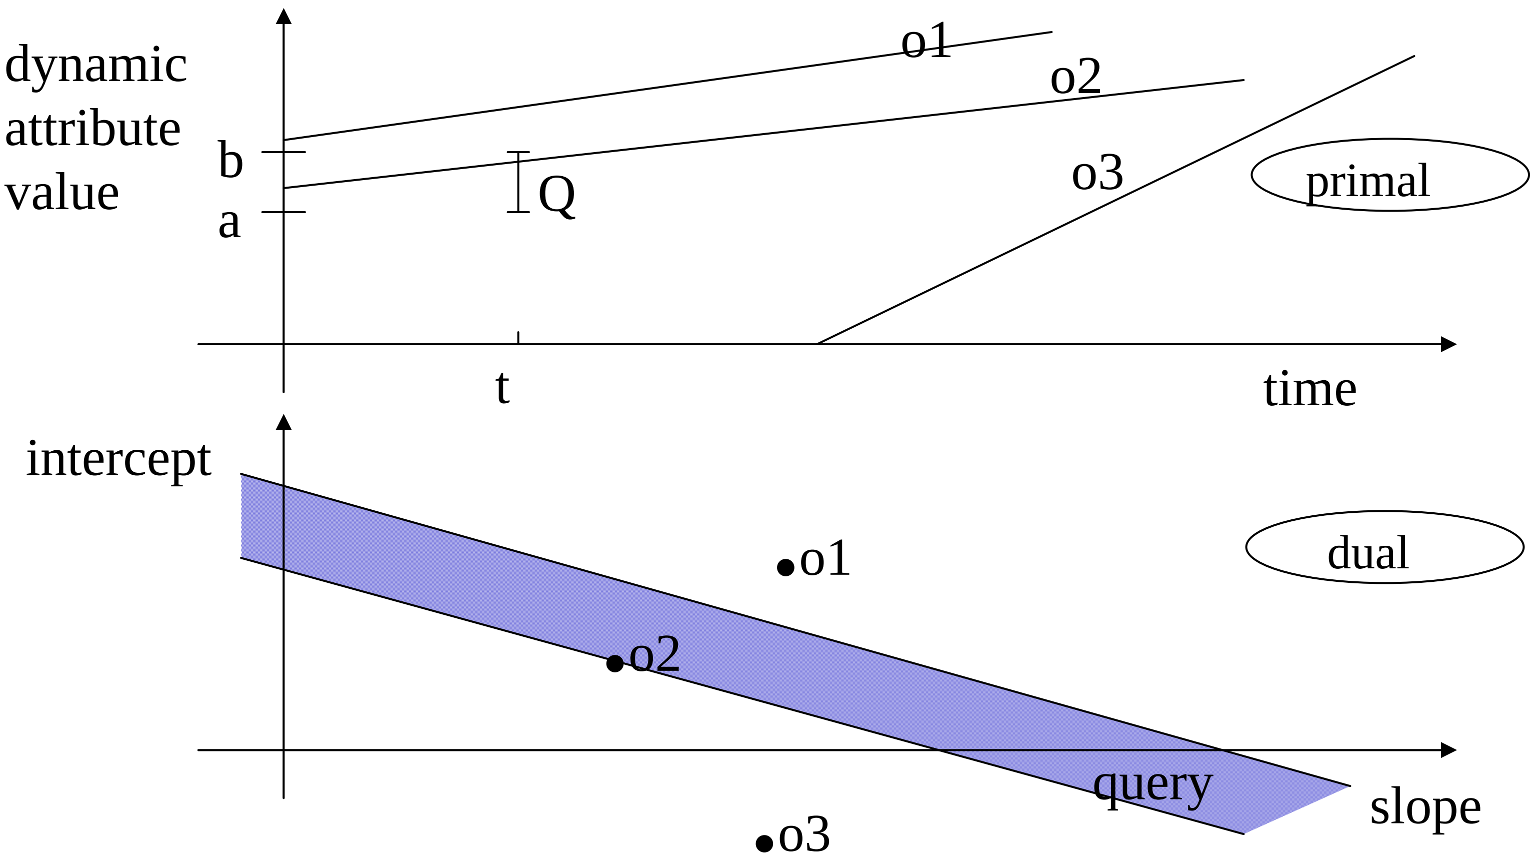

Spatial indexing to find objects that satisfy (intersect) query.

Primal Space Methods

First Method: (Sistla, Wolfson et al 97, Tayeb et al 98)

Divide time into periods of length T. For each period construct a multi-dimensional index using quad-trees.

•Divide primary space recursively into cells

•Store an object in a cell that its trajectory intersects

•The index can get large as an object may appear in more than one cell.

Indexing (cont.) Primal Space Examples (TUW98)

assuming a bucket size of 2 for the leaf nodes:

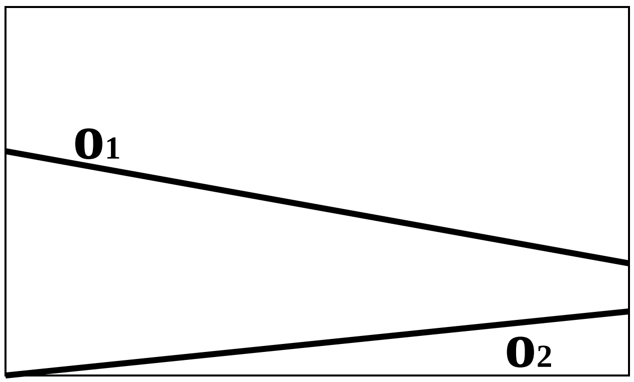

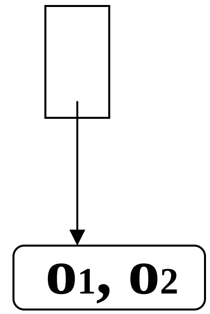

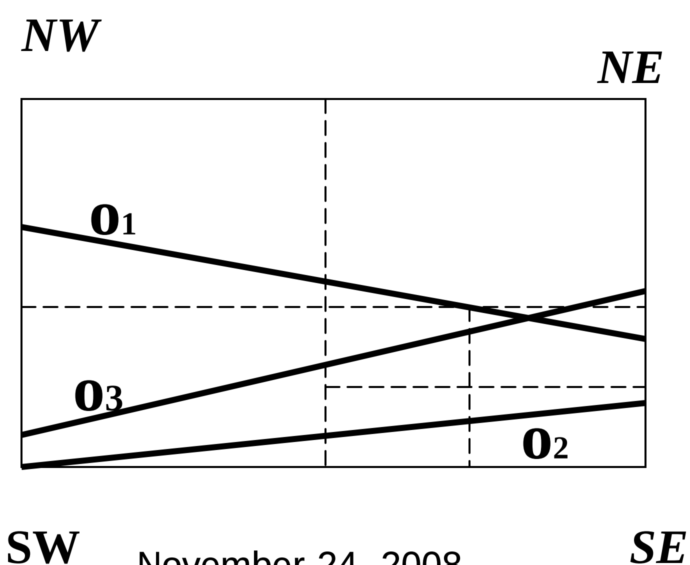

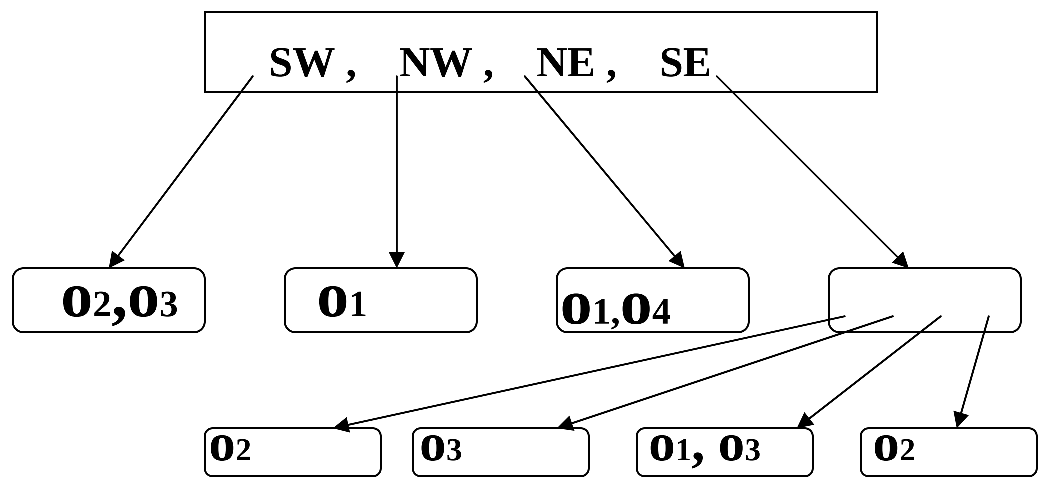

Performance analysis of primal plane representation using quadtrees. Tayeb, Ulusoy, Wolfson; Computer Journal, 1998

Example of implication: $\Longrightarrow$

30,000 objects 3 I/O's per range query

Primal Space Methods--- Contd

Second Method: (kinetic data structure**)** only for one dimensional motion.

•At any point in time we can linearly order objects based on their location. The ordering changes at those times when object trajectories cross.

•Fix a time period T and determine the orderings at the beginning and at T.

•Find the crossing points t1,t2,…,tm of the objects.

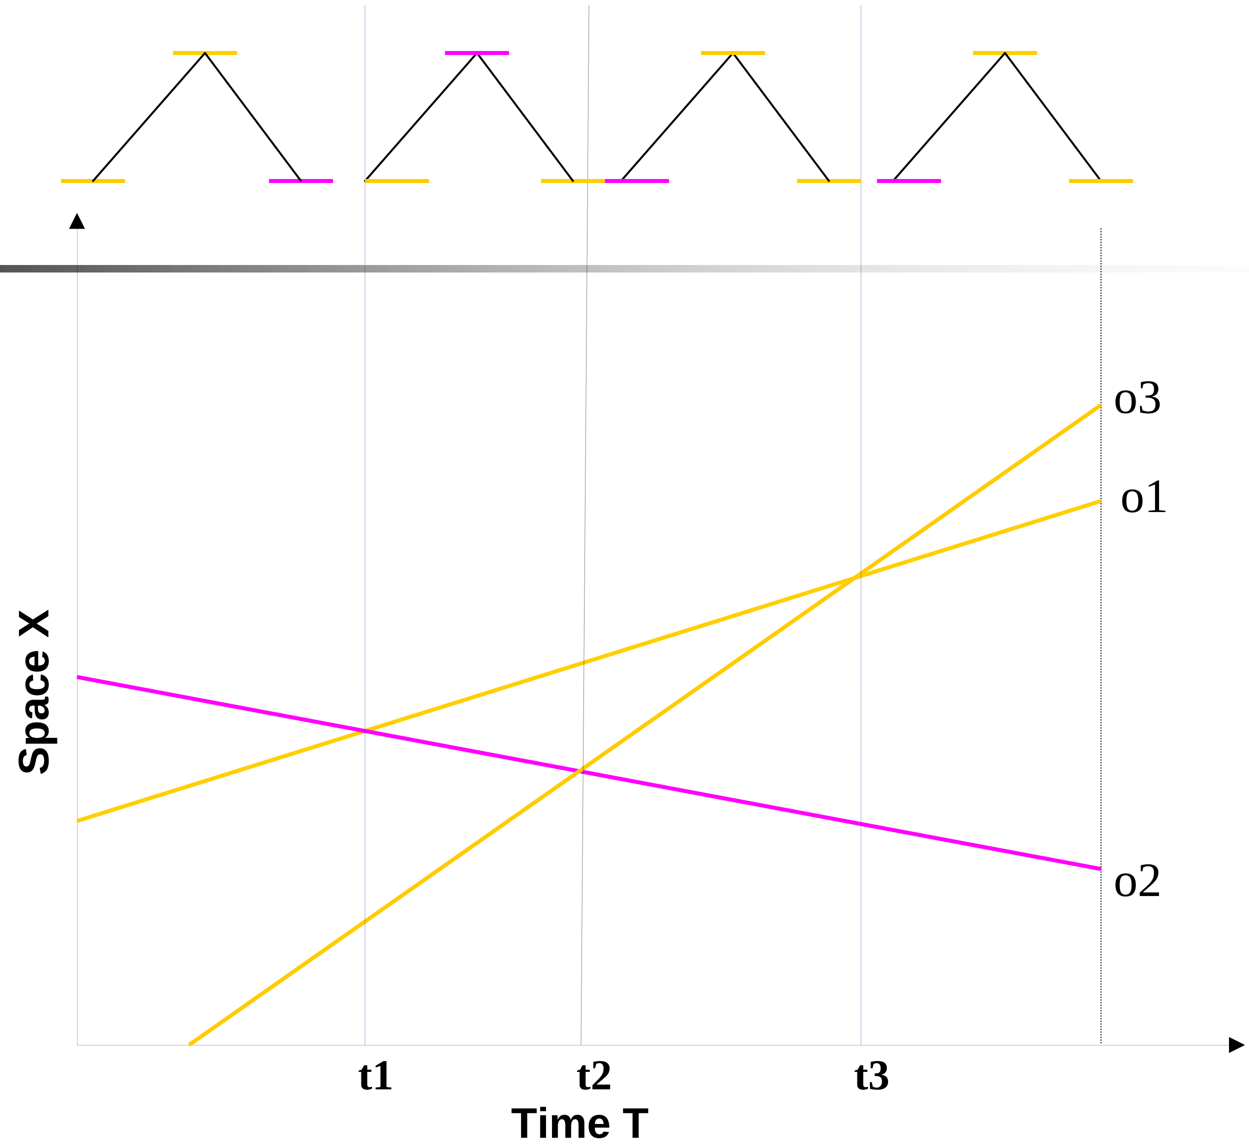

Determine ordering between crossing points

•Obtain the orderings between the crossings. Build binary search trees T1,…,Tm based on them.

•**To retrieve objects in the space interval I at time t do as follows**:

•Find a value j such that time t falls in the jth time interval.

•Use the search tree Tj to find all objects in the space interval I at that time.

Has space complexity O(n+m) and time complexity O(log (n+m)) where m is the number of crossings and n = N/B; N is the number of moving point objects, and B is the block size.

Primal Space Methods (continued)

- **Third method:** Saltenis et. al. 2000
- Time parameterized R*-trees (TPR)
  - works for motion in any number of dimensions.
  - Similar to R*-trees except that the MBRs are time parameterized.
- Objects are clustered and grouped into MBRs.
  - MBRs are enclosed into bigger MBRs.
  - They are arranged into a R*-tree.
- Each MBR has the following information.
  - Its coordinates
  - the min and max velocities of objects (in each direction).

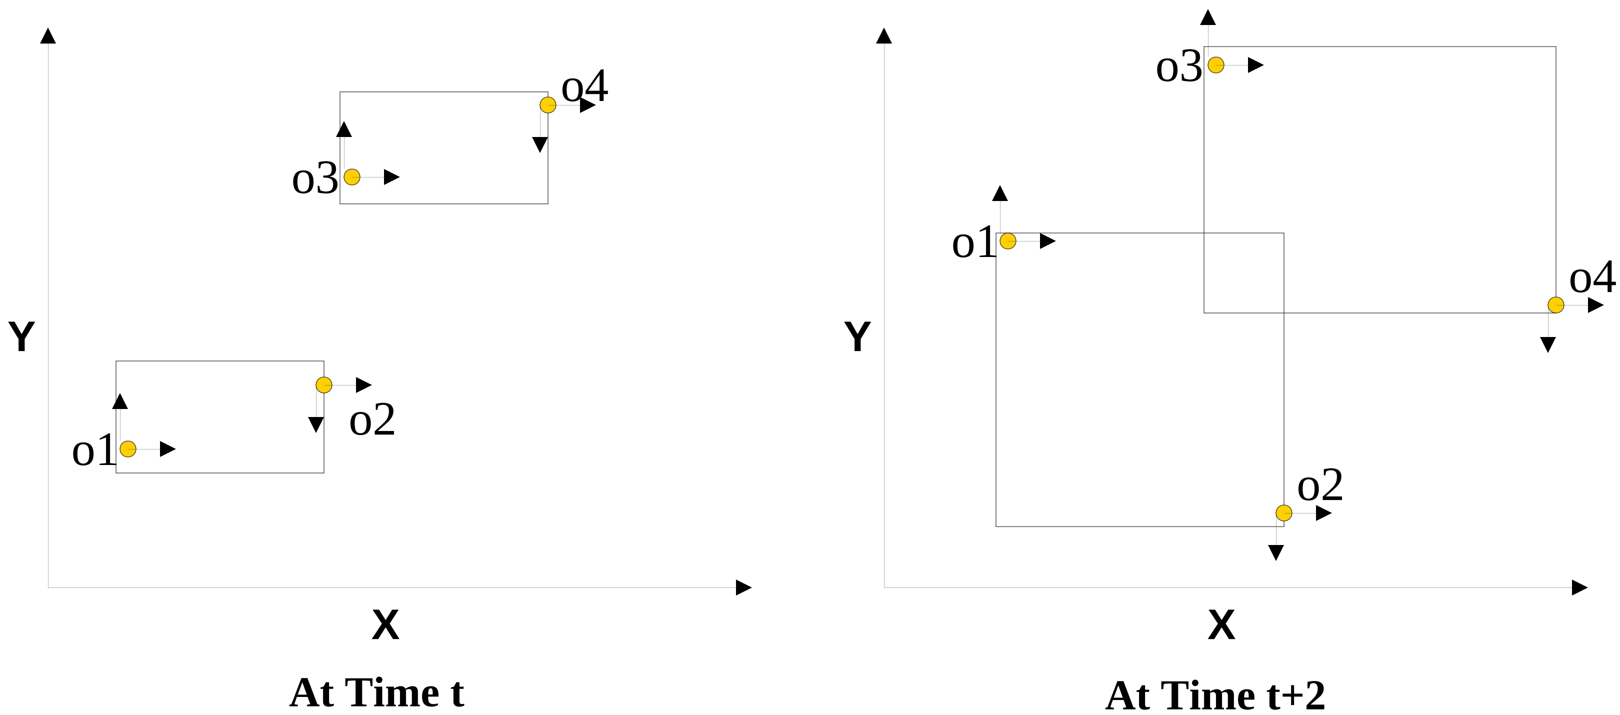

Leaf level MBRs overtime

- The position and sizes (i.e. the coordinates) change with time. The actual values can be computed at any time.

- Searching is performed as in R*-trees except that whenever an MBR is used its actual coordinates at that time are computed.

- The tree is reconstructed periodically

- Tree construction and insertions are processed so as to reduce the average area of the MBRs over the time period.

Dual Space Methods

[Kollios et al 99 (1-dimension), Agarwal et al 00 (2-dimensions)]

• similar to quad trees except:

- Employs Partition trees (used in computational geometry)
- Partition trees use simplicial partitions of sets of points.
- A simplicial partition of S is a set of pairs (S1,D1),…,(Sr,Dr) such that S1, …Sr is a partition of S. Di is a triangle enclosing points in Si.

For the given set of moving point objects, a partition tree is constructed satisfying the following properties.

- The leaves are blocks containing moving objects.
- A triangle is associated with each node in the tree.
  - The vertices of the triangle are stored in the node.
  - This triangle contains all the object points of the sub-tree.
- The sets of points and triangles associated with the children of a node form a ***balanced*** simplicial partition of the set of nodes of the parent.
- The size of the tree is O(n) and the height O(log n).

Searching for the points enclosed in a time/space region X:

Recursively search starting from the root.

- If the triangle at a node is contained in X then output all points in the subtree.
- Otherwise, recursively search along the sub-trees of the children whose triangles intersect X.
- At a leaf node, output all points in the node that are in X.

Range query Complexity: approximately O(k+ Sqrt(n)); k is output size

Insertion/deletion--- $O(\log^2(n))$ amortized complexity.

Solution Paradigm

- Geometric Problem Representation in Multidimensional Time-Space
- Spatial Indexing of Geometric Representation

Indexing with Uncertainty

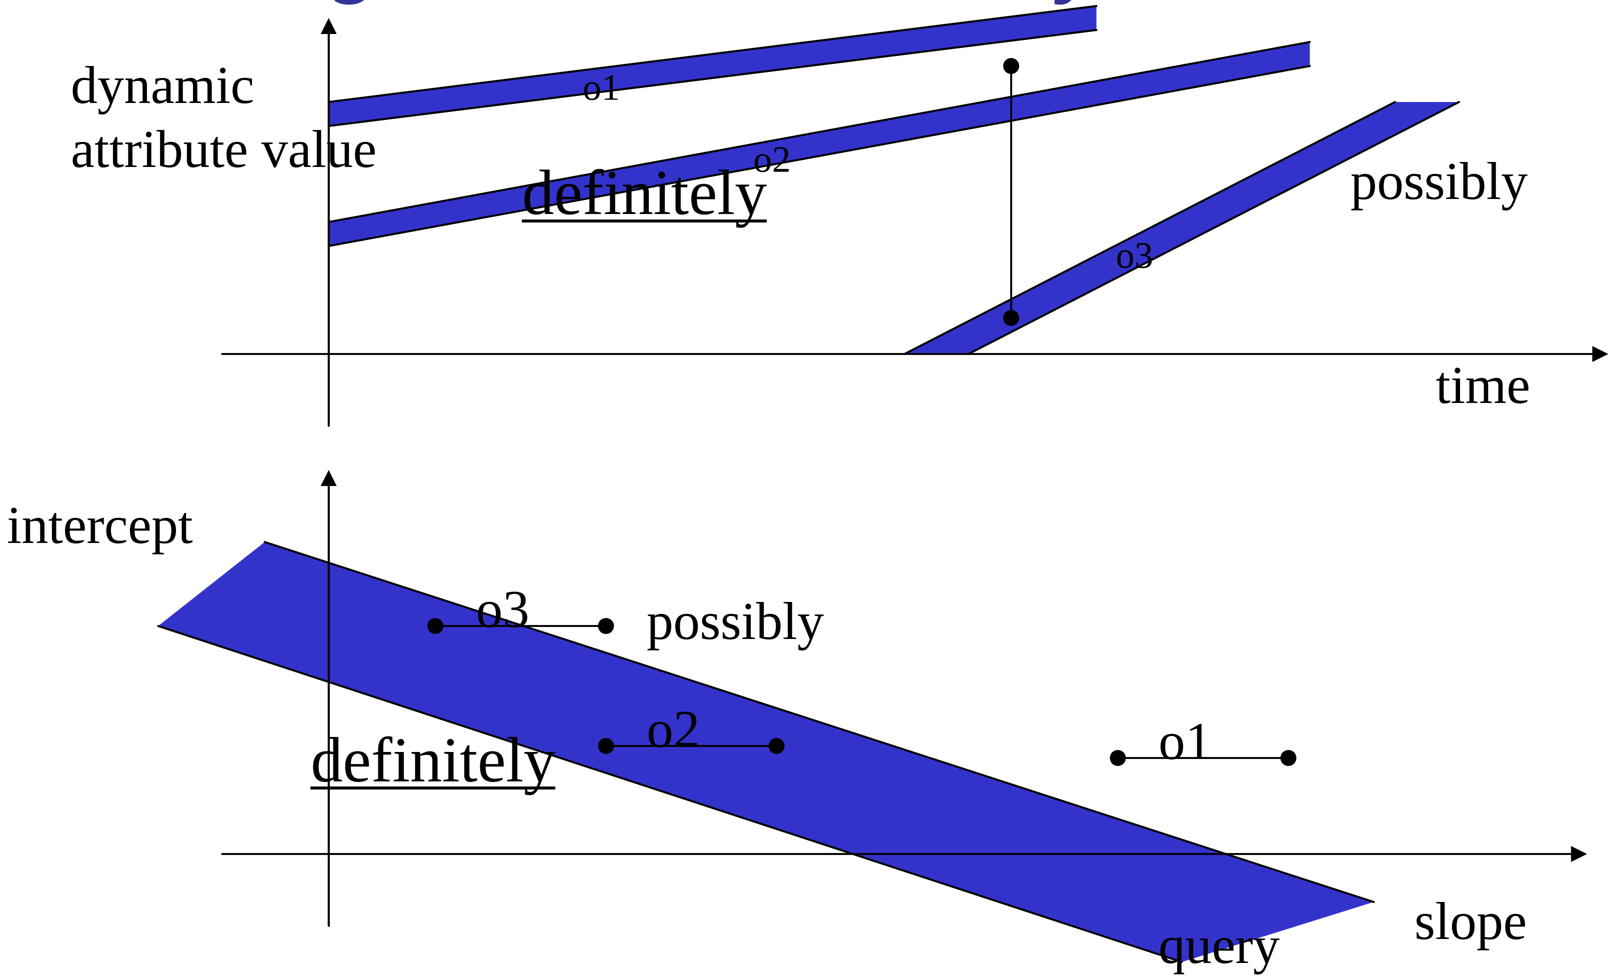

Conclusion

- Query processing
  - Basics of MP2P search
  - Data push and pull (state/ful/less, topologies, flooding, negotiation, store/forward
  - Power, memory, bandwidth management: transmission freqency:size tradeoff
  - Ranking of information: economic model
  - Experimental/commercial projects, Performance metric, and comparison
  - Backchannel exploitation
- Dissemination analysis (pattern, coverage)
- Value of information, comparison with Client/Server
- Participation incentives
- Remote Querying
- Data Integration of sensor and higher level information (maps, trip plans, ride-sharing profiles)
- Related approaches

Add

- DTN's,
- wake-up schemes,
- Incentives: Madria (store to maximize revenue), negotiation (eg bidding) 1 hop and multihop
- grassroots,
- vary transmission range (Santi)

Query Processing outline

- Basics of MP2P search
- Data push and pull
- Power, memory, bandwidth management
- Ranking of information
- Experimental/commercial projects, Performance metric, and comparison
- Backchannel exploitation

Conclusion

- Mobile P2P – new interesting problem
- Approach to data dissemination:
  - Flooding modified with
    - Unicast and broadcast
    - data push/pull
    - Energy bound
    - Ranking
    - Adaptive-size broadcast based on
    - Backchannel exploitation
- Formula for bandwidth/energy optimization
  - Broadcast-size = f(broadcast-frequency)

Conclusion

sensor-rich-environment + short-range wireless ➔ mobile p2p

- Query processing
- Dissemination analysis
- Value of information, comparison with Client/Server
- Participation incentives
- Remote Querying
- Integration of Mobile p2p and MOD's

Future Challenges

- Prolong network lifetime
  - Employ redundancy and node density
- Sparse networks
  - Many algorithms do not perform well if the network is not dense
- Rapid topology changes
  - Self-configuration and reconfiguration
- Global behavior from local knowledge
  - Achieve desired global goal using adaptive localized algorithms
- (Self-) Localization techniques
  - GPS is not available indoors
- Integration of MANET & WSN with wired or cellular networks
  - Connecting wireless sensor networks to the Internet

RANDI

When two peers meet they conduct a two-phase exchange:

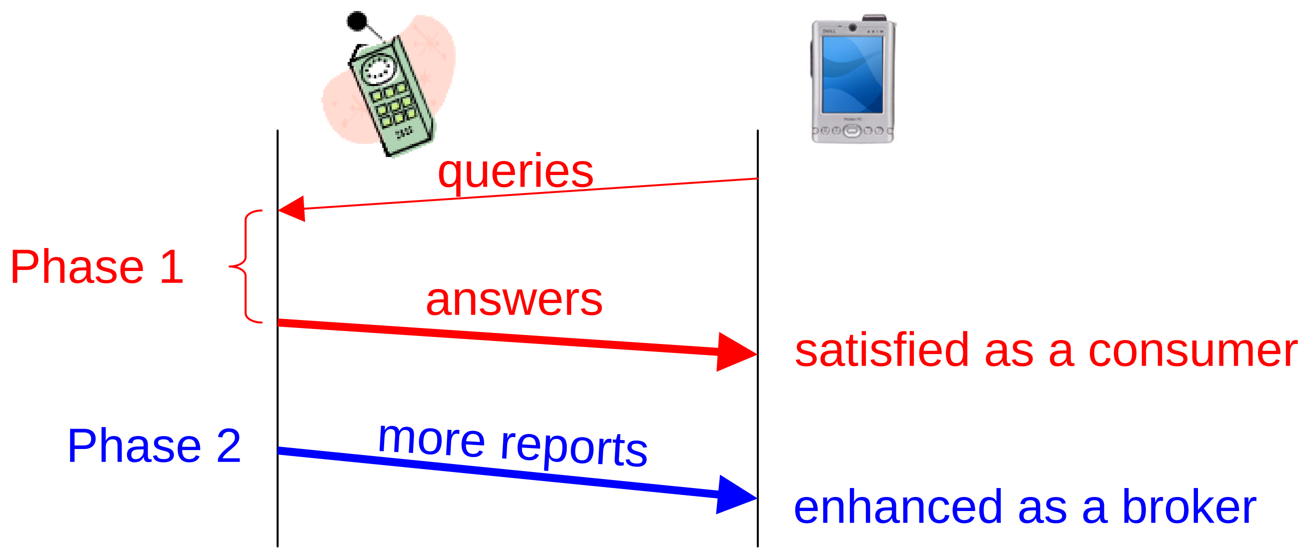

Phase 1: Exchange queries and receive answers.
Phase 2: Exchange more reports using available energy/bandwidth
Combination of:
    unicast (thin line) and
    broadcast (thick lines) to enable overhearing.

RANDI (Cont'd)

Two interaction modes which combine pull and push

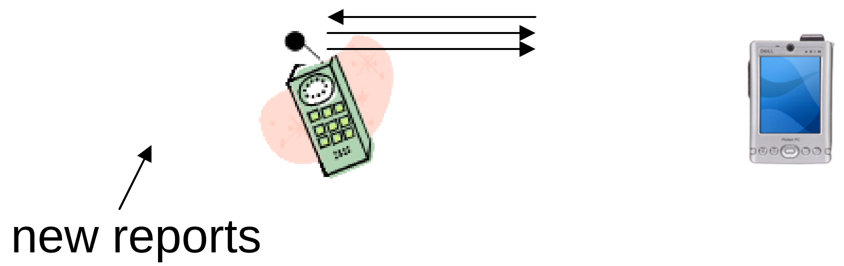

- Query-response: triggered by discovery of new neighbors
- Relay: triggered by receipt of new reports
  - Disseminate to existing neighbors

How much faster is IGS than BS?

In a Mobile Opportunistic p2p system, the answer depends on:

- Transmission (rendevous) range -- *r*
- Traffic speed -- *v*
- Broker density -- *g*
- Resource density -- *s*
- Relevance threshold -- H_0
- (Resource available time) / (resource unavailable time) -- *k*

Outline

- Introduction
- Resource-capture strategies
- Performance evaluation

Simulation Environment for OP2P

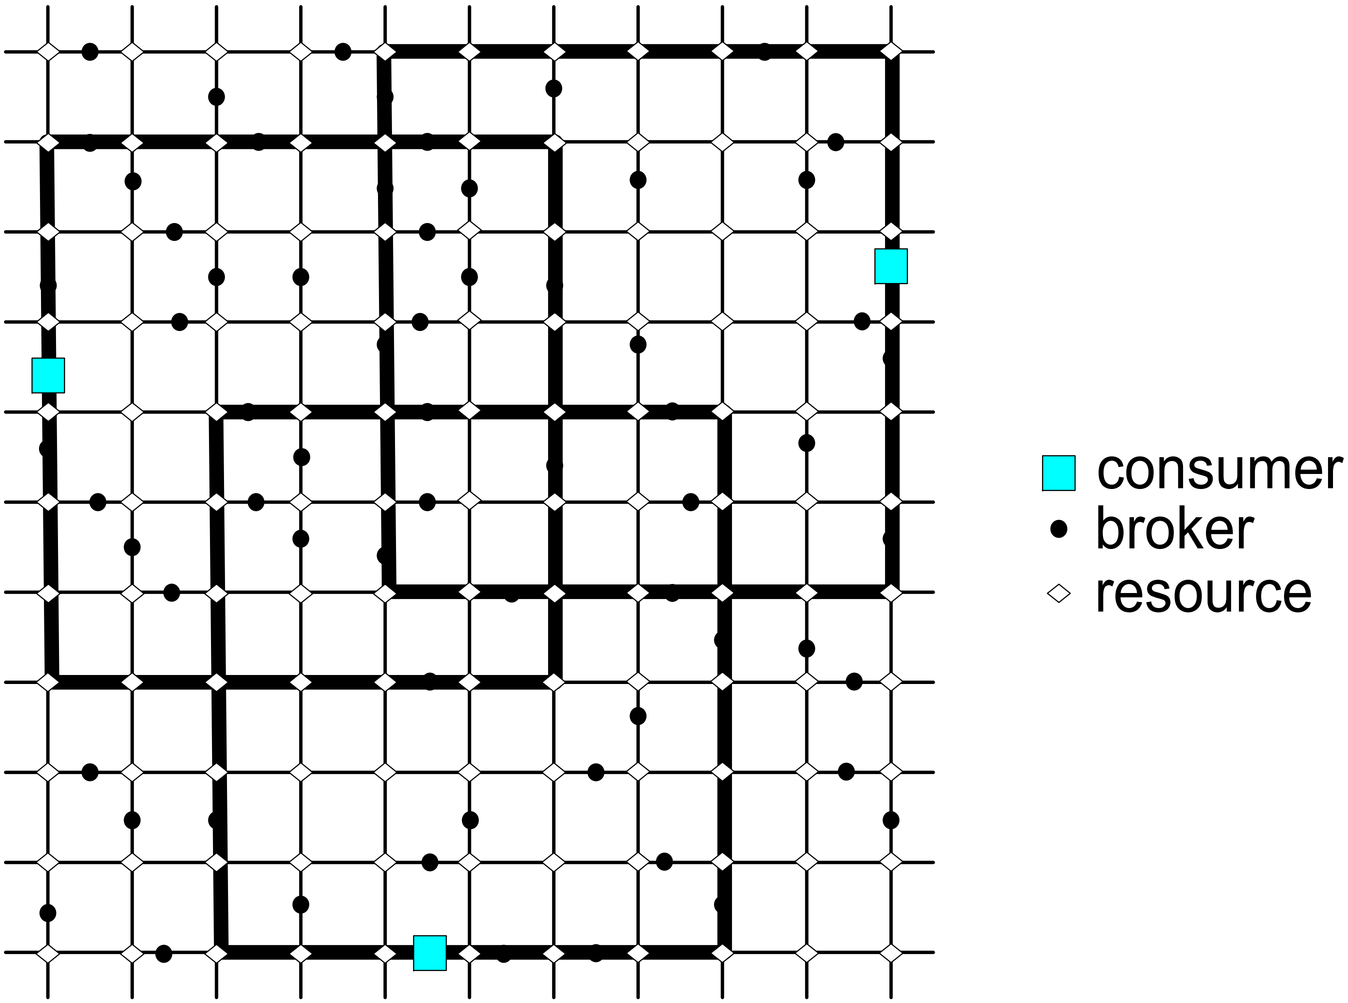

Values of Parameters

Parameter	Symbol	Unit	Value
Transmission range	r	meter	50, 100, 150, 200
Motion speed	v	miles/hour	10, 20, 30, 40, 50, 60
Broker density	g	brokers/mile^2	0, 50, 100, 150, 200
unavailable/available time	k		10, 20, 30, 40, 50
Relevance threshold	H_0		0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9
Hotspot density	s	hotspots/mile^2	17, 36, 100

Simulation Results
-- Relevance Threshold

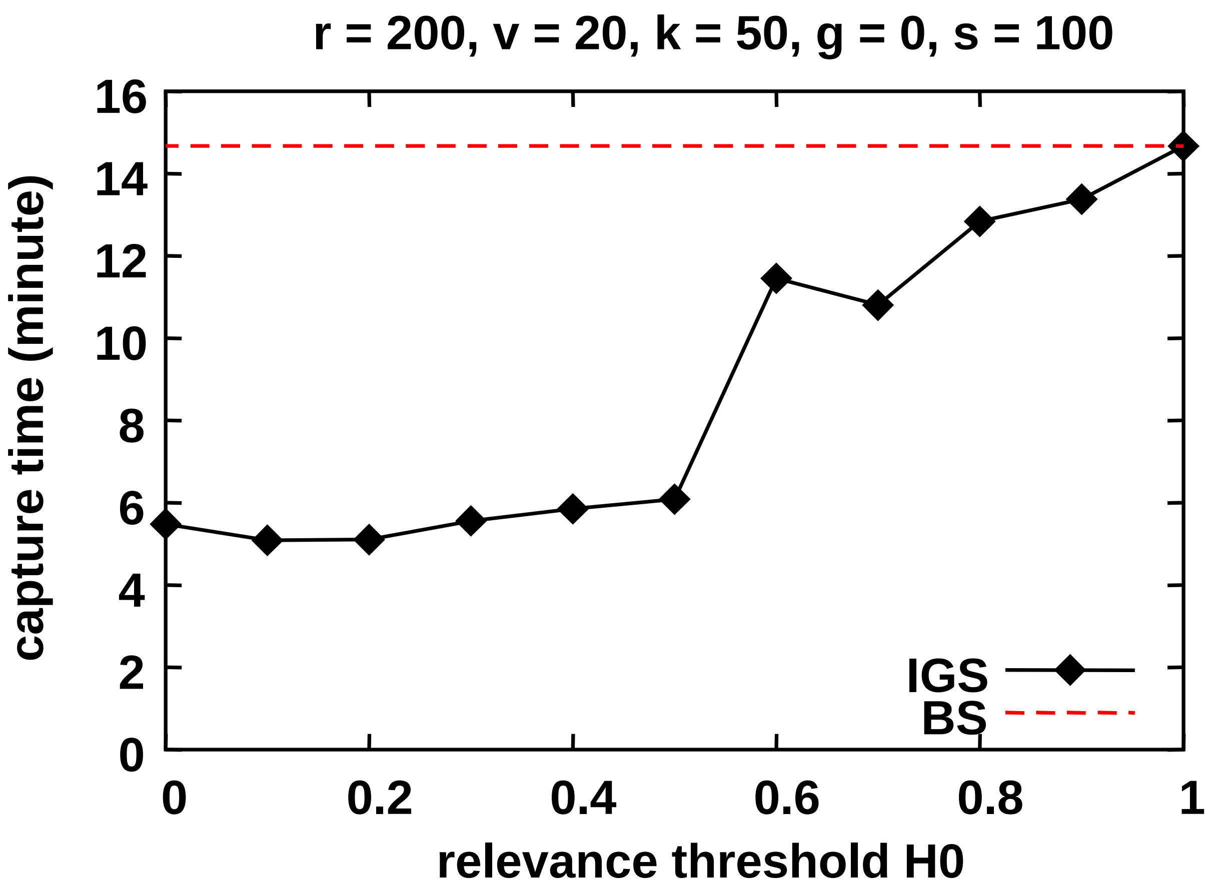

Even reports with a low relevance are better than no reports at all.

Simulation Results
-- Broker Density & Transmission Range

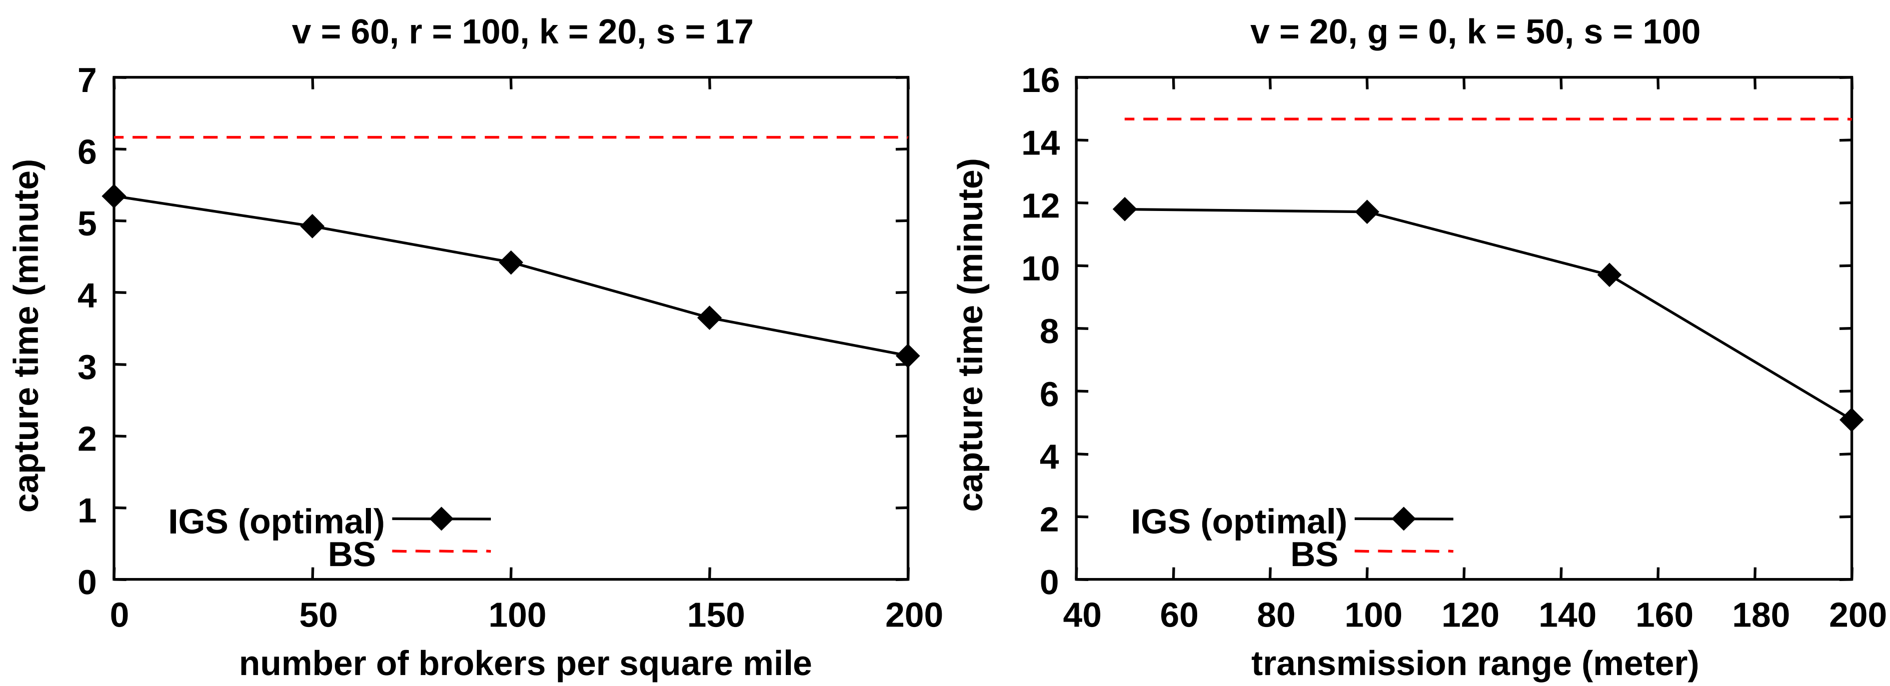

- The discovery time of BS is not affected by the transmission range and the broker density.
- speed of information propagation ↑ ➔ IGS discovery time ↓

Simulation Results
-- unavailable-to-available Ratio

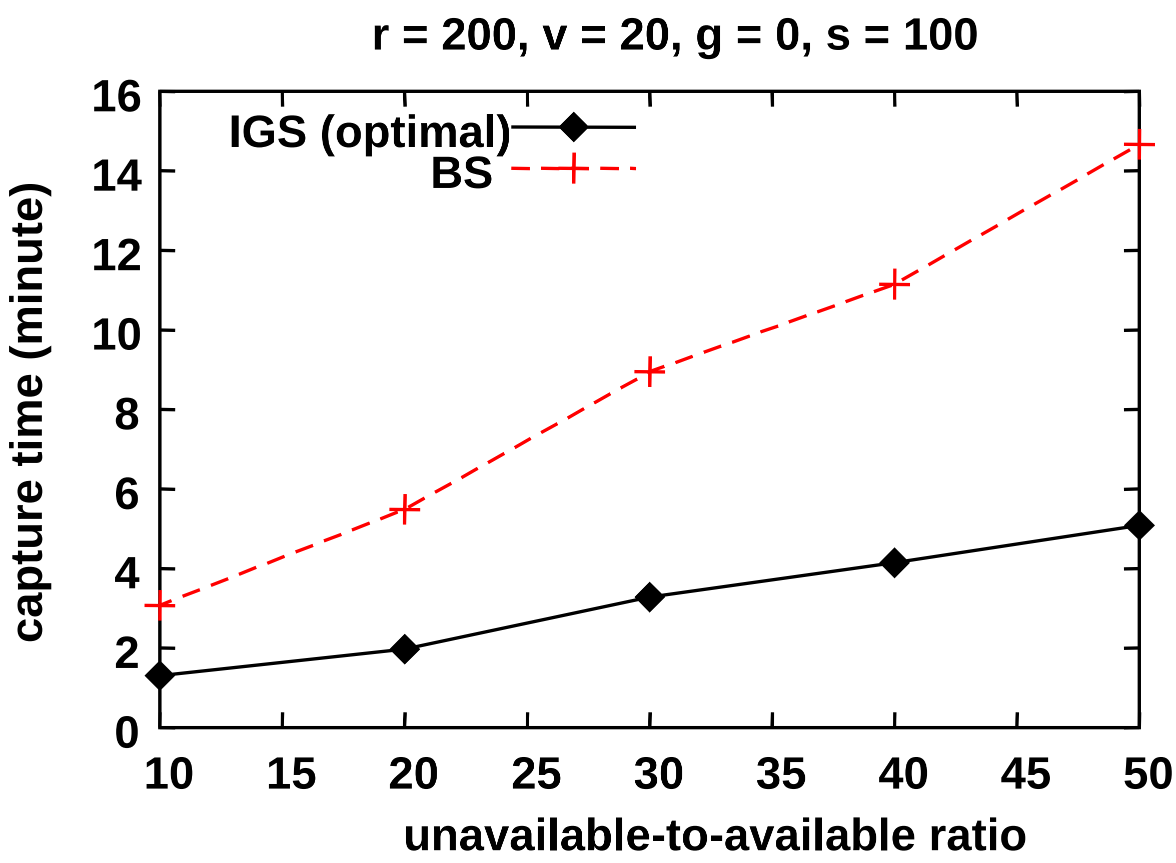

Time(BS) - Time(IGS) ↑ as unavailable/available ↑ ➔

IGS advantage ↑ when competition ↑.

Simulation Results -- Resource Density

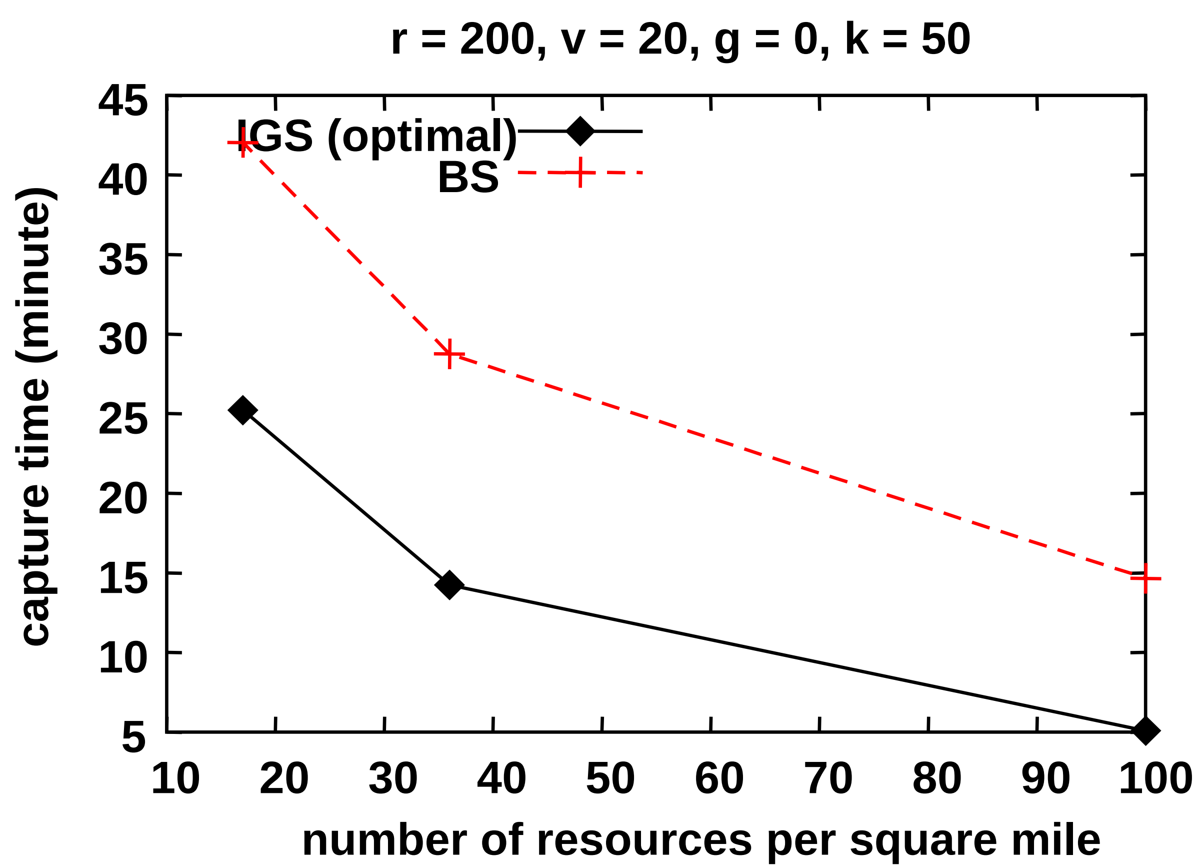

- Resource density ↑ ➔ time(BS) - time(IGS) ↓ ➔ IGS advantage ↑ as competition ↑

Simulation Results -- Motion Speed

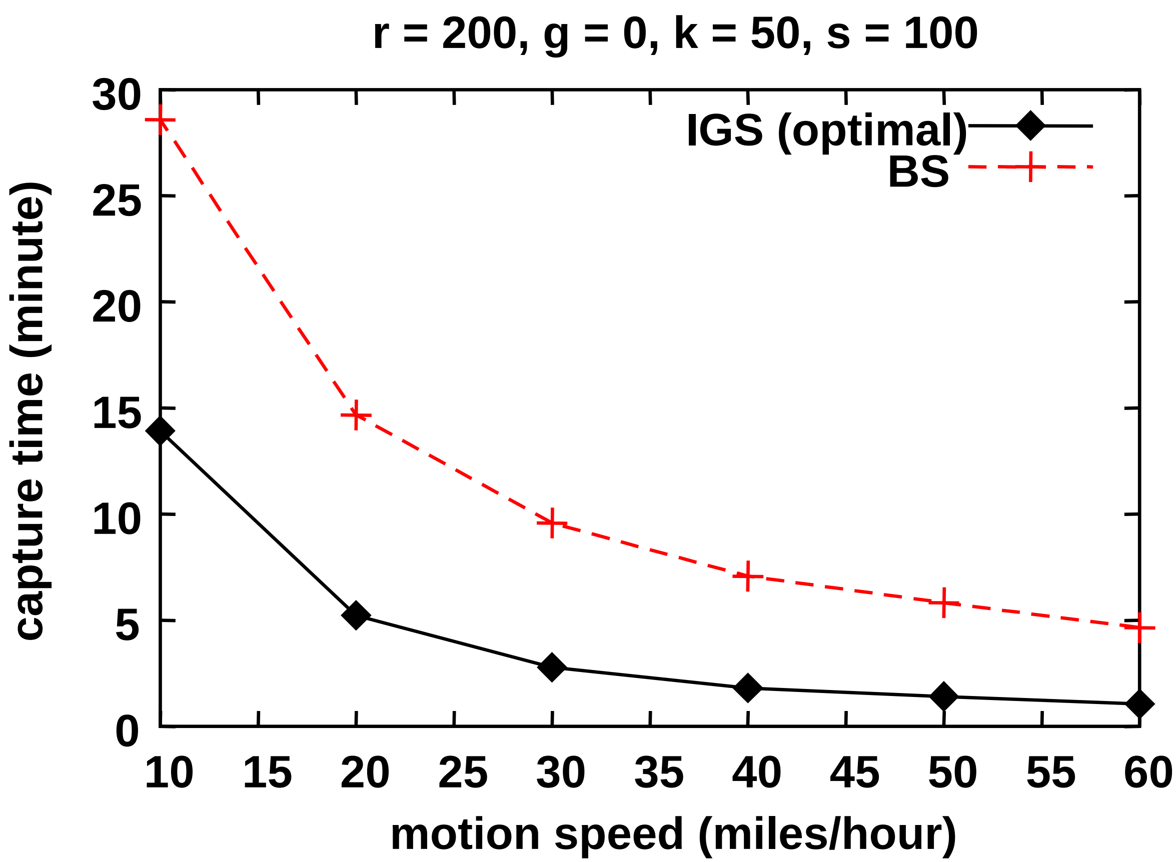

- speed ↑ ➔ Time(BS) - time(IGS) ↓ ➔ IGS advantage ↑ with ↓ of speed.

Justification of exponential relevance function

$$\text{relavance of report a}(R) = e^{-\alpha \bullet t - \beta \bullet d}$$

Theorem:

Assume: (1) The length of the valid duration of *R* is a random variable with an exponential distribution having mean *u*. (2) The speed of the consumer is *v*.

If: $\alpha = 1/u$ and $\beta = 1/(u \bullet v)$

Then: the relevance of a report *a*(*R*) is the probability that the resource *R* is available when the consumer reaches *R*.

Observations concerning (semi-) competitive resources

- Information by itself is not sufficient to capture resource
- If move to obsolete resources may waste time compared to blind search

Mp2p vs. client-server

- Mp2p advantages
  - Zero cost
    - Unregulated communication
    - No central database to maintain
  - Independent of infrastructure
  - Higher reliability
  - Privacy preservation
- Mp2p disadvantages
  - Weaker answer-completeness guarantees

Relevant research work

Relationship to work on Mobile Ad Hoc Networks

- Work mainly concerned with sending a message to an ip-address

- In contrast, MP2P focuses on dissemination among a group interested peers

Future work

- Improve ranking by:
  - Reliability
  - Size
- Processing for other types of queries (e.g. min, joins)
- Privacy/security issues

Future work: Reducing Duplicate Backchannel Transmissions

- Negotiation

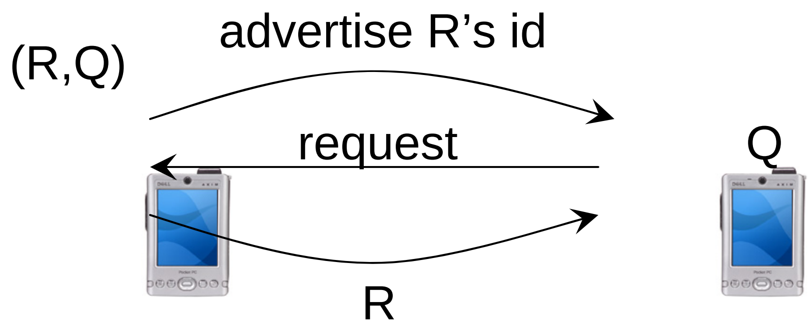

- Visit-trace

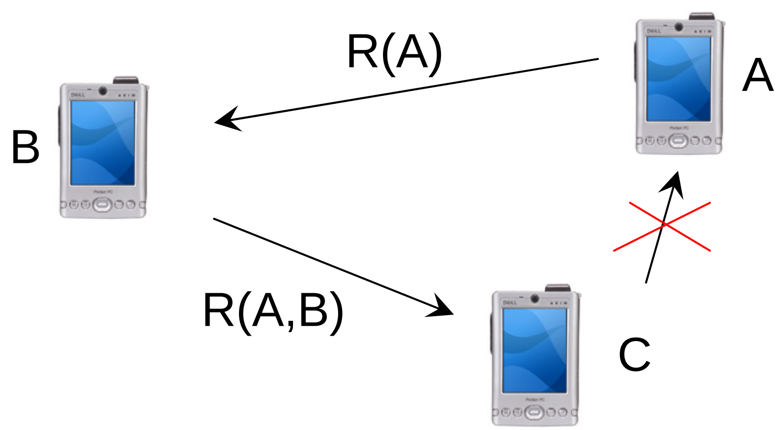

- Novelty estimation (machine-learning)