DRIVE - Disseminating Resource Information in VEhicular and other mobile peer-to-peer networks

Bo Xu
Ouri Wolfson

University of Illinois at Chicago

wolfson@cs.uic.edu

DRIVE objective

- Enable dramatic improvement of the travel experience – based on information
- Real-time information to traveler has not changed much in 40 years

Sensor-networked Transportation

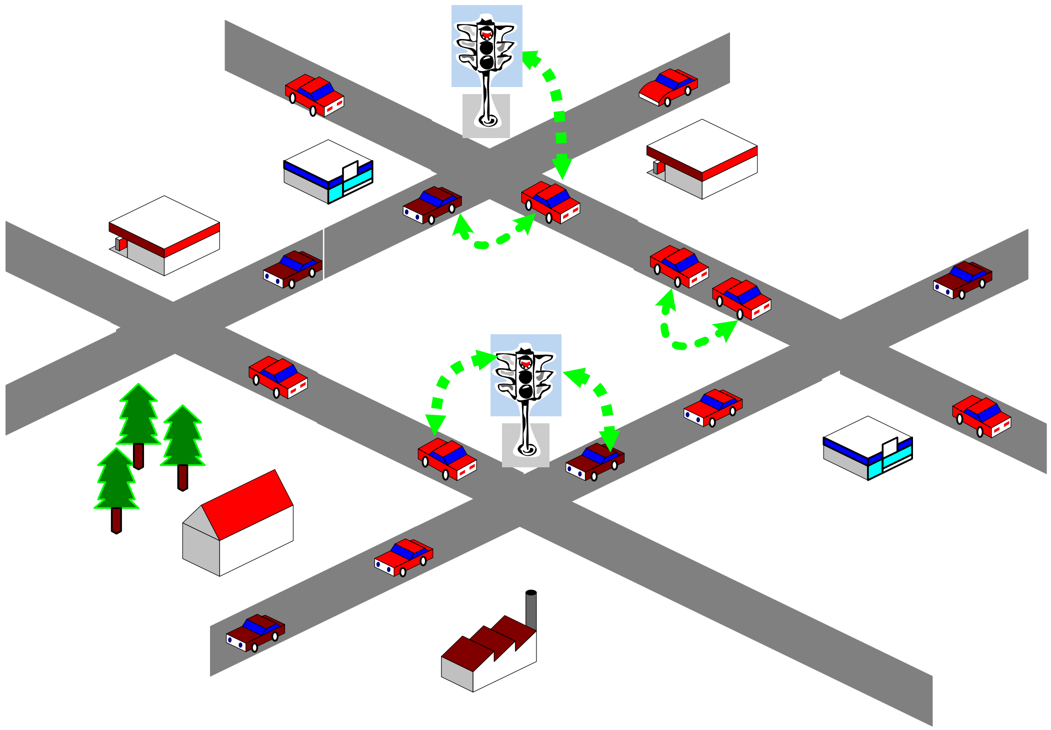

Vehicle sensors: speed, fuel, cameras, airbag, anti-lock brakes
Infrastructure sensors: speed detectors on road, parking slots, traffic lights, toll booth
Wireless Networking: tens Mbps, 50-100 meters (802.11, UWB, Bluetooth, CALM)

Application examples

- 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
  - Replay accident based on sensor traces
  - Infrastructure transmits speed-limit; dependent on vehicle type (works in France)

Application examples (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 sharing opportunities

Ride sharing – untapped potential

- 4% increase in ridesharing – offset 2000 congestion increase
- Example – most arriving airport passengers go downtown
- Initial efforts
  - Washington DC “slugging”
  - Illinois ride-sharing program at UIC, Prof. Nelson’s lab
- Wireless/short-range Peer-to-Peer communication enables real-time matchmaking
  - Eliminates need for 3rd party mediation, business model

Application examples (cont.)

- Beyond transportation:
  - Sighting of enemy vehicle in downtown Mosul in last hour?
  - Cockroach robots in disaster areas
  - Disseminate ticket-availability before a sporting event

How to enable these applications?

- Develop product that performs them
- Develop standards to build them
- 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
- Software tools:
  - Data mining
  - Travel-time prediction
  - Trip planning
  - Regional planning
  - ……

Research issues in data management

- Sensor data acquisition, modeling, fusion, dissemination
- Data usage strategies
- Participation incentives
- Remote Querying
- Data Integration of sensor and higher level information (maps, trip plans, ride-sharing profiles)

The players

- Moving/stationary objects with processing and communication power
  - Personal digital assistants (pda's)
  - Computers in vehicles
  - Processors embedded in the infrastructure

Collect, Organize, Disseminate, information about resources

- Resources -- examples
  - Gas stations
  - Parking slots
  - Cabs
  - Ride-share partners
  - Malfunctioning brake-light
  - Accident at a time/location

Resource reports are generated by infrastructure or moving objects sensors

Spatial and Temporal Resources

- Spatial resources
  - Examples: gas station at 342 State st., patch of ice at milepost 97, Italian restaurant at 300 Morgan St.
  - The importance/relevance of information decays with distance
  - Possible relevance function: $-\beta \cdot d$
- Temporal resources
  - Examples: Price of IBM stock at 2pm, DJI average at 10am
  - The importance/relevance of information decays with age
  - Possible relevance function: $-\alpha \cdot t$

Spatio-temporal Resources

Spatio-temporal resources: specific to time and location

- Traffic conditions, available parking spaces, occurrence of car accidents, taxi cab customers, ride-share partners

- The importance/relevance of a resource-availability report decays with age and distance
- Possible relevance function: $-\alpha \cdot t - \beta \cdot d$
- Each resource-availability report includes *create-time* and *home-location ---* sensor fusion tool

Relevance-ranked resource-type lists

Moving Object Memory:

Hazards and alerts

time	location

Parking Information

time	location

Traffic Conditions

time	location

Taxi cab customers

time	location

Each resource list keeps top K resources

Opportunistic Resource Dissemination (ORD)

- Each vehicle has an *interest profile*:
  - ranked list of resource-types
  - relevance-threshold in each type
- Two vehicles exchange local database information when they encounter each other (i.e. come within transmission range)

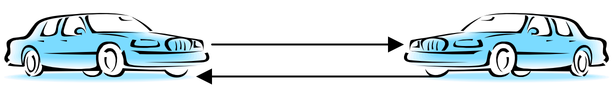

- Least relevant resources that do not fit in allocated memory are purged out

Exchanging and purging resources

Cab customers

Before exchange

Sears Tower (NE), 10:30am
Halstead & Taylor, 10:24am

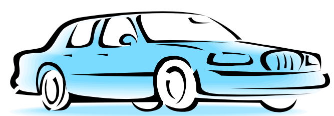

Navy Pier, 10:20am
Madison & Clark, 10:25am

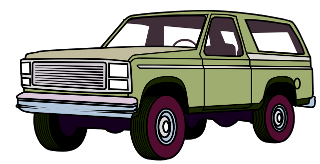

After exchange

Sears Tower (NE), 10:30am
Madison & Clark, 10:25am

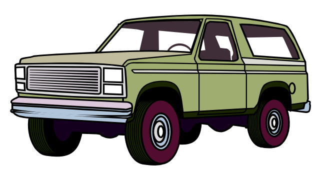

Sears Tower (NE), 10:30am
Madison & Clark, 10:25am

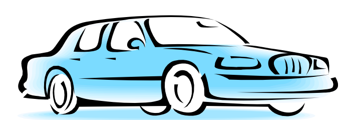

Localized spatio-temporal diffusion

Ensured by relevance-ranking and limited memory allocation

How fast/far a resource is disseminated?

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

- Memory allocation to the resource type
- Relevance threshold
- Transmission (randevous) range
- Traffic speed
- Vehicle density
- Resource density
- Average resource availability time

Other possible relevance functions

- Nonlinear
- Other factors
  - Travel Direction (gas station, malfunctioning brake-light)
  - *Transmit*-time, in addition to *create*-time (analogous to transaction/valid time)

Advertising spatial resources

- Gas stations, restaurants, ATM's, etc., announce continuously
- An announced resource item is acquired by the vehicles within the wireless coverage of the stationary site
- Different location-based-services paradigm than
  - Cellular-service provider database
  - Geographic web searching

Further research in data dissemination – mathematical model

- Spread resembles epidemiological models of (Bailey 75) but there are important differences
  - Spatio-temporal relevance function
  - Interaction of multiple infectious-diseases (resources)
- Should answer: given resource report generated at (0,0,0), what is the probability that a vehicle at (x,y,t) receives it

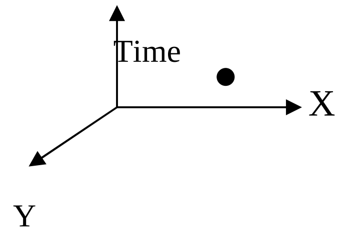

Further research in data acquisition(2)

- Data granularity/aggregation-level of sensor-data
  - Raw: multiple applications, more b/w
  - Abstractions/aggregations: less b/w, application specific
- Sensor fusion
  - fuse sensors of same kind from different vehicles
  - fuse different sensor-data, e.g. computer vision -- laser range-finding
- Resource-exchange modalities
  - Broadcast vs. 1:1
  - Push vs. pull

Research issues in data management

- Sensor data acquisition, fusion, dissemination
- Data usage strategies
- Participation incentives
- Remote Querying
- Data Integration, Moving Objects Databases

Another resource classification

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

Problem

- Information by itself is not sufficient to capture resource

- If move to obsolete resources may waste time compared to blind search

Strategies for capturing (semi-) competitive resources

- Example (Threshold Driven)– Go to the resource if its availability-report relevance is higher than a threshold *th*

- How much does TD save compared to Blind Search ?

Information Guided Resource Discovery

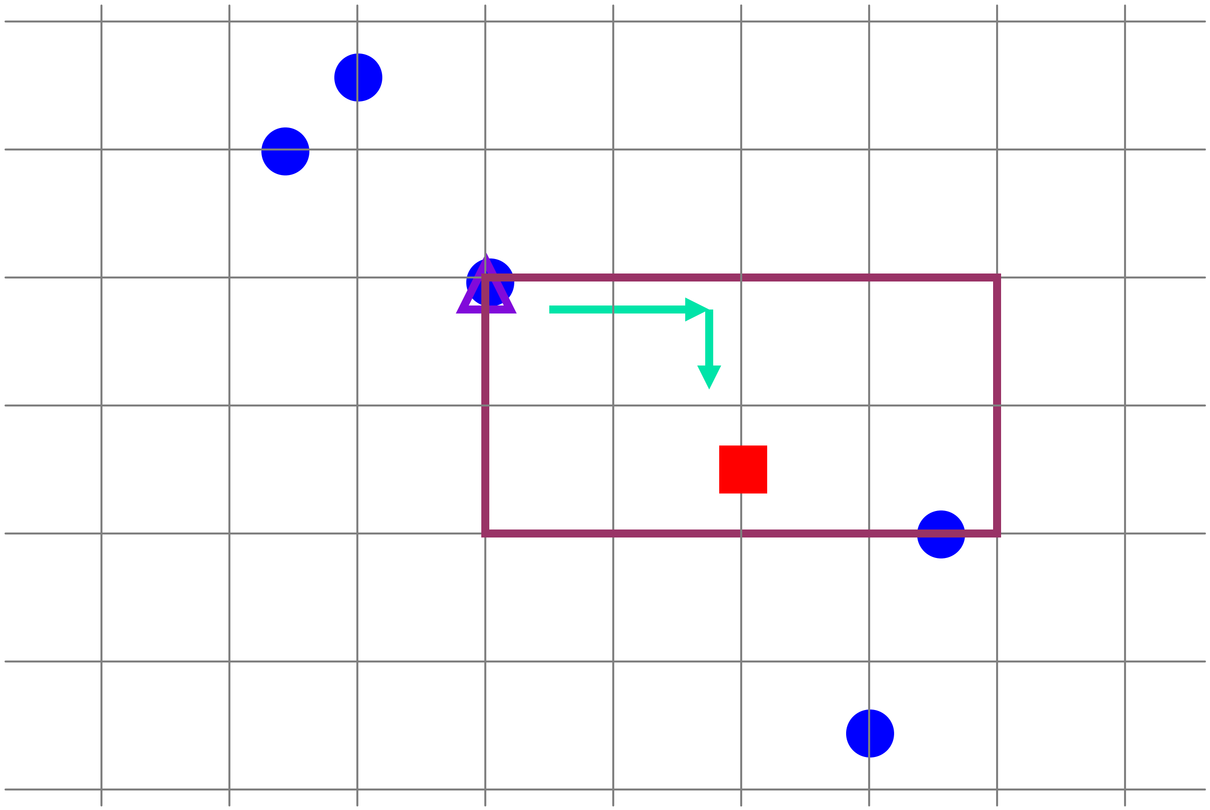

On average, TD captures the resource up to twice as fast as BS

Another strategy example

- Consider spatial-clustering of resources

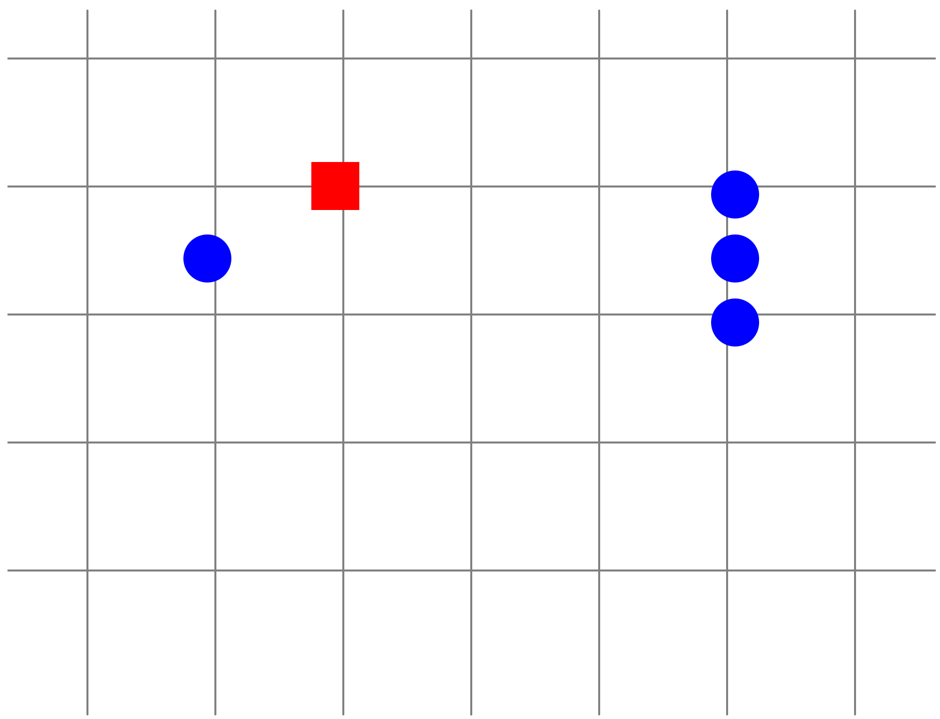

Further research in Spatio-temporal resource-capture strategies

- Develop and analyze information-guided spatio-temporal strategies (game theoretic approach?)
- How much does information improve resource utilization?
- Do invalidation messages help?
- If so, how should they be treated w.r.t. availability-reports?

Research issues in data management

- Sensor data acquisition, fusion, dissemination
- Data usage strategies
- Participation incentives
- Remote Querying
- Data Integration, Moving Objects Databases

Problem

- The mobile opportunistic p2p scheme heavily depends on wide participation

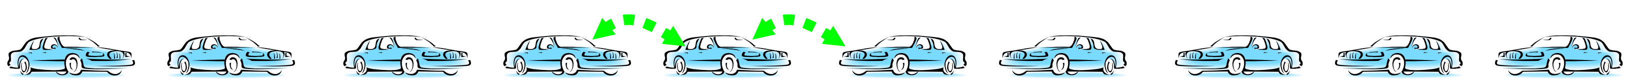

- Why should a vehicle/pda provide and transfer resources?

Possible incentive model

- Virtual currency -- tokens
- Producer-paid resources (road-emergency call, gas station)
  - Each report (advertisement) sent has a token-budget
  - On transfer between vehicles:
    - Carrier withdraws flat commission
    - Rest of budget split equally
- Consumer-paid resources (parking slots, cab customer, traffic-incident). 2 modes:
  - Consumer mode – pays amount proportional to relevance
  - Broker mode – cannot view resource, speculative
- Tamper-resistant security module
  - Stores resource-reports and tokens
  - Executes p2p protocol

Research in incentive models

- Other virtual currency models
- Pricing and negotiation
- Cost optimizations in such models
  - For example, minimize advertisement cost per potential customer
- Distributed reputation models
- Transactions and atomicity issues
- Security
  - eavesdropping
  - fake resources
  - tampering to gain unfair advantage, create havoc

Research issues in data management

- Sensor data acquisition, fusion, dissemination
- Data usage strategies
- Dissemination incentives
- Remote Querying
- Data integration, Moving Objects Databases

Spatio-temporal resource query modes

- Moving object queries local database
- **Moving object queries a region R, i.e. all the moving objects in R**

Examples and Issues

- Queries that find all the resources within a particular geographic area
  - find all the available parking spaces within the UIC eastern campus
  - find all the cab requests within five blocks of the Sears Tower
- How to determine the set of objects to which the query is sent?
- How to disseminate the query?
- How to collect the answers?

Determination of Query Destination Area – Possible answer

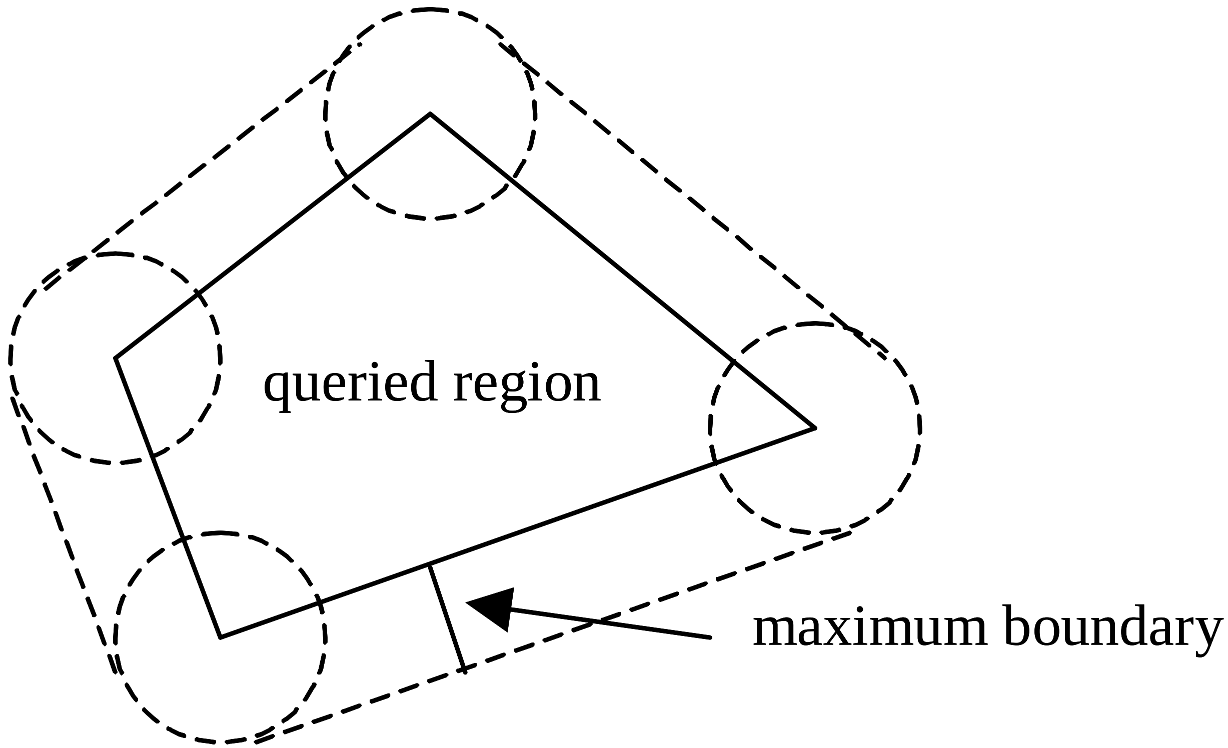

Remote Query Approach

- Query dissemination
  - Query originator sends the query into the destination area.
  - The query is flooded to all the moving objects within the area.
- Answer delivery
  - Each object in the destination area sends the answer back to the query originator
  - Query originator consolidates the answers.

How is query originator v found?

- Via the infrastructure using node-id
  - May be costly
- In p2p mode
  - v sends future trajectory in query

Two Cases

- Each object knows the trajectories of each other object
  - Trajectories exchanged as resources
- Each object does not know the trajectories of other objects except that of the querying object

Known Trajectories

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

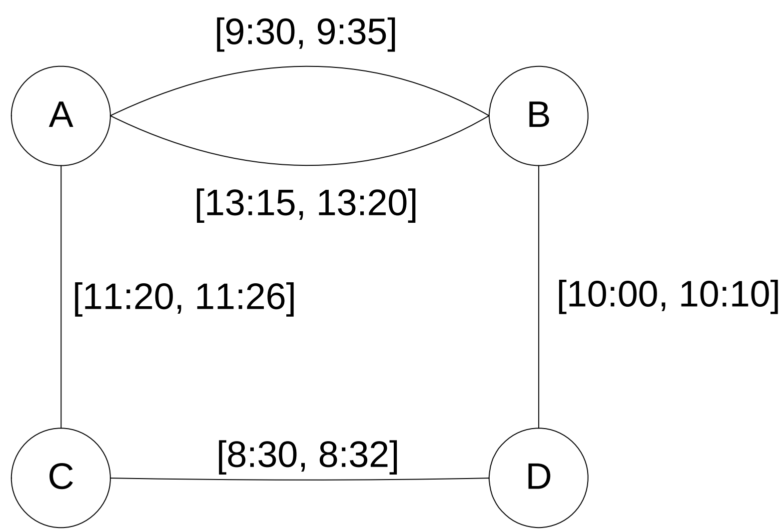

Known Trajectories

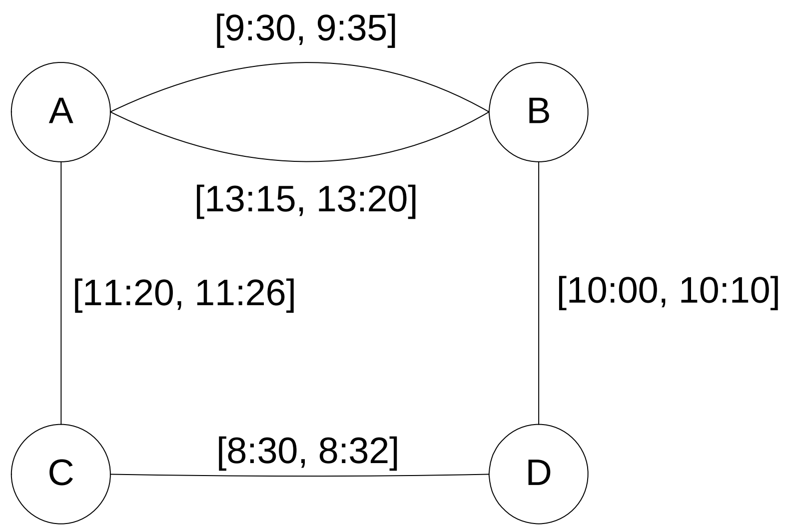

- A revised Djikstra algorithm is used to find
  - the shortest path between the querying moving object and the query destination area (for query dissemination)
  - The shortest path between an object in the query destination area and the querying moving object (for answer delivery)

Unknown Trajectories

- Question: How does a moving object decide whether or not to forward a message to its encountered neighbor?

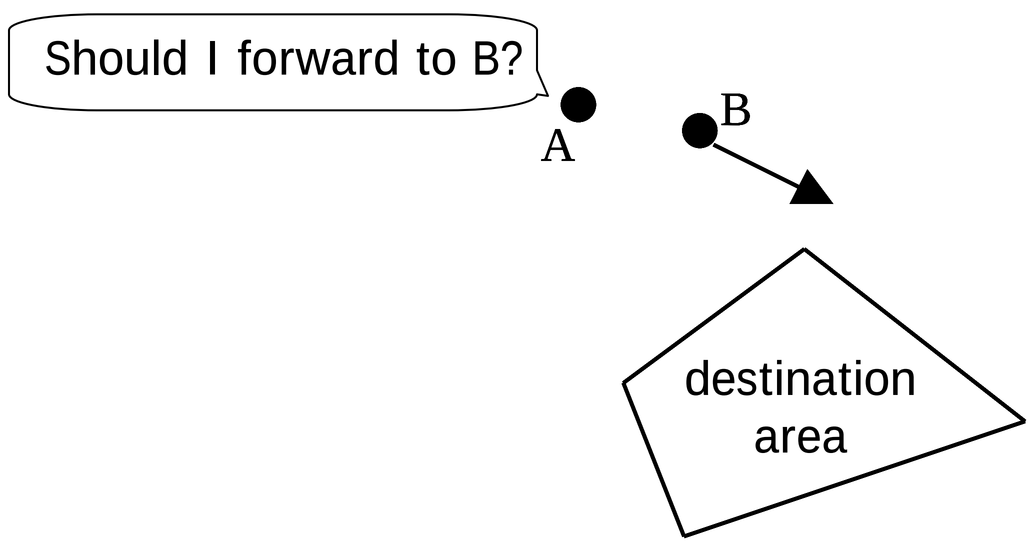

Unknown Trajectories

- Answer: Forward iff θ is smaller than a certain threshold (*critical angle*)

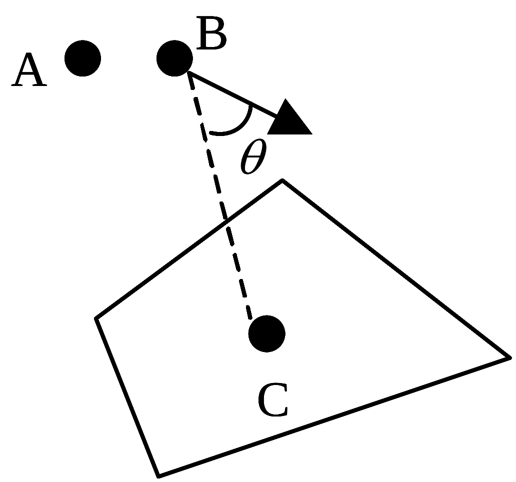

Choosing the Critical Angle

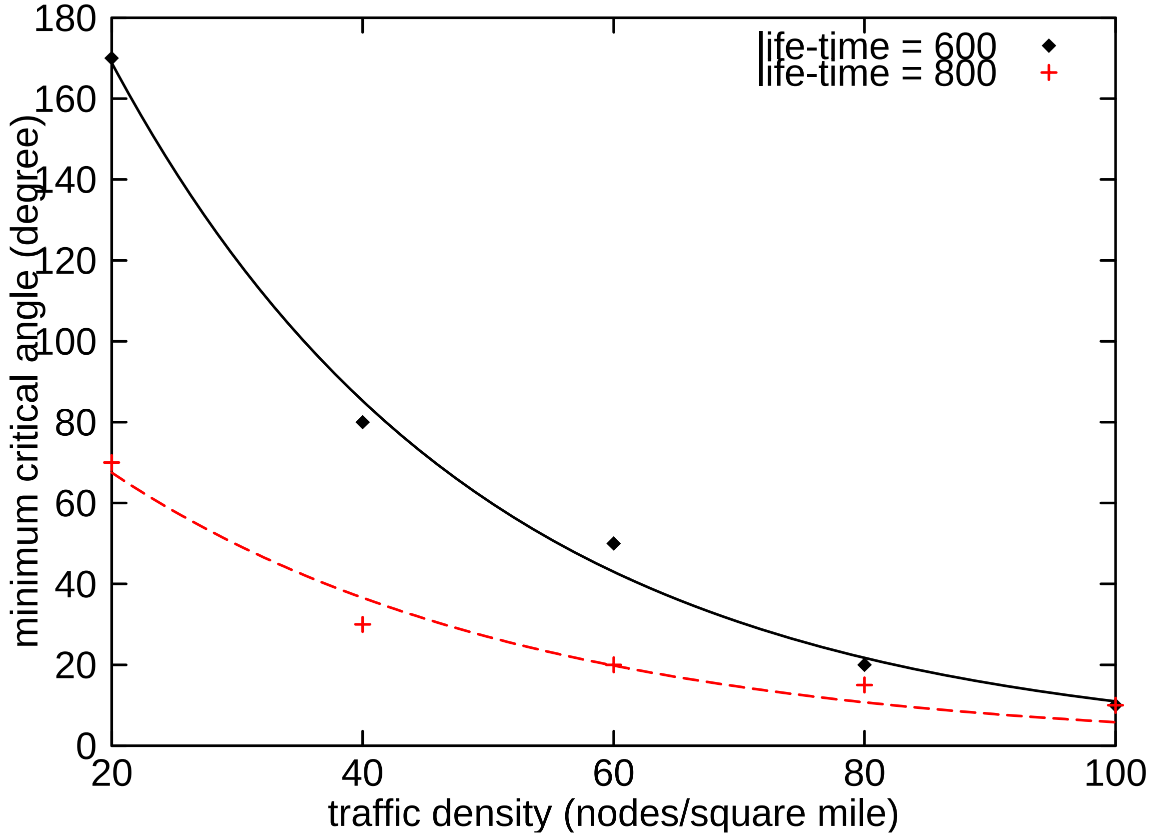

Query Processing Modes (1)

- Response to originator by each queried vehicle

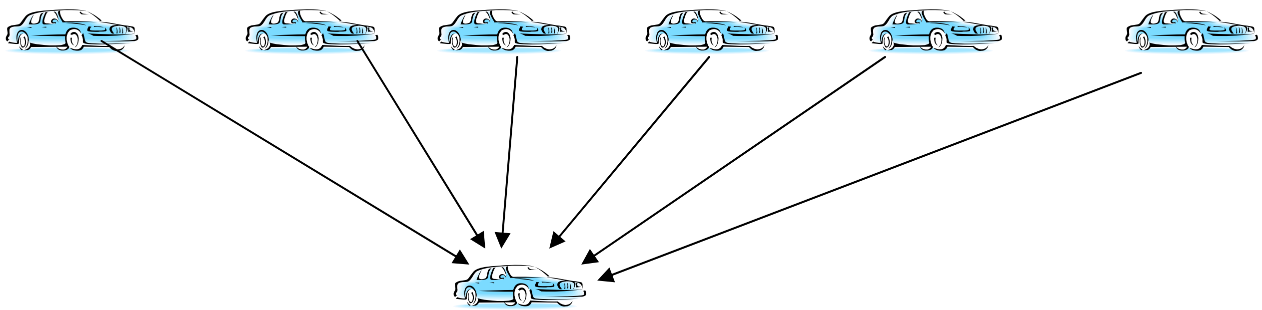

Query originator/
consolidates

Query Processing Modes (2)

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

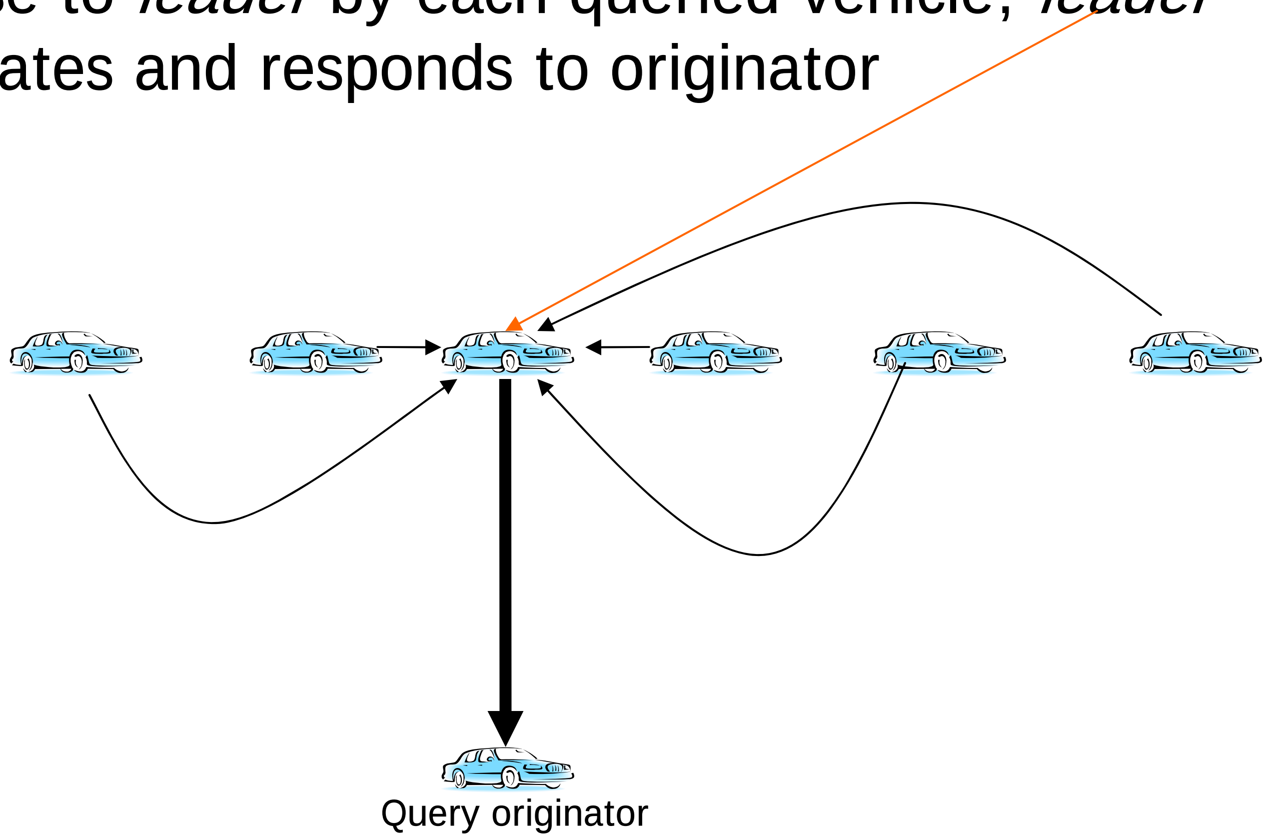

Query Processing Modes (3)

Hierarchical solution

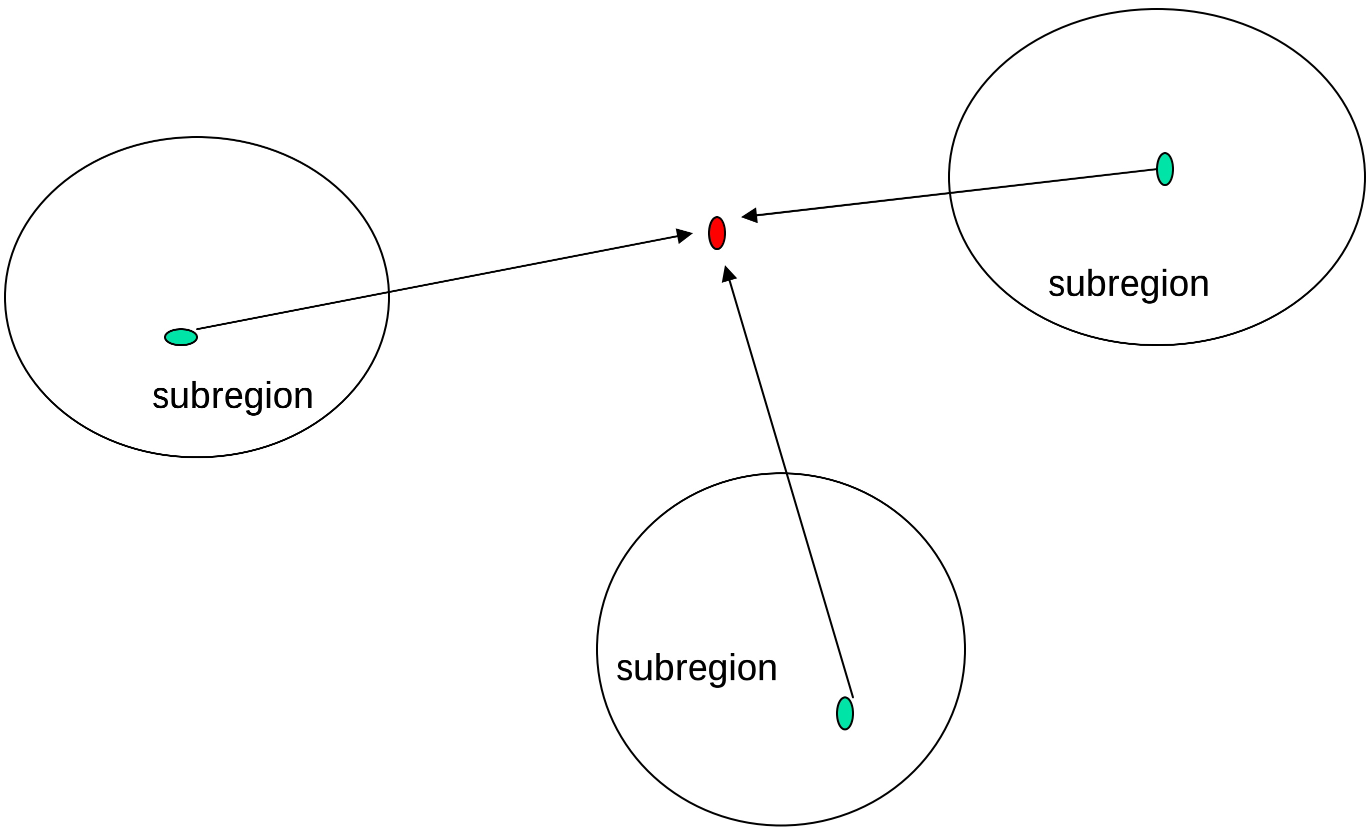

Further research in Remote Querying

- Comparison of query processing modes; coping with high mobility
- Other query types, aggregate/imprecise (average speed a mile ahead)
  - How to determine the set of objects to which the query is sent?
  - How to disseminate the query?
  - How to collect the answers?
- How/when to use cellular/infrastructure in communication of queries and answers?

Research issues in data management

- Sensor data acquisition, fusion, dissemination
- Data usage strategies
- Dissemination incentives
- Remote Querying
- Integration of sensor and higher level data

Moving Objects Database Technology

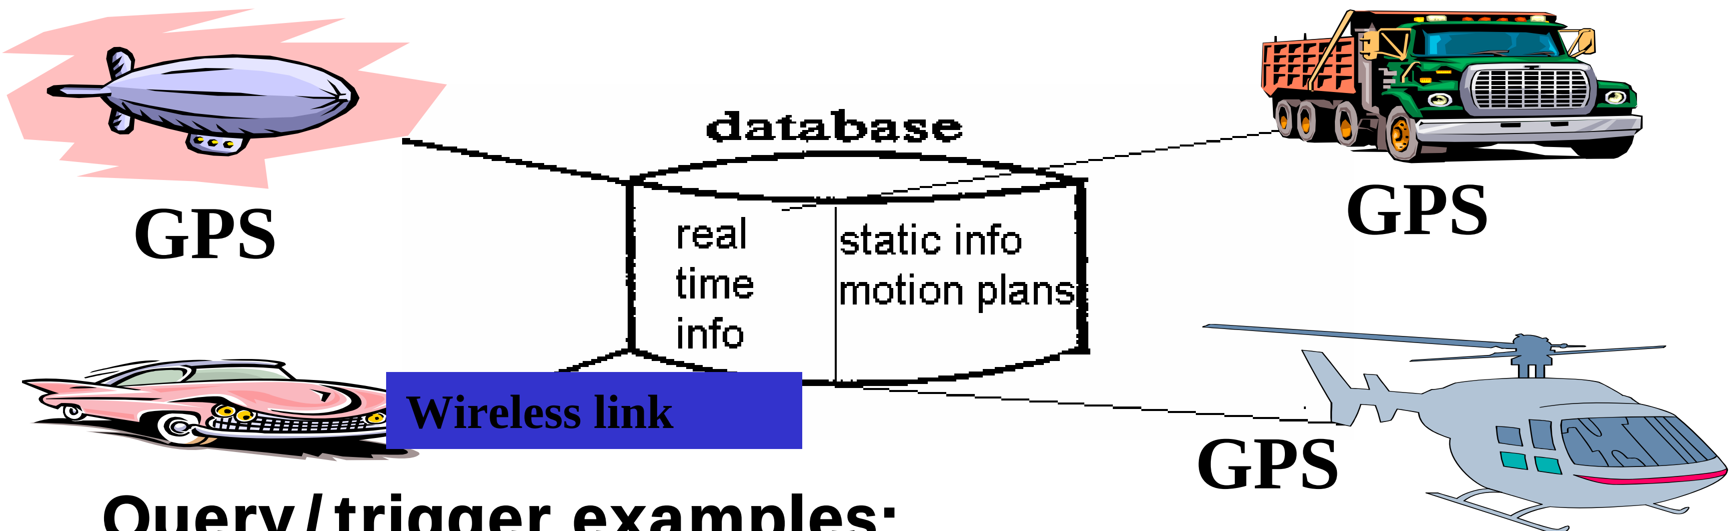

Query/trigger examples:

- 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)
- Taxi cabs within 1 mile of my location (present query)
- Average speed on highway, one mile ahead
- Tracking for “context awareness”

Applications

- Location Based Services e.g., “Closest gas station”
- Digital Battlefield
- Transportation (taxi, courier, emergency response, municipal transportation, traffic control)
- Supply Chain Management, logistics
- Context-awareness, augmented-reality, fly-through visualization
- Location- or Mobile-Ecommerce and Marketing
- Mobile workforce management
- Air traffic control (www.faa.gov/freeflight)
- Dynamic allocation of bandwidth in cellular network

Currently built in an ad hoc fashion

Further research in Moving Objects Databases

- Location modeling/management
- Linguistic issues
- Uncertainty/Imprecision
- Indexing
- Synthetic datasets
- Compression/data-reduction
- Joins and data mining (similarity of trajectories)

Relevant Work

- Resource discovering protocols
  - SLP, Jini, Salutation, UPnP
  - Rely on a dedicated directory server
  - Not suitable for high mobility environments
- Epidemic replication/routing (Demers 87, Vahdat 00, Khelil 02)
  - Regular data/messages, not spatial-temporal
- Sensor networks (Bonnet 00, Intanagonwiwat 00, Mandden 02)
  - Sensors are stationary
- Epidemiology (Bailey 75)

Conclusion

sensor-rich-environment + short-range wireless

Computer Science research issues:

- Sensor data acquisition/fusion/dissemination
- Data usage strategies
- Dissemination incentives
- Remote Querying
- Integration of sensor and higher level data

Future Work

- Privacy/security considerations
- Experiments based on a road network and Monarch/ns-2