CS 525M Mobile Computing

Emmanuel Agu

Database Management in Mobile Environments

- Last week, talked about applications:
  - wireless web, messaging, MPEG, application-aware adaptation
- Today focus on data management issues
- First, I'll motivate problem and recurrent themes, then 3 student presentations
- Summarize main ideas from paper:
- *Imielinksi and B. R. Badrinath "Mobile Computing: Challenges in Data Management," Communications of the ACM, October 1994"*
- Imielinski and Badrinath:
  - pioneers in mobile data research
  - Dataman group at Rutgers
  - 10-year best VLDB paper award

Data Management

- Envisioned user computing scenario: nomadic and mobile computing
- Actual mobility management will be a network level issue. E.g. mobile IP
- However, data-centric issues exist
- Mobile users will be both ***producers*** and ***consumers*** of data
- Some data will be location-dependent or time-sensitive: need update as user location changes
- Two main data access options:
  - Server broadcasts information ***(push)***
  - Users query for specific information ***(pull)***

Data Management

- Broadcast (push):
  - Air charged with “hot” (frequently requested) information
  - User PDA acquires, filters and presents required info
  - Issues: optimal broadcast patterns (metrics: latency, real time guarantees, etc), caching at client
- Searching (pull):
  - Small user groups share information e.g. salesmen
  - Issues: partitioning database, location-dependent queries
- Hybrid approaches: some push, some pull?

Data Management

- Mobility-related definitions:
  - **Location management:** moving users (e.g. handoff in cellular)
  - **Configuration management:** mobility of resources
- **Location management paradigms:** searching Vs. informing:
  - A wants to locate B
  - Should A search entire network for B?
  - Should A look only in predefined locations (e.g. using B's statistics and movement history)?
  - Should B inform A anybody of his or her moves?

Data Management

- Mobile user may be very far from server
- Servers may also move far away
- This leads to varying spatial distribution of clients
- Dynamic data replication required
- Basic question: should the data “follow” the user (static Vs. mobile)
- Other issues: disconnection, reintegration, consistency
- **Bandwidth constraints:** wireless links cannot sustain large amount of update messages
- **Energy constraints:** wireless terminals need to sleep to conserve energy, wake up at pre-defined times
- Today’s papers:
  - Broadcast disk: seminal work, multiple disk abstraction
  - Database survey: nice overview with valuable pointers
  - Data on Air: Indexing and structured data dissemination
- **Interesting aside:** Digital Fountain, company using broadcast disks *(courtesy of Mark Figura)*

Projects

- At this point, should be thinking of a few potential areas
- Suggestion on how to proceed:
  - Start with 1 paper in chosen section, which interests you
  - Make sure you understand the issues and specific problem being solved
  - Look at references, especially “Related Work” section
  - Chase up and skim papers (abstract, intro, results) in related work section
    - Start a bibliography
    - Take note of experimental tools
  - Resources to search:
    - ACM digital library: www.acm.org/dl
    - Citeseer: http://citeseer.nj.nec.com/cs
    - IEEE Xplor: ieeexplore.ieee.org/
  - **Good indicator:** Can I reproduce paper experimental setup and results?