

Distributed Systems

Few current developments, ethical
and social issues

Rik Sarkar

University of Edinburgh
Fall 2015

More advanced things

- With advancing technology
- New features are cropping up in mobile nets
 - And distributed systems
- New issues are appearing

Localization in wireless networks

- Can be done based on signal strengths
 - Decreases with distance
- Trilateration – three distances suffice to locate a point
 - Signal strength from three or more wireless transmitters with known location
 - Approximate localization due to variation in signal propagation, noise etc

Location in mobile nets

- Can be done using three cellular towers
 - Either at the mobile phone
 - Or at cell network
- New possibilities
 - Predict handoffs from motion, make preparations early
 - Learn/predict user behavior
 - Give location based recommendations etc...

Privacy issues

- Location is considered private information
- Tracking location all the time is considered intrusive
- No way to prevent cell companies from doing it...
- At the mobile phone end, efforts are being made to fuzz locations
 - Give some information for location based services to run
 - Avoid some other information, or avoid precise locations
- Context/activity information is also private...

Femto-cell networks

- Small personalized cell antennas (~10m range)
 - Plug into your wireless router
- Even better spatial reuse
- Less organized, trickier to coordinate channel assignment etc
- Possibility of channel clash with macro cells

Beamforming directed communication

- Sends signals in a particular direction
 - Uses multiple antennas together transmitting at different phases
 - Destructive interference ensures that signal does not travel in other directions
 - Easy to change directions quickly
- Angle can range from 20 or 30 degrees to a few degrees (narrow beam)

Beamforming directed communication

- Usually around 60GHz frequency
- High supported bandwidth
- Good for hi def video, large data volumes etc
- At this range, signal has high attenuation from air
 - Short range (few meters to 100 meters)
- Together with directionality implies very effective spatial reuse

Beamforming directed communication

- Challenges
 - Easily affected by obstacles
 - Both transmitter and receiver need to be “looking” at each-other at the right time
 - Medium access becomes harder
 - Mobility can create a challenge of “tracking” a device

Mobile computing

- Mobile wireless devices are going to be even more popular
- More apps/services/media
- Great need for wireless bandwidth
 - Current infrastructure is not sufficient
- New technologies
- Simultaneously, detection of collective context, groups etc for better adaptive services

Privacy, policy and ethical issues

- Tracking location and activities of people
- Mobile health
 - Enhancing healthcare with mobile apps
 - Potentially useful, but has privacy issues
- Distributed storage and services
 - Easy to lose privacy (you don't know who is holding your data, who has access to it)
- P2P systems : Efficient distributed distribution
 - Continuously faces legal (piracy) issues

Legal issues

- Globalization: Extreme distribution of computing
- You have a UK company
- Your server is in china
- Your customer is in US

- What is legal and what is illegal?
- Which set of laws apply?

- In general, legal and administrative systems are far behind in adapting to distributed globalized computing

Other ethical issues: Energy and heat

- Large cloud data centers generate huge amounts of heat
 - Heat management/cooling a research topic on its own
 - Problems with climate change etc..
 - Facebook locating servers in the arctic circle
- Overall many computers, mobile, cell towers
- Consume huge amounts of energy

Cluster computing

- Google operates on large number of ordinary computers (instead of few powerful servers)
 - Rise of “cluster” computing
- When the dataset is large, one computer of any size does not suffice
 - Many processors/computers are necessary
 - Better off using distributed computing
- Better redundancy/fault tolerance

Distributed and cluster computing

- Problem:
 - Many of the tasks are hard to distribute
 - Standard algorithms assume working on a single computer
 - E.g. Optimization (machine learning), Analytics, data mining etc
 - How do you split the task to operate on many machines?

Other things?

Course

- What topics did you like?
- What did you not like?
- What else could have been done?