Tongji: Mobile Computing Systems
Sample Solution to Final Exam, Summer 2007

Question 1. Data in Cellular Networks (10 marks)

Outline the evolution of data services in cellular networks, from AMPS to 3rd generation cellular systems and beyond.

Answer (10 marks):
First generation cellular systems such as AMPS were designed for voice communication only. Data services, at best, could be added by using modems. In some countries, including USA and Canada, CDPD was developed and deployed as an overlay network to AMPS to provide a packet data service. Achievable bandwidths in either case were low, somewhere around 10 kbps.

Second generation cellular systems such as GSM were designed to provide voice and circuit-switched data services at various data rates (up to 9.6 kbps originally, later extended to 56 kbps). Subsequent evolutions of these networks added even higher rate circuit-switched data services and packet data services (GPRS).

Third generation systems are designed from the beginning with the goal of providing high-rate data services (up to 2 Mbps), as future revenues are expected to be driven mostly by data services. While 3rd generation cellular systems such as UMTS provide all services (voice, data, SMS, etc.) over a common, standard air interface, 4th generation (or beyond 3rd generation) systems envision networks where mobile users will have access to a range of wireless access technologies with widely differing coverage, mobility support, and data rates. In addition, these networks are expected to move to all-IP networks, so that even voice will/may be offered by essentially a data network (VoIP), though that will only happen if IP networks provide much better resource management and QoS support.

Question 2. Bluetooth (10 marks)

One common problem in mobile computing is that devices are battery-operated and therefore very power-constrained. Explain all the ways in which energy conservation is taken into account in Bluetooth's design. Does Bluetooth's energy conservation mechanism differ from what IEEE 802.11 offers?

Answer (10 marks):

For starters, Bluetooth was designed to be a short-range transmission technology only (up to 10 meters), which significantly reduces the transmission power requirements. In addition, great care is taken in the MAC protocol design to allow nodes to be in a low-power standby mode as much as possible, only waking up very sporadically (every 1.28 seconds) to check whether someone wants to connect to them. And even when a device participates in a piconet, it can switch into one of multiple low-power sleep modes instead of staying in the more energy-consuming connect state. Finally, the MAC protocol allows the piconet Master to schedule transmission of packets by piconet slaves, avoiding collisions and costly retransmissions.

IEEE 802.11 was designed for a much longer communication distance, with significantly higher transmission power. There is no protocol in place to have devices being paged periodically to allow them to switch into a low-power sleep/standby mode while nobody wants to communicate with them (as IEEE 802.11 is designed to packet data networks, there is no notion of being “connected”, similar to Bluetooth's notion of being a member of a piconet). While the infrastructure deployment allows the AP to schedule transmission by other nodes, this mode (CFP) is rarely implemented and used. IEEE 802.11 does provide for sleep mode operation to conserve power, but it does not scale well in ad-hoc mode.

Question 3. Mobile IP (10 marks)

Describe Route Optimization in Mobile IP for IPv4. Explain the problem with routing in the standard Mobile IP solution, sketch the operation of the route optimization protocol, and list potential problems with the route optimization solution discussed in class.

Answer (10 marks):

The original Mobile IP for IPv4 solution suffers from triangular routing: while data from the MN to a CN is routed on the most direct path (according to the deployed Internet routing protocol), data from a CN is first routed to a HA and then from the HA, via the FA, to the MN.

Route Optimization Solution: update the CN with the current "mobility binding", i.e., the CoA (care-of-address or dynamic, topologically correct IP address). This allows the CN to send the data packets (via a tunnel) directly to the MN (via the FA). CN aquires CoA from the HA when sending first packet to MN using the static IP address. When MN moves, it will set up forwarding pointer from old FA to new FA to allow packets sent to the old CoA to be forwarded to the new location. Packets arriving at the oldFA from a CN will cause the old FA to send an invalidation message to the CN. CN can then either continue by using the static IP address of MN or inquire about new CoA from MN's HA. The following message sequence chart summarizes the relevant protocol operations:

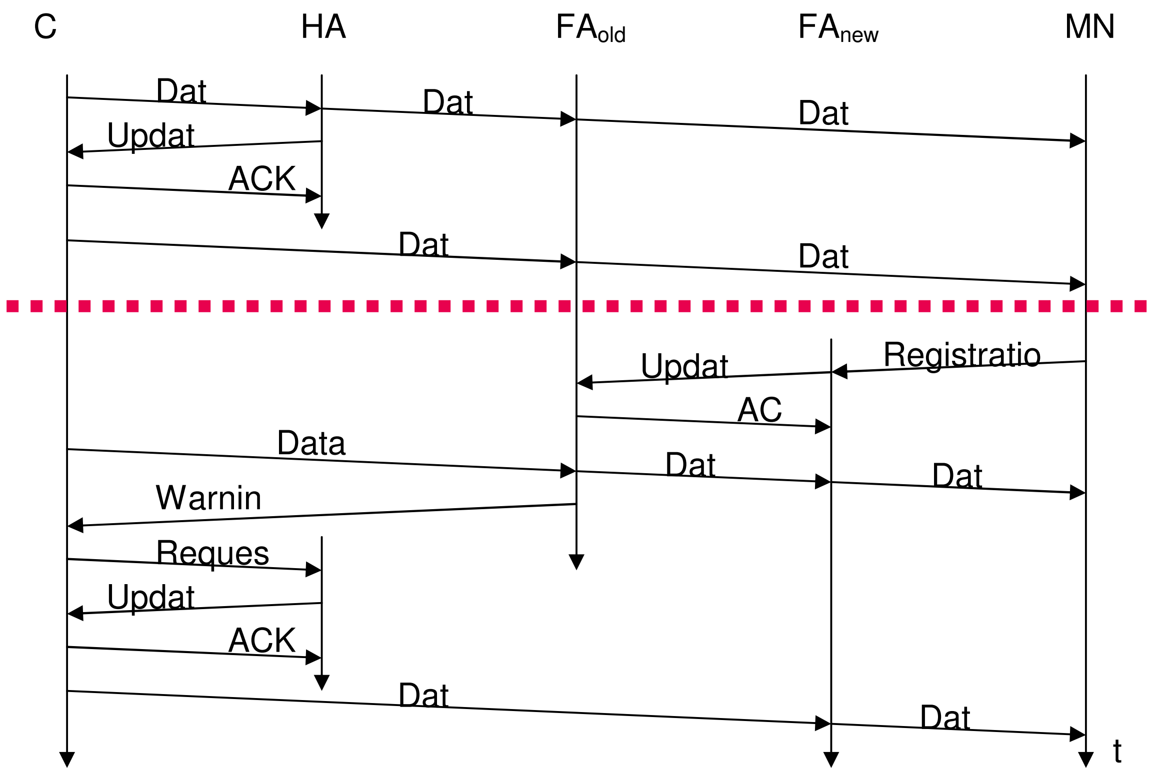

Problems:

1) CN needs to understand new protocol messages (binding updates/invalidations)
2) CN need to manage cache of binding updates
3) CN needs to be able to build a tunnel
4) Security: how does CN trust binding updates it receives from a node claiming to be a HA

Question 4. MANETs (10 marks)

Describe the various categories of MANET routing protocols, as standardized by the IETF. Explain the relative strengths and weaknesses of each category.

Proactive, Table-driven Approach

- Based on traditional link-state and distance-vector routing protocols.
- Continuously update the topological view of the network by periodically exchanging appropriate information among the nodes.
- Determine routes independent of traffic pattern
- Examples: DSDV, OLSR (Optimized Link State), TBRPF etc.

Reactive, On-demand Approach

- Discover and maintain routes only if needed
- Do not continuously maintain the overall network topology
- The network is flooded with “route request” control packets when a new route is required.
- Examples: DSR, AODV, LAR, etc.

Proactive protocols: lower route acquisition latency, higher protocol overhead (typically)
On-demand protocols: higher route acquisition latency, lower protocol overhead (typically)

Question 5. TCP over Wireless Links (10 marks)

As discussed in class, TCP performs typically poorly in networks where packet loss is not due to congestion. In a mobile environment, where packets could get lost due to poor wireless link quality, handoffs, and route failures in a MANET, TCP needs to be complemented with additional mechanisms, as discussed in the course. Summarize the strength and weaknesses of link-layer retransmission schemes.

Answer (10 marks):

When is a reliable link layer beneficial to TCP performance?

- if it provides *almost in-order* delivery
- TCP retransmission timeout large enough to tolerate additional delays due to link level retransmits

Problems:

- Not all connections benefit from retransmissions or ordered delivery
    - audio
- Need to be able to specify requirements on a per-packet basis
    - Should the packet be retransmitted? How many times?
    - Enforce in-order delivery?
- Even if we are successful in retransmitting a packet before TCP sender times out, we will increase RTT for this packet, increasing the RTO value at the sender and therefore decreasing the chance that the sender will react to congestion in the fixed network early on.