CSCI 4900/6900 Ubiquitous Computing

Mid term exam

March 1, 2001

Rules

- Open book, open notes, individual effort
- **USE OF WIRELESS OR OTHER ELECTRONIC TECHNOLOGIES TO SHARE YOUR ANSWERS IS STRICTLY PROHIBITED**
- Provide clear, thorough and concise answers. **Explain your assumptions (if any) briefly**.
- Answers without brief explanations will not receive any credit (for example, a vague and general answer such as "Use a cache" will not receive any credit, unless you explain specifically how you use the cache to solve the problem)
- Throwing a bunch of terms that we used in this course (e.g. anti-entropy direct mail peer-to-peer central server) will not get any credit. You have to specifically answer the question asked.
- Grading policy The number of questions that you need to answer depends on the level that you have registered for this course. You are free to attempt more than the required number of questions. I will choose the best (2 or 3 or 4) of your answers.

4900 **Students taking this course for 4900 credit should answer either Part 1 or Part 2 completely.**

6900 **Students taking this course for 6900 credit should answer any three of the four sub-parts (1.a, 1.b, 2.a and 2.b).**

8xxx **Students taking this course for 8000 level credit will answer both the parts completely (including all four sub parts).**

1 Part 1

Napster is a popular music exchange service that uses a central server approach to hold information about music files available on users desktops. The actual music transfer happens between the desktops without involving the Napster central service.

This model is being contested by the copyright holders through the court system. Suppose the Chief Technical Officer (CTO) of Napster reachs an agreement with the music labels to pay them for using copyrighted materials. Since the actual file transfers happen between the clients, Napster offers the following model to charge (and reward) the end users.

- Each song downloaded through Napster would cost $x.
- If some user downloads information from your desktop, Napster will pay you 0.1*$x for providing the storage and network resources.

Question 1.a

Since it pays to serve music to others, a class of entrepreneurs (referred as NAPERS) evolve who make money using Napster to serve music for others. Napster still expects these NAPERS to pay for the first copy. How would you, as a NAPER, develop a peer-to-peer approach , that works on top of Napster, to disseminate information about how often a song is being requested to Napster? This information can be collated by the individual machines to create the “most popular music” list. Fellow NAPERS will use this list to download popular songs (with the expectation that enough users will download the song from you to offset the one time cost to Napster).

Note: You can assume that your fellow NAPERS are not malicious, i.e. they do not intentionally lie about a song’s popularity to make you go broke!!

Question 1.b

How would you make sure the users get the right amount of credit? Since the users desktops are unreliable and may be disconnected for long periods of time, how do you make sure that the credits for serving music and costs for downloading music are propagated in a timely and reliable fashion?

2 Part 2

Package delivery services such as FedEx and UPS offer a number of mechanisms for their customers to track their packages. Users can access the tracking information on their web site, using custom software, using telephones etc. None of these services allow the customer any control over the routing to the last mile. For example, if a customer sees that a package is being delivered to their home and they are at work; they cannot ask for the package to be held in the depot so that they could pick it up. The customer has to wait for the package delivery to be attempted, the package brought back the depot before they can change the routing mechanisms.

Suppose you are the Chief Technical Officer (CTO) for UPS. You want to provide a service that would allow your customers to alter the routing dynamically. In the above scenario, you want the customers to be able to re-route the packet so as to deliver the package at work (instead of home). You want the trucks to exchange items on the road so that the package can be moved from one truck to the other without passing through the depot. Each of the trucks will be fitted with wireless access points with a range of 10 miles so that they can talk to each other while on the road.

Question 2.a

In this UPS scenario, how would you disseminate new package routing information to the UPS trucks? (remember, the trucks only carry a wireless access point with a range of 10 miles, enough to talk to other trucks but not necessarily back to the depot)

Question 2.b

You decide to charge the customer a fixed cost of $x for a routing change request and a cost of $y per package if a package was actually re-routed. Your truck computing devices are unreliable to hold such cost information. How do you reliably make sure that the customer is charged the correct fee? (not less and not more)