CS 422 (Park) **Assignment II** **Due: Feb. 18 (Wed.), 2004**

Submission instructions: Please type your answers and submit electronic copies using `turnin` *by 11:59pm on the due date. You may use any number of word processing software (e.g., Framemaker, Word, LaTeX), but the final output must be in pdf or ps format that uses standard fonts (a practical test is to check if the pdf/ps file prints on a CS Department printer without problem). For experiments and programming assignments that involve output to terminal, please use* `script` *to record the output and submit the output file. Use* `gnuplot` *to plot graphs. Use* `ps2gif` *to convert a eps/ps plot to gif format (e.g., for inclusion in Word).*

PROBLEM 1 (20 pts)

Read Chapters 5, 6, and 7 from Comer. Solve Problems 4.7, 5.4, 6.9, and 7.1.

PROBLEM 2 (20 pts)

Assume you are the new kid in the digital radio market. Suppose you have permission from the FCC to transmit signals at exactly 100 GHz. Being averse to multiple sine waves (you can tolerate at most one), you have decided to use TDM to broadcast 10,000 channels of digital radio (a sure-fire way to go bankrupt but never mind that). Each channel carries super-quality audio in the 0–20,000 Hz range, whose magnitude is quantized using a somewhat excessive 64 bits.

• Design a frame format that accommodates all 10,000 channels and their audio spec. Per frame, use 8 bits as start bits for synchronization. What is the bit rate (bps) of your design?

• You are launching a single GEO satellite orbiting at ~22,200 miles to broadcast bits at 100 GHz that will cover West Lafayette and surrounding states. Since you are averse to using FDM and have opted to use TDM instead, assuming you are employing AM with two amplitude levels to represent 0 or 1, what is the maximum bit rate you can achieve using your single 100 GHz sine wave? How much of the total bandwidth are you utilizing when broadcasting your TDM frame carrying 10,000 audio channels?

• Suppose that when it rains the signal reception from your satellite degrades—the amplitude of your 100 GHz sine wave gets mangled—so that for every 3 consecutive bits there is likely to be a bit that has flipped (0 is misread as 1, or vice versa). Given that you have some idle bandwidth available, how can you use it to correct these 1-in-3 bit flips so that your customer receives crystal clear reception rain or shine? Extend your frame specification to incorporate this added error correcting feature. Calculate the new bit rate of the extended frame design.

PROBLEM 3 (30 pts)

After operating the digital radio network for a couple of years, suppose the inevitable happens and you go broke. Having learned a hard lesson—I should not be so averse to FDM—you pick up the pieces and relaunch the same system, with slight modification, as a satellite-based Internet access provider competing with telcos and wireline cable providers. Your slogan is "100 Mbps to every home in WL and surrounding states" at a flat subscription fee of \$1.25/mo. Being stubborn you decide to stick with TDM (same old AM with two levels at frequency 100 GHz).

• Up to how many users max can you accommodate over your satellite system while guaranteeing 100 Mbps access? You may ignore IP encapsulation issues. What if you consider your own access to the Internet? As an access provider, you are only a conduit for your customers unless you have cached all the information that your customers will ever need on board the satellite's disk. Assuming 1 Mbps access from a tier-1 ISP costs you $\$X$/mo, but cost is not a concern and you are guaranteeing a 100 Mbps private tunnel to every customer to the Internet through your tier-1 ISP, how many customers max can you accommodate over the satellite TDM system? To break even—revenue from your customers at \$1.25 for 100 Mbps equals cost paid to your tier-1 ISP (ignoring all other costs like operating the satellite, electricity, staff, etc.)—what must your cost $\$X$/mo that you negotiate with your tier-1 provider be? Considering \$50-per-Mbps is the minimum going price these days for "bottom feeder ISPs," with tier-1 ISPs charging \$100+, is your \$1.25 price economically viable? At rock bottom \$50-per-Mbps cost, what price must you charge your customers to break even? Is this a cost that a typical consumer can afford?

• Does increasing the number of magnitude levels used in AM from 2 to, say, 8, change anything? Explain. What if you were to convince your tier-1 ISP to charge you nothing for traffic coming into your satellite from their wireline network for paying more, say, $10,000-per-1 Mbps for traffic going out to their network from your satellite network? Should you go for it if you find an agreeable tier-1 ISP (ignoring all other cost factors)? What may still cause you pause in today's post-Napster Internet environment?

• Before launching the new Internet access service, you have come to your senses and determined that it's a short-cut to bankruptcy. Instead, you contemplate the recent popularity of Voice-over-IP (VoIP) where digitized voice is carried over packet networks, including the commodity Internet (again in the news due to Skype, a creation of the Kazaa duo, which facilitates VoIP). Assuming the standard 4 KHz bandwidth for voice with 8-bit quantization where each 8-bit voice sample is encapsulated as an IP packet that adds 20 bytes due to the IP header, how many simultaneous calls can you support over your TDM satellite system? Show your frame design and operational specification (i.e., how it is to be used). Make sure to consider the fact that telephone calls are two-way, i.e., full-duplex. Assuming a $50-per-Mbps cost for your own access to the Internet via a tier-2 provider, how much can you charge your customers to break even while guaranteeing continuous VoIP telephone call access to all (in the worst case, every customer may continue his/her call, without sleep, for a whole month but your system still has to work)? Is this cost schedule economically viable? Note, in practice, compression can reduce the data rate required for voice (half-duplex) to around 10 kbps.

PROBLEM 4 (20 pts)

Considering the VoIP service that uses 100 GHz satellite TDM, we realize that not all users will be making telephone calls all the time. Hence the calculations carried out in Problem 3 were under a worst-case scenario, necessitated by the notion of *guaranteeing* a 24×7 service to each customer. When a telephone call by a customer is not being made, the same customer may use his/her allotted slot(s) to send data traffic, for example, when browsing the Web. In a similar vein, the fact that not all users will be using the network at the same time can be exploited to admit more customers than the worst-case calculation in Problem 3 allows—a principle used by banks to keep only a fraction of their deposits at hand (if all bank customers were to withdraw their savings at the same time, the banking system would be in trouble). Assuming a data packet can be up to 64 kB bytes long, how can "full-fledged" data networking coexist with the IP-over-TDM frame for voice in Problem 3? Propose a design that would allow variable length data packets to be sent over the TDM VoIP system advanced in Problem 3. Discuss what the key issues are, how your approach overcomes them, and what the pros/cons of your solution are.

PROBLEM 5 (20 pts)

TDM, by definition, dictates that time slices or slots be allocated to each user in a round-robin, periodic manner once-and-for-all. If a user does not use his/her slot at a given moment, it goes to waste. Can you advance an "on-demand" TDM system that, while preserving the notion of fixed-size slots, allocates each slot in an on-demand fashion in real-time? Ground your solution by implementing it as an extension of the TDM design in Problem 4. What are the pros/cons of your on-demand TDM (sometimes called STDM for statistical TDM) when compared to the plain old (deterministic) TDM?

PROBLEM 6 (30 pts)

(a) You will find a client/server application under ~park/pub/cs422; the client program is `client.c` and the server program is `server.c`. Reverse-engineer the code and explain what the server and client are doing. Give a line-by-line explanation of the code—you may group several lines into one unit, as you see fit, and give a modular description. The important criterion is that your answer makes clear that you understand all aspects of the code. If you find inefficiencies, indicate what and why, and suggest possible changes. Include in your answer a clear explanation of what UNIX system programming trick is being used to direct the output of the server back to the client and how this works.

(b) The client/server code will not compile as is. Specify what changes need to be made, institute them, compile the modified program and test the client/server application by first running the server process `server.bin` in the background followed by the client process `client.bin` in the foreground. Use the UNIX `script` facility to record the run-time session. Submit a copy of the session script along with your write-up.