

Bolo – A Simple Audioconference

Project 2

CS525u Multimedia Computing

Due date:



Overview

- Audioconference
 - Been able to do this (well) for at least 10 years
 - Web and Internet made Internet telephony possible, therefore now popular
- Basic
 - Two-person
 - System parameters to evaluate quality
 - Speech detection
 - + Simpler version than Project 1
 - Minimal interface
 - + Load time options for connection, parameters ...



Details

- Linux
 - Two Linux boxes, actually
- Internet sockets
 - Specify host (and port)
 - TCP and UDP
- Variable sample sizes
 - 20, 40, 60ms ... up to 1000ms



More Details

- Basic speech detection
 - Modified from project 1
 - 250ms search for zero crossing too big!
 - Detect based on energy level only
 - (Samples provided)
 - Can have on or off
- Packet "Loss"
 - Drop before/after sending
 - Percent, randomly distributed



Hints

- Many different architectural possibilities
 - Can design how you want
- Provide sample code ("tools")
- Remote development (not at console)
 - "Pre-Record" some conversation
 - Read from file or device
 - Write to file or device
 - Allow development of system code
 - Also good for one person testing



Sample Code

- Speech detection: [computeEnergy.c](#) and [getThresh.c](#)
- Basic TCP sockets: [talk-tcp.c](#) and [listen-tcp.c](#)
- Basic UDP sockets: [talk-udp.c](#) and [listen-udp.c](#)
- Setting a timer: [setitimer.c](#)
- POSIX threads: [add2.c](#)
- Multiple interrupts: [select.c](#)
- Parsing command line parameters: [get-opt.c](#)



Evaluation

- “Pilot studies” with group or a friend
- Use Bolo, provide PQ scores (1-10)
- Cases (done individually, not in combo):
 - Loss: 0%, 1%, 5%, 10%, 20%
 - Latency (by increasing the sample size): 40ms, 100ms, 250ms, 500ms, 1000ms
 - Connection type: TCP, UDP
 - Speech detection: on, off
- “User study” with 1-3 people outside group
- Use Bolo, provide PQ scores plus comments
- Loss: 0%, 5%, 20%



Questions

1. What is a major detriment of using a TCP connection on a lossy network without speech detection?
2. Does having a TCP vs. UDP connection matter on a LAN? Why or why not?
3. Would a 480ms end-to-end network latency with a 20ms sample rate have the same impact on perceptual quality as a 500ms sample rate on a LAN? Why or why not?



Hand-In

- Email a tar ball
- Include
 - All source code and Makefile
 - README file with instructions on running
 - Tables or graphs of the PQ scores and comments
 - Answers to the questions
 - Group information

