status report for

A Sun Sparcstation-Based Speech Data Collection Platform

LDC Subagreement 5-24431-C
ISIP Project No. 02-95

for the period of May 15, 1995 to June 30, 1995

submitted to:

Linguistic Data Consortium

441 Williams Hall
University of Pennsylvania
Philadelphia, PA 19104-6305

submitted by:

Joseph Picone, Ph.D., Associate Professor

Institute for Signal and Information Processing
Department of Electrical and Computer Engineering
Mississippi State University
Box 9571
216 Simrall, Hardy Rd.
Mississippi State, Mississippi 39762
Tel: 601-325-3149
Fax: 601-325-3149
email: picone@isip.msstate.edu

EXECUTIVE SUMMARY

This portion of the contract was largely dedicated to the development of infrastructure required for the project. All computer hardware related to the project has been purchased, delivered, and installed, except for the T1 telephone lines and interface board. The equipment includes two Sun Sparcstation 5's, one two-processor Sparcstation 20, 30 Gbytes of disk, an 8 mm tape library, and several NCD X terminals. The hardware was configured to run Sun's newest operating system, Solaris 2.4. After resolving several hardware and software issues, the computer hardware has been verified to be stable and robust.

A pure Sun Solaris 2.4 environment was developed to support the contract. We miscalculated the amount of effort required to install this operating system. Adjustments in staffing, at no cost to the contract, were made so that this phase of the project could be brought to a timely completion. The environment is based on MIT X and the Free Software Foundation's (FSF) GNU software, and is built on top of several other key pieces of free software that were used to replace defective software from Sun. Most of the shareware acquired was initially found to have major problems under Solaris 2.4, and required significant effort to install. The bottom line is that Solaris 2.4 is significantly different from its predecessor, Sun OS 4.1.3 and Solaris 2.3, and that most software providers have yet released Solaris 2.4-compatible software.

For these reasons, we wanted to delay the acquisition of the T1 hardware and software. We wanted the vendor, Linkon, to have as much time as possible to debug the system. We also did not want to waste money on a T1 line while we could not use it. Now that our environment is stable, we will begin the software evaluation and planning stages of the project. The first step will be to investigate Sun's telephone services API (XTL), and understand its architecture and functionality. Since the Linkon system claims to support XTL, this step is important in that it will give us one method of evaluating and debugging the Linkon hardware (we anticipate problems). Hopefully, it will also give us insight into hidden features of the system.

At the same time, we will now begin negotiations with Linkon to acquire the T1 system. The product should have completed its beta test phase (scheduled for 2Q'95) recently and should be ready for application development. Prior to the start of the contract, Linkon had agreed to provide a unit for unlimited evaluation and co-development, and to potentially provide a small number of additional units at minimal cost to LDC for use in speech research applications. Linkon was very supportive of the development of an industry-standard platform using their system. Such an agreement is important because the list price of the system is $15K, much too high for the intended market. We need to formalize this agreement and acquire the first system.

In parallel, we will begin implementing a basic GUI using Tcl. Tcl was yet another tool that had problems building properly under Solaris 2.4. Now that we have a stable X and Tcl installation, we can proceed with the GUI development. Sun's XTL software will serve as an example of one such interface. We will also discuss the merits of the Intervoice GUI's with LDC, TI, and others who have experience with the system, and try to assimilate the useful features into our software engineering model. We will release an overview of the proposed design for comments in our next status report, which will be provided on August 31, 1995.

1. HARDWARE DEPLOYMENT

We experienced significant delays in the procurement and delivery of our computing equipment required for this project. Our original plan was to deploy all equipment prior to the start of the contract, in the hope that we could accelerate the schedule. Instead, the MS State purchasing process took many months to clear. In the end, we were able to successfully circumvent most of the process, and acquire the planned equipment at the quoted prices — but in several cases, unanticipated red tape cost us time. Once the purchases we approved and processed, a quality control problem with 1 Gbyte internal drives on the Suns further delayed shipment. Fortunately, our local Sun office pulled strings to get our machines delivered as soon as possible. We received our shipment of Suns on May 31, almost six months after the purchasing cycle was initiated.

An overview of the ISIP computing environment is shown in Appendix A. At the heart of the network are three Sun Sparcstations dedicated to separate tasks. A Sparcstation 5 with 32 Mbytes of memory performs fileserving and hosts approximately 30 Gbytes of data. A Sparcstation 20 with two CPUs and 192 Mbytes of memory supplies interactive and compute services. Most importantly, the third machine, a Sparcstation 5, is designated as the demo machine and will serve as the platform for the development of the T1 interface.

In addition to the basic computing hardware shown, audio and backup systems to support the project have been installed and verified. These systems are integral pieces of the data collection task, in that they will be used to archive and certify data collected via the T1 link. The backup system performs nightly incremental backups and monthly full backups, using a tape rotation scheme that will minimize the chance of losing any important data or software.

2. GENERAL SOFTWARE INSTALLATION

A major consideration in this project has been the development of an environment that minimizes the number of third-party software packages required. In the past, dependence on proprietary software has proven to create portability problems of the system proposed in this project (research organizations don't want to have to buy all sorts of software to use someone else's data collection platform). This consideration becomes increasingly important when dealing with foreign organizations for obvious reasons (export problems, lack of foreign-language support, lack of a local reseller, cultural differences to name a few).

Hence, we have based our system on freely available software such as MIT X, and the Free Software Foundation's (FSF) GNU tools. This proved to be a major headache in most software available for Solaris 2.4 seems to rely on Sun's compilers for the default installation mode. Further, MIT has not released a certified version of X for Solaris 2.4 (only 2.3 to date), and is indirectly dependent on Sun's compilation tools. Building X11R6 for Solaris 2.4 turns out to be quite a challenge. We had similar problems with the GNU software (which requires modification to Sun's standard Solaris include files).

To further complicate the picture, Solaris 2.4 is a non-trivial departure from Solaris 2.3. Hence, many vendors' software will compile for 2.3, but not for 2.4 without source code modifications. Worse yet, Solaris 2.5 is beginning beta-testing, and appears to be a departure from 2.4. Hence, in terms of developing software that is portable and stable, we appear to be shooting at moving targets. Nevertheless, we expect things to stabilize by the end of the year, and will be able to ship users whatever tools they need to work with the code. In addition, we will provide on-line support as needed, in order that we minimize the porting task.

Fortunately, we have overcome these problems, and have now developed a stable installation of many of the tools required for the project: MIT X (R5 and R6), GNU's gcc, Tcl, and perl. We have also assimilated a large database of infromation about what patches and modifications are generally required to get code to port to Solaris. With this important step behind us, we are now ready to begin looking seriously at the telecommunications hardware and software.

3. TELECOMMUNICATIONS-BASED INFRASTRUCTURE

The Linkon T1 hardware interface, a two-board S-bus system capable of interfacing a Sun to a T1 digital telephone line represents our last hardware obstacle. Sun's telephone services API, known as XTL, is the last major piece of software which we need to port to Solaris 2.4. XTL, though attractive in concept, currently is priced at $1000 per license, and seems to be overkill for this project. We also feel the package is too expensive for the project. Hence, we plan to evaluate the product, understand it, and assimilate its best features, but not require it for the LDC delivery system. We have acquired a copy of XTL courtesy of our local Sun sales office, and have installed it, but have not spent much time as of yet studying the software.

At the same time, we will now begin negotiations with Linkon to acquire the T1 system. The product should have completed its beta test phase (originally scheduled for 2Q'95) recently and should be ready for application development. Prior to the start of the contract, Linkon had agreed to provide a unit for unlimited evaluation and co-development, and to potentially provide a small number of additional units at minimal cost to LDC for use in speech research applications. Linkon was very supportive of the development of an industry-standard platform using their system. Such an agreement is important becuase the list price of the system is $15K, much too high for the intended market. We need to formalize this agreement and acquire the first system.

4. NEAR-TERM ISSUES

Our major concern thus far has been the stability of the operating system, and the degree to which Linkon will be able to provide support of the most recent version of the OS. It appears the Sun operating system will change at least twice in the next calendar year. Thus far, these changes have had a major impact on software — much more so than operating systems revisions in the past. Hence, it will make it increasingly difficult to provide an environment that is byte-for-byte compatible, because many of the users of the LDC system are not going to be running the latest revision of the OS. Yet, because so many things are changing so rapidly, older releases of the system are likely to be obsoleted quickly. Fortunately, our university has a very close relationship with Sun (our unviersity is currently beta-testing Solaris 2.5), and we can leverage this relationship to stay abreast of the OS issues.

5. NEAR-TERM PLANS

Over the next two months of the contract, we expect to accomplish three major tasks:

- acquisition and installation of the Linkon hardware;
- installation of the T1 link by MCI;
- implementation of low-level interfaces to the Linkon system.

Our goal will be to be able to demonstrate some minimal call processing capabilities by the end of August. Though we experienced delays in deploying our equipment, we are working overtime to get back on schedule. A demonstration of minimal capabilities will go a long way towards generating enthusiasm for this platform.

6. APPENDIX A: The ISIP Computing Environment

The following three pages contain graphics depicting the ISIP computing environment as it has been implemented as of June 30, 1995. It is physically located in three separate rooms, and includes a temperature-controlled computer room, a room dedicated to demonstrations and high-quality recording, and a student work area.

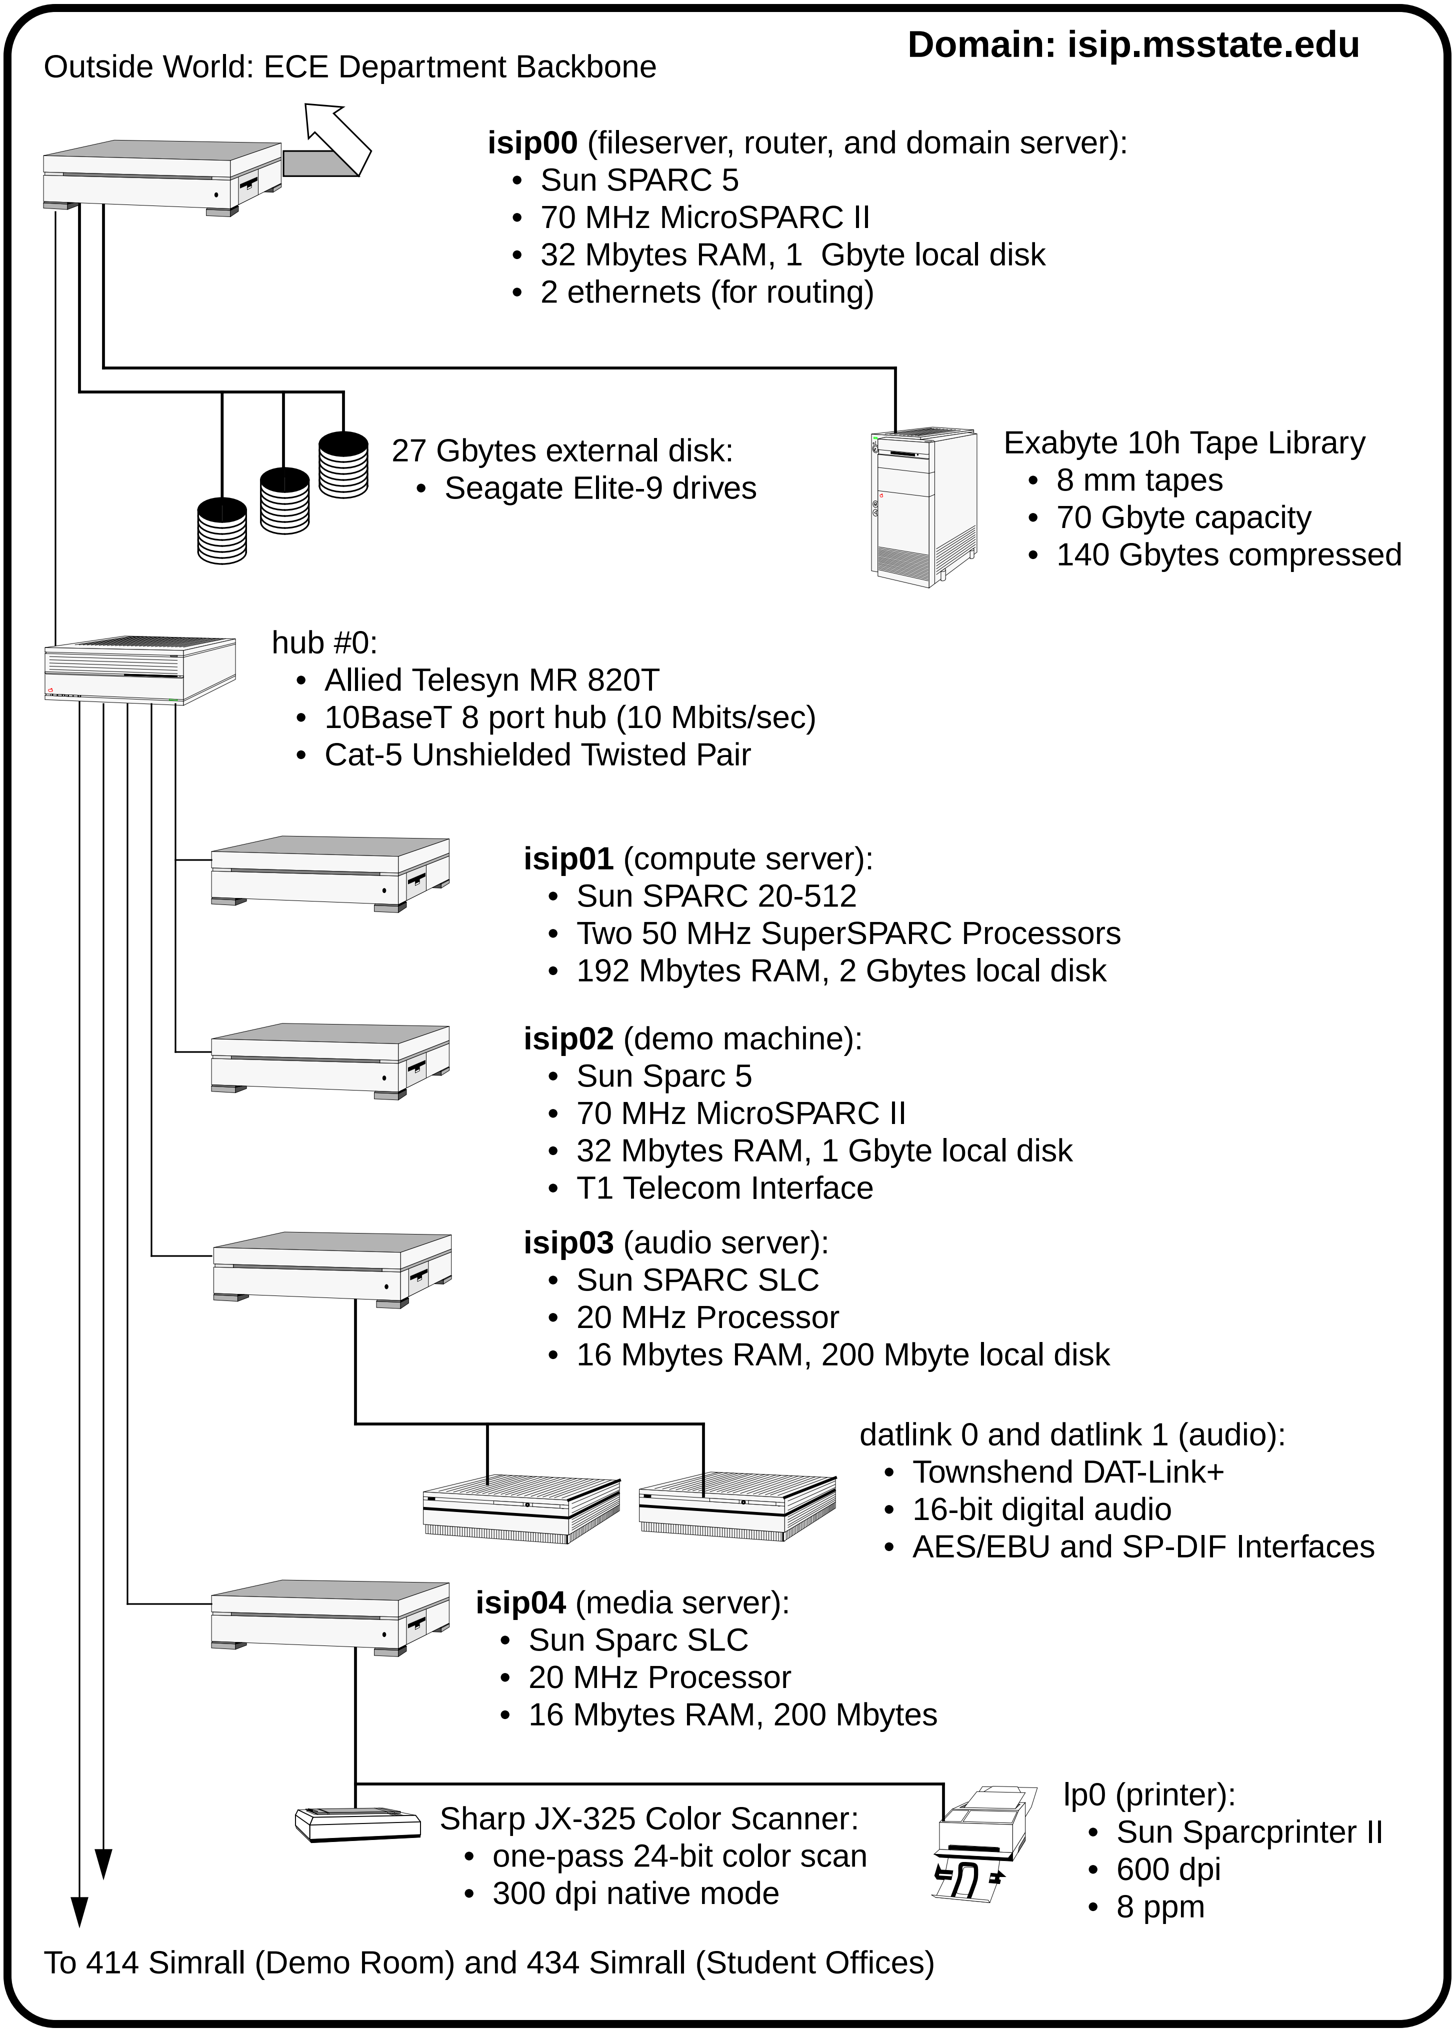
Domain: isip.msstate.edu
Outside World: ECE Department Backbone
isip00 (fileserver, router, and domain server):
• Sun SPARC 5
• 70 MHz MicroSPARC II
• 32 Mbytes RAM, 1 Gbyte local disk
• 2 ethernets (for routing)
27 Gbytes external disk:
• Seagate Elite-9 drives
Exabyte 10h Tape Library
• 8 mm tapes
• 70 Gbyte capacity
• 140 Gbytes compressed
hub #0:
• Allied Telesyn MR 820T
• 10BaseT 8 port hub (10 Mbits/sec)
• Cat-5 Unshielded Twisted Pair
isip01 (compute server):
• Sun SPARC 20-512
• Two 50 MHz SuperSPARC Processors
• 192 Mbytes RAM, 2 Gbytes local disk
isip02 (demo machine):
• Sun Sparc 5
• 70 MHz MicroSPARC II
• 32 Mbytes RAM, 1 Gbyte local disk
• T1 Telecom Interface
isip03 (audio server):
• Sun SPARC SLC
• 20 MHz Processor
• 16 Mbytes RAM, 200 Mbyte local disk
datlink 0 and datlink 1 (audio):
• Townshend DAT-Link+
• 16-bit digital audio
• AES/EBU and SP-DIF Interfaces
isip04 (media server):
• Sun Sparc SLC
• 20 MHz Processor
• 16 Mbytes RAM, 200 Mbytes
Sharp JX-325 Color Scanner:
• one-pass 24-bit color scan
• 300 dpi native mode
lp0 (printer):
• Sun Sparcprinter II
• 600 dpi
• 8 ppm
To 414 Simrall (Demo Room) and 434 Simrall (Student Offices)

Demo Room: 414 Simrall

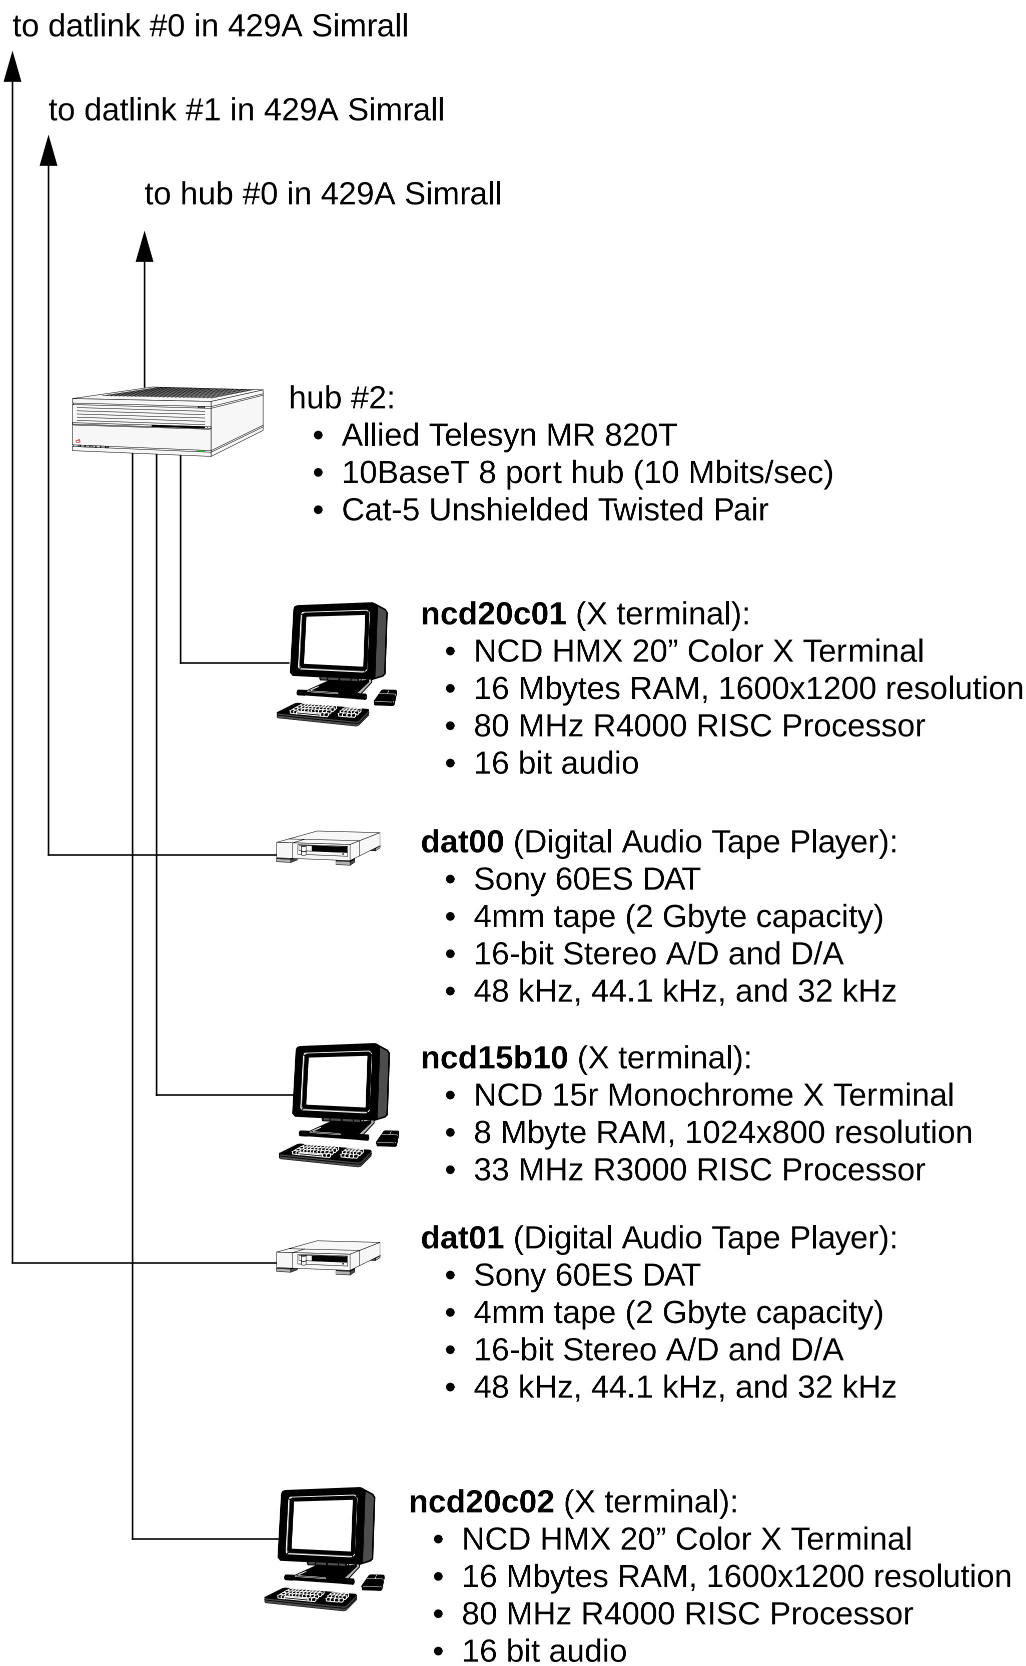

Student Office: 434 Simrall

to hub #0 in 429A Simrall

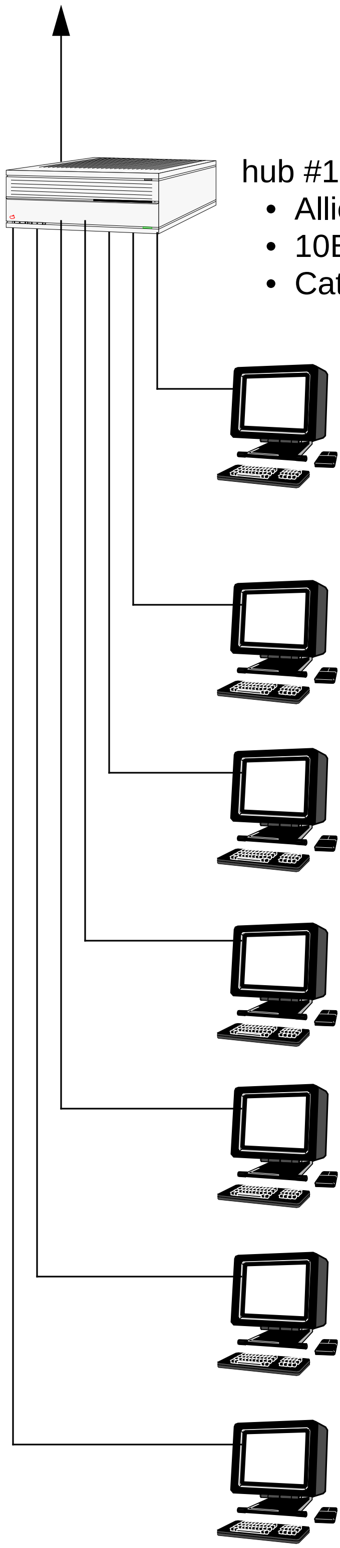

hub #1:
- Allied Telesyn MR 820T
- 10BaseT 8 port hub (10 Mbits/sec)
- Cat-5 Unshielded Twisted Pair

ncd19c00 (X terminal):
- NCD MCX 19” Color X Terminal
- 6 Mbytes RAM, 1280x1024 resolution
- 20 MHz 88100 RISC Processor
- 16 bit audio

ncd15b04 (X terminal):
- NCD 15r Monochrome X Terminal
- 8 Mbyte RAM, 1024x800 resolution
- 33 MHz R3000 RISC Processor

ncd15b05 (X terminal):
- NCD 15r Monochrome X Terminal
- 8 Mbyte RAM, 1024x800 resolution
- 33 MHz R3000 RISC Processor

ncd15b06 (X terminal):
- NCD 15r Monochrome X Terminal
- 8 Mbyte RAM, 1024x800 resolution
- 33 MHz R3000 RISC Processor

ncd15b07 (X terminal):
- NCD 15r Monochrome X Terminal
- 8 Mbyte RAM, 1024x800 resolution
- 33 MHz R3000 RISC Processor

ncd15b08 (X terminal):
- NCD 15r Monochrome X Terminal
- 8 Mbyte RAM, 1024x800 resolution
- 33 MHz R3000 RISC Processor

ncd15b09 (X terminal):
- NCD 15r Monochrome X Terminal
- 8 Mbyte RAM, 1024x800 resolution
- 33 MHz R3000 RISC Processor