1. S2Knet Committee Meeting Minutes

1.1. Summary of Sequoia Network Committee Meeting on 6/20/91 at UC Berkeley,

Present: Gary Darling (DWR), Joel Fine (UCB student), Jim Frew (UCSB), Paul Love (SDSC), Ken Lutz (UCB), Joe Pasquale (UCSD, acting chair), Mike Shannon (OP), Mike Stonebraker (UCB)

Absent: Domenico Ferrari (UCB, chair)

1.1.1. The Sequoia Network Committee

The Sequoia network committee has the following duties:

1. to evaluate Sequoia network design proposals, most likely generated by the Sequoia network researchers but also by the committee itself

2. to produce an implementation plan

3. to allocate resources (e.g., money) for the implementation

4. to make policy decisions regarding global networking issues

Each committee member should make a point of understanding his site's characteristics, such as number of users, applications, bandwidth and delay requirements, so that such information can be provided to the committee. Finally, each committee member is responsible for carrying out action items for his local site.

1.1.2. The Sequoia Network

The Sequoia Network should be competitive with the other major high-speed networking projects. In particular, we should strive for a network providing gigabit bandwidth, thus not necessarily limiting ourselves to a T3 backbone as described in the proposal.

The Sequoia Network is distinguished from other network projects in that research will focus on at least the following items:

1. performance guarantees: transmission of data which has high-bandwidth and delay constraints requires protocols which provide performance guarantees. These protocols will provide both statistical and deterministic guarantees for levels of bandwidth and bounds on delay.

2. compression: this may be a fundamental service provided by the network. Since information such as a time sequence of images requires such high bandwidth, it will be transmitted in a compressed form, and will be decompressed at the last possible moment (e.g., just before being displayed). This will require close interaction between the OS I/O system and the network.

3. variable packet sizes: given the wide range of information types that will be transmitted, a wide range of packet sizes (particularly large ones) will be needed. And since the amount of compression that can be done on an information stream can vary producing variable bandwidth requirements over time, variable packet sizes may be desirable.

4. fast inter-device data paths: a major performance problem is how to get data from a source device to the network interface device or from the network interface device to a sink device. An approach is to build fast data paths between these devices by minimizing data copying and protection boundary crossings, and to do some sort of inter-device DMA if possible.

The Sequoia Network will be an experimental network for the 3 year duration of the project. It will be an environment where the network researchers will have the ability to evaluate new protocols, and periodically take down the network for controlled experiments. These new protocols will support distributed Global Change applications, such as workstations visualizing images obtained from a large object server or from remote sensors. After the termination of the project, the network would be considered by the OP to provide production services for the UC system.

Since we already have internet services, the Sequoia Network will only be useful if it can provide new functionality (e.g., real-time guarantees) and/or better performance (e.g., higher bandwidth). Since these depend on progress by the network research group and acquisition of industrial support to build a high bandwidth backbone, the Sequoia Network will not be expected to be functional for approximately one year. However, as soon as we get support to build the network, we will at least provided basic TCP/IP service while protocol research goes on (this may occur within a year).

The Sequoia Network will be hierarchical: there will be a high-bandwidth backbone connecting the Sequoia campuses, and then some set of local-area networks providing intra-campus connectivity. The backbone termination points will be within the labs of the Sequoia researchers. The OP will be mainly responsible (in conjunction with industrial help) for providing this backbone. Mike Shannon offered support (i.e., one of his staff people) for the monitoring of the backbone. Each campus will be responsible for intra-campus connectivity.

1.1.3. Industry Support

Clearly, to meet the objective of building a gigabit network, we will need major industrial support. Specifically, we need to find a major communication company willing to donate, or provide at low cost, high bandwidth between the major Sequoia sites (Berkeley, Los Angeles, Santa Barbara, San Diego). This is because the funding provided by the OP (300K for year 1, 600K for year 2, 900K for year 3) will quickly be consumed if we tried to build a T3 backbone. Mike Shannon prepared some figures for T3 charges, which totalled approximately 67K for installation and 1.5M/year leasing charges for UCB, UCSB, UCSD connectivity. Rough estimates for a T1 backbone providing connectivity for UCB, UCLA, UCSB, UCSD, amounted to approximately 300K/year.

We will approach a number of companies for sponsorship, including AT&T (Domenico and Joe), MCI (Mike Shannon), Sprint (Mike Shannon and Paul). The names in parentheses are the committee members who will investigate these possibilities.

We will also see if DWR would allow us to use bandwidth provided by existing fibers running from Stockton to Riverside. Gary suggested that we first approach General Services for this (Mike Shannon and Joe). If this doesn't work, then Gary will help us approach DWR directly.

Other industry support will also be sought. In particular, we will approach Cisco (Ken), since they make a good routing machine, to determine whether we can have access to their protocol source code and software development tools. By an existing agreement with DEC, we will be able to buy Cisco equipment using our DEC money. However, Ken indicated that this agreement may only last for a year; Ken will look into whether this is indeed the case. Keeping this in mind, we will also approach DEC (Paul and Joe) to see if they have any routing machine in development that we may be able to use.

We will also seek support from IBM (Paul, who will ask for help from Hans-Werner Braun, and Jim who has IBM contacts) on their modified RS6000 hardware for routing. Mike Stonebraker indicated that IBM has already expressed an interest in supporting Sequoia, and would help explore this issue.

Finally, Mike Shannon agreed to look into providing a 56Kb/s link to connect UCSB and the Mammoth research station.

1.1.4. Schedule

In the short-term, we will meet 12:00 on July 31 at UCB to report on initial findings regarding industrial support. Each of us should have made contact and started a dialog with the various companies listed above. In approximately 3 months (end of September, exact date to be determined), we need to have a fairly clear picture on who our major sponsor(s) will be.

In the long-term, we expect the network to be physically existent (with at least T1 bandwidth) by the end of one year, with at least TCP/IP running. By the end of two years, we expect that the network support at least T3 bandwidth, and that protocols developed by the network research group be usable by others in the

Sequoia project.

1.2. Summary of Sequoia Network Committee Meeting on 7/31/91 at UC Berkeley,

Present: Fabrice Cuq (UCLA), Gary Darling (DWR), Joel Fine (UCB), Jim Frew (UCSB), Paul Love (SDSC), Ken Lutz (UCB), Cliff Lynch (OP), Ira Machefsky (DEC), Joe Pasquale (UCSD, acting chair), Mike Shannon (OP), Mike Stonebraker (UCB)

Absent: Domenico Ferrari (UCB, chair)

1.2.1. Status of Industrial Partnerships

Joe on AT&T: contacts have been made with Brian Quock (AT&T, San Diego), Larry Hughes (AT&T, San Ramon), Val Santo (PacBell, San Diego). They have been given the proposal. Larry Hughes responded after reading the proposal that AT&T is interested in pursuing the matter; he gave the proposal to people at PacBell, and said they are expressing interest also, and will get back to us. Joe said that Domenico Ferrari met with Sandy Fraser, who heads the systems research group at AT&T Bell Labs and currently supports the Blanca Gigabit network project which Domenico is involved, and said Sandy expressed interest in hearing about Sequoia; Joe sent the proposal to him.

Mike Shannon on PacBell, MCI, and US Sprint: Mike has spoken to Rick Albertson (PacBell), Lee Assenmacker (MCI), and Joe Goth (US Sprint). All have expressed enough interest to warrant a formal letter which is specific in what we are requesting, along with a general 2-3 page description of the project which will be prepared by Mike Stonebraker. Mike Shannon will prepare the letter, and send it out within the next week or two.

Mike Shannon and Gary Darling on DWR fiber: Mike has spoken with Department of General Services, and they are not willing to donate DWR fiber. Gary says that the next realistic strategy is to see about obtaining bandwidth at cost. This may be possible not only regarding DWR fiber, but other links that would directly connect the various campuses. Mike, with help from Gary, will investigate this further after we know more about the status of our industrial partnerships (although if anything can be pursued in parallel, it should; Mike will make a judgement on this).

Mike Shannon on 56Kb link to Mammoth: The remoteness of this location makes obtaining this link non-trivial, and Mike is still working on this. He expects that something should be available within a week or so.

Ken Lutz on Cisco: Ken has made some inquiries with Cisco regarding using their routers, and obtaining source code. The general feeling he got was that we need to make a careful strong concerted effort in obtaining source code, as these are the "family jewels" and won't be easy to come by. Ken will get the name of someone at Cisco who could make such a decision (such as VP of research), and Joe will begin planning a meeting between this person and Sequoia researchers led by Domenico. Given the importance of identifying routers for the Sequoia network, this meeting should happen as soon as possible.

1.2.2. DEC Networking

Ira Machefsky expressed that DEC should be given a first "right of refusal" regarding any networking equipment. If DEC has appropriate equipment, it is to our advantage to use it since we will presumably have access to the people who designed and built it, and perhaps may be able to affect future design decisions. Ira has identified two groups at DEC which are of potential interest to the project: Al Kirby's group in Maynard, and Chuck Thacker's group at SRC. Communication between the two groups and various Sequoia investigators (Hans-Werner Braun, George Polyzos, Joe and Paul) have already started.

Ira mentioned that Al Kirby's group is currently working on a router that will support FDDI and T3; this product is 18 months down the road before general public availability. However, Ira said we should be able to get tested versions before that time, and will look into exactly what is possible.

Visit to DEC SRC

Joe reported on a visit on July 30 to DEC SRC by George, Ira, Joe, and Paul. We met with a fairly distinguished group of about 10 DEC researchers. Michael Schroeder and Chuck Thacker talked about Autonet I and II. Autonet I is a 100Mb local-area network; it is a point-to-point network, and not a ring like FDDI, so there is more potential for effective use of the 100Mb bandwidth. Also, it is a reliable network in that many connections are redundant, and the network can automatically reconfigure its routing based on sensing what links are available. A demo of was given where the power cord was pulled out of one of the network's switching devices (of the production network being used by people working there), and automatic reconfiguration took place; very impressive.

The DEC researchers were particularly interested in Sequoia in that it is rare to find applications which can actually use up large amounts of network bandwidth. In fact, their first question to us was "why do you think you need so much bandwidth?" After describing the applications, they seemed convinced that indeed, a Gb/s or more is justified. They said they would be interested in hearing more about the project, in particular the specific requirements of the applications, as this gives them a better idea on how to direct their research also.

Autonet I would be an interesting network on which to base the Sequoia network, especially since Auotonet II (to appear in 1-2 years) will support 1Gb bandwidth. However, switches are not available in any quantity, and it is not expected that they will be made in the near future. One possibility is to use DecStation 5000s as switches, since Autonet interfaces will be soon available (by the end of this year). However, this is risky in light of one of our goals, which is to provide a fast reliable network with at least TCP/IP support to other Sequoia researchers after a year. Another possibility is to look into using Autonet while relying on Cisco routers or the DEC routers from Al Kirby's group in the interim period. We need to consider the overall costs of this approach (Each Cisco router will run 30-50K), and whether network research based on the use of Autonet has any significant advantages over using FDDI, which is currently available from DEC. Possible advantages are more effective use of bandwidth (i.e., better performance guarantees), reliability, possibility of extending Autonet over a wide region (the DEC researchers said there is no inherent reason why Autonet could not span the distances needed to connect the Sequoia sites), and maybe most important, collaborations between Sequoia and DEC researchers. Disadvantages are limited availability of hardware, non-robustness since it is yet untested for a large production setting over large distances, and non-standardness. These are not necessarily absolute disadvantages, but relative ones to a FDDI/T3 network.

1.2.3. Research Problems

The network researchers have been meeting approximately once a week to discuss research problems. Some of the topics have been the effectiveness of simple priority-based vs. protocols that provide explicit performance guarantees, in light of what Sequoia applications will expect from the network (Domenico, who is away this summer, will probably have alot to say about this since he has focused on performance guarantees). Another topic has been how to avoid one of the major problems affecting high-bandwidth I/O, which are the multiple buffer copies that happen in current I/O system software. One possibility is to build provide support for direct device-to-device communication, although the DecStation Turbochannel won't support this without at least one copy to memory. Another possibility is to plan I/O such that advantages that come from data being page-aligned are taken advantage of.

One very clear problem is that we don't fully understand the performance requirements of the Sequoia applications. This will be a major goal at the August 8 Sequoia investigator meeting: we need to talk with the scientists to better understand their applications, and what the current and future I/O requirements are.

Gary Darling mentioned that PacBell is implementing an experimental ISDN network at Davis, and that we should consider whether ISDN would be useful (despite its low bandwidth).

1.2.4. Plan of Action

Mike Shannon will send out letters to his contacts at AT&T, MCI, US Sprint, indicating what we need from them. He should report on any developments as soon as possible, so that we are best prepared to discuss strategy at the next network committee meeting at the end of September. This strategy will involve what to do about DWR fiber (do we need it, if so, do we buy it, at what cost) and how to refocus our energies and resources. Mike will also look into NET (Network Equipment Technologies), which supplies the T1 muxes currently used for the inter-campus network, on obtaining information on their latest technology, which includes a switch supporting T3 and Sonet. The contact person is Tim Schaffer (415) 546-3009.

Domenico and Joe will continue to pursue support from AT&T. Domenico should be back by August 23. Joe will keep Mike Shannon informed on any developments so that our efforts are coordinated.

Mike Stonebraker will prepare the 2-3 page general description of the Sequoia project, which can be used by Sequoia researchers in soliciting support from industry.

Ira will continue developing contacts at DEC, and will determine availability of the DEC FDDI/T3 router.

Cliff will look into Proteon network equipment. Networks which have interfaces for the DecStation 5000, and any routing and switching equipment are of interest. Cliff will also talk to Richard West since he is on the board of directors of ANS (Advanced Networks and Services, a not-for-profit organization set up by IBM and MCI), so see what kind of help they might be able to provide.

Paul will look into NSC, Network Systems Corp. According to Paul, "They used to make HyperChannel 50meg/sec 15 years ago the best you gould get. Now they make T1/T3/FDDI/Hippi gear. Also speciallized interfaces to such machines as Cray's, IBM's, etc as well as norma; ones to VME, UniBus, QBus, BIBus, MultiBus, etc." Paul should check if they support the DecStation 5000 Turbochannel.

Ken will determine the name of a high-level person at Cisco. Joe will plan a meeting with this person and a group led by Domenico.

Ken and Cliff will begin developing a plan for integration of the Sequoia network with other networks. The goal is to come up with a model of addressing, routing, and interconnectivity between a local Sequoia site and local points of contact with other networks. An email discussion of this was encouraged by various participants.

Fabrice and Jim will talk to the scientists at their respective sites to find out what their network requirements are. Here are some initial questions (feel free to add to the list):

> Can you give me a list of applications that you use regularly?
>
> If any are interactive, what is the nature of the interactivity (e.g., pressing a key to fast forward or stop)?
>
> How big are the data files that are accessed?
>
> Is the data accessed sequentially or randomly?
>
> Is the data accessed all at once, or in chunks?
>
> If in chunks, how large are the chunks? Do they get accessed at a regular rate (periodically), and if so, is this important to the application (a time-sequence of images requires periodic updates to simulate continuity)?
>
> The whole point of this is to determine network I/O requirements, in terms of required bandwidth, bounds on latency, reliability, etc. In general, can you characterize the I/O requirements of the applications in these terms?

1.2.5. Next Meeting

The next meeting will be help on Monday September 30th at Berkeley. We did not have time to discuss whether this was a good date for everyone (although those I had a chance to mention it said it was OK), so if it is bad for you, let me know. 12:30 seemed to be a good time, so let's try that again.

1.3. Summary of Sequoia Network Committee Meeting on 10/8/91 at UC Berkeley,

Present: Fabrice Coq (UCLA), Gary Darling (DWR), Domenico Ferrari (UCB, chair), Jon Forrest (UCB), Jim Frew (UCSB), Ken Lutz (UCB), Cliff Lynch (OP), Ira Machefsky (DEC), Joe Pasquale (UCSD), George Polyzos (UCSD), Mike Shannon (OP), Mike Stonebraker (UCB).

Absent: Paul Love (SDSC)

1.3.1. Bandwidth

The Committee approved the following tentative schedule for providing bandwidth to the Sequoia 2000 Project:

- T1 connections among the 6 Sequoia sites - March 1992
- T3 connection between UCB and UCSD - ASAP
- T3 connection among the 6 Sequoia sites - December 1992

The strategy for bandwidth procurement will be as follows:

(1) Gary will organize a meeting between a Sequoia delegation (Mike Stonebraker, Mike Shannon, and Domenico Ferrari) and Archie Noriega and J. Kiliany in Sacramento, to discuss the possibility that the Sequoia network use as a backbone the General Services fibers connecting Stockton to Rialto, and the price the project woul have to pay for such use. The meeting will take place as soon as possible. If the results of the meeting are unsatisfactory, we will ask Richard West to meet with David Kennedy, Director of DWR.

(2) In the meantime, Mike Shannon will intensify his discussions with MCI and U.S.Sprint, and try to convince either company to donate access to fibers connecting the 6 Sequoia sites (DWR, UCB, UCSB, UCLA, UCSD, Scripps) according to the above schedule. If strong resistance is encountered, and the General Services fibers can be used, Sequoia will ask MCI and/or U.S.Sprint for connections between its sites and Stockton (for northern sites) or Rialto (for southern sites).

1.3.2. Routers

The strategy for router procurement will be as follows:

(1) Each site will buy with DECdollars a 3MAX+ workstation to be used as the local router. Each site will also receive (either free of charge or for a suitable amount of DECdollars) one or more T1 boards from David Boggs, and a copy of NSL's router software. These will be the routers for the T1 Sequoia network, which will have to be operational by March 1992 (see the above schedule).

(2) Ira, with help from Domenico, Joe, and Mike Stonebraker, will exert pressure on David Boggs to convince him to design, implement, and have someone to produce, T3 Turbochannel boards. To be useful to Sequoia, these boards should be ready to be distributed to the Sequoia sites by December 1992, and sooner for the UCB-UCSD T3 experiments.

(3) If solution (2) will turn out not to be feasible, we will consider procuring a "DS3 box" and an UltraNet Hub from UltraNetwork Technologies; the box will provide the DS3 interface, while the DEC workstation that will have served as T1 router will be connectable to the Hub via a VME interface obtained by the use of DEC's Turbochannel-to-VME adaptor board. Other alternatives to be considered in that case will be the Sonoma router, being developed by DEC at Littleton, MA, and the Hastings router, being developed by DEC in England. We should be able to decide whether solution (2) is viable around the end of 1991 or the beginning of 1992.

1.3.3. Next Meeting

The next meeting will be held at UC Berkeley on Tuesday, December 2, at 2 pm.

By November 15 we will decide whether to invite the networking manager(s) from each of the sites to the December 2 meeting or to the meeting after that. The main topic we need to discuss with them is the design of the local connections to the Sequoia network, including addressing and routing issues. Each member of this committee will be asked to designate, in consultation with Mike Shannon, and to contact the person(s) to be invited.

1.4. Summary of the 4th Sequoia Network Committee Meeting on 12/17/91 at UC Berkeley, (prepared by D. Ferrari, 12/18/91)

Present: Lee Assenmacher (MCI), Fabrice Coq (UCLA), Domenico Ferrari (UCB, chair), Jon Forrest (UCB), Amit Gupta (UCB), Paul Love (SDSC), Ken Lutz (UCB), Ira Machefsky (DEC), Joe Pasquale (UCSD), Mike Shannon (OP), Mike Stonebraker (UCB), Fred Templin (DEC).

Absent: Gary Darling (DWR), Joel Fine (UCB), Jim Frew (UCB), Cliff Lynch (UCOP).

1.4.1. Bandwidth procurement: the meeting in Sacramento

Mike Shannon and Domenico reported on the meeting held in Sacramento on December 10. The Sequoia delegation (Gary, Domenico, Mike Shannon, and Mike Stonebraker) met with Jerry Kiliany of GSA and Archie Noriega of DWR to discuss the possibility that GSA and DWR provide fibers to be operated at T3 speed for the Sequoia Network. The prices such facilities can be offered by the two organizations were found to be too high to be affordable for the project.

1.4.2. Bandwidth procurement: the MCI deal

Mike Shannon gave background information about his negotiations with MCI, mentioning the strong ties that have existed since 1986 between MCI and UC, and then invited Lee Assenmacher to describe the deal being proposed by MCI to the Sequoia 2000 project. The essentials of the MCI proposal can be summarized as follows:

- MCI will connect UCB, UCLA, and UCSD to UCSB with T1 links from March 1992 through March 1993, provide a T3 link between UCB and UCSD from December 1992 through March 1993, and connect the same three campuses to UCSB with T3 links from March 1993.
- The (heavily discounted) price to the project of these MCI connections will be $164,112.00. This price does not include local-access charges by the local telephone companies. The installation prices for local access facilities are $1127.00 (T1) and $23,850.00 (T3) for each campus. The per-campus monthly charges will be approximately $750 (T1) and $3800 (T3). Over the lifetime of the (first phase of the) Sequoia 2000 project, each campus will have to pay about $10,900 for local access to the T1 network, and $88,500 for local access to the T3 network.
- Starting in July 1994, UC will pay 80% of the standard tariffed rates for backbone facilities for two years, with a minimum undiscounted monthly rate of $87,536.00.

The proposal has still to be approved by MCI's top management, but Lee expects it will soon be. The Committee approved the deal unanimously, and thanked Lee Assenmacher and Mike Shannon for negotiating it. A contract will have to be formulated and signed very soon to make the agreement final.

As for the subdivision of the bandwidth charges between UCOP and the Sequoia 2000 project, Mike Shannon and Mike Stonebraker agreed that:

- UCOP will pay all the charges mentioned in the MCI letter and its attachments, including those for local loops and connections;
- the Sequoia 2000 project will pay all hardware costs at each site, including the possible $50K exposure for T3 modems at UCSB.

1.4.3. Router procurement: the T3 board problem

Ira and Joe reported on their meeting with David Boggs of DEC WRL. They were able to convince Boggs that his T3 boards are crucial for the success of Sequoia, and got him to agree that he will try to have a few boards ready by December 1992. Boggs will be able to make a firmer commitment to this project in about two months from now. Joe proposed that the use of IBM RS/6000 workstations and their T3 boards be considered as a fallback alternative. Fred suggested that some of the networking people at DEC Littleton be asked about the existence of any T3 board development projects within DEC, as well as about the possibility that some of them might be given the task of designing and developing such a board. Ira said he has heard rumors about a T3 board that Chuck Thacker would be building at DEC SRC. The Committee decided that a DEC solution would be preferable, and asked Fred to contact Bill Hawe and Chuck Thacker to inquire about current or future T3 board development projects. The IBM solution will be kept as a less desirable but possibly viable alternative.

1.4.4. Router procurement: experiences and configuration

Joe reported on the experiences recently made at UCSD with a DS5000/200-based router and David Boggs's T1 boards, as well as his two-day experimentation with a T3-speed connection between the UCSD campus and the Educom Conference site. He also proposed the following configuration for the router at each Sequoia site:

DECstation 5000/240, 32 MB

inexpensive monitor (e.g., CRT)

FDDI board

665MB disk drive

The total price of this configuration is about DEC$29,000, plus the price of an extended warranty option for 2 additional years (note that the local system at each site will have to include an FDDI concentrator, whose price is about DEC$11K plus DEC$3K for 2 years of extended warranty). Since three T1 or T3 boards and one FDDI board do not fit into the three available TURBOchannel slots, UCSB will have to install two routers, hence to buy two machines with the above configuration. If the topology of the network were modified, another campus would have to do so, unless the topology were reduced to a chain of sites. This would cause the local loop at UCLA to become substantially more expensive. Mike Shannon will study the problem and recommend a final topology within the next few days.

1.4.5. Routing problems

Jon and Ken presented their views on the routing loops that can potentially be created by connecting the same machines to the Sequoia Network and the Internet through each campus's networks. After extensive discussion, the Committee asked Jon and Ken to circulate, some time around mid January, the guidelines they would recommend each campus to follow so that unpleasant routing phenomena can be avoided. The guidelines should cover all the networking configurations that are likely to be found in the various campuses.

1.4.6. Review of the Sequoia Network time-table

Domenico summarized the current state of the time-table as follows:

- **December 1991**: Mike Shannon recommends the best topology, and determines whether and where T3 modems will be needed; Joe finds out how many T1 modems will be needed, if any; campuses order equipment for their router(s); agreement with MCI is finalized and signed, and T1 circuits (and T1 modems, if necessary) are ordered
- **January 1992**: Ira procures T1 boards from David Boggs; experiments are conducted (primarily at UCSD and, if possible, by Fred) on NSL's router software and Boggs's T1 boards; Jon and Ken circulate their proposed guidelines

- **February 1992**: David Boggs announces his plans; decisions are made about the possible T3 board alternatives
- **March 1992**: T1 network is deployed; router equipment is delivered and installed; local FDDI rings are set up
- **April-November 1992**: Protocols and operating system enhancements are developed at UCB and UCSD
- **December 1992**: T3 boards are given to UCB and UCSD; T3 link between UCB and UCSD is installed
- **January-February 1993**: T3 experiments between UCB and UCSD
- **March 1993**: T3 network is deployed; T1 network is discontinued; router and host software developed at UCB and UCSD is distributed to all campuses.

1.4.7. Next meeting

The next meeting will be held on **Tuesday February 4, 1992** at 2 pm. Each campus representative is expected to identify the right networking person within his campus or department or institute, and invite that person to the meeting, where the local configurations and the guidelines to be proposed by Jon and Ken will be discussed.

1.5. Summary of the 5th Sequoia Network Committee Meeting on 2/5/92 at UC Berkeley, (prepared by A. Gupta, 2/12/91)

Present: Gary Boyles (Digital), Fabrice Cuq (UCLA), Gary Darling (DWR), Domenico Ferrari (UCB), Jon Forrest (UCB), Jim Frew (UCB), Cliff Frost (UCB), Amit Gupta (UCB), Ken Lutz (UCB), Ira Machefsky (Digital), Jim Madden (UCSD), Steve Miley (UCSB), Joe Pasquale (UCSD), Vince Sefcik (UCSB), Jerry Shannon (UCSD), Mike Shannon (UCOP), David Shaughnessy (UCOP), Mike Stonebraker (UCB), Fred Templin (UCB and Digital).

1.5.1. Bandwidth procurement and network topology

Mike Shannon reported that the decision to base the hub of the S2K network at Santa Barbara was taken because of the extra $240K that a hub in San Diego would have required. He also announced that the order for the T1 lines was given to MCI, and he distributed copies of the order to all the members present. The T1 links will be in place around mid-March, and the local loops should be installed in the first half of that month. He added that there is a "schedule slippage" by about 1-2 weeks, due to the fact that his office had to do a lot more "homework" than anticipated.

David Shaughnessy informed the Committee that the monitoring of the T1 lines for diagnostics purposes will be external to the lines, as the T1 Turbochannel boards have an interface that does not permit the use of the IDNX boxes on the links themselves.

The Committee reviewed the topology of the network, and briefly discussed the connections to Scripps and to DWR. While Scripps can be connected "any time" using an existing broadband link to UCSD, DWR is still debating how to connect. The Committee suggested that DWR could lease a T1 link to UCB as soon as possible for this purpose. Such a link would cost $2.8K per month.

Joe expressed his disappointment at the decision not to locate the hub at San Diego, but said that he understood the economic and financial reasons behind the decision. He stated that the project needs a plan for handling possible network failures at the hub. There should be some person at UCSB who can be contacted at any time if there is a connectivity problem. Steve said that he could do this with the help of another person. Joe also proposed that 9.6Kbps modems and reverse terminal servers be procured for remote diagnosis and therapy of the hub routers. Mike Stonebraker asked Joe to submit a "wish list".

1.5.2. Sequoia Network evolution and the Stratacomm proposal

Ira reported on a meeting that he and other members of the Committee had with Dr. Hanafy Meleis, the new DEC Technical Director of Corporate Networks. Dr. Meleis is very interested in Sequoia 2000, and has convinced two companies to design and implement T3 Turbochannel boards. Both the boards should be ready by mid-July. The promise of these product-quality boards from two different sources drastically reduces our dependence on David Boggs's project. Brian Reid has promised to have the NSL write the driver code for the T3 boards.

Ira stated that Dr. Meleis has suggested that Sequoia 2000 procure (through an External Research grant from Digital) the switches made by Stratacomm and marketed by Digital, which would allow channelization and facilitate the transition of the S2K network to a faster technology such as SMDS.

Domenico took the opportunity to point out the need for a long-term plan. In particular, he asked what Sequoia will do at the end of the first 3 years of the project, when the T3 links will be incorporated into the UC administrative network. The need for such a plan was recognized by the members of the Committee, and it was agreed that the discussions (and perhaps negotiations with MCI) should be started soon.

David Shaughnessy pointed out that the DSUs for T3 are very expensive ($20K per DSU), and that each link will need two of them. He also asked what interfaces the T3 boards will have. The answer was that they will have an HSSI interface.

1.5.3. Router and T3 board procurement

Joe reported on his interactions with David Boggs, who indicated that he would like to have a student assistant for the summer and possibly beyond (perhaps an M. S. student doing his project under David's supervision). He also said that he would prefer an EE/hardware student, rather than a CS/software one. Due to the geographical factor, it was decided that the student will be recruited by the Berkeley contingent.

Joe also reported on the work being done at UCSD on testing/improving Ultrix networking code performance.

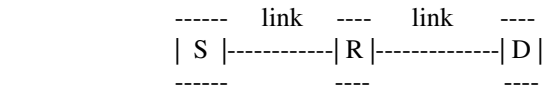

The setup in the picture, consisting of a source, a router and a destination, was the one used by the group at UCSD. Running the standard Ultrix TCP/IP code, the T1 link was saturated (at about 1.3Mbps out of the rated 1.5 Mbps). However, the results with FDDI were disappointing, as only 11% of the rated FDDI bandwidth (100Mbps) could be exploited. Further investigation showed that the sending CPU was the bottleneck. The sender CPU time was spent as follows:

Checksumming	40%
Data copying	40%
Mbuf free/alloc	10%
Other	10%

The UCSD group modified the Ultrix kernel and reduced the CPU consumption for checksumming to 15%. Even then, the throughput did not go beyond a disappointing 13-15 Mbps.

1.5.4. Campus plans for the Sequoia 2000 network

Ken Lutz requested the various campuses to send him a map of their local network configuration. This information will be kept in a database, and will be used for several purposes, e.g., fault handling and repair.

Following Ken's presentation of the addressing alternatives, the Committee decided to get one Class C IP group of addresses per site for the S2K machines. It was also decided that external traffic will not be carried on the S2K network.

Ken advised the members of the Committee to avoid "dual homing", i.e., assigning two IP addresses to any S2K machine. He showed some sample network configurations and topologies that can meet the needs of the S2K researchers without requiring dual homing.

The Committee agreed that the additional routers needed in Santa Barbara should be paid for with project-wide DEC$, not with DEC$ from UCSB's allocation.

The local configuration at UCSD was presented by Jim Madden and Jerr Shannon, the one at UCSB by Vince Sefcik.

Joe said that UCSD will need at least $10K for local fiber connections to the T1 (and, later, T3) network. The total amount required by UCSD for this purpose has been estimated to be around $27K. Fabrice indicated that UCLA will also need some for fiber modems. Domenico volunteered to be the recipient of these requests (which he urged the various sites to estimate more precisely and as rapidly as possible) and to convey them to Mike Stonebraker. Domenico also asked the site representatives whether they had ordered their routers. All the answers were in the affirmative, except that from UCSD, which was uncertain, as SDSC is responsible for issuing the orders, and Paul Love could not attend the meeting.

1.5.5. Review of the Sequoia Network time-table

Domenico stated that he did not see any need for updating the time-table he had proposed during the Dec 17 meeting. The Committee agreed that the time-table should not be changed at this time.

1.5.6. Next meeting

The next meeting will be held on **Wednesday March 11, 1992** at 2 pm on the Berkeley campus.

1.6. Summary of the 6th Sequoia Network Committee Meeting on 3/11/92 at UC Berkeley, (prepared by A. Gupta, 3/18/91)

Present: Fabrice Cuq (UCLA), Domenico Ferrari (UCB), Jon Forrest (UCB), Jim Frew (UCB), Cliff Frost (UCB), Amit Gupta (UCB), Ken Lutz (UCB), Ira Machefsky (Digital), Steve Miley (UCSB), Joe Pasquale (UCSD), Vince Sefcik (UCSB), Jerry Shannon (UCSD), Mike Shannon (UCOP), David Shaughnessy (UCOP), Mike Stonebraker (UCB), Fred Templin (UCB and Digital).

1.6.1. Update on T1 link installation

Mike Shannon reported that the installation of the T1 fibers and of the CSUs was completed by March 5. David Shaughnessy informed the Committee that automatic CSU and line testing will start once the fibers are properly terminated. The lines will be continuously tested (by sending repeated bit patterns) until the routers come up.

1.6.2. Campus LANs, Routers and Tail Circuits, and Network Project Planning

Items 2 and 3 of the agenda were discussed together, since the detailed plan for the S2K network presented by Mike Shannon and David included LANs, routers, and tail circuits.

The local wiring connecting the CSU to the router had not yet been done at Berkeley and Los Angeles. Los Angeles expected its wiring to be completed in about one week. Fred Archibald is handling the wiring in Berkeley.

David requested the creation of a separate mailing list for those involved in the day-to-day operation of the Sequoia network. The list should include at least one person at each site. That person should be responsible for taking deliveries, handling equipment, organizing regular maintenance, and taking care of connectivity problems. Jon Forrest volunteered to set up the mailing list. The alias will be called "sequoia-net-ops@postgres.Berkeley.EDU" (or something similar).

The deadline for the installation of all routers (or their temporary replacements) was set at Wednesday March 18. The Committee decided that a big effort will be made to demonstrate roter-to-router connectivity between at least some of the Sequoia sites before the Retreat.

Ken informed the Committee that he had already submitted a request for Sequoia addresses for Berkeley and Santa Barbara. He will get the desired addresses within 5 days. For the rest of the sites, he needs some more information, in particular, the name of the contact person at the site, to be supplied to the NIC. He is asking for two class-C addresses for each site. The need for consistent router and workstation names was also stressed. Jon volunteered to take care of the naming problem.

The deadline for router-to-FDDI connections was set at April 1 for UCB, UCLA, UCSD, SDSC, and Scripps. UCSB should get its connections up by March 26 plus shipping delays. UCLA, SDSC, and Scripps have not obtained their FDDI concentrators yet.

Joe presented to the Committee some of the results of the work being carried out at UCSD. The peak FDDI throughput obtained was 2 MBps at the user process level, and 7 MBps at the device driver level, while the maximum throughput of an FDDI ring is 12.5 MBps.

Joe also said that the driver code for the router boards is available, and promised to send an Ultrix kernel image including the driver code to all the sites via the Internet. He also requested that Ultrix4.2a sources be made available as soon as possible. A problem in San Diego is that the fiber terminators are missing. David promised to provide as many 30 feet cables as he will be able to get from his inventory.

Fred will set up a source code control system for Sequoia.

1.6.3. The Network for the Second Phase of the Project

Domenico restated his conviction that it is time to start thinking about the Sequoia network for the second phase of the project (the 1994-96 period). Mike Shannon and David indicated that UCOP will not exclude the Sequoia researchers from the T3 network as of the expiration date of the first phase. There was, however, widespread consensus that a higher-speed replacement should be considered anyway.

Domenico discussed SMDS, and said that it would be attractive, especially if we could be given some control over the network instead of being sold a black-box service. Ira announced Elysia Tan's seminar on SMDS (Berkeley, April 2, 4-5 pm), and said that Ms. Tan, who is with Bellcore, seems interested in discussing possible collaborations, and even experimental modifications of SMDS by Sequoia researchers.

1.6.4. Review of the Sequoia Network Time-table

Domenico, Ira and Fred reported on the results of their recent meeting, in which they discussed a plan for Sequoia networking research. UCB's Tenet real-time protocols will be ported to the Sequoia network in the Summer and Fall of 1992, under Fred's leadership. The protocols should be running on the Sequoia network by December 1992, when the T3 link between UCB and UCSD will become available for experimentation. UCB and UCSD researchers are planning to meet soon to coordinate their respective networking and operating system activities.

1.6.5. Next Meeting

The date and time of the next meeting will be discussed at the Sequoia Retreat.