

AIM Executive Committee Telecon
Transcript - 10/6/77

7) SUMEX Equipment Upgrade

Rindfleisch: We have the new tape and file system and the additional 256K of memory on line. That has provided quite a bit of additional filespace, almost a factor of 2, that has partially been allocated among the various projects, and the tapes have made a substantial improvement in the dumping time. --We're no longer taking the system down to users on Sunday mornings since we have the tapes now that are fast enough. We just do the dumps with the system running. Every month we do the dump with the users excluded so that we get all of the files. --If the user is working and has a file open then it wouldn't be dumped, so just to be sure we get all of them on a monthly basis, and we exclude users for that morning. The memory has made a substantial qualitative difference in the feel of the system in the sense that there are more jobs in memory that are runnable now, and I think that the predicted things have happened: a) the load average has come down because jobs are getting run faster; and b) when the load average gets up around 6 or so the feel of the system is better in that responsiveness is faster. We're still collecting quantitative data to measure the improvement in I/O wait time and that sort of thing.

Myers: Here at Pittsburgh, the addition has been tremendously valuable so that we can work now at times when we couldn't at all work before from the standpoint of speed and whatnot, and we've found it tremendously helpful.

Rindfleisch: We've had similar comments from other users, and that was the main thing that we were after in making these augmentations.

8) Rutgers Site Visit/Resource Renewal
(impressions on site visit and where resource stands in general picture)

Amarel: We had a site visit on 9/9, and it was very good, very thorough, an excellent group chaired by...and including Harry Pople, I heard...a few days that the council approved continuation of the resource for 3 years, and I have some...guidelines about the level of the resource. It's higher than the one we have now, less than what I proposed. I'm still at the point where I want to study some of the measures that come from the study section... I don't have yet the particulars but a fairly good idea. The indication is to continue most of the scientific work we're doing and the collaboration and a certain amount more specifically in the amount of computer science that we have, but the study section didn't approve the upgrade of the current system... That's where we stand at this point. I still want to discuss that in some detail with Bill to see exactly what it means...

Lederberg: Does that mean that there will not be a substantial

increment of computer capability that could be used by the AIM?

Amarel: The way I read the scientific and technical discussion, it looks as if there are not going to be substantial increases in the computer capability to the general community but rather this is going to be mainly focused on the environment for our own research activities at Rutgers plus collaborators that are working very closely with us rather than general groups of users in the AIM community...

Baker: That's not the way I read it. (4 on agenda) The recommendation was to go ahead and buy out the Rutgers machine. This leaves a whole lot of maintenance costs to be borne by the resource activity. Rutgers would still have to maintain some interest in supporting the machine in terms of maintenance. The implication of going out and buying out the machine I interpret as adding additional capacity to the overall AIM community, and it's going to be up to Saul to decide what that means in terms of connect time from people outside his own set of collaborators and activities there. The Stanford connect time is approaching 150K hours/year, so the way I would like to make the assessment is some sort of picture from Saul of how much connect time he sees he can put out to the best of the community with the kind of level we'll be supporting him at. I think you figure it's about 5% at the present time; I think it should be more than that under the situation where we're going towards the payoff of the hardware. That's the difference that Saul and I have right now. I have no idea of what Saul sees as an increment from 5%; he'll have to determine that and tell us.

Amarel: I'm not sure there is a difference because I'd like to have at least another day to look into the details and another day to talk with Bill, and I think this particular discussion would have been much more productive a week from now. My own goal is to try to have some additional capacity available here both for those directly involved in the Rutgers research and collaborative activity but also some...outside...for the community. The way I see this, it doesn't read this way...too strong a constraint...

Lederberg: I see that things are still somewhat fluid there, and I would agree that you're going to need a little while to let that shake down. I hope it will be possible to get some clear line of policy from Rutgers that we can use as the premise of our own discussions at the site visit. That's something we just must not be confused about. I think we can live with a wide variety of alternative formulations of it, but we've got to have a clear picture of what it is that we're going to present to the site visit.

Amarel: Both during the site visit and in the material I've received so far that it's very clear that the group here at Rutgers is perceived as being a very important integral part of the national community around SUMEX-AIM and that this is a single enterprise and whatever we're going to come up with in terms of either collaborative mechanisms or cycles, it's got to be within the framework of the national AIM project... In terms of collaboration, dissemination of information and AIM workshops, the thing that still has to be worked is exactly how we're going to handle the question of the computing facility...capacity going to be available for local users...and even more general use.

2) SUMEX Site Visit Plans

Lederberg: I think you've touched on the two main issues and it's quite possible that the questions of functional collaboration will be more important than the issues of cycle availability. We've talked about that before, and I think we'd better have some more conversations among ourselves in the next couple of weeks and get that clarified further. If it's possible for you (Saul) to be at the site visit, I think it will be easier to explain these matters. In any case, I think we will have to have some more conversations about that. It's terribly important that we have a clear and unified policy on these matters; that's what AIM's all about and I'm sure that's what the study section will be looking for.

I really would like to minimize the burden (AIMEXEC attendance at site visit). I know the site visit needs information and so forth, but the notion that every user has to go all around the country every time some question comes up about the resource is not very satisfactory either. I will try to get one or two other representative users that are convenient. Wipke would be an obvious instance of that which would be no great trouble to bring here and would round out what we have on the clinical side with the membership of the present group. If any of you have any other thoughts about that or any particular users that we should try to highlight for the site visit, I would be glad to have your opinion about that.

Baker: If Jack could possibly come, I think that would be very valuable.

Lederberg: ...I think your project is going to be quite instrumental in the examination of the network utilities and what its benefits will be.

3) Distribution of Projects between SUMEX and Rutgers (included 1 above)

Baker: Saul has to deal with me and Rutgers. I think we ought to have another Executive Committee meeting before the site visit so that everyone knows how firm we are and what our future potential is in terms of distribution of projects between SUMEX and Rutgers.

10) Next AIMEXEC Telecon (3 and some 2 from 10/6 agenda)
(Thursday, 10/27, 10 a.m. PDT x approx. 1 hour)

3) (continued)

Feigenbaum: In preparation for talking about 6, I wonder if we could get an update on the distribution between Stanford and the AIM projects on the machine. Tom, do you have a status report on that. Also, does anyone know the results of the other computer requests that were site-visited last spring having to do with Jack's project and MIT's project and how they affect our AIM community?

Baker: Jack's project was funded. The MIT project is funded, but we have no idea how they relate to SUMEX-AIM.

Feigenbaum: Do you think we would get any embarrassing questions about that at the site visit? (Baker=no, but I might.)

Lederberg: I think we do have to have some sort of answers, particularly in respect to the Pittsburgh effort.

Myers: We've been funded and we expect to carry on in much the same fashion as we have in the last couple years.

Lederberg: So the implication is that you will have a facility which will be used for...

Myers: No, we were not approved for a separate facility. We were approved on a continuation basis much as we've been going now.

Lederberg: Sorry...I knew you'd put in to try to get something you could use for a field demonstration, and we certainly agreed that that was not appropriate for SUMEX.

Myers: We're dickering with other ways to get that accomplished, but for the foreseeable future, I hope our relation to SUMEX will be exactly as it has been.

Lederberg: Warner, Jack. There may be some more things that we can do for you now than we were able to do until the recent upgrade, so maybe we ought to talk about that a little further. If you want to collect some ideas on how that might work, let's hear about them.

Baker: I'll dig into the MIT thing and see what exists there. When I talked to Martin very early on in his developing of his application, he was not interested at all in reaching into the SUMEX-AIM capacity to get his work done. The thing that has bothered me is that they are 100% dependent upon a machine owned by the Navy, and I don't know how stable...don't know anything about that environment.

Lederberg: What you are telling me is that NIH is not funding resource capacity on their part, so in that sense, there's not a coordination problem between different computer resources in this area. That point on something we didn't need to know about is quite relieving to me because I had an image in mind that was quite different from that and that they might indeed be getting funding for computer capacity and if so obviously it would be a little awkward if we didn't know anything about it. Well, that puts my mind to rest on those two points.

Rindfleisch: This is the data (usage split) based on the proposal which was last spring. I haven't done anything more recent lately. The total Stanford CPU time, for example, for last year was 2233 hours, and for the AIM community was 1035 hours.

Feigenbaum: The question I'm getting at is how much room is there on the AIM side for more projects? Do you feel we're heavily loaded on the AIM side or not?

Lederberg: Plainly, we're not. We have a policy of equal access to the machine and the Stanford guys get there and use that privilege to the fullest. We have a scheduling algorithm now that does not allow that to crowd out the AIM people, so if we're talking about

a fair aliquot on the AIM side, the machine should look half as full to an incremental AIM user as it does to an incremental Stanford user; and it will be at the expense of Stanford use and convenience, but I think we have an obligation to try to reach out to that, and that's implemented now in terms of the scheduling algorithm.

Rindfleisch: That's exactly the scheme that we tried to put together, ...The division is 40% for Stanford, 40% for AIM and 20% for Staff, and MAINSAIL is subsumed in the Staff part which includes all system work, administrative and MAINSAIL development.

Lederberg: Those are the access privileges, I think Staff is about on par with Stanford in terms of the extent to which they use it, maybe a little bit behind it.

Rindfleisch: Staff last year used 903 CPU hours. That's less than the AIM community. (AIM + Staff almost = Stanford)

Amarel: What is the prognosis for being able to get some new projects such as PUFF, for instance, from the point of view of technical response as augmented now?

Lederberg: I think there's room on the system to mount them, and as I say, to them the system will look half as crowded as it does to the marginal user from the Stanford side, and I think that's not bad. Only under the worst conditions will it appear to be particularly crowded. They will have the obstacles of long-distance communications, which of course the computer cycle doesn't help very much, but in terms of throughput on the programs and so forth, I think any incremental AIM user will find himself in pretty good shape. He'll certainly be in better shape than the typical Stanford user is at the present time given the loading and giving the schedule algorithm that we have in mind. We have a couple of projects waiting to be looked at. There's the PUFF thing from S.F. There's one you'll see that I don't know whether we want to act on finally or not now from Alan Lesgold at Pittsburgh. There are one or two others that I would put at the second tier that are coming up, and they will add more on the AIM side. I feel that at this point we should still be actively recruiting for additional AIM projects but also making sure that the ones we've got are adequately served.

...I think it's the Stanford side that knows how to get what it wants out of the system. That's going to have to be held in a queue in order to get a fair allocation. The other will happen and that will be a sign of our continued maturity. Right now we have gotten an augmentation and I think that the justification/rationale for that was to give us some room for expansion on the AIM side as well as to make life possible for the rest of the system. There are one or two projects that we have on AIM that don't use much resources and, I feel, are of marginal value, but I don't feel we have to go after them because we're not in that much of a crunch right now. It would be nice to be at a place where we had to make harder choices among projects. The question of administrative controls on usage is terribly important too, and we're going to have to learn how to monitor that. The fact is, no AIM project is living

over its budget of allocated resource, and the only place that's problematical is filespace. Certainly no AIM project was using more than its aliquot of computer cycles.

Feigenbaum: I think we ought to consider the quality of life argument in interaction as well as the total usage statistics. The quality of life on the AIM side right now is extremely good and on the Stanford side is quite difficult. I would not like to see the AIM community get itself into the same state that the Stanford community has gotten itself into with respect to quality of life on the machine. So we should be quite careful in letting new projects get admitted to the AIM side because pretty soon, as Tom points out, the load average on the AIM side will look very bad.

Lederberg: We've always been living on a tightrope trying to...those considerations. I still think we have to look for good projects and then make our decisions about what we take or don't take in the light of the quality of the effort that wants to come on. I certainly don't accept the proposition that we've got it good so find your own resources. You may have to disappoint some applicants but only after very careful and deliberate consideration of their merits.

Baker: Do you feel like the quality of the projects that are being supported are all real good projects--nothing weak in this whole group?

Lederberg: I think there are several weak projects. I don't think they consume much by way of resources so I haven't felt it's worthwhile to make an issue of them, but that's just one opinion. The whole point of this managerial structure was that it was not a question of do it the way Lederberg says; it was a question of allocating that responsibility to an executive and advisory committee, so I have followed their lead with respect to what does/doesn't go on. I'm not going to say any more about it unless we have a specific project that has to be acted on right now.

Baker: Are there any projects on the SUMEX-AIM system, in or out of Stanford, that have questionable quality?

Lindberg: ...you're also doing an experiment in community building and promoting collaboration, so I don't see any reason to take after any of those projects right now when they don't consume substantial resources and they have those other virtues.

Myers: Obviously, there's quite a wide range of quality of the projects. Some are much better than others and some might even approach the mediocre. From what Josh says, there hasn't been any reason to be oppressive against the poorer ones at the moment, but I would thoroughly agree with him that we don't want to adopt a stand pat attitude that where there are good projects we ought to receive their applications, examine them carefully, be objective... I think there's room for improvement and we ought to receive applications and test them very carefully as in any academic pursuit.

Lederberg: Let me expand just a little bit further. I've been watching the community usage statistics quite closely. Tom has generated beautiful project control information. I think we've got better here than in almost any other system of any kind I know about in

terms of being able to manage the resource utilization. I don't see any significant amount of the resource being wasted. There are projects that are given theoretical access that if I say them burning up an awful lot of time I would start raising some questions as to whether it was worth it but I don't see them doing that so that's why I've been quite modest in pressing on those points and that's exactly what I would tell the site visit. The actual allocation of usage is not in really terribly bad shape. In fact, it's in excellent shape except that I think the AIM side does deserve a little more...and facilitation and help to bring it up to par with what the Stanford guys can do. The Stanford group obviously knows very well what it wants, has access to a lot of information--it's right there--and even with the scheduling algorithm it's a larger group, a larger critical mass, people talking to one another all the time...new ideas, etc. It's no...plot; it's the fact that we have a very strong working group that is responsible for that degree of utilization. It's not as specialized as the other groups outside so a lot of new things are getting started and so on. I'm very proud of that. I make no apologies for that at all. I would like to encourage a wider variety of users outside to bring it up to par, and I think on the AIM side there is still room to do that.

Amarel: In the Advisory Committee meeting that we had this summer at the AIM Workshop, there was some discussion about various projects quality...and relevance to the entire business of AIM. It was felt there may be one or two projects that are sort of close to the uncertainty level as to whether they should continue, but on the whole, the main projects are very good.

Lederberg: The official finding of the Advisory Committee was let the status quo go for the time being. Don raised one question about Bob Lindsay's project. He was perfectly correct that that did require some re-focussing, and we've communicated with him about that to get something that might be tighter. The fact is, his resource consumption has been quite minimal so that illustrates the point that I was making earlier, but we will either get a more satisfactory/more concrete statement from him for further review or we'll shut him off altogether. I don't think that case represents a severe problem. That was the resolution of the Advisory Committee.

Lederberg: ...I agree with what I think you said, Saul, but I would have phrased it a little differently. It's the good projects that are on AIM being compared to projects that may or may not be as good from the outside, but I think we have enough resources to accept one or two other significant good projects and that we all will be better off if we can manage to do that. If we do, we may/may not have to bump one or two of the marginal ones. It's not clear to me that we have to do that either. In my own opinion, I do not think that there's a significant less than excellent quality consumer of resources who's getting in the way of somebody else coming on the AIM side. There are lots of problems on the Stanford side.--The scheduling algorithm protects the interests of the AIM group, and I think that's working very well from everything I'm able to see, that it's doing what it's supposed to do in that regard, maybe even leaning over backwards a little bit. Any new user who happens to come on can sometimes get essentially full ownership of the machine if he happens to be isolated at that time regardless of the

importance of what he's doing at that particular moment, but that's fine; that's just the way it should be. The way we designed the system was that as far as resource access was concerned that it was going to look like three different machines with a 40/40/20 split and then whatever's left over after there's been complete access on that basis can be used first come first serve so it isn't wasted. I think it is working that way, and I think that's an ideal arrangement.

Baker: ...It seems to me that the quality of the projects on SUMEX-AIM is good, and we don't need to be concerned about it at this time.

Lederberg: The basis of that was the mailing out of the annual reports with the detailed project summary. We jogged each of the members of the advisory committee to look at them again. I think that they did spend at least half an hour doing that just prior to the meeting and many of them before. We did not go over each project one at a time and get a vote on it and try to make...inquiry. In fact, we discussed whether to do it and the consensus was this is not a good time particularly since there's no operational requirement for it that is in terms of the competition for resources at that time. The discussion then was on the basis of that general overview--was there anything that did stand out. I tried hard to not just have a perfunctory blanket cover...the one thing that was brought up was this thing about Bob Lindsay. We do have resolution in minutes. ...I don't see any reason except the appearance of spinning wheels not to adhere to the procedure that we've had so far. At some time and I think preferably after we've gotten straightened out about our renewal and so forth maybe a more focussed re-examination of one or two projects at a time would be in order but more from the view of rather paternalistically trying to help them along and kind of relate them to what other people are doing and so forth than a terribly rigorous or critical examination. But we obviously do have to lay on some protocols for careful review, (approach agreeable to all)

5) Critique of '77 AIM Workshop

Amaral: I received a few letters, mostly positive...I think it was very good. I think there was a lot of information communicated to people, an excellent get together about what's going on scientifically in AIM. We had 5 working sessions that were good. We had the whole thing videotaped and at this point quite a bit of work has been done... One of the important things we decided on the last day of the workshop was collaboration between Rutgers and Stanford groups on looking at possibly editing the material and preparing the proceedings both in a form that would be on SUMEX-AIM and distributed hardcopy. We'll have to work that out but I think things are moving along this line.

Baker: The feedback I have is that by keeping it small it was much more effective than the previous workshops.

Lederberg: That's my reaction too. The only negative side of that is that a lot of our younger people on the projects and graduate students missed the opportunity to participate. The people at Rutgers, obviously, were very welcome there, and that did raise the question of whether we should/shouldn't try to alternate the location for it for essentially that reason from the East Coast to

the West Coast. I think that's something that's still up in the air. I've heard some enthusiasm around here that maybe next year we should have it on the West Coast and I suppose that's something we have to settle before too long.

Feigenbaum: Saul and I have spoken briefly about it. There was quite a bit of dissatisfaction here at Stanford with the limited attendance of Stanford people caused by the fact that many...to send them to Rutgers and the fact that Rutgers students and junior researchers get an unfair advantage in that they're able to attend all of these sessions. Many feel that the AIM Workshop is their annual meeting and they're being excluded from their annual meeting, so they volunteered to run a workshop next year. Jonathan King, a graduate student here on leave of absence working for the CS Dept. for a year, was willing to undertake the arrangements for the workshop so that is a possibility; it's at least feasible for Stanford to take on the great amount of work that Rutgers has been doing over the last few years on the workshops, but we have to decide pretty quickly mostly because dates are important and we have to reserve rooms/reservations/people's schedules.

Amarel: I don't know that I would call the Rutgers students attending unfair... There's a lot of effort to putting on the workshop both at the technical level and management and organization. One person is not enough; it requires a few people on some continuing basis working/interacting...various parts of the AIM community. I feel it's a good idea to make it a little more practical...documents. One of main responsibilities of that document is the area of the workshop and AIM dissemination. This is the major responsibility, not that of providing a site. ...it comes down to the question of how we're going to organize ourselves and divide responsibility. From our point of view, we are again prepared to have the workshop here. We are funded for that...for next year at the same level as now for a small workshop, and we have an excellent place, the Continuing Education Center at Rutgers, which is a good one both in terms of facilities and a good place for people to interact. The CEC is available to us June 25-30, 1978. We have mandate/funds to do it, but I'd like to discuss further with Bill.

Lederberg: I've certainly had in mind the same view that you've indicated about the appropriateness of the sharing of responsibility. One thing is that I feel we're trying to put together a national effort, and one of the prices of that is that you don't take on every responsibility and every burden at one place, so I'm delighted at the enormous effort that you and colleagues have been putting into it, and I think that plans having gone this far, let's agree that we will have the next workshop at Rutgers on the dates that you've mentioned and leave open the question for the future of an occasional alternation. The only other thing is that if we can just raise another \$5,000 that we could be able to cover another dozen graduate students maybe on some austere basis like they go on a charter flight or something of that sort that would go a long way to solving part of the problem and would still not be a major increment in the cost for the effort since I think there is some justice to their concern about it. We could explore, although the way the telecons are going is pretty marginal, is whether we can get a good long line communication and have at least the audio part of the discussions available at Stanford. We might look into that a little bit...

Amarel: You mean instantaneous, on line, with the possibility of people asking questions and so on?

Lederberg: Sure. Questions maybe ought not to be verbal; maybe we ought to set up the teleconferencing over the network for that kind of dialogue. That would also mean that we ask users to at least in some cases put their text files corresponding to their slides so they could be called up. This just occurred to me and I don't want to be too much of a zealot about that but it just possibly might work. If we can get the cooperation of the speakers, we can demonstrate that we don't always have to travel to do everything that's necessary here. Obviously it's much more intimate and there are more unplanned encounters and that's terribly important at the meeting. We'll make some further inquiries about facilitating graduate student participation.

Amarel: Excellent idea...but we should not increase the total number of invitees to the workshop. It was very important to keep it at a reasonable size. Even under these conditions we have to be a little careful... But, certainly if we have a little more money to bring graduate students to this area that would be excellent.there is another question. This year we decided to sort of separate the entire question of internal discussion and interaction and more dissemination, and last year we made an effort to have some activity especially a session on applications of AI...

b) New SUMEX User Proposal Review

b) PUFF

Lederberg: ...but it is an off campus centered activity and, therefore, belongs on the AIM agenda, ...about status of application and funding for it.

Feigenbaum: Application was site visited in June. There is supposed to be some council action about now... I suspect there's a fairly good chance of its getting funded. In any event, the technical work is proceeding at a slow pace. There are two parts to the project. One has to do with pulmonary function diagnosis from the use of spymeter data. The other has to do with suggestions to the attending physician about interventions in the use of a ventilating machine. There's a Ph.D. student in CS working on each side of that project. Both students are proceeding to work on their theses with the use of some of the private expertise of the San Francisco people even before there's been any funding. That's the current state.

Lederberg: I have a few comments. I recall now a little more precisely why it was not formally acted on. I felt some embarrassment about discussing this proposal in Ed's presence. I commented that I thought it would be better procedure to permit any member of the Advisory Committee who had any other reactions to communicate that either to me or Bill Baker. I have not had any such communication. It was in order to give opportunity for what might be more critical and more objective comment. That was one reason for this delay.

Baker: Regardless of the outcome of the application, I think it should be supported as a pilot project. If the students really do make progress on it then become a full-fledged project.

Lederberg: I have one question about that. If the proposal has been approved then I would suggest we approve access to SUMEX independently of whether it's funded or not which I think raises other kinds of questions. If there's any chance of its being disapproved, I would be quite loathe to contravene the affirmatively negative action of a study section without a closer inquiry.

Baker: This was a computer study section, and I think it represents again a review where two peers were not picked to go along on the site visit. I'm not really concerned about whether the study section approves it or not.

Myers: I have been considering this on the study section. I don't know all the details of priorities, but the project received very favorable consideration. My recollection was that there was no discussion of disapproval; the only question was how high up the priority list to place it.

Lederberg: Well, I suggest we go ahead and approve it, and if a disapproval comes through then I think we should reconsider it in the light of all the circumstances.

Amarel: Why should we consider this project as a pilot? All of our information about the project...numerous discussions already... it covers all the things we like to see in the AIM community in terms of applications of AI in medicine and relevance and possible contribution to the AIM community. Why do we have to wait for whatever kind of formal statement? It may be an issue of policy not to handle things like that, but I feel we should consider it for a regular project and I'd like to see it as part of the community.

Lederberg: I appreciate that comment and I agree with you. I think the reason it was put in this form was that at the time the application had just gone in for NIH support and so was a project pending from that standpoint. Now that it's becoming a definitive one, we should delete the word pilot in our consideration of it. I've heard unambiguously positive statements about it, so if there's no dissent, that one has been accepted.