

GRAPHICS MEETING REPORT

The second Network Graphics Group Meeting was convened at the Stanford Artificial Intelligence Lab at 6:00p.m. Sunday, November 21st. (Attendees are listed in the Appendix.) Jim Michener served as chairman, and I either volunteered or was volunteered to serve as recording secretary, with Karl Kelly's assistance in keeping notes.

An agenda was agreed upon for the meeting, covering three major topics: 1) reports on the experiments which had been set up at the July meeting, 2) prepared talks by attendees who had general points to raise about Network Graphics, and 3) specification of a "first-pass" graphics protocol. Before the reports were given, some general discussion was held on two important topics: the "context" problem (just how, in the Network sense, are graphics connections established, and who is supposed to do what for whom), and what might be called the "console types" problem (should there be a separate protocol for inherently static storage tube type devices and one for inherently interactive refresh type devices which have their own processors, or can we come up with some sort of continuous -- or layered -- single protocol which covers both). Both points were noted as being necessary to keep in mind for the protocol specification phase of the meeting, an apparent consensus emerged that a single protocol would be preferable, and the reports on experiments were turned to.

REPORTS ON EXPERIMENTS

RAND - UCSB

Eric Harslem of RAND and Ron Stoughton of UCSB reported on their experiment, which entailed use of the UCSB On-Line System (OLS) from RAND Videographics terminals. As demonstrated by a videotape which was shown, the experiment was successful. An RFC describing the simple protocol they used is forthcoming. As noted in their discussion and in the RFC, the experimental protocol is not being proposed as a Network standard. In addition to using OLS from RAND, a subsidiary experiment tested the sensitivity of the hook-up to variations in the size of the allocations (in the Host-to-Host Protocol sense) given at the RAND end. It seemed clear from the videotape of the same pictures being drawn at various allocation levels that larger allocations allow for noticeably smoother

"drawing" at maximum allocation, the picture essentially appeared all at once, whereas at minimum allocation, NCP-NCP overhead was sufficiently large that the picture appeared a portion at a time.

SDC - DMCG

An experiment intended to input tablet data collected at MIT Project MAC's Dynamic Modeling/Computer Graphics Group's PDP-10 to a character recognizer package at SDC was reported on by Jean Saylor of SDC and Jim Michener of DMCG. Problems ranging from hardware/software difficulties at both ends (and in the middle) to time zone-induced system availability conflicts retarded the experiment's progress, although some transmission of data has been achieved.

ILLINOIS MULTICS

Also plagued with problems was the attempt to drive a console at U. of Ill. from the Multics Graphics System. This experiment was reported on by Jack Bouknight (Illinois) and Ed Meyer (Multics). An NCP bug at the Multics end and a machine switch at the Illinois end combined to prevent the carrying out of the experiment.

DIGRESSION

During his report, Bouknight expressed concern as to whether the Network as a whole is as yet sufficiently reliable to support graphics work. As the ensuing discussion focused on the frequent unavailability of a host other than Multics, I feel that it is within my province to draw the curtain of anonymity over it without prejudice. However, I feel that mention of the discussion need not be suppressed as well, in view of the fact that most of the attendees shared Jack's concern. The apparent consensus, reached after considerable conversation, is that the present reliability level of the Network server hosts is not crippling to graphics work, but can be quite hampering.

SEX - NIC

Jon Postel (UCLA) and John Melvin (SRI) gave the last experiment report, on an attempt to make an IMLAC on the SEX system look like a local NLS console at the Network Information Center. The experiment has not yet been performed, but UCLA has ordered the necessary equipment to modify their IMLAC.

PRESENTATIONS

Most of the speakers who gave prepared talks responded favorably to my plea for abstracts, probably out of kindness, but perhaps out of fear of (threatened) garbling. Authors' abstracts are in quotation marks in the following section.

PLASMA PANEL DISPLAY - Dave Liddle

"The Owens - Illinois DS-1 terminal will be available to Network users who request them through ARPA. The display module is the OI 512X512 line plasma panel; the processor is a 16 bit, 4K machine with modem; ASCII keyboard, and optional tape cassette. Simple software (character and vector generators, etc.) will be provided. If orders can be assembled by 1 January, deliveries will begin this summer."

LANGUAGES FOR GRAPHICS ATTENTION HANDLING - Ira Cotton

"Available languages for programming the processing of operator inputs to a computer graphic system were organized into functional classes and briefly surveyed. Some of the problems associated with providing this facility in a multi-computer graphics system (such as the Network) were discussed, and a new approach was presented. This system, implemented by Univac for one of its systems, employs an interpretively executed command language to direct attention-handling in the remote graphics controller. The commands of the language were outlined, and some program fragments illustrated."

"INTERACTIVE" GRAPHICS ISSUES - Ken Pogran

"The purpose of this talk was to raise a number of significant issues we must face in the development of a Network protocol for interactive graphics. While the bulk of the work at this second graphics meeting dealt with a protocol for "static" or "storage-tube" graphics, it is appropriate that we begin to think about the problems we will encounter in the development of an interactive graphics protocol."

"The issues raised included: 1) the nature of graphical interaction, 2) various possible hardware/software configurations which might be employed, 3) computational capabilities required at the server and user host sites, 4) the nature of a graphical data structure suited to a wide range of applications, and 5) the nature and treatment of graphic inputs for a generalized interactive graphics system."

PROTOCOL FOR THE OLS EXPERIMENT - Ron Stoughton, Eric Harslem

"A short presentation was given describing a graphics protocol used to interface the RAND Videographics System to the USCB On-Line System. A video tape of alive demonstration of the experiment [had also been] presented. An RFC describing the experiment and protocol in detail will be issued in the near future."

CONNECTION CONSIDERATIONS - Andy Moorer [Abstracted by M.A.P.]

The topic was started succinctly as "how this thing should work." It was proposed to use the Telnet Protocol for simple graphics (i.e., when device dependent codes are being transmitted), with the addition of Telnet control codes for Enter graphics Mode, Leave Graphics Mode, and Console Type being necessary. For complex graphics (i.e., when an intermediate form is being transmitted) it was proposed that an additional socket pair be employed.

CONNECTION TYPES - Jim Michener [Abstracted by M.A.P.]

There are at least three types of graphics devices which may be connected over the Network: "simple" (ARDS-like), "smart" (IMLAC-like), and "powerful" (E&S-like). There are three kinds of processing involved: applications packages (A), graphics packages (G), and conversion to device-specific codes (C), potentially from an intermediate form such as the "Network Standard Graphics Stream" discussed in earlier RFC's. There are also three places where each kind of processing may be performed: at the graphics device itself, at the local host (which may not be able to help if it's a TIP), and at a remote host (OR HOST). This should lead neatly to some sort of 3X3X3 matrix which depicts the sorts of connections we want to support, but I don't know how to draw it.

The talk leaned heavily on blackboard pictures of specific connections, but for purposes of this report, I'll try to summarize the situation in words. For all simple devices, C must be performed "elsewhere"; if the simple device is on the Net via a TIP, C apparently must be performed either at the remote host (RH1) where A and G are, or at some other remote host (RH2) (which offers, say, the Data Reconfiguration Service). Further, negotiations for C may have to be performed by RH1 on the TIP's behalf. Still more complications result from the possible desirability of including an additional application (A') and/or an additional graphics package (G') on RH2. If the simple device is on the Net via a full-fledged local host (LH), then A, G, and C can each potentially be performed at LH or RH1 -- or RH2 for that matter ("ship it to an E&S for clipping").

In the case of smart devices, C can potentially be performed at the device itself - - although the TIP may not be able to furnish the extra socket pair which one would want in order to handle such cases cleanly. Finally, powerful devices can do G internally but we may well wish to do A and G over the Net. (Again, how the TIP would handle such cases was not clear.)

Jim had presented this discussion for the expressed purpose of getting attention focused on the "ends" of the protocol pipeline before the meeting became totally concerned with the contents of that pipeline. We responded in the only possible manner:

CONNECTION PROTOCOL COMMITTEE

A committee was designated to formulate a Graphics Connection Protocol, the protocol to play an analogous role to that of the Initial Connection Protocol with respect to the Telnet Protocol. There was a clearcut consensus that only device-specific codes should be transmitted over Telnet connections unless the committee uncovered overwhelmingly convincing arguments to the contrary. The committee consists of Michener, Bouknight, Harslem, and me. Will Crowther of BBN will be invited to join the committee to furnish TIP representation and expertise.

GRAPHICS RESOURCE DOCUMENTATION

Before turning to the protocol specification, it should be pointed out that most attendees felt that Resource Notebook-like documentation on Graphics should be prepared. Postel volunteered to coordinate this effort. Hosts should have drafts submitted to him, and he will see to getting them published as new portion of the Resource Notebook. Format considerations were not discussed, but assumedly the format should imitate that of the main Resource Notebook sections. Call Jon if you have questions (213-825-2368).

THE PROTOCOL

At the outset of the main protocol discussion, it was agreed that a committee would be established to resolve those issues on which a consensus could not be reached at the meeting, and to prepare a draft of the protocol for distribution to the NGG by year's end. Members of the committee are Michener, Meyer, Kelly, Cotton, and Liddle.

ASSUMPTIONS

The following assumptions were agreed upon:

1. There shall be a "virtual screen" and a Standard Graphics Stream.
2. The origin is in the center.
3. Coordinates are signed, 2's complement fractions (-.5 to +.499).
4. The Standard Graphics Stream will consist of 8-bit bytes initially, coordinates are two bytes. (A "set coordinate size" operator will be introduced if and when needed.)
5. Network ASCII will be used for text output, with default to upper case where necessary. Control characters are, for the time being, site specific.
6. Where appropriate, operators shall have "absolute," "relative," and "local" (to a subpicture) modes.
7. The protocol will be organized on a "levels of complexity" basis, with level 0 comprising operators for simple picture drawing, level 1 comprising operators for one level of subpicture definition ("macros", or loosely, "subroutines") and level 2 comprising "viewport" and "window" type operators.

Note that the discussion dealt specifically with graphics OUTPUT. The Protocol Committee was also empowered to prepare recommendations for an input-side protocol, but first priority is to be attached to the formulation of an acceptable output-side protocol.

OPERATORS

As the Protocol Committee's draft is not immediately available, the following list of low-level operators (the syntax and semantics of which were discussed at length during the meeting) may be of interest here:

1. Erase and reset to origin. This operator causes the screen to be erased and the beam to be positioned at the 0,0 (virtual screen center) point. A new picture is started.
2. Move. No line is drawn the beam is positioned to the specified x, y position. There are specific operators for "move relative", "move absolute" and "move local" modes.

3. Draw. A line (of the current "linetype" -- see 5, below) is drawn from the present beam position to the specified x, y position. Modes are as with move. Treatment of the "off-screen" condition is at the displaying host's option.

4. Point. Display a point at the specified position. Modes are as with move.

5. Line type. Draw lines of the specified type until further notice. Currently defined types are solid (0), dashed (1), dotted (2). If a requested type is not implemented, default to the next-lower-valued type. After an "erase", type is solid until changed.

6. Line intensity. Requests line intensity to be as follows: 0 = off, 128 = normal, 255 = brightest, intermediate values = map appropriately. After an "erase", intensity is normal until changed.

7. Text. Cause display of a specified number of specified (Net ASCII) characters. There are specific operators for "return beam" after last character (to position before text display) and "leave beam" (wherever it ends up). Size is to be whatever the displaying host considers "normal". Treatment of "right-hand margin" and ASCII controls is host-specified at present. (A character size operator may be specified later.)

8. Escape. If the console is of specified type, pass a specified number of bytes directly to it.

Operators for viewports and subpictures were also discussed. Bouknight and Kelly prepared an BNF treatment of all points discussed, which will appear in the Protocol Committee's draft.

OTHER BUSINESS

The remaining technical discussion dealt with graphic input, on a rather general level.

Michener extended the attendees' thanks to Andy Moorer for having hosted the meeting.

Cotton volunteered to host the next meeting at Mitre, Washington, in mid-April, at which time we hope to have had enough experience with the connection protocol and first-pass output protocol to agree on a "final" statement of them, and to have done enough thinking about the input side to specify a first-pass protocol for it (unless the Protocol Committee manages to do so first)

APPENDIX - LIST OF ATTENDEES

Marshall Abrams, Ntl. Bureau of Stds.

Jack Bouknight, U. of Ill.

Jackson T. Cole, Rome Air Development Ctr.

Ira Cotton, MITRE

Daniel Debrosse, UTAH

Eric Harslem, RAND

Karl Kelly, U. of Ill.

David Liddle, Owens Illinois

John Melvin, SRI

Ed Meyer, MAC

James Michener, MAC

James Moorer, SAIL

Hamid Naficy, UCLA

Mike Padlipsky, MAC

Ken Pogran, MAC

Jon Postel, UCLA

Jerry Powell, MITRE

Jean Saylor, SDC

Ron Stoughton, UCSB

Elaine Thomas, BBN

Howard Wactlar, Carnegie-Mellon

Bill White, SUHP

[This RFC was put into machine readable form for entry]
[into the online RFC archives by Kelly Tardif, Viagénie 10/99]

