

SOME THOUGHTS ON SRI'S PROPOSED REAL TIME CLOCK

Re: NWG/RFC's #28 and 29.

The addition of a clock in one or more of the network HOST's seems to be very desirable since it (or they) would allow user-oriented message delay measurements. Our present network measurement facilities do not include any internal HOST delays, and these delays may be an appreciable portion of the total delay encountered by a HOST-to-HOST message transmission. We may find that an extension of our "Trace" capabilities to include internal HOST delays would be an appropriate mechanism for utilizing such a clock. Such usage would require a clock at both the source and the destination of the message, although such clocks would not have to be particularly accurate nor synchronized. Other tests, such as the absolute overall message delay from HOST A to HOST B would require synchronization of the two clocks.

A reasonable specification for the SRI real-time clock would seem to include a resolution of about 1 msec., an accuracy of about 1 part in 10^7 (so that two such clocks could maintain reasonable relative accuracies over periods of many hours), and a range of about 24 hours. A crystal controlled clock should easily meet these requirements at a moderate cost.

The choice of the mechanism by which the HOST can read the clock appears to be of concern also. The 1 msec. resolution may require that the clock be entirely hardware (as opposed to a core location which would be incremented at each clock pulse), and therefore the clock may require some rather complicated interface circuitry.

At UCLA, we presently have two clocks on the Sigma 7, and one of these has a resolution of about 2 msec. which might be usable for some internal HOST measurements. However, it does not have the long term accuracy for the absolute measurements mentioned above.

[This RFC was put into machine readable form for entry]
[into the online RFC archives by Richard Ames 1/98]

