USER-ORIENTED NETWORKS: A SERIES
PART II. INTER-USER COMMUNICATION AND ITS IMPLICATIONS FOR NETWORK TOPOGRAPHY

Wm. Orchard-Hays

July 1975

WP-75-82

TABLE OF CONTENTS

The Problem Area 1

The Problem as Seen in an Extension of Conventional Networks 2

Restructuring of the Network; A Node Identification Scheme 3

Functions of GRPs and CONs; Protocols and Identifiers 5

Extension to Broadcast and Mail-Box Service 9

Figure 1. An Initial Set of Symbols 11

Figure 2. Nodes in a Partial Network 12

Figure 3. A Hypothetical Network 13

THE PROBLEM AREA

In Part I-B, three types of networks were identified: commercial, cooperative, and user-oriented. It is the last which is of long-range interest to IIASA. A possible fourth type--data base networks--although of great interest to IIASA, need not be considered separately for the problem discussed here.

The usual arrangement for a commercial network permits many users to connect to a single central computing system and to utilize whatever services it provides. One of these services is the ability for one user to send messages to any other user who is on-line. This is frequently extended to provide a "mail-box" function for users not on-line at the time, and also to broadcast messages to all users or to leave "mail" for all users, or for all users of a group. The important point to note is that these services are functions of the central computing system.

In a cooperative network, the user first accesses the network and then, with network protocols, connects to a particular central computing system. From then on, the situation is no different from a commercial system. The network's central control system may have a limited ability to broadcast emergency messages, but users can only communicate with other users on the same central computing system.

In a user-oriented network, the above arrangement negates one of the principal purposes of the network, i.e. the ability for users to communicate easily with one another even though they may be using different central computing systems or even none at all. In order to analyse this problem easily, we need to invent some succinct symbology. Abbreviations and symbolic geometrical forms will serve to begin with, as shown in Figure 1.

THE PROBLEM AS SEEN IN AN EXTENSION OF CONVENTIONAL NETWORKS

Figure 2 shows part of a possible network. For the present, assume only one central computing system, SYS I, which has a high-capacity modem MOD I. There may be users connected locally as shown by S1 and S2. S1 and S2 can communicate with each other through SYS I in standard fashion.

One of the high-capacity ports of MOD I connects to a distant concentrator, CON 1, in Region 1, via high-capacity, dedicated lines. CON 1 serves users in two nearby areas, Area A and Area B.

In a conventional commercial network, the grouping computers GRP A and GRP B would not exist; direct telephone connections would be made on demand from users A1, A2, A3, B1, B2, B3 to CON 1. The printer at A3 and the tape unit at A2 would probably be run off separate lines when needed. Printers are often connected over long distances via standard long-distance telephone calls. (At least this is true in the US.) However, frequently, dedicated lines are used and, for high speed operation, higher capacity lines are needed.

In a cooperative network, something like the grouping computers are used when a using organization has a number of terminals in a single building or complex. However, the purpose of these units is to localize certain network functions and not to handle either computing tasks for the attached terminals or communication between them.

In a user-oriented system, the grouping computers must handle more complex network tasks and also duplicate (or better, replace) certain functions of a central computing system, particularly inter-user communication. Conceivably, this could be done at the concentrator level, but, in either case the concentrators must have additional switching logic over what is required in a commercial network. In a conventional arrangement, CON I is only concerned with identifying messages between each of its attached terminals and the central system, and routing them accordingly. The function of the grouping computers changes this considerably.

Suppose GRP A has the ability to forward messages between any of its terminals and SYS I (i.e. via CON 1) and also, intermixed with these, messages among its terminals. GRP B can do the same for B1, B2, B3. Now, if A1 wants to send a message to B2, GRP A must recognize this and route the message to GRP B but again via CON 1. Hence, CON 1 must know the difference between a message from GRP A for SYS I and from GRP A for GRP B.

Now suppose there is another concentrator CON 2 in a Region 2 with an attached group GRP C: C1, C2, C3. If A1 wants to send a message to C1, GRP A must route this to CON 1. If lines exist between CON 1 and CON 2, then CON 1 must recognize the routing from CON 1 to CON 2 for forwarding to GRP C and C1. If such lines do not exist, the message must be sent to SYS I for retransmission. However, since the message originated as an inter-terminal message (A1 and/or C1 may not even be logged in to SYS I), this creates imponderable questions with respect to SYS I, CON 1 and CON 2. Note that some of the same problems exist with respect to S1 and S2.

RESTRUCTURING OF THE NETWORK; A NODE IDENTIFICATION SCHEME

Considering the above analysis, the most sensible arrangement appears to be to have lines connecting all concentrators (not in all combinations) and to regard the concentrators as belonging to the network, not the central system. This has the further advantage that not all concentrators need to have lines to the central system, and conversely, different concentrators can connect to different central computing systems. Furthermore, the CONs and GRPs are inherently the same kinds of devices, differing only in line and switching capacity.

Figure 3 shows a user-oriented network such as IIASA might create. A new type of node is introduced, the gateway:

GWY 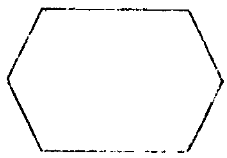

This is a system which is a SYS on two different networks. Connections between users on different networks are possible via a GWY but not with the same flexibility as within a network.

Every message must now carry "to" and "from" addresses. For simplicity and illustration only, let us use the following identification scheme.

Maximum of 26 CONs: A, B,...,Z
Maximum of 26 GRPs/CON: a, b,...,z
Maximum of 99 (TER,PRT,TAP,RDR)s/GRP: 01,02,...,99
Maximum of 9 SYSs/CON: 1,2,...,9

Thus all terminals have a 4-character address which identifies the GRP to which they are attached and the CON to which the GRP is attached, for example:

Ac21 is TER 21 on GRP c of CON A.

Zero can be used for special situations among CONs and GRPs, e.g.:

Ab00 is GRP b on CON A
A000 is CON A itself

Systems (including GWYs) require only a 2-character address since they are always attached to CONs. However, a SYS might include several software systems or operating modes which can be identified by the extra two characters. Thus,

B100 is SYS 1 attached to CON B
B1CM might be the conversational monitor system on B100

The distinction between a CON and a GRP may sometimes evaporate. This can occur in two ways: either TERs (or PRTs, TAPs, RDRs) may be attached directly to a CON; or, a GRP may play the role of a CON. In the first case, either 0 or some special letter can be used for the second address position. In the second case, the GRP must be elevated to the rank of a CON and then the convention above applied.

A more difficult situation exists for terminal equipment attached directly to a SYS. In this case, it may be necessary to go through SYS to communicate with them, probably at a different level of protocol. Something similar to this will almost certainly be required for communication to a unit on a different network via a GWY.

FUNCTIONS OF GRPs AND CONs; PROTOCOLS AND IDENTIFIERS

We can now describe the required functions of GRPs and CONs and the nature of network protocols and message identifiers. No attempt will be made to describe these down to the transmission or line-signal level. It is assumed that this has been worked out. Likewise, there is no consideration here of routing--which is assumed fixed--or of line-switching versus packet-switching. In fact, packet switching is assumed in the sequel but the usual form of packet-switching analysis appears incompatible with interactive network usage, or, at least, extremely inefficient and frustrating. A message unit of, say, 256 characters is extremely long for interactive use. Frequently, responses consist of no more than half a dozen characters (sometimes only one!) and it is virtually impossible to input more than 80 characters or get back more than 120 at a terminal.

Assume GRP-a attached to CON-b is in operation. When first brought on-line GRP-a must query CON-B to see if it is in operation and remember the answer. (Whenever CON-B comes up or goes down, all its GRPs which are in operation must be signal.) Now suppose TER Ba12 is turned on (dialed-up or whatever). It may be necessary for Bal2 to identify itself since it may be connected to a port at GRP-a in some random fashion. However, we will assume this has been done and that GRP-a knows that Bal2 is active and on which line.

The first thing that GRP-a must do is require Bal2 to give a userid, password, and possibly an account number. Note that these do not belong to the terminal Bal2 but to the human user at the terminal. GRP-a must record this current relationship in a list of active users. Note also that this decentralizes access to the network since user identification is at a local level. The implications of this arrangement need considerable study since it is decidedly different from that of conventional networks in which a user is recognized no matter where and from what equipment his log-in originates.

Once the identification protocols are satisfactorily completed, if CON-B is not up, then GRP-a should immediately notify Bal2 since it severely limits network usage. If Bal2 persists, then he (i.e the user at Bal2) can only communicate with other Bann terminals who may be on-line.

Suppose now CON-B is up and Bal2 wishes to communicate with a user on a different GRP. Suppose this second user is actually on-line at Bc15. Neither the user at Bal2, nor GRP-a, nor (probably) CON-B knows this; only GRP-c knows it. But to get to GRP-c we need the address Bc00. Hence it is clear that the userid's must be coded in the same style as TERs but distinguishable therefrom. If we use two letters in the last two positions--AA to ZZ--then as many as 676 human users may have accounts at any GRP, which should be sufficient. Hence the user at Bal2 might be identified as BaJS. The user he wishes to contact might be BcHB. Only at the local GRP level does BcHB get translated to Bc15 and BaJS to Bal2. This also indicates a way to relax the identification protocol restrictions.

Suppose BaJS is visiting BcHB and wants to log-in at Bc15. GRP-c can query GRP-a (if CON-B and GRP-a are up) to see if BaJS is an authorized user. If this is allowed, then a GRP's active list may contain userid's which do not belong to it and terminal assignments which are not hooked to it. The latter occurs because, in the example above, GRP-a should be notified that BaJS is at Bc15 so any messages from other users may be forwarded. The extent to which such temporary exchanges can be permitted may have to be limited, since it may use up extra storage at the GRPs. However, this is not the case if each GRP has room to record a terminal against each of its users and a user against each of its terminals. Processing time may increase since, in some situations, a "foreign" user must be searched for in the list of terminals rather than directly addressed in the list of local users. Still, this extra work is relatively trivial. More serious is the extra work in retransmitting messages and the extra loads on transmission lines. However, this situation only arises with inter-terminal messages and not with TER-SYS messages (discussed in Part III).

One other requirement must be noted in connection with the above. If a GRP receives a message already addressed to a terminal, rather than a user, it must simply pass it through, perhaps after some checking. For if GRP-a receives a message for BaJS who is temporarily at Bc15, GRP-a must retransmit it with the address Bc15 and not BaJS (which would cause it to be sent back again). This can lead to confusion in cases where GRP-c has gone down and come back up and BaJS has moved to another terminal. Such confusion can be avoided by putting secondary addresses on retransmitted messages, like a "care of." For example suppose AaWM working at Aa11 sent a message to BaJS who is temporarily at Bc15. The following sequence of identifications and transmissions would take place.

1. Aa00 receives a message from Aa11 addressed to BaJS. It looks up Aa11 and finds that TER in use by AaWM. It then constructs the "to/from" identification BaJS/AaWM and sends the message to A000.

2. A000 receives the message from Aa00 addressed to BaJS. A000 must have a routing table to all other CONs. It looks up B000 and finds it is a neighboring CON with a direct connection. Hence, it forwards the message to B000.

3. B000 receives the message from A000 addressed to BaJS. Since this is a user in one of its own GRPs, it forwards the message to Ba00.

4. Ba00 receives the message and looks up JS. This user exists but is currently on Bc15. (If BaJS were unknown or not on any TER, Ba00 should send an appropriate message back to AaWM.) Ba00 changes the to/from identification to Bc15/BaJS/AaWM and sends the message back to B000.

5. B000 receives the message now addressed to Bc15. Since this is a TER on one of its own GRPs, it forwards the message to Bc00.

6. Bc00 receives a message for Bc15. Since this is a TER, not a user, it does two things:

   a) It looks at the next section of the address and finds BaJS.

   b) It looks up Bc15 and finds it in use by BaJS.

   Since the userid's match, Bc00 sends the message to Bc15 as coming from AaWM.

   If Bc15 is not in use by BaJS, Bc00 must look through its TER list to see if BaJS is there.

If it is, say at Bc12, then Bc00 does two things:

a) Sends a message to Bc12 as coming from AaWM.

b) Sends a message to Ba00 to correct its list for BaJS.

If BaJS is not at any TER hooked to Bc00, Bc00 again does two things:

a) Sends an appropriate message to AaWM.

b) Sends a message to Ba00 to correct its list for BaJS.

Note that all the above actions are completely deterministic and can be readily flow-charted for CON and GRP control programs.

Although not previously mentioned, it will become evident in Part III that inter-terminal (user-user) messages must be specially flagged at the GRP level, that is, something like an ATTN key from TER to GRP is required to initiate such messages. The reason is that, in normal interactive use between a TER and a SYS, all unflagged messages from the TER must be assumed destined for the SYS. The reverse is unnecessary.

EXTENSION TO BROADCAST AND MAIL-BOX SERVICE

The extension of the above scheme for broadcasting messages is trivial. One obvious expedient is to reserve, say, the identification z and 99 to mean "all GRPs" and "all TERs," respectively. This is not appropriate at the CON level unless there is a predetermined sequencing since the message might circulate through the network indefinitely. However, either such a sequencing can be defined or a message for the entire network can be duplicated for each CON.

A more important consideration is that of authority to broadcast messages. It may be inappropriate to allow any user to broadcast to all others. Perhaps a user can broadcast to all users in his group. It may be necessary to have a monitoring user assigned to each GRP and each CON. Then the network might allow a GRP monitor to broadcast to all GRPs under the same CON, and a CON monitor to broadcast to any other CON. Such policies need not be formulated here; it is clear the identification scheme provides sufficient flexibility to implement them.

The matter of mail-box service has two requirements. First, an additional addressing flag is needed to indicate that a message which is not immediately deliverable should be held some predetermined length of time. If the user logs in during this period, the message is delivered. Second, storage is required for held messages. The obvious place to do this is in the GRP to which the user belongs. The feasibility of this depends on the storage capacity of the grouping computers. The options for grouping computers and other non-network functions which they might perform will be the subject of a later discussion.

TER	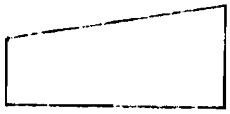	A remote terminal with keyboard input and either typed or video-display output. TER is also used generically for RDR, PRT or TAP.
RDR		A remote card reader. It may also be a card punch. If this distinction is important, it is denoted by RDR/PUN.
PRT	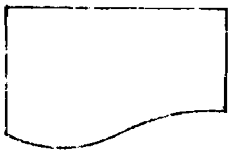	A remote line printer.
TAP	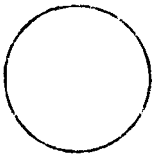	A remote tape unit.
MOD	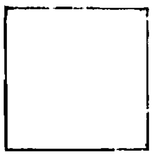	A telecommunication modem. For any remote unit, a MOD is implied, even if not explicityly shown (as it usually will not be). In a portable terminal, it is built into the terminal.
CON	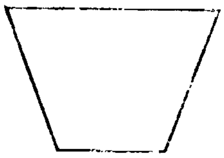	A concentrator.
SYS	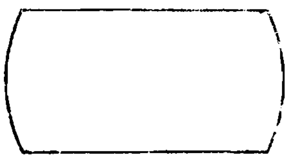	A central computing system
GRP	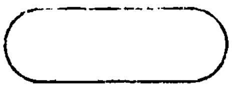	A grouping computer. May also imply the terminals hooked to it, when their distinction is unimportant.

FIGURE 1. AN INITIAL SET OF SYMBOLS

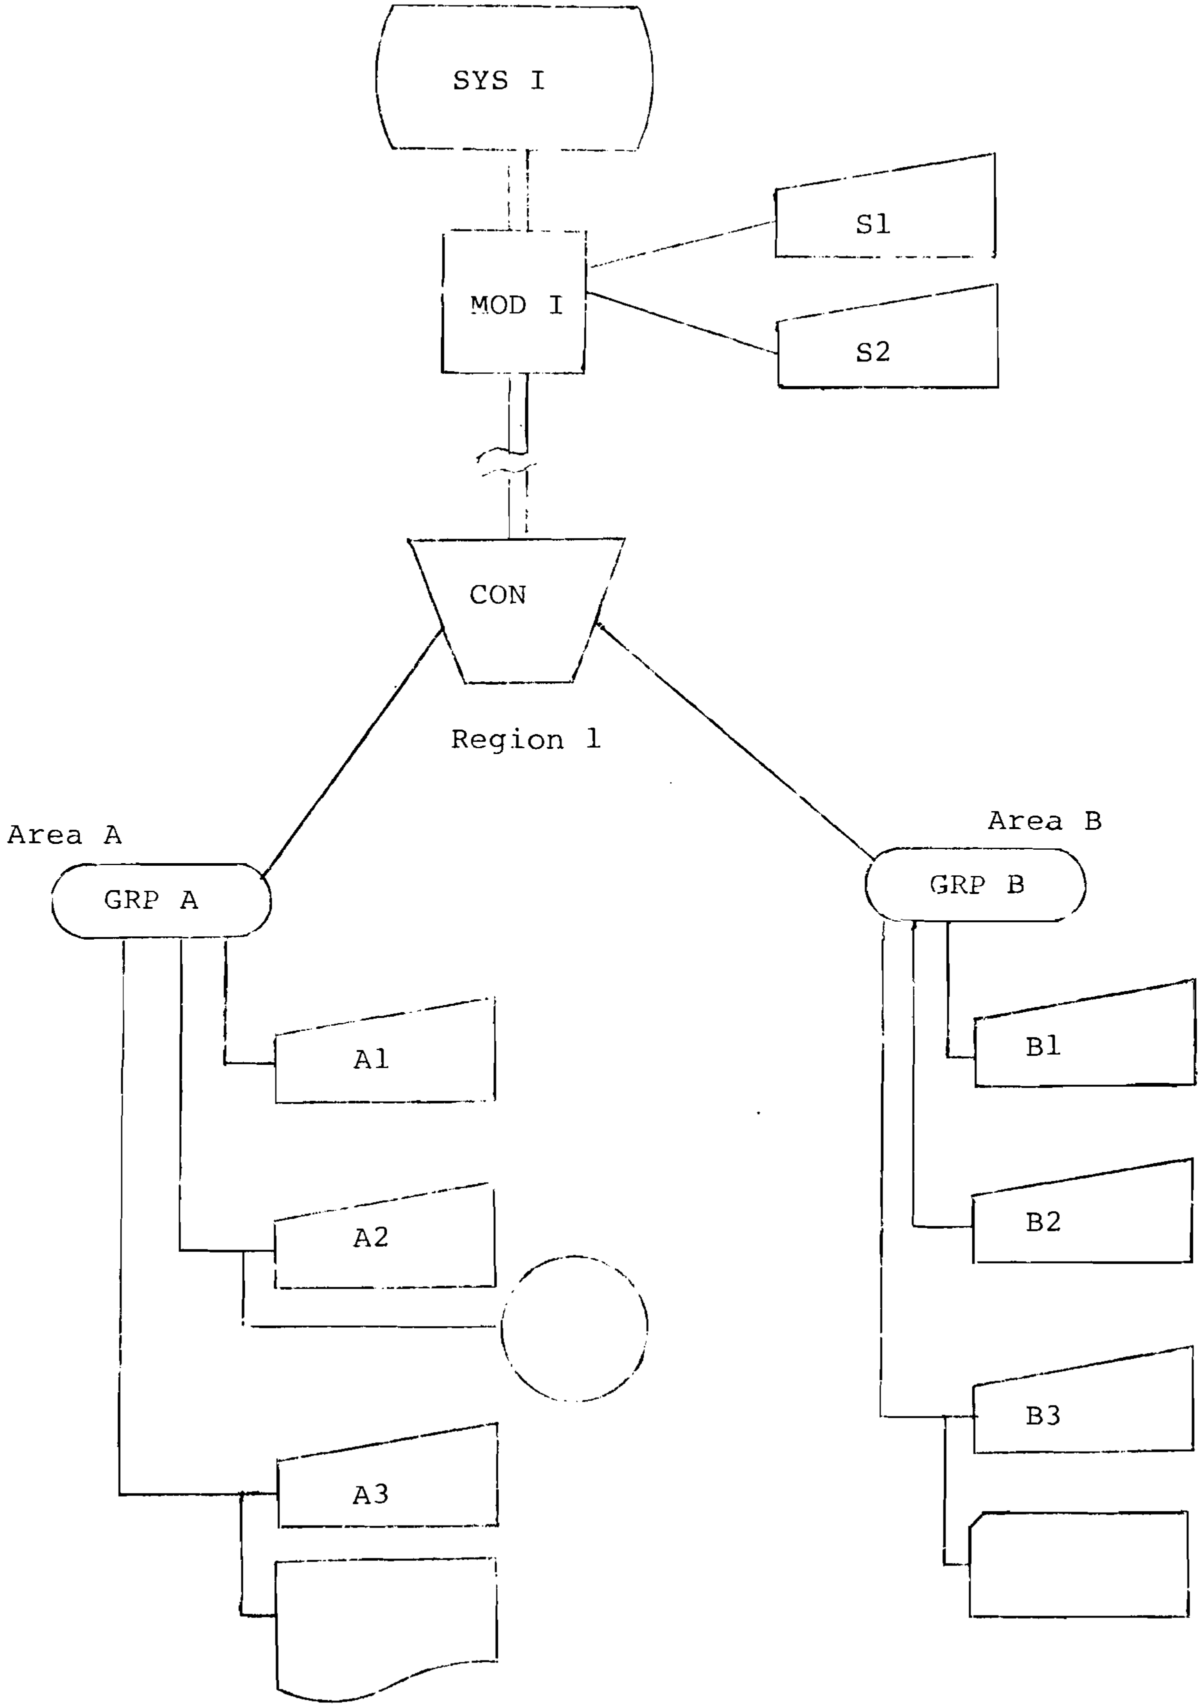

(Remote modems not shown)

FIGURE 2. NODES IN A PARTIAL NETWORK

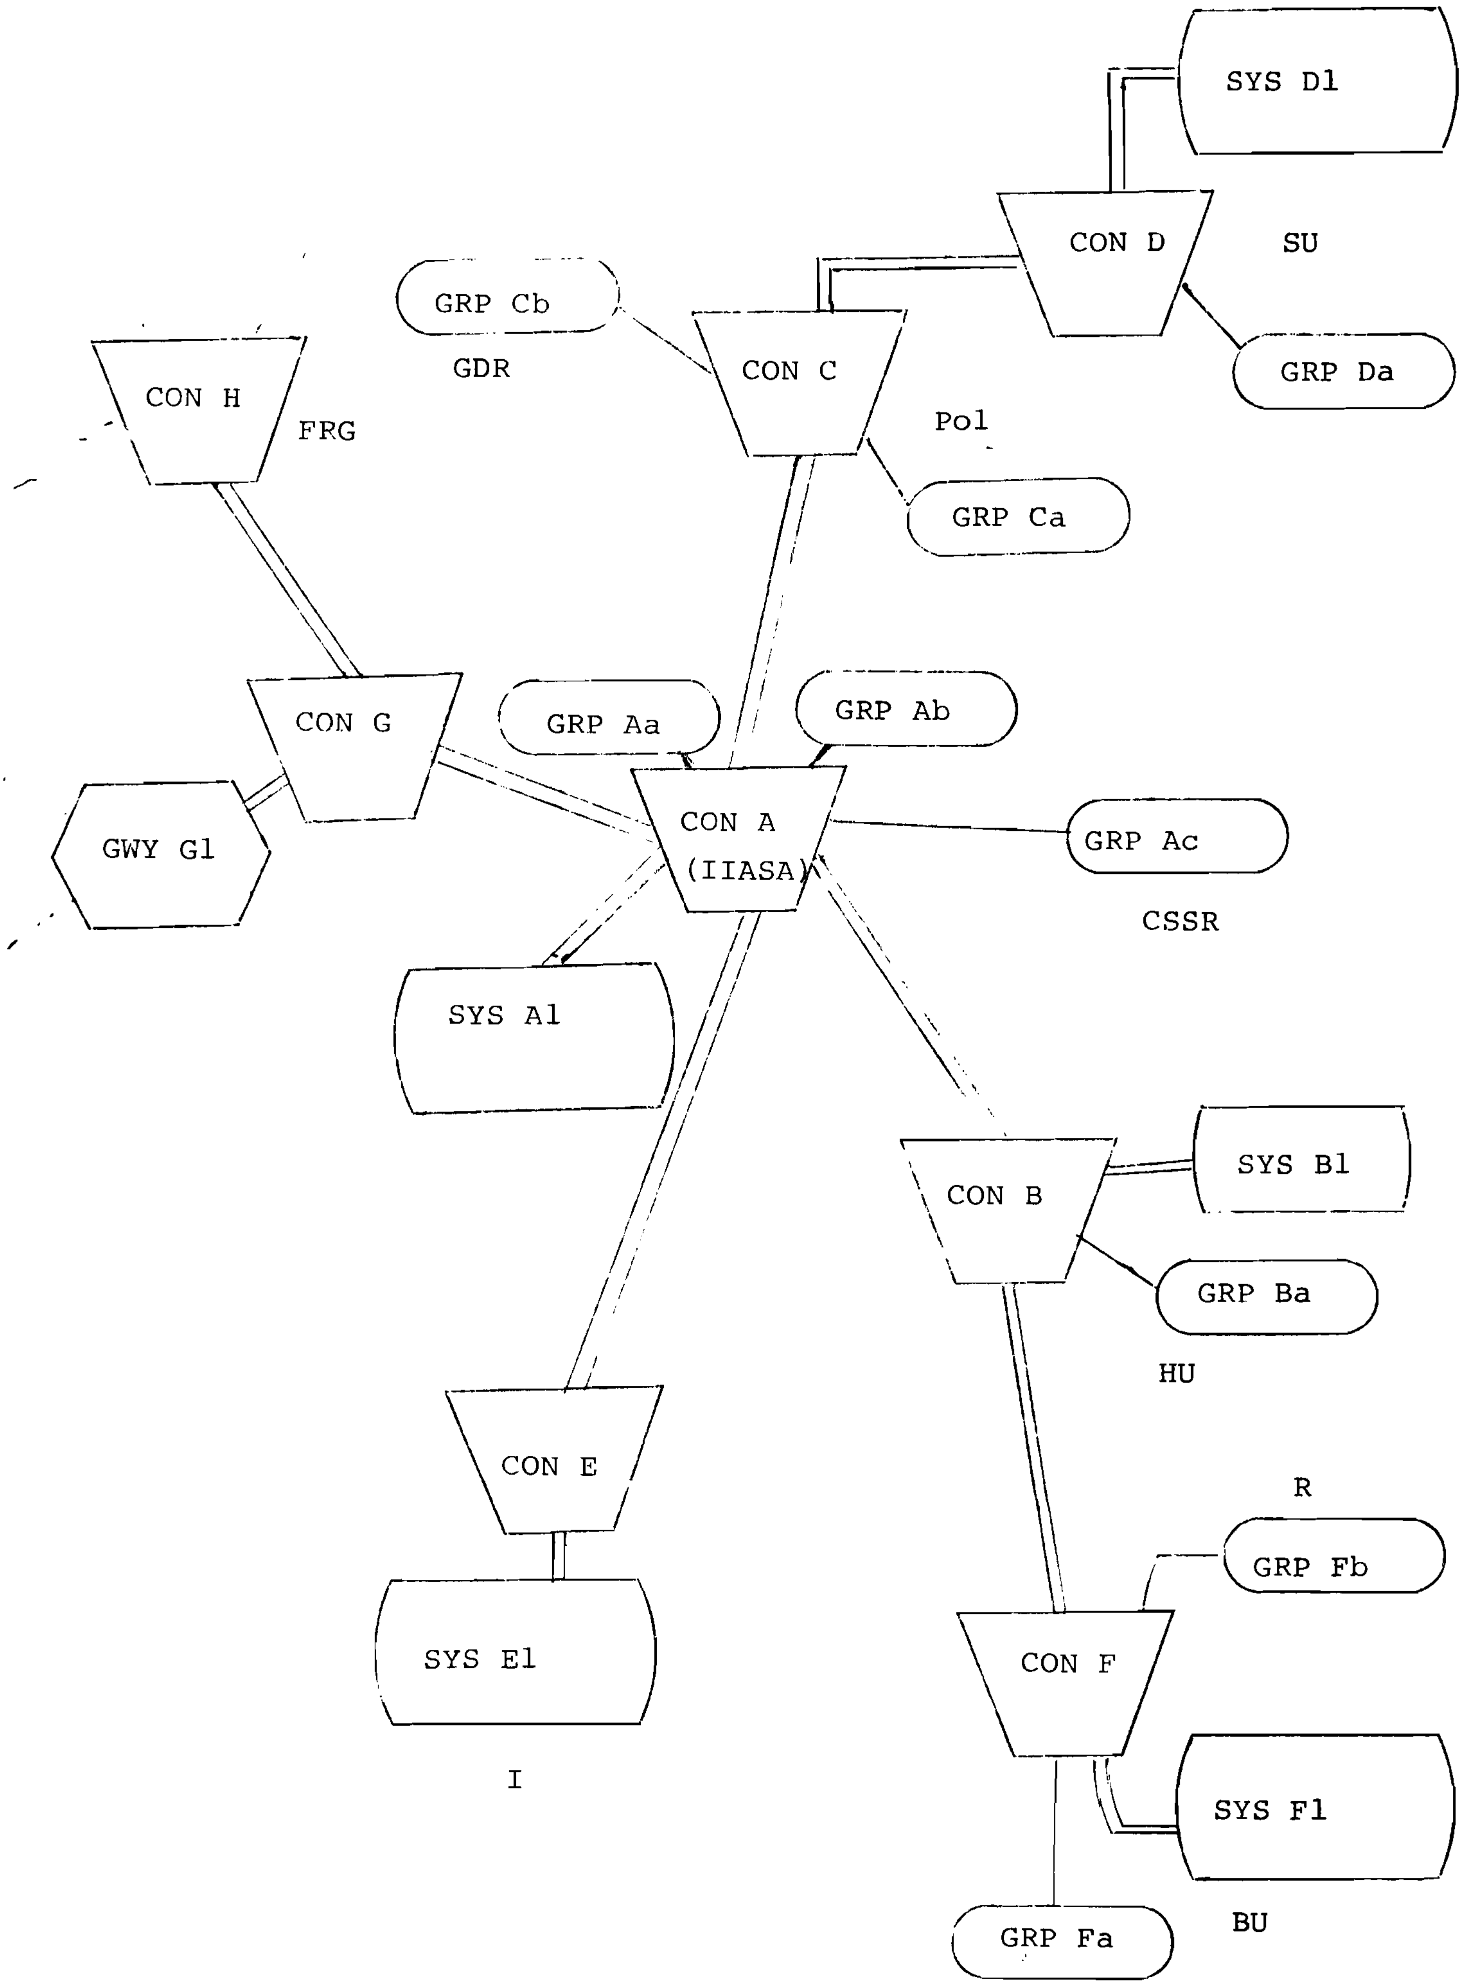

FIGURE 3. A HYPOTHETICAL NETWORK