

On Names, Addresses and Routings (II)

Dan Cohen
ISI
28 April 1978

2.3.3.11 On Names, Addresses and Routings (II)

SUMMARY

This note deals with internetwork addressing and routing.

It suggests the following:

- (1) Some systems have a tree-like hierarchical Universal-Address (UA) space.
- (2) The communication connectivity is a superset of this tree.
- (3) The postal system, the telephone and the ARPANET are such systems.
- (4) The address of a process tells where it is located (or connected to) by specifying the route to it from the root of the universal addressing tree.
- (5) The default routing to any address (unless a better one is specifically known) is up the UA-tree, from the source, and down the tree to the destination.

In the case of networks like the ARPANET, the set of all the IMPs (the subnet) is considered as a single process, known as (ARPANET).

- (6) Since the set of all networks is not connected, it cannot be tree structured. However, for ease of name-management it is possible to introduce (e.g., administratively) any arbitrary hierarchy to the address space of all networks. This hierarchy is "artificial" because it is not a subset of the communication connectivity.
- (7) Since there is no hierarchical structure to the space of all networks, there is no tree-like hierarchical internetwork Universal-Addressing scheme.

In particular, the notion of extending addresses like

(NET)-(HOST)-(PORT) or (NET)-/IMP/-(HOST)-(PORT)

upward to include METANETS, SUPERNETS and GIGANETS suffers from the lack of corresponding underlying communication structure.

- (8) Since the set of all networks is too big to be captured in local tables, and since routing cannot be derived in general from the addresses, it should always either be apriori known, or supplied by the source. This apriori knowledge does not include every step (e.g., the sequence of intermediate IMPs or PRUs). It has to include only a sequence of addresses such that the routing between them is locally known.

- (9) The corollary of this note is

OPTIONAL SOURCE ROUTING SHOULD BE IMPLEMENTED.

This note explains in great detail several aspects of the addressing schemes used by the postal and telephone systems. It also mentions the ARPANET addressing.

~~These examples are used to support the arguments which are summarized in paragraphs (1)-(9) above.~~

Unless one is very interested in the details of these systems and their relevance to the internetwork environment, or in the details of the arguments, there is no need to continue reading this note beyond this point.

INTRODUCTION

Before discussing the internet addressing/routing issues let us summarize the basic concepts:

- Processes (not hardware pieces) have names and addresses.
- A process may have more than one name, and one address, but addresses and names correspond uniquely to processes.
- The name tells WHAT the process is.
- The address tells WHERE the process is.
- The route tells HOW-TO-GET-THERE.

The last three beautiful definitions are borrowed (with appreciation) from John Shoch.

More detailed discussion of names and addresses may be found in "Internetwork Naming, Addressing, and Routing" [Internet Notebook Section 2.3.3.5, IEN 19] by John Shoch, and in "On Names, Addresses and Routings" [Internet Notebook Section 2.3.3.7, IEN 23] by Danny Cohen.

Generally the concept of address is well understood, but the concept of routing is much more complex. Who performs the address-to-routing mapping? How is it performed? These and similar problems are the topic of this note.

Most of us are familiar with the postal, the telephone, and the ARPANET addressing schemes. We also have a very good understanding of the routing processes performed for these networks. In this note we discuss the similarity of these addressing schemes, and argue that the internetwork environment violates the basic axiom which is common to them, and therefore the internetwork environment requires a different addressing and routing philosophy.

This is why we have so much trouble with internet addressing -- it is different. Our experience cannot be used as a model.

ON OTHER ADDRESSING/ROUTING SCHEMES

Three addressing/routing schemes are discussed: the postal, the telephone, and the ARPANET schemes.

The postal addressing scheme is a UA-scheme. It is defined for human processing, and therefore may tolerate a fair amount of redundancy which improves the robustness of the scheme with respect to errors and to partial losses (like stains on envelopes).

At the top level of the hierarchy there is the country name, and underneath it there are as many addressing schemes as there are countries. Some countries use ZIP codes which identify major post offices, and require more information to identify the terminal addressee. Other countries use the ZIP code (or its equivalent) to identify the individual letter-carriers, and require less additional information to identify the terminal addressee. In some cases, a street address is sufficient. In others, suite number and names are required.

In summary, the postal addressing tree is of variable depth (or: postal addresses are of variable length). Its top level (the country level) has a complete connectivity, since every country knows how to route mail to any other one.

Letters are routed, either directly to the destination or to one of its ancestors. This is performed by POs either directly or via their ancestors.

Since the address and routing processing is not fully automated, human intelligence is used for resolving ambiguities, for coping with unknown addresses, for redundancy handling and the like. Missing information is usually supplied by using common sense and defaults. As a result, a great amount of variability in addresses can be tolerated.

For example, while at Tech, I received letters addressed to

Danny (I was the only Danny there)
 256-80 (The Computer Science Mail-stop)
 91125 (The zip code for Caltech)

I also received letters addressed to:

Dr. Dan Cohen
 Computer Science, Mail-Stop 256-80
 California Institute of Technology
 Pasadena, California 91125

The first address contains all the needed information, but requires delicate handling. If any of the digits is mistyped, there is very little chance that the letter would be delivered to the right destination. The second address, which has about six times more characters, is more robust, and can cope with the multi-Danny situation, which (even though undesirable) is still very probable. The terse address has to be modified if another Danny joins this Mail-Stop.

The phone addressing is considered next. Like the postal system, each telephone station has a UA. It always starts with the country code, and usually continues with a Numbering-Plan-Area (NPA, which is the familiar Area-Code, AC), continues through a Central Office (CO) number and terminates with the station number. Even though the above fields are of variable length, the total length never exceeds 12 digits. However, when private networks are connected to the universal phone system, the entire address length may exceed this limit.

For example, my current phone number is 12138221511105, which is 14 digits long. The first "1" is the USA code, and the last "105" is my station (extension) number.

At the top level of the hierarchy there is the country code, and underneath it there are as many addressing schemes as there are countries. All countries know how to communicate with any other country. In both systems addresses are of variable length, and by looking at an arbitrary address one cannot parse it into fields without knowing the specifics of each national addressing scheme.

Here are some examples of telephone addresses and their correct parsing as COUNTRY-NPA-CO-station:

1-213-822-1511 (USA, LA)
 44-1-387-3400 (UK, London)
 44-31-332-2424 (UK, Edinburgh)
 44-745-58-3301 (UK, Wales)
 972-4-25-2690 (Israel, Haifa)
 972-67-4-0777 (Israel, Kiryat Shmona)

Obviously, the number of fields to be dialed depends on "distance" from the destination.

Since the telephone routing is performed by (relatively) simple minded automated equipment, no variability in dialing a number is allowed (except in very unusual situations).

Like the postal addresses, phone dialing sequences are of variable length, as a function of the distance to the destination. While adding the country code ("USA") would not hurt either of the postal addresses shown above, adding the local AC and country code is not allowed for local phone calls.

The reason is that one does not dial the address (telephone numbers), but dials the routing! The purpose of the telephone number is to be used for accounting and for deriving routing, but not for verbatim-dialing.

The actual telephone routing (i.e., hop-by-hop) is very efficient and deserves attention. The key to it is the existence of more communication lines than branches in the UA tree.

The USA is divided into 10 regions, which subsequently are divided into sections and areas. The grouping of areas into sections and regions cannot be simply deduced from the ACs but has to be found by a table lookup operation. This is because the ACs were most cleverly assigned to areas by population size and not geographically like the ZIP codes.

The routing is performed by each center assigning each call to a line known to be connected (or enroute) to the destination central-office. This is performed by looking at the first six digits of the address (AC+CO).

If such a line does not exist, then the call is assigned to a line known to be connected to the principal city of the destination NPA. If such a line does not exist either, the call is forwarded to the center above this one. Since at the top all regions are interconnected, this process is guaranteed to terminate.

In addition to the SIX digits recognition, the system is designed such that in many cases CO numbers do not conflict across NPAs. This eliminated the need to dial an AC of a neighboring town, across a state line, which is necessarily in a different NPA. This allows, for example, dialing from Washington, D.C., (AC=202) to Alexandria, Va., (AC=703) and to Potomac, Md., (AC=301) without dialing the AC.

The third addressing scheme is the one used in the ARPANET.

Addresses on the ARPANET are of processes which are either NCP-like in actual hosts or of other types in "fake" hosts. It is logical to extend the address notion "down" to include the port, too.

Conceptually, one can consider the ARPANET as a single process or as a star network. The fact that this single process is implemented in a very clever way by a multitude of INPs is irrelevant, from a functional point of view. This allows treating this entire network as a single addressable process, the (ARPANET), if so required.

Each IMP knows how to forward messages to any other, and therefore all IMPs constitute the top-level (and the only level) for routing.

In the ARPANET all the connections are between centers (nodes) of the same (and only) level. This is in contrast to the telephone network, which has several levels of centers, partially connected at all levels (except the top, where they are fully connected) and also partially connected between levels.

An ARPANET address is of the form {ARPANET}-{HOST}-{PORT}, which one may consider as {ARPANET}/IMP/{HOST}-{PORT}. Adding the /IMP/ field may help in the ARPANET situation, though some generality will be lacking.

This is a valid address, since the {ARPANET}-process (which is the set of all IMPs) can forward messages to all hosts, and hosts can give messages to PORTs. Therefore, routing every message up to the {ARPANET} and then down through the host to the destination port is a good default routing strategy.

INTERNETWORK ADDRESSING

After this (very) long introduction, let's return to the Internet Addressing situation.

But first, let's introduce some more formality:

- * An address is a string (i.e., an ordered set) of symbols taken from a given alphabet (e.g., ASCII, {0,1}, {1,2,...9,0}).
- * In a UA tree the level of a given address is its depth in the tree. The level of the address A is denoted by L(A).
- * Address concatenation (extended to the right) will be used. The concatenation is denoted by a "-".
- * If both A1 and A2 are addresses of the same process P, then
 - (1) L(A1) and L(A2) may differ, and
 - (2) The addresses (A1)-(X) and (A2)-(X) are necessarily addresses of the same process.
- * Addresses should always be decodable in a strictly left-to-right sequential manner (prefix coding).

What is the ARPANET address of my mailbox?

In the TENEX environment it is [ISI]-(MAILBOX)-<DANNY>. But in another environment, in the host [X], it could be [X]-<DANNY>-(MAILBOX). Or maybe in the form [X]-/TCP/-(MAIL.DEPO)-<DANNY>.

Obviously we want to expand it upwards, to allow for other networks. We could simply add in front of these addresses a network field and get

[ARPANET]-[ISI]-[MAILBOX]-<DANNY>

and similar addresses. In other networks the value of the NETWORK-field may be {PRNET-SF}, {PRNET-BOS}, {SATNET}, etc.

What is the relation between these nets? Do they all belong to the same parent like USA or USA-DoD?

This could be a solution, and if adopted my mail address may be upward expanded in the UA scheme to be like:

{GALAXY-573}-[SOLAR.SYS]-[EARTH]-<USA>-[ARPANET]-[ISI]-[MAILBOX]-<DANNY>

With a clever use of defaults, the first several fields may be omitted from most of the intraglobal communication.

The extension of this address upward, to include METANETS, SUPERNETS and GIGANETS, is very elegant. This is the prevailing popular approach expressed in a series of notes and papers, such as Ken Harrenstien's note, and various other communications between the members of [SRI-KL]<NETINFO>FIELD-ADDRESS.LIST. The intelligent reader probably has noticed by now that this note does not subscribe to the same philosophy.

This solution suffers from several problems. What is a network? Is every bus to which several processes are connected a network? If not, why? What is the relation between networks?

In order to stress the difference between the structure of the internet address and of the other universal addressing schemes, let's consider the following example. One of the ARPANET hosts, [PARC-MAXC], is tied to a private network. This network is actually a very rich internetwork environment, with about 14 networks and about 400 hosts, but for simplicity let us consider it functionally as a single {XEROX-net}. One of the hosts on this net is [RIG], the University of Rochester Intelligent Gateway to the internal {U-of-R-net}. This description is not accurate, since [RIG] is actually connected directly to the {ARPANET}, not to {XEROX-net}. But for the sake of the argument let us assume this connectivity. We preferred not to use other actual examples for several reasons. A poet's-license is nice to have....

One of the hosts on the {U-of-R-net} is, say, [NOVA-3]. What is its address? Obviously it is

[ARPANET]-[PARC-MAXC]-{XEROX-Net}-[RIG]-{U-of-R-Net}-[NOVA-3].

By the way, whenever an intermediate host in such a specification (i.e., a gateway) is between two networks only, there is no need to specify the destination network, since it is uniquely defined by context. However, this is not a good practice, since addresses have to be changed when this host is connected to more networks.

What is [ISI]'s address? Obviously (ARPANET)-[ISI].

"Not so!" scream the U-of-R people. The addresses of [NOVA-3] and of [ISI] are quite different from the point of view of [NOVA-2]. According to it, the address of [NOVA-3] is simply (U-of-R-Net)-[NOVA-3] but of [ISI] is (U-of-R-Net)-(RIG)-(XEROX-Net)-(PARC-MAXC)-(ARPANET)-[ISI].

Who is right? Neither!! Either approach is equally wrong. Neither of these addresses is above the other in the UA scheme. All the networks involved in the interconnection of these networks are of equal level, unless we decide otherwise for administrative reasons. The internet communication environment does not have up-and-down relations, except in the eyes of some users, which may be very subjective.

INTERNETWORK ADDRESSES ARE DIFFERENT

Telephone stations are always connected to the system network, and their position in it dictates their addresses. Not so with computer networks which spring into an independent existence until they are interconnected to others, if ever. Therefore, their addresses cannot be deduced from their positions (geographically or connection-wise) and vice versa.

Therefore, the network ID is an arbitrary string. Who assigns it and makes sure that no conflicts occur? Is it NBS? Jon Postel? Another Czar?

At this point we suggest that:

- * There is no universal hierarchy of networks, in contrast to the telephone and the postal systems.
- * There are too many networks to be named and/or addressed in a single flat name/address space.
- * Therefore, some naming/addressing hierarchy has to be introduced. However, this addressing tree does not serve as a basis (or underlying structure) for the communication connectivity.
- * Routing cannot be computed from any point to any point from the addresses alone.

How should internetwork routing be performed? There are obviously several possibilities. It could be performed entirely by the networks involved (i.e., the communication system), by the source, or by any combination of the communication system and the source.

It is always desirable to distribute the knowledge about possible destinations to the various centers (gateways?), such that their "sphere of knowledge" is as large as possible, though uniformity should not be required. More knowledge should be kept about frequent destinations than about less frequent ones.

Since this information must be limited due to practicalities (such as finite storage and updating procedures), it is impossible that all sources always know about all possible destinations.

What should be done about unknown destinations?

Several possibilities may be considered.

- o Having a 'supernet' of default sub destinations with the hope that they know how to find a way to the terminal destination (like the phone and the postal systems),
- o providing internetwork-wide directory services, or
- o refusal of service.

Let's consider each of these three possibilities.

Having both <USA>- (ARPANET) and <USA>- (XEROX-Net-3) does not guarantee the existence of the path (ARPANET)==<USA>==(XEROX-Net-3). Therefore, the supernet (metanet?) is not a part of the underlying communication structure as in the phone and postal situations.

In order for this approach to work, one has to create this supernet and keep it updated about all changes of the entire internetwork environment. Both the storage and the updating procedures seem to be impractical.

The second approach, the network-wide-directory-help, is a very reasonable one. It has a major drawback in the necessity to maintain centers with indefinite knowledge radius. Note that the telephone directory services are structured according to the addressing hierarchy of NPA-CO, and are not flat as may be suggested for the internetwork environment. In essence, this second approach has all the problems of the first one.

The third approach, namely, refusal of service to unknown destinations, is consistent, to say the least. Admittedly, it seems like an inconvenience to users. It should be supplemented by ways of "learning" about the unknown destinations, such that refusal should never occur, or at most, only in very rare situations. Directory services could be helpful.

Hence, this approach supports serving only destinations to which the routing is either apriori known or supplied. Note how similar this is to the telephone system philosophy.

In summary: The internet routing problem is different from the routing problems of other systems because the internet environment does not have a communication connectivity which supports a UA scheme, and therefore the addresses cannot support a direct address-to-routing mapping by using only a definite amount of knowledge.

I HATE TO ADMIT IT, BUT ...

At the beginning of this note, and in an earlier note, I used a great line telling that "names tell what the processes are, and addresses tell where they are." It continues by "routings tell how to get there."

I hate to admit that by now I have some reservations about this definition. My name is "Danny." My address is "ISI." When I was at Tech, my name was the same, but the address was different. This supports the definition. How about the addresses in a broadcasting media network? When a host changes its position (location) on the same Ethernet, its address does not change. Well, maybe these addresses are not real addresses, according to the definition. Admittedly, this is an uncomfortable thought.

I believe that there is a better explanation. I suggest that an address is "the canonic routing from the root of the addressing-tree." It sounds recursive, doesn't it?

To be more precise, an addressing scheme is a hierarchical organization of elements, with code assignment such that each element has a unique set of codes, corresponding to its position in the hierarchy.

The notion that the address tells how-to-get-there from the root of the tree is very similar to the notion that absolute coordinates are really relative, with respect to the origin.

Since we know (by default) how to get from the source to the UA root, and since the address tells how to get to the destination from the root, the address tells how to get from the source to the destination.

Hence, by definition, addresses are routings.

This leaves us only with names and routings. This should not surprise us now, since we already discovered that the telephone system has only names and routings.

Since the general internet environment does not have a hierarchy, the notion of addresses suffers. Since the addresses are "routing from the root," and since there is no root of the entire system, our conclusion is that there are no addresses, only names and routing. In other words, what we are used to call an address is actually a routing (even though it is of higher level than the hop-by-hop routing). In a well defined (and tightly controlled) environment, such as the (ARPANET), this address/routing is a well defined string. In general, it may be of indefinite length and structure.

If the destination is in the neighborhood, such as the same network, the same nets-cluster, the same agency (even on different net), the system may have a built-in knowledge of how to get to it. Otherwise the routing information needed should be supplied by the sender.

PROPOSAL FOR ADDRESSING AND ROUTING

Our proposal for addressing and routing is as follows:

- * Establish a UA scheme, of variable level structure.
- * Disseminate as much knowledge to each participating node as deemed practical.
- * Allow the option of routing to be included in the headers of the messages.
- * Refuse delivery of messages to a destination with unknown routing
- * Establish internet-directory-assistance service.

The proper use of the optional routing is to supply a set of subdestinations which may be as far apart as the networks can handle without help. This is very much like source routing for telephone connections, where a sequence of switching-centers is designated by a user, but the communication subsystem is free to optimize the hop-by-hop routing between these centers.

As long as the number of participating networks in the internetwork environment is small, it is possible to have each of them know about routing to all the others. However, we already have a large community of networks, including several PRNETs, networks in universities, about 15 networks at Xerox, the commercial and the national ones, many in DoD, all the DEC-Nets, and many others. In addition, each big modern substantial computing system is a network.

One may be advised to expect the number of networks to grow, and the internet connectivity to get more and more obscure.

Therefore, optional source routing seems to be the most sensible, and the only, alternative. This does not exclude the notion that internet clusters of any size optimize their internal routing by any scheme, for example by ARPANET-like dynamic routing scheme.

We recommend that the optional source routing will be composed of self level-identifying fields. Note that this is like the telephone dialing sequence.

The self-identifying fields could be implemented either by codes-exclusion (like the telephone system) or by identification subfields. Obviously these two schemes are equivalent, and the choice between them is just a matter of convenience.