Mail addresses are based on the ARPANET protocol listed at the end of this manual page. These addresses are in the general format where a domain is a hierarchical, dot-separated list of subdomains. For example, the address is normally interpreted from right to left: the message should go to the ARPA name tables (which do not correspond exactly to the physical ARPANET), then to the Berkeley gateway, after which it should go to the local host When the message reaches it is delivered to the user Unlike some other forms of addressing, this does not imply any routing. Thus, although this address is specified as an ARPA address, it might travel by an alternate route if that were more convenient or efficient. For example, at Berkeley, the associated message would probably go directly to over the Ethernet rather than going via the Berkeley ARPANET gateway. Under certain circumstances, it may not be necessary to type the entire domain name. In general, anything following the first dot may be omitted if it is the same as the domain from which you are sending the message. For example, a user on could send to without adding the since it is the same on both sending and receiving hosts. Certain other abbreviations may be permitted as special cases. For example, at Berkeley, ARPANET hosts may be referenced without adding the as long as their names do not conflict with a local host name. Certain old address formats are converted to the new format to provide compatibility with the previous mail system. In particular, is allowed and is converted to in order to be consistent with the command. Also, the syntax is converted to: This is normally converted back to the form before being sent on, for compatibility with older UUCP hosts. The current implementation is not able to route messages automatically through the UUCP network. Until that time you must explicitly tell the mail system which hosts to send your message through to get to your final destination. Domain names (i.e., anything after the sign) may be given in any mixture of upper and lower case with the exception of UUCP hostnames. Most hosts accept any combination of case in user names, with the notable exception of MULTICS sites. Under some circumstances it may be necessary to route a message through several hosts to get it to the final destination. Normally this routing is done automatically, but sometimes it is desirable to route the message manually. Addresses which show these relays are termed These use the syntax: This specifies that the message should be sent to from there to and finally to This path is forced even if there is a more efficient path to Route-addrs occur frequently on return addresses, since these are generally augmented by the software at each host. It is generally possible to ignore all but the part of the address to determine the actual sender. Every site is required to have a user or user alias designated to which problems with the mail system may be addressed. Some other networks can be reached by giving the name of the network as the last component of the domain. and may be supported at all sites. For example, messages to CSNET or BITNET sites can often be sent to or respectively. The RFC822 group syntax is not supported except in the special case of because of a conflict with old berknet-style addresses. Route-Address syntax is grotty. UUCP- and ARPANET-style addresses do not coexist politely. Crocker, D. H., RFC822, Standard for the Format of Arpa Internet Text Messages