Chapter 12

Epilogue

Just to confirm the old saying that there is nothing new under the sun, I offer the following epilogue.

As stated in Chapter 1, there was a flurry of creative activity in the mid to late 1970's that resulted in the internetwork architecture, and in particular Pup [9] and IP [86, 87]. These two protocols, however, do not in my mind adequately reflect the depth of understanding that the early internet architects had of routing and addressing.

Specifically, one RFC and three IENs (Internet Engineering Notes) [85, 25, 22, 18] document some of the debate that took place during the year from mid-1977 to mid-1978 and that led to the IP protocol. (Of course, I discovered these documents well after I had done the basic work on Pip :-). These four documents show that 1) one of the central ideas of Pip, namely that of putting the fields of the address in individual header fields, had been proposed by Jon Postel [85], and that 2) most of what the internet community is currently "discovering" about routing and addressing was already thought of by Postel, Sunshine, Cohen, Clark, Cerf, and perhaps others who did not bother putting out IENs (or whose IENs or other publications I have not bothered to read).

In what follows, the aspects of that debate that relate to the findings of this thesis are summarized.

The opening paragraph in Postel's May 1977 RFC730, "Extensible Field Addressing", says:

> This memo discusses the need for and advantages of the expression of addresses in a network environment as a set of fields. The suggestion is that as the network grows the address can be extended by three techniques: adding fields on the left, adding fields on the right, and increasing the size of individual fields. Carl Sunshine has described this type of addressing in a paper on source routing [102].

This, in a nutshell, is Pip's (or SPip's) route sequence, which allow adding fields (RSEs) on the left, the right, and in the middle (which Postel's can do too, though he does not mention it above).

Later in the RFC, Postel says:

> The prospect of interconnections of networks to form a complex multinetwork system poses additional addressing problems. The new Host-IMP interface specification has reserved fields in the leader to carry network addresses. There is experimental work in progress on interconnecting networks. We should be prepared for these extensions to the address space.

Talk about understatement!

And later still:

> A problem with simple field addressing is the desire to specify only the fields that are necessary given the local context. A program interpreting the address is then unsure what the first field represents. Some clues are needed in the address specification for correct parsing to be possible. Dave Crocker has described a syntax for a similar problem in network access of data.

Trying to do this (only including the fields of the address relevant to local context in the header) with SPip turned out to be problematic. Still, the "clue" that Postel refers to is useful for efficient processing of the header, as the router only needs to parse the relevant fields. This "clue" is SPip's Active RSE field.

Postel gets to the meat of the thing in the following excerpt:

> Specifically I suggest that we adopt a field based extensible address scheme where each field is separated from its neighbors by a delimiter character and each field has a name. When an address is specified the name of the most general field must also be indicated.
>
> Definitions:
> ⟨address⟩ ::= ⟨field-name⟩ ":" ⟨fields⟩
> ⟨field-name⟩ ::= "NET" — "IMP" — "HOST" — "MESSAGE-ID"
> ⟨fields⟩ ::= ⟨field⟩ — ⟨field⟩ "/" ⟨fields⟩
> ⟨field⟩ ::= a decimal number

SPip has a few differences from this, but the basic idea is there. Something to give the context of the field to be parsed (SPip's Active RSE = Postel's field-name), and a series of fields (SPip's are binary not decimal, but had Postel's found their way into a packet, I'm sure they would have been binary).

In April of 1978, Danny Cohen shed more light on the nature of addresses in his IEN31, "On Names, Addresses and Routings (II)" [25]:

> 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, does not 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 last conclusion, that addresses are routes, is a key "finding" of the taxonomy section (Section 2.2) of this thesis, and is the basis of SPip's route sequence.

Later on in the IEN, Cohen makes a proposal:

> Our proposal for addressing and routing is as follows:
>
> - Establish a UA (Universal Address) 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.

This last point is crucial. "Internet-directory-assistance" (now known as DNS) must advertise the "route" from the root of the hierarchy to the leaves. In particular, if the packet format is a string of fields (or addresses, as in SIPP), then DNS should advertise that string.

So, at this point in the discourse (April 1978), Postel has provided the address format, and Cohen the architectural underpinning from which to understand that address. So, what happened? Why did we end up with IP and not something more like SPip?

We find a clue to the answer from IEN46, written by Clark and Cohen in June of 1978 called "A Proposal for Addressing and Routing in the Internet" [22]. After discussing several problems with routing and addressing, they make the following statement:

> The solution which has been proposed in the past to cope with this is to replace the address in the packet with a route, called a source route since it is provided by the source of the packet. The disadvantage of having a route in a packet instead of an address is that the concept of an address is very useful one. For example, for accounting purposes it may be necessary to note the source and destination of a packet as it passes through a transit net. Clearly, it is desirable that the source and destination be uniquely identified for this purpose, something not easily done if the source and destination are specified only by a route. Thus, we propose that the address continue to be the primary piece of information in the packet, but that it be possible to include, in addition, an optional source route.

So, here they recognize the need for a compact, simple, fixed length *something* to identify the source and destination of a packet. But, this is nothing more than the EID of SPip. So, the need for both an identifier and an "address" (still at that time arguable to be a route) was clearly recognized. However, they added the source route to handle the routing bit, and kept the address as the primary piece of information.

I think this would have been fine (indeed, this is SIPP's approach) except for the crucial thing mentioned by Cohen in IEN31:

- Establish internet-directory-assistance service.

Well, DNS was of course established, but it did not contain the source route, just the address. So, the "routing" information in the packet was effectively limited to a single 32 bit field.

I was interested to find the following in the Clark/Cohen paper:

> 5. Migration
>
> What is the relationship between the scheme proposed here and the current internet header with a fixed size address field? Happily, adoption of the addressing strategy involving regions together with the optional internet source route implies no immediate upheaval to packet formats or gateway code. Currently, every network is a region, and every gateway thus contains code for doing inter-region routing. Eventually, gateways will want to be modified as follows. When a region finally is defined which contains more than one network, then gateways inside that region will need to understand an additional component of the internet address. Thus, unless gateway code is rewritten

> for different regions, it will be necessary to write code which can deal, eventually, with a variable size component of the address. The address itself, however, can reasonably be a fixed size, since it is merely an address and not a route. In fact, it seems that the field as specified for the current internet header is sufficient in size, although perhaps marginally so.

Well, something happened here. An argument was put forth that 32 bits is enough because the address does not have to do routing–the source route can handle the rest. Clearly it was recognized that a variable length *something* was needed, but the source route was deemed sufficient for that, and the 32-bit address won out in the end. So, perhaps what killed IP is not that the address is too short (though probably it is), but that the ability for DNS to hand a host a source route (which it could then put in the header so that the right thing could happen in the network) was not created.

So, indeed SPip's routing sequence is a combination of Postel's Extensible Field Addressing (EFA) and Clark and Cohen's "address", though with SPip the "routing" part of the "address" has been largely moved over to the EFA (route sequence in SPip), and the "address" (EID in SPip) is left with the identification function.

An IEN from Cerf the following month (July 1978) seems to meld with Clark/Cohen (IEN48, "The Catenet Model for Internetworking" [18]). It generally confirms the Clark/Cohen proposal. It, however, makes some additional interesting statements:

> In order to limit the overhead of address fields in the header, it was decided to restrict the maximum length of the host portion of the internet address to 24 bits. The possibility of true, variable-length addressing was seriously considered. At one point, it appeared that addresses might be as long as 120 bits each for source and destination. The overhead in the higher level protocols for maintaining tables capable of dealing with the maximum possible address sizes was considered excessive.

Not only is it interesting that a longer address (120 bits, almost as long as an NSAP), was seriously considered, but the reason for not going with it (memory overhead to "upper layer protocols") really shows how times have changed.

Finally, Cerf's IEN seems to delegate source routing to its current, and very limited, role:

> One of the major arguments in favor of variable length "addressing" is to support what is called "source-routing." The structure of the information in the "address" really identifies a route (e.g., through a particular sequence of networks and gateways). Such a capability could support ad hoc network interconnections in which a host on two nets could serve as a private gateway. Though it would not participate in catenet routing or flow control procedures, any host which knows of this private gateway could send "source-routed" internet datagrams to that host.

It is interesting that the original ideas of Postel and Cohen (very SPip-like) evolved into the source route, which was then limited to a "special service" role (i.e., routing a packet through a private host on two nets).

To conclude, I must say that when I read these four documents, I found it fascinating and delightful to discover that my work, with the considerable aid of hindsight, was able to confirm, and put in a modern context, the early thinking of the internet architects.

Bibliography

[1] A. Autolitano, F. Bernabei, M. Ciampi, and M. Listanti. Application of Generalized Parallel Delta Networks to a Hybrid Broadband Switch. *International Conference on Communications*, 1:123–127, June 1989.

[2] A. Ballardie, P. Francis, and J. Crowcroft. An Architecture for Scalable Inter-Domain Multicast Routing. *Proceedings of ACM SIGCOMM 93*, pages 85–95, September 1993.

[3] A. Bar-Noy and M. Gopal. Topology Distribution Cost vs. Efficient Routing in Large Networks. *Proceedings of ACM SIGCOMM 90*, pages 242–252, September 1990.

[4] F. Bernabei and M. Listanti. A Hybrid Switching Exchange for Broadband Communications. *Proceedings of the Ninth International Conference on Computer Communication Technologies for the 90's*, pages 61–65, October 1988.

[5] D. Bertsekas and R. Gallager. *Data Networks*. Prentice-Hall, 1987.

[6] L. Bhuyan and D. Agrawal. Generalized Hypercube and Hyperbus Structure for a Computer Network. *IEEE Transactions on Computing*, C-33:323–333, 1984.

[7] big internet@munnari.oz.au. *Big Internet Mailing List*. archive at munnari.oz.au:big-internet/list-archive/*.

[8] K. Birman, A. Schiper, and P. Stephenson. Lightweight Causal and Atomic Group Multicast. *ACM Transactions on Computer Systems*, 9(3):272–314, August 1991.

[9] D. Boggs, J. Shoch, E. Taft, and R. Metcalfe. Pup: An Internetwork Architecture. *IEEE Transactions on Communications*, COM-28(4):612–624, April 1980.

[10] R. Braden and J. Postel. Requirements for Internet Gateways. Request For Comments 1009, University of Southern California Information Sciences Institute, June 1987.

[11] L. Breslau and D. Estrin. Design of Inter-Administrative Domain Routing Protocols. *Proceedings of ACM SIGCOMM 90*, pages 231–241, September 1990.

[12] J. Case, M. Fedor, M. Schoffstall, and J. Davin. A Simple Network Management Protocol (SNMP). Request For Comments 1157, University of Southern California Information Sciences Institute, May 1990.

[13] Steve Casner. *Private Communications.* ISI.

[14] CCITT. *CCITT E.163, Numbering Plan for the International Telephone Service*, Blue Book, 1988.

[15] CCITT. *CCITT X.121, International Numbering Plan for Public Data Networks*, Blue Book, 1988.

[16] CCITT. *CCITT X.25, Interface Between Data Terminal Equipment (DTE) and Data Circuit Terminating Equipment (DCE) for Terminals Operating in the Packet Mode on Public Data Networks*, Blue Book, 1988.

[17] CCITT. *Draft Recommendation Q.931, ISDN User-Network Interface Layer 3–General Aspects*, Blue Book, 1988.

[18] V. Cerf. The Catenet Model for Internetworking. Internet Engineering Note 48, University of Southern California Information Sciences Institute, July 1978.

[19] D. Cheriton. Sirpent: A High-Performance Internetworking Approach. *Proceedings of ACM SIGCOMM 89*, Sepember 1989.

[20] I. Cidon and I. Gopal. Control Mechanisms for High-Speed Networks. *Proceedings of IEEE International Conference on Communications '90*, April 1990.

[21] D. Clark. Policy routing in Internet Protocols. Request For Comments 1102, University of Southern California Information Sciences Institute, May 1989.

[22] D. Clark and D. Cohen. A Proposal for Addressing and Routing in the Internet. Internet Engineering Note 46, University of Southern California Information Sciences Institute, June 1978.

[23] D. Clark, V. Jacobson, J. Romkey, and H. Salwen. An Analysis of TCP Processing Overhead. *Transactions on Communications*, 27(6):23–29, June 1989.

[24] D. Clark and D. Tennenhouse. Archictectural Considerations for a New Generation of Protocols. *Proceedings of ACM SIGCOMM 90*, pages 200–208, September 1990.

[25] D. Cohen. On Names, Addresses and Routings (II). Internet Engineering Note 31, University of Southern California Information Sciences Institute, April 1978.

[26] J. Crowcroft and K. Paliwoda. A Multicast Transport Protocol. *Proceedings of ACM SIGCOMM 88*, pages 247–256, August 1988.

[27] S. Deering. Multicast Routing in Internetworks and Extended LANs. *Proceedings of ACM SIGCOMM 88*, August 1988.

[28] S. Deering. Host Extensions for IP Multicasting. Request For Comments 1112, University of Southern California Information Sciences Institute, August 1989.

[29] S. Deering. *Multicast Routing in a Datagram Internetwork*. PhD thesis, Stanford University, Palo Alto, California, 1991.

[30] S. Deering. SIP: A Simple Internet Protocol. *IEEE Network*, 7(6):16–28, May 1993.

[31] S. Deering, D. Waitzman, and C. Partridge. Distance Vector Multicast Routing Protocol. Request For Comments 1075, University of Southern California Information Sciences Institute, November 1988.

[32] R. Droms. Dynamic Host Configuration Protocol. Request For Comments 1531, University of Southern California Information Sciences Institute, October 1993.

[33] W. Edmond, K. Seo, M. Leib, and C. Topolcic. The DARPA Wideband Network Dual Bus Protocol. *Proceedings of ACM SIGCOMM 90*, pages 79–89, September 1990.

[34] D. Estrin. Policy Requirements for Inter Administrative Domain Routing. Request For Comments 1125, University of Southern California Information Sciences Institute, November 1989.

[35] D. Estrin and K. Obraczka. Connectivity Database Overhead for Inter-Domain Policy Routing. *Proceedings of IEEE Infocom 91*, April 1991.

[36] D. Estrin, Y. Rekhter, and S. Hotz. Scalable Inter-Domain Routing Architecture. *Proceedings of ACM SIGCOMM 92*, pages 40–52, August 1992.

[37] P. Francis. A Near-term Architecture for Deploying Pip. *IEEE Network*, 7(6):30–37, May 1993.

[38] P. Francis. Pip Header Processing. Request For Comments 1622, University of Southern California Information Sciences Institute, May 1994.

[39] P. Francis. Pip Near-term Architecture. Request For Comments 1621, University of Southern California Information Sciences Institute, May 1994.

[40] P. Francis and R. Govindan. Flexible Routing and Addressing For a Next Generation IP. *Submitted to ACM SIGCOMM 94*, 1994.

[41] G. Fredereckson and R. Janardan. Designing Networks with Compact Routing Tables. *Algorithmica*, 3:171–190, 1988.

[42] V. Fuller, T. Li, J. Yu, and K. Varadhan. Classless Inter-domain Routing (CIDR): an Address Assignment and Aggregation Strategy. Request For Comments 1519, University of Southern California Information Sciences Institute, September 1993.

[43] G. Finn. Routing and addressing problems in large metropolitan-scale internetworks. Research Report ISI/RR-87-180, Information Sciences Institute, Marina del Rey, California, March 1987.

[44] G. Finn. Reducing the vulnerability of dynamic computer networks. Research Report ISI/RR-88-201, Information Sciences Institute, Marina del Rey, California, June 1988.

[45] P. Gross and P. Almquist. IESG Deliberations on Routing and Addressing. Request For Comments 1380, University of Southern California Information Sciences Institute, November 1992.

[46] J. Hagouel. *Issues in Routing for Large and Dynamic Networks.* PhD thesis, Columbia University, New York City, 1983.

[47] Joel Halpern. *Private Communications.* Network Systems Corp.

[48] E. Huizer and D. Crocker. IETF Wroking Group Guidelines and Procedures. Request For Comments 1603, University of Southern California Information Sciences Institute, March 1994.

[49] idmr@cs.ucl.ac.uk. *Inter-Domain Multicast Routing Mailing List.* ftp cs.ucl.ac.uk:darpa/idmr-archive.Z.

[50] IEEE. *Logical Link Control*, IEEE 802.2, 1988.

[51] IEEE. *Token Ring Access Method and Physical Layer Specifications*, IEEE 802.5, 1989.

[52] International Organization for Standardization. *Data Interchange - Structures for the Identification of Organizations*, ISO 6523, 1984.

[53] International Organization for Standardization. *End System to Intermediate System routeing exchange protocol for use in conjunction with the Protocol for providing the connectionless-mode network service (ISO 8473)*, ISO 9542, 1988.

[54] International Organization for Standardization. *Information Processing Systems - Data Communications - Network Service Definition - Addendum 2: Covering Network Layer Addressing*, ISO 8348 AD2, March 1988.

[55] International Organization for Standardization. *Protocol for providing connectionless-mode network service*, ISO 8473, 1988.

[56] International Organization for Standardization. *Intermediate System to Intermediate System routeing exchange protocol for use in conjunction with the Protocol for providing the connectionless-mode network service (ISO 8473)*, ISO 10589, 1991.

[57] International Organization for Standardization. *Information Technology - Telecommunications and Information Exchange Between Systems - Network Service Definition for Open Systems Interconnection: Amendment 5: Addition of Group Network Addressing*, ISO/IEC 8348/Am. 5, 1993.

[58] J. Ioannidis, D. Duchamp, and G. Maguire Jr. IP-Based Protocols for Mobile Internetworking. *Proceedings of ACM SIGCOMM 91*, pages 235–245, 1991.

[59] J. Liu and W. Hsu. Routing and broadcasting algorithms for a large family of interconnection topologies. Technical Report CPS-91-05, Dept. of Computer Science, Michigan State University, October 1991.

[60] V. Jacobson. Compressing TCP/IP Headers for Low-Speed Serial Links. Request For Comments 1144, University of Southern California Information Sciences Institute, February 1990.

[61] J. Jaffe. Hierarchical clustering with topology databases. *Computer Networks and ISDN Systems*, 15(5):329–339, 1988.

[62] F. Kamoun and L. Kleinrock. Hierarchical Routing for Large Networks: Performance Evaluation and Optimization. *Computer Networks*, 1(3):155–174, January 1977.

[63] H. Katseff. Incomplete Hypercube. *IEEE Transactions on Computing*, C-37:604–607, 1988.

[64] D. Katz and P. Ford. TUBA: Replacing IP with CLNP. *IEEE Network*, 7(6):38–47, May 1993.

[65] D. Knuth. *The Art of Computer Programming, Vol. 3 (Sorting and Searching).* Addison Wesley, 1973.

[66] J. Kolar and H. Wu. A Study of Survivability Versus Cost for Several Fiber Network Architectures. *IEEE International Conference on Communications '88*, pages 61–66, vol. 1, June 1988.

[67] G. Lauer. Hierarchical routing design for SURAN. *IEEE International Conference on Communications (ICC) '86*, 1:93–102, June 1986.

[68] D. Lynch and M. Rose. *Internet System Handbook.* Addison Wesley, 1993.

[69] N. Maxemchuck. The Manhattan Street Network. *IEEE GLOBECOM '85*, pages 255–261, vol. 1, December 1985.

[70] A. McAuley and P. Francis. Fast Routing Table Lookup Using CAMs. *Proceedings of INFOCOM 93*, pages 1382–1391, March 1993.

[71] P. Mockapetris. Domain Names - Implementation and Specification. Request For Comments 1035, University of Southern California Information Sciences Institute, November 1987.

[72] R. Moore, N. Geer, and H. Graf. GRIDNET- An Alternative Large Distributed Network. *IEEE Computer Magazine*, pages 57–66, April 1984.

[73] J. Moy. OSPF Version 2. Request For Comments 1131, University of Southern California Information Sciences Institute, July 1991.

[74] R. Newman, Z. Budrikis, and J. Hullett. The QPSX Man. *IEEE Communications Magazine*, 26(4):20–28, April 1988.

[75] C. Partridge. A Proposed Flow Sepcification. Request For Comments 1363, University of Southern California Information Sciences Institute, September 1992.

[76] C. Partridge, T. Mendez, and W. Milliken. Host Anycasting Service. Request For Comments 1546, University of Southern California Information Sciences Institute, November 1993.

[77] C. Partridge and S. Pink. A Faster UDP. *Transactions on Networking*, 1(4):429–440, August 1993.

[78] T. Pei and C. Zukowski. VLSI Implementation of Routing Tables: Tries and CAMs. *Proceedings of INFOCOM 91*, April 1991.

[79] R. Perlman. A Protocol for Distributed Computation of a Spanning Tree in an Extended LAN. *Ninth Data Communications Symposium*, 1985.

[80] R. Perlman. Hierarchical Networks and the Subnetwork Partition Problem. *Computer Networks and ISDN Systems*, 5:297–303, 1985.

[81] R. Perlman. *Network Layer Protocols with Byzantine Robustness*. PhD thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, 1988.

[82] R. Perlman and G. Varghese. Pitfalls in the Design of Distributed Routing Algorithms. *Proceedings of ACM SIGCOMM 88*, pages 43–54, August 1988.

[83] D. Pitt and R. Dixon. Addressing, Bridging, and Source Routing (LAN Interconnection). *IEEE Network*, 2(1):25–32, January 1988.

[84] D. Plummer. Ethernet Address Resolution Protocol: Or Converting Network Protocol Addresses to 48.Bit Ethernet Address for Transmission on Ethernet Hardware. Request For Comments 826, University of Southern California Information Sciences Institute, November 1982.

[85] J. Postel. Extensible Field Addressing. Request For Comments 730, University of Southern California Information Sciences Institute, May 1977.

[86] J. Postel. Internetwork Protocol Approaches. *IEEE Transactions on Communications*, COM-28(4):604–611, April 1980.

[87] J. Postel. Internet Protocol. Request For Comments 791, University of Southern California Information Sciences Institute, September 1981.

[88] J. Postel. Transmission Control Protocol. Request For Comments 793, University of Southern California Information Sciences Institute, September 1981.

[89] D. Pradhan. Dynamically Restructurable Fault Tolerant Processor Network Architectures. *IEEE Transactions on Computing*, C-34:434–447, 1985.

[90] B. Rajagopalan. Reliability and Scaling Issues in Multicast Communication. *Proceedings of ACM SIGCOMM 92*, pages 188–198, August 1992.

[91] K. Ramakrishnan. Performance Considerations in Designing Network Interfaces. *Journal on Selected Areas in Communications*, 11(2):203–219, February 1993.

[92] J. Reynolds and J. Postel. Assigned Numbers. Request For Comments 1340, University of Southern California Information Sciences Institute, July 1992.

[93] Benny Rodrig. *Private Communications.* RAD Network Devices, Ltd.

[94] V. Rutenburg and R. Ogier. Fair Charging Policies and Minimum-Expected-Cost Routing in Internets with Packet Loss. *Proceedings of IEEE Infocom 91*, April 1991.

[95] H. Saltzer, D. Reed, and D. Clark. Source Routing for Campus-Wide Internet Transport. *Proceedings of the IFIP WG 6.4 Workshop on Local Networks*, pages 25–32, August 1980.

[96] J. Saltzer. On the Naming and Binding of Network Destinations. *Local Computer Networks*, pages 311–317, 1982.

[97] N. Santoro and R. Khalib. Routing Without Routing Tables. Technical Report SCS-TR-6, Carlton University School of Computer Science, Ottawa, 1982.

[98] N. Shacham. Organization of Dynamic Radio Network by Overlapping Clusters: Architecture Considerations and Optimization. *Proceedings of the Tenth International Symposium PERFORMANCE '84*, pages 435–447, December 1984.

[99] J. Shoch. Inter-Network Naming, Addressing, and Routing. *Proceedings 17th IEEE Computer Society Internation Conference*, pages 72–79, September 1978.

[100] T. Starling. Network Reconfiguration Saves Society Corp. $1 Million a Year. *Bank System Equipment*, 25(11):48–49, 1988.

[101] M. Steenstrup. An Architecture for Inter-Domain Policy Routing. Request For Comments 1478, University of Southern California Information Sciences Institute, June 1993.

[102] C. Sunshine. Source Routing and Computer Networks. *Proceedings of ACM SIGCOMM 77*, January 1977.

[103] C. Sunshine. Addressing Problems in Multi-Network Systems. *Proceedings of INFOCOM 82*, pages 12–18, March 1982.

[104] F. Teraoka, Y. Yokote, and M. Tokoro. A Network Architecture Providing Host Migration Transparency. *Proceedings of ACM SIGCOMM 91*, pages 209–220, 1991.

[105] D. Train and A. Carlton. Error Handling and Other Implementation Details of Cut-Through Switching. *Pacific Rim Conference on Communications, Computers and Signal Processing*, pages 125–127, June 1987.

[106] P. Tsuchiya. The Landmark Hierarchy: a New Hierarchy for Routing in Very Large Networks. *Proceedings of ACM SIGCOMM 88*, pages 35–42, August 1988.

[107] P. Tsuchiya. Efficient and Robust Policy Routing Using Multiple Hierarchical Addresses. *Proceedings of ACM SIGCOMM 91*, September 1991.

[108] P. Tsuchiya. On the Assignment of Subnet Numbers. Request For Comments 1219, University of Southern California Information Sciences Institute, April 1991.

[109] P. Tsuchiya. Internet Routing over Large Public Data Networks using Shortcuts. *Proceedings of ACM SIGCOMM 92*, October 1992.

[110] U.S. Department of Commerce, National Institute of Standards and Technology, Gaithersburg, MD, USA. *Government Open Systems Interconnection Profile (GOSIP) Version 2*, federal information processing standard 146-1 edition, April 1991.

[111] J. Van Leeuwen and R. Tan. Interval Routing. *The Computer Journal*, 30(4):298–307, 1987.

[112] D. Wall. *Mechanisms for Broadcast and Selective Broadcast*. PhD thesis, Stanford University, Palo Alto, California, 1980.

[113] J. Westcott and G. Lauer. Hierarchical Routing for Very Large Networks. *IEEE Military Communications Conference MILCOM '84*, 2:214–218, October 1984.

[114] L. Zhang, S. Deering, D. Estrin, S. Shenker, and D. Zappala. RSVP: A New Resource ReSerVation Protocol. *IEEE Network*, September 1993.