MIRAGE:

A Model for Latency in Communication [1]

A Proposal for Dissertation Research[2]

Joseph D. Touch

July 15, 1990

Department of Computer and Information Science
University of Pennsylvania
Philadelphia, PA 19104-6389
touch@cis.upenn.edu

---

1 This work is supported by the Information Science and Technology Office of the Defense Advanced Research Projects Agency, under contract NAG-2-639.

2 Submitted in partial fufillment of the requirements of the Dissertation Proposal, towards the degree of Ph.D. in Computer and Information Science.

ABSTRACT: Mirage is a formal model for the design and analysis of high speed wide area network protocols. It attempts to extend Shannon's communication theory by accounting for the effects of latency, just as Shannon's accounts for communication errors.

Current protocols are expected to become inefficient if used at speeds in excess of 1 Gigabit per second, in wide area networks. These potential inefficiencies would arise since the speed of communication is increased, without a corresponding decrease in latency; only now, in the gigabit, wide area domain, are these effects becoming noticable.

This is a proposal for the development of a model for understanding communication under noticable latency, which may also explain the deficiencies of existing protocols in this new domain. We want to construct a framework for examining these issues, to determine if and why existing protocols may fail.

Mirage is motivated by concepts from quantum physics, and attempts to introduce virtual communication akin to quantum virtual particles. Existing protocols are akin to classical mechanics; 1Gigabit/second is the speed near which relativistic effects emerge. In order to account for these effects, we need to express knowledge at a distance, latent measurement, and uncertainty as real entities, not negligible estimates. The result is a model which can express existing protocols, and may contribute to a better understanding of the Gigabit communications domain.

1. Introduction

Mirage is a model for the analysis and design of high speed wide-area protocols [ToFa89]. The primary goal of this model is to understand the effects of moving to the gigabit domain in wide area networks, verifying or disproving the predicted failure of existing protocols, and anticipate potential solutions.

One characteristic of gigabit wide area networks which differentiates them from their slower or more proximal counterparts is that latency becomes comparable to bandwidth and node size. It is predicted that expected inefficiencies will be attributed to the effects of fixed latency, as bandwidth and network sizes scale. As such, a derivative and more fundamental goal of this research is *providing a framework for understanding the effects of fixed propagation delay and bandwidth increases on communication*.

This approach is a highly explorative effort which will set the basis for understanding the characteristics of high speed wide-area protocols by investigating the effects of latency. While it is hoped it will also provide a basis for design, the model was not created with a solution in mind.

Mirage denotes the difficulty with high-speed, wide area network protocols, in that by the time requested information arrives, it may no longer be accurate. Nodes in a high speed network never really "see" each other; rather, they work with (and around) the *mirages* which high-speed and latency conjure before them.

First, a brief discussion of some of the changing network characteristics which necessitate a fresh view of communication protocols is required. The primary influence of these changes is a decrease in the large gap between node size and the amount of data caught in network round trip latency. This decrease causes conventional protocols to decrease in channel utilization, many of which are designed for file transfer based on sliding window flow control.

One possible reason for the deterioration of efficiency in existing protocols is that they operate in a point model, based on Shannon's communication theory [ShWe63]. This theory accounts for channel error by sequence encoding; higher channel errors requiring encoding over longer sequences, resulting in a trade of error for latency.

In the high speed, wide-area domain, network inefficiencies are caused by the combined effect of increased channel capacity, without a corresponding decrease in communication latency (due to finite propagation delays). Mirage proposes a view where latency can be compensated by accepting information imprecision (a controlled form of error), thus inverting the problem. In the inversion, the information of remote nodes, formerly precise points in state space, become imprecise volumes in state space. Mirage defines forms of communication as transformations on these state space volumes. It also incorporates the effects of time in its transformations on the volumes as well.

This research is based on suggestions derrived from analogies in physics. Communication theory already incorporates physics analogs, most notably that between information and negative entropy. Here we investigate other analogies as well. The Mirage model, of state space volume transformations, is an attempt to incorporate the imprecision evident in quantum models into communication protocol analysis.

2. What has changed?

We have previously built networks where the bandwidth*delay product was small compared to the size of the nodes participating in the protocols. In existing, window-based protocols, utilization is optimal only where the buffer space is larger than the round trip data delay; when smaller, the utilization is proportional to the ratio. Graphing the three factors of node size, which determines the number of buffers available, round trip time, and channel bandwidth, we see that we have previously been operating on the plateau of a curve, and are rapidly approaching the cliff, where utilization drops dramatically even for subtle variation in parameter values [Figure 1].

Existing protocols and models specifically do not address the issue of latency, because in existing networks it hasn't been a problem. Here we want to extend communication theory to give a formal underpinning, to create a framework for understanding latency which is both analytic and predictive. In addition, existing protocols and models should be expressible in this model, where latency is considered negligible. One of the fist considerations is to examine the most fundamental model of communication, Shannon’s, with respect to its treatment of latency.

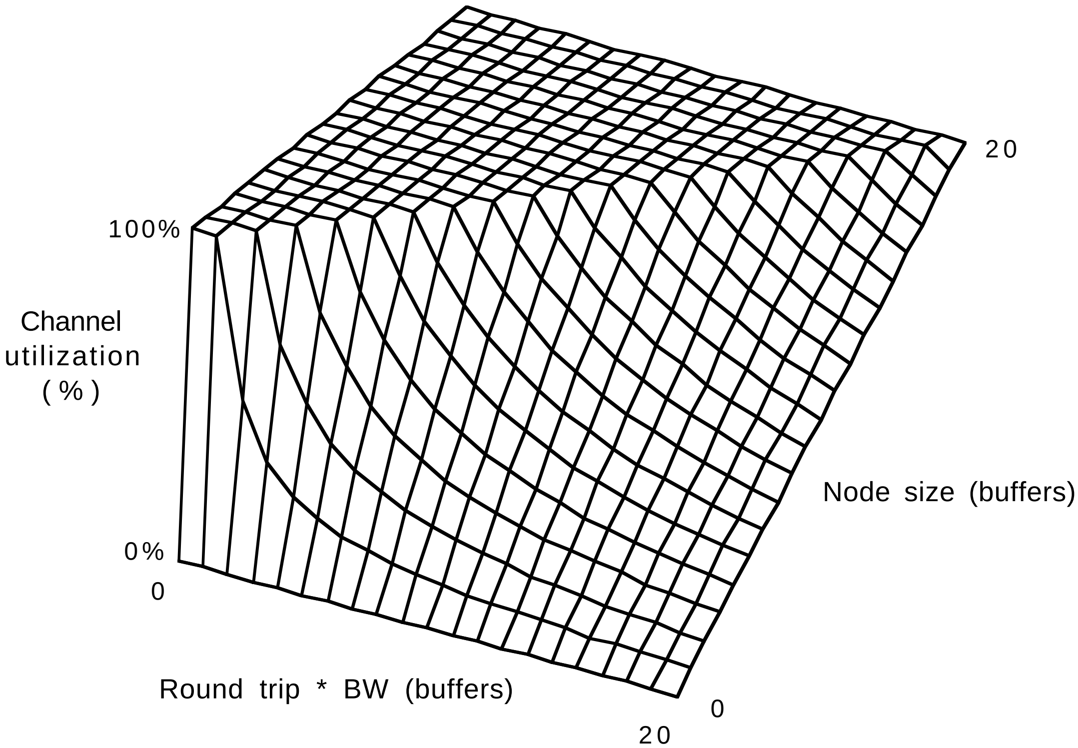

Figure 1: Utilization v.s. existing protocol parameters

2.1. Shannon's Model - Extensions for High Speed

Shannon's mathematical model of communication defines channel bandwidth and capacity, and analyzes the capacity of the channel under the constraint of transmission error [ShWe63]. In his model, the channel is viewed as a pipe between the communicating nodes [Figure 2]. Bandwidth is a unit of volume of flow in this pipe - bits wide x signal duration. Note that the propagation (latency) of this volume as it traverses the pipe is ignored - Mirage will add this factor, in its extension of this model.

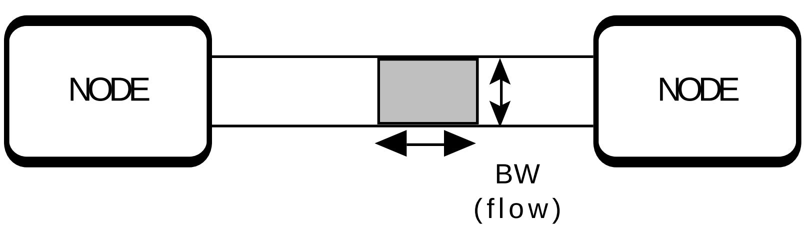

Figure 2: Shannon's communication channel

Shannon's model is based on a denoting the state of a node as a point in state space, implying that the values at the node are known precisely at remote nodes. This is implicit in the communication model, which attempts to emulate the transitions of the transmitter by equivalent transitions in the receiver [Figure 3].

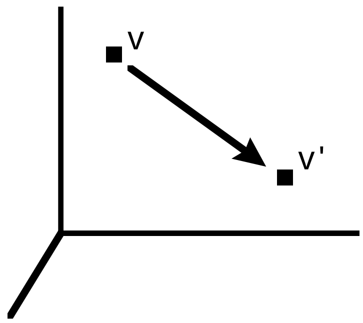

Figure 3: State space point transformation (translation of a point)

Here we are interested in extending Shannon's theory to it's application in high speed, wide area networks. As latency is one characteristic of these networks, we consider how to extent Shannon's model to account for latency, as it already accounts for error. One constraint of our extension is that, where latency is negligible, it should reduce to the original model; thus it will be an extension to the model. Other constraints are that the model be useful, i.e. that it describe the new domain effectively and that it enable the derivation of protocols which account for this increased latency. We also prefer the model to exhibit these characteristics by an extension which is minimal.

One of the fundamental results of Shannon's theory is that any amount of channel error (below 100%) can be removed by sufficient encoding. The effect of error compensation and reduction is to require encoding, which requires delaying the symbol stream by the length over which encoding is performed. As such, error reduction is traded for an increase in propagation delay. Here we examine the complement of this, where we will reduce latency by increasing the error across the channel; the error will be exhibited by the imprecision of information about remote nodes in the network.

3.The Mirage model

Mirage is a model for protocol analysis and design, in high speed, wide-area network environments [ToFa89,ToFa90]. It is based on representing remote nodes as volumes in state space, where data transmission and reception, as well as time evolution, are modelled as transformations on those subspace volumes.

Inherent in the Mirage model is the notion of latency as a measurable entity. We can extend the Shannon model of a communication channel, by including the latency measurement [Figure 4]. Here we assume that the latency will be either constant or predictable. If we describe flow in the Shannon model as a volume along the communication pipe, then latency is a measure of the length of the pipe. As such, incorporation of latency into the model reveals a spatial aspect to the formerly topographic Shannon model.

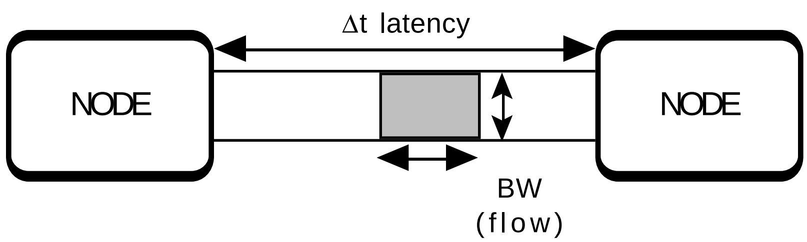

Figure 4: Mirage communication channel

In Shannon's model, information about remote nodes is modelled as a point in state space, and any operations which affect this information translate the point in space [Figure 3]. Here information is modelled as a volume in state space, where operations become transformations of that space. Time expands the volume of a space, transmission logically OR's the volume with a transformed copy, and reception collapses the volume to a sub-volume [Figure 5].

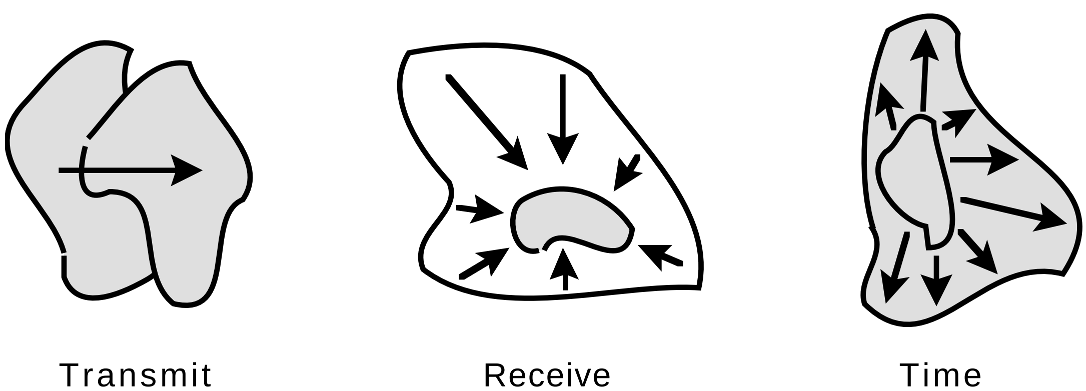

Figure 5: Mirage state space volume transformations

Mirage provides distinct transformations on the state space volumes for each of transmission, reception, and time lapse, as contrasted to the conventional single transformation of a point (i.e. translation). In each transform, there exist constraints related to link bandwidth and transmission latency related to the analysis of the imprecision of the volume as entropy. Mirage also indicates a stability condition, where network information is conventionally or entropically stable. This stability criterion forms the contraint which the model optimizes communication over [ToFa90]

Mirage works with three types of information, introducing the third to protocol models. The first two are Real Direct and Virtual Direct - the former is explicit communication, the latter is common knowledge, i.e. information derived from communication to others in a group when global group constraints are known [HaMo84]. Here we also utilize Virtual Indirect communication, information known from a combination of prior constraints and the absence of explicit communication [ToFa89].

3.1. Guarded messages

Also described in Mirage is the use of guarded messages, similar to the guarded commands, as used in programming languages. Since the perception of the remote node can be a volume in state space, we can send messages labelled with various regions of that volume, emitting multiple messages for a single desired action. The remote node compares its current local state to the label of the incoming message, and acts on the received information only if they match. It thus becomes possible to send information redundant messages. In current protocols, it is common to send identical copies of data, to reduce loss due to corruption. In this model, there is an indication that two messages with null-intersecting guards are information redundant, thus constructed to have equivalent results on the receiving node. They are data distinct, but thus information redundant [ToFa89].

Guarded messages permit the expansion of the state space of a transmitting node to be constrained sufficiently to remain tractable. The information redundancy used by these replicate communication paths also provides a use of the additional bandwidth provided by the high speeds of the problem domain. In effect, we then utilize the phenomenon which causes the effect - the high bandwidth*delay product - to compensate for it, by using the latent transmission to store replicates of the same information content.

3.2. Petri Net Representation

We need to recast an existing protocol model, based on the Shannon point model, to use state space volumes. We also require a notation for describing an instance, in order to discuss the model’s implications in the definite. For this purpose, we have selected

Timed Petri Nets [MeFa76], a variation of Petri Nets [Pete77, Petr62]. We extend these nets to describe the state transformations of Mirage, being careful to preserve the graphical/formal properties of the original nets.

First, we begin with the Timed Petri Net (TPN) of a protocol. We are specifically interested in the set of markings of this TPN, and the valid transitions between these markings, caleed the Token Machine (TM) of the TPN. Now consider the Timed Petri Net whose places correspond to the states of the TM, and whose transitions correspond to the arcs of the TM. This is is also a valid model of the protocol; we call this a Meta-Petri net, or MPN. Note that the MPN has a single token, which begins in the initial marking of the TPN, and moves along the MPN to represent the TPN's current (single) marking. We consider here transformations on the MPN which will enable the MPN to model multiple PN markings, under certain conditions. The number of tokens in the marking of a MPN reflects the entropy of the state of the node being modelled.

We consider the MPN where nondeterminism exists, and an advantage can be gained by running the MPN into the 'future', with virtual tokens. A *virtual token* is one of a set created which causes nondeterminism in a MPN. Rather than having a one token continue on a single path, a virtual token is created for each path possible. These tokens then belong to a *codependant* set. Later, if any of these tokens is to be considered real, *all codependants of that tokens set, and all ancestors of all tokens in that set*, must be destroyed. A token is *real* if it is the lone token in a MPN. These operations may appear similar to the virtualization which occurs in Feynman duagrams, which is appropriate, as they were patterned after concepts from quantum physics. Token virtualization and realization can be introduced by a graph transformation in the MPN. The messages emitted in the transformed MPN are guarded messages, where the conditional label on the message indicates which of the virtual tokens caused that message.

4.Comparisons to other models

Some existing protocols or protocol variations incorporate some aspects of the Mirage model, but none is as complete. Mirage attempts to unify aspects of several types of models, from distributed operating systems, partitioned databases, and general feedback and control systems. The following is a preliminary analysis of the Mirage components in relation to existing work.

4.1. Cybernetics / control theory

Mirage contains constraints of stability derrived from fundamental work in cybernetics [As64, Wi61]. These notions are also related to feedback and control theory, which relate to method in which Mirage attempts to describe the extent to which remote processes can communicate and share information (i.e. be synchronized), given feedback latency.

4.2. Models of communication

As previously stated, Mirage is an attempt to extend Shannon's communication theory to account for latency, and to compare the conjugate spaces of error and latency. In addition, there are aspects of existing protocol analysis which Mirage utilizes, specifically aspects of common and distributed knowledge, extending it to account for Virtual Indirect communication as well [HaMo84, Go88].

Constraints on communication links have been examined in the determination of optimal window sizes in windowed flow control schemes, especially where the interdependence between local link windows and global (overall network input/output) network windows is examined [Ak88, Lu88]. Mirage attempts to extend this search for optimal interdependency values to multi-dimensional systems. In particular, the issue of transmitter/receiver feedback as anticipating state space expansion has been used successfully in windowed flow control schemes, albeit in a discrete and restricted fashion, with respect to the level proposed here. The concept of expansion of state space upon data transmission, and collapsing upon message receipt, has been examined in the design of buffer "barriers", a modified flow control which attempts to equalize the uncertainty of communication among transmitter and receiver [Fr88].

The Mirage model is not suitable for direct computation, but equivalent substructures of the model may be, by a coarse-grained partitioning of the state space, via equivalence relations, examined as protocol projections [La82, Sh82], and partitions [ChMi86].

The incorporation of time into protocols is limited to two methods, where time is modelled either by boundaries or by finite time-steps. In the former, actions occur when these boundaries are exceeded, as in Timed Petri Nets, temporal logic, or time-out timers in more conventional models [MeFa76, Sc82]. In Mirage, time is a fully parametric value, a continuous entity, over which other entities vary.

Similarly, the modelling of time as an aging variable [Sh82] is not as general a concept. These time markers age, but other entities do not vary with time; aging is considered a static process. In the Mirage model, time is allowed to transform the subspaces arbitrarily (the transformation is known at the time it is invoked, but the model does not restrict the transformation a-priori). Incorporating time as a valid period for each state of a protocol machine is akin to denoting the interval over which the expansion of the subspace is well-defined [Ag83]. The extension of this method, which presents hold times of protocol states as cumulative distribution functions, is very much similar to a time extrapolation of this state space.

There also exists a wealth of information on state space models of communication protocols [Bo78, Da80]. These models are based on determinism, modelling a point in the state space, while Mirage is aimed at the a constrained nondeterministic version of this, where a volume of the state space is modelled.

4.3. Models of distributed systems / databases

The notion of restricting a machine to operate only within the valid subspace is an extension of distributed/replicated database techniques, most notably read/write quorum strategies [He86]. In addition, there have existed designs for external entities, which maintain the operation of a system to within some desired constraints [Fa76]. The use of these environments or supplemental programs to warn of dead-ends, maintain locality, and restrict other programs to within some valid subspace is similar to the methods used here. Mirage differs in that it appears that these notions are central to the operation of the protocol, not external, supplemental constraining devices.

4.4. Physics analogies

Finally, as noted earlier, Mirage exploits some analogies between latent communication and quantum interactions. The splitting of a point in state space into a volume is similar to the actions described by Feynman diagrams, and the deterimation of aggrgate characteristics of a system based on the distribution in the state space volume is similar to a path integral. In addition, there are analogues in thermodynamics and its relationship to quantum mechanics, beyond the existing analogue of information as negative entropy [ToFa89, ToFa90]

5. PLAN OF ACTION

The research proposed for Mirage has several facets, each complementing and extending the others. As an extension to Shannon's communication theory, it will provide a concrete underpinning for futher research in protocols where latency is a distinguishing factor. Mirage also provides several avenues for instantiation as a particular model, most notably in the further formalization of the Petri Net extension, i.e. proving the extensions preserve prior properties of the model, and characterizing the properties of the extensions. We will also apply the Mirage model in extensions of other protocol models, with an emphasis on determining whether existing models are special cases of Mirage. These include the state space models discussed earlier. In this way, a relationship can be determined, as to whether Mirage is a novel extension or, more importantly, whether it provides a model where unification of other existing models can be described.

The use of guarded messages in Mirage is analagous to the use of guarded commands in programming languages, and this will be investigated. There may be useful logical properties in these constructs, or in their equivalence to each other. Similarly, the bounds on computation expressed by the stability requirements permits formal analysis.

Mirage will also be applied to particular instances of protocols, specifically in the analysis of the new flow protocols [Zh??]. While these protocols are novel and useful, they often incorporate mechanisms for accumulating unused bandwidth, and arbitrarily resetting this accumulation. It is hoped Mirage will explain the need for such resetting mechanisms, and provide a method for a less arbitrary mechanism to implement them.

The implications of this research may also help develop new methods of communication. Currently, most protocols provide file transfer capabilities. The constraints which Mirage indicates may be required to exceed the bounds imposed by latency could indicate different types of communication, that which is more tightly-coupled. This may yield new protocols for real-time system management and distributed systems development, specifically in the area of alternate cache mechanisms.

We will also be examining other similarities between these models and quantum physics, to exploit analogues which might exist, specifically in the thermodynamic properties of quantum mechanics, and path integral constraints.

References

[Ag83] Aggarwal, S., and Kurshan, R. P., Modelling Elapsed Time in Protocol Specification, in: Rudin, H., and West, C. H. (eds.), Protocol Specification, Testing, and Verification, III (1983) pp. 51-62.

[Ak88] Akyildiz, I. F., Performance Analysis of Computer Communication Networks with Local and Global Window Flow Control, in: Proceedings IEEE Infocom (1988) pp. 401-410.

[As64] Ashby, W.R., An Introduction to Cybernetics, Methuen & Co. Ltd., London, 1964.

[Bo78] Bochmann, G.V., “FInite State Description of Communication Protocols,” Computer Networks 2, North-Holland, Amsterdam, 1978, pp. 361-372.

[ChMi86] Choi, T.Y. and Miller, R.E., “Protocol Analysis and Synthesis by Structured Partitions,” Computer Networks and ISDN Systems II, North-Holland, Amsterdam, 1986, [[. 367-381.

[Da80] Danthine, A., “Protocol Representation with Finite-State Models,” IEEE Transactions on Communications, Vol. COM-28, No. 4, Apr. 1980, pp. 632-643.

[Fa76] Farber, David J., and Pickens, John R., The Overseer: A Powerful Communications Attribute for Debugging and Security in Thin-Wire Connected Control Structures, in: Proceedings of IEEE International Communications Conference (1976)

[Fr88] Fraser, A. G., Current Research on Local Area and Wide Area Networks, in: Report on the Advanced Computer Communication Workshop Mar. 30-31, 1987 (Information Sciences Institute, CA, 1988) pp. 27-29.

[Go88] Gotzhein, R., “Knowledge-oriented consideration of communication protocols,” Protocol Specification, Testing, and Verification VIII, Elsevier, Amsterdam, 1988, pp. 295-306.

[HaMo84] Halpern, J.Y. and Moses, Y., “Knowledge and Common Knowledge in a Distributed Environment,” Tech. Report RJ-4421, IBM San Jose Research Laboratory, 1984.

[He86] Herlihy, M., “A Quorum-Consensun Replicaion Method for Abstract Data Types,” ACM Transactions on Computer Systems, Vol. 4, No. 1, Feb. 1986, pp. 32-53.

[La82] Lam, Simon S., and Shankar, A. Udaya, An Illustration of Protocol Projections, in: Sunshine, C. (ed.), Protocol Specification, Testing, and Verification (North-Holland, 1982) pp. 343-360.

[Lu88] Luan, Daniel T. D., and Lucantoni, David M., Throughput Analysis of a Window-Based Flow Control Subject to Bandwidth Management, in: Proceedings IEEE Infocom 1988, pp. 411-416.

[MeFa76] Merlin, P.M. and Farber, D.J., "Recoverability of Communication Protocols - Implications of a Theoretical Study," IEEE Transactions on Communications, Sept. 1976, p1036-1043.

[Pete77] Peterson, J.L., "Petri Nets," ACM Computing Surveys, Vol. 9, No. 3, Sept. 1977, p223-252.

[Petr62] Petri, C.A., "Communication with automata," thesis, Darmstadt Inst. of Technology, Bonn, Germany, 1962.

[Sc82] Schwartz, Richard L., and Melliar-Smith, P. M., From State Machines to Temporal Logic: Specification Methods for Protocol Standards, in: Sunshine, C. (ed.), Protocol Specification, Testing, and Verification (North-Holland, 1982) pp. 3-19.

[Sh82] Shankar, A. Udaya, and Lam, Simon S., On Time-Dependent Communication Protocols and Their Projections, in: Sunshine, C. (ed.), Protocol Specification, Testing, and Verification (North-Holland, 1982) pp. 215-233.

[ShWe63] Shannon, C.E. and Weaver, W., The Mathematical Theory of Communication, Univ. of Ill. Press, Urbana IL, 1963.

[ToFa89] Touch, J.D. and Farber, D.J., "Mirage: A Model for Ultra-High-Speed Protocol Analysis and Design," Protocols for High-Speed Networks, North-Holland, Amsterdam, 1989, p115-134.

[ToFa90] Touch, J.D. and Farber, D.J., "Mirage: An Introduction to the Formal Model," submitted to the IFIP WG 6.1/6.4 Workshop on Protocols for High Speed Networks, Nov. 28, 1990.

[Wi61] Wiener, N., Cybernetics: or Control and Communication in the Animal and Machine, MIT Press, Cambridge MA, 1961.

[Zh??] Zhang, L. ??????????