Eric Jonsson School of Engineering and Computer Science

1999-5-26

Using Power Control to Build Neighborhoods in Ad-hoc Networks: Shouting!

Don Montgomery and Sarvesh Kulkarni

Using Power Control to Build Neighborhoods in Ad-hoc Networks: Shouting!

Don Montgomery & Sarvesh Kulkarni
Department of Computer Science
The University of Texas at Dallas

Technical Report UTDCS-03-99
May, 1999

Using Power Control to Build Neighborhoods in Ad-hoc Networks: Shouting!

Don Montgomery* Sarvesh Kulkarni

May 26, 1999

Computer Science Program
Erik Jonsson School of Engineering and Computer Science
The University of Texas at Dallas
{`kulkarni,montgo`}@utdallas.edu

Abstract

The past few years have seen a tremendous surge of interest in wireless mobile (ad-hoc, or nomadic) networks. In many cases, mobile nodes are powered by batteries, therefore their uptime depends on how efficiently they can utilize the limited energy. Some current research in ad-hoc networks does address power control per se, but we have not found work that centers on the use of dynamic power control to reduce the power expenditure at each node during the setup and maintenance of communication links.

We present a new mechanism, *Shouting*, as a tool to progressively build one-hop neighborhoods in ad-hoc networks in an energy-efficient manner. Additionally, we present a fully distributed and scalable link-layer protocol, based on Shouting, to build and maintain such energy-efficient neighborhoods. Our protocol has the added advantage that it can respond to higher-level protocols for more effective route discovery and maintenance, enhancing reliability.

1 Introduction

Much thought is now being given to ad-hoc (nomadic) networking environments, consisting entirely of independent identical mobile nodes, that is, mobile networks with no base stations, which must cooperate to route messages across a dynamically changing multi-hop topology [7][8][18]. The most frequently noted scenarios are military (battlefield, tactical) applications and rescue/disaster relief. Clearly, though, the envisioned functionality extends to any kind of logistic enterprise involving scattered and/or semi-isolated units at work in the field, pursuing a common goal: military, public service, or commercial. Conferences, markets, and festivals will also benefit from easily deployed, self-configuring wireless mobile networks. For example, each mobile host could be a notebook PC, and various hosts could exchange site-maps and use voice communication over the network.

In the following paper, we first look at available literature in Section 2. We present our model and assumptions in Section 3. The problem and goals are outlined in Section 4, followed by definitions in Section 5, and the protocol in Section 6. A simple upper bound for the Shouting mechanism is provided in Section 7. Conclusions are presented in Section 8. Future directions are considered in Section 9. Three appendices contain supplementary material.

---

*Work supported in part by a Texas Telecommunications Engineering Consortium (TxTEC) Fellowship.

2 Survey of existing literature

Since mobile nodes have limited battery-supplied energy for operation, the uptime for such nodes is expected to be small, usually a few hours at most before recharging. This limited energy supply is one of the chief constraints on most ad-hoc networks. A survey of current literature in the field of mobile networking reveals that there has been considerable effort directed toward dynamically controlling the transmission power of mobile hosts in cellular networks [11][17]. This has allowed cellular systems to function with less interference between hand-held units, and resulted in enhanced use of available bandwidth. However, for ad-hoc environments, not much effort has been made to use dynamic power-control methods to establish and maintain a reliably operating set of network links, actively seeking to reduce power consumption with the objective of increasing the average uptime of each node across the system.

Some concern for power-control has been embodied in schemes to reduce interference in CDMA (Code Division Multiple Access) in ad-hoc environments by means of artificial power capture [12], and in efforts to characterize network capacity of CDMA networks with power control [20][21]. CDMA schemes typically try to balance transmission signal strength among nodes transmitting on the same channel. One method of controlling topology [9] has been advanced which uses power control in pursuit of enhanced link layer reliability. It is based on an assumption of available directional/geographical information associated with each link, and is primarily oriented toward robustness in the face of frequent node failure. Some attention has been given to the benefits of "spatial reuse of bandwidth" accruing to systems using fixed limited radius of transmission across the network [7][13][14]. Certainly it is unrealistic to presume that all nodes will stay sufficiently far apart to avoid interfering with neighboring transmissions in a scheme that limits the number of channels or frequencies available. But, to the authors' knowledge, the dynamic control of transmission signal strength, and thus transmission radius, has thus far not been considered as it specifically relates to efficient reuse of bandwidth across an ad-hoc wireless system.

Further (perhaps because the hiding of complexity within a network protocol layer from adjacent network protocol layers is universally acknowledged to be useful and necessary), the idea of making the link layer more accessible and responsive to the network layer, and specifically, to routing routines, seems not to have been discussed. Especially in ad-hoc networks, where links are tenuous, both in quality and existence, the fragile subnet can benefit from increased ability to control the physical infrastructure in response to routing needs. Finally, as useful as the assumption of synchronous time is in theory, in practice it may be difficult to achieve over wireless networks, as is tacitly acknowledged in discussions surrounding the proposed IEEE 802.11 wireless LAN standard [3][4][6][10]. Indeed, a progression away from synchronous schemes may be discerned in [7][13][14].

In this paper, we present a solution to these problems using a novel tool: the Shouting mechanism. We present a link-layer protocol which incorporates this Shouting mechanism in a fully distributed and greedy fashion, attempting to reduce the power consumption at each node without affecting the reliability of the network. We anticipate that for a given power budget, an average node in the system can show increased uptime as compared to the system with no power-control. Also our link-layer protocol is fully scalable both in the number of nodes and in the density and the geographical distribution of the nodes. Additionally, it provides modular procedures which interface seamlessly with higher-level routines useful in route searching and maintenance. Lastly, the protocols are tailored to be effective and energy-efficient in an asynchronous network.

Our solution is based on the premise that if each node cooperatively reduces its power requirement by using fewer and less power-hungry connections (without disconnecting parts of the established network), we get power savings across the entire network, which is beneficial as long as no node (or set of nodes) is unduly burdened, any more than it would have been under fixed-power constraints.

3 System model and assumptions

We assume that all nodes have a unique ID, are fully mobile and are identically equipped. Computational time at a node is deemed insignificant. Each node can measure its own transmission signal strength (TSS) as well as the Signal-to-Interference Ratio (SIR) of a received signal within some degree of accuracy [17] and encode this measure in its packets. The TSS at the sender can be adjusted according to messages coming from the receiver informing the sender of the required TSS. (For further discussion, see appendix A)

We further assume that there are no base stations, and that each node operates as a host and a switch in a cooperative environment. There is no global synchronization, though measured time intervals are assumed to be of equal duration for all nodes. Two common channels exist: a Shout channel and a Control channel. The Shout channel is common across the network whereas the Control channel is common within a locality or neighborhood. Both the Shout and the Control channels are asynchronous contention channels employing a form of CSMA/CA (modification of proposed IEEE 802.11 wireless LAN standard). Data channels are simplex, and are negotiated in pairs by the nodes on either end of the link. We refer to such a pair of data channels as a 'link.' Note that if there are not too many nodes in a 'neighborhood,' and there exist some reasonable number of channels, then the channels may be allocated in such a way that there is minimal interference over the data links.

Each node can vary its signal strength independently on each channel (i.e., half-link) from zero to some maximum. Since the nodes are battery-powered, their energy supply is limited. (This is a reason that they are subject to frequent downtime.) Transmission is very expensive in terms of energy expended, as compared to reception. Our core assumption is that multi-hop routing is preferable to single-hop routing in that it is more power-efficient, since the TSS required increases polynomially as transmission signal radius increases linearly. In addition, multi-hop routing promotes frequency reuse by reducing transmission signal radius [14]. At this time, we do not deal with contention, interference, capture, or hidden stations. Finally, we assume the existence of efficient channel allocation and routing algorithms.

4 Problem formulation and goals

Given a set of n nodes, we would like to establish some given level of connectedness at the link-layer of the ad-hoc network, expressed in the number of links (degree) per node. Should certain links be lost for some reason, say, due to node movement, we would like to reactively establish new links, and restore the same level of connectedness as far as possible. In addition, we prefer to employ power-frugal methods to establish new links, and make power-frugal choices regarding which links we use for data communications, without limiting network scalability. In doing so, we intend to reduce average power consumption at the nodes, and, in the process, provide a mechanism for the dynamic guidance of our link-layer protocol by higher level protocols for routing and reliability.

5 Definitions

First we define the terms used in our protocol. The 'weight' of a link is the optimal signal strength at the transmitter. The 'expense' of a link is a function of its weight, traffic and lossiness, the last two parameters being unknown for a new link.

- Let α, β, and γ be system constants, where α and β are given in time units, and γ is an iterator
- Let α be the length of a Shout-type
- Let β be the period of a Shout within a Shout-burst
- Let γ be the number of Shouts in a Shout-burst
- Let $t_{LS} = \alpha + \beta$ be the length of a Listening Slot (LS)
- Let $t_{SB} = (\gamma - 1)\beta + \alpha$ be the length of a Shout-burst
- Let $\gamma\beta$ be the period of a Listening Slot
- Let $t_I = (\gamma - 1)\beta - \alpha$ be the interval between Listening Slots

A Shout-type can either be a Shout or an ACK. Refer to figs. 1, 2, and 3 for illustrations of the Shout-burst and listening slot.

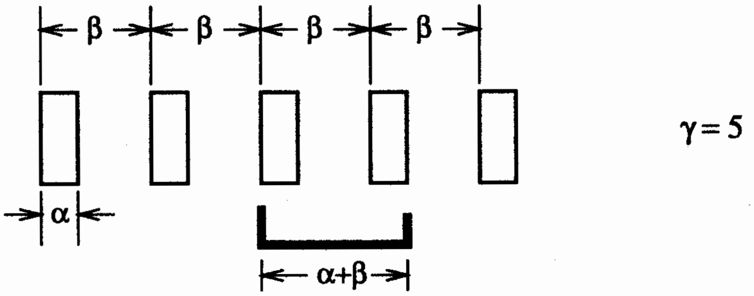

Figure 1: A single Shout burst

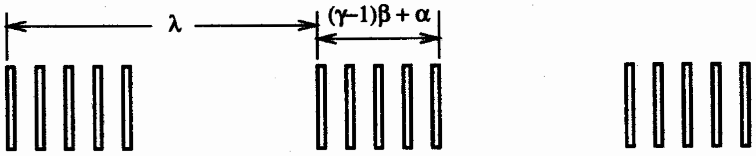

Figure 2: A string of Shout bursts

6 The protocol

6.1 Top level description

At the top level, our protocol can be summarized as follows.

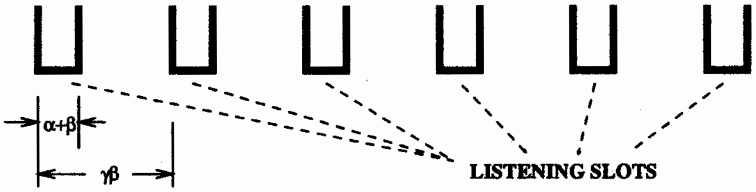

Figure 3: Listening slots

1. Node i Shouts on the Shout channel with its power-level encoded in each Shout.
2. Node j receives a Shout, measures the received SIR, and ACKs on the Shout channel specifying the frequency of the Control channel on which to rendezvous, and the estimated ideal transmission power needed from node i.
3. Node i responds by moving to the Control channel and transmitting the estimated ideal transmission power it wants from node j within its Control packet.
4. Node j moves to the Control channel, receives node i's Control packet and negotiates with node i.
5. Negotiations conclude as the nodes agree on the status of the link between them.

6.2 One-hop neighborhood construction and maintenence

The state diagrams of figs. 7 and 8 in appendix C show the state transitions. The protocol itself is outlined by flowcharts. The top level diagram of the algorithm is shown in fig. 4. The details of the loop, search, link addition, and link deletion modules are shown in appendix C in figs. 9, 10, 11, and 12.

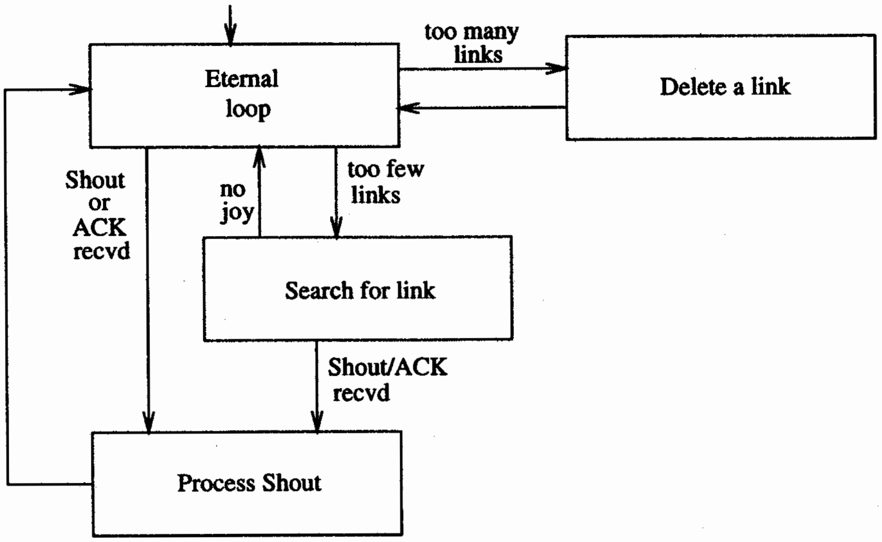

Figure 4: Top-level view of protocol

The components of a Shout and an ACK are now given.

Shout()

- node ID#, TSS, Shout flag, urgent flag

ACK()

- node ID#, TSS, ACK flag, Shouter ID#, SIR, CTL channel

6.3 Remarks

6.3.1 Avoiding livelock

In the Safe-to-Delete() procedure (fig. 12), we assume that if either of the concerned nodes insists on retaining a costly link, the other has to comply. This gives rise to a pathological scenario in which we may have a circular dependency among some nodes. For instance, consider the case where A wants to delete its link to B, B wants to delete its link to C and C wants to delete its link to A. If all of them refuse one another that privilege, we have a circular dependency. Such a condition could certainly lead to flailing, and calls for the use of some method, perhaps an adaptation of an edge-chasing algorithm for deadlock detection and resolution, e.g., a priority-based probe scheme [1][2][5][16][19]. Parenthetically, we note that this remedy would have to be carried out on what is normally considered the network level, and so we envision it as a mandatory routine embedded in the Safe-to-Delete() function, and thus transparent to the network layer.

6.3.2 'flag'

The 'flag' in figs. 9 and 10 is set when the transmission radius is equal to the maximum radius, i.e., the last Shout was at full power. This flag is used to allow for the case where a node is isolated and of low degree. We don't want such a node to be Shouting in a futile (energy-wasting) or dangerous (e.g., in a military scenario) manner.

6.3.3 Assigning values to system constants

We note that the Shouting mechanism itself is designed to place the onus of continual activity upon the Listener, and every node is a Listener during its execution of the Eternal Loop (figure 4a). Only when additional links are required is Shouting performed. Since the system constants α, β, and γ determine the actual rate at which the distributed parts of the Shouting/Listening routine are iterated, they may be initialized across the system at levels that are convenient for a given situation. Especially note that the 'sleep' phase of the 'Eternal Loop' given in figure 4a is where a node with only one tuner would actually perform its normal (non-link-layer) activities. The 'sleep' interval is of duration t_I, which is entirely dependent upon the settings of the system constants.

7 Proof of the Shouting mechanism

- Main Assertion: Given a Shout-burst, any Listener within 'hearing range' must 'hear' at least one Shout in that burst.

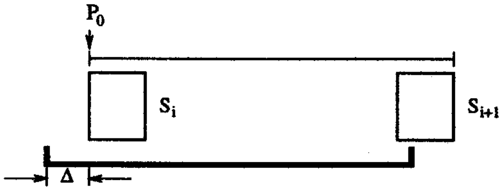

Figure 5: Proof, part 1

Proof of correctness of the Shouting mechanism is given as follows.

I. Assertion A: if a Listening Slot (LS) occurs within a Shout-burst, a Listener must 'hear' at least one Shout. (Refer to fig. 5.)

case 1) $\Delta = 0$: LS straddles two Shouts exactly, and both are 'heard'
case 2) $0 < \Delta < \beta =$ period of Shout: LS 'hears' S_{i+1}
case 3) $-\beta < \Delta < 0$: LS 'hears' S_i

The three cases are exhaustive.

II. Assertion B: a Shout cannot slip 'unheard' between two Listening Slots. (Refer to fig. 6.)

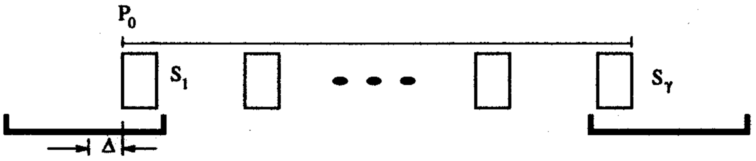

Figure 6: Proof, part 2

case 1) $\Delta = 0$: Shout straddles two LSs by α on each end,
therefore S_1 is 'heard' by LS_i and S_γ is 'heard' by LS_{i+1}
case 2) $0 < \Delta < \gamma\beta =$ period of LS: LS_{i+1} 'hears' one of $S_1, \dots, S_\gamma$
case 3) $-\gamma\beta < \Delta < 0$: LS_i 'hears' one of $S_1, \dots, S_\gamma$

The three cases are exhaustive.

III. Since Assertion A and Assertion B are exhaustive for the Main Assertion, the Main Assertion is proved. Given two nodes within 'hearing range' of each other, each following the algorithm, if one Shouts, the other will 'hear' within a maximum time of $(\gamma - 1)\beta + \alpha$.

8 Conclusion

We have provided a distributed algorithm based on an original idea, the Shouting mechanism. This algorithm provides a way for nodes in an ad-hoc network to establish and maintain links with a lower power expenditure than the case in which there is no power-control. If two nodes are within hearing range of each other, if one Shouts, the other will hear within a maximum time of $(\gamma-1)\beta+\alpha$ in the no-contention scenario. Should a node lose any or all of its links on account of its motion or the failure of its neighbors, the routing algorithm on top of our protocol can induce the node to Shout for more (new) connections through our protocol. Since our protocol is fully distributed and assumes no knowledge of the overlying network (logical or routing) topology, it is flexible, scalable, and transparent. The protocol is just short of implementation since it already has the requisite granularity of detail.

9 Future directions

The next logical step is to verify this algorithm for the contention scenario on the Shout and the Control channels. Simulation is the obvious approach to test the existing protocol. Additionally, the model must be extended to take into account the phenomena of hidden stations, capture, and interference. Simulation is also required to determine reasonable values for the system constants, and to estimate throughput, stability, and actual power savings for specific offered system loads in contrast to conventional systems.

References

[1] K. M. Chandy, J. Misra, "A distributed algorithm for detecting resource deadlocks in distributed systems," Proc. ACM SIGACT-SIGOPS Symp. Principles of Dist. Computing, (Ottawa, Canada, August 18-20, 1982), ACM, New York, 1982, pp. 157-164.

[2] K. M. Chandy, J. Misra, L. M. Haas, "Distributed deadlock detection," ACM Trans. Comput. Syst., vol. 1, May 1983, pp. 144-156.

[3] H.S. Chhaya, S. Gupta, "Throughput and fairness properties of asynchronous data transfer methods in the IEEE 802.11 MAC protocol," 6th IEEE Int'l. Sym. Personal, Indoor, Mobile Computing., PIMRC, 1995, pp. 613-617.

[4] H.S. Chhaya, "Performance modeling of asynchronous data transfer methods of IEEE 802.11 MAC protocol," Wireless Networks, v. 3, 1997, pp. 217-234.

[5] A. L. Choudhary, W. H. Kohler, J. A. Stanovic, and D. Towsley, "A modified priority-based probe algorithm for distributed deadlock detection and resolution," IEEE Trans. Software Eng., vol. 15, Jan. 1989, pp. 10-17.

[6] B.P. Crow, I. Widjaja, J.G. Kim, P.T. Sakai, "IEEE 802.11 wireless local area networks," IEEE Comm., Sep. 1997, pp. 116-126.

[7] Gerla, M., J.T.-C. Tsai, "Multicluster, mobile, multimedia radio network," Wireless Networks, v. 1, 1995, pp. 255-265.

[8] Z.J. Haas, "A new routing protocol for the reconfigurable wireless networks," http://www.ee.cornell.edu/ haas/wnl.html, 1997.

[9] L. Hu, "Topology control for multihop packet radio networks," IEEE Trans. Comm., v. 41, no. 10, Oct. 1993, pp. 1474-1481.

[10] S. Jiang, D.H.K. Tsang, "On architectures for broadband wireless systems," IEEE Comm., Oct. 1997, pp. 132-140.

[11] T. Lee, J. Lin, "A Fully Distributed Power Control Algorithm for Cellular Mobile Systems," IEEE Journal on Selected Areas in Comm., vol. 14, No.4, May 1996, pp.

[12] Leung, Y.-W., "Mean power consumption of artificial power capture in wireless networks," IEEE Trans. Comm., vol. 45, Aug. 1997, pp. 957-964.

[13] Lin, C.R., M. Gerla, "Multimedia transport in multihop dynamic packet radio networks," IEEE GlobeCom, (1995), pp. 209-216.

[14] Lin, C.R., M. Gerla, "Adaptive clustering for mobile wireless networks," IEEE JSAC, vol. 15, Sep. 1997, pp. 1265-1275.

[15] Pagani, E., G.P. Rossi, "Reliable broadcast in mobile multihop packet networks," ACM Mobicom, 1997, pp. 34-42.

[16] M. Roesler, W. A. Burkhard, "Resolution of deadlocks in object-oriented distributed systems," IEEE Trans. Computers, Aug. 1989, pp. 1212-1224.

[17] J. Rulnick, N. Bambos, "Mobile Power Management for Wireless Communication Networks," Wireless Networks, vol.3, No.1, March 1997, pp.3-14

[18] Scott, K., N. Bambos, "The self-organizing wireless network (SWAN) protocol for communication among mobile users," IEEE GlobeCom, 1995, pp. 355-359.

[19] M. K. Sinha, and N. Natarajan, "A priority-based distributed deadlock detection algorithm," IEEE Trans. Software Eng., vol. SE-11, Jan. 1985, pp. 67-80.

[20] C.W. Sung, and W.S. Wong, "Power control for multirate multimedia CDMA systems," IEEE Infocom, Mar. 1999.

[21] J. Zhang, and E.K.P. Chong, "CDMA systems with Random spreading in fading channels: network capacity and power control," IEEE Infocom, Mar. 1999.

A Calculation of desired TSS

The basis for a calculation of desired TSS (transmission signal strength is as given as follows. At time t, let P_i be the TSS at node i, and let SIR_{ij} be the SIR received at node j for a signal from node i. Then, for some node m and some node n, if m shouts and n receives the shout, we calculate

$$f(P_m, SIR_{mn}) = \rho_{mn},$$

where ρ_{ij} is the attenuation factor over link ij. In simpler terms, $P_m/SIR_{mn} = \rho_{mn}$.

Let SIR'_{ij} be the desired SIR at node j for a TX from any node i. Then the estimated desired TSS at m for a transmission to n is estimated by node n as:

$$P'_m = g(\rho_{mn}, SIR'_{mn}) = \rho_{mn} * SIR'_{mn}.$$

B Usage of Shout and Control channels

We expect the traffic on the Shout channel to be sparse, due to (1) relatively infrequent link acquisition, (2) the short time required for its use, and (3) the probable natural distribution of need. Therefore, we anticipate no great problems arising from the use of contention on the Shout channel. Nevertheless, the implementation of something akin to a binary exponential backoff algorithm should be considered provident.

Although we do not require it, we envision the Control channel being used not only for the original link negotiations and channel allocation, but also for making reservations to send data along the links among the nodes, much as the control channel in a fiber-optic MAN. The reason we consider reservations is that it is unlikely that a node would be equipped with as many tuners as it would have available channels, and therefore the node would need to be listening to the correct channel to receive a transmission. The Control channel traffic should be relatively sparse in any given neighborhood, due to the usage of the Shout protocol and its inherent tendency to limit signal radii, and hence the amount of local transmission coverage overlap. The RTS/CTS protocol put forth in the proposed IEEE 802.11 standard should be readily adaptable to this situation. This constitutes an area for further work.

C Supplementary Figures

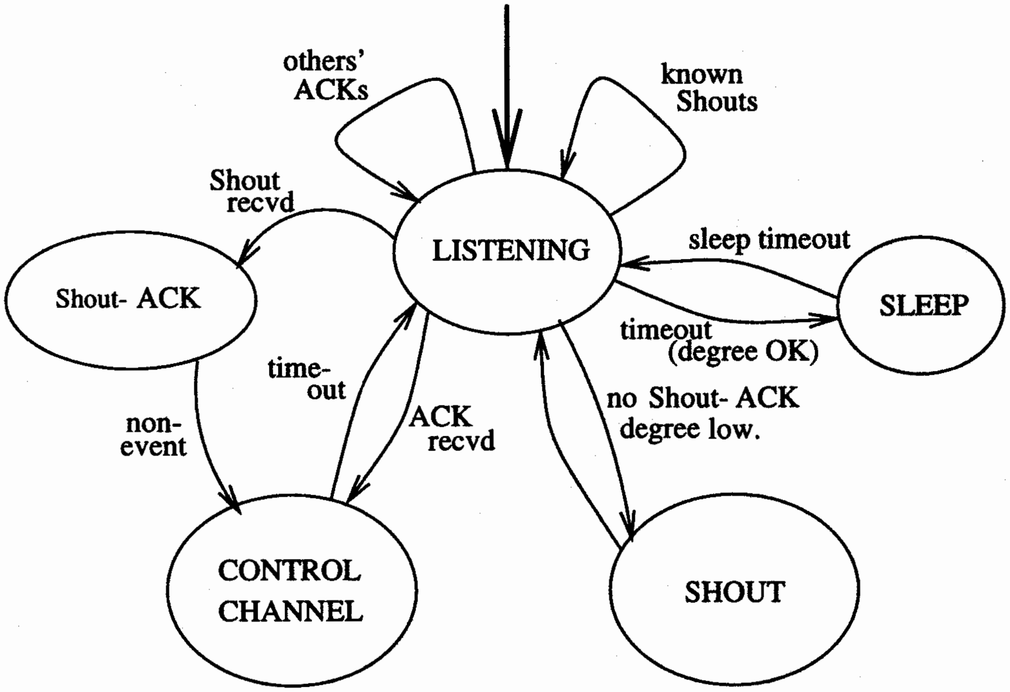

Figure 7: Shout channel state transitions

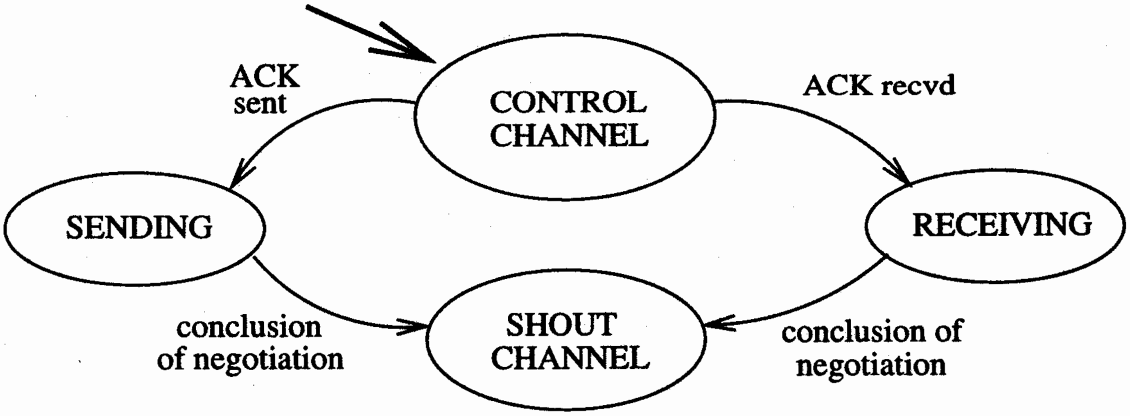

Figure 8: Control channel state transitions

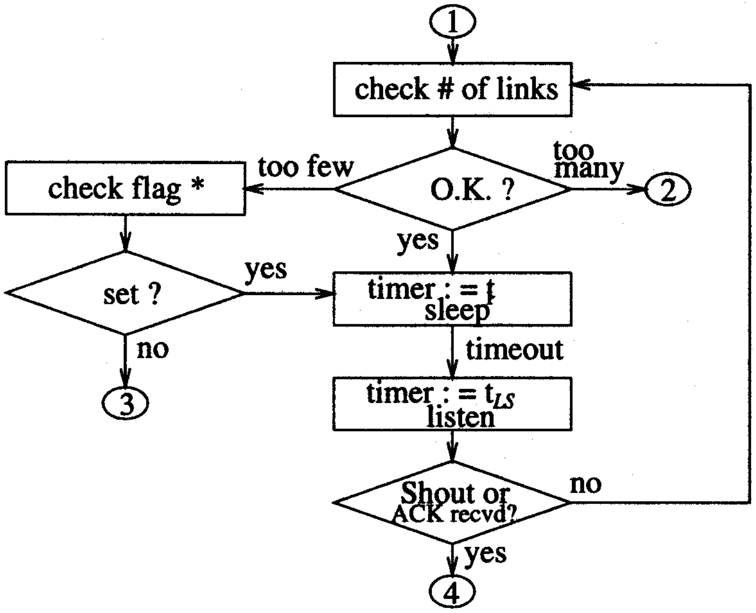

Figure 9: The "Eternal Loop"

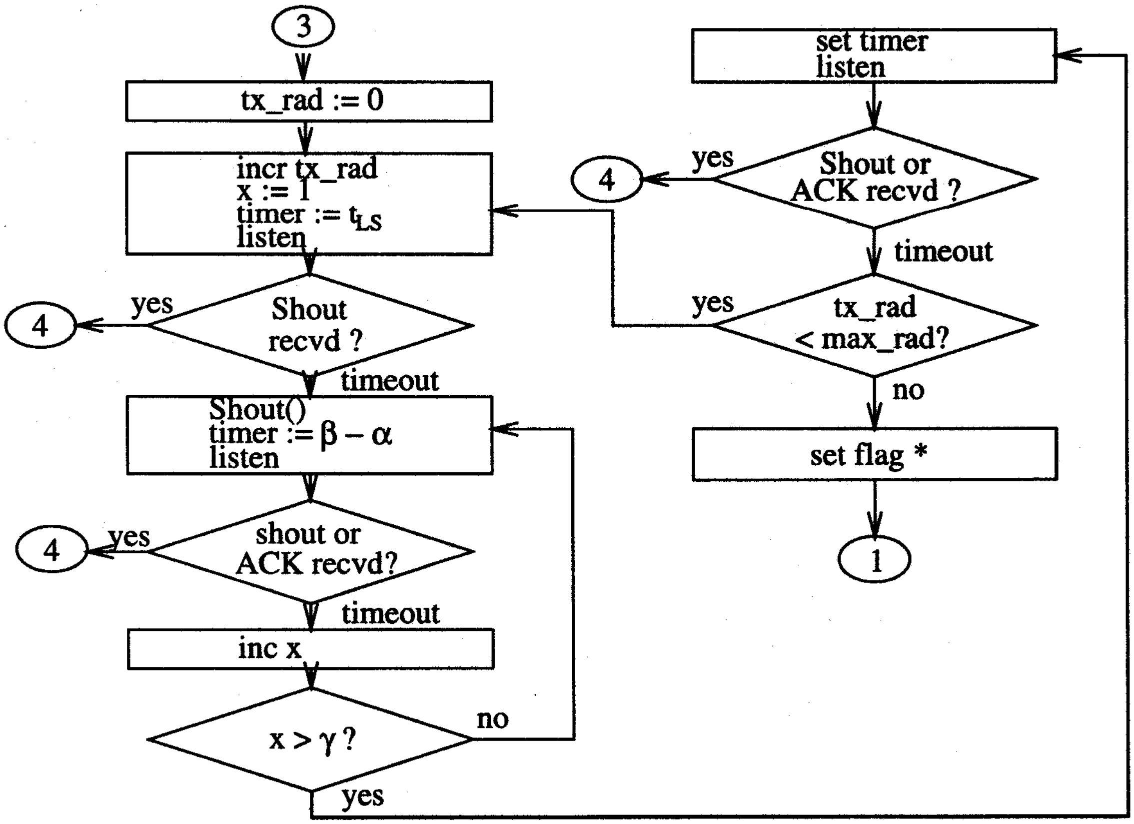

Figure 10: Search for link

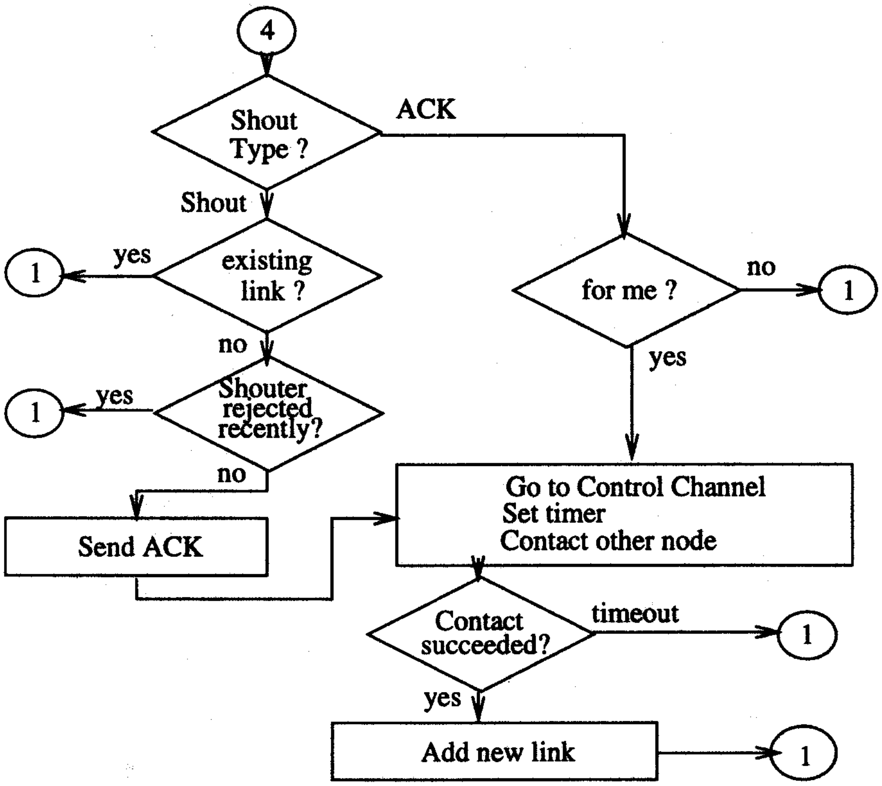

Figure 11: Process a Shout

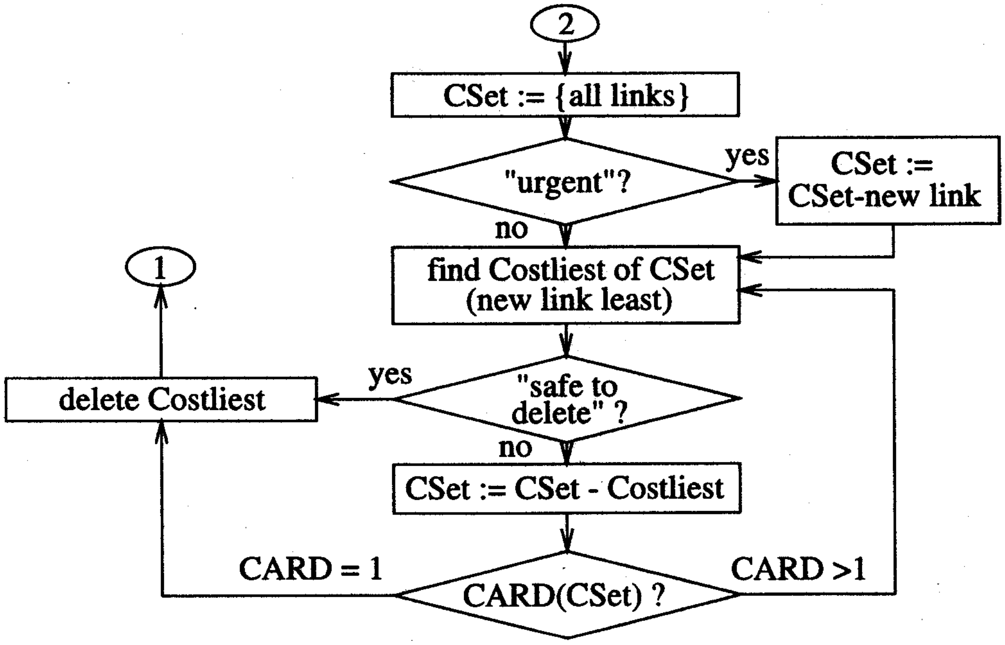

Figure 12: "Delete a link"