

AD-A110 948 MASSACHUSETTS INST OF TECH CAMBRIDGE LAB FOR INFORMA--ETC F/G 12/1
A RELIABLE BROADCAST ALGORITHM, (U)
UNCLASSIFIED JAN 82 A SEGALL, B AWERBUCH N00014-75-C-1183
LIDS-P-1177 NL

1 OF 1
ADA
1-09-82

END
DATE
FILED
03-82
DTIC

LEVEL II

12

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
	AD-A110 948	
4. TITLE (and Subtitle)	5. TYPE OF REPORT & PERIOD COVERED	
A Reliable Broadcast Algorithm	Paper	
6. AUTHOR(s)	7. PERFORMING ORG. REPORT NUMBER	
Adrian Segall and Baruch Awerbuch	IDS-P-1177	
	8. CONTRACT OR GRANT NUMBER(s)	
	ARPA ORDER No. 3045/5-75 ONR/N00014-75-C-1183	
9. PERFORMING ORGANIZATION NAME AND ADDRESS	10. PROGRAM ELEMENT PROJECT TITLE AREA & WORK UNIT NUMBERS	
MIT Laboratory for Information and Decision Systems Cambridge, MA 02139	Program Code No. 5T10 ONR Identifying No. .049-383	
11. CONTROLLING OFFICE NAME AND ADDRESS	12. REPORT DATE	
Defense Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, Virginia 22209	January 1982	
	13. NUMBER OF PAGES	
	21 pages	
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)	15. SECURITY CLASS. (of this report)	
Office of naval Research Information Systems Program Code 437 Arlington, Virginia 22217	Unclassified	
	15a. DECLASSIFICATION DOWN GRADING SCHEDULE	
16. DISTRIBUTION STATEMENT (of this Report)		
Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
E		
18. SUPPLEMENTARY NOTES		
This paper has been submitted to the IEEE Transactions on Communications.		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number)		
Broadcast in a communication network is the delivery of copies of messages to all nodes. A broadcast algorithm is reliable if all messages reach all nodes in finite time, in the correct order and with no duplicates. The present paper presents an efficient reliable broadcast algorithm.		

AD A110948

DTIC FILE COPY

DTIC
ELECTE
FEB 16 1982
E

DD FORM 1473 1 JAN 73 EDITION OF 1 NOV 65 IS OBSOLETE

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

82 02 16 001

A Reliable Broadcast Algorithm

Adrian Segall and Baruch Awerbuch

Department of Electrical Engineering
Technion - Israel Institute of Technology
Haifa, Israel

ABSTRACT

Broadcast in a communication network is the delivery of copies of messages to all nodes. A broadcast algorithm is reliable if all messages reach all nodes in finite time, in the correct order and with no duplicates. The present paper presents an efficient reliable broadcast algorithm.

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By _____	
Date _____	
Distribution Codes	
and/or	
Dist	Special
A	

The work of A. Segall was performed on a consulting agreement with the Laboratory for Information and Decision Systems at MIT, Cambridge, Mass., and was supported in part by the Office of Naval Research No. ONR/N00014-77-C-0532 and in part by the Advanced Research Project Agency of the US Department of Defense (monitored by ONR) under Contract No. N00014-75-C-1183.

This paper has been submitted to the IEEE Transactions on Communications.

1. INTRODUCTION

Broadcast multipoint communication is the delivery of copies of a message to all nodes in a communication network. In a network with mobile subscribers, for example, the location and connectivity to the network of such subscribers may change frequently and this information must be broadcast to all nodes in the network, so that the corresponding directory list entry can be updated. Broadcast messages are used in many other situations, like locating subscribers or services whose current location is unknown (possibly because of security reasons), updating distributed data bases or transmitting information and commands to all users connected to the communication network.

There are certain basic properties that a good broadcast algorithm must have and the most important are: a) reliability, b) low communication cost, c) low delay, d) low memory requirements. Reliability means that every message must indeed reach each node, duplicates, if they arrive at a node, should be recognizable and only one copy accepted, and messages should arrive in the same order as transmitted. Communication cost is the amount of communication necessary to achieve the broadcast and consists of, first, the number of messages carried by the network per broadcast message, (broad-
cast communication cost), second, the number of control messages necessary to establish the broadcast paths, (control communication cost), and, third, the overhead carried by each message (overhead cost). Low delay and memory are basic requirements for any communication algorithm, and broadcasts are no exception.

The broadcast communication cost is minimized if the algorithm uses spanning trees, but normally there is need for a large control communication cost in order to establish and maintain these trees. However, the control cost can be reduced considerably provided that the routing mechanism in the network constructs routing paths that form directed trees towards each destination, in which case these trees can be used in the reverse direction for broadcast purposes. This general idea is presented in [1], but the authors show that the proposed algorithms named reverse path forwarding and extended reverse path forwarding are not reliable when the routing algorithm is dynamic, since in this case nodes may never receive certain messages, duplicates may be received and accepted at nodes, and the order of arriving messages may not be preserved. As said before, in order to be efficient, the above mentioned algorithms require that the routing paths to each destination are directed trees. An adaptive routing algorithm that maintains at all times spanning directed trees rooted at the destination has been proposed in [2] and throughout the present paper we assume that the protocol of [2] is the underlying routing algorithm in the network. However, for the reasons stated before, namely the fact that the routing paths are dynamic, the broadcast algorithm of [1] is unreliable even if applied to the routing procedure of [2].

The purpose of the present paper is to propose and validate an algorithm whose main property is that the broadcast propagating on the tree provided by the routing protocol of [2] is reliable. It is convenient for the purpose of our discussion to separate the property of reliability into two parts: completeness means that each node accepts broadcast messages in the order released by their origin node, without duplicates or messages missing, while finiteness is the property that each broadcast message is indeed accepted at each node in finite time after its release.

As mentioned and exemplified by the authors, the algorithms of [1] are neither complete nor finite. In the algorithm of the present paper, completeness is achieved by requiring nodes to store broadcast messages in the memory for a given period of time and by introducing counter numbers at the nodes. Finiteness is obtained by attaching a certain impeding mechanism to the routing protocol. We may mention here that it is relatively easy to make a broadcast algorithm reliable if one allows infinite memory, unbounded counter numbers and infinite overhead in the broadcast messages. However, the properties that make our algorithm tractable are: bounded memory, bounded counter numbers, no overhead carried by broadcast messages (in form of counter numbers or any other kind) and the fact that the impeding mechanism is not activated most of the time.

In the rest of the paper we proceed as follows: Sec. 2.1 contains a brief description of the routing algorithm of [2]. Sec. 2.2 and 2.3 build the reliable broadcast protocol step by step, while its final form and main properties are given in Sec. 3. The proofs of the main theorems are included in the Appendix.

2. THE BROADCAST PROTOCOL

2.1 The Routing Protocol

The underlying routing protocol considered in this paper is The Basic Protocol of [2]. In summary, this protocol proceeds in updating cycles triggered and terminating at the destination node named SINK. An updating cycle consists of two phases: a) control messages propagate uptree from SINK to the leaves of the current tree and each node i performs this phase whenever it receives a control message MSG from its current preferred neighbor p_i ; b) control messages propagate downtree, while new preferred neighbors are selected and this phase is performed at node i upon detecting receipt of MSG from all neighbors.

In the following sections we shall need to identify the updating cycles and it is convenient to attach to each cycle a counter number α . For the time being α will be unbounded, but later we shall show that a binary variable is sufficient. The routing protocol is specified by the following algorithm.

Routing Algorithm for node i (RA)

1. For MSG (ℓ, α)
2. $N_i(\ell) + 1$
3. if $\ell = p_i$, then: $\alpha_i + \alpha$; send MSG(α_i) to all $\ell \in G_i$, except p_i
4. if $\forall \ell' \in G_i$, holds $N_i(\ell')=1$, then: send MSG(α_i) to p_i ; select new p_i ; $\forall \ell' \in G_i$, set $N_i(\ell')=0$.

Here G_i denotes the set of neighbors of i and "For MSG (ℓ, α)" means "when receiving MSG(α) from neighbor ℓ , perform.....". Subscript i indicates variables at node i and corresponding variables without subscript indicate variables in the received message. We have deliberately suppressed from the algorithm of [2] all variables that are not directly relevant to the

broadcast (like d_i , d , $D_i(\ell)$) and have not explicitly indicated the procedure for selecting the new p_i because it is not important for our purpose, except for the property that it maintains at all times a directed spanning tree rooted at SINK. For simplicity p_i will be called the father of i . The algorithm is indicated for a given SINK that is not specified explicitly (again for simplicity of notation) and that becomes the source of the broadcast messages. The SINK performs the following algorithm (lines are numbered to match equivalent instructions in the Routing Algorithm):

3. Start new cycle by $\alpha_{\text{SINK}} + \alpha_{\text{SINK}} + 1$, send $\text{MSG}(\alpha_{\text{SINK}})$ to all $\ell \in G_{\text{SINK}}$.
(Note: $\langle 3 \rangle^*$ can be performed only after $\langle 4 \rangle$ of the previous cycle has been performed).
1. For $\text{MSG}(\alpha)$
2. $N_{\text{SINK}}(\ell) + 1$,
4. if $\forall \ell' \in G_{\text{SINK}}$, holds $N_{\text{SINK}}(\ell')=1$, then cycle α completed,
 $\forall \ell' \in G_{\text{SINK}}$, set $N_{\text{SINK}}(\ell')=0$.

In principle, the routing tree can be used for broadcast purposes as follows: a node i accepts only broadcast messages received from its father p_i and forwards them to all nodes k whose father is i . Observe that we distinguish between receiving a broadcast message and accepting it. In general, a broadcast message received at a node may be either accepted or rejected, depending on the specific algorithm.

The first problem that one encounters with the above procedure is that in the routing algorithm a node i knows only its father p_i , but does not know the nodes k for which $p_k=i$. Consequently, we need an addition to the routing algorithm, so that whenever a node i changes its father p_i , (line $\langle 4 \rangle$ in the Routing Algorithm) it sends two special messages: DCL (declare) to the new father and CNCL (cancel) to the old father.

* A specific line in an algorithm will be indicated in angular brackets $\langle \rangle$. The algorithm we refer to will either be clear from the context or indicated explicitly.

Each node i will have a binary variable $z_i(k)$ for each neighbor k that will take on the value 1 if i thinks that $p_k=i$ and 0 otherwise. Receipt of DCL at node k from i shows that at the time DCL was sent, node i selected k as p_i , so that $z_k(i)$ is set to 1. The nodes i for which $z_k(i)=1$ are called sons of k . Observe that because of link delays, if i is a son of k it does not mean that at the same time k is the father of i . We can now write in our notation the combination of the above routing algorithm and the Extended Reverse Path Forwarding (ERPF) Broadcast Algorithm of [1], where B denotes a broadcast message:

ERPF Broadcast

1. For MSG (ℓ, α)
2. $N_i(\ell) \leftarrow 1$
3. if $\ell = p_i$, then: $\alpha_i \leftarrow \alpha$, send MSG to all $\ell \in G_i$, except p_i
4. if $\forall \ell' \in G_i$, holds $N_i(\ell')=1$, then:
 - 4.a. select new p_i ;
 - 4.b. if new $p_i \neq$ old p_i , then send DCL(α) to new p_i and CNCL to old p_i ;
 - 4.c. send MSG(α) to old p_i ; $\forall \ell' \in G_i$, set $N_i(\ell') \leftarrow 0$.
5. For CNCL(ℓ) set $z_i(\ell) \leftarrow 0$.
6. For DCL(ℓ, α), set $z_i(\ell) \leftarrow 1$
7. For B(ℓ)
8. if $\ell = p_i$, then accept B, send copy of B to all ℓ' 's.t. $z_i(\ell')=1$

Note: It is worth pointing out that line $\langle 8 \rangle$ means that if $\ell = p_i$, then B is accepted, while if $\ell \neq p_i$, then it is rejected.

2.2 Completeness

As mentioned in the introduction, the above broadcast protocol is noncomplete and nonfinite. The purpose of this section is to show that completeness can be achieved by using memory and counter numbers at the nodes. Observe that we achieve our goal without requiring that the counter numbers will be carried by the broadcast messages, so that the algorithm has no overhead cost. For purposes of illustration, it is best to impose for the time being no bounds on the memory or on the counters and also to describe the protocol as if completeness was already proved. After indicating the formal algorithm we shall show that it is indeed complete and in the following sections we shall introduce features that will make the memory and the counters finite.

Suppose we require each node i to have a $LIST_i$ where every accepted broadcast message is stored in the received order and also to keep a counter IC_i , counting the accepted messages. Completeness of the broadcast protocol means that for any value of IC_i , the list $LIST_i$ contains all messages sent by the source SINK up to counter number IC_i , with no duplicates and in the correct order. In other words if IC_i^B denotes the value of IC_i after broadcast message B was accepted at node i , we have $IC_i^B = IC_{SINK}^B$ for all B and all i . In the algorithm we also require that every DCL message sent by node k will have the format $DCL(\alpha, IC)$ where $IC = IC_k$ at the time DCL is sent. In this way when a node i receives DCL from k , it will have updated information about the "state of knowledge", denoted by $IC_i(k)$, of its new son k . Only broadcast messages B with $IC_i^B > IC_i(k)$ need to be sent by i to k .

The formal algorithm is now

The Complete Routing - Broadcast Algorithm (CRB) for node i

1. For $MSG(\ell, \alpha)$
2. $N_i(\ell) \leftarrow 1$
3. if $\ell = p_i$, then : $\alpha_i \leftarrow \alpha$; send $MSG(\alpha_i)$ to all $\ell' \in G_i$ except p_i
4. if $\forall \ell' \in G_i$, holds $N_i(\ell') = 1$, then:
 - 4a. select new p_i
 - 4b. if new $p_i \neq$ old p_i , then send $DCL(\alpha_i, IC_i)$ to new p_i and
CNCL to old p_i
 - 4c. send $MSG(\alpha_i)$ to old p_i ; $\forall \ell' \in G_i$, set $N_i(\ell') \leftarrow 0$
5. For $CNCL(\ell)$, set $z_i(\ell) \leftarrow 0$
6. For $DCL(\ell, \alpha, IC)$ set $z_i(\ell) \leftarrow 1$,
 - 6a. if $IC < IC_i$, then send to ℓ contents of $LIST_i$ from $IC+1$ to IC_i
while incrementing $IC_i(\ell)$ up to IC_i
 - 6b. else $IC_i(\ell) \leftarrow IC$
7. For $B(\ell)$
 - 7a. if $\ell = p_i$, then $IC_i \leftarrow IC_i + 1$, include B in $LIST_i$,
 - 7b. $\forall j \in G_i$ s.t. $z_i(j) = 1$, $IC_i(j) < IC_i$, then
 - 7c. send B to j , $IC_i(j) \leftarrow IC_i(j) + 1$

Our basic assumption is that all messages sent on a link arrive in arbitrary but finite time after their transmission, with no errors and in the correct order (FIFO). Observe that this does not preclude channel errors provided there is an acknowledgement and retransmission protocol on the link. The proof that under this condition the CRB protocol is indeed complete appears in Appendix A. Here we only mention that the important property leading to completeness is the statement of Lemma A1, that will be called the session property. Broadcast protocols associated with other routing algorithms can be made to have this property, but several additions to the algorithm are necessary. It is a special feature of the routing protocol of [2] that the

session condition holds with no extrainstructions. Also observe that as will be seen in Lemma A2 and Theorem A1, completeness is achieved without requiring messages to carry their counter number.

2.3 Finiteness

Completeness means that broadcast messages are accepted at nodes in the correct order and with no duplicates or messages missing. However, it does not ensure that all messages are indeed accepted at all nodes. The following scenario shows that, since we allow arbitrary propagation time for messages on each link, there may be a situation in the CRB algorithm where a node i accepts no messages from a certain time on. Considering Fig. 1, recall that t_D is the time when $\langle 4 \rangle$ of cycle α is executed at node i , but observe that $\text{MSG}(\alpha)$ was sent from i to j when i executed $\langle 3 \rangle$. Therefore $\text{MSG}(\alpha)$ may have arrived to j well before t_D and if the propagation time of $\text{DCL}(\alpha, \text{IC})$ is long enough, j may have completed cycle α and entered cycle $(\alpha+1)$ before time t_D . When it entered cycle $(\alpha+1)$, node j has sent $\text{MSG}(\alpha+1)$ to i , so that node i could have performed $\langle 3 \rangle$ and $\langle 4 \rangle$ of cycle $(\alpha+1)$ before time t . Now if this happens and at the time of performing $\langle 4 \rangle$ of cycle $(\alpha+1)$ node i changes its p_i , then $p_i[t] \neq j$ and B is not accepted. In principle, this scenario can be repeated indefinitely, so that B and the broadcast messages following it keep arriving at node i , but will never be accepted.

In order to correct the situation and achieve finiteness, we introduce an "Impeding Mechanism" in the CRB algorithm. Control messages $\text{MSG}(\alpha)$ sent from j to i will carry in addition a variable $z = z_j(i)$, and $\text{MSG}(\alpha, z)$ such that $\alpha = \alpha_i + 1$, $z = 0$ received from $j = p_i$ will be ignored. Moreover, if node j receives $\text{DCL}(\alpha, \text{IC})$ with $\alpha < \alpha_j$ (and then by Lemma A3 we have $\alpha = \alpha_j - 1$) node j transmits again $\text{MSG}(\alpha_j, z)$ with $z = 1$. In this way, node i postpones execution of $\langle 3 \rangle$ until it receives acknowledgement from $j = p_i$ (in the form of $\text{MSG}(\alpha_j, z = 1)$), that the last DCL message has been received at j .

For each broadcast message accepted at a node i , it is convenient at this point to indicate explicitly the cycle during which it was accepted. To do so we replace $LIST_i$ by a set of buffers $LIST_i(\alpha)$, $\alpha=1,2,\dots$ (for the meantime an infinite number of unbounded buffers) and all broadcast messages accepted while i was in cycle α are stored in $LIST_i(\alpha)$. Also, counters $C_i(\alpha)$ are used, counting messages accepted during cycle α . Out of the messages corresponding to cycle α , those that have been accepted at neighbor z as far as i knows are counted in $C_i(z)(\alpha)$. Consequently, the counter IC is redefined as the pair $IC = (\alpha, C(\alpha))$, where $IC' < IC''$ means that either $\alpha' < \alpha''$ or $\alpha' = \alpha''$ and $C'(\alpha') < C''(\alpha')$.

The resulting algorithm is given below and the proof that it is complete and finite appears in the Appendix.

The Reliable Routing-Broadcast Algorithm (RRB) for node i

1. For $MSG(\ell, \alpha, z)$
2. if $\ell \neq p_i$, then $N_i(\ell) + 1$
3. if $\ell = p_i$ and $z=1$, then : $N_i(\ell)+1, \alpha_i+\alpha_i+1$, send $MSG(\alpha_i, z_i(\ell))$ to all $z' \in G_i$ except p_i
4. if $\forall z' \in G_i$, holds $N_i(z') = 1$, then:
 - 4a. select new p_i
 - 4b. if new $p_i \neq$ old p_i , then send $DCL(\alpha_i, C_i(\alpha_i))$ to new p_i and $CNCL$ to old p_i
 - 4c. send $MSG(\alpha_i)$ to old p_i ; $\forall z' \in G_i$, set $N_i(z') + 0$
5. For $CNCL(\ell)$, set $z_i(\ell) + 0$
6. For $DCL(\ell, \alpha, C)$, set $z_i(\ell) + 1$
 - 6a. if $C < C_i(\alpha)$, then send to z contents of $LIST_i(\alpha)$ from C to $C_i(\alpha)$ while incrementing $C_i(z)(\alpha)$ to $C_i(\alpha)$
 - 6b. if $\alpha = \alpha_i - 1$, then send $MSG(\alpha_i, z_i(\ell))$ to z
send to z contents of $LIST_i(\alpha_i)$ from 1 to $C_i(\alpha_i)$ while incrementing $C_i(z)(\alpha_i)$ to $C_i(\alpha_i)$
 - 6c. else, if $C \geq C_i(\alpha)$, then $C_i(z)(\alpha) + C$

7. For $B(z)$

7a. if $z=p_i$, then $C_i(\alpha_i) + C_i(\alpha_i) + 1$, include B in $LIST_i(\alpha_i)$,

7b. $\forall j \in G_i$, s.t. $z_i(j)=1$, $C_i(\alpha_i)(j) < C_i(\alpha_i)$, then

7c. send B to j, $C_i(\alpha_i)(j) + C_i(\alpha_i)(j) + 1$

Before proceeding, we note here that the Impeding Mechanism slows down the routing algorithm, but only in extreme situations. This is because the Impeding Mechanism is in fact activated only in the case when $DCL(\alpha, C)$ sent by a node i to j arrives there after node j has performed $\langle 3 \rangle$ of cycle $(\alpha+1)$. Since such a DCL message is sent by i when it performs $\langle 4 \rangle$ of cycle α , this means that propagation of DCL on link (i, j) takes more time than propagation of the routing cycle α from i all the way to SINK plus propagation of cycle $(\alpha+1)$ all the way from SINK to node j . This may indeed happen if we allow arbitrary delays on links, but the chances are small.

3. THE RELIABLE BROADCAST PROTOCOL

The final form of the broadcast protocol will be obtained from the RRB algorithm after making several observations.

- a) The broadcast messages accepted by node i while it is in cycle α are exactly those broadcast messages released by SINK while it is in cycle α (follows from Corollary A1).
- b) If node i is in cycle α , it will never be required to send to neighbors messages accepted prior to cycle $(\alpha-1)$ and therefore it needs to store only messages accepted during the present and the previous cycles.

From a) and b) follows that we can make significant simplifications in RRB. The variables α, α_i can be binary; only two lists $LIST_i(0)$ and $LIST_i(1)$ need to be stored; if SINK is allowed to send no more than M broadcast messages per cycle, those LIST's can have finite size M ; only counters $C_i(0)$, $C_i(\ell)(0)$, $C_i(1)$, $C_i(\ell)(1)$ are needed and all those are bounded by M ; control messages MSG need not carry the variable α . The resulting broadcast algorithm has the following properties:

Properties of RRB (network has N nodes and E link)

- 1) Reliability
- 2) Finite memory and counters
- 3) No overhead cost
- 4) Control communication cost: the routing protocol requires $2E$ messages MSG per cycle whether broadcast is operating in the network or not. Broadcast requires no new MSG messages, except in the peculiar situation described at the end of Section 2.3. In addition we need at most N DCL messages and N CNCL messages per cycle.
- 5) Broadcast communication cost: most of the time broadcast messages propagate on spanning trees. The only situation when two copies of the same message arrive at a node (and one is ignored) is when a broadcast message "crosses paths" with a CNCL message. This means that CNCL is

sent by i to j and the broadcast message is sent by j before CNCL has arrived and is received by i after CNCL was sent. The worst case gives $2(N-1)$ messages in the net per broadcast message, but in most cases this situation will not occur, especially if the propagation time of CNCL is small, so that the average is very close to $(N-1)$ copies per message, which is the minimal broadcast communication cost.

6) Delay: the routing algorithm tends to find paths with small total weight (sum of link weights from nodes to SINK). The delay of broadcast messages will be small if the weights are link delays and the traffic is symmetric on links or if the weights of link (i,j) contain a measure of the delay on link (j,i) .

Appendix A

Here we prove that the CRB Protocol of Section 2.2 is indeed complete and that the RRB Protocol of Section 2.3 is indeed complete and finite. First we recall several properties of the routing protocol of [2] indicated in Section 2.1 and introduce several definitions:

- a) in each cycle α , the routing protocol requires each node i to send exactly one $MSG(\alpha)$ to each neighbor
- b) cycle α starts when SINK sends $MSG(\alpha)$ to all neighbors (<3> in the algorithm for SINK) and ends when SINK receives $MSG(\alpha)$ from all neighbors (line <4>).
- c) a node i is said to be in cycle α while $\alpha_i = \alpha$, i.e. from the time it performs <3> with $\alpha_i = \alpha$ and until it performs <3> with $\alpha_i = \alpha + 1$.
- d) just before node i performs <3>, holds $\alpha = \alpha_i + 1$, so that α_i always increases by 1.
- e) whenever we need to indicate the value of a variable, say p_i , at a certain time t we shall write $p_i[t]$.

Lemma A1 (Session Property)

Consider the CRB Protocol of Section 2.2. If a broadcast message B is received at time t at node i from j and it is accepted, then B was sent by j after receiving the last DCL message sent by i until time t .

Proof

Let $\tau < t$ be the time B was sent by j . Since broadcast messages are accepted only from fathers (see <9> of CRB) and sent only to sons (see <7> and <10>), we have $p_i[t] = j$ and $z_j(i)[\tau] = 1$. Thus the last DCL message sent by i before time t (at time t_0 say) was indeed sent to j and we want to show that it was received by j (at time τ_0 say) before time τ , or in other words i is the son of j at time τ as a result of this last DCL and not of some previous DCL's. This is exactly the session property. The timing diagram is given in Fig. 1. Consider also the last CNCL sent by i before t

to j and let t_C, τ_C, α be respectively the time it was sent, the time it was received and the cycle number of i at time t_C . Clearly $t_C < t_D$ and by FIFO we also have $\tau_C < \tau_D$. In order to prove the lemma we need to show that $\tau_D < \tau$. Observe now that $z_j(i) = 0$ between τ_C and τ_D and since $z_j(i)[\tau] = 1$, time τ cannot be between τ_C and τ_D . It is sufficient therefore to show that $\tau_C < \tau$. Observe that <4b> shows that CNCL is sent after receiving MSG(α) from all neighbors, in particular j and before sending MSG(α) to j and therefore $\alpha_j[\tau_C] = \alpha$, where α_j is the cycle number of node j . Suppose now that $\tau_C > \tau$. Then $\alpha_j[\tau] \leq \alpha$ and B was sent (and received, by FIFO) from j to i before MSG($\alpha+1$), so that i could not have performed <4> of cycle $\alpha+1$ before t . Since p_i changes only in <4>, it follows that $p_i[t] = p_i[t_C^+] \neq j$ which is a contradiction. This proves the session property of the Routing-Broadcast Protocol of Section 2.2. Observe that the proof relies heavily on the Properties of the Routing Protocol of [2].

Lemma A2

If broadcast message B is received at node i from j and is accepted, then $IC_i^B = IC_j^B$. (Recall that IC_i^B denotes the value of the counter IC_i just after node i has accepted B).

Proof

Consider the notations of Lemma A1 and of Fig. 1. From line <4b> in the CRB algorithm follows that the DCL(α, IC) message carries the counter number $IC = IC_i[t_D]$. Since $p_i = j$ on the interval $(t_D, t]$, node i accepts during this time broadcast messages only from j , and by the Session Property, those are sent only after time τ_D at which j performs <6>, <7>. Now it is easy to check (see <7>, <9>-<11> for node j) that in both cases, $IC < IC_j[\tau_D^-]$ and $IC \geq IC_j[\tau_D^-]$, node j will consecutively send to i after τ_D the broadcast messages corresponding to counter numbers $IC+1, IC+2$, etc. When they will be received and accepted at i , the counter IC_i will be increased respectively to $IC+1, IC+2$, etc.

Theorem A1

The CRB algorithm of Section 2.2 is indeed complete, namely $IC_i^B = IC_{SINK}^B$ holds for every node i and every broadcast message B .

Proof

If the above relation does not hold, let i and B be the node and broadcast message for which it is violated for the first time throughout the network, and let t be the time B was accepted at i . If B was received from j , then lemma A2 implies $IC_j^B = IC_i^B$ so that $IC_j^B \neq IC_{SINK}^B$. But B was accepted at j before being accepted at i , violating the fact that the statement of the Theorem held throughout the network until time t .

For future reference we need

Lemma A3

If $DCL(\alpha, IC)$ arrives at node j , then $\alpha = \alpha_j$ or $\alpha_j - 1$.

Proof

Consider the notations of Lemma A1 and of Fig. 1. Then $\alpha_i[t_D] = \alpha$ and therefore $MSG(\alpha+1)$ will be sent from i to j after the DCL message. Consequently $\langle 4 \rangle$ of cycle $(\alpha+1)$ can be performed at j only after τ_D , hence $\alpha_j[\tau_D] \leq \alpha+1$. On the other hand t_D is the time i performs $\langle 4 \rangle$ of cycle α and hence $MSG(\alpha)$ has been received at i from j before or at t_D , so that $\alpha_j[\tau_D] \geq \alpha$.

We next proceed to the proof that the RRB Protocol of Section 2.3 is complete and finite.

Lemma A4

In the RRB Protocol, if a $MSG(\alpha', z=0)$ arrives at i from $j=p_i$, (and by $\langle 2 \rangle$, $\langle 3 \rangle$ is ignored), then $MSG(\alpha', z=1)$ will arrive at i in finite time from j and then j will still be the father p_i of i .

Proof

With the notations of Fig. 1, where B is replaced by $MSG(\alpha', z=0)$, holds $\tau < \tau_D$ (since $z=0$) and $t > t_D$ (since $p_i=j$). Now

$\alpha_j[\tau_D] \geq \alpha_j[\tau] = \alpha' = \alpha_i[t]+1 \geq \alpha_i[\tau_D]+1 = \alpha+1$, where the second equality follows from property d) at the beginning of the Appendix. From Lemma A3 follows that $\alpha_j[\tau_D] = \alpha+1$ and hence j will send to i at time τ_D control message $MSG(\alpha', z=1)$ according to line <6b> in RRB.

Definition

A control message $MSG(\alpha, z=1)$ is said to be "accepted" at node i if it triggers performance of <3> in RRB at node i. Also, define the counter number associated with an accepted message $MSG(\alpha, z=1)$ as $IC_i(MSG(\alpha, z=1)) = (\alpha, C_i(\alpha)=0)$.

Lemma A5

With the above definitions, control messages with $z=1$ propagate in RRB as if they were regular broadcast messages.

Proof

Broadcast messages are accepted at i only if they arrive from p_i and are sent to sons, either when they are accepted or in response to DCL with $IC < IC_i$. Control messages $MSG(\alpha, z=1)$ are accepted only if they arrive from p_i and are sent to sons, either when they are accepted (<3> in RRB) or in response to DCL with $IC < IC_i$ (<6b> in RRB). Moreover, $MSG(\alpha, z=1)$ is accepted at i before all broadcast messages B with $IC_i^B = (\alpha, C_i(\alpha))$, since node i enters cycle α as a result of accepting $MSG(\alpha, z=1)$ from p_i and broadcast messages with IC_i^B as above are all accepted while i is in cycle α . Now, $MSG(\alpha, z=1)$ is sent to any node before all such broadcast messages (see <3> and <6b>), so that the order is preserved as well. Hence the statement of the Lemma.

Corollary A1

The combination of broadcast messages and control messages with $z=1$ performs a jointly complete algorithm, i.e. all such messages are accepted in the order released by the source node SINK, with no duplicates and no messages missing.

Theorem A2

The RRB protocol is complete and finite.

Proof

From Lemma A4 and the fact that every routing cycle of the algorithm of [2] propagates in finite time, follows that the propagation of control messages with $z=1$ is finite, namely every node enters every cycle in finite time. By Corollary A1, all broadcast messages released by SINK while SINK is in cycle α are accepted at each node while the node is in cycle α , and since each node enters cycle $(\alpha+1)$ in finite time, all such broadcast messages are accepted at each node in finite time.

References

1. Y.K. Dalal and R.M. Metcalfe, Reverse Path Forwarding of Broadcast Packets, Communications ACM, Vol. 21, No. 12, pp. 1040-1048, Dec. 1978.
2. P.M. Merlin and A. Segall, A Failsafe Distributed Routing Protocol, IEEE Trans. on Comm., Vol. COM-27, No. 9, pp. 1280-1287, Sept. 1979.
3. Y.K. Dalal, Broadcast Protocols in Packet Switched Computer Networks, Ph.D. Thesis, Stanford Univ., April 1977, DSL Technical Report 128.




Fig. 1. Timing Diagram for Lemma A1

DATE
FILMED
-8