

Application Layer Multicast Algorithm

Sergio Machado

*Universitat Politècnica de Catalunya
Castelldefels*

Javier Ozón

*Universitat Politècnica de Catalunya
Castelldefels*

Abstract

This paper presents a multicast algorithm, called MSM-s, for point-to-multipoint transmissions. The algorithm, which has complexity $O(n^2)$ in respect of the number n of nodes, is easy to implement and can actually be applied in other point-to-multipoint systems such as distributed computing. We analyze the algorithm and we provide some upper and lower bounds for the multicast time delay.

1 Introduction

Multicast is a point-to-multipoint means of transmitting data in which multiple nodes can receive the same information from one single source. The applications of multicast include video conferencing, multiplayer networking games and corporate communications. The lack of deployment of IP Multicast has led to considerable interest in alternative approaches at the application layer, using peer-to-peer architectures[1]. In an application layer multicast approach, also called overlay multicast, participating peers organize themselves into an overlay topology for data delivery. In this topology each edge corresponds to an unicast path between two end-systems or peers (also called nodes) in the underlying IP network. All multicast-related functionality is implemented by peers instead of routers, with the goal of depicting an efficient overlay network for multicast data transmission.

In this work we present an algorithm suitable for peer-to-peer multicast transmissions, although the high degree of abstraction of its definition makes it also suitable for its implementation in other layers and in general message-passing systems. The main contribution of this proposal is that the operation of our algorithm is simple, with a complexity of $O(n^2)$, where

n is the number of peers and thus, it may adapt dynamically to the characteristics of the source and the network in order to improve the multicast time delay. Algorithm execution may be computed by a single multicast group member, usually the one which multicasts the message, or by all the members after the complete network status has been broadcasted.

2 Single Message Multicast Algorithm

Bar Noi *et al.* introduced in [2] the *Message Passing System Model MPS* which characterizes systems that use packet switching techniques at the communication network. In this work, we extend the Bar Noi model to $EMPS(n, \lambda, \mu)$, which consists of a set of n full-duplex nodes $\{p_0, \dots, p_{n-1}\}$ such that each node can simultaneously send and receive a message. The term *message* refers to any atomic piece of data sent by one node to another using the protocols of the underlying layers. For each node p we define the *transmission time* μ_p as the time that requires p to transmit a message M . Moreover, for each pair of nodes p and q in a message-passing system we define *communication latency* λ_{pq} between p and q . If at time t node p starts to send a message M to node q , then node p sends message M during the time interval $[t, t + \mu_p]$, and node q receives message M during the time interval $[t + \lambda_{pq} - \mu_p, t + \lambda_{pq}]$. Thus λ_{pq} is the sum of transmission time μ_p of node p , and propagation delay between p and q . For simplicity sake, we assume that the communication latency is constant, and we consider multicast as a broadcast problem, since we can isolate the receiving nodes of a multicast communication, form with them a complete overlay graph, and then depict a routing table through a broadcast algorithm.

Let p_0 be a node in $EMPS(n, \lambda, \mu)$ model which has a message M to multicast to the set of receiving nodes $R = \{p_1, p_2, \dots, p_{n-1}\}$, we find an algorithm that minimizes the *multicast time*, that is, the time at which all nodes in R have received the message M . Though the result of $EMPS$ is a multicast spanning tree, in Fig. 1 we show that this problem is different from the well known minimum spanning tree problem.

The algorithm that we propose, called SMM *Single Message Multicast*, operates as follows: at each step SMM algorithm chooses the node which has not yet received the message and has the lowest cost, that is, the unvisited node that can be reached at minimum time from the nodes which have already received the message. Once the message has been received by this

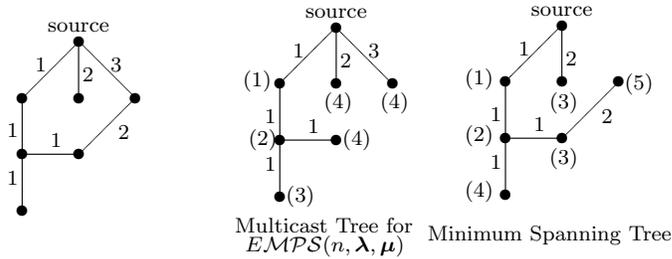


Figure 1: Example of a network where $EMPS(n, \lambda, \mu)$ model does not correspond to the MST problem. We show in parenthesis the multicast delay for each node in the case that $\mu = 1$ for any node.

node, the algorithm recalculates the arrival times of the remaining nodes, searches the next node at which the message must be forwarded, and so forth. We assume that when a sending node finishes the retransmission of the message to another node, it begins immediately with another destination node. SMM algorithm is very similar to Dijkstra's shortest path algorithm with the difference that in this case the time delay between two nodes p and q is not constant. Actually, in $EMPS(n, \lambda, \mu)$ this delay is equal to λ_{pq} plus μ_p multiplied by the number of previous retransmissions of node p .

The multicast time achieved by the algorithm SMM is minimum when $\mu_p = 0$ for all the nodes. In this case, the time delay between two nodes p and q is always the weight λ_{pq} of the edge which joins them, and thus the SMM algorithm corresponds to the optimal algorithm Dijkstra of complexity $O(n^2)$. In a general case, however, the SMM algorithm is not always optimal. Fig. 2 shows a network where SMM is not optimal.

Proposition 1 *Algorithm SMM for $EMPS(n, \lambda, \mu)$ has complexity $O(n^2)$.*

Proof At each step SMM searches the node which has not yet received the message with lowest cost. As the maximum number of unvisited nodes is $n - 1$ this operation requires at most $n - 2$ comparisons. Moreover, the algorithm executes one step for every node which receives the message. Thus, we have $n - 1$ steps and at each step we perform at maximum $n - 2$ comparisons resulting in a complexity of $O(n^2)$.

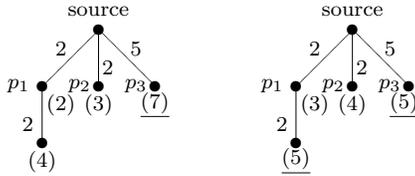


Figure 2: Example of a network where SMM is not optimal. Transmission time is 1 for all the nodes. On the left we apply SMM and on the right another multicast transmission order with a better result.

2.1 Message Stream Multicast Algorithm

The SMM algorithm has been defined for the multicast of a single message. For a set of messages we can repeat indefinitely the routing table obtained with the SMM algorithm, multicasting each message independently of the others. That means that when one message finally arrives at all the nodes, the message source would proceed to multicast the next message, and so forth. The total delay multicast time of the stream would be in this case the total number of messages M multiplied by the multicast SMM delay for one single message. The main inconvenience of this solution is that the source can not send the next message until the previous one has been received by all the group members and this could slow down the rate of communication.

Next, we consider a new possibility. Before the first message has arrived to all nodes, the source could stop sending it and begin with the second message. With this restriction, the multicast time of the first message will increase, but we will begin to send before the second message. This saving of time between the sending of two consecutive messages will be progressively accumulated, and if the number of messages is large enough it will compensate the increase of the multicast time for one single message. The modified algorithm, that we call *MSM-s Message Stream Multicast*, works as SMM and applies the same multicast scheme for every message with the particularity that it stops the transmission of any message once a node has already sent it s times. Then it will begin to send the next message and so forth. Since the restriction on the number of retransmissions could isolate some nodes of the network, MSM-s should choose a minimum restriction number s to guarantee full-connectivity. As SMM, MSM-s algorithm has complexity $O(n^2)$.

In next sections we prove that under certain conditions it is possible to calculate a minimum number M_σ in such a way that if the number of messages is equal or larger than M_σ then MSM- σ is better than MSM- $(\sigma + 1)$. Moreover, when restricting the number of transmissions for each node, MSM-s has to take into account the transmission rate of the source. That is, if the source sends at most s times the first message and then, after $s \cdot \mu_r$ time units, stops the transmission of the first message to begin with the second one, we must be able to assume that the source has the second message ready to forward. Otherwise, the source would stop sending the first message before having the second one and would remain unnecessarily idle, with the consequent loss of efficiency.

2.2 Message Stream Multicast Algorithm with Time Restriction

Let p be a node which forwards the message to node q , and let $s_p(s), s_q(s) \leq s$ be the times that p and q forward the message for MSM-s, respectively. In this case the second message will be received at q with a delay of $s_p(s) \cdot \mu_p$ in respect to the first message, since the second message follows the same path but with a source delay of $s_p(s) \cdot \mu_p$. When the forwarding period $s_q(s) \cdot \mu_q$ of node q is higher than the forwarding period $s_p(s) \cdot \mu_p$ of node p , then successive messages may have higher delays than former messages. In this context, the second message could arrive at node q before it has finished forwarding the first message and then the second message would have to be buffered, with the consequent time delay. This buffering delay would be also accumulated by the third message, and so forth. Nevertheless, this situation may be avoided by limiting the time period $s_q(s) \cdot \mu_q$ at which each node forwards a message, that is, by assuring that the forwarding rate $1/(s_q(s) \cdot \mu_q)$ of any node q is higher than the rate $1/(s_p(s) \cdot \mu_p)$ of any node p which is in the path from the source to node q , including the source. Therefore, the delay of the first message will be always the same as the time delay of any other message, an issue which has great importance in Section 3. Note also that we do not want $1/(s_q(s) \cdot \mu_q) \gg 1/(s_p(s) \cdot \mu_p)$ since in this case q would stop forwarding the first message long before receiving the second one, and then the communication would lose efficiency.

3 Analysis of MSM-s

3.1 Stream Multicast Delay

Let τ_s be the multicast time delay for a single message when the number of transmissions of each node is established up to s , M the number of messages of the stream and μ_r the transmission time of the source, also called root. We assume that s is large enough to arrive at all the nodes of the network. In this case, the total stream multicast delay τ_{Ms} for MSM-s is $\tau_{Ms} = (M - 1) \cdot s \cdot \mu_r + \tau_s$. That is, the root sends the first message s times and then, $s \cdot \mu_r$ time units later, it begins with the second and so forth. At moment $(M - 1) \cdot s \cdot \mu_r$ the root finishes to send the $(M - 1)$ th message and it begins with the last message, that will arrive at the last node τ_s time units later. Remember that, as shown in Section 2.2, under certain restrictions, the delay τ_s for the last message is the same as the delay for any other message.

Equation for τ_{Ms} is only valid when the root sends each message s times. When s is large the message may be received by all nodes before the root has sent it s times. Though in this case the node could remain idle and wait until $s \cdot \mu_r$ and then begin to send the second message, this would mean a loss of efficiency. So, for MSM-s, when the message is received for all nodes before the root has sent it s times, we will allow the root to send the second message immediately, without an interval of silence. In this particular case the parameter s should be replaced by the actual number of times $s_r(s) \leq s$ that the root sends each message for MSM-s, and then $\tau_{Ms} = (M - 1) \cdot s_r(s) \cdot \mu_r + \tau_s$.

Proposition 2 *Given the algorithm MSM-s for EMPS, the delay of a single message is such that $\tau_{\sigma+\Delta} \leq \tau_\sigma \forall \sigma, \Delta > 0$.*

Proof By construction of the algorithm. When bounding up to $\sigma + \Delta$ the transmissions of each node, MSM- $(\sigma + \Delta)$ will depict a better multicast tree than MSM- $(\sigma + \Delta - 1)$ for any message only if there exists a better solution. Otherwise MSM- $(\sigma + \Delta)$ will depict the same multicast tree depicted by MSM- $(\sigma + \Delta - 1)$. Thus $\tau_{\sigma+\Delta} \leq \tau_{\sigma+\Delta-1}$. Repeating the argument for $\tau_{\sigma+\Delta-1}$ and $\tau_{\sigma+\Delta-2}$ and so forth, we obtain $\tau_{\sigma+\Delta} \leq \tau_\sigma \forall \sigma, \Delta > 0$.

Theorem 3 *Given the algorithm MSM-s for EMPS(n, λ, μ), we may obtain the conditions such that MSM- σ is faster than MSM- $(\sigma + 1)$.*

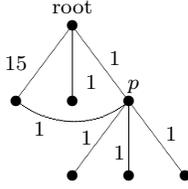


Figure 3: Example of network where $s_r(\sigma) \geq s_r(\sigma + 1)$. In particular we have $s_r(3) = 3$ and $s_r(4) = 2$.

Proof First we define, in the case that MSM- σ could be better than MSM- $(\sigma + 1)$, the minimum number M_σ of messages from which MSM- σ is better than MSM- $(\sigma + 1)$. We begin with $\sigma = 1$. The value of M_1 can be easily obtained once we have computed τ_1 , τ_2 and $s_r(2)$ by executing MSM-1 and MSM-2. Remember that $s_r(2)$ is the number of times that the root sends each message in MSM-2 and that, by Theorem 2, $\tau_1 \geq \tau_2$. Then $\tau_{M_1} = (M - 1) \cdot \mu_r + \tau_1$ and $\tau_{M_2} = (M - 1) \cdot s_r(2) \cdot \mu_r + \tau_2$. In this case $s_r(2)$ may be equal to either 1 or 2. In the unusual first case where $s_r(2) = 1$, since $\tau_1 \geq \tau_2$, MSM-2 will be equal or better than MSM-1 for any number of messages. In the more usual case where $s_r(2) = 2$ we establish the restriction $\tau_{M_1} \leq \tau_{M_2}$ and then $M \geq (\tau_1 - \tau_2) / \mu_r + 1 = M_1$.

For a general case, the number M_σ of messages from which the total stream multicast delay is better for $s = \sigma$ than for $s = \sigma + 1$ may be obtained repeating the arguments for M_1 . First we obtain by implementing MSM- σ and MSM- $(\sigma + 1)$ the following expressions $\tau_{M_\sigma} = (M - 1) \cdot s_r(\sigma) \cdot \mu_r + \tau_\sigma$ and $\tau_{M_{\sigma+1}} = (M - 1) \cdot s_r(\sigma + 1) \cdot \mu_r + \tau_{\sigma+1}$. By Theorem 2 we have $\tau_\sigma \geq \tau_{\sigma+1}$. Thus, in the unusual case that $s_r(\sigma) \geq s_r(\sigma + 1)$, MSM- $(\sigma + 1)$ will be equal or better than MSM- σ for any number of messages. In other case when, $s_r(\sigma) < s_r(\sigma + 1)$ we establish $\tau_{M_\sigma} \leq \tau_{M_{\sigma+1}}$ and then we obtain $M \geq (\tau_\sigma - \tau_{\sigma+1}) / ((s_r(\sigma + 1) - s_r(\sigma)) \cdot \mu_r) + 1 = M_\sigma$.

Though it is not a usual case, in Fig. 3 we depict a network where $s_r(\sigma)$ is greater than $s_r(\sigma + 1)$. In particular we have $s_r(3) > s_r(4)$, and thus MSM-4 will be faster than MSM-3 for any number of messages. In order to accomplish the restrictions discussed in section 2.2, we suppose $\mu_r = 2$ and $\mu_p = 1$.

3.2 Analytical Bounds for M_1

In this section, we obtain an analytical bound for M_1 , that is, for the number of messages from which the total stream multicast delay is better for $s = 1$ than for $s = 2$. As explained in former section we assume $s_r(2) = 2$. In other case, when $s_r(2) = 1$, MSM-2 will be equal or faster than MSM-1. Let M be the number of messages; τ_{M1} and τ_{M2} the multicast delay for MSM-1 and MSM-2 respectively; μ_r the transmission time of the root; and λ_{min} and λ_{max} the minimum and maximum latency between any pair of nodes, respectively. First, we find an upper bound for τ_{M1} and a lower bound for τ_{M2} which we denote respectively by T_1 and t_2 . If we force T_1 to be lower or equal than t_2 , then MSM-1 will be better than MSM-2:

$$\tau_{M1} \leq T_1 \leq t_2 \leq \tau_{M2} \quad (1)$$

To find T_1 and t_2 , we modify slightly the MSM-s performance. First we have:

$$\tau_{M1} \leq (M - 1) \cdot \mu_r + (n - 1)\lambda_{max} = T_1 \quad (2)$$

Remember that for MSM-1 each node sends each message only once, so MSM-1 depicts a linear tree with $n - 1$ links. In this case it is clear that $\tau_{M1} \leq T_1$ since Eq. 2 corresponds to the worst case where a message has to cross the $n - 1$ links with the maximum latency λ_{max} . To find a lower bound for τ_{M2} we consider an algorithm with a lower delay than MSM-2. First we assume that the latency for any pair of nodes is the minimum latency λ_{min} . Moreover, in the new algorithm we consider that a node can send the same message simultaneously to two different nodes, that is, that μ_p is equal to 0 for all the nodes. Note that, though this is physically impossible, the new algorithm would be better than MSM-2. Let $N(t), t \in \mathbb{Z}^+$, be the number of nodes that have received the message at step t according to the new algorithm, then $N(t) = 1 + 2 + 4 + \dots + 2^t = 2^{t+1} - 1$. If we equal $N(t)$ to the number n of nodes we will obtain the number of steps that we need to arrive at all the network $t = \lceil \log_2(n + 1) - 1 \rceil$. In this case the new algorithm could send the single message to all the other nodes in $\lceil \log_2(n + 1) - 1 \rceil \cdot \lambda_{min}$ time units and then for all the messages we have:

$$\tau_{M2} \geq (M - 1) \cdot 2 \cdot \mu_r + \lceil \log_2(n + 1) - 1 \rceil \lambda_{min} = t_2 \quad (3)$$

Finally, if according to Eq. 1 we force T_1 to be lower or equal than t_2 then τ_{M1} will be also lower or equal than τ_{M2} and MSM-1 will be better than MSM-2. From Eq. 2 and Eq. 3 we have $M \geq ((n-1)\lambda_{max} - \lceil \log_2(n+1) - 1 \rceil \lambda_{min}) / \mu_r + 1$. And since we have considered tighter cases than MSM-1 and MSM-2:

$$M_1 \leq \frac{(n-1)\lambda_{max} - \lceil \log_2(n+1) - 1 \rceil \lambda_{min}}{\mu_r} + 1 \quad (4)$$

Note that when $s_r(2) = 2$ there is always a number of messages from which MSM-1 is better than MSM-2. From Eq. 4 we see that this minimum number of messages is linear respect to the number n of nodes. So we can conclude that for the general case that $s_r(2) = 2$, MSM-1 is in general better than MSM-2, provided that the number of messages is usually larger than the number of nodes.

The bound obtained in Eq. 4 can be improved by recalculating t_2 , that is, by comparing MSM-2 to a tighter algorithm and by using the same lower bound T_1 for MSM-1. We assume that $s_r(2) = 2$. First, we define an algorithm such that, at each step, each node sends the message to one node and such that each node can send the message only twice. We do not consider by the moment time delays. We call $N(t)$ the number of nodes which have received the message at step t . Note that from step $t-1$ to next step t , only the $N(t-1) - N(t-3)$ nodes which have not yet forwarded the message twice can forward it. Thus we have, at step t , the $N(t-1)$ nodes of the last step plus the $N(t-1) - N(t-3)$ nodes that have just received the message one or two iterations before. Thus, $N(t) = N(t-1) + (N(t-1) - N(t-3)) = 2N(t-1) - N(t-3)$. In our case we have also $N(0) = 1$, $N(1) = 2$ and $N(2) = 4$. It is easy to see that $N(t) = F(t+3) - 1$ where $F(t)$ is the well known Fibonacci serie for $F(0) = 0$ and $F(1) = 1$. Hence, considering $\phi_1 = (1 + \sqrt{5})/2$ and $\phi_2 = (1 - \sqrt{5})/2$ we have:

$$N(t) = F(t+3) - 1 = (\phi_1^{t+3} - \phi_2^{t+3}) / \sqrt{5} - 1 \quad (5)$$

As we are determining a lower bound, in order to calculate the number t of steps as a function of the number n of nodes we define $N'(t)$ which is a little faster than $N(t)$ as $N'(t) = (\phi_1^{t+3} + 1) / \sqrt{5} - 1$. Observe that from Eq. 5 we have $-1 < \phi_2 < 0$ and then $N'(t) > N(t)$. Hence, if we calculate for $N'(t)$ the number of steps necessary to visit n nodes, we will

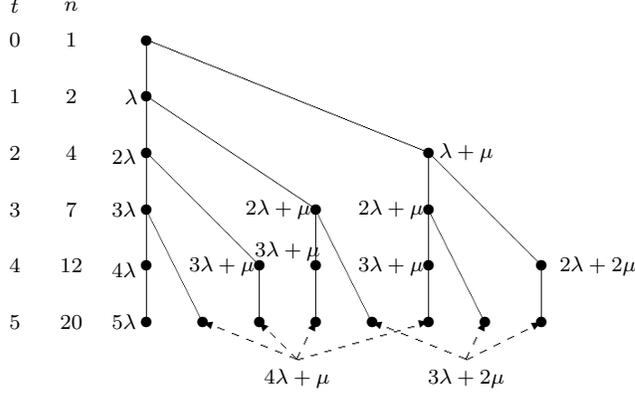


Figure 4: The Fibonacci Tree. Each node forwards the message twice.

obtain a lower bound for $N(t)$. For $t \gg 1$ the term ϕ_2^{t+3} is close to 0 and then we have a very accurate bound. If we equal $N'(t)$ to n we obtain $t = \lceil \log_{\phi_1}((n+1)\sqrt{5}-1) - 3 \rceil$. We can also prove from Fig. 4 that at each step t we have a minimum delay of $t \cdot (\lambda_{min} + \mu_{min})/2$ and then:

$$\tau_{M2} \geq (M-1) \cdot 2 \cdot \mu_r + \left\lceil \log_{\phi_1}((n+1)\sqrt{5}-1) - 3 \right\rceil \frac{\lambda_{min} + \mu_{min}}{2} = t_2$$

Hence, considering the new bound of t_2 with $\phi_1 = (1 + \sqrt{5})/2$ and repeating the arguments from the former section with the same value of T_1 , we have:

$$M_1 \leq \frac{(n-1)\lambda_{max} - \left\lceil \log_{\phi_1}((n+1)\sqrt{5}-1) - 3 \right\rceil \frac{\lambda_{min} + \mu_{min}}{2}}{\mu_r} + 1 \quad (6)$$

This bound is tighter than the bound of Eq. 4 depending on the value of μ_{min} . Actually, if μ_{min} is close to λ_{min} the new bound is better than the former, whereas if $\mu_{min} \ll \lambda_{min}$ then we must consider Eq. 4. In a practical case we can calculate both bounds and consider the best one.

3.3 An Upper and a Lower Bound for Time Delay in MSM-s

Let τ_{Ms} be the multicast delay for MSM-s and T_s an upper bound of τ_{Ms} . We obtain first a bound for $s = 2$ and then we generalize the result. With

this purpose we consider again the algorithm $N(t)$. As we determine an upper bound, in order to calculate the number t of steps as a function of the number n of nodes we define $N''(t)$, which is a little slower than $N(t)$, as $N''(t) = (\phi_1^{t+3} - 1)/\sqrt{5} - 1$. From Eq. 5 we have $-1 < \phi_2 < 0$ and then $N''(t) < N(t)$. If we calculate for $N''(t)$ the number of steps that we need to visit n nodes we will obtain an upper bound for $N(t)$. If we equal $N''(t)$ to n we obtain $t = \lceil (\log_{\phi_1}((n+1)\sqrt{5} + 1)) - 3 \rceil$. Considering from Fig. 4 that at each step t the maximum delay is $t \cdot \lambda_{max}$, we obtain $\tau_2 \leq \lceil (\log_{\phi_1}((n+1)\sqrt{5} + 1)) - 3 \rceil \lambda_{max}$. Note that, if we want to guarantee that $N''(t) < N(t)$, we have to suppose that each node can send the message twice for MSM-2. Remember also that in section 2.2 we have seen that this is not always possible (in order to avoid congestion). However, we now assume that for MSM-s every node can send the message s times (twice for $s = 2$), even if its rate is lower than the root rate, and then we force the root to send any message with a lower rate than the slowliest node in the graph. In this case, we will not have congestion problems, as referred in section 2.2, and all the messages will have the same delay. Then, since $\tau_s \leq \tau_2 \quad \forall s \geq 2$ and the root will begin to send a message at most $s \cdot \mu_s$ time units later than the previous one (being μ_s the maximum transmission time of any node in the graph), we obtain:

$$\tau_{Ms} \leq (M-1) \cdot s \cdot \mu_s + \left\lceil \log_{\phi_1}((n+1)\sqrt{5} + 1) - 3 \right\rceil \lambda_{max} = T_s \quad \forall s \geq 2 \quad (7)$$

To find a lower bound t_s for τ_{Ms} we consider, as we did for MSM-2 in section 3.2 but now in a general case, an algorithm such that a node can forward the same message simultaneously to s different nodes. Moreover, we assume that the latency for any pair of nodes is the minimum latency λ_{min} . The new algorithm is therefore better than MSM-s. Let $N(t), t \in \mathbb{Z}^+$, be the number of nodes that have received the message at step t , then $N(0) = 1$ and $N(t) = 1 + s + s^2 + \dots + s^t = (s^{t+1} - 1)/(s - 1)$ for $t > 0$. If we equal $N(t)$ to the number of nodes n we obtain that the number of steps that we need to arrive at all the network for $s \geq 2$ is $t = \lceil (\log_s(n(s-1) + 1)) - 1 \rceil$. And thus for $s \geq 2$ we obtain t_s :

$$\begin{aligned} \tau_{Ms} &\geq (M-1) \cdot s_r(s) \cdot \mu_r + \left\lceil \log_s(n(s-1) + 1) - 1 \right\rceil \lambda_{min} \\ &\geq (M-1) \cdot \mu_r + \left\lceil \log_s(n(s-1) + 1) - 1 \right\rceil \lambda_{min} \end{aligned} \quad (8)$$

For $s = 2$ we obtain the expression in Eq. 3 and for $s = 1$ we have $\tau_{M1} \geq (M - 1) \cdot \mu_r + (n - 1)\lambda_{min}$.

3.4 A General Bound for M_σ

Taking the bounds of the former section and repeating the arguments of section 3.2 for M_1 , we obtain a bound for the minimum number M_σ of messages from which MSM- σ is better than MSM- $(\sigma + 1)$. As in former sections, this bound has only sense when $s_r(\sigma) < s_r(\sigma + 1)$. In other case, MSM- $(\sigma + 1)$ is always equal or better than MSM- σ . From Eq. 7 we have an upper bound T_σ for $s = \sigma$ and from Eq. 8 we obtain a lower bound $t_{\sigma+1}$ for $s = \sigma + 1$. Forcing $T_\sigma \leq t_{\sigma+1}$ we will have $\tau_{M\sigma} \leq T_\sigma \leq t_{\sigma+1} \leq \tau_{M(\sigma+1)}$ and then MSM- σ will be better than MSM- $(\sigma + 1)$. This results in the next bound for $\sigma \geq 2$:

$$M \geq \frac{\left\lceil \log_{\phi_1}((n+1)\sqrt{5}+1) - 3 \right\rceil \lambda_{max} - \left\lceil \log_{(\sigma+1)}(n\sigma+1) - 1 \right\rceil \lambda_{min}}{s_r(\sigma+1) \cdot \mu_r - \sigma \cdot \mu_s} + 1$$

Note that we assume $s_r(\sigma + 1) \cdot \mu_r \geq \sigma \cdot \mu_s$ and thus $s_r(\sigma + 1) = \sigma + 1$ (since $\mu_r \leq \mu_s$). Otherwise, we should find tighter values for T_σ and $t_{\sigma+1}$ and then recalculate the bound for M . Finally, since we have a pessimistic case, we obtain for $\sigma \geq 2$:

$$M_\sigma \leq \frac{\left\lceil \log_{\phi_1}((n+1)\sqrt{5}+1) - 3 \right\rceil \lambda_{max} - \left\lceil \log_{(\sigma+1)}(n\sigma+1) - 1 \right\rceil \lambda_{min}}{(\sigma+1) \cdot \mu_r - \sigma \cdot \mu_s} + 1$$

References

- [1] J. H. Saltzer and D. P. Reed and D. D. Clark. End-to-end arguments in system design. *ACM Transactions on Computer Systems*, 2(4):277–278, 1984.
- [2] Amotz Bar-Noy and Shlomo Kipnis. Designing broadcasting algorithms in the postal model for message-passing systems. *SPAA '92: Proceedings of the fourth annual ACM symposium on Parallel algorithms and architectures*, 13–22, New York, NY, USA, 1992. ACM Press.