

A Complete Characterization of an Optimal Timer Based Selection Scheme

Virag Shah

Neelesh B. Mehta
Senior Member, IEEERaymond Yim
Member, IEEE

Abstract—Timer-based mechanisms are often used in several wireless systems to help a given (sink) node select the best helper node among many available nodes. Specifically, a node transmits a packet when its timer expires, and the timer value is a function of its local suitability metric. In practice, the best node gets selected successfully only if no other node's timer expires within a 'vulnerability' window after its timer expiry. In this paper, we provide a complete closed-form characterization of the optimal metric-to-timer mapping that maximizes the probability of success for any probability distribution function of the metric. The optimal scheme is scalable, distributed, and much better than the popular inverse metric timer mapping. We also develop an asymptotic characterization of the optimal scheme that is elegant and insightful, and accurate even for a small number of nodes.

I. INTRODUCTION

Selection is an attractive solution that has been proposed in many wireless communication schemes. For example, in cooperative communication systems, it helps exploit spatial diversity and avoids synchronization problems among multiple transmitting relays [1]–[6]. In cellular systems that exploit multiuser diversity, the base station transmits to the mobile station with the highest instantaneous channel. Selection finds applications in sensor networks [5], [7], [8] and vehicular ad-hoc networks (VANETs) [9], [10]. By accounting for the average throughput or average energy consumed, even fairness in selection can be ensured [11], [12].

Common to all the above systems is the following abstraction: *each node i maintains a local suitability metric μ_i , and the system attempts to select the 'best' node with the highest metric.* The mechanism that physically selects the best node is, therefore, an important component in many wireless systems. We consider a timer-based mechanism in this paper, which is popular because of its simplicity and distributed nature. It requires no feedback during the selection process. Each node sets its timer as a function of its metric value such that the timer of the node with the best metric expires the earliest. The timer mechanism is markedly different from the centralized polling mechanism, in which the sink polls each node about its metric and then chooses the best one, and from distributed splitting algorithms, which require slot-by-slot feedback [13], [14].

Ad hoc timer schemes have been proposed in the literature in [7], [15]. In [15], the timer is set as c/μ , where c is a

Virag Shah was at the Indian Institute of Science (IISc), Bangalore, India, during the course of this work. He is now with the Dept. of Electrical Eng. at the Indian Institute of Technology (IIT), Bombay, India. Neelesh B. Mehta is with the Dept. of Electrical Communication Eng. at IISc. Raymond Yim is with Mitsubishi Electric Research Labs (MERL), Cambridge, MA, USA.

constant. In [7], a piecewise linear function is used. In general, to ensure that the best node transmits first, the mapping is a deterministic monotone non-increasing function [7], [15]. In general, the timer scheme works by ensuring that the best node transmits first. However, for successful selection in practical systems, it is necessary that no other timer expires within a time window, called the *vulnerability window* Δ [16], of the expiry of the best node's timer. Otherwise, a selection failure occurs since two or more packets collide at the receiver. While scaling the timer values and increasing the total selection duration reduces the collision probability, it is undesirable because it reduces the time available to the selected node to transmit data. It also reduces the ability of the system to handle larger Doppler spreads. Therefore, optimizing over the mapping itself so as to maximize the probability of selection of the best node given T_{\max} and Δ is a problem of significant theoretical and practical interest.

In this paper, we consider the most general timer scheme in which the metric-to-timer function is monotone non-increasing, and provide a complete characterization of the optimal mapping in terms of T_{\max} and Δ . We show that to maximize the probability of selecting the best node, the nodes transmit only at finite discrete time instants, viz. $\{0, \Delta, \dots, N\Delta\}$, where $N = \lfloor \frac{T_{\max}}{\Delta} \rfloor$ and $\lfloor \cdot \rfloor$ is the floor function. We then provide a complete characterization of the timer scheme that specifies which node should transmit when. An expression for the optimal probability of success is also derived. In the asymptotic regime, where the number of nodes is large, we show that the characterization of the optimal scheme simplifies considerably. We show that the optimal scheme performs substantially better than the ad hoc mappings used in the literature, and is easy to implement in practice. Our results hold for all real-valued metrics with arbitrary probability distribution functions.

The paper is organized as follows. The system model and the general timer scheme are described in Sec. II. The optimal scheme is characterized in Sec. III. Section IV presents numerical simulations and comparisons, and is followed by our conclusions in Sec. V. The Appendix shows the key steps in the proofs; detailed proofs are available in [17].

II. TIMER-BASED SELECTION: SYSTEM BASICS

We consider a system with k nodes and a sink as shown in Figure 1. The sink represents any node that is interested in the message transmitted by the k nodes; it need not conduct any coordinating role. Each node i possesses a suitability metric




Fig. 1. A system consisting of a sink and k nodes. Each node has a metric μ_i and sets its timer as $T_i = f(\mu_i)$. The sink needs to select the node with the highest metric.

μ_i that is known only to that specific node. The metrics are assumed to be independent and identically distributed across nodes. The probability distribution is assumed to be known by all nodes. The aim of the selection scheme is to make the sink determine which node has the highest μ_i , which we henceforth call the ‘best’ node.

Each node i , based on its local metric μ_i , sets a timer $T_i = f(\mu_i)$, where $f(\cdot)$ is called the metric-to-timer function. When the timer expires (at time T_i), the node immediately transmits a packet to the sink. As mentioned, $f(\mu)$, in general, is a monotone non-increasing function to ensure that the best node’s timer expires first. The selection process has a maximum selection duration, T_{\max} , after which nodes do not start a transmission.

The sink can decode the packet from the best node, if the timers of the best and second best node, denoted by $T_{(1)}$ and $T_{(2)}$ respectively, expire such that $T_{(2)} - T_{(1)} \geq \Delta$. The expiry of timers of other nodes, which occurs after $T_{(2)}$, does not matter since the sink is only interested in the best node. The value of Δ depends on system capabilities. For example, Δ typically includes the maximum propagation and detection delays between all nodes. Carrier sensing is optional but beneficial; if it is not used, Δ needs to include maximum transmission time of packets. Δ may also include the receive-to-transmit switching times in half-duplex nodes [15].

Henceforth, we will abuse the above general definition of Δ and will say that a *collision* occurs when the timers of the best and the second best nodes expire within a duration Δ . Thus, the best node is selected successfully if $T_{(1)} \leq T_{\max}$ and $T_{(2)} - T_{(1)} \geq \Delta$. Otherwise, the selection process fails.

A. Discussion: Alternate Models and Extensions

In this paper, an inability to select the best node is treated as a failure or an outage. In fact, if a sink is available, it may respond to a selection failure in multiple ways. For example, it may resolve the nodes whose packets collided during the selection process, using extra feedback. If a sink is not available, then repeated transmission can be used to improve the overall reliability of broadcast messages. The details of how the system deals with a selection failure are beyond the scope of this paper.

In addition to maximizing the probability of successful selection, one can also choose to minimize the expected time

taken by the selection scheme to stop. The latter, in the presence of a constraint on the probability of success, is characterized in [17]. The scheme considered in this paper, which maximizes the probability of success, provides a feasibility criterion for the above scheme.

To keep notation simple, we first consider the case where the metrics are uniformly distributed over the interval $[0, 1]$. Thereafter, the results are generalized to all real-valued metrics with arbitrary probability distribution functions.

III. OPTIMAL TIMER SCHEME’S STRUCTURE

We first consider the case where each metric is uniformly distributed between 0 and 1. This is generalized to metrics with arbitrary probability distribution functions in Sec. III-B.

We shall use the following notation henceforth. $\mathbf{E}[X]$ denotes the expected value of a random variable (RV) X . Using order statistics notation, the node with the i th largest metric is denoted by (i) . Consequently, $\mu_{(1)} \geq \mu_{(2)} \geq \dots \geq \mu_{(k)}$ and $T_{(1)} \leq \dots \leq T_{(k)}$. For notational convenience, the summation $\sum_{i=l_1}^{l_2}$ equals 0 whenever $l_1 > l_2$. We use the superscript $*$ to denote an optimal value; for example, optimal value of x is x^* . $\Pr(A)$ denotes the probability of an event A , and $\Pr(A|B)$ denotes the conditional probability of A given B .

The following lemma shows that an optimal $f^*(\mu)$ maps the metrics into discrete timer values. Let $N = \lfloor \frac{T_{\max}}{\Delta} \rfloor$.

Lemma 1: An optimal metric-to-timer mapping $f^*(\mu)$ that maximizes the probability of success within a maximum time T_{\max} maps μ into $(N + 1)$ discrete timer values $\{0, \Delta, 2\Delta, \dots, N\Delta\}$.

Proof: The proof is given in Appendix A. ■

Note that the above discrete mapping, while optimal, need not be unique. Intuitively, this result seems to be in sync with the fact that time slotted multiple access protocols have better throughput than the unslotted ones. However, there is an important difference. While time slotting reduces the vulnerability window in multiple access protocols, in our selection problem the vulnerability window remains unchanged.

Implications of Lemma 1: We have reduced an infinite-dimensional problem of finding $f(\mu)$ over the space of all positive-valued monotone non-increasing functions to one over just $N + 1$ real values that lie between 0 and T_{\max} , as illustrated in Figure 2. Only the contiguous metric intervals in $[0, 1)$ that get assigned to the timer values $0, \Delta, \dots, N\Delta$ remain to be determined. As shown in Figure 2, all nodes with metrics in the interval $[1 - \alpha_N[0], 1)$, of length $\alpha_N[0]$, set their timers to 0. Nodes with metrics in the next interval $[1 - \alpha_N[1] - \alpha_N[0], 1 - \alpha_N[0])$, of length $\alpha_N[1]$, set their timers to Δ , and so on. In general, nodes with metrics in the interval $[1 - \sum_{j=0}^i \alpha_N[j], 1 - \sum_{j=0}^{i-1} \alpha_N[j])$, of length $\alpha_N[i]$, set their timer to $i\Delta$. Any node with metric less than $(1 - \sum_{j=0}^N \alpha_N[j])$ does not transmit at all.

The following theorem provides optimal values for $\alpha_N[i]$, for $0 \leq i \leq N$, and, thus, along with Lemma 1 provides a complete characterization of the optimal scheme. Let $N = \lfloor \frac{T_{\max}}{\Delta} \rfloor$.




Fig. 2. Illustration of the optimal metric-to-timer mapping $f^*(\mu)$. A user with metric in the interval $[1 - \alpha_N[0], 1)$ transmits at time 0, a user with metric in the interval $[1 - \alpha_N[0] - \alpha_N[1], 1 - \alpha_N[0])$ transmits at time Δ , and so on. Users with metrics less than $1 - \sum_{i=0}^N \alpha_N[i]$ do not transmit.

Theorem 1: The probability of selecting the best node among k nodes, subject to a maximum selection duration of T_{\max} , is maximized when the metric-to-timer mapping is

$$f^*(\mu) = \begin{cases} i\Delta, & 1 - \sum_{j=0}^i \alpha_N^*[j] \leq \mu < \\ & 1 - \sum_{j=0}^{i-1} \alpha_N^*[j], \text{ for } 0 \leq i \leq N, \\ T_{\max} + \epsilon, & \text{otherwise} \end{cases} \quad (1)$$

where ϵ is any arbitrary strictly positive real number. The $N+1$ interval lengths $\alpha_N^*[0], \dots, \alpha_N^*[N]$ are recursively given by

$$\alpha_N^*[j] = \begin{cases} \frac{1 - P_N^*}{k - P_N^*}, & j = 0 \\ (1 - \alpha_N^*[0]) \alpha_{N-1}^*[j-1], & 1 \leq j \leq N \end{cases}, \quad (2)$$

where $\alpha_0^*[0] = 1/k$. P_N^* is the maximum probability of success and is given by

$$P_N^* = k \sum_{l=0}^N \alpha_N^*[l] \left(1 - \sum_{j=0}^l \alpha_N^*[j] \right)^{k-1}. \quad (3)$$

Proof: The proof is given in Appendix B. ■

The discrete nature of the optimal scheme also makes it amenable to practical implementation since each node only needs to store an unwrapped version of the above recursion in the form of a lookup table that has $N+1$ entries $\{\alpha_N^*[0], \dots, \alpha_N^*[N]\}$. Notice that the interval lengths depend on k . Such a dependence is also characteristic of splitting-based selection schemes [13], [14].

A. Asymptotic Analysis of Optimal Scheme as $k \rightarrow \infty$

We now provide asymptotic expressions for the optimal timer scheme as the number of nodes $k \rightarrow \infty$. The maximum selection duration, T_{\max} , or equivalently N , is kept fixed. As we shall see, the recursions simplify to a simple and elegant form. This is because in the asymptotic regime, a scaled version of the metric follows a Poisson process [18].

From (2), it can be seen that $\alpha_N^*[j]$ tends to zero as $k \rightarrow \infty$. Therefore, for node i , consider a scaled metric $y_i = k(1 - \mu_i)$, and normalize the interval lengths to

$$\beta_N^*[j] = k\alpha_N^*[j]. \quad (4)$$

Thus, selecting a node with the highest μ_i is equivalent to selecting the node with the lowest y_i . Let $y_{(1)} \leq y_{(2)} \leq \dots \leq y_{(k)}$ denote an ascending ordering on the y_i s. Define the point process $M(z) \triangleq \sup \{k \geq 1 : y_{(k)} \leq z\}$. $M(z)$ is simply the number of nodes whose $y_i = k(1 - \mu_i)$ is less than z .

Lemma 2: $M(z)$ forms a Poisson process as $k \rightarrow \infty$. ■ This result enables the following use of the *independent increments property* of Poisson processes [18, Chp. 2], which states that the number of points that occur in disjoint intervals are independent of each other.

Theorem 2: The optimal $\beta_N^*[j]$ that maximize the probability of success are given by

$$\beta_N^*[j] = \begin{cases} 1, & j = N \\ 1 - e^{-\beta_N^*[j+1]}, & 0 \leq j \leq N-1 \end{cases}. \quad (5)$$

Also, the optimal probability of success is $P_N^* = e^{-\beta_N^*[0]}$.

Proof: The proof is omitted to conserve space and is given in detail in [17]. ■

Theorem 2 leads to several key insights about the optimal timer scheme, which are formally stated as corollaries below.

Corollary 1 (Scalability): Even when $k \rightarrow \infty$, $P_N^* \geq 1/e$, with equality occurring only for $N = 0$.

Corollary 2 (Independence): $\beta_N^*[N-r]$ depends only on r , and is independent of N .

Corollary 3 (Monotonicity): $\beta_N^*[0] < \dots < \beta_N^*[N]$.

The monotonicity property can be intuitively understood as follows. As the time available decreases, the risk of selection failure due to the best node not transmitting increases. This is counteracted by increasing the risk of collision.

B. Extension to Arbitrary Real-Valued Metrics

We now generalize the optimal solution to the general case where the metric is not uniformly distributed. Let the cumulative distribution function (CDF) of a metric be denoted by $F_c(x) = \Pr(\mu \leq x)$, where $-\infty < x < \infty$.

The optimum mapping when the CDF of the metric is $F_c(\cdot)$ is $f^*(F_c(\mu))$, where $f^*(\cdot)$ is given by Theorem 2. This follows because $F_c(\cdot)$ is a monotonically non-decreasing function, and the RV $Y = F_c(\mu)$ is uniformly distributed between 0 and 1.¹ The problem has, therefore, been reduced to the one considered earlier. This also shows that the performance for the optimal mapping does not depend on $F_c(\cdot)$. Note here that we assume that the nodes know $F_c(\cdot)$, as was also assumed in [13], [14]. Practically, this is justified because $F_c(\cdot)$, being a statistical property, can be computed over time.

IV. RESULTS AND PERFORMANCE EVALUATION

We now study the performance of the optimum timer scheme. We also compare it with the popular inverse metric timer mapping that uses $f(\mu) = c/\mu$ [2], [3], [15], [19].² In

¹The CDF needs to be continuous to ensure this. The case where the CDF is not continuous can be easily handled by a technique analogous to *proportional expansion* that was proposed in [14] for splitting algorithms. In it, each node generates a new continuous metric such that at least one of the nodes with the highest metric still remains the best node.

²We do not plot the piece-wise linear mapping of [7] because its performance needs to be numerically optimized over several parameters.

order to ensure a fair comparison with the optimal scheme, we optimize the inverse timer scheme as well. For each pair of T_{\max} and k values, the value of c that maximizes the probability of success is numerically determined from extensive Monte Carlo simulations. Since it also depends on $F_c(\cdot)$, we show results for the following two metric distributions: (i) A unit mean Rayleigh distribution with CDF $F_c(\mu) = 1 - e^{-\mu^2/2}$, and (ii) A unit mean exponential distribution with CDF $F_c(\mu) = 1 - e^{-\mu}$.

Figure 3 plots the maximum probability of success, P_N^* , as a function of $N = \lfloor \frac{T_{\max}}{\Delta} \rfloor$. Also plotted are results from Monte Carlo simulations, which match very well with the analytical results. Notice that the asymptotic curve is close to the actual curve even for $k = 5$. The asymptotic curve shows a rather remarkable result: *regardless of k and without the use of any feedback, the best node gets selected with a probability of over 75% when N is just 5. When N increases to 17, the success probability exceeds 90%. $N = 17$ may seem like a large number compared to the 2.467 time slots required on average by the splitting algorithm to select the best node [13], [14]. However, the slot duration in the splitting algorithm is much larger than Δ because of the slot-by-slot transmission and feedback required by the splitting algorithm. A detailed comparison based on the IEEE 802.11 standard's parameters shows that a slot in the splitting algorithm can be eleven times larger than Δ [17].*

Also plotted in Figure 3 are the results for the inverse timer mapping scheme, when c optimized for each parameter set and when c is kept fixed. We see that the optimal scheme significantly outperforms the inverse metric mapping, even when the latter's parameters are optimized. For example, for $N = 10$ and $N = 30$, the probability that the system fails to select the best node for the inverse timer scheme is respectively 3.4 and 5.0 times greater than that of the optimum scheme for the Rayleigh CDF. The factors increase to 3.7 and 7.2 for the exponential CDF. Interestingly, even though the exponential RV is a square of the Rayleigh RV and the squaring operation preserves the metric order, the performance of the inverse timer scheme changes.

The optimal scheme's parameters are studied in Figure 4, which plots $\alpha_N^*[j]$ for $N = 10$ when the metric is uniformly distributed between 0 and 1. (The parameters for arbitrary distributions can be obtained using Sec. III-B.) We see that $\alpha_N^*[j]$ increases with j , which is in line with Corollary 3.

V. CONCLUSIONS

We considered a distributed timer-based selection scheme, in which each node maps its priority metric to a timer value, and transmits when its timer expires. It works by ensuring that the best node's timer expires first. We showed that the optimal mapping that maximizes the probability of selection maps the metrics into $N + 1$ discrete timer values, where $N = \lfloor T_{\max}/\Delta \rfloor$. Thus, a smaller vulnerability window Δ or a larger maximum selection duration T_{\max} improves performance.

In the asymptotic regime, where the number of nodes is large, the occurrence of a Poisson process led to a consid-

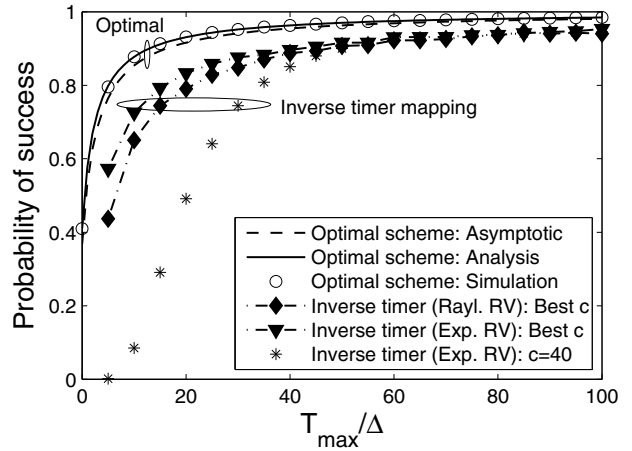


Fig. 3. Optimum P_N^* as a function of $\frac{T_{\max}}{\Delta}$ for $k = 5$. Also plotted is the benchmark performance of the ad hoc inverse metric mapping both when c is optimized and when c is kept fixed.




Fig. 4. $\alpha_N^*[j]$ as a function of j and the number of nodes, k , for $N = 10$.

erably simpler characterization of the optimal scheme. The optimal schemes' performance was significantly better than the inverse metric mapping. Unlike the latter mapping, the optimal mapping's performance did not depend on the probability distribution of the metric.

The optimal timer in view of its scalability and superior performance is likely to find applications in several areas. Examples of these include cooperative systems that need to find the best relay node [2], [15], wireless network coding to find the best relays that will combine the signals transmitted by multiple sources [19], mobile multi-hop networks [6], VANETs to determine which vehicle should rebroadcast an emergency message [9], [10], wireless local area networks [20] to enable opportunistic channel access, and sensor networks to improve network lifetime [3], [7].

APPENDIX

Given the space constraints, the key steps in the proofs are highlighted below. Detailed proofs are given in [17].

A. Proof of Lemma 1

The key idea behind the proof is to successively refine $f(\cdot)$, by making parts of it discrete, and show that this can only improve the probability of success.

Let $f(\mu)$ be an optimal monotone non-increasing mapping. If $T_{\max} < \Delta$ (i.e., $N = 0$), the problem is trivial since setting all timers to 0 does not change the probability of success. When $\Delta \leq T_{\max} < 2\Delta$, let:

$$f_1(\mu) = \begin{cases} 0, & 0 \leq f(\mu) < \Delta \\ f(\mu), & \text{else} \end{cases}. \quad (6)$$

It can be shown that $f_1(\mu)$ is monotone non-decreasing and its probability of success is greater than or equal to that of $f(\mu)$. Since $T_{\max} - \Delta < \Delta$, setting all the remaining timers that lie in $[\Delta, T_{\max}]$ to Δ does not change the probability of success. Thus, the discrete mapping is optimal for $T_{\max} < 2\Delta$.

When $T_{\max} \geq 2\Delta$, consider the new mapping $f_2(\mu)$ such that $f_2(\mu) = \begin{cases} \Delta, & \Delta \leq f_1(\mu) < 2\Delta \\ f_1(\mu), & \text{else} \end{cases}$. It can again be shown that the probability of success of $f_2(\cdot)$ is greater than or equal to that of $f_1(\cdot)$. A successive application of this argument shows that an optimal mapping takes values in the discrete set $\{0, \Delta, \dots, N\Delta\}$.

B. Proof of Theorem 1

Success occurs at time $l\Delta$, for $l = 0, \dots, N$, if $\mu_{(1)}$ lies in $\left[\left(1 - \sum_{j=0}^l \alpha_N[j]\right), \left(1 - \sum_{j=0}^{l-1} \alpha_N[j]\right) \right)$ and the remaining $k-1$ metrics lie in $\left[0, \left(1 - \sum_{j=0}^l \alpha_N[j]\right)^{k-1}\right)$. This occurs with probability $k\alpha_N[l] \left(1 - \sum_{j=0}^l \alpha_N[j]\right)^{k-1}$ since the metrics are independent and identically distributed (i.i.d.) and uniform over $[0, 1)$. Summing over l gives (3).

We henceforth denote the probability of success by $P_N(\alpha_N[0], \dots, \alpha_N[N])$ to show its dependence on $\{\alpha_N[i]\}_{i=0}^N$. Let the maximum probability of success, P_N^* , occur when $\alpha_N[i] = \alpha_N^*[i]$, $0 \leq i \leq N$. Note that $\sum_{i=0}^N \alpha_N[i] \leq 1$. A key step in the proof is that, for $N \geq 1$, the probability of success can be written as follows:

$$P_N(\alpha_N[0], \dots, \alpha_N[N]) = \Pr(\mu_{(1)} \in [1 - \alpha_N[0], 1)) \Pr(\text{success} | \mu_{(1)} \in [1 - \alpha_N[0], 1)) + \Pr(\mu_{(1)} \notin [1 - \alpha_N[0], 1)) \Pr(\text{success} | \mu_{(1)} \notin [1 - \alpha_N[0], 1)).$$

This helps reduce the number of variables that need to be optimized from $N+1$ to N as follows. Conditioned on $\mu_{(1)} \notin [1 - \alpha_N[0], 1)$, the k metrics are i.i.d. and uniformly distributed over the interval $[0, 1 - \alpha_N[0])$. Therefore, $\Pr(\text{success} | \mu_{(1)} \notin [1 - \alpha_N[0], 1)) \leq P_{N-1}^*$. Hence,

$$P_N(\alpha_N[0], \dots, \alpha_N[N]) \leq k\alpha_N[0](1 - \alpha_N[0])^{k-1} + (1 - \alpha_N[0])^k P_{N-1}^*. \quad (7)$$

The bound is achieved when $\frac{\alpha_N[1]}{1 - \alpha_N^*[0]} = \alpha_{N-1}^*[0], \dots, \frac{\alpha_N[N]}{1 - \alpha_N^*[0]} = \alpha_{N-1}^*[N-1]$, for any $0 \leq \alpha_N[0] < 1$. Therefore,

$$P_N^* = k\alpha_N^*[0](1 - \alpha_N^*[0])^{k-1} + (1 - \alpha_N^*[0])^k P_{N-1}^*. \quad (8)$$

Using the first order condition, we get $\alpha_N^*[0] = \frac{1 - P_{N-1}^*}{k - P_{N-1}^*}$. Similarly, for $N = 0$, we can show that $P_0^* = (1 - 1/k)^{k-1}$ and $\alpha_0^*[0] = 1/k$.

The value of $f^*(\mu)$ when it exceeds T_{\max} can be left unspecified because such a node will not transmit. This is ensured by setting $f^*(\mu)$ to $T_{\max} + \epsilon$, where $\epsilon > 0$.

REFERENCES

- [1] Z. Lin and E. Erkip, "Relay search algorithms for coded cooperative systems," in *Proc. Globecom*, 2005.
- [2] E. Beres and R. Adve, "Selection cooperation in multi-source cooperative networks," *IEEE Trans. Wireless Commun.*, vol. 7, pp. 118–127, Jan. 2008.
- [3] Z. Zhou, S. Zhou, J.-H. Cui, and S. Cui, "Energy-efficient cooperative communication based on power control and selective single-relay in wireless sensor networks," *IEEE Trans. Wireless Commun.*, vol. 7, pp. 3066–3078, Aug. 2008.
- [4] Y. Jing and H. Jafarkhani, "Single and multiple relay selection schemes and their achievable diversity orders," *IEEE Trans. Wireless Commun.*, vol. 8, pp. 1414–1423, Mar. 2009.
- [5] W.-J. Huang, Y.-W. P. Hong, and C.-C. J. Kuo, "Lifetime maximization for amplify-and-forward cooperative networks," *IEEE Trans. Wireless Commun.*, vol. 7, pp. 1800–1805, May 2008.
- [6] J. Kim and S. Bohacek, "A comparison of opportunistic and deterministic forwarding in mobile multihop wireless networks," in *Proc. 1st Intl. MobiSys Workshop Mobile Opportunistic Netw.*, pp. 9–16, Jun. 2007.
- [7] Q. Zhao and L. Tong, "Opportunistic carrier sensing for energy-efficient information retrieval in sensor networks," *EURASIP J. Wireless Commun. and Networking*, pp. 231–241, May 2005.
- [8] Y. Chen and Q. Zhao, "An integrated approach to energy-aware medium access for wireless sensor networks," *IEEE Trans. Signal Process.*, vol. 7, pp. 3429–3444, Jul. 2007.
- [9] M. Nekovee and B. B. Bogason, "Reliable and efficient information dissemination in intermittently connected vehicular adhoc networks," in *Proc. VTC (Spring)*, pp. 2486–2490, 2007.
- [10] R. Yim, J. Guo, P. Orlik, and J. Zhang, "Received power-based prioritized rebroadcasting for V2V safety message dissemination," in *Int. Transport. Sys. World Congr.*, Sept. 2009.
- [11] D. Tse and P. Vishwanath, *Fundamentals of Wireless Communications*. Cambridge University Press, 2005.
- [12] D. S. Michalopoulos and G. K. Karagiannidis, "PHY-layer fairness in amplify and forward cooperative diversity systems," *IEEE Trans. Wireless Commun.*, vol. 7, pp. 1073–1083, Mar. 2008.
- [13] X. Qin and R. Berry, "Opportunistic splitting algorithms for wireless networks," in *Proc. INFOCOM*, pp. 1662–1672, Mar. 2004.
- [14] V. Shah, N. B. Mehta, and R. Yim, "Splitting algorithms for fast decentralized cooperative relay selection," *To appear in IEEE Trans. Wireless Commun. (Manuscript available online at: http://arxiv.org/abs/0911.4357)*, 2010.
- [15] A. Bletsas, A. Khisti, D. P. Reed, and A. Lippman, "A simple cooperative diversity method based on network path selection," *IEEE J. Sel. Areas Commun.*, vol. 24, pp. 659–672, Mar. 2006.
- [16] L. Kleinrock and F. Tobagi, "Packet switching in radio channels: Part I – Carrier sense multiple access modes and their throughput delay characteristics," *IEEE Trans. Commun.*, vol. 23, pp. 1400–1416, Dec. 1975.
- [17] V. Shah, N. B. Mehta, and R. Yim, "Optimal timer based selection schemes," *To appear in IEEE Trans. Commun. (Manuscript available online at: http://arxiv.org/abs/0909.1241)*, 2010.
- [18] R. W. Wolff, *Stochastic Modeling and the Theory of Queues*. Prentice Hall, 1989.
- [19] Z. Ding, T. Ratnarajah, and K. K. Leung, "On the study of network coded AF transmission protocol for wireless multiple access channels," *IEEE Trans. Wireless Commun.*, vol. 7, pp. 4568–4574, Nov. 2008.
- [20] C. S. Hwang, K. Seong, and J. M. Cioffi, "Throughput maximization by utilizing multi-user diversity in slow-fading random access channels," *IEEE Trans. Wireless Commun.*, vol. 7, pp. 2526–2535, Jul. 2008.