Feedback Overhead-Aware Fast Distributed Selection Scheme for Multi-node Wireless Systems

Rajat Talak, *Student Member, IEEE*, and Neelesh B. Mehta, *Senior Member, IEEE*

***Abstract*—Opportunistic selection is a practically appealing technique that is used in multi-node wireless systems to maximize throughput, implement proportional fairness, etc. However, selection is challenging since the information about a node's channel gains is often available only locally at each node and not centrally. We propose a novel multiple access-based distributed selection scheme that generalizes the best features of the timer scheme, which requires minimal feedback but does not always guarantee successful selection, and the fast splitting scheme, which requires more feedback but guarantees successful selection. The proposed scheme's design explicitly accounts for feedback time overheads unlike the conventional splitting scheme and guarantees selection of the user with the highest metric unlike the timer scheme. We analyze and minimize the average time including feedback required by the scheme to select. With feedback overheads, the proposed scheme is scalable and considerably faster than several schemes proposed in the literature. Furthermore, the gains increase as the feedback overhead increases.**

I. Introduction

Opportunistic selection is a practically effective technology in diverse multi-node wireless systems such as cellular systems [1], relay-based cooperative communication systems [2]–[4], and sensor networks [5]. In general, in opportunistic selection, the source or destination or base station (henceforth referred to as 'sink') needs to select a node based on its channel gains. For example, selecting the user with the highest instantaneous rate maximizes throughput. On the other hand, selecting the user with the highest ratio of instantaneous rate to average rate ensures fairness but also exploits spatial diversity.

However, information about the channel gains is available only locally at the nodes. Based on its local knowledge, each node maintains a local preference metric. The sink finds the node with the highest metric by running a selection scheme. For example, in an amplify-and-forward relay, the metric, μ, as a function of the source-to-relay gain, h_{sr}, and relay-to-destination channel gain, h_{rd}, is given by $\mu = \frac{h_{\text{sr}} h_{\text{rd}}}{h_{\text{sr}} + h_{\text{rd}}}$ [3]. In a cellular downlink, the metric is the channel power gain, h, from the base station to the node and is given by $\mu = h$ [1].

A simple, but centralized, selection scheme is polling. In it, each node sequentially reveals its metric to the sink, which then selects the best node. However, the time taken by polling grows linearly with the number of nodes. Distributed selection schemes successfully tackle this problem [2], [5]–[8]. One prominent example is the time-slotted splitting-based selection scheme [6], [7]. In it, all nodes whose metrics lie in between two thresholds, independently transmit in a slot. At the end of every slot, the sink feeds back one of three possible outcomes – idle, success, or collision, depending on whether 0, 1, or multiple nodes transmitted. The thresholds are updated based on this feedback. The splitting scheme is attractive because it provably requires less than 2.52 time slots, on average, to select the best node even with infinitely many nodes [6].

Another interesting and simple alternative is the timer-based selection scheme [2], [5], [8]. Here, every node sets a timer as a function of its metric. When a node's timer expires, it transmits a small packet to the sink. The mapping is such that the best node's timer always expires first. However, a collision occurs if another node's timer expires within a vulnerability window, Δ, from the time of expiry of the best node's timer [9], and the scheme fails to select the best node. It has been shown in [8] that the optimal timer mapping that maximizes the probability of success for a given maximum selection time, $T_{\max}$, is a simple staircase function in which the timers expire only at $\{0, \Delta, 2\Delta, \ldots\}$.

While the timer and splitting schemes accomplish the same goal, they are fundamentally different. The time required by the splitting scheme depends on the metric realizations, but the best node is guaranteed to get selected. On the other hand, the timer scheme runs for a fixed duration, $T_{\max}$, but may fail to select the best node [2], [5], [8]. To resolve this problem, doubling $T_{\max}$ every time a selection failure occurs has been considered in [10]. In [11], [12], metrics are instead randomized. However, such randomization degrades overall system performance since the best node may not get selected.

The timer scheme does not require any feedback, except at the end to inform the node that has been selected. On the other hand, the splitting scheme uses slot-by-slot feedback, whose overhead can be significant. For example, in an 802.11 wireless local area network (WLAN), the feedback duration is at least $F = 131$ μsec [13, Table 17-15], while Δ, which characterizes the timer function, is just 13 μsec [8]. Since the slot duration must accommodate feedback, the total time required by the splitting scheme to select can be large.

In this paper, we propose a novel selection scheme that combines the best features of the splitting and timer schemes and generalizes both of them. The feedback duration is explicitly accounted for. The proposed scheme runs in stages, and each stage is followed by an idle/success/collision feedback from the sink. In each stage, nodes whose metrics lie within two transmission thresholds, set timers and transmit a small packet when their timers expire. *Thus, the timer scheme is run in each stage.* Unlike the splitting scheme, the time of occurrence of

The authors are with the Dept. of Electrical Communication Eng. in the Indian Institute of Science (IISc), Bangalore, India.

This work was partially supported by a research grant from the DRDO-IISc program on mathematical engineering.

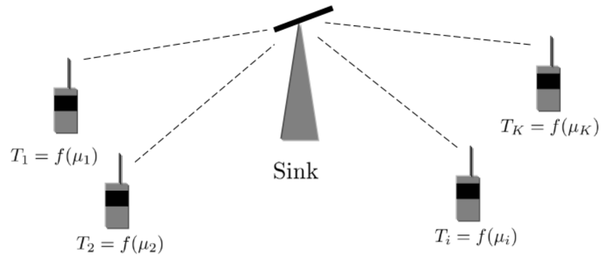

Fig. 1. A wireless system consisting of K nodes and a sink. Each node i tracks a local metric μ_i. The sink has to select the best node with the highest metric. In the timer scheme, a node i sets its timer, T_i, as $T_i = f(\mu_i)$.

the first collision is also fed back with a resolution of Δ, which is easily facilitated by the staircase nature of the optimal timer mapping. Notably, the timer scheme's parameters in different stages depend on the feedback. Unlike the timer scheme, the best node always gets selected.

We analytically characterize the average time required by the scheme to select the best node, and minimize it. The proposed scheme significantly outperforms the conventional splitting and timer schemes, and other schemes in the literature. Further, the splitting scheme is shown to be a special case of the proposed scheme and is sub-optimal unless $F = 0$.

This paper is organized as follows. Section II describes the system model and basic timer-based selection scheme. We develop the proposed scheme for selection in Sec. III. In Sec. IV, we benchmark its performance with other selection schemes. We conclude in Sec. V.

II. System Model and Timer-Based Selection

As shown in Figure 1, we consider a system consisting of K nodes and a sink. Each node i computes a metric, μ_i, which is not known to any other node. The goal is to find the best node with the highest metric, $\max_{i \in \{1,2,\ldots,K\}} \mu_i$. The metrics are assumed to be independent and identically distributed (i.i.d.) across nodes, as has been assumed in [2]–[6], [10]–[12]. The nodes are assumed to know the continuous cumulative distribution function (CDF) of μ_i, $C(\cdot)$, and K, as has also been assumed in [6], [10]–[12].

In conventional single-stage timer-based selection, each node i sets a timer $T_i = f(\mu_i)$, where $f : [0, \infty) \to [0, T_{\max}]$ is the metric-to-timer mapping. When the timer expires at time T_i, the node transmits a small packet to the sink. The mapping is a monotonically non-increasing deterministic function to ensure that the node with the highest metric transmits first.

For the sink to successfully decode the packet from the best node, no other node's timer must expire within a vulnerability window, Δ, of the expiry of the timer of the best node [8]. The physical layer capabilities of the wireless system determine Δ. It includes the maximum propagation delay, the maximum delay spread in the channels seen by the nodes, and time synchronization errors, if any, among nodes. When the nodes do not use carrier sensing, Δ includes the packet duration. However, carrier sense multiple access with collision avoidance (CSMA/CA) makes Δ smaller as the nodes avoid transmitting and colliding when the channel is busy. However, in this paper, the nodes are not required to have the CSMA/CA capability.

A. Single-Stage Optimal Timer Scheme

Without loss of generality, let the metrics be uniformly distributed over $[0, 1]$.[1] Let, $\lceil \cdot \rceil$ and $\lfloor \cdot \rfloor$ denote the ceil and floor function respectively. Given a $T_{\max}$ and Δ, the optimal timer mapping that maximizes the probability of selecting the best node is given by [8]

$$f(\mu) = \begin{cases} 0, & 1 - \alpha_N[0,K] \le \mu < 1, \\ \Delta, & 1 - \sum_{j=0}^{1} \alpha_N[j,K] \le \mu < 1 - \alpha_N[0,K] \\ \cdots & \\ N\Delta, & 1 - \sum_{j=0}^{N} \alpha_N[j,K] \le \mu < 1 - \sum_{j=0}^{N-1} \alpha_N[j,K] \\ N\Delta + \epsilon, & \text{otherwise} \end{cases},$$

where $N = \left\lfloor \frac{T_{\max}}{\Delta} \right\rfloor$ is called the *number of levels* and $\epsilon > 0$ is an arbitrary number to indicate that no transmission occurs after $T_{\max}$. Here, $\alpha_N[i,K]$ is called the ith step length. Note that nodes whose metrics lie in the interval $\left[0, 1 - \sum_{j=0}^{N} \alpha_N[j,K]\right)$ do not transmit.

The optimal step lengths that maximize probability of success are given by the following recursion [8]:

$$\alpha_N[j,K] = \begin{cases} \frac{1 - P_s(N-1,K)}{K - P_s(N-1,K)}, & j = 0 \\ (1 - \alpha_N[0,K])\alpha_{N-1}[j-1,K], & 1 \le j \le N \end{cases}, \quad (1)$$

where, for $N = 0$, $\alpha_0[0,K] = 1/K$ and $P_s(N,K)$, which denotes the probability of success, is given by $P_s(N,K) = K \sum_{i=0}^{N} \alpha_N[i,K] \left(1 - \sum_{j=0}^{i} \alpha_N[j,K]\right)^{K-1}$. Note that the step lengths depend on the number of levels and the number of nodes that participate in the single-stage timer scheme.

The single-stage optimal timer scheme may fail to select the best node if no timer expires within $T_{\max}$, which we term as *idle*, or if the best node's timer expires in the same slot as that of the second best node's, which we term as *collision*.

III. Proposed Selection Scheme

Before we present the proposed scheme for the general case of $K \ge 2$ nodes, we first explain and optimize it for the special, but insightful, case of $K = 2$ nodes. The notation $X \sim U[a, b)$ shall indicate that the random variable (RV) X is uniformly distributed over the interval $[a, b)$.

A. System with Two Nodes ($K = 2$)

Let an N-level single-stage optimal timer scheme be run for a system with 2 nodes. We make the following two important observations.

Observation *1:* Let $\mathcal{E}_{C(s)}$ denote the event in which a collision occurs at $s\Delta$, $0 \le s \le N$, in the single-stage optimal timer scheme. Then, given $\mathcal{E}_{C(s)}$, we know

[1] In general, if μ_i has a CDF $C(\cdot)$, then $\widetilde{\mu}_i = C(\mu_i)$ is uniformly distributed over $[0, 1]$. Thus, the ith node can use $\widetilde{\mu}_i$ as its metric.

that $\mu_i \sim U\left[1-\sum_{i=0}^{s}\alpha_N[i,2]\,, 1-\sum_{i=0}^{s-1}\alpha_N[i,2]\right)$, for both nodes, $i = 1, 2$.
Therefore, the rescaling of μ_i to $R_{C(s)}(\mu_i)$, where

$$R_{C(s)}(\mu_i) = \frac{\mu_i - 1 + \sum_{j=0}^{s} \alpha_N[j,2]}{\alpha_N[s,2]}, \tag{2}$$

ensures that $R_{C(s)}(\mu_i) \sim U[0,1)$, for both nodes, $i = 1, 2$. Also, note that $R_{C(s)}(\mu_1)$ and $R_{C(s)}(\mu_2)$ are independent RVs. Moreover, this rescaling preserves order, i.e., if $\mu_1 > \mu_2$ then $R_{C(s)}(\mu_1) > R_{C(s)}(\mu_2)$.

Observation *2:* Let $\mathcal{E}_I$ denote the event of an idle. Then, given $\mathcal{E}_I$, $\mu_i \sim U\left[0, 1-\sum_{i=0}^{N}\alpha_N[i,2]\right)$, for both nodes, $i = 1, 2$.
Therefore, the rescaling of μ_i to $R_I(\mu_i)$, where

$$R_I(\mu_i) = \frac{\mu_i}{1 - \sum_{j=0}^{N} \alpha_N[j,2]}, \tag{3}$$

ensures that $R_I(\mu_i) \sim U[0,1)$, for both nodes, $i = 1, 2$. As before, this rescaling preserves independence and order.

These observations show how to construct a new set of rescaled metrics, which are again independent and uniformly distributed over $[0,1)$ with the order preserved. We now propose a multi-stage selection scheme for $K = 2$ nodes. It is based on the above two observations, which help set up the next timer stage in the event of a failure, i.e., idle or collision.

- *Start:* An N-level single-stage optimal timer scheme that is designed for 2 nodes is run.
- *Success:* The scheme terminates after $(N + 1)\Delta$ sec and the best node is notified about its selection with a broadcast message of duration F sec.[2] All nodes then know that the node that transmitted in the success slot has been selected.
- *Collision:* The sink feeds back the slot, s, in which the collision occurred. This requires F sec. Thereafter, node i, for $i = 1, 2$, rescales its metric to $R_{C(s)}(\mu_i)$, and participates in another run of the N-level single-stage optimal timer scheme.
- *Idle:* The sink feeds back an idle outcome. This again requires F sec. Thereafter, node i, for $i = 1, 2$, rescales its metric to $R_I(\mu_i)$, and participates in another run of the N-level single-stage optimal timer scheme.

The above procedure is repeated until a success occurs. Since the rescalings in (2) and (3) preserve order, it is guaranteed that the first successful transmission to the sink is from the best node.

We now obtain an expression for the average selection time.

Theorem *1:* For a system with two nodes, the average time, $T_2(N)$, required by the proposed scheme to select the best node when the number of levels in each stage is N equals

$$T_2(N) = (N + 2)\Delta + \left(\frac{N+2}{N+1}\right) F. \tag{4}$$

Proof: The proof is relegated to Appendix A. ■

[2]We assume that the feedback is error-free, as is often been assumed in the literature [2], [6], [10]–[12]. This is reasonable given its low payload.

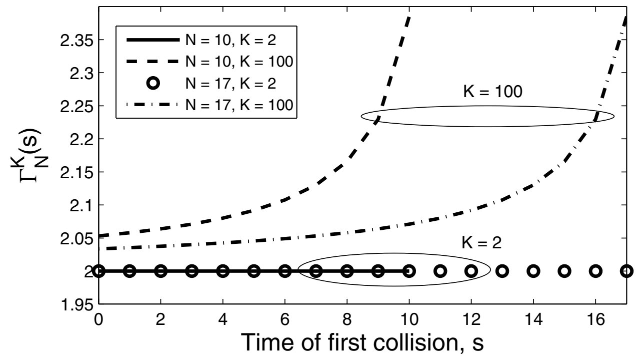

Fig. 2. Plot of the average number of nodes, $\Gamma_N^K(s)$, that collide at $s\Delta$ given that the first attempted transmission was at $s\Delta$.

B. General Case of $K \geq 2$ Nodes

Let an N-level single-stage optimal timer scheme be run. Then, along the lines of Observation 2, we can again show that in the event of an idle, for node i, $1 \leq i \leq K$, $\mu_i \sim U\left[0, 1-\sum_{j=0}^{N}\alpha_N[j,K]\right)$. Hence, the following rescaling

$$R_I^{(N,K)}(\mu_i) = \frac{\mu_i}{1 - \sum_{j=0}^{N} \alpha_N[j,K]}, \tag{5}$$

ensures that $R_I^{(N,K)}(\mu_i) \sim U[0,1)$, for all the K nodes. Clearly, all the K nodes participate again in the next stage, which, therefore, also uses an N-level single-stage optimal timer.

In the case of a collision we can show the following along the lines of Observation 1. If a node i collided in the first collision slot s, $0 \leq s \leq N$, then $\mu_i \sim U\left[1-\sum_{j=0}^{s}\alpha_N[j,K]\,, 1-\sum_{j=0}^{s-1}\alpha_N[j,K]\right)$. Hence, the following rescaling

$$R_{C(s)}^{(N,K)}(\mu_i) = \frac{\mu_i - 1 + \sum_{j=0}^{s} \alpha_N[j,K]}{\alpha_N[s,K]}, \tag{6}$$

ensures that $R_{C(s)}^{(N,K)}(\mu_i) \sim U[0,1)$. As before, the rescalings in (5) and (6) preserve order and independence.

However, in the case of a collision, no node knows how many nodes transmitted in the first collision slot. Therefore, it is not clear at first sight how many nodes to design the next timer stage for and how many levels to use for it. The following observation proves invaluable in this regard.

In an N-level single-stage optimal timer, designed for K nodes, let X_s denote the number of nodes that set their timer values to $s\Delta$, for $0 \leq s \leq N$. Let Y_s denote the number of nodes that set their timers to $s\Delta$ given that $s\Delta$ is the first transmission slot in the event of a collision. Mathematically,

$$Y_s = [X_s | X_s \geq 2, X_{s-1} = \cdots = X_0 = 0]\,, \text{ for } s = 0, 1, \ldots, N.$$

The expected value of Y_s, $\Gamma_N^K(s)$, is plotted in Figure 2 as a function of K and s.

Observation *3:* In an N-level single-stage optimal timer the average number of nodes, $\Gamma_N^K(s)$, that set their timers to

$s\Delta$, given that at least two timers have expired at $s\Delta$ and no timer has expired before $s\Delta$, lies in a narrow range:

$$2 \leq \Gamma_N^K(s) \leq \frac{e-1}{e-2} < 2.40, \quad \text{for } s = 0, 1, \ldots, N. \qquad (7)$$

Explanation: The lower bound on $\Gamma_N^K(s)$ is easy to see since at least 2 nodes must be involved in a collision. Thus, $\Gamma_N^K(s) \geq 2$. Secondly, from Figure 2, we observe that $\Gamma_N^K(s)$ increases with s, for any N and K, and with K, for any N and $0 \leq s \leq N$. Hence, $\Gamma_N^K(s) \leq \Gamma_N^K(N) \leq \lim_{K\to\infty} \Gamma_N^K(N)$. Using the asymptotic analysis of the single-stage optimal timer scheme [8], it can be shown that $\lim_{K\to\infty} \Gamma_N^K(N) = \frac{e-1}{e-2} < 2.40$.

Thus, Observation 3 implies that, in the case of a collision, the number of collided nodes on an average will be very near to 2. Therefore, one can design all the subsequent timer stages for 2 nodes.

Therefore, the scheme for the general case of selecting among K nodes is similar to that described in Section III-A for 2 nodes, except for the following important differences.

- *Start:* An N_I-level single-stage optimal timer scheme that is designed for K nodes is run.
- *Collision:* Let a node i be involved in the first collision that occurs, say, in slot s. If this is the first time a collision outcome has occurred, the node rescales its metric to $R_{C(s)}^{(N_I,K)}(\mu_i)$. For subsequent collision outcomes, it rescales its metric to $R_{C(s)}^{(N_C,2)}(\mu_i)$. An N_C-level single-stage optimal timer that is designed for 2 nodes is then run in the next stage. All nodes that did not transmit in slot s do not participate in subsequent stages.
- *Idle:* All the nodes that participated in the current stage, participate again, by rescaling their metrics to $R_I^{(N_C,2)}(\mu_i)$, if a collision outcome has occurred previously, or to $R_I^{(N_I,K)}(\mu_i)$, otherwise. In either case the same single-stage optimal timer is used in the next stage.

As shown in Appendix B, the average selection time of the proposed scheme, $T_{\text{sel}}(K)$, is approximately

$$T_{\text{sel}}(K) \approx \frac{(N_I+1)\Delta + F + P_c(N_I,K)(N_C+2)\left(\Delta + \frac{F}{N_C+1}\right)}{1 - P_I(N_I,K)}, \qquad (8)$$

where $P_I(N_I,K)$ and $P_C(N_I,K)$ denote the probabilities of an idle and a collision of an N_I-level optimal timer (designed for K nodes), respectively. They are given by $P_I(N_I,K) = \left(1 - \sum_{j=0}^{N_I} \alpha_{N_I}[j,K]\right)^K$ and $P_C(N_I,K) = \sum_{r=2}^{K} \sum_{j=0}^{N_I} \binom{k}{r} \alpha_{N_I}[j,K]^r \left(1 - \sum_{l=0}^{j} \alpha_{N_I}[l,K]\right)^{K-r}$. The following result provides an optimal choice for N_C, which we denote by N_C^{opt}.

Lemma *1:* The optimal value of N_C, which minimizes $T_{\text{sel}}(K)$ in (8) is given by, $N_C^{\text{opt}} = \sqrt{\frac{F}{\Delta}} - 1$ if $\sqrt{\frac{F}{\Delta}} - 1$ is an integer, else, $N_C^{\text{opt}} = \left\lfloor \sqrt{F/\Delta} - 1 \right\rfloor$ or $N_C^{\text{opt}} = \left\lceil \sqrt{F/\Delta} - 1 \right\rceil$. Furthermore, for $K = 2$, $N_I^{\text{opt}} = N_C^{\text{opt}}$.

Proof: The proof is relegated to Appendix C. ■

Thus, for $K > 2$, only N_I^{opt} needs to be determined numerically. Notice that N_C^{opt} increases as the feedback overhead increases. Thus, for larger $\frac{F}{\Delta}$, more time gets allocated to each stage to improve its probability of success.

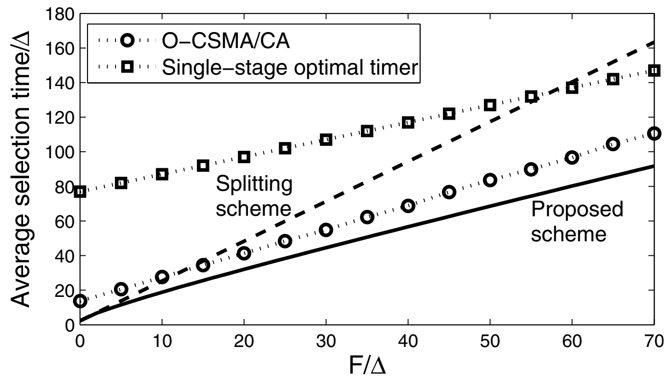

Fig. 3. Comparison of the average selection time vs. normalized feedback overhead ($K = 5$). Success probability of the single-stage optimal timer is set to 98%.

IV. Simulations and Results

We now evaluate the performance of the proposed scheme using Monte Carlo simulations that uses $30,000$ samples.

Figure 3 compares the average selection time of the proposed scheme with that of the splitting scheme as a function of the normalized feedback overhead, $\frac{F}{\Delta}$, for $K = 5$.[3] Also shown is the average time required to select by the O-CSMA/CA [10] and the single-stage optimal timer [8]. Since the single-stage timer scheme cannot guarantee success, its parameters are chosen to ensure a success probability of 98%. In O-CSMA/CA, a timer mapping with M levels and equal step lengths is used in the first stage. In case of a failure, the number of levels is doubled, till it reaches $M_{\max}$, and the collided nodes participate in the next stage with probability $\frac{1}{2}$. Even O-CSMA/CA does not guarantee selection of the best node. We use $M_{\max} = 2^7 M$ and $M = 7$ [10], [13]. Its success probability turns out to be 81%.

We observe that the proposed scheme outperforms all these schemes and ensures a 100% success probability. For example, for $\frac{F}{\Delta} = 20$ and $K = 5$, the splitting scheme, O-CSMA/CA, and the single-stage optimal timer take 50%, 28%, and 200% more time, respectively, than the proposed scheme to select.[4] Moreover, the gains increase as $\frac{F}{\Delta}$ increases. For larger K, e.g., 100, the performance gap between the proposed scheme and the benchmark schemes increases even further. The corresponding plot is not shown due to space constraints.

In order to understand the proposed scheme better, Figure 4 plots N_I^{opt} and N_C^{opt} as a function of $\frac{F}{\Delta}$. We see that N_I^{opt} and N_C^{opt} increase with $\frac{F}{\Delta}$. Thus, as the feedback overhead

[3] For the proposed scheme, for all F/Δ less than 60, the feedback requires $\left\lceil \log_2(\max(N_I^{\text{opt}}, N_C^{\text{opt}}) + 4) \right\rceil \leq 4$ bits, which is marginally greater than the 2 bit feedback of the splitting scheme. This causes no increase in F in a WLAN type network, for example, since it is much less than the minimum payload. Therefore, we use the same F for all the schemes.

[4] The results for the single-stage inverse timer scheme of [2] are not shown since it is considerably slower than the single-stage optimal timer scheme for the parameters chosen.

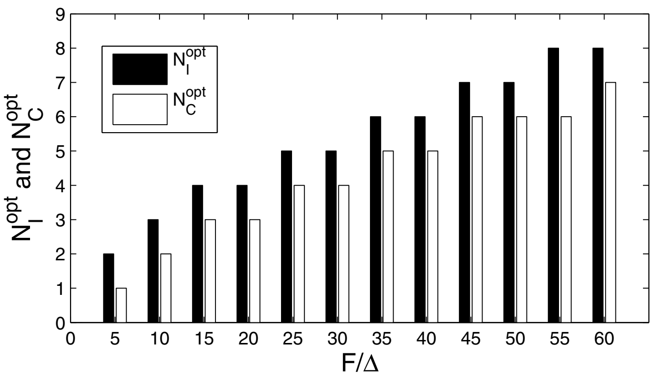

Fig. 4. Plot of optimized N_I and N_C as a function of normalized feedback overhead, $\frac{F}{\Delta}$, for the proposed scheme ($K = 5$ nodes).

increases, the duration of each timer stage increases in order to increase the probability of success in each stage. Further, $N_I^{\text{opt}} \geq N_C^{\text{opt}}$. Finally, only when $F = 0$, $N_I^{\text{opt}} = N_C^{\text{opt}} = 0$, which is nothing but the splitting scheme.

V. Conclusions

We proposed a novel distributed selection scheme that inherits the best features of timer and splitting schemes and guarantees best node selection. The scheme runs in stages with the timer scheme used in every stage. The duration of each stage takes one of two values, $N_I^{\text{opt}}\Delta$ or $N_C^{\text{opt}}\Delta$, and depends on whether a collision has occurred or not. These durations tend to be larger, by design, for a system with larger feedback overheads so as to improve the success probability of each stage and reduce the need for subsequent feedback.

We saw that the proposed scheme is much faster than the splitting scheme, single-stage optimal timer and inverse timer schemes, and the O-CSMA/CA, which doubles the window duration in case of a collision. Interestingly, in our scheme, the window size actually shrinks once a collision occurs.

Appendix

A. *Brief Proof of Theorem 1*

For the N-level single-stage optimal timer that is designed for 2 nodes the success probability can be shown to equal

$$P_s\left(N,2\right) = \frac{N+1}{N+2}. \tag{9}$$

The first timer stage and the feedback that follows it require a total time of $(N+1)\Delta + F$. In case of a success outcome, the scheme terminates and requires no more time. In case of an idle, it can be seen from the rescaling in (3) that the scheme requires an additional time of $T_2(N)$ to select, on average. Similarly, in case of a collision also, the rescaling in (2) implies that, on average, $T_2(N)$ more time is required to select.

Let $P_s\left(N,2\right)$, $P_I\left(N,2\right)$, and $P_C\left(N,2\right)$ be the probabilities of success, idle, and collision, respectively. Then,

$$T_2(N) = (N+1)\Delta + F + \left(P_C\left(N,2\right) + P_I\left(N,2\right)\right) T_2(N). \tag{10}$$

However, $P_C\left(N,2\right) + P_I\left(N,2\right) = 1 - P_s\left(N,2\right) = 1 - \frac{N+1}{N+2}$, from (9). Substituting this in (10) results in (4).

B. *Brief Derivation of* (8)

The first stage and its feedback will require a time of $(N_I+1)\Delta + F$. The scheme terminates after the first stage in case of a success, which occurs with probability $1 - P_I\left(N_I,K\right) - P_C\left(N_I,K\right)$.

From Observation 3, we know that once a collision occurs, the number of nodes that have collided is very likely two. Therefore, in case of a collision, an additional time of approximately $T_2(N_C)$ is required to select the best node, on average. Here, $T_2(N_C)$ is given by Theorem 1. Similarly, in case of an idle, an additional time of $T_{\text{sel}}(K)$ is required to select the best node, on average. Therefore, $T_{\text{sel}}(K) \approx (N_I+1)\Delta + F + P_C\left(N_I,K\right) T_2(N_C) + P_I\left(N_I,K\right) T_{\text{sel}}(K)$. Substituting (4) in the above equation gives the desired result.

C. *Proof of Lemma 1*

In (8), we see that $T_{\text{sel}}(K)$ is an affine function of $T_2(N_C)$ for any given N_I. Thus, the value of N_C that minimizes $T_{\text{sel}}(K)$, for a given N_I, is the one that minimizes $T_2(N_C)$.

From Theorem 1, it is easy to see that $T_2(N_C)$ is a convex function in N_C, if $N_C \in \mathbb{R}$. It, thus, has a unique minimum, which can be shown to occur at $\sqrt{\frac{F}{\Delta}} - 1$ using the first order condition. Upon restricting N_C to the set of natural numbers, the desired result in Lemma 1 follows. Furthermore, for $K = 2$, $N_I = N_C$; thus, $N_I^{\text{opt}} = N_C^{\text{opt}}$.

References

[1] A. J. Goldsmith, *Wireless Communications*. Cambridge Univ. Press, 2005.

[2] 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.

[3] Y. Zhao, R. Adve, and T. J. Lim, "Improving amplify-and-forward relay networks: optimal power allocation versus selection," in *Proc. ISIT*, pp. 1234–1238, Jul. 2006.

[4] 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.

[5] Q. Zhao and L. Tong, "Opportunistic carrier sensing for energy-efficient information retrieval in sensor networks," *EURASIP J. Wireless Commun. Netw.*, vol. 2005, pp. 231–241, May 2005.

[6] X. Qin and R. Berry, "Opportunistic splitting algorithms for wireless networks," in *Proc. INFOCOM*, vol. 3, pp. 1662–1672, Mar. 2004.

[7] V. Shah, N. B. Mehta, and R. Yim, "Splitting algorithms for fast relay selection: generalizations, analysis, and a unified view," *IEEE Trans. Wireless Commun.*, vol. 9, pp. 1525–1535, Apr. 2010.

[8] V. Shah, N. B. Mehta, and R. Yim, "Optimal timer based selection schemes," *IEEE Trans. Commun.*, vol. 58, pp. 1814–1823, Jun. 2010.

[9] 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.

[10] C.-S. Hwang and J. M. Cioffi, "Using opportunistic CSMA/CA to achieve multi-user diversity in wireless LAN," in *Proc. Globecom*, pp. 4952–4956, Nov. 2007.

[11] C.-S. Hwang, K. Seong, and J. M. Cioffi, "Opportunistic p-persistent CSMA in wireless networks," in *Proc. ICC*, pp. 183–188, Jun. 2006.

[12] T. Tang and R. Heath, "Opportunistic feedback for downlink multiuser diversity," *IEEE Commun. Lett.*, vol. 9, pp. 948–950, Oct. 2005.

[13] "Part 11: Wireless LAN medium access control (MAC) and physical layer (PHY) specifications," Tech. Rep. IEEE Std 802.11-2007, IEEE Computer Society, Jun. 2007.