Relay-Assisted Scheduling in Wireless Networks with Hybrid-ARQ

Caleb K. Lo[1], John J. Hasenbein[2], Sriram Vishwanath[1] and Robert W. Heath, Jr.[1]

[1]Wireless Networking and Communications Group

Department of Electrical and Computer Engineering

The University of Texas at Austin

1 University Station C0803

Austin, TX 78712-0240

Phone: (512) 471-1190

Fax: (512) 471-6512

Email: {clo, sriram, rheath}@ece.utexas.edu

[2]Graduate Program in Operations Research and Industrial Engineering

Department of Mechanical Engineering

The University of Texas at Austin

Austin, TX 78712-1063

E-mail: jhas@mail.utexas.edu

Abstract

This paper studies the problem of relay-assisted user scheduling for downlink wireless transmission. The transmitting base station or access point employs hybrid automatic-repeat-request (ARQ) transmission with the possible assistance of a set of fixed relays to serve a set of mobile users. By minimizing a cost function that depends on the queue lengths at the base station and the number of retransmissions of the head-of-line packet for each user, the base station can determine an appropriate set of users to be serviced in each time slot. In the system model of interest, though, minimizing the cost function at the base station may entail selecting a set of relays to service the scheduled user set. Three cost functions are considered, including a linear function with packets arriving according to a Poisson process, a discounted version of this function with rewards, and an increasing convex function with packets that need to be drained from the queues at the base station. It is shown that a priority-index policy is optimal in all cases. Performance comparisons illustrate the benefits of relay-assisted scheduling with hybrid-ARQ transmission.

Keywords - Relays, scheduling policies, hybrid automatic-repeat-request, priority-index rules.

1 Introduction

Relay-assisted communication will yield increased system throughput and extended coverage in both local area and metropolitan area networks [1, 2]. One type of relay-assisted communication is a two-hop system, where a source transmits to one or more relays that process the source transmission and forward the processed output to its intended destination. Two-hop communication, which relies on half-duplex signaling, is especially applicable to cellular networks, as well-placed relays can extend coverage to mobile users that do not have strong links to the base station. Another type of relay-assisted communication involves direct transmission between a source and a destination along with one or more relays that can collaborate to assist the source transmission. The key aspect of this system is that the relays assist the source only if the destination cannot decode the direct transmission.

The benefits of relay-assisted communication are reduced by the fact that transmissions occur over inherently lossy wireless links. In particular, random signal fading severely decreases the received signal quality for each destination node. To partially offset fading, the relays and the destination can all employ intelligent reception strategies, including hybrid-ARQ strategies such as Chase combining and parity forwarding based on incremental redundancy [3]. Prior work on this topic has considered placing relays between the source and the destination who can forward parity information to the destination in the event of uncorrectable decoding errors, which yields significant gains in terms of throughput [3–8]. Allowing the relays to forward reliable parity information decreases the number of retransmissions that are needed for successful decoding, which decreases the delay observed by all destination nodes in the network. Note that higher throughput values are generally obtained if the relays forward reliable parity information, as opposed to having each relay retransmit the entire packet from the source.

Since relay-based wireless transmission improves performance via intelligent reception strategies, we now focus on the relaying-enabled performance gains in wireless networks. For downlink transmission, this raises the key question of how the presence of fixed relays impacts the user scheduling decisions at the base station or access point. As described in [12], one of the key benefits of relay-assisted scheduling in a cellular network is that multiple relays can simultaneously transmit due to intelligent frequency reuse planning in the base station. Another key benefit of considering relaying in a cellular network is that synchronizing the transmissions from multiple relays is relatively straightforward and controlled by the base station, which avoids the difficulty of synchronizing transmissions that is inherent to ad hoc networks [13–15]. Relay-assisted scheduling has been studied in cellular networks [9, 12] and also for more general networks [21–25]. In particular, a throughput-optimal multi-link activation policy based on backpressure principles was applied to cellular networks [12].

We reiterate that the benefits of relay-assisted transmission are more fully realized when intelligent reception strategies such as hybrid-ARQ are employed in the network of interest. The importance of hybrid-ARQ is evidenced by the support for Chase combining and incremental redundancy in the IEEE 802.16 standard [2]. Prior work on relay-assisted scheduling, though, does not account for hybrid-ARQ. Relays can be employed to decrease the number of hybrid-ARQ transmissions that are required to serve a particular user. Given a hybrid-ARQ transmission framework, the relay-assisted scheduling problem is not merely a function of user queue lengths at the base station and the relays [12]. It is now important to also consider the number of hybrid-ARQ transmissions that have occurred for each mobile user's head-of-line (HoL) packet, which directly impacts the decoding delay that each mobile user incurs. Decoding delay also depends on queue lengths via Little's theorem [19] and has influenced a significant amount of work in the scheduling domain [10, 11, 18].

In this paper we derive cost-minimizing user scheduling policies for relay-assisted downlink transmission under a hybrid-ARQ transmission framework. We apply the approach in [16] to our relay-assisted system model and consider a generic hybrid-ARQ transmission strategy. In [16], packets for several mobile users arrive at the base station, and in each time slot one of the users is scheduled to transmit. Each user has a cost function at the base station that depends on both its queue length at the base station and the number of retransmissions that have been used for its HoL packet. The objective is to schedule a user to minimize the long-term average expected cost at the base station. It is shown in [16] that the optimal scheduler is a fixed priority-index policy, where the users are ranked and the highest-ranked user with a nonempty queue is serviced in that time slot.

For our relay-assisted system model, if a user cannot decode its HoL packet, the base station can possibly select a set of decoding relays to assist that user. In particular, a set of mobile users can be simultaneously scheduled along with a set of relays to service them. It is now non-obvious as to whether a priority-index policy is optimal for our relay-assisted system model. Ranking users is more difficult than in [16], since the priority index for each user should now depend on the decoding status of each relay. If relays are not allowed to assist the base station, then the scheduling problem simplifies to the no-relay problem of [16].

To illustrate the performance impact of relaying in hybrid-ARQ scheduling, we consider relay-assisted variants of the two key problems considered in [16]. One problem entails minimizing a linear cost function with Poisson arrivals at the base station (LPA), while the other problem entails minimizing a convex increasing function of queue length without any new arrivals at the base station (DC). We prove that the optimal scheduler for the relay-assisted variants of the LPA and DC problems is actually a priority-index policy as in [16], which is intuitive and amenable to practical implementation. We also formulate a relay-assisted variant of the DC problem as a Markov decision process to gain additional insights on computation of the optimal priority-index policy.

In addition, we consider a discounted version of the relay-assisted variant of the LPA problem, where the wireless network operator obtains a reward in each time slot where a user decodes its HoL packet. The cost function for each user is multiplied by a discount factor that monotonically decreases with time. Note that discount factors are typically used to yield finite costs over an infinite time horizon [26]. We apply the main result from [17] to this problem and prove that a priority-index policy minimizes the discounted cost.

This paper is organized as follows. In Section II we present the relay-assisted scheduling problem in a downlink wireless system. We present the LPA problem from [16] in Section III and obtain the optimal scheduler for the relay-assisted extension of this problem. We then present the DC problem from [16] in Section IV and obtain the optimal scheduler for its relay-assisted extension. Next we present the discounted cost problem from [17] in Section V and obtain the optimal scheduler for its relay-assisted extension. Simulation results are presented in Section VI and we conclude the paper in Section VII.

2 System Model

First, we introduce the notation used throughout the paper. $|z|^2$ denotes the absolute square of a complex number z. $\mathbb{E}$ denotes the mathematical expectation operator. $\|\mathcal{A}\|$ denotes the cardinality of a set $\mathcal{A}$.

Consider the downlink transmission model in Fig. 1. The network consists of a single base station, M fixed relay stations, and N mobile users. Packets for each mobile user arrive at the base station, and each packet is placed in a queue for its intended mobile user. The packet arrival processes are mutually independent. Let $h_{i,n}$ denote the channel between nodes i and n.

In the first time slot, the base station selects a packet for one of the users and broadcasts this packet to the entire network. If the selected user successfully decodes the packet, the base station flushes the packet from its queue and prepares to select another packet for transmission during the second time slot. Each relay would also flush the packet from its queue in this case, assuming that it has decoded the packet after the first time slot.

2.1 Key Assumptions

We make the following critical assumptions in this paper:

- The base station, all of the relays, and all of the mobiles are synchronized.
- All channels are narrowband and cause transmitted signals to undergo flat fading.
- Before the first time slot, the base station knows its channel gain $|h_{t,i}|^2$ to each user $i \in \{1, 2, \ldots, N\}$.
- The base station also knows the channel gain $|h_{a,i}|^2$ from each relay a to each user i.
- Time is slotted, and each channel gain $|h_{i,n}|^2$ remains constant over a single hybrid-ARQ retransmission sequence, which consists of a finite number of slots. This assumption is reasonable in a slow fading environment.
- Each channel gain $|h_{i,n}|^2$ varies independently from a single hybrid-ARQ retransmission sequence to the next retransmission sequence. This is a block fading assumption.

2.2 User Scheduling Problem

Given the transmission model, it is apparent that if each selected user can decode its packet after it is initially transmitted, the relays do not need to play a role in the scheduling policy at the base station. The relays are only considered in the scheduling policy when a selected user makes a decoding failure and its packet needs to be retransmitted.

Each relay station can store one packet for each mobile user. The only packets that are stored at each relay are those that 1) the base station has transmitted, 2) the relay has decoded and 3) the mobile has not decoded. Thus, we do not need to consider packet arrival processes at each relay station.

Let $\mathbf{S}(n)$ be the state vector for the base station at time slot n, where $\mathbf{S}(n)$ is the same as in [16]. Thus, $\mathbf{S}(n)$ includes the number of transmission attempts for the current HoL packet for each user and the queue length for each user at the base station.

Let $\mathbf{S}_a(n)$ be the state vector for relay a at time slot n. Here, $\mathbf{S}_a(n) = \{R_{a,1}(n), R_{a,2}(n), \ldots, R_{a,N}(n)\}$. We have $R_{a,i}(n) = 1$ if a packet has been selected for user i, relay a has decoded it and user i has not decoded it. Otherwise, $R_{a,i}(n) = 0$.

Let $\mathcal{M} = \{BS, 1, 2, \ldots, M\}$ and $\mathcal{N} = \{1, 2, \ldots, N\}$ denote the set of allowed transmitters and mobile users in the network, respectively. In this paper, our objective is to design a scheduling policy $\pi(\cdot, \cdot, \cdots, \cdot) \in \Pi_R$ such that $\pi(\mathbf{S}(n), \mathbf{S}_1(n), \mathbf{S}_2(n), \cdots, \mathbf{S}_M(n)) = (Q, W, f(\cdot))$, where $Q \in \mathcal{N}$ is the set of scheduled users and $W \in \mathcal{M}$ is the set of scheduled transmitters. Note that $f(w) \in Q$ is a rule for each scheduled transmitter $w \in W$ that assigns a scheduled user $f(w)$ to w.

Note that $\mathbf{S}(n)$ and $\mathbf{S}_a(n)$ for $a \in \{1, 2, \ldots, M\}$ evolve with n according to $\pi(\cdot, \cdot, \cdots, \cdot)$. In particular, the number of transmissions of the HoL packet for user i evolves as

$$r_i^{HoL}(n+1) = \begin{cases} 0 & i \in Q, \text{ user } i \text{ decodes its HoL packet} \\ r_i^{HoL}(n) + 1 & i \in Q, \text{ user } i \text{ cannot decode its HoL packet} \\ r_i^{HoL}(n) & i \notin Q. \end{cases} \quad (1)$$

Also, let $A_i(n)$ be the number of packets for user i that arrive at the base station in time slot n. Then, the queue length for user i evolves as

$$x_i(n+1) = \begin{cases} x_i(n) + A_i(n) - 1 & i \in Q, \text{ user } i \text{ decodes its HoL packet} \\ x_i(n) + A_i(n) & \text{otherwise.} \end{cases} \quad (2)$$

In addition, the decoding status of the HoL packet for user i at relay a evolves as

$$R_i(n+1) = \begin{cases} 0 & i \in Q, a \in W, \text{ decoding success at user } i \text{ or} \\ & i \in Q, a \notin W, \text{ decoding failures at user } i \text{ and relay } a \\ 1 & i \in Q, a \notin W, \text{ decoding failure at user } i \text{ but success at relay } a \\ R_i(n) & i \in Q, a \in W, \text{ decoding failure at user } i \text{ or} \\ & i \notin Q. \end{cases} \quad (3)$$

3 Relay-Assisted Linear Poisson Arrivals Problem

In this section, we consider the relay-assisted linear Poisson arrivals (RLPA) problem, which is a variant of the LPA problem introduced in [16].

3.1 Single-Relay Selection

We initially consider a system where a single relay out of M relays can assist the base station in transmitting the packets for the N users. We provide a recap of the basic formulation of the LPA problem from [16].

Packets for each of the N users arrive at their corresponding queues at the base station. The arrival process for the packets of user i is Poisson with rate λ_i, where $i \in \{1, 2, \ldots, N\}$. Let c_{i,r_i} denote the cost of storing a packet for user i that has already been transmitted r_i times. The base station computes a cost function U_i for user i that is linear in the queue length $x_i(n)$, where

$$U_i(x_i(n), r_i^{HoL}(n)) = \begin{cases} c_{i,0}(x_i(n) - 1) + c_{i,r_i^{HoL}(n)} & x_i(n) > 0 \\ 0 & x_i(n) = 0. \end{cases} \tag{4}$$

In particular,

$$0 \leq c_{i,r_i} \leq c_{i,r_i'}, \quad r_i < r_i' \tag{5}$$

which implies that the storage cost for any packet is a nondecreasing function of the number of transmission attempts.

The LPA problem entails determining the scheduling policy $\pi \in \Pi$ that minimizes the long-run average expected cost J_{LPA}, where

$$J_{LPA} = \lim_{\tau \to \infty} \frac{1}{\tau} \mathbb{E}_\pi \left[\sum_{n=1}^{\tau} \sum_{i=1}^{N} U_i(x_i(n), r_i^{HoL}(n)) \right].$$

The RLPA problem entails determining the scheduling policy $\pi \in \Pi_R$ that minimizes the long-run average expected cost J_{RLPA}, where

$$J_{RLPA} = \lim_{\tau \to \infty} \frac{1}{\tau} \mathbb{E}_\pi \left[\sum_{n=1}^{\tau} \sum_{i=1}^{N} U_i(x_i(n), r_i^{HoL}(n)) \right].$$

It turns out that the optimal policy for the RLPA problem is based on the priority index policy that is optimal for the LPA problem [16]. For the RLPA problem, knowledge of $\mathbf{S}_R(n)$ at the base station is useful in deciding which users can be served more quickly than others. This is especially useful for the RLPA problem, since cost increases with the number of retransmission attempts or delay.

Theorem 1. *The optimal scheduling policy for the RLPA problem is a priority-index rule, where the HoL packet with the highest priority index over all nonempty base station queues is selected. The transmitter that yields the highest priority index transmits the selected HoL packet.*

Proof. The proof is in Appendix A. □

We provide an intuitive justification of Theorem 1 in the following. In [16], each user i is assigned a fixed priority index $c_{i,r_i^{HoL}}/T_{i,r_i^{HoL}}$, where $c_{i,r_i^{HoL}}$ is the holding-cost rate for the HoL packet of user i that has undergone r_i^{HoL} transmission attempts. Also, $T_{i,r_i^{HoL}}$ is the expected service time for the HoL packet of user i, where

$$T_{i,r_i^{HoL}} = 1 + \sum_{a=r_i^{HoL}}^{r_i^{max}-1} \prod_{l=r_i^{HoL}}^{a} g_i(l) \tag{6}$$

and $g_i(l)$ is the probability of a decoding failure by user i given that its HoL packet has been transmitted l times. The optimal policy from [16] is to select the HoL packet of the user with the highest priority index.

Since $T_{i,r_i^{HoL}}$ is a function of the decoding failure probability, a lower value of $T_{i,r_i^{HoL}}$ implies that user i achieves a higher priority index. Now consider the two-user system depicted in [16, Figure 4], where a new packet arrival for the first user has priority over a packet retransmission for the second user. No relays are present, which is equivalent to the RLPA problem if $\mathbf{S}_R(n) = (0, 0, \ldots, 0)$. Assume that only relay i has decoded the HoL packet for user 2, so relay i can decrease $T_{2,r_2^{HoL}}$. This increases user 2's priority, assuming that $|h_{r,2}|^2 > |h_{t,2}|^2$. Thus, if $|h_{r,2}|^2$ is above a threshold value, user 2 can actually end up with a higher priority index than user 1, and so a retransmission for user 2 would have priority over a new arrival for user 1. In particular, $\pi(\mathbf{S}(n), \mathbf{S}_R(n)) = (2, i)$ as opposed to $\pi(\mathbf{S}(n)) = 1$ for the LPA problem, where no relays are present.

Thus, the introduction of relays for the RLPA problem results in a slight modification to the optimal priority index policy in [16]. Users are still sorted according to their priority indices and the highest priority user with a nonempty queue is scheduled. In this case, though, each relay i can inform the base station of its ability to improve the priority indices of some subset of the users by reporting $\mathbf{S}_{R,i}(n)$ to the base. The base station can calculate an improved priority index for each user a such that $R_{i,a}(n) = 1$. Then, all priority indices including any revised indices are sorted, and the highest priority user with a nonempty queue is scheduled along with the transmitter that yields that highest priority index.

Note that the result in [16, Corollary 1] on the monotonicity of the optimal policy with respect to the number of transmission attempts can be extended to the RLPA problem, with a key caveat. Based on the discussion of [16, Figure 4], scenarios exist where a new packet arrival to an empty queue has priority over a retransmission of a previously scheduled HoL packet.

Corollary 1. *Assume that it is optimal to transmit to user i when $\boldsymbol{R} = (r_1^{HoL}, \ldots, r_i^{HoL}, \ldots, r_N^{HoL})$ and that for each empty queue $a \neq i$, no new packet arrivals occur after the HoL packet of user i has been scheduled. Then it is also optimal to transmit to user i if r_i^{HoL} increases and r_a^{HoL} is held fixed for $a \neq i$.*

Proof. The proof follows directly from the discussion in [16, Appendix] and [16, Section 3.B]. □

3.2 Multiple-Relay Selection

We now consider the RLPA problem where multiple relays can assist a scheduled user by simultaneously transmitting to it in a given time slot. In particular, let $g_{i,m}(l)$ denote the probability of decoding failure at user i if multiple relays assist it for retransmission l, and let $g_{i,s}(l)$ denote the probability of decoding failure at user i if a single relay assists it for retransmission l. The multiple relays should yield $g_{i,m}(l) < g_{i,s}(l)$ using a generic multiple-transmitter hybrid-ARQ strategy. For example, each of the assisting relays can retransmit the data for the scheduled user, possibly using different modulation/code-rate pairs. This implies that the scheduled user can implement a strategy such as Chase combining to recover the transmissions from the assisting relays.

If each assisting relay retransmits the data for the scheduled user using the same modulation/code-rate pair, the assisting relays must employ a cooperative transmission approach. For example, in a relay-assisted cellular network, the base station can synchronize the assisting relays so that their transmissions will add coherently in both time and phase at the scheduled user. On the other hand, if the assisting relays transmit distinct information, the scheduled user must employ multiuser detection to recover those transmissions.

It turns out that the optimal policy for the RLPA problem with multiple-relay transmission is similar to that for the RLPA problem with single-relay transmission. The key is to define the priority indices associated with having multiple nodes simultaneously transmit to a user i, which depends on their respective channel gains to user i.

Corollary 2. *The optimal scheduling policy for the RLPA problem with multiple-relay transmission is a priority-index rule, where the HoL packet with the highest priority index over all nonempty base station queues is selected. The set of transmitters that yields the highest priority index transmits the selected HoL packet.*

Proof. This follows in a straightforward manner from the analysis in Appendix A. □

Thus, allowing multiple relays to simultaneously assist a user in the RLPA problem does not change the essence of the optimal scheduling policy. The necessity of considering all possible sets of transmitters, though, increases the complexity of determining the optimal scheduled user.

4 Relay-Assisted Draining Convex Problem

In this section, we consider the relay-assisted draining convex (RDC) problem, which is a variant of the DC problem introduced in [16].

4.1 Single-Relay Selection

As in Section 3.1, we initially consider a system where a single relay out of M relays can assist the base station in transmitting the packets for the N users. We provide a recap of the basic formulation of the DC problem from [16].

The DC problem is a draining problem where no new packets arrive at the base station, so $A_i(n) = 0$ for all users i and time slots n. In the DC problem, the base station wants to empty all of the user queues. The base station computes a cost function U_i for user i, where U_i is an arbitrary increasing function of the queue length $x_i(n)$ and is independent of the number of transmission attempts of the HoL packet of user i. Thus,

$$U_i(x_i(n), r_i^{HoL}(n)) = U_i(x_i(n)).$$

The base station initially has a set of packets $(x_1(1), x_2(1), \ldots, x_N(1))$ for all N users. The DC problem entails determining the scheduling policy $\pi \in \Pi$ that minimizes the total expected draining cost J_{DC}, where

$$J_{DC} = \mathbb{E}_\pi \left[\sum_{n=1}^{\tau} \sum_{i=1}^{N} U_i(x_i(n)) \right].$$

The RDC problem entails determining the scheduling policy $\pi \in \Pi_R$ that minimizes the total expected draining cost J_{RDC}, where

$$J_{RDC} = \mathbb{E}_\pi \left[\sum_{n=1}^{\tau} \sum_{i=1}^{N} U_i(x_i(n)) \right].$$

As in Section 3.1, it turns out that the optimal policy for the RDC problem is based on the priority index policy that is optimal for the DC problem [16]. For the RDC problem, knowledge of $\mathbf{S}_R(n)$ at the base station is useful in deciding which users can be served more quickly than others. This is especially useful for the RDC problem, since cost increases with the delay.

Theorem 2. *The optimal scheduling policy for the RDC problem is a priority-index rule, where the HoL packet with the highest priority index over all nonempty base station queues is selected. The transmitter that yields the highest priority index transmits the selected HoL packet.*

Proof. The proof is in Appendix B. □

The intuitive justification of Theorem 2 is similar to that of Theorem 1. Again, each relay i can inform the base station of its ability to improve the priority indices of some subset of the users by reporting $\mathbf{S}_{R,i}(n)$ to the base. The base station can calculate an improved priority index for each user a such that $R_{i,a}(n) = 1$. Then, all priority indices including any revised indices are sorted, and the highest priority user with a nonempty queue is scheduled along with the transmitter that yields that highest priority index.

4.2 Multiple-Relay Selection

We now consider the RDC problem where multiple relays can assist the base station by simultaneously transmitting the packet for the scheduled user in each time slot. As in Section 3.2, the scheduled user can implement a strategy such as Chase combining to recover the transmissions from the assisting relays.

As in Section 3.2, the optimal policy for the RDC problem with multiple-relay transmission is similar to that for the RDC problem with single-relay transmission. The key is to define the priority indices associated with having multiple nodes simultaneously transmit to a user i, which depends on their respective channel gains to user i.

Corollary 3. *The optimal scheduling policy for the RDC problem with multiple-relay transmission is a priority-index rule, where the HoL packet with the highest priority index over all nonempty base station queues is selected. The set of transmitters that yields the highest priority index transmits the selected HoL packet.*

Proof. This follows in a straightforward manner from the analysis in Appendix B. □

As in the RLPA problem, allowing multiple relays to simultaneously assist a user in the RDC problem does not change the essence of the optimal scheduling policy. Again, the necessity of considering all possible sets of transmitters increases the complexity of determining the optimal scheduled user

4.3 Optimal Policy Computation

As noted in Appendix B, closed-form expressions cannot be found for the optimal priority indices for the RDC problem. To facilitate the computation of the optimal policy, we adopt the approach of [16] and formulate the RDC problem as a Markov decision process. For the purposes of simplicity, we consider a two-user system with a single assisting relay. In this case, the system state space is $S = \{(r_1, r_2, x_1, x_2, a_1, a_2) : 0 \leq r_i \leq r_i^{max}, 0 \leq x_i \leq A_i, a_i \in \{0, 1\}, i \in \{1, 2\}\}$, where $a_i = 1$ if the relay has decoded the HoL packet for user i and user i has not decoded it, and $a_i = 0$ otherwise.

The action space is $V = \{v_0, v_{10}, v_{11}, v_{20}, v_{21}\}$, where v_0 occurs if both the base station and the relay are idle, while v_{i0} occurs if the base station serves user i and v_{i1} occurs if the relay serves user i.

In Appendix C we utilize Bellman's equation [26] to obtain necessary conditions for the optimal cost function J_{RDC}. These conditions motivate the following result.

Proposition 1. *Let $J_{DC}(r_1, r_2, x_1, x_2)$ denote the optimal cost when starting from the initial state (r_1, r_2, x_1, x_2) in the DC problem from [16]. Then, for $i \in \{1, 2\}$, $r_i \in \{0, 1, \ldots, r_i^{max}\}$, $x_i \in \{0, 1, \ldots, A_i\}$ and $a_i \in \{0, 1\}$,*

$$J_{RDC}(r_1, r_2, x_1, x_2, a_1, a_2) \quad \leq \quad J_{DC}(r_1, r_2, x_1, x_2).$$

Proof. From [16, Section 5.C], it is clear that the action space V_{DC} for the DC problem is a strict subset of the action space V for the RDC problem. The result follows immediately from this fact. □

Note that if the base station-to-relay and relay-to-user channel gains are stronger than the base station-to-user channel gains, the optimal cost for the RDC problem should be strictly less than that of the DC problem. This simple result motivates our simulations of a single-relay, two-user instance of the RDC problem.

5 Relay-Assisted Discounted Costs and Rewards Problem

In this section, we consider the relay-assisted variant of the discounted cost problem from [17], which we term the DCR problem.

5.1 Single-Relay Selection

The basic formulation of the discounted cost problem from [17] is the same as that of the LPA problem in [16], with two key differences. First, a reward is obtained in each time slot where a user successfully decodes its HoL packet. Second, the cost function for each user is multiplied by a discount factor that monotonically decreases with time.

As in Section 3.1, the arrival process for the packets of user i at the base station is Poisson with rate λ_i, where $i \in \{1, 2, \ldots, N\}$. Let a_i denote the reward obtained by servicing user i. The base station computes a cost function U_i for user i that is linear in the queue length $x_i(n)$, where

$$U_i(x_i(n), r_i^{HoL}(n)) = \begin{cases} -a_i(x_i(n) - A_i(n-1) - x_i(n-1)) + c_{i,0}(x_i(n) - 1) + c_{i,r_i^{HoL}(n)} & x_i(n) > 0 \\ 0 & x_i(n) = 0. \end{cases} \quad (7)$$

Let β denote the discount factor for each user i. The DCR problem entails determining the scheduling policy $\pi \in \Pi$ that minimizes the long-run average expected discounted cost J_{DCR}, where

$$J_{DCR} = \lim_{\tau \to \infty} \frac{1}{\tau} \mathbb{E}_\pi \left[\sum_{n=1}^{\tau} \sum_{i=1}^{N} \beta^n U_i(x_i(n), r_i^{HoL}(n)) \right].$$

The RDCR problem entails determining the scheduling policy $\pi \in \Pi_R$ that minimizes the long-run average expected cost J_{RDCR}, where

$$J_{RDCR} = \lim_{\tau \to \infty} \frac{1}{\tau} \mathbb{E}_\pi \left[\sum_{n=1}^{\tau} \sum_{i=1}^{N} \beta^n U_i(x_i(n), r_i^{HoL}(n)) \right].$$

It turns out that the optimal policy for the RDCR problem is based on the priority index policy that is optimal for the DCR problem [17]. As in the RLPA problem, knowledge of $\mathbf{S}_R(n)$ at the base station is useful in deciding which users can be served more quickly than others.

Theorem 3. *The optimal scheduling policy for the RDCR problem is a priority-index rule, where the HoL packet with the highest priority index over all nonempty base station queues is selected. The transmitter that yields the highest priority index transmits the selected HoL packet.*

Proof. The proof follows by applying the transformation of the RLPA problem to the RLPAK problem in Appendix A along with the analysis in [17, Section 1]. In particular, the RDCR problem is transformed into a Klimov-type RDCRK problem via an appropriate state space expansion as in Appendix A.

The analysis in [17, Section 1] states that the optimal priority-index rule for the DCR problem is also optimal for the DCR problem where the storage cost for each queue is set to zero and all rewards are modified as in [17, (1)]. In particular, using the notation from Appendix A, the cost of storing a packet in queue (i, r_i, a) is $c_{i,r_i,a} = 0$. All queue transition probabilities are the same as in Appendix A. Also, the reward for moving from queue (i, r_i, a) to queue (k, r_k, m) is

$$a_{(i,r_i,a),(k,r_k,m)} = \begin{cases} a_i & i = k, r_k = 0 \text{ and } m = 0 \\ 0 & \text{otherwise.} \end{cases} \tag{8}$$

The rest of the proof follows in a straightforward manner from Appendix A. □

As in the RLPA and RDC problems, the value of the relays for the RDCR problem lies in their ability to modify the queue transition probabilities. This results in a lower expected service time for any user that is serviced by a relay. Note that the optimal priority-index policy can be computed via the algorithm in [17, Section 3].

5.2 Multiple-Relay Selection

We now consider the RDCR problem where multiple relays can assist the base station by simultaneously transmitting the packet for the scheduled user in each time slot. As in Section 3.2, the scheduled user can implement a strategy such as Chase combining to recover the transmissions from the assisting relays.

As in Section 3.2, the optimal policy for the RDCR problem with multiple-relay transmission is similar to that for the RDCR problem with single-relay transmission. Again, the key is to define the priority indices associated with having multiple nodes simultaneously transmit to a user i, which depends on their respective channel gains to user i.

Corollary 4. *The optimal scheduling policy for the RDCR problem with multiple-relay transmission is a priority-index rule, where the HoL packet with the highest priority index over all nonempty base station queues is selected. The set of transmitters that yields the highest priority index transmits the selected HoL packet.*

Proof. This follows in a straightforward manner from the analysis in Appendix A. □

As in the RLPA problem, allowing multiple relays to simultaneously assist a user does not change the essence of the optimal scheduling policy. In this case, determining which set of relays should transmit in each time slot will depend on the reward a_i for each user i and the discount factor β.

6 Simulation Results

In this section we evaluate the performance of relaying in the RLPA and RDC problems. First, we consider the RLPA problem.

Fig. 2 evaluates the performance impact of employing $M = 1$ relay in the RLPA problem. We consider a system with $N = 2$ users and set the arrival rates as $\lambda_1 = \lambda_2 = 0.3$. We also set the maximum number of retransmissions $r_1^{max} = r_2^{max} = 2$, and we set the cost rates $c_1 = [0.98\ 1\ 1.02]$. In addition, we set the channel parameters from the base station to users 1 and 2 as $\eta_1 = \eta_2 = 0.45$. The user probability of decoding failure is computed as

$$g_i(r_i) = \begin{cases} \eta_i \cdot 0.5^{r_i} & 0 \le r_i < r_i^{max} \\ 0 & r_i = r_i^{max}. \end{cases} \tag{9}$$

To gain an initial understanding of the performance impact of relaying in the RLPA problem, we perform three actions. First, we vary the channel parameter from the relay to user 2, as this determines the performance benefits of relaying for user 2. Second, we fix the channel parameter from the relay to user 1 as $\eta_{1,1} = 0.5$, as this prevents the relay from assisting user 1 and allows us to focus on how the relay can assist user 2. Third, as in [16], we set a constant storage cost $c_{2,r_i} = 1.5$ for user 2, as this diminishes the impact of the particulars of the storage cost function on our performance evaluation. The relay-to-user and base station-to-relay decoding probabilities follow the same rule as in (9). Setting $\eta_{1,2} = 0.5$ models the case where the relay is essentially absent, as a transmission from the relay yields no further information than a random coin flip.

We see that as the relay is placed in a more desirable location relative to user 2, the long-term cost is significantly reduced, which is primarily due to the relay's strong channel gains to both the base station and user 2. In particular, the cost decreases by about 28.8% from that with $\eta_{1,2} = 0.5$ to that with $\eta_{1,2} = 0.9$. This demonstrates the strong performance gains that are obtained by employing a relay, and especially an intelligently deployed relay.

Fig. 3 shows how the optimal policy for the RLPA problem behaves as a function of the relay channel gains to the users. We adopt many of the same parameters as in Fig. 2. We vary the relay channel parameters $\eta_{1,1}$ and $\eta_{1,2}$ to users 1 and 2, respectively and we set $\eta_{1,1} = \eta_{1,2}$.

We consider three values of the cost rate for user 2, namely $c_{2,r_2} \in \{1, 2, 3\}$ for each value of r_2. We see that the performance of the optimal policy deteriorates as the relay channel gains to the users decrease. In particular, if we fix the cost rate at $c_2 = 3$, we see that the long-term cost increases by about 29.2% from that with $\eta_{1,1} = 0.01$ to that with $\eta_{1,1} = 0.3$. We have also observed that the long-term cost continues to increase when simulating $\eta_{1,1}$ up to 0.5, where the relay is essentially absent. In addition, Fig. 3 shows that the long-term cost increases at roughly the same rate for each value of c_{2,r_2}. This demonstrates that 1) as in Fig. 2, strong performance gains are obtained by employing a relay and 2) the particulars of the cost function at hand are not as important as the intelligent placement of relays between the base station and the mobile users.

Fig. 4 shows how the optimal policy for the RLPA problem behaves as a function of the base station channel gains to the users. We adopt many of the same parameters as in Fig. 2. In this case we set $c_{2,r_2} = 1.5$ for each value of r_2, and we vary the channel parameter η_2 from the base station to user 2.

We set the channel parameters from the relay to users 1 and 2 as $\eta_{1,1} = \eta_{1,2} = 0.15$. We see that the performance of the optimal policy deteriorates as the base station channel gain to user 2 decreases. We also consider a case where no relay is present, and it can be seen that the performance of the optimal policy decreases at an even faster rate than the case where $M = 1$ relay assists the base station. This example further highlights the inherent challenges in a cellular network of servicing cell-edge users, and it is apparent that relay-assisted signaling is an important method for overcoming these difficulties.

We now consider the RDC problem. Fig. 5 compares the performance via employing $M = 1$ relay with that of employing no relays. We consider a system with $N = 2$ users and set the initial queue lengths as $(A_1, A_2) = (10, 10)$. We set the cost function for each user as $U_1(x) = U_2(x) = x^{1.1}$. In addition, we set the channel parameters from the base station to users 1 and 2 as $\eta_2 = \eta_1$ and allow for $(r_1^{max}, r_2^{max}) = (3, 3)$. The user probability of decoding failure is computed as

$$g_i(r_i) = \begin{cases} \eta_i^{r_i+1} & 0 \leq r_i < r_i^{max} \\ 0 & r_i = r_i^{max}. \end{cases} \tag{10}$$

We vary the base station-to-user decoding probabilities and observe that as the probability of decoding failure increases, the average cost per packet that has been in the system increases for both the no-relay and single-relay cases. We consider two single-relay sub-cases, namely $(\eta_{1,1}, \eta_{1,2}) = (0.01, 0.05)$ and $(\eta_{1,1}, \eta_{1,2}) = (0.001, 0.005)$. Interestingly, employing a single relay with strong channel gains to both the base station and the users does not yield a significant gain over not employing any relays until the base station-to-user channel gains decrease appreciably.

To further investigate this phenomenon, we consider Fig. 6 where we employ most of the same parameters as in Fig. 5. We again consider the single-relay sub-cases $(\eta_{1,1}, \eta_{1,2}) = (0.01, 0.05)$ and $(\eta_{1,1}, \eta_{1,2}) = (0.001, 0.005)$, and we compare their performance with that where no relays are employed. We apply a more stringent cost function $U_1(x) = U_2(x) = x^{2.1}$ for each user.

Again, we observe that employing a single relay with strong channel gains to both the base station and the users does not yield a significant gain over not employing any relays until the base station-to-user channel gains decrease appreciably. This is most likely due to our choice of the decoding probability function in (10). Since all of the decoding failure probabilities decrease rapidly, the cost advantage of employing a relay is mitigated. Thus, the specific hybrid-ARQ combining strategy that is used by all receiving nodes must be considered for both the RDC and the RLPA problems.

In this section, we do not perform any simulations for the RDCR problem. Note that when performing Monte Carlo simulations of discounted cost problems, the initial simulation steps have a greater impact on the output cost function value than the subsequent steps due to the discount factor. Thus, random values of J_{RDCR} and J_{DCR} are output for different simulation runs. Sophisticated strategies can be applied to reduce the variance of the output cost function values [27].

7 Conclusion

We have considered the problem of user scheduling in a downlink wireless system with hybrid-ARQ retransmissions. By allowing fixed relays to assist the base station or access point in servicing a set of scheduled users, a cost function that depends on the user queue lengths at the base station and the number of retransmissions of the HoL packet for each user can be minimized. We have studied the relay-assisted extensions of three problems presented in [16, 17] where either 1) packets arrive at the base station according to independent Poisson processes, 2) rewards are obtained when users decode their HoL packets given Poisson arrivals at the base station, or 3) packets are drained from their queues without any new arrivals. For all three problems, we prove that the optimal scheduler is a fixed priority-index rule.

Relay-assisted scheduling entails selecting a set of fixed relays to assist the base station, and so relay-assisted scheduling is actually a relay selection problem. It is apparent that relay selection is a difficult cross-layer problem, so a more comprehensive approach to the scheduling problem would consider factors such as the specific type of hybrid-ARQ being employed at the relays and the users along with more general arrival processes for packets at the base station. For example, if one user is downloading multimedia content while another user is sending text messages, this could be used to design an appropriate cost function for each user at the base station. By considering relay selection in the context of a cellular system, it is possible to minimize many of the key problems that arise in ad hoc networks, including timing synchronization between the base station and the selected relays.

A Proof of Theorem 1

We follow the same approach as in [16] for proving the optimality of our proposed scheduler for the RLPA problem. This entails transforming the RLPA problem into an instance of the multiclass queueing problem of Klimov [20]. Thus, the transformed problem, which we refer to as the RLPAK problem, has an optimal priority index policy. We then show that that the optimal policy for the RLPAK problem is optimal for the RLPA problem.

A.1 RLPAK Scheduling Problem

The RLPAK problem is similar to the LPAK problem in [16]. The key differences are as follows. For each user i, the M relays are sorted according to their channel gains to user i as $\{d_{i,1}, d_{i,2}, \ldots, d_{i,M}\}$, where

$|h_{d_{i,1}}|^2 < |h_{d_{i,2}}|^2 < \cdots < |h_{d_{i,M}}|^2$. Then, each user i has $(M+1)(r_i^{max}+1)$ queues, and each queue is labeled as (i, r_i, a). If a packet is in queue (i, r_i, a), it has been transmitted r_i times, relay $d_{i,a}$ has decoded it, and relay $d_{i,m}$ has not decoded it for $a < m \leq M$. In particular, a packet in the queue $(i, r_i, 0)$ has not been decoded by any of the M relays. Fig. 7 shows an example of the RLPAK problem for user 1 where $r_1^{max} = 2$.

Thus, there are a total of $K = \sum_{i=1}^{N}(M+1)(r_i^{max}+1)$ queues in the RLPAK problem. Each arriving packet is assigned to queue $(i, 0, 0)$ with probability $p_{i,0} = \lambda_i/\lambda$. Each queue (i, r_i, a) has a deterministic service time of $b_{i,r_i,a} = 1$ time slot.

The transition probabilities for the queues in the RLPAK problem are determined as follows. Let $g_{i,a,k}(r_i)$ denote the probability that relay $d_{i,a}$ cannot decode the HoL packet of user i after its transmission attempt r_i by relay d_k. In particular, $g_{i,a,0}(r_i)$ corresponds to a transmission by the base station. Also, let $g_{i,a}(r_i, 1)$ denote the probability that user i cannot decode its HoL packet after relay $d_{i,a}$ has transmitted it. Then

$$p_{(i,r_i,0),(i,r_i+1,0)} = g_i(r_i)g_{i,1,0}(r_i)g_{i,2,0}(r_i)\cdots g_{i,M,0}(r_i)$$

and

$$p_{(i,r_i,0),(i,r_i+1,a)} = g_i(r_i)(1 - g_{i,a}(r_i))g_{i,a+1}(r_i)g_{i,a+2}(r_i)\cdots g_{i,M}(r_i)$$

so the packet departs the system from queue $(i, r_i, 0)$ with probability $1 - g_i(r_i)$. In addition

$$p_{(i,r_i,a),(i,r_i+1,n)} = g_{i,a}(r_i, 1)(1 - g_{i,n,a}(r_i))g_{i,n+1,a}(r_i)g_{i,n+2,a}(r_i)\cdots g_{i,M,a}(r_i), n > a$$

$$p_{(i,r_i,a),(i,r_i+1,a)} = g_{i,a}(r_i, 1)g_{i,a+1,a}(r_i)g_{i,a+2,a}(r_i)\cdots g_{i,M,a}(r_i),$$

and

$$p_{(i,r_i,a),(i,r_i+1,n)} = 0, n < a$$

so the packet departs the system from queue (i, r_i, a) with probability $1 - g_{i,a}(r_i, 1)$. Note that after a packet from queue (i, r_i^{max}, a) has been served, it departs the system with probability 1, where $a \in \{0, 1, \ldots, M\}$.

Thus, for any set $M \subset \Omega = \{1, 2, \ldots, K\}$ and any queue $(i, r_i, a) \in M$, the average total service time is

$$T_{i,r_i,a}^{(M)} = 1 + \sum_{k,r_k,m} p_{(i,r_i,a),(k,r_k,m)} T_{k,r_k,m}^{(M)}.$$

The cost of storing a packet in queue (i, r_i, a) is $c_{i,r_i,a}$ and the number of packets in queue (i, r_i, a) at the beginning of the nth time slot is $x_{i,r_i,a}(n)$.

From the above discussion, it can be concluded that the RLPAK problem, which is a transformed version of the RLPA problem, is an instance of the multiclass queueing problem of [20]. This conclusion also relies on the simple queueing dynamics of the relay: 1) a packet only arrives at the relay if the base station has transmitted it, and 2) the relay automatically flushes a packet once it has been decoded by its intended user. In addition, the base station automatically flushes a packet once it has been decoded by its intended user. It should be noted that the state space of the RLPAK problem is an expanded version of that in the LPAK problem.

Now, the objective is to find a policy $\pi \in \Pi_R$ that minimizes the time-averaged expected cost J_{RLPAK}, where

$$J_{RLPAK} = \lim_{\tau \to \infty} \frac{1}{\tau} \mathbb{E}_\pi \left[\sum_{n=1}^{\tau} \sum_{(i,r_i,a)\in\Omega} c_{i,r_i,a} x_{i,r_i,a}(n) \right].$$

A.2 Optimal Policies for RLPAK and RLPA Scheduling Problems

The optimal policies for both the RLPAK and RLPA problems are priority-index rules. First, we state the following result.

Lemma 1. *Let M_k, $k = 1, 2, \ldots, K$ be the sets of queues generated by the Klimov algorithm in [16, Section 3] for the RLPAK problem. For each $k = 1, 2, \ldots, K$ and for all $(i, r_i, a) \in M_k$, the following are true:*

1) $(i, r_i', m) \in M_k$ for all $r_i' > r_i$ and for all $m \geq a$.

2) $T^{(M_k)}_{i,r_i,0} = 1 + \sum_{a=r_i}^{r_i^{max}-1} \prod_{l=r_i}^{a} g_i(l) = T^{(\Omega)}_{i,r_i,0}$.

3) $T^{(M_k)}_{i,r_i,m} = 1 + \sum_{a=r_i}^{r_i^{max}-1} \prod_{l=r_i}^{a} g_{i,m}(l,1) = T^{(\Omega)}_{i,r_i,m}, m > 0$.

4) $T^{(M_k)}_{i,r_i',m} \leq T^{(M_k)}_{i,r_i,a}$ for all $r_i' > r_i$ and for all $m \geq a$.

5) $\alpha_k = \arg\min_{(i,r_i,a)\in M_k} (c_{i,r_i,a}/T^{(\Omega)}_{i,r_i,a})$.

Proof. This result follows in a straightforward manner from [16, Lemma 1]. □

By combining [16, Theorem 1] and Lemma 1, it follows that the optimal scheduling policy for the RLPAK problem is a priority-index rule where the priorities $\alpha_1, \alpha_2, \ldots, \alpha_K$ satisfy

$$\frac{c_{\alpha_1}}{T^{(\Omega)}_{\alpha_1}} \geq \frac{c_{\alpha_2}}{T^{(\Omega)}_{\alpha_2}} \geq \cdots \geq \frac{c_{\alpha_K}}{T^{(\Omega)}_{\alpha_K}}.$$

Since the optimal scheduling policy for the RLPAK problem is a priority-index rule, we employ [16, Corollary 1] to conclude the the optimal scheduling policy for the RLPA problem is also a priority-index rule. The HoL packet with the highest priority index of $c_{i,r_i^{HoL}}/T_{i,r_i^{HoL}}$ along with the transmitter that yields that index are selected over all nonempty queues at the base station.

B Proof of Theorem 2

We follow the same approach as in [16] for proving the optimality of our proposed scheduler for the RDC problem. As in Section A, we transform the RDC problem into an instance of Klimov's multiclass queueing problem [20]. Thus, the transformed problem, which we refer to as the RDCK problem, has an optimal priority index policy. We then show that that the optimal policy for the RDCK problem is optimal for the RDC problem.

B.1 RDCK Scheduling Problem

The RDCK problem is similar to the DCK problem in [16]. The key differences are as follows. For each user i, the M relays are sorted according to their channel gains to user i as $\{d_{i,1}, d_{i,2}, \ldots, d_{i,M}\}$, where $|h_{d_{i,1}}|^2 < |h_{d_{i,2}}|^2 < \cdots < |h_{d_{i,M}}|^2$. Assume that in the RDC problem, user i initially has A_i packets at the base station.

Then, each user i has $K_i = (M+1)A_i(r_i^{max}+1)$ queues, and each queue is labeled as (i, r_i, x_i, a). If a packet is in queue (i, r_i, x_i, a), it has been transmitted r_i times, relay $d_{i,a}$ has decoded it, relay $d_{i,m}$ has not decoded it for $a < m \leq M$ and the queue length of user i at the base station is $x_i(n)$. In particular, a packet in the queue $(i, r_i, x_i, 0)$ has not been decoded by any of the M relays. Note that in the RDCK problem, if a packet is in queue (i, r_i, x_i, a), the other $K_i - 1$ queues for user i are empty.

Thus, there are a total of $K = \sum_{i=1}^{N} K_i$ queues in the RDCK problem. Each queue (i, r_i, x_i, a) has a deterministic service time of $b_{i,r_i,x_i,a} = 1$ time slot.

The transition probabilities for the queues in the RDCK problem are determined as follows. Let $g_{i,a,k}(r_i)$ denote the probability that relay $d_{i,a}$ cannot decode the HoL packet of user i after its transmission attempt r_i by relay d_k. In particular, $g_{i,a,0}(r_i)$ corresponds to a transmission by the base station. Also, let $g_{i,a}(r_i, 1)$ denote the probability that user i cannot decode its HoL packet after relay $d_{i,a}$ has transmitted it. Then

$$p_{(i,r_i,x_i,0),(i,r_i+1,x_i,0)} = g_i(r_i) g_{i,1,0}(r_i) g_{i,2,0}(r_i) \cdots g_{i,M,0}(r_i)$$

and

$$p_{(i,r_i,x_i,0),(i,r_i+1,x_i,a)} = g_i(r_i)(1 - g_{i,a}(r_i)) g_{i,a+1}(r_i) g_{i,a+2}(r_i) \cdots g_{i,M}(r_i)$$

so $p_{(i,r_i,x_i,0),(i,0,x_i-1,0)} = 1 - g_i(r_i)$. In addition

$$p_{(i,r_i,x_i,a),(i,r_i+1,x_i,n)} = g_{i,a}(r_i,1)(1 - g_{i,n,a}(r_i)) g_{i,n+1,a}(r_i) g_{i,n+2,a}(r_i) \cdots g_{i,M,a}(r_i), n > a$$

$$p_{(i,r_i,x_i,a),(i,r_i+1,x_i,a)} = g_{i,a}(r_i,1)g_{i,a+1,a}(r_i)g_{i,a+2,a}(r_i)\cdots g_{i,M,a}(r_i),$$

and

$$p_{(i,r_i,x_i,a),(i,r_i+1,x_i,n)} = 0, n < a$$

so $p_{(i,r_i,x_i,a),(i,0,x_i-1,0)} = 1 - g_i(r_i,1)$. Note that after a packet from queue $(i, r_i^{max}, 1, a)$ has been served, it departs the system with probability 1, where $a \in \{0, 1, \ldots, M\}$.

Thus, for any set $M \subset \Omega = \{1, 2, \ldots, K\}$ and any queue $(i, r_i, a) \in M$, the average total service time is

$$T^{(M)}_{i,r_i,x_i,a} = 1 + \sum_{k,r_k,m} p_{(i,r_i,x_i,a),(k,r_k,x_k,m)} T^{(M)}_{k,r_k,x_k,m}.$$

From the above discussion, it can be concluded that the RDCK problem, which is a transformed version of the RDC problem, is an instance of the multiclass queueing problem of [20]. As for the RLPAK problem, this conclusion relies on the simple queueing dynamics of the relay: 1) a packet only arrives at the relay if the base station has transmitted it, and 2) the relay automatically flushes a packet once it has been decoded by its intended user. In addition, the base station automatically flushes a packet once it has been decoded by its intended user. Again, we note that the state space of the RDCK problem is an expanded version of that in the DCK problem.

Now, the objective is to find a policy $\pi \in \Pi_R$ that minimizes the time-averaged expected cost J_{RDCK}, where

$$J_{RDCK} = \mathbb{E}_\pi \left[\sum_{n=1}^{\infty} \sum_{(i,r_i,x_i,a)\in\Omega} \mathbf{1}_{i,r_i,x_i,a}(n) U_i(x_i) \right]$$

where

$$\mathbf{1}_{i,r_i,x_i,a}(n) = \begin{cases} 1 & (i, r_i, x_i, a) \text{ is nonempty in slot n} \\ 0 & \text{otherwise.} \end{cases} \tag{11}$$

B.2 Optimal Policies for RDCK and RDC Scheduling Problems

The optimal policies for both the RDCK and RDC problems are priority-index rules. First, we state the following result.

Lemma 2. *Let M_k, $k = 1, 2, \ldots, K$ be the sets of queues generated by the Klimov algorithm in [16, Section 3] for the RCK problem. For each $k = 1, 2, \ldots, K$, for all $(i, r_i, x_i, a) \in M_k$ and for all $r_i' > r_i$, the following are true:*

1) $(i, r_i', x_i, m) \in M_k$ for all $m \geq a$.

2) $T^{(M_k)}_{i,r_i',x_i,m} \leq T^{(M_k)}_{i,r_i,x_i,a}$ for all $m \geq a$.
3) $C^{(M_k)}_{i,r_i',x_i,m} \geq C^{(M_k)}_{i,r_i,x_i,a}$ for all $m \geq a$.

Proof. This result follows in a straightforward manner from [16, Lemma 2]. □

By combining [16, Theorem 2] and Lemma 2, it follows that the optimal scheduling policy for the RDCK problem assigns queue (i, r_i', x_i, m) higher priority than queue (i, r_i, x_i, a) for all i, x_i, $r_i' > r_i$ and $m \geq a$.

Since the RDCK problem is a special case of Klimov's problem, the optimal scheduling policy for the RDCK problem is a priority-index rule. We then transform the RDCK problem back to the RDC problem to conclude that the optimal scheduling policy for the RDC problem is also a priority-index rule. The HoL packet with the highest priority index along with the transmitter that yields that index are selected over all nonempty queues at the base station. Note that a key difference between the RDC problem and the RLPA problem is that the priority indices for the RDC problem do not admit closed-form expressions.

We also have the following result for the RDC problem.

Corollary 5. *Once the optimal RDC scheduler initially selects a packet in the queue for user i at the base station, it will continue to select that packet along with the appropriate transmitter until user i successfully decodes it.*

Proof. This result follows in a straightforward manner from [16, Corollary 2]. The proof of [16, Corollary 2] should be modified to consider a packet that is initially in queue $(i, 0, x_i, 0)$ and enters queue $(i, 1, x_i, a)$ if user i fails to decode it, where $0 \leq a \leq M$. □

C Necessary Conditions for Optimal Cost Function

From [26], we know that the RDC problem can be recast as a stochastic shortest path problem over an infinite time horizon. Let $J_{RDC}(r_1, r_2, x_1, x_2, a_1, a_2)$ denote the optimal cost when starting from the initial state $(r_1, r_2, x_1, x_2, a_1, a_2)$, and this cost function must be a solution of Bellman's equation. Thus, following the definitions in Appendix B, the following conditions must be satisfied:

- $J_{RDC}(0, 0, 0, 0, 0, 0) = 0$

- for $x_1 > 0$, $J_{RDC}(r_1, 0, x_1, 0, 0, 0) = U_1(x_1) + (1 - g_1(r_1))J_{RDC}(0, 0, x_1 - 1, 0, 0, 0)$
  $+ g_1(r_1)g_{1,0}(r_1)J_{RDC}(\min(r_1 + 1, r_1^{max}), 0, x_1, 0, 0, 0)$
  $+ g_1(r_1)(1 - g_{1,0}(r_1))J_{RDC}(\min(r_1 + 1, r_1^{max}), 0, x_1, 0, 1, 0)$

- for $x_1 > 0, J_{RDC}(r_1, 0, x_1, 0, 1, 0) = U_1(x_1) + \min((1 - g_1(r_1))J_{RDC}(0, 0, x_1 - 1, 0, 0, 0)$
  $+ g_1(r_1)J_{RDC}(\min(r_1 + 1, r_1^{max}), 0, x_1, 0, 1, 0), (1 - g_1(r_1, 1))J_{RDC}(0, 0, x_1 - 1, 0, 0, 0)$
  $+ g_1(r_1, 1)J_{RDC}(\min(r_1 + 1, r_1^{max}), 0, x_1, 0, 1, 0))$

- for $x_2 > 0, J_{RDC}(0, r_2, 0, x_2, 0, 0) = U_2(x_2) + (1 - g_2(r_2))J_{RDC}(0, 0, 0, x_2 - 1, 0, 0)$
  $+ g_2(r_2)g_{1,0}(r_2)J_{RDC}(0, \min(r_2 + 1, r_2^{max}), 0, x_2, 0, 0)$
  $+ g_2(r_2)(1 - g_{1,0}(r_2))J_{RDC}(0, \min(r_2 + 1, r_2^{max}), 0, x_2, 0, 1)$

- for $x_2 > 0, J_{RDC}(0, r_2, 0, x_2, 0, 1) = U_2(x_2) + \min((1 - g_2(r_2))J_{RDC}(0, 0, 0, x_2 - 1, 0, 0)$
  $+ g_2(r_2)J_{RDC}(0, \min(r_2 + 1, r_2^{max}), 0, x_2, 0, 1), (1 - g_2(r_2, 1))J_{RDC}(0, 0, 0, x_2 - 1, 0, 0)$
  $+ g_2(r_2, 1)J_{RDC}(0, \min(r_2 + 1, r_2^{max}), 0, x_2, 0, 1))$

- for $x_1 > 0$ and $x_2 > 0, J_{RDC}(r_1, r_2, x_1, x_2, 0, 0) = U_1(x_1) + U_2(x_2)$
  $+ \min((1 - g_1(r_1))J_{RDC}(0, r_2, x_1 - 1, x_2, 0, 0) + g_1(r_1)g_{1,0}(r_1)J_{RDC}(\min(r_1 + 1, r_1^{max}), r_2, x_1, x_2, 0, 0)$
  $+ g_1(r_1)(1 - g_{1,0}(r_1))J_{RDC}(\min(r_1 + 1, r_1^{max}), r_2, x_1, x_2, 1, 0),$
  $(1 - g_2(r_2))J_{RDC}(r_1, 0, x_1, x_2 - 1, 0, 0) + g_2(r_2)g_{1,0}(r_2)J_{RDC}(r_1, \min(r_2 + 1, r_2^{max}), x_1, x_2, 0, 0)$
  $+ g_2(r_2)(1 - g_{1,0}(r_2))J_{RDC}(r_1, \min(r_2 + 1, r_2^{max}), x_1, x_2, 0, 1))$

- for $x_1 > 0$ and $x_2 > 0, J_{RDC}(r_1, r_2, x_1, x_2, 1, 0) = U_1(x_1) + U_2(x_2)$
  $+ \min((1 - g_1(r_1))J_{RDC}(0, r_2, x_1 - 1, x_2, 0, 0) + g_1(r_1)J_{RDC}(\min(r_1 + 1, r_1^{max}), r_2, x_1, x_2, 1, 0),$
  $(1 - g_1(r_1, 1))J_{RDC}(0, r_2, x_1 - 1, x_2, 0, 0) + g_1(r_1, 1)J_{RDC}(\min(r_1 + 1, r_1^{max}), r_2, x_1, x_2, 1, 0),$
  $(1 - g_2(r_2))J_{RDC}(r_1, 0, x_1, x_2 - 1, 1, 0) + g_2(r_2)g_{1,0}(r_2)J_{RDC}(r_1, \min(r_2 + 1, r_2^{max}), x_1, x_2, 1, 0)$
  $+ g_2(r_2)(1 - g_{1,0}(r_2))J_{RDC}(r_1, \min(r_2 + 1, r_2^{max}), x_1, x_2, 1, 1))$

- for $x_1 > 0$ and $x_2 > 0, J_{RDC}(r_1, r_2, x_1, x_2, 0, 1) = U_1(x_1) + U_2(x_2)$
  $+ \min((1 - g_1(r_1))J_{RDC}(0, r_2, x_1 - 1, x_2, 0, 1) + g_1(r_1)g_{1,0}(r_1)J_{RDC}(\min(r_1 + 1, r_1^{max}), r_2, x_1, x_2, 0, 1)$
  $+ g_1(r_1)(1 - g_{1,0}(r_1))J_{RDC}(\min(r_1 + 1, r_1^{max}), r_2, x_1, x_2, 1, 1),$
  $(1 - g_2(r_2))J_{RDC}(r_1, 0, x_1, x_2 - 1, 0, 0) + g_2(r_2)J_{RDC}(r_1, \min(r_2 + 1, r_2^{max}), x_1, x_2, 0, 1),$
  $(1 - g_2(r_2, 1))J_{RDC}(r_1, 0, x_1, x_2 - 1, 0, 0) + g_2(r_2, 1)J_{RDC}(r_1, \min(r_2 + 1, r_2^{max}), x_1, x_2, 0, 1))$

- for $x_1 > 0$ and $x_2 > 0, J_{RDC}(r_1, r_2, x_1, x_2, 1, 1) = U_1(x_1) + U_2(x_2)$
  $+ \min((1 - g_1(r_1))J_{RDC}(0, r_2, x_1 - 1, x_2, 0, 1) + g_1(r_1)J_{RDC}(\min(r_1 + 1, r_1^{max}), r_2, x_1, x_2, 1, 1),$
  $(1 - g_1(r_1, 1))J_{RDC}(0, r_2, x_1 - 1, x_2, 0, 1) + g_1(r_1, 1)J_{RDC}(\min(r_1 + 1, r_1^{max}), r_2, x_1, x_2, 1, 1),$
  $(1 - g_2(r_2))J_{RDC}(r_1, 0, x_1, x_2 - 1, 1, 0) + g_2(r_2)J_{RDC}(r_1, \min(r_2 + 1, r_2^{max}), x_1, x_2, 1, 1),$
  $(1 - g_2(r_2, 1))J_{RDC}(r_1, 0, x_1, x_2 - 1, 1, 0) + g_2(r_2, 1)J_{RDC}(r_1, \min(r_2 + 1, r_2^{max}), x_1, x_2, 1, 1))$

References

[1] ESS Mesh Networking. http://www.ieee802.org/11/Reports/tgs_update.htm.

[2] Relay Task Group. http://www.ieee802.org/16/relay/index.html.

[3] B. Zhao and M.C. Valenti. Practical relay networks: a generalization of hybrid-ARQ. *IEEE J. Select. Areas Commun.*, 23(1):7–18, January 2005.

[4] P. Razaghi and W. Yu. Parity forwarding for multiple-relay networks. In *Proc. of the Intl. Symp. on Inform. Theory*, 1:1678–1682, Seattle, WA, July 2006.

[5] P. Popovski and E. de Carvalho. Spectrally-efficient wireless relaying based on superposition coding. In *Proc. of the IEEE Veh. Technol. Conf.*, 1:2936–2940, Dublin, Ireland, April 2007.

[6] B. Rankov and A. Wittneben. Spectral efficient protocols for half-duplex fading relay channels. *IEEE J. Select. Areas Commun.*, 25(2):379–389, February 2007.

[7] D. Gunduz and E. Erkip. Opportunistic cooperation by dynamic resource allocation. *IEEE J. Trans. Wireless Commun.*, 6(4):1446–1454, April 2007.

[8] M. Yuksel and E. Erkip. Broadcast strategies for the fading relay channel. In *Proc. of the IEEE Military Commun. Conf.*, 2:1060–1065, Monterey, CA, October 2004.

[9] Y. Liu, R. Hoshyar, X. Yang and R. Tafazolli. Integrated radio resource allocation for multihop cellular networks with fixed relay stations. *IEEE J. Select. Areas Commun.*, 24(11):2137–2146, November 2006.

[10] S. Shakkottai and R. Srikant. Scheduling real-time traffic with deadlines over a wireless channel. *Wireless Netw.*, 8(1):13–26, January 2002.

[11] P. Liu, R.A. Berry and M.L. Honig. A fluid analysis of a utility-based wireless scheduling policy. *IEEE Trans. Inform. Theory*, 52(7):2872–2889, July 2006.

[12] H. Viswanathan and S. Mukherjee. Performance of cellular networks with relays and centralized scheduling. *IEEE Trans. Wireless Commun.*, 4(5):2318–2328, September 2005.

[13] A.S. Hu and S.D. Servetto. On the scalability of cooperative time synchronization in pulse-connected networks. *IEEE Trans. Inform. Theory*, 52(6):2725–2748, June 2006.

[14] S. Wei, D.L. Goeckel and M.C. Valenti. Asynchronous cooperative diversity. *IEEE Trans. Wireless Commun.*, 5(6):1547–1557, June 2006.

[15] W.C. Lindsey, F. Ghazvinian, W.C. Hagmann and K. Dessouky. Network synchronization. *Proc. of the IEEE*, 73(10):1445–1467, October 1985.

[16] J. Huang, R.A. Berry and M.L. Honig. Wireless scheduling with hybrid-ARQ. *IEEE Trans. Wireless Commun.*, 4(6):2801–2810, November 2005.

[17] D.W. Tcha and S.R. Pliska. Optimal control of single-server queueing networks and multi-class M/G/1 queues with feedback. *Oper. Res.*, 25(2):248–258, March-April 1977.

[18] Y. Lu and D.D. Yao. Optimal control of a fluid network with side constraints. *IEEE Trans. Autom. Control*, 48(10):1865–1870, October 2003.

[19] D. Bertsekas and R. Gallager. *Data Networks*. Prentice Hall, Upper Saddle River, NJ, 1991.

[20] G.P. Klimov. Time-sharing service systems. I. *Theor. Probab. Appl.*, 19(3):532–551, 1974.

[21] H. Chen and D.D. Yao. Dynamic scheduling of a multiclass fluid network. *Oper. Res.*, 41(6):1104–1115, November-December 1993.

[22] D. Bertsimas, I.C. Paschalidis and J.N. Tsitsiklis. Optimization of multiclass queueing networks: polyhedral and nonlinear characterizations of achievable performance. *Ann. Appl. Probab.*, 4(1):43–75, 1994.

[23] M.J. Neely, E. Modiano and C.E. Rohrs. Dynamic power allocation and routing for time-varying wireless networks. *IEEE J. Select. Areas Commun.*, 23(1):89–103, January 2005.

[24] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. Autom. Control*, 37(12):1936–1948, December 1992.

[25] E.M. Yeh and R.A. Berry. Throughput optimal control of cooperative relay networks. *IEEE Trans. Inform. Theory*, 53(10):3827–3833, October 2007.

[26] D.P. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, Belmont, MA, 1995.

[27] P. Glasserman and J. Staum. Resource allocation among simulation time steps. *Oper. Res.*, 51(6):908–921, November-December 2003.

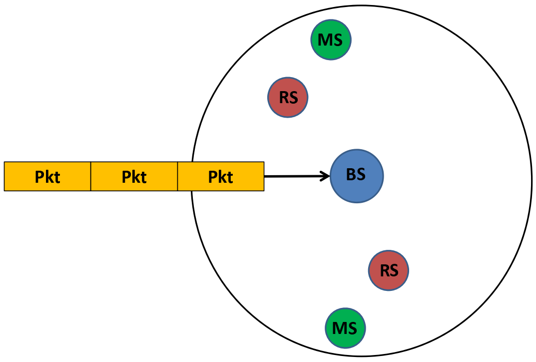

Figure 1: Wireless network with relay-assisted scheduling.

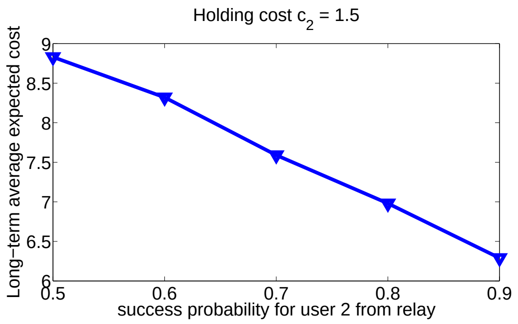

Figure 2: Long-term average expected cost of relaying for RLPA as function of channel from relay to user 2.

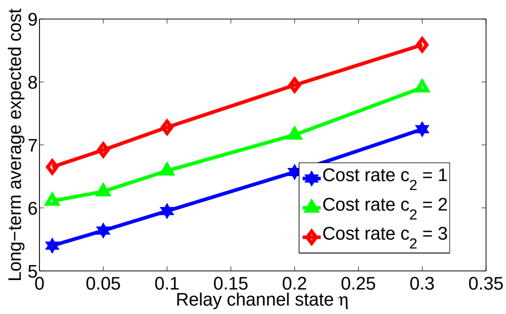

Figure 3: Long-term average expected cost of relaying for RLPA as function of channel from relay to users 1 and 2.

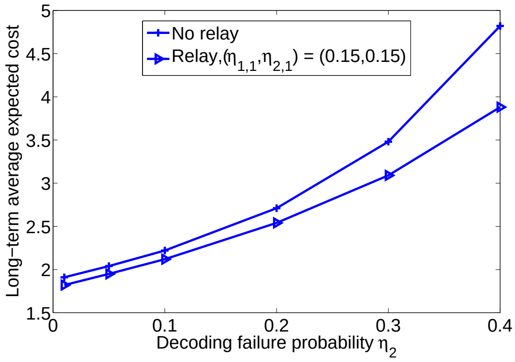

Figure 4: Long-term average expected cost of relaying for RLPA as function of channel from base station to user 2.

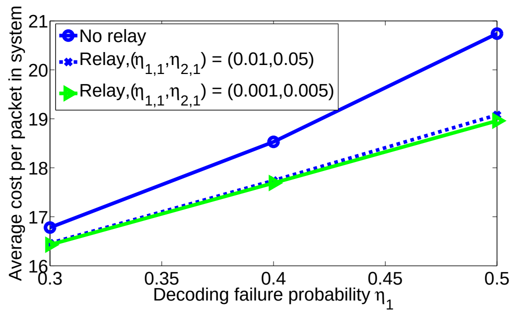

Figure 5: Average cost per packet for RDC as function of channel from base station to both users.

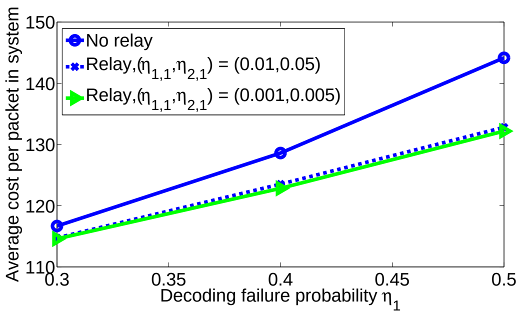

Figure 6: Average cost per packet for RDC with cost function $U(x) = x^{2.1}$.

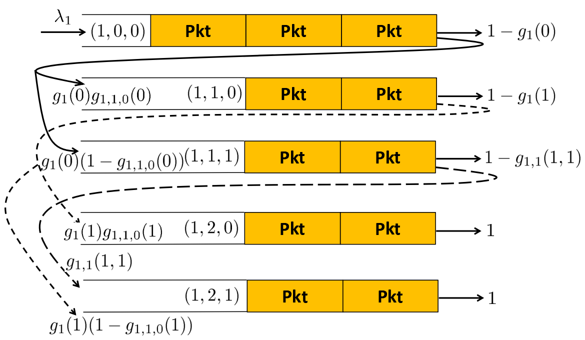

Figure 7: System model for RLPAK problem with queues for user 1.