Distributed Strategies for Channel Allocation and Scheduling in Software-Defined Radio Networks

Bo Han*, V. S. Anil Kumar†, Madhav V. Marathe†, Srinivasan Parthasarathy‡, and Aravind Srinivasan§
*Department of Computer Science, University of Maryland, College Park, MD 20742, USA
†Department of Computer Science and Virginia Bioinformatics Institute, Virginia Tech, Blacksburg, VA 24061
‡IBM T.J. Watson Research Center, Hawthorne, NY 10532
§Dept. of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742

***Abstract*—Equipping wireless nodes with multiple radios can significantly increase the capacity of wireless networks, by making these radios simultaneously transmit over multiple non-overlapping channels. However, due to the limited number of radios and available orthogonal channels, designing efficient channel assignment and scheduling algorithms in such networks is a major challenge. In this paper, we present provably-good distributed algorithms for simultaneous channel allocation of individual links and packet-scheduling, in Software-Defined Radio (SDR) wireless networks. Our distributed algorithms are very simple to implement, and do not require any coordination even among neighboring nodes. A novel *access hash function* or *random oracle* methodology is one of the key drivers of our results. With this *access hash function*, each radio can know the transmitters' decisions for links in its interference set for each time slot without introducing any extra communication overhead between them. Further, by utilizing the *inductive-scheduling technique*, each radio can also backoff appropriately to avoid collisions. Extensive simulations demonstrate that our bounds are valid in practice.**

I. Introduction

Significant advances in wireless technology have resulted in cheaper, reliable and adaptive wireless devices. While wireless networks are much easier to deploy in comparison with wire-line networks, the phenomenon of wireless interference poses a major challenge in the task of operating a wireless network close to its optimal throughput capacity. One of the techniques used to improve the performance of wireless networks is to design multi-channel multi-radio (MCMR) networks in which each node is equipped with multiple radios that can operate on multiple (non-overlapping) channels. Additionally, recent advances in radio technology have led to the design of *Software-Defined Radios (SDR)* [18], in which packet transmissions on a radio can be switched from one channel to another dynamically. While these advances have led to an increase in network capacity, they also introduce a host of difficult algorithmic challenges, as nodes need to dynamically make decisions at a per-packet level about which radios and which channels they will employ for communication at any time.

In this work, we deal with a fundamental algorithmic issue that arises in such networks. The problem we address is the following: given a network formed by a set of nodes V, and a collection of source-destination pairs $(s_1, t_1), \ldots, (s_k, t_k)$, what is the maximum throughput capacity of the system, *i.e.*, the maximum rate at which packets can be sent from the sources to their corresponding destinations?[1] In an SDR network, a node could have multiple radios, each of which can transmit or receive on multiple channels (called frequency bands in SDR terminology). Wireless interference places constraints on which pairs of radios can communicate simultaneously on the same channel. Thus, the throughput optimization presented above decomposes into the following sub-problems: (i) determining the end-to-end throughput r_i achieved by connection i; (ii) choosing routes for each (s_i, t_i) pair; and (iii) determining which pairs of radios would communicate at each time step and on which channels. Thus we have a cross-layer optimization problem involving constraints from the transport (end-to-end rate control), routing, and MAC (channel allocation and link scheduling) layers.

Our central focus is the design and analysis of channel assignment & scheduling strategies. Given an SDR network, consider the *utilization matrix* $\mathbf{X}$ whose rows correspond to ordered pairs of radios in the network, columns correspond to channels, and each entry of $\mathbf{X}$ specifies the *fraction of the time* the corresponding pair of radios communicate in the associated channel. The set of utilization matrices achievable by any channel allocation and link-scheduling scheme is the *capacity region* of the SDR network. A channel allocation and scheduling scheme $\mathcal{A}$ is *α-competitive*, if for any $\mathbf{X}$ in the capacity region, $\mathcal{A}$ can achieve (component-wise) $\frac{1}{\alpha}$-fraction of $\mathbf{X}$. It is well-known that this α essentially determines the performance ratio of a solution to the overall throughput-optimization problem [1]. In fact, as shown in [8], it is also possible to convert an α-competitive *distributed* channel allocation and scheduling scheme into an α-competitive distributed strategy for the cross-layer optimization problem using distributed flow mechanisms [5], [2]. Motivated by these observations, we present near-optimal and provably-competitive distributed schemes for joint channel allocation and scheduling in SDR wireless networks.

A. Our Contributions

Our first contribution is PLDS, a Purely Localized Distributed Scheme for joint channel assignment and link schedul-

[1]Our algorithms solve the general problem of maximizing $\sum_i U_i(r_i)$, where r_i is the throughput for connection i and U_i is any concave function.

ing; we prove that it is $(\Delta + 2) \cdot e$-competitive (e denotes the base of natural logarithms). Here, Δ is the *independence number* of the network which is defined as the maximum number of links that interfere with a given link and which can transmit simultaneously without mutual interference. This parameter is of importance since, for several interference models, this parameter can be upper bounded by a *fixed* constant independent of the network [1], [11]. The key innovation behind PLDS is the notion of an *access hash function*, a binary function that is parameterized by (i) an ordered pair of radios; (ii) a channel; and (iii) the index of the current time slot. The binary outcome of this function (probabilistically) determines whether or not the pair of radios will communicate over the channel during this time slot. A unique property of the access hash function is that it carefully introduces dependencies in the random choices made by the end-points of a link with the aim of increasing their probability of choosing the same channel, while *simultaneously* minimizing the probability of other conflicting radios choosing the same channel. Access hash functions may be viewed as a non-trivial generalization of random-access scheduling for SDR networks; here, a stochastic process not only arbitrates the accesses to the wireless medium over time, but also yields the channel assignment. In order to appreciate its significance, consider the following natural generalization of p-persistent MAC for slotted ALOHA, a well-studied random-access protocol for SCSR networks [3]: during each time slot, for a given link, each end-point chooses a channel *independently* at random and then communicates on this channel with some pre-defined probability. *Unlike PLDS, this protocol can be shown to have a competitive ratio which is arbitrarily far from optimal.*

Our second contribution is CFDS, a Collision-Free Distributed Scheme which combines the access-hash function methodology of PLDS along with the inductive-scheduling techniques of Kumar *et al.* [10]. CFDS achieves a competitive ratio of $(\lambda+2)\cdot e$, and as its name indicates, it is collision-free (i.e., interfering links never transmit simultaneously) although channel assignment and scheduling decisions are made in a completely distributed manner. Here, λ is the *inductive number* of the network which is defined as the maximum number of "larger" links which interfere with a given link and which can transmit simultaneously without mutual interference (we define this formally in §II-C). As in the case of the independence number Δ, for a large class of geometric interference models, the inductive degree of a network is also upper bounded by a *fixed* constant independent of the size or topology of the network [1], [9], [10], [11], [20]. The best known algorithms for link scheduling in *single-channel single-radio* (SCSR) wireless networks essentially achieve a competitive ratio of λ [10], [11], [20]. The *catch* in this scheme is that it presupposes a protocol initialization phase in which each link communicates a single value (its utilization) to other links which interfere with it. Despite this overhead, we believe that the collision-free property of CFDS renders it particularly attractive for energy-starved scenarios such as sensor networks. In CFDS, nodes can spend a large fraction of their time in *sleep* mode and only need to wake up during (locally) pre-computed time slots for communication, with the guarantee of no loss due to collisions.

To the best of our knowledge, PLDS and CFDS are the first purely local-control distributed algorithms for joint channel assignment and scheduling in SDR networks, which do not require any signaling amongst nodes for individual (per-packet) transmissions. A salient aspect of our work is that the competitive factors we derive for all our schemes are independent of the topological properties of the network such as size, degree, the number of radios available at a node, and the number of channels available at a radio; instead, they are fixed constants that depend only upon the specific interference model we assume. Variations in the physical-layer transmission technologies and link-layer schemes have resulted in a large variety of interference models being studied in the literature. Our work presents a unified framework for SDR throughput optimization across these broad spectrum of interference models.

II. Background

A. Network model

We use the disk-graph model for the physical layer [1], [10] and model the wireless network as a *directed* graph $G = (V, E)$, where V denotes the set of nodes in the network and E denotes all ordered pairs of nodes across which direct communication is possible. Each node $v \in V$ is equipped with a collection of radios denoted by $Radios(v)$; each radio ρ is associated with a collection of non-interfering channels denoted by $Channels(\rho)$. We assume that radios can dynamically switch their channels by using SDR technology. Unlike the usual model of MCMR networks where the entire system has a fixed collection of common channels, we make the more general assumption that each radio has its own collection of channels. As mentioned above, for a successful transmission to occur on a link ℓ, the transmitting and receiving radios of ℓ need to use the same channel.

For simplicity, we assume that all the radios at a node u transmit at the same power level, and there is an edge (u, v) if node v can be reached by a transmission from node u at this power level. We will assume a synchronous model of time: time is divided into equi-sized slots that are indexed $0, 1, 2, \ldots$. We note that this is a common assumption of almost all the existing works for time-slotted scheduling algorithms [1], [8], [20]. We assume that the channel switching latency is negligible compared with the duration of a time slot. An (edge, channel)-pair (w, ψ) has a capacity of $cap(w, \psi)$ bits/slot. This is the maximum number of bits that can be transmitted across edge w on channel ψ in a single time slot. If ψ is *not* a channel that is available at both the end-points of edge w, then we assume w.l.o.g. that $cap(w, \psi)$ is zero (since transmitting on channel ψ is not an option for edge w).

B. Interference in SCSR Networks

We first describe our interference models in the context of SCSR networks and extend them in §II-D to SDR networks.

We will consider edge-conflict-based interference models which specify interference as a binary relation between pairs of edges in G. An interference model specifies, for each edge $w \in E$, a subset of edges $I(w) \subseteq E \setminus \{w\}$. A transmission on edge w during some time slot is successful if and only if no other edges in the set $I(w)$ are active during the same slot. In a single-radio network, each node can be involved in at most one transmission during any time slot. Hence, we will assume w.l.o.g. that all edges other than w that are incident on the end-points of w belong to $I(w)$. In reality, whether two edges in a network interfere or not is determined essentially by their relative locations in space and physical laws of radio propagation. Hence, we will lay special focus on *geometric* interference models that lay down geometric conditions under which one link interferes with another. Thus, an interference model is not defined w.r.t. a specific network, but is a set of rules for determining conflicting link-pairs in any network. Several geometric interference models have been studied in the literature due to variations in the physical layer hardware of wireless networks as well as differences in physical layer transmission technologies. In all the models below, the nodes are assumed to be embedded in a two-dimensional plane. Each node u has a transmission range $r_{tx}(u)$. A necessary condition for edge $w = (u, v) \in E$ to be present is that node v is within a distance of $r_{tx}(u)$ from u. We consider seven models: (a) Node-exclusive model [13]; (b) Non-uniform RTS-CTS model [9], [20]; (c) Uniform RTS-CTS model with parameter q [1]; (d) Tx-model [10], [22]; (e) fPrIM model [20]; (f) Protocol model [6], [10]; and (g) K-hop model [17].

In all of these models, if two edges $w_1 = (u_1, v_1)$ and $w_2 = (u_2, v_2)$ share a common end-point, then they are in the interference sets of each other: a node can be involved in at most one transmission on any link during a time slot. As is standard convention [15], we define $I_{pri}(w_1)$ to be the primary interference set of w_1, which contains all other edges w_2 that share an end-point with edge w_1. We also define $I_{sec}(w_1) = I(w_1) \setminus I_{pri}(w_1)$ to be the secondary interference set of w_1; this contains all edges which interfere with w_1 without sharing any end-point with w_1.

C. Independence and Induction in SCSR networks

We now define the key notions of Δ-*independence* and λ-*induction* for SCSR networks; we then extend them in §II-D to SDR networks. Given a network $G = (V, E)$ and an associated interference model, we say that the independence number of G is $\Delta(G)$ if $\Delta(G)$ is the maximum number of links that are in the interference set of some specific link w, but are mutually interference-free amongst themselves:

$$\Delta(G) \stackrel{\text{def}}{=} \max_{w \in E} \max_{J \subseteq I(w) \,:\, \nexists (j_1, j_2) \in J \;s.t.\; ((j_1 \neq j_2) \wedge (j_2 \in I(j_1)))} |J|.$$

Given $G = (V, E)$ and an interference model, the induction number of G is $\lambda(G)$ if there exists a total ordering $\succ$ of the links such that $\lambda(G)$ is the maximum number (taken over all links) of links that are (i) $\succ w$, and (ii) in the interference set of link w, but are mutually interference-free amongst themselves. Formally, given $\succ$, let $I_{\succ}(w)$ denote the set of links that are greater-than w but interfere with w. Then,

$$\lambda(G) \stackrel{\text{def}}{=} \max_{w \in E} \max_{J \subseteq I_{\succ}(w) \,:\, \nexists (j_1, j_2) \in J \;s.t.\; ((j_1 \neq j_2) \wedge (j_2 \in I(j_1)))} |J|.$$

Given an interference model and a *finite constant* Δ, we say that it is Δ-independent if for *any* network $G = (V, E)$, under this interference model, G's independence number $\Delta(G)$ is at most Δ. Given an interference model and a *finite constant* λ, we say that it is λ-inductive if for *any* network $G = (V, E)$, under this interference model, there exists a total ordering $\succ$ of G's links such that the induction number of G under this ordering is at most λ. In this case, we will also call the ordering $\succ$ as the *greater-than* ordering of edges in E which achieves the λ-induction property. Clearly, while $\Delta(G)$ and $\lambda(G)$ are properties of a given network G, the Δ-independence and λ-induction are properties of an interference model and are *independent* of the size of the network and the network topology. Also, observe that the requirement of independence is *stronger* than the notion of induction: an interference model which is Δ-independent is always Δ-inductive. In particular, networks with nodes of heterogeneous (transmission) ranges can have much higher values of Δ than λ.

Remark: using Δ and λ. The utility of Δ and λ owes to the fact that several geometric interference models considered in the literature have constant induction or constant independence number [1], [9], [10], [11], [20]. For instance, λ is at most 5 for the Tx-model, at most 4, 8 and 12 for the Tx-Rx model with parameters 1, 2, and 2.5 [1], [9], [10], [11], [20]. Also, the corresponding ordering $\succ$ of the edges is usually a geometric function: e.g., descending order of the sum of the transmission radii of the two end-points of the edge [9].

D. Interference and Scheduling in SDR Networks

It is convenient to employ the notion of an *induced radio network*, in order to extend our interference models to SDR networks. Given a network $G = (V, E)$, the induced radio network $\mathcal{G} = (\mathcal{V}, \mathcal{L})$ is defined as follows. The set of nodes in $\mathcal{V}$ is the set of all *radios* which belong to the nodes of V. If radio $\rho \in \mathcal{V}$ belongs to the node $u \in V$, then we say that node u is the parent of radio ρ. A link $\ell = (\rho, \rho')$ lies in $\mathcal{L}$ if and only if: (i) the two radios ρ and ρ' have at least one channel in common in their channel sets; and (ii) the edge $w = (u, v) \in E$ exists, where u is the parent of ρ and v is the parent of ρ'. In this case, we will say that edge w is the parent of link ℓ.[2] The capacity $cap(\ell, \psi)$ of a link $\ell \in \mathcal{L}$ when it transmits on channel ψ, is equal to $cap(w, \psi)$, where w is the parent edge of ℓ. We now describe a natural extension of our interference models to an induced radio network. If an end-point of link ℓ_1 is the same as an end-point of link ℓ_2 (*i.e.*, they have a radio in common), then they belong to the primary interference sets of each other. For a fixed link ℓ, we let $Pri(\ell)$ denote the set of all links that are in the primary interference set of ℓ. Let edges $w_1 \in E$ and $w_2 \in E$ denote the parents of links ℓ_1 and

[2]We will use the terms *nodes* and *edges* in the context of the parent graph G, and *radios* and *links* for $\mathcal{G}$.

ℓ_2; if the end-points of w_1 and w_2 have a node in common, then ℓ_1 and ℓ_2 belong to the secondary interference sets of each other. Further, if w_2 is in the secondary interference set of w_1, then ℓ_2 is in the secondary interference set of ℓ_1 (and vice-versa). We let $Sec(\ell)$ denote the set of links in the secondary interference set of link ℓ. The interfering links that share a node in common are called Type I secondary interfering links and the other secondary interfering links are called Type II secondary interfering links. Analogous to the parent network, if two links ℓ_1 and ℓ_2 in the induced radio network interfere with each other, then they cannot both transmit successfully at the same time slot *using the same channel*.

Induction and independence in SDR networks: We show how to export these two notions from the parent network to the induced radio network, in § III.

Given an induced radio network $\mathcal{G} = (\mathcal{V}, \mathcal{L})$, a schedule $\mathcal{S}$ describes the specific times at which data is moved over the links of the network, and the channel-assignment decisions made for each link during each of its transmissions. Let $Z_{(\ell,\psi),t}$ be a scheduling variable which is indexed by a link-channel pair (ℓ, ψ) and time slot t: $Z_{(\ell,\psi),t} = 1$ if link ℓ transmits successfully at time t on channel ψ, and $Z_{(\ell,\psi),t} = 0$ otherwise. A schedule $\mathcal{S}$ is an assignment of values to the variables $Z_{(\ell,\psi),t}$.

Given an induced radio network and a corresponding schedule, we define the *utilization value*, $x(\ell, \psi)$, of each link-channel pair (ℓ, ψ) as follows: $x(\ell, \psi) \stackrel{\text{def}}{=} \liminf_{t\to\infty} \frac{\sum_{j=1}^{t} Z_{(\ell,\psi),j}}{t}$. The utilization $x(\ell, \psi)$ of the pair (ℓ, ψ) is the *fraction of the time* during which link ℓ is successfully transmitting on channel ψ in a schedule. A *utilization matrix* $\mathbf{X}$ specifies the required utilization for each link-channel pair in the network. The rows of $\mathbf{X}$ correspond to the network links, while the columns correspond to channels. The entry $x(\ell, \psi)$ corresponding to row ℓ and column ψ denotes the required utilization for the link-channel pair (ℓ, ψ). This entry can be non-zero only if both the radios comprising link ℓ have ψ as one of their common channels. $\mathbf{X}$ is *stable* iff there exists a schedule which meets the utilization requirements for all link-channel pairs, as specified by $\mathbf{X}$. Note that $\mathbf{X}$ is determined by the routing protocols and that our scheduling algorithms assume that it is given in advance. Let Ψ be the set of all available channels; Table I lists all of our notation.

Notation	Meaning
ρ	radio
ℓ	link
$\mathcal{L}$	the set of links
ψ	channel
Ψ	the set of all available channels
$\mathcal{S}$	schedule
$\mathbf{X}$	utilization matrix
G	directed graph (network)
$\mathcal{G}$	induced radio network
$\Delta(G)$	independence number of G
$\lambda(G)$	induction number of G
$I(e)$	the interference set of edge e
$Pri(\ell)$	the primary interference set of link ℓ
$Sec(\ell)$	the secondary interference set of link ℓ
H	access-hash function
η	parameter of protocol interference model
q	parameter of fPrIM interference model

Table I
NOTATION USED IN THIS PAPER.

III. Necessary conditions for scheduling

We now develop two necessary conditions that are satisfied by any stable utilization matrix $\mathbf{X}$. Given a λ-inductive model and a parent network $G = (V, E)$, we first extend its greater-than ordering $\succ$ to its induced radio network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ as follows. Given links ℓ_1 and ℓ_2, if $parent(\ell_1) \succ parent(\ell_2)$, then we let $\ell_1 \succ \ell_2$. Similarly, if $parent(\ell_2) \succ parent(\ell_1)$, we let $\ell_2 \succ \ell_1$. However, if $parent(\ell_1) = parent(\ell_2)$, then we break the tie between ℓ_1 and ℓ_2 arbitrarily, and order them in some fixed manner. Define $Pri_{\succ}(\ell)$ as the set of links in $Pri(\ell)$ that are greater-than ℓ; define $Sec_{\succ}(\ell)$ analogously. In Theorem 1, we work under an interference model that is λ-inductive, a parent network $G = (V, E)$, its induced radio network $\mathcal{G} = (\mathcal{V}, \mathcal{L})$, and the greater-than ordering of links in $\mathcal{L}$. In Theorem 2, we work with an interference model that is Δ-independent, and the networks G and $\mathcal{G}$. Due to space limitation, we omit the proof of Theorem 1 and 2 and refer interested readers to [7].

Theorem 1: Consider a network $G = (V, E)$ which is λ-inductive; let $\mathbf{X}$ be a utilization matrix which is defined on $\mathcal{G}$. Matrix $\mathbf{X}$ can be stably scheduled only if the following condition holds $\forall(\ell, \psi) \in \mathcal{L} \times \Psi$:

$$x(\ell,\psi) + \sum_{\rho\in\Psi\setminus\{\psi\}} x(\ell,\rho) + \sum_{\chi\in\Psi}\sum_{f\in Pri_{\succ}(\ell)} x(f,\chi) + \sum_{g\in Sec_{\succ}(\ell)} x(g,\psi) \le \lambda + 2 \quad (1)$$

Theorem 2: Consider a network $G = (V, E)$ which is Δ-independent; let $\mathbf{X}$ be a utilization matrix which is defined on $\mathcal{G}$. Matrix $\mathbf{X}$ can be stably scheduled only if the following condition holds $\forall(\ell, \psi) \in \mathcal{L} \times \Psi$:

$$x(\ell,\psi) + \sum_{\rho\in\Psi\setminus\{\psi\}} x(\ell,\rho) + \sum_{\chi\in\Psi}\sum_{f\in Pri(\ell)} x(f,\chi) + \sum_{g\in Sec(\ell)} x(g,\psi) \le \Delta + 2 \quad (2)$$

Theorem 1 (Theorem 2) essentially states that if the parent network G is λ-inductive (Δ-independent), then the induced radio network is $\lambda + 2$-inductive ($\Delta + 2$-independent). In §IV, we will use the bounds yielded by Theorems 1 and 2 respectively to establish the performance of our scheduling and channel-assignment algorithms.

IV. Distributed channel assignment and scheduling

We now present two distributed algorithms for channel assignment and scheduling. The first needs only local information but is not collision-free, and the second completely avoids collisions but needs to exchange interference and link-utilization information during an initial setup phase.

A. Purely Localized Distributed Algorithm

Our basic approach is as follows. During each slot, each radio in the network chooses *at random*, one of the links incident on it. It also assigns a channel at random to this link. If the radio is the transmitting end-point of the link, it transmits on the chosen link on the chosen channel. Otherwise, it tries to receive data on the chosen link on the chosen channel. A successful transmission occurs on the link-channel pair (ℓ, ψ) during slot t, if both the receiving and transmitting end-points of ℓ choose link ℓ, and assign the same channel ψ to ℓ at slot t, *as well as* no transmissions are attempted during this slot on other link-channel pairs that interfere with (ℓ, ψ). It is clear that if all the radios were to make their choices in a completely uncoordinated manner, each attempted transmission in this scheme will have an abysmally low probability of succeeding; on the other hand, perfect coordination is undesirable and/or expensive in a distributed setting. Our key innovation is the use of an *access-hash function*, which is a parameterized binary hash function that is used by all the nodes in the network, and which carefully introduces dependencies between the various random choices made by the nodes in the network. Such functions are also known as *random oracles* in cryptography. Before we describe its use in our algorithm, we first describes a few key properties of the function. The access-hash function H takes as input three parameters: (i) a link ℓ, (ii) a channel ψ, and (iii) the index of the current time slot t. An invocation of $H(\ell, \psi, t)$ returns a value 1 with probability $1 - e^{-e \cdot x(\ell,\psi)}$, and a value 0 with probability $e^{-e \cdot x(\ell,\psi)}$. Once the input parameters to H are fixed, the value returned by H does not change. For instance, given that $H(\ell, \psi, t)$ returned 1 when invoked by radio ρ_1, with probability 1, $H(\ell, \psi, t)$ returns value 1 when invoked by another radio ρ_2. Finally, the random variables $\{H(\ell, \psi, t)\}$ are independent of each other. Popular hashing techniques can be used to implement H: we use SHA-1 in our simulations.

We are now ready to describe our distributed algorithms. During slot t, a radio ρ computes the values $H(\ell, \psi, t)$ for all the links ℓ that are incident on it (*i.e.*, both outgoing and incoming links). It then randomly selects a pair (ℓ, ψ) such that $H(\ell, \psi, t) = 1$ (if no such pair exists, u sleeps during time t). If ℓ is an outgoing link, it transmits data across ℓ on channel ψ during this slot. Otherwise, if ℓ is an incoming link, it tunes to channel ψ and awaits an incoming transmission from the transmitting end-point of ℓ, on channel ψ during this slot. The pseudo-code for the distributed algorithm is presented in Algorithm 1; **the actions taken by this algorithm are by a specific radio ρ at a specific time slot t - each radio executes this distributed algorithm during each time slot**.

We now present a sufficient condition under which Algorithm 1 is guaranteed to achieve the given link utilization matrix $\mathbf{X}$. The stability condition is *asymptotic boundedness of queue-sizes* as in [13]. Theorem 3 basically shows that the probability of successful transmission on pair (ℓ, ψ) at any time t is at least $x(\ell, \psi)$. This theorem and Theorem 4, technically require an additional slack of ϵ that can be arbitrarily small but positive: this slack is useful in proving stability by a standard Chernoff-Hoeffding large-deviations approach, which is omitted here for lack of space.

Algorithm 1 PURELY LOCALIZED DISTRIBUTED SCHEME (PLDS)(Matrix $\mathbf{X}$)

Require: Access hash function $H(\ell, \psi, t)$ such that:

$$H(\ell, \psi, t) = \begin{cases} 1 & \text{with probability } 1 - e^{-e \cdot x(\ell,\psi)} \\ 0 & \text{with probability } e^{-e \cdot x(\ell,\psi)} \end{cases} \quad (3)$$

Require: Given a fixed triplet (ℓ, ψ, t), every invocation of $H(\ell, \psi, t)$ yields the same result

Require: The random variables $\{H(\ell, \psi, t)\}$ are independent

1: For all links $\ell \in \mathcal{L}_{out}(\rho) \bigcup \mathcal{L}_{in}(\rho)$ (*i.e.*, for all links incident on ρ), and for all $\psi \in \Psi$, compute $H(\ell, \psi, t)$

2: Randomly pick a pair (ℓ, ψ) such that $H(\ell, \psi, t) = 1$; if no such pair exists, sleep during time t

3: If the selected link $\ell \in \mathcal{L}_{out}(\rho)$, then schedule an outgoing transmission across ℓ on channel ψ at time t; else, if $\ell \in \mathcal{L}_{in}(\rho)$, then tune to channel ψ and await an incoming transmission across ℓ on channel ψ at time t

Theorem 3: Let $\epsilon > 0$ be an arbitrary constant. Consider an induced radio network $\mathcal{G} = (\mathcal{V}, \mathcal{L})$; let $\mathbf{X}$ be a utilization matrix which is defined on $\mathcal{G}$. Matrix $\mathbf{X}$ can be stably scheduled by Algorithm 1 if

$$\forall (\ell, \psi),\ x(\ell, \psi) + \sum_{\rho \in \Psi \setminus \{\psi\}} x(\ell, \rho) + \sum_{\chi \in \Psi} \sum_{f \in Pri(\ell)} x(f, \chi) + \sum_{g \in Sec(\ell)} x(g, \psi) \leq \frac{1}{e} - \epsilon \quad (4)$$

Proof: Let $y(\ell, \psi)$ denoted the expected utilization of the pair (ℓ, ψ) in the schedule obtained by PLDS: this is the expected fraction of the time link ℓ is *successfully* active on channel ψ. We now prove that $\forall (\ell, \psi),\ y(\ell, \psi) \geq x(\ell, \psi)$, which will suffice. Let $\mathcal{A}(\ell, \psi, t)$ denote the event that link ℓ is successfully active on channel ψ during time-slot t (*i.e.*, the pair (ℓ, ψ) is chosen for transmission during slot t but no other interfering pair is chosen for transmission during t). Since the random process which occurs during every time slot is identical, it follows that $y(\ell, \psi) = \Pr[\mathcal{A}(\ell, \psi, t)]$ for an arbitrary fixed t. Let $\mathcal{B}(\ell, \psi, t)$ denote the following event: $H(\ell, \psi, t) = 1$ *and* $\forall \rho \in \Psi \setminus \{\psi\},\ H(\ell, \rho, t) = 0$ *and* $\forall (f, \chi) \in Pri(\ell) \times \Psi,\ H(f, \chi, t) = 0$ *and* $\forall g \in Sec(\ell),\ H(g, \psi, t) = 0$. Clearly, event $\mathcal{A}(\ell, \psi, t)$ occurs whenever $\mathcal{B}(\ell, \psi, t)$ occurs. Since $y(\ell, \psi) = \Pr[\mathcal{A}(\ell, \psi, t)] \geq \Pr[\mathcal{B}(\ell, \psi, t)]$, we have:

$$\begin{aligned} y(\ell, \psi) \quad \geq \quad & \Pr[H(\ell, \psi, t) = 1] \cdot \Pi_{\rho \in \Psi \setminus \{\psi\}} \Pr[H(\ell, \rho, t) = 0] \\ & \Pi_{(f,\chi) \in Pri(\ell) \times \Psi} \Pr[H(f, \chi, t) = 0] \\ & \Pi_{g \in Sec(\ell)} \Pr[H(g, \psi, t) = 0] \\ & / * \text{ since the } H(\cdot, \cdot, \cdot) \text{ are independent} * / \\ = \quad & \left(1 - e^{-e \cdot x(\ell,\psi)}\right) \Pi_{\rho \in \Psi \setminus \{\psi\}} e^{-e \cdot x(\ell,\rho)} \\ & \Pi_{(f,\chi) \in Pri(\ell) \times \Psi} e^{-e \cdot x(f,\chi)} \end{aligned}$$

$$
\begin{aligned}
& \Pi_{g\in Sec(\ell)} e^{-e\cdot x(g,\psi)} \\
\geq\ & \left(1-e^{-e\cdot x(\ell,\psi)}\right)\cdot e^{-e\cdot\left(\frac{1}{e}-x(\ell,\psi)\right)} \quad \text{(from (4))} \\
=\ & \left(1-e^{-e\cdot x(\ell,\psi)}\right)\cdot e^{e\cdot x(\ell,\psi)-1} \\
=\ & \frac{e^{e\cdot x(\ell,\psi)}-1}{e} \geq x(\ell,\psi).
\end{aligned}
$$

The slack ϵ is useful in proving stability, which we will show in the full version of this paper. ■

B. Collision-free Distributed Scheduling Algorithm

PLDS (Algorithm 1) is very simple to implement: each radio can make its own transmission decision locally. However, it does not guarantee that there are no collisions due to simultaneous transmissions of interfering links. By utilizing the inductive-scheduling technique of Kumar *et al.* [10], we further improve PLDS and develop a collision-free distributed scheme (CFDS) which is presented in Algorithm 2. PLDS cannot avoid collisions because in the second step of PLDS, each radio makes its transmission decision randomly, without any coordination with other radios. To avoid collisions, the transmitter radio ρ of link ℓ needs to know the final transmission decisions for the transmitter radios of ℓ's interfering links. In CFDS, each link exchanges its utilization value with its interfering links during the protocol initialization phase. Therefore, the transmitter radio of link ℓ knows the value of $H(\ell',\psi,t)$ for all $\ell' \in I(\ell)$ based on their utilization value $x(\ell',\psi)$. We also require that for all links $\ell \in \mathcal{L}_{out}(\rho)\bigcup\mathcal{L}_{in}(\rho)$ (*i.e.,* for all links incident on ρ), and for all $\psi\in\Psi$, after computing $H(\ell,\psi,t)$, radio ρ will pick a pair (ℓ,ψ) such that $H(\ell,\psi,t)=1$ using SHA-1 as follows. The input of SHA-1 is $H(\ell,\psi,t)$ for all links incident on ρ and for all $\psi\in\Psi$. Suppose among these link-channel pairs, there are n pairs with $H(\ell,\psi,t)=1$. The output of SHA-1 is x and the maximal output of SHA-1 is y. Let $k=\lceil\frac{n\times(x+1)}{y+1}\rceil$. Radio ρ will choose the kth ($1\leq k\leq n$) link-channel pair with $H(\ell,\psi,t)=1$. Note that for the selected link ℓ, radio ρ can also know the final transmission decisions for the transmitter radios of ℓ's interfering links using the same procedure, because link ℓ knows the utilization value of its interfering links during the protocol setup. After that, ρ can make its proper transmission-decision using the inductive-scheduling technique as described in Algorithm 2.

Using inductive scheduling naively may not completely avoid collision, due to the *asymmetric* nature of Type II secondary interfering links (defined in Section II-D). For example, suppose link ℓ_1 is in the secondary interference set of link ℓ_2 but ℓ_2 is not in the secondary interference set of ℓ_1, and ℓ_2 has higher order than ℓ_1. As a result, ℓ_1 does not yield to ℓ_2 because ℓ_2 is not in its interference set and ℓ_2 does not yield to ℓ_1 due to its higher order. Hence, collisions may occur if both of them transmit. In this case, we force ℓ_2 to backoff, no matter what its inductive order is (compared with its Type II secondary interfering links). Theorem 4 provides a sufficient condition for Algorithm 2 to achieve desired link utilizations. Its proof is very similar to that of Theorem 3.

Algorithm 2 COLLISION-FREE DISTRIBUTED SCHEME (CFDS)(Matrix $\mathbf{X}$)

Require: Each link communicates its utilization value to its interfering links during the protocol initialization phase

Require: Access hash function $H(\ell,\psi,t)$ defined in (3)

Require: Given a fixed triplet (ℓ,ψ,t), every invocation of $H(\ell,\psi,t)$ yields the same result

Require: The random variables $\{H(\ell,\psi,t)\}$ are independent

1: For all links $\ell\in\mathcal{L}_{out}(\rho)\bigcup\mathcal{L}_{in}(\rho)$, and for all $\psi\in\Psi$, compute $H(\ell,\psi,t)$

2: Pick a pair (ℓ,ψ) such that $H(\ell,\psi,t)=1$ using SHA-1, based on the value of $H(\ell,\psi,t)$ for all links $\ell\in\mathcal{L}_{out}(\rho)\bigcup\mathcal{L}_{in}(\rho)$ and for all $\psi\in\Psi$; if no such pair exists, sleep during time t

3: If the selected link $\ell\in\mathcal{L}_{in}(\rho)$, then tune to channel ψ and await an incoming transmission across ℓ on channel ψ at time t

4: Else schedule an outgoing transmission across ℓ on channel ψ at time t, if $\ell\in\mathcal{L}_{out}(\rho)$ and there are no interfering links with higher inductive order which also decide to schedule a transmission on channel ψ

Theorem 4: Let $\epsilon>0$ be an arbitrary constant. Consider an induced radio network $\mathcal{G}=(\mathcal{V},\mathcal{L})$ and its greater-than ordering $\succ$ defined on $\mathcal{L}$; let $\mathbf{X}$ be a utilization matrix which is defined on $\mathcal{G}$. $\mathbf{X}$ can be stably scheduled by Algorithm 2 if:

$$
\begin{aligned}
\forall(\ell,\psi),\ x(\ell,\psi) &+ \sum_{\rho\in\Psi\setminus\{\psi\}} x(\ell,\rho) \\
&+ \sum_{\chi\in\Psi}\sum_{f\in Pri_{\succ}(\ell)} x(f,\chi) \\
&+ \sum_{g\in Sec_{\succ}(\ell)} x(g,\psi) \leq \frac{1}{e}-\epsilon \qquad (5)
\end{aligned}
$$

V. IMPLICATIONS FOR GEOMETRIC INTERFERENCE MODELS

Recall that the interference models which we introduced in §II-B have bounded independence/induction. Thus, Theorem 3 yields the following constant-factor performance guarantees (recall some sample values for Δ and λ from § II-C):

Corollary 5: Let $\epsilon>0$ be an arbitrary constant. Given any interference model that is Δ-independent, Algorithm PLDS yields a $(\Delta+2)\cdot(e+\epsilon)$-competitive ratio for joint channel assignment and link scheduling. Specifically, when the interference model is the uniform $Tx-model$, the $Tx-Rx$ model with parameters 1, 2, and 2.5, the K-hop model on unit disk graphs, and the K-hop model on (r,s)-civilized graphs, algorithm PLDS yields a $7(e+\epsilon)$-factor, $6(e+\epsilon)$-factor, $10(e+\epsilon)$-factor, $14(e+\epsilon)$-factor, $51(e+\epsilon)$-factor, $O(\frac{r^2}{s^2})$-factor competitive ratio respectively, for end-to-end utility maximization with multiple radios and multiple channels.

We omit the proof of this corollary here, but refer interested readers to [10] for a similar proof.

VI. Simulation Results

In this section, we show the feasibility of the proposed distributed and localized scheduling algorithm and evaluate its performance through extensive simulations on random networks. A custom simulator developed in C is used for performance evaluation. There are two main goals of our simulations: (1) study the feasibility of our randomized algorithm using SHA-1 as the *access hash function*, under fPrIM interference model [20]; and (2) compare the probability of events $\mathcal{A}(\ell, \psi, t)$ and $\mathcal{B}(\ell, \psi, t)$ with $x(\ell, \psi)$ for all the link-channel pairs, under fPrIM interference model.

A. Network Settings and Simulation Setup

We assume that each node has two radios as this is the typical setting for most of current IEEE 802.11 mesh network deployments. We also assume that there are 3 available independent channels in the system as it is the case for IEEE 802.11b/g networks (*i.e.*, channels 1, 6 and 11). Note that under this setting, there are totally 24 available link-channel pairs in the induced radio network for an edge (u, v), *i.e.*, 4 links from node u to v and verse vice and each link can operate on 3 different channels. For the random networks, we randomly generate n wireless nodes uniformly in a 500×500 units region. The transmission range for each node is randomly chosen from the set $\{60, 80, 100, 120, 140\}$ units. We perform simulations on random networks with 20, 50 and 100 nodes. Given a network, we first randomly generate its utilization matrix $\mathbf{X}$ which satisfies (4) for all link-channel pairs and then run the simulation for T time slots. To make $\mathbf{X}$ more realistic, we randomly choose $m\%$ of link-channel pairs in $\mathbf{X}$ and offer no traffic on them. Note that in practice, $\mathbf{X}$ is essentially determined by routing algorithms and our scheduling algorithm does not rely on any particular routing algorithms. Therefore, here we use a random $\mathbf{X}$ for simplicity.

The popular hash function SHA-1 is used as the access-hash function H in our current simulation. Suppose all the probabilities are well-approximated by rationals of the form $a/2^b$, where 2^b is some large integer. Then, we choose a random subset R of size b from SHA-1's 160-bit output and compare the new integer c formed by these bits in R with $a = 2^b(1 - e^{-e \cdot x(\ell,\psi)})$ to determine the value of $H(\ell, \psi, t)$ (*i.e.*, $c < a$ means $H(\ell, \psi, t) = 1$). In our simulation, R is pre-distributed to all nodes and we take $b = 40$.

For the randomly generated utilization matrix $\mathbf{X}$, we report the range (i.e, the min and max values) and average values of the sum in the LHS of (4) for all the link-channel pairs. During the generation of $\mathbf{X}$, we try to make the maximal sum as close to $1/e$ as possible. The reason is that if the sums are much smaller than $1/e$, which means the traffic loads are low and thus the contention probabilities are low, we do expect good performance of the proposed randomized algorithm. We report the worst case (minimal) ratios between the probabilities of events $\mathcal{B}(\ell, \psi, t)$, $\mathcal{A}(\ell, \psi, t)$ and the utilization value $x(\ell, \psi)$ for all the link-channel pairs, respectively. For each given topology, we run the simulation three times with different random seeds and report the result with the lowest worst case ratio for event $\mathcal{B}(\ell, \psi, t)$.

Topology	min	max	$avg.$	$\mathcal{B}$	$\mathcal{A}$
R-20 (0.5, 1×10^5)	0.03	0.3607	0.18	1.23	1.39
R-20 (1.0, 1×10^5)	0.03	0.3677	0.18	1.20	1.22
R-20 (1.5, 1×10^5)	0.04	0.3516	0.22	1.10	1.27
R-50 (0.5, 2×10^5)	0.03	0.3677	0.22	1.05	1.13
R-50 (1.0, 2×10^5)	0.09	0.3559	0.21	1.08	1.15
R-50 (1.5, 3×10^5)	0.06	0.3552	0.22	1.00	1.18
R-100 (0.5, 3×10^5)	0.03	0.3614	0.17	1.26	1.28
R-100 (1.0, 3×10^5)	0.02	0.3670	0.18	1.16	1.18
R-100 (1.5, 3×10^5)	0.02	0.3650	0.20	1.22	1.23

Table II
SIMULATION RESULTS FOR FPRIM INTERFERENCE MODEL.

Topology	min	max	$avg.$	$\mathcal{B}$	$\mathcal{A}$
2, 5, 0.95, 2×10^5	0.05	0.3660	0.22	1.02	1.11
2, 8, 0.95, 2×10^5	0.05	0.3643	0.23	1.01	1.13
3, 5, 0.96, 3×10^5	0.05	0.3604	0.22	1.02	1.09
3, 8, 0.97, 3×10^5	0.05	0.3651	0.21	1.02	1.09
4, 5, 0.98, 4×10^5	0.07	0.3531	0.22	1.00	1.05
4, 8, 0.98, 4×10^5	0.05	0.3669	0.21	1.05	1.10
5, 5, 0.99, 5×10^5	0.04	0.3629	0.21	1.01	1.03
5, 8, 0.99, 5×10^5	0.04	0.3639	0.21	1.04	1.10

Table III
SIMULATION RESULTS FOR SDR NETWORKS WITH MORE RADIOS AND CHANNELS.

B. Validation of the Developed Bounds

Table II presents the simulation results for fPrIM interference model. In this table, the first column shows the topology of simulated networks. R-n stands for random network with n nodes. The first parameter in the parenthesis is q. The value of q in fPrIM model determines the interference range of each radio. Since in this model, nodes may have non-uniform transmission ranges and interference ranges, we only conduct performance evaluation on random networks with $q =$ 0.5, 1 and 1.5. The link-channel pair idle probabilities for random networks with 20, 50 and 100 nodes are 0.9, 0.95 and 0.99, respectively. In one of our randomly generated networks with 100 nodes, there are still about 650 active link-channel pairs even if the idle probability is 0.95. The results from Table II demonstrate that our randomized scheduling algorithm works well under the fPrIM model. Note that we also run the simulations on other random networks and for node-exclusive interference model [13] and protocol interference model [6], [10] and got similar results which are omitted due to space limitation. We refer interested readers to [7] for other results.

Table III summarizes the simulation results for SDR networks with 2, 3, 4 and 5 radios and with 5 and 8 channels, respectively. The results are for an example network with 100 nodes, using fPrIM interference model ($q = 1.0$). For the first column, its first parameter is the number of radios; the second parameter is the number of channels; the third parameter is the percentage of idle link-channel pairs; and the last parameter is the number of time slots for each run of the simulation. The results from Table III show that the worst case ratios between the probability of event $\mathcal{B}(\ell, \psi, t)$ and the utilization value are larger than 1 for all these combinations of parameters.

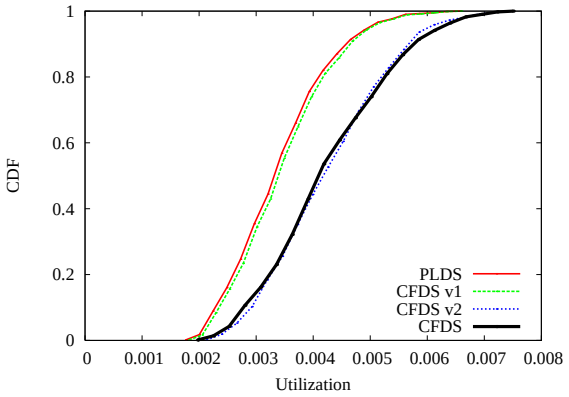

Figure 1. CDF for the ratio of successful transmissions for all the active link-channel pairs.

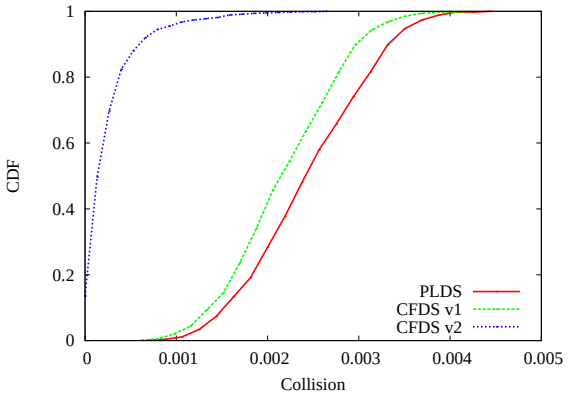

Figure 2. CDF for the ratio of collisions for active link-channel pairs.

C. Comparison Between the Distributed Algorithms

To compare the performance of the PLDS and CFDS algorithms, Figure 1 and Figure 2 show the cumulative distribution function (CDF) for the ratio of successful transmissions and collisions for all the active link-channel pairs, respectively. The results are for an example network with 100 nodes, using fPrIM interference model ($q = 1.5$). The number of time slots for this simulation is 30,000. CFDS v1 and CFDS v2 are two variances of the CFDS algorithm. For the first one, when radio ρ makes transmission decision for link ℓ, it only considers the decisions made by other transmitter radios of the links in ℓ's primary interference set. For the second one, ρ will consider the transmitter radios of the links in ℓ's primary and secondary interference set. These two variances obey the inductive-scheduling technique strictly, *i.e.*, a link will schedule its transmission if there is no links with higher inductive order schedule their transmissions. As we can see from Figure 1, compared with coordination only among primary interfering links, the coordination among both primary and secondary interfering links can improve the ratio of successful transmission. The CFDS v2 algorithm cannot completely avoid collisions which is shown in Figure 2.

VII. Related Work and Conclusion

A large body of research decomposes cross-layer throughput maximization and treats channel assignment and link scheduling *in isolation*. The works [21], [4], [19], [16], [14] assume the availability of a scheduling protocol such as IEEE 802.11 or maximal scheduling, and focus on channel assignment under which the sub-network induced by each channel has desirable properties; [12] assumes the availability of a channel-assignment protocol and proposes link-layer scheduling protocols for MCMR networks. While this modular approach has clear advantages, it is less desirable for throughput maximization: none of the above works analyze how close the throughput region achievable by their schemes is, to the actual region. From this perspective, the most relevant works are [1], [8], [13], which we survey next.

Alicherry, Bhatia and Li [1] study the joint channel assignment and link scheduling problem under the uniform RTS-CTS model with parameter q; they proposed a centralized algorithm under the assumption that the network is homogeneous: *i.e.*, each node has the same number of radios ν, and each radio has the same set of κ channels, and for a given link, each channel has the same capacity. For q equal to 1, 2 and 2.5, they prove that the throughput region yielded by their techniques is at most a factor of $\frac{4\kappa}{\nu}$, $\frac{8\kappa}{\nu}$ and $\frac{12\kappa}{\nu}$ respectively (these are the respective competitive factors for their scheme). In contrast, our algorithmic results are derived under a generic model of interference, and the performance guarantees we present are in terms of λ and Δ, which are properties specific to a given interference model. Our algorithms and guarantees also apply for arbitrary heterogeneous networks. Most significantly, we present two distributed schemes whose guarantees improve upon those of [1]; for the special-case of the uniform RTS-CTS model studied by [1], for values of q equal to 1, 2, and 2.5 respectively, our distributed scheme yields competitive factors of $6e$, $10e$, and $14e$ respectively. *These factors are independent of any parameter determined by the network topology*.

Kodialam and Nandagopal [8] propose two centralized heuristics – a greedy heuristic, and a packing based heuristic for the joint channel allocation and scheduling problem. Their schemes are applicable to arbitrary link-conflict based interference models. However, they do not present any guarantees for the competitiveness of their algorithms. Indeed, it is possible to construct a family of geometric network topologies where the admissible throughput region of their schemes are a factor of $\Omega(n)$ away from the optimal joint channel assignment and scheduling. Since n is the number of network nodes, the competitiveness of their scheme is unbounded – even under geometric models of network interference. In contrast, we present two distributed schemes with provably-good competitive guarantees; this translates to constant-factor competitiveness for several geometric models of interference.

The work of Lin and Rasool [13] is most similar in spirit to our results. They present a "distributed" algorithm for joint channel assignment and link scheduling whose competitive ratio is guaranteed to be at most $\Delta + 2$, where Δ is the independence number. In their scheme, each "link" in the network makes the channel assignment and scheduling decisions for *each time slot* by examining the instantaneous queue length of the links in its interference set during the slot. We contend that, while the scheme of Lin and Rasool is amenable to distributed

implementation, it is not fully distributed in its present form. A "link" in a wireless network is a logical entity; *assigning* a channel to a link implies that the sender of the link decides to transmit on the channel, and the receiver simultaneously decides to listen on that channel. The packet queues for the link are maintained only at the sending end-point of the link. Thus, in order for the receiver to make channel assignment and scheduling decisions based on queue sizes, the sender and receiver for each network link need to exchange this information prior to each transmission. Exchanging queue-size information prior to each transmission is a significant overhead; how this information exchange is performed and how this affects the competitive ratio are critical issues that are left unaddressed in [13]. While it is in principle possible to listen in on neighbors' "current state" information in every step of the protocol, this poses challenges. *First, the frequency at which a node broadcasts this information should be available at all of its neighbors, a necessary condition for which is that any pair of distance-2 nodes share a channel – a condition may not hold. Second, even if this condition holds, a node may have to listen simultaneously on many different channels to get its neighbors' queue-state information, which may be infeasible.* In contrast, observe that our protocol requires essentially no such coordination after the initial discovery of distance-two neighborhoods. The assumptions made in our distributed algorithm come with far less overhead: the end-points of each link only need to know the *long-term rate* which needs to be sustained by the link as opposed to instantaneous queue sizes. In our scheme, the end-points of a link make channel assignment decisions through the use of access-hash functions. This completely eliminates the need for information exchange on a per-transmission basis, and thus makes our channel assignment and scheduling strategy truly distributed.

Thus, we have developed two novel provably-good (distributed) algorithms for channel allocation and scheduling in SDR networks. In the first (PLDS), each radio makes its transmission decision purely locally and does not need to exchange any information with links in its interference set. The second (CFDS) is collision-free through the use of inductive scheduling. Each radio only needs to exchange the utilization values between interfering links during the protocol setup phase. After this phase, the radios can make their decisions locally for each time slot and do not need to exchange information with interfering links any more. Simulation results show that our bounds are valid in practice and that the second distributed algorithm can enable smart backoff-decisions to avoid unnecessary transmissions, which can save valuable energy. We plan to investigate further applications of our access hash function methodology in the future.

Acknowledgment The authors wish to thank the anonymous reviewers for their insightful comments and Jonathan Katz for useful discussions. Aravind Srinivasan and Bo Han were supported in part by NSF ITR Award CNS-0426683 and NSF award NeTS-NBD 0626636. Madhav V. Marathe and V. S. Anil Kumar were supported in part by NSF NeTS Grant CNS-0626964, NSF HSD Grant SES-0729441, CDC Center of Excellence in Public Health Informatics Grant 2506055-01, NIH-NIGMS MIDAS project 5 U01 GM070694-05, DTRA CNIMS Grant HDTRA1-07-C-0113 and NSF NeTS Grant CNS-0831633.

REFERENCES

[1] M. Alicherry, R. Bhatia, and L. E. Li. Joint channel assignment and routing for throughput optimization in multi-radio wireless mesh networks. In *Proceedings of MobiCom 2005*, pages 58–72, 2005.

[2] B. Awerbuch and R. Khandekar. Greedy distributed optimization of multi-commodity flows. In *Proceedings of PODC*, 2007.

[3] R. Bruno, M. Conti, and E. Gregori. Optimal capacity of p-persistent CSMA protocols. *IEEE Commun. Lett.*, 7:139–141, 2003.

[4] A. Brzezinski, G. Zussman, and E. Modiano. Enabling distributed throughput maximization in wireless mesh networks: a partitioning approach. In *Proceedings of MobiCom 2006*, pages 26–37, 2006.

[5] N. Garg and J. Koenemann. Faster and simpler algorithms for multi-commodity flow and other fractional packing problems. In *Proceedings of FOCS 1998*, pages 300–309, 1998.

[6] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388–404, 2000.

[7] B. Han, V. S. A. Kumar, M. V. Marathe, S. Parthasarathy, and A. Srinivasan. Distributed strategies for channel allocation and scheduling in software-defined radio networks. Technical report, Univ. of Maryland, 2008. Available at http://hdl.handle.net/1903/8717.

[8] M. Kodialam and T. Nandagopal. Characterizing the capacity region in multi-radio multi-channel wireless mesh networks. In *Proceedings of MobiCom 2005*, pages 73–87, 2005.

[9] V. S. A. Kumar, M. V. Marathe, S. Parthasarathy, and A. Srinivasan. End-to-end packet-scheduling in wireless ad-hoc networks. In *Proceedings of SODA 2004*, pages 1021–1030, 2004.

[10] V. S. A. Kumar, M. V. Marathe, S. Parthasarathy, and A. Srinivasan. Algorithmic aspects of capacity in wireless networks. In *Proceedings of SIGMETRICS 2005*, pages 133–144, 2005.

[11] V. S. A. Kumar, M. V. Marathe, S. Parthasarathy, and A. Srinivasan. Provable algorithms for joint optimization of transport, routing and MAC layers in wireless ad hoc networks. In *Proceedings of Dial M-POMC 2007*, 2007.

[12] P. Kyasanur and N. H. Vaidya. Routing and link-layer protocols for multi-channel multi-interface ad hoc wireless networks. *SIGMOBILE Mob. Comput. Commun. Rev.*, 10(1):31–43, 2006.

[13] X. Lin and S. Rasool. A distributed joint channel-assignment, scheduling and routing algorithm for multi-channel ad-hoc wireless networks. In *Proceedings of IEEE INFOCOM 2007*, pages 1118–1126, 2007.

[14] K. N. Ramachandran, E. M. Belding-Royer, K. C. Almeroth, and M. M. Buddhikot. Interference-aware channel assignment in multi-radio wireless mesh networks. In *Proceedings of IEEE INFOCOM 2006*, pages 1–12, 2006.

[15] S. Ramanathan and E. L. Lloyd. Scheduling algorithms for multihop radio networks. *IEEE-ACM Transactions on Networking (ToN)*, 1:166–177, 1993.

[16] A. Raniwala, K. Gopalan, and T.-C. Chiueh. Centralized channel assignment and routing algorithms for multi-channel wireless mesh networks. *SIGMOBILE Mob. Comput. Commun. Rev.*, 8(2):50–65, 2004.

[17] G. Sharma, R. R. Mazumdar, and N. B. Shroff. On the complexity of scheduling in wireless networks. In *Proceedings of MobiCom 2006*, pages 227–238, 2006.

[18] Y. Shi and T. Hou. Optimal power control for multi-hop software defined radio networks. In *Proceedings of IEEE INFOCOM*, 2007.

[19] A. P. Subramanian, H. Gupta, and S. R. Das. Minimum interference channel assignment in multi-radio wireless mesh networks. In *Proceedings of SECON 2007*, pages 481–490, 2007.

[20] W. Wang, X.-Y. Li, O. Frieder, Y. Wang, and W.-Z. Song. Efficient interference-aware TDMA link scheduling for static wireless networks. In *Proceedings of MobiCom 2006*, pages 262–273, 2006.

[21] K. Xing, X. Cheng, L. Ma, and Q. Liang. Superimposed code based channel assignment in multi-radio multi-channel wireless mesh networks. In *Proceedings of MobiCom 2007*, pages 15–26, 2007.

[22] S. Yi, Y. Pei, and S. Kalyanaraman. On the capacity improvement of ad hoc wireless networks using directional antennas. In *Proceedings of MobiHoc 2003*, pages 108–116, 2003.