Wireless Netw (2013) 19:1121–1133
DOI 10.1007/s11276-012-0522-4

Broadcasting in multi-radio multi-channel wireless networks using simplicial complexes

Wei Ren · Qing Zhao · Ram Ramanathan · Jianhang Gao · Ananthram Swami · Amotz Bar-Noy · Matthew P. Johnson · Prithwish Basu

Published online: 27 November 2012

Abstract We consider the broadcasting problem in multi-radio multi-channel ad hoc networks. The objective is to minimize the total cost of the network-wide broadcast, where the cost can be of any form that is summable over all the transmissions (e.g., the transmission and reception energy, the price for accessing a specific channel). Our technical approach is based on a simplicial complex model that allows us to capture the broadcast nature of the wireless medium and the heterogeneity across radios and channels. Specifically, we show that broadcasting in multi-radio multi-channel ad hoc networks can be formulated as a minimum spanning problem in simplicial complexes. We establish the NP-completeness of the minimum spanning problem and propose two approximation algorithms with order-optimal performance guarantee. The first approximation algorithm converts the minimum spanning problem in simplical complexes to a minimum connected set cover (MCSC) problem. The second algorithm converts it to a node-weighted Steiner tree problem under the classic graph model. These two algorithms offer tradeoffs between performance and time-complexity. In a broader context, this work appears to be the first that studies the minimum spanning problem in simplicial complexes and weighted MCSC problem.

Keywords Broadcast · Multi-radio multi-channel · Ad hoc network · Simplicial complex · Minimum spanning · Minimum connected set cover

Part of this work was presented at IEEE MASS 2011.

W. Ren (✉)
Microsoft Corporation, Redmond, WA, USA
e-mail: wren@microsoft.com

Q. Zhao · J. Gao
Department of Electrical and Computer Engineering, University of California, Davis, CA, USA
e-mail: qzhao@ucdavis.edu

R. Ramanathan · P. Basu
Raytheon BBN Technologies, Cambridge, MA, USA

A. Swami
Army Research Laboratory, Adelphi, MD, USA

A. Bar-Noy
Department of Computer Science, City University of New York, New York City, NY, USA

M. P. Johnson
Department of Computer Science and Engineering, Pennsylvania State University, Philadelphia, PA, USA

1 Introduction

Multi-radio multi-channel (MR-MC) wireless networking arises in the context of wireless mesh networks, dynamic spectrum access via cognitive radio, and next-generation cellular networks [12]. By the use of multiple channels, spatially adjacent transmissions can be carried over non-overlapping channels to avoid mutual interference. Furthermore, each node, equipped with multiple radios, is capable of working in a full-duplex mode by tuning the transmitting and receiving radios to two non-overlapping channels.

The increasing demand for high data rate and the persistent reduction in radio costs have greatly stimulated research on MR-MC networks. Considerable work has been done on capacity analysis, channel and radio assignment [11, 17, 19, 22], and routing protocols [11, 19]. In this paper, we consider the broadcasting problem in MR-MC ad hoc networks.

1.1 Broadcasting in single-radio single-channel networks

Broadcasting is a basic operation in wireless networks for disseminating a message containing, for example, situation awareness data and routing control information, to all nodes. For a single-radio single-channel (SR-SC) network, a key question for the network-wide broadcast is: which set of nodes should be selected to transmit such that the total cost (such as energy consumption or the number of transmissions) is minimized. In contrast to the wireless broadcast problems for minimizing the energy consumption and the number of transmissions which are shown to be NP-complete in [2], their counterparts in wired networks have polynomial solutions. A heuristic algorithm, referred to as broadcast incremental power (BIP) algorithm, is proposed in [23].

The complexity of the problem arises from the broadcast nature of the wireless medium: a single transmission from one node can reach all the other nodes within the transmission range of this node, but it may cause interference to other nearby transmissions. This "node-centric" nature of the wireless broadcasting problem along with the mutual interference between concurrent transmissions complicates the design of efficient broadcasting algorithms.

1.2 Broadcasting in MR-MC networks

In an MR-MC ad hoc network, such as the DARPA wireless network after next (WNaN) [20], each node is equipped with multiple radios, each operating on a different channel. The introduction of multiple channels and multiple radios further complicates the design of an efficient broadcasting scheme. Since the number of radios at a node is usually smaller than the number of channels, the broadcast scheme should decide not only which nodes act as relays but also for those relay nodes, which channel(s) should be assigned to the transmitting radio(s). Given the selection of the relay nodes, two simple broadcast schemes are: (1) transmitting multiple copies of the message on all channels, (2) transmitting a single copy of the message on a common channel dedicated to broadcasting. Both schemes are inefficient. For the latter one, if the broadcast load is high, the common channel will be overwhelmed, even while plenty of other channels are free.

One subtle issue is the complication of the wireless broadcast advantage. In an MR-MC network, if the radios of the neighboring nodes are tuned to different channels, a single transmission on one channel cannot reach all the neighboring nodes simultaneously. In other words, only the neighboring nodes on the same channel can share the wireless broadcast advantage. More precisely, the concept of neighborhood must be defined both by radio range and channel. Another subtle issue is channel heterogeneity.

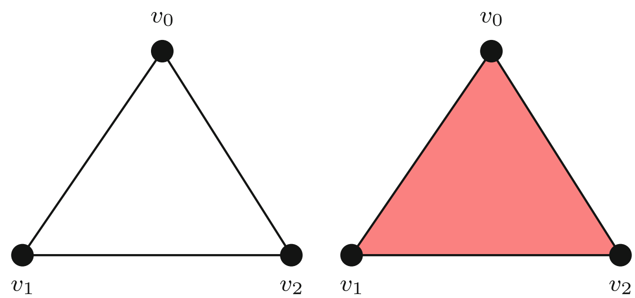

(a) Graph **(b)** Simplicial Complex

Fig. 1 *Graph* and simplicial complex: $V = \{v_0, v_1, v_2\}$, $\mathcal{S}_{(a)} = \{(v_0, v_1), (v_0, v_2), (v_1, v_2)\}$, $\mathcal{S}_{(b)} = \{(v_0, v_1, v_2), (v_0, v_1), (v_0, v_2), (v_1, v_2), (v_0), (v_1), (v_2)\}$

Channels may have different bandwidth, fading condition, and accessing cost, leading to different implications for the total broadcast cost.

Broadcasting in MR-MC networks is thus a multi-faceted problem, involving channel assignment, relay node selection, and channel selection for the source and relay nodes. In this paper, we focus on the latter two issues by assuming a given channel-to-radio assignment. To avoid the hidden channel problem [17], two nodes that are two-hops away from each other are assigned two distinct sets of channels. Our design objective is to minimize the total broadcast cost, where the cost can be of any form that is summable over all the transmissions, including, for example, the transmission and reception energy,[1] the price for accessing each channel.

1.3 A simplicial complex model for broadcasting in MR-MC networks

Our technical approach is based on a simplicial complex model of the broadcasting problem in MR-MC networks. A simplicial complex is a collection of nonempty sets with finite size that is closed under the subset operation. In other words, if a set s belongs to the collection, all subsets of s also belongs to the collection. An element of the collection is called a *simplex* or *face*. This subset constraint is often satisfied in the network context. For example, subsets of a broadcast/multicast group are broadcast/multicast groups, subsets of a clique are cliques. A simple example of graph and simplicial complex is given in Fig. 1. While the concept of simplicial complex has been around since the 1920s, many well-solved fundamental problems in graph remain largely open under this more general model.

We use a simplicial complex model rather than a graph because the simplicial complex more naturally captures the broadcast channel, and the distinction and disjointness between broadcasting on different channels. For instance, it

[1] The 'reception energy' denotes the energy consumed by the radio in reception mode.

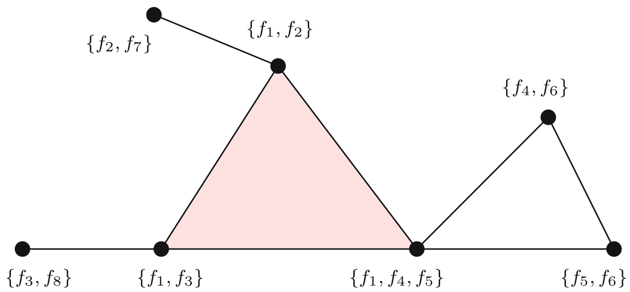

Fig. 2 An illustration of an MR-MC network and the constructed simplicial complex. The parameters within the *braces* are the channels which each node can access. In the communication *graph* derived from the network, a link exists between two nodes if and only if two nodes are within each other's transmission range and they share at least one common channel. Notice that a clique in the communication graph may not be a clique in the MR-MC network (correspondingly, a simplex in the simplicial complex), e.g., the three nodes of the *right empty triangle* (they do not share a common channel)

is not possible to discern by only looking at the graph in Fig. 1(a) whether v_0, v_1, and v_2 can communicate simultaneously over a shared broadcast channel or if they have dedicated channels and so each can only talk to one other node at a time. Further, costs can be attached to faces (simplices) in a way not easily possible with graphs. A more detailed discussion on the motivation of adopting the simplicial complex for the broadcast problem in multi-hop wireless networks is provided in [18].

Consider an example MR-MC network where nodes are all stationary. As shown in Fig. 2, after the channels are assigned, the network is partitioned into cliques of nodes. Such clique assignment scheme is implemented in the DARPA WNaN system and was used in an exercise at Ft. Devens [20]. A clique consists of the nodes which are within each other's transmission range and share at least one common channel, and two cliques are spliced (i.e., connected) via nodes operating on multiple channels shared in common by the two cliques. Within each clique, depending on the cost function, the transmitter decides which dimension simplex (i.e., a subclique or the clique itself) in a clique complex to activate. The message for the network-wide broadcast is thus propagated through a sequence of cliques, possibly of different dimensions. Note that the unicast case corresponds to a clique of dimension 1 (an edge). This example could also apply to the case where nodes may have multiple radios, perhaps of different modality (e.g., RF and optical); in this case, there may also be a cost associated with switching modes.

The network-wide broadcast problem can be formulated as the minimum spanning problem in simplicial complexes. A clique in the MR-MC network is modeled as a simplex in the simplicial complex (see Fig. 2), and since a subset of a clique is still a clique, the constructed simplicial complex meets the requirement of being closed under the subset operation. The minimum spanning problem in a simplicial complex is to find a connected subset of simplices that covers all the vertices with the minimum total weight, i.e., the minimum connected spanning subcomplex (MCSSub).[2] Then the solution to the network-wide broadcast problem can be obtained by solving the MCSSub problem.

1.4 Minimum spanning problem in simplicial complexes

The minimum spanning problem in a graph is to find a connected subgraph that covers all the vertices with minimum total weight. The solution must be a tree for graphs with nonnegative weights [hence called the minimum spanning tree (MST)]. There are several polynomial-time algorithms for MST, e.g., Kruskal's algorithm and Prim's algorithm [10]. The MST problem has many applications in network planning, broadcasting in communication networks, touring problems, and VLSI design [1].

With the addition of high dimensional simplices, the minimum spanning problem in a simplicial complex is fundamentally different and much more difficult than its counterpart in a graph. First, unlike the case in a graph, the MCSSub of a simplicial complex may not be a "tree".[3] As illustrated in Fig. 3, the MCSSub of the simplicial complex is the three filled triangles which form a cycle. Second, while simple greedy-type polynomial-time algorithms exist for finding the MST in a graph, the minimum spanning problem in a simplicial complex is NP-complete as established in this paper (see Sect. 3.1).

We develop polynomial-time approximation algorithms for the minimum spanning problem in simplicial complexes. We propose two algorithms: one reduces this problem to a " (MCSC) problem, and the other reduces the problem to a node-weighted Steiner tree problem in a graph derived from the original simplicial complex. We also establish the approximation ratios of the two algorithms. Both are shown to be order-optimal. The time-complexity of these two algorithms is also analyzed, illustrating the tradeoff between performance and complexity offered by these two algorithms. In a broader context, this work appears to be the first that studies the minimum spanning problem in simplicial complexes and weighted MCSC problem.

[2] Strictly speaking, a subcomplex should also be closed under the subset operation, but without loss of generality, we do not include this condition in the definition of minimum connected spanning subcomplex, which is also more relevant to the broadcasting problem at hand.

[3] Although there is no unified definition of tree in simplicial complexes, a couple of definitions can be obtained by generalizing those equivalent definitions of tree in a graph. For example, simplicial trees can be defined based on the universal existence of leaves in any subgraph, or the uniqueness of simplicial facet paths.

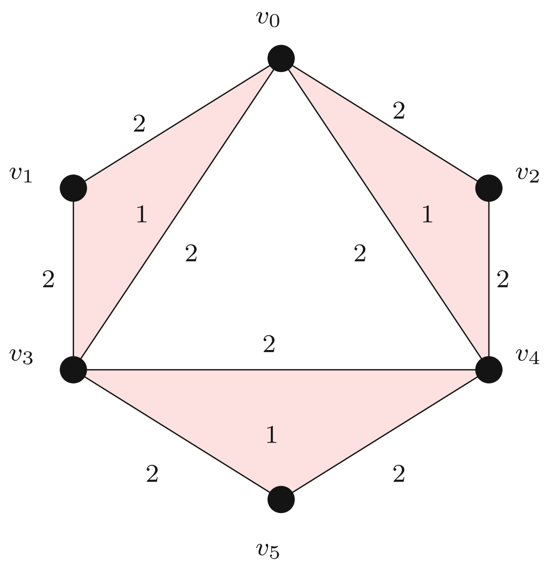

Fig. 3 A simplicial complex whose MCSSub is the *three filled triangles*, and is not a "tree" (the integers are the weights of the simplices)

1.5 Related work

Broadcasting in MR-MC networks, mostly in the context of wireless mesh networks, has been studied for different optimization objectives (see [3, 4, 13, 17] and references therein). Different from the previous work, the optimization objective in our work can be any cost function which is summable over all the transmissions, thus taking into account channel heterogeneity (e.g., transmissions on different channels may consume different amounts of energy, due to different bandwidths or different propagation characteristics or some other factor). We point out that neither minimizing the total number of transmissions nor minimizing the total number of radios used in the broadcast is, in general, equivalent to minimizing the total energy consumption. The reception energy is ignored if the total number of transmissions is minimized, while the transmission energy and the reception energy are equated if the total number of radios is minimized. More importantly, channel heterogeneity is not addressed if these two objectives are optimized.

Furthermore, to our best knowledge, our work is the first to adopt simplicial complexes to model and solve the broadcast problem in wireless ad hoc networks. For a more detailed discussion on the potential applications of simplicial complexes in communication and social networks, readers are referred to [18].

2 Basic concepts in simplicial complexes

In this section, we introduce several basic concepts in simplicial complexes [16].

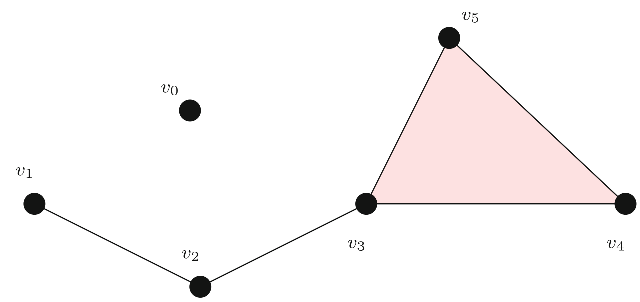

Fig. 4 A simplicial complex Δ with 6 vertices (0-dimensional simplices: $\{v_0\}$, $\{v_1\}$, $\{v_2\}$, $\{v_3\}$, $\{v_4\}$, $\{v_5\}$), 5 edges (1-dimensional simplices: $\{v_1, v_2\}$, $\{v_2, v_3\}$, $\{v_3, v_4\}$, $\{v_3, v_5\}$, $\{v_4, v_5\}$), and 1 filled triangle (2-dimensional simplex: $\{v_3, v_4, v_5\}$)

An *(abstract) simplicial complex* is a collection Δ of nonempty sets with finite size such that if $A \in \Delta$, then $\forall B \subseteq A$, $B \in \Delta$, i.e., Δ is closed under the operation of taking subsets. The element A of Δ is called a *simplex* of Δ; its *dimension* (denoted by $\dim A$) is one less than the number of its elements. Each nonempty subset of A is called a *face*[4] of A. The *dimension* of Δ is the maximum dimension over all its simplices, or is infinite if the maximum does not exist. The *vertex set* V of Δ is the union of the one-point elements of Δ. Figure 4 shows an example of a 2-dimensional simplicial complex. A subcollection of Δ that is itself a simplicial complex is called a *subcomplex* of Δ. A subcomplex of Δ is the *p-skeleton* of Δ, denoted by $\Delta^{(p)}$, if it is the collection of all simplices of Δ with dimension no larger than p. Thus, the 1-skeleton is the underlying graph of Δ.

A *facet* of a simplicial complex Δ is a maximal face of Δ, i.e., it is not a subset of any other face. A simplicial complex is *connected* if its 1-skeleton (i.e., the underlying graph) is connected in the graph sense.

A weighted simplicial complex (WSC) Δ is a triple $(V, \mathcal{S}, w)$[5], where V is the set of vertices, $\mathcal{S}$ the set of faces of Δ, and $w : \mathcal{S} \to \{\mathbb{R}^+ \cup \{0\}\}$ a nonnegative weight function defined for each face in $\mathcal{S}$ with $w(v) = 0$ for all $v \in V$. We define the *facet-only weight* $W_F(\Delta)$ of a WSC Δ as

$$W_F(\Delta) = \sum_{F_i \in \{\text{facet of } \Delta\}} w(F_i).$$

3 Minimum connected spanning subcomplex

In this section, we show that the MCSSub problem is NP-complete, and we propose two approximation algorithms based on connected set cover and node-weighted Steiner tree. We also establish the approximation ratios of the two algorithms and analyze their time complexity.

[4] Notice that since A is its own nonempty subset, the simplex A is also a face of A.

[5] $(\mathcal{S}, w)$ suffices to denote the WSC since $V \subseteq \mathcal{S}$, but we use the redundant $(V, \mathcal{S}, w)$ for convenience.

3.1 NP-completeness

The decision version (D-MCSSub) of the MCSSub problem is stated as follows: let $V(\Delta)$ denote the vertex set of a WSC Δ and $W_F(\Delta)$ the facet-only weight of Δ. Given a WSC $\Delta = (V, \mathcal{S}, w)$ and $K > 0$, is there a connected subcomplex Δ^{sub} of Δ such that $V(\Delta^{sub}) = V$ and $W_F(\Delta^{sub}) \leq K$? Then we have the following theorem.

Theorem 1 *The D-MCSSub problem is NP-complete.*

To prove the NP-completeness, we reduce a classic NP-complete problem—the unweighted set cover problem to the MCSSub problem.

Proof To check a solution to the D-MCSSub problem, we only need to verify the following points: (1) compute the facet-only weight of the solution and compare the weight with K; (2) check whether the underlying graph of the solution is connected or not; (3) check whether all the vertices of the original simplicial complex are covered by the solution. Since all these can be done within polynomial time, the D-MCSSub problem is NP.

Given is an unweighted set cover instance I, i.e., a universe of elements U and a family of subsets F of U. For each element $u \in U$, we introduce a corresponding vertex in MCSSub instance I'. We introduce one additional vertex d. For each set $f \in F$, we introduce a corresponding face $f' = f \cup \{d\}$, of weight 1. Being a simplcial complex, all subsets of f' are also introduced, all also of weight 1. In addition to covering all vertices, a solution to I' must be connected.

Given any solution SOL to I of cost c, we can construct a solution to I' also of cost c. For every set f in SOL not including d, replace f with $f \cup \{d\}$.

On the other hand, given any solution SOL' to I' of cost c, a solution to I of the same cost can be constructed as follows: for any face $f \in SOL'$ such that f does not appear as a set in F, replace f with *any* superset of $f - \{d\}$ appearing in F. Note that there must exist at least one such superset. □

This proof is for general D-MCSSub problems. It can be shown that even if the weight function of the WSC is monotone or strictly monotone,[6] the D-MCSSub problem is NP-hard. But it is still possible that the D-MCSSub problem under some special structured weight function is P.

In the following, we present two approximation algorithms for the MCSSub problem both with performance guarantee $O(\ln n)$, where n is the number of vertices in the WSC. Since the best possible approximation ratio for the set cover problem is $\ln n$ [6], these two algorithms are order-optimal. Note that to achieve the best possible approximation ratio $O(\ln n)$, the algorithm based on connected set cover requires that the ratio R_w of the maximum weight to the minimum weight be bounded. As shown in Theorem 4, if R_w is unbounded, then the approximation ratio can be $\Omega(n)$.

[6] We say that the weight function satisfies the *monotone* property if for any two faces $S_1 \subseteq S_2$, $w(S_1) \leq w(S_2)$, i.e., the weight is monotone non-decreasing with respect to the dimension of the face.

3.2 Algorithm based on connected set cover

Let A be a set with finite number of elements, and $\mathcal{B} = \{B_i \subseteq A : i = 1, \ldots, n\}$ a collection of subsets of A where each B_i is associated with a weight $w(B_i) \geq 0$. Let G be a connected graph with the vertex set $\mathcal{B}$. A connected set cover (CSC) $\mathcal{S}_C$ with respect to $(A, \mathcal{B}, w, G)$ is a set cover of A such that $\mathcal{S}_C$ induces a connected subgraph of G. The MCSC problem is to find the CSC with the minimum weight, where the weight of a CSC $\mathcal{S}_C$ is defined as

$$w(\mathcal{S}_C) = \sum_{B_i \in \mathcal{S}_C} w(B_i).$$

From a WSC $\Delta = (V, \mathcal{S}, w)$, we derive an auxiliary undirected graph G_Δ in the following way: let $\mathcal{S} \setminus V$ be the vertex set of G_Δ, and connect two vertices (non-vertex faces in Δ) S_1 and S_2 if and only if $S_1 \cap S_2 \neq \emptyset$ (i.e., S_1 and S_2 have at least one element of V in common). Then we have the following theorem on the relation between the MCSSub problem and the MCSC problem.

Theorem 2 *Let Δ^* be the MCSSub of a WSC $\Delta = (V, \mathcal{S}, w)$ and $\mathcal{S}_C^*$ the MCSC of $(V, \mathcal{S} \setminus V, w, G_\Delta)$. Then we have*

$$w_F(\Delta^*) = w(\mathcal{S}_C^*).$$

Proof The proof is based on the following lemma. □

Lemma 1 *Let $\mathcal{S}_C^*$ be the MCSC of $(V, \mathcal{S} \setminus V, w, G_\Delta)$. For any face $S \in \mathcal{S}_C^*$ with $w(S) > 0$, we have that there does not exist a face $S' \in \mathcal{S}_C^*$ such that $S \subset S'$.*

Proof *(Proof of Lemma 1)* Suppose that for some face $S \in \mathcal{S}_C^*$ with $w(S) > 0$, $\exists\, S' \in \mathcal{S}_C^*$ such that $S \subset S'$. Let $\mathcal{S}_C' = \mathcal{S}_C^* \setminus S$. Obviously, $\mathcal{S}_C'$ is a set cover, and $w(\mathcal{S}_C') = w(\mathcal{S}_C^*) - w(S) < w(\mathcal{S}_C^*)$. On the other hand, since $S \cap S'' \neq \emptyset$ implies $S' \cap S'' \neq \emptyset$ for any face $S'' \in \mathcal{S}_C^*$, it follows from the connection rule of the auxiliary graph G_Δ that any path via S has an alternative path via S'. Thus, $\mathcal{S}_C'$ is a CSC, leading to a contradiction.

Given the MCSC $\mathcal{S}_C^*$ of $(V, \mathcal{S} \setminus V, w, G_\Delta)$, we can obtain a connected spanning subcomplex Δ_C^* by mapping each element of $\mathcal{S}_C^*$ to a face in Δ. Since the facet-only weight $w_F(\Delta_C^*)$ of Δ_C^* only counts facets in Δ_C^*, it follows that $w_F(\Delta_C^*) \leq w(\mathcal{S}_C^*)$. Based on Lemma 1, we have that every element of $\mathcal{S}_C^*$ with positive weight is a facet in Δ_C^*, and thus

$$w_F(\Delta^*) \leq w_F(\Delta_C^*) = w(\mathcal{S}_C^*),$$

where Δ^* is the MCSSub of Δ.

On the other hand, the facets of Δ^* leads to an CSC $\mathcal{S}_\Delta^*$, and $w(\mathcal{S}_\Delta^*) = w_F(\Delta^*)$. It implies that

$w(\mathcal{S}_C^*) \leq w(\mathcal{S}_\Delta^*) = w_F(\Delta^*)$.

Thus, $w_F(\Delta^*) = w(\mathcal{S}_C^*)$. □

3.2.1 Algorithm

Based on Theorem 2, we can reduce the MCSSub problem of a WSC $\Delta = (V, \mathcal{S}, w)$ to the MCSC problem $(V, \mathcal{S} \setminus V, w, G_\Delta)$. We obtain the following set cover based algorithm (SCA) for the MCSSub problem.

Zhang et al. [24] propose a greedy approximation algorithm for the unweighted MCSC problem, i.e., $w(B_i) = 1$ for all i. The original algorithm in [24] has a flaw and the established approximation ratio is incorrect. In [21], the flaw is corrected and a stronger result on the approximation ratio is shown. By generalizing their greedy approach, we develop a greedy algorithm for the weighted MCSC problem.

Before stating the algorithm, we introduce the following notations and definitions. For two sets $S_1, S_2 \in \mathcal{S}$, let $\text{dist}_G(S_1, S_2)$ be the length of the shortest path between S_1 and S_2 in an auxiliary graph G, where the length of a path is given by the number of edges; S_1 and S_2 are said to be *graph-adjacent* if they are connected via an edge in G (*i.e.*, $\text{dist}_G(S_1, S_2) = 1$), and they are said to be *cover-adjacent* if $S_1 \cap S_2 \neq \emptyset$. Notice that in a general MCSC problem, there is no connection between these two types of adjacency. The *cover-diameter* $D_C(G)$ is defined as the maximum distance between any two cover-adjacent sets, i.e.,

$$D_C(G) = \max\{\text{dist}_G(S_1, S_2) | S_1, S_2 \in \mathcal{S} \quad \text{and} \quad S_1 \cap S_2 \neq \emptyset\}.$$

For the MCSC problem derived from the MCSSub problem of a WSC Δ, we have that $D_C(G_\Delta) = 1$.

At each step of the algorithm, let $\mathcal{R}$ denote the collection of the subsets (faces of Δ) that have been selected, and U the vertex subset of Δ that has been covered. Given $\mathcal{R} \neq \emptyset$ and a set $S \in \mathcal{S} \setminus \mathcal{R}$, an $\mathcal{R} \to S$ path is a path $\{S_0, S_1, \ldots, S_k\}$ in G such that (1) $S_0 \in \mathcal{R}$; (2) $S_k = S$; (3) $S_1, \ldots, S_k \in \mathcal{S} \setminus \mathcal{R}$. We define the weight ratio $r(P_S)$ of P_S as

$$r(P_S) = \frac{w(\mathcal{S}(P_S) \setminus \mathcal{R})}{|V_N(P_S)|} = \frac{\sum_{S \in \mathcal{S}(P_S) \setminus \mathcal{R}} w(S)}{|V_N(P_S)|}, \tag{1}$$

where $\mathcal{S}(P_S) \setminus \mathcal{R}$ is the subsets (faces in $\mathcal{S}$) of P_S that are not in $\mathcal{R}$, and $|V_N(P_S)|$ is the number of vertices of Δ that are covered by P_S but not covered by $\mathcal{R}$.

Algorithm 1 *SCA for MCSSub:*
INPUT: A WSC $\Delta = (V, \mathcal{S}, w)$.
OUTPUT: An approximate MCSSub Δ_C *of* Δ.

1. *Derive the auxiliary graph* G_Δ.
2. *Find an approximate MCSC* $\mathcal{S}_C$ *of* $(V, \mathcal{S} \setminus V, w, G_\Delta)$ *by using the greedy algorithm for MCSC (Algorithm 2).*
3. *Transform* $\mathcal{S}_C$ *to a connected spanning subcomplex* Δ_C *by mapping each element of* $\mathcal{S}_C$ *to a face in* Δ.

Algorithm 2 *A Greedy Algorithm for MCSC.*
INPUT: $(V, \mathcal{S} \setminus V, w, G_\Delta)$
OUTPUT: A CSC $\mathcal{R}$.

1. *Choose* $S_0 \in \mathcal{S} \setminus V$ *such that the weight ratio* $r(S_0)$ defined in (1) *is the minimum, and let* $\mathcal{R} = \{S_0\}$ *and* $U = S_0$.
2. ***WHILE*** $V \setminus U \neq \emptyset$ ***DO***
   - 2.1. *For each* $S \in \mathcal{S} \setminus (V \cup \mathcal{R})$ *which is cover-adjacent or graph-adjacent with a set in* $\mathcal{R}$*, find a shortest*[7] $\mathcal{R} \to S$ *path* P_S.
   - 2.2. *Select* P_S *with the minimum weight ratio* $r(P_S)$ defined in (1)*, and let* $\mathcal{R} = \mathcal{R} \cup P_S$ *(add all the subsets of* P_S *to* $\mathcal{R}$*) and* $U = U \cup V_N(P_S)$.

   END WHILE
3. ***RETURN*** $\mathcal{R}$.

3.2.2 Approximation ratio

The approximation ratio of SCA is determined by Step 2, i.e., the approximation ratio of the greedy algorithm for the MCSC problem. First, we establish the following lemma.

Lemma 2 *Given a weighted MCSC problem* $(V, \mathcal{S} \setminus V, w, G)$ *with* $D_C(G) = 1$, *let*

$$R_w = \frac{\max_{S \in \mathcal{S}}\{w(S)\}}{\min_{S \in \mathcal{S}}\{w(S)\}}. \tag{2}$$

Then the approximation ratio of the greedy algorithm for MCSC is at most $R_w + H(\gamma - 1)$, *where* $\gamma = \max\{|S| \mid S \in \mathcal{S} \setminus V\}$ *is the maximum size of the subsets in* $\mathcal{S}$ *and* $H(\cdot)$ *is the harmonic function.*

Proof *(Proof of Lemma 2)* The proof is based on the classic charge argument. Let $\mathcal{S}^*$ be an optimal solution to the weighted set cover problem $(V, \mathcal{S} \setminus V, w)$, and $\mathcal{R}$ the solution returned by the greedy algorithm for the weighted MCSC problem $(V, \mathcal{S} \setminus V, w, G)$ with $D_C(G) = 1$. Let $w(\mathcal{S}^*)$ and $w(\mathcal{R})$ denote the total weight of the subsets included in $\mathcal{S}^*$ and $\mathcal{R}$, respectively. In the following, we will show that

$$\frac{w(\mathcal{R})}{w(\mathcal{S}^*)} \le R_w + H(\gamma - 1). \tag{3}$$

Let $\mathcal{S}_C^*$ be an optimal solution to the weighted MCSC problem $(V, \mathcal{S} \setminus V, w, G)$ with $D_C(G) = 1$. Since $w(\mathcal{S}^*) \le w(\mathcal{S}_C^*)$, Lemma 2 follows immediately from (3).

To prove (3), we apply the classic charge argument. Each time a subset S_0 (at step 1) or a shortest $\mathcal{R} \to S$ path P_S^* (at step 2) is selected to be added to $\mathcal{R}$, we charge each of the newly covered elements $\frac{w(S_0)}{|S_0|}$ (at step 1) or $r(P_S^*)$ defined in (1) (at step 2). Notice that when $D_C(G) = 1$, the shortest $\mathcal{R} \to S$ path P_S^* is only a single edge connecting some subset in $\mathcal{R}$ and S, and

$$r(P_S^*) = \frac{w(\mathcal{S}(P_S^* \setminus \mathcal{R}))}{|V_N(P_S^*)|} = \frac{w(S)}{|S \setminus U|}.$$

During the entire procedure, each element of V is charged exactly once. Assume that step 2 is completed in $K - 1$ iterations. Let P_{Si}^* be the shortest $\mathcal{R} \to S$ path selected by the algorithm at iteration i. Let $C(v)$ denote the charge of an element v in V. Then we have that

$$\begin{aligned}\sum_{v \in V} C(v) &= \sum_{i=0}^{K-1} \sum_{v \in V_N(P_{Si}^*)} C(v) \\ &= \sum_{i=0}^{K-1} \sum_{v \in S_i \setminus U} \frac{w(S_i)}{|S_i \setminus U|} \\ &= \sum_{i=0}^{K-1} w(S_i) = w(\mathcal{R}),\end{aligned} \tag{4}$$

where $P_{S0}^* = \{S_0\}$.

Suppose that $\mathcal{S}^* = \{S_1^*, \ldots, S_N^*\}$ is a minimum weighted set cover for $\{V, \mathcal{S} \setminus V, w\}$. Since an element of V may be contained in more than one subset of $\mathcal{S}^*$, it follows that

$$\sum_{v \in V} C(v) \le \sum_{i=1}^{N} \sum_{v \in S_i^*} C(v). \tag{5}$$

Next we will show an inequality which bounds from above the total charge of a subset in $\mathcal{S}^*$, i.e., for any $S^* \in \mathcal{S}^*$,

$$\sum_{v \in S^*} C(v) \le [R_w + H(|S^*| - 1)] w(S^*). \tag{6}$$

Let n_i $(i = 0, 1, \ldots, K)$ be the number of elements of S^* that have not been covered by $\mathcal{S}$ after iteration $i - 1$, where step 1 is considered as iteration 0. Let $\{i_1, \ldots, i_k\}$ denote the subsequence of $\{i = 0, 1, \ldots, K - 1\}$ such that $n_i - n_{i+1} > 0$. For each element a covered at iteration i_1, if $i_1 = 0$, based on the greedy rule at step 1, we have that

$$C(v) = r(P_{S_0}^*) \le \frac{w(S^*)}{n_{i_1}}; \tag{7}$$

Otherwise,

$$C(v) = r(P_{S_{i_1}}^*) = \frac{w(S_{i_1})}{|S_{i_1} \setminus U|} \le \frac{w(S^*) R_w}{n_{i_1} - n_{(i_1+1)}}. \tag{8}$$

The inequality in (8) is due to the fact that S_{i1} covers at least $n_{i1} - n_{(i1+1)}$ elements of V, i.e., $|S_{i_1} \setminus U| \ge n_{i_1} - n_{(i_1+1)}$. Summing up (7) and (8),

$$C(v) \le \frac{w(S^*) R_w}{n_{i_1} - n_{(i_1+1)}}. \tag{9}$$

Consider two cases:

1. If all the elements of S^* have been covered after iteration i_1, *i.e.*, $n_{(i_1+1)} = 0$, then

$$\sum_{v \in S^*} C(v) \le \sum_{v \in S^*} \frac{w(S^*) R_w}{n_0} = w(S^*) R_w. \tag{10}$$

2. If not all the elements of S^* have been covered by $\mathcal{R}$ after iteration i_1, S^* becomes cover-adjacent with $\mathcal{R}$ and thus a candidate for being selected at the following iterations. At each iteration, for each element $v \in S^*$ covered at iteration i_j $(j = 2, \ldots, k)$, the greedy rule at step 2 still yields

$$\begin{aligned}C(v) &= r(P_{S_{i_j}}^*) \le r(P_{S^*}^*) \\ &= \frac{w(S^*)}{|S^* \setminus U|} = \frac{w(S^*)}{n_{i_j}}.\end{aligned} \tag{11}$$

It follows from (9,11) that

$$\sum_{v\in S^*} C(v) \le w(S^*)(n_{i_1} - n_{(i_1+1)})\frac{1}{n_{i_1} - n_{(i_1+1)}} + w(S^*)\sum_{j=2}^{k}(n_{i_j} - n_{(i_j+1)})\frac{1}{n_{i_j}} = w(S^*)\left(1 + \sum_{j=2}^{k}\frac{n_{i_j} - n_{i_{(j+1)}}}{n_{i_j}}\right). \quad (12)$$

Here we have used the fact that $n_{(i_j+1)} = n_{i_{(j+1)}}$. It is because between iteration i_j and iteration $i_{(j+1)}$, no elements of S^* are covered.

For the summation term in (12), we have the following inequality:

$$\sum_{j=2}^{k}\frac{n_{i_j} - n_{i_{(j+1)}}}{n_{i_j}} \le \sum_{j=2}^{k}\frac{1}{n_{i_j}} + \cdots + \frac{1}{n_{i_{(j+1)}}+1} = H(n_{i_2}) \le H(|S^*| - 1). \quad (13)$$

The last inequality is due to the fact that $n_{i2} \le n_{i1} - 1 = |S^*| - 1$.

Equationn (6) is a direct consequence of (10), (12), and (13). Thus, using (4–6),

$$w(\mathcal{R}) = \sum_{v\in V} C(v) \le \sum_{i=1}^{N}\sum_{v\in S_i^*} C(v) \le \sum_{i=1}^{N}[R_w + H(|S_i^*| - 1)]w(S_i^*) \le [R_w + H(\gamma - 1)]w(\mathcal{S}^*).$$

□

Then, as a direct consequence of Lemma 2, we have the following theorem on the approximation ratio[7] of the greedy algorithm for the MCSC problem with $D_C(G) = 1$.

Theorem 3 *Let Δ^* be the MCSSub of a WSC $\Delta = (V, \mathcal{S}, w)$ and Δ_C be the solution returned by Algorithm 1. Let R_w be defined as in (2). Then we have*

$$\frac{w_F(\Delta_C)}{w_F(\Delta^*)} \le R_w + H(dim\Delta),$$

where $dim\Delta$ is the dimension of Δ and $H(\cdot)$ is the harmonic function.

From Theorem 3, we see that the approximation ratio depends on the ratio R_w of the maximum weight to the minimum weight. It is shown in the following theorem that if R_w is unbounded, then the scaling order of the approximation ratio can be as bad as linear with respect to the number of vertices in the simplicial complex.

Theorem 4 *Let n be the number of the vertices in a WSC $\Delta = (V, \mathcal{S}, w)$, and R_w defined as in (2). If R_w is unbounded, then the approximation ratio of Algorithm 1 for the MCSSub problem of Δ is $\Omega(n)$.*

Proof Consider a specific example: Δ is a $(n-1)$-dimensional simplex with the vertex set $V = \{v_1, \ldots, v_n\}$, and all the weights of the faces are infinite except for the following five faces:

$$\begin{aligned} w(S_1) &= w\left(\left\{v_1, \ldots, v_{\frac{n}{2}}\right\}\right) = \frac{1}{2}, \\ w(S_2) &= w\left(\left\{v_1, \ldots, v_{\frac{n}{4}}, v_{\left(\frac{n}{2}+2\right)}\right\}\right) = \frac{1}{2}, \\ w(S_3) &= w\left(\left\{v_{\left(\frac{n}{4}+1\right)}, \ldots, v_{\left(\frac{n}{2}+1\right)}\right\}\right) = \frac{1}{2}, \\ w(S_4) &= w\left(\left\{v_{\frac{n}{2}}, \ldots, v_n\right\}\right) = \frac{n}{8}, \\ w(S_5) &= w\left(\left\{v_{\left(\frac{n}{2}+1\right)}, \ldots, v_n\right\}\right) = 1. \end{aligned}$$

For ease of presentation, we have assumed that n is a multiple of 4. By applying Algorithm 1, we reduce the MCSSub problem for Δ to the MCSC problem $(V, \mathcal{S}\setminus V, w, G_\Delta)$. Due to the weight assignment, it suffices to only consider the subgraph of G_Δ induced by the above five faces, as shown in Fig. 5.

The optimal solution Δ^* to the MCSSub problem is given by

$$\Delta^* = \{S \in \mathcal{S} \mid S \subseteq S_2 \text{ or } S_3 \text{ or } S_5\},$$

and

$$w_F(\Delta^*) = w(S_2) + w(S_3) + w(S_5) = 2.$$

On the other hand, the solution Δ_C returned by Algorithm 1 is given by

$$\Delta_C = \{S \in \mathcal{S} \mid S \subseteq S_1 \text{ or } S_4\},$$

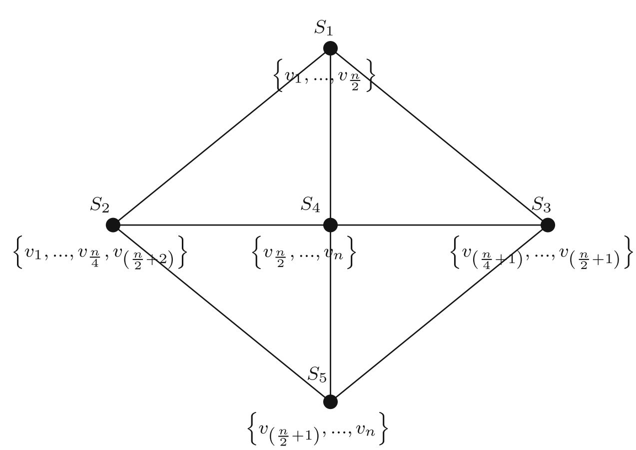

Fig. 5 The *subgraph* of G_Δ induced by the five faces S_1, S_2, S_3, S_4, and S_5 with finite weights

[7] The approximation ratio of the greedy algorithm for general weighted MCSC problem is still an open problem.

and

$$w_F(\Delta_C) = w(S_1) + w(S_4) = \frac{1}{2} + \frac{n}{8}.$$

Specifically, S_1 is firstly selected, and then S_4. Thus,

$$\frac{w_F(\Delta_C)}{w_F(\Delta^*)} = \frac{n}{16} + \frac{1}{4} = \Theta(n).$$

It follows that the approximation ratio of Algorithm 1 is $\Omega(n)$. □

From Theorem 4, we see that Algorithm 1 is not suitable for the MCSSub problem of a WSC Δ if its weight function has a relatively wide range. As shown next in Sect. 3.3, the other approximation algorithm based on the Steiner tree does not have this issue: its approximation ratio does not depend on the range of the weight function.

3.3 Algorithm based on Steiner tree

From a WSC $\Delta = (V, \mathcal{S}, w)$, we derive an undirected graph H_Δ with the vertex set $\mathcal{S}$: for each face $S \in \mathcal{S} \setminus V$ (i.e., the faces that are not the vertices of Δ), we replace it by a vertex v_S in H_Δ and connect v_S to all the vertices of S. The weight $w(v_S)$ assigned to the vertex v_S is the weight $w(S)$ of the face S. Notice that the weight of vertices in H_Δ corresponding to the vertices in Δ (i.e., V) is zero. Fig. 6 shows an example of the derivation of the graph from a 2-simplex. We have the following theorem on the relation between the MCSSub of Δ, the Steiner tree of H_Δ that spans the vertex set V of Δ and the minimum connected dominating set[8] of H_Δ.

Algorithm 3 *STA for MCSSub:*
INPUT: A WSC $\Delta = (V, \mathcal{S}, w)$.
OUTPUT: An approximate MCSSub Δ_C of Δ.

1. *Derive the graph H_Δ from Δ.*
2. *Obtain an approximate Steiner tree T of H_Δ by using the algorithms given in [8, 9].*
3. *Transform T to a connected spanning subcomplex Δ_C of Δ by mapping each vertex of T to a face of Δ.*

Theorem 5 *Let Δ^* denote the MCSSub of a WSC $\Delta = (V, \mathcal{S}, w)$, T^* the Steiner tree of H_Δ that spans the vertex set V of Δ, and D_C^* the minimum CDS of H_Δ. Then we have that*

$$w_F(\Delta^*) = w(T^*) = w(D_C^*).$$

Proof First we show that $w_F(\Delta^*) = w(T^*)$. Since every connected spanning subcomplex Δ' of Δ corresponds to a connected subgraph of H_Δ which only contains the vertices of Δ and the vertices representing the facets of Δ', it follows that $w(T^*) \le w_F(\Delta^*)$. On the other hand, since by contradiction, there is a one-to-one mapping between the vertices of the Steiner tree of H_Δ and the vertices plus the facets of a connected spanning subcomplex of Δ, it follows that $w_F(\Delta^*) \le w(T^*)$.

Next we show that $w(T^*) = w(D_C^*)$. Notice that the vertex set V of Δ is a dominating set of H_Δ. Since the Steiner tree T^* of H_Δ spans the vertex set V, T^* is a CDS of H_Δ. Thus, $w(D_C^*) \le w(T^*)$. On the other hand, given the minimum CDS D_C^* of H_Δ, since each vertex v in the vertex set V is either in D_C^* or a neighbor of some face in D_C^* and the weights of the vertices in V are all zero, the combination of V and D_C^* yields a connected subgraph of H_Δ that spans V with the same weight as D_C^*. Thus, $w(T^*) \le w(D_C^*)$. □

Based on Theorem 5, we propose the following Steiner tree based algorithm (STA) for the MCSSub problem.

Since approximation only occurs in Step 2, the approximation ratio of STA is equal to that of the algorithm for the node-weighted Steiner tree problem. The best approximation ratio is known to be $(1.35 + \epsilon) \ln n$ for any constant $\epsilon > 0$, where n is the number of vertices of Δ and is also the number of terminals in the Steiner tree of H_Δ [8]. Here we do not try to find the CDS D_C^* of H_Δ at step 2, because the best known approximation ratio for the CDS problem is $(1.35 + \epsilon) \ln n_{(H_\Delta)}$ where $n_{(H_\Delta)}$ is the number of vertices of H_Δ [7, 8]. Since $n_{(H_\Delta)} \gg n$, the latter approximation ratio is much worse than the former one.

3.4 Time complexity analysis

Here we analyze the time complexity of SCA and STA for the MCSSub problem. Given a WSC $\Delta = (V, \mathcal{S}, w)$, let $n = |V|$ denote the number of vertices in Δ, $m = |\mathcal{S} \setminus V|$

[8] A dominating set of a graph is a subset of vertices such that every vertex of the graph is either in the subset or a neighbor of some vertex in the subset, and a connected dominating set (CDS) is a dominating set where the subgraph induced by the vertices in the dominating set is connected. The CDS problem asks for a CDS with the minimum total weight, and it is shown to be a special case of the MCSC problem [21].

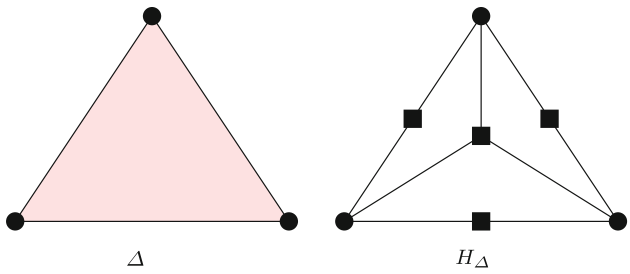

Fig. 6 The derived *graph* of a 2-simplex (*squares* in H_Δ represent the faces that are not vertices of Δ)

the number of non-vertex faces in Δ, and d the dimension of Δ. Since all the information of the auxiliary graph G_Δ for SCA and the derived graph H_Δ for STA can be easily retrieved from the WSC Δ, Step 1 in both algorithms can be skipped in the implementation and the time complexity of both algorithms is determined by their Step 2.

Step 2 of SCA is to apply the greedy algorithm to the MCSC problem $(V, \mathcal{S} \setminus V, w, G_\Delta)$. It takes $O(m)$ time to complete Step 1 of the greedy algorithm and $O(n)$ iterations to complete Step 2 of the greedy algorithm. At each iteration of Step 2, the weight ratios of at most m faces are computed and the weight ratio of each face is done in constant time. Thus, the running time of SCA is $O(m + nm) = O(nm)$. Since the derived graph H_Δ has $n + m$ vertices and $O(dm)$ edges and the Steiner tree has n terminals to cover, it follows from [14] that the running time of Step 2 of STA is $O(dnm^2 + nm^2 \log m)$. From the above, we see that the time complexity of STA is significantly higher than that of SCA. This is mostly because the approximation algorithm for the Steiner tree requires the computation of the shortest paths between all vertex pairs.

We point out that while the Steiner tree based algorithm has a higher complexity, it can offer better performance in a WSC with a large weight range. In a simulation example of random simple complexes, we consider a case where each face weight takes only two values w_{min} and w_{max} with equal probability. With $w_{min} = 1$, $w_{max} = 10{,}000$, and 1,000 Monte Carlo runs for a 200-vertex random simplicial complex,[9] [15] we find that the total weight of the solution returned by the set cover based algorithm can be 1.7 times that of the solution returned by the Steiner tree based algorithm. These two algorithms thus offer a tradeoff between performance and complexity.

4 Simulation results

In this section, we present simulation results on the performance of the two approximation algorithms (SCA and STA) for the broadcast problem in an MR-MC network. We assume that the communication rate is determined by the channel bandwidth and its radio propagation characteristics, and does not depend on the radio. We consider a dense MR-MC network, where all the nodes are within each others transmission range, and thus whether two nodes can communicate with each other depends solely on the channel assignment. As discussed in Sect. 1.2, the cost to be minimized can be of any form that is summable over all the transmissions, and here we aim to minimize the total energy consumption of the broadcast.

There are C non-overlapping channels f_i $(1 \le i \le C)$, possibly with different communication rates r_i, available for the MR-MC network, and each node is equipped with R radios $(R \le C)$. We assume that the communication rate is determined by the channel bandwidth and its radio propagation characteristics, and does not depend on the radio. At the beginning of the broadcast, each node randomly selects R of the C channels for its R radios. As discussed in Sect. 1.3, the nodes which share at least one common channel form a clique, and there is a one-to-one correspondence between the cliques and the faces of the derived WSC. The weight of the face is defined as the energy consumption of the broadcast within the

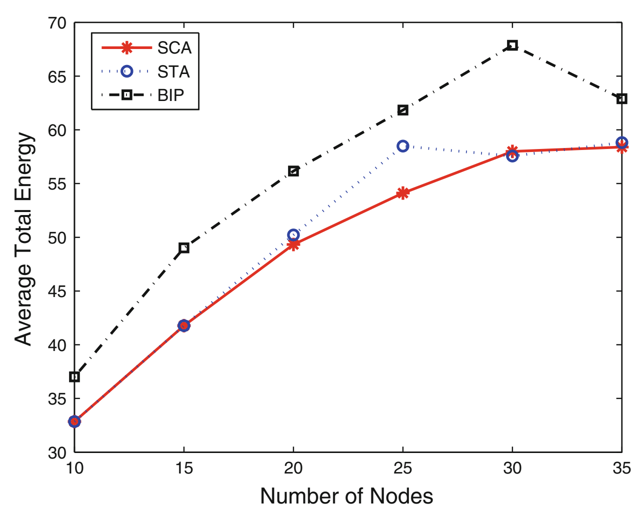

Fig. 7 Average total energy versus number of nodes. Parameters: $C = 12$, $R = 4$, $P_{tx} = 1$, $P_{rx} = 0.01$, $L = 100$, $r_i = i$ for $1 \le i \le 12$

[9] A random simplicial complex $\Delta(n, D, \mathbf{p})$ with n vertices, dimension at most D, and a D-dimensional probability vector $\mathbf{p} = \{p_1, p_2, \ldots, p_D\}$ is constructed in a bottom-up manner: first n vertices are fixed, which are the 0-simplices of Δ, and then higher-dimensional simplices are generated inductively. Specifically, for each $1 \le i \le D$, after all the simplices with dimension lower than i have been generated, consider every i-tuple of vertices: if they have formed all the lower dimensional simplices, then an i-simplex consisting of them is generated with probability p_i. Notice that a random simplicial complex $\Delta(n, 1, p)$ is the random graph introduced by Erdős and Rényi [5].

corresponding clique, i.e., the sum of the transmission energy and the reception energy. Let S be a face containing $k+1$ nodes and $\{f_{Sj} : j = 1, 2, \ldots, q\}$ the q ($1 \le q \le C$) common channels shared by the $k+1$ nodes. Assume that if a node in the clique is selected as relay, it will choose the common channel with the maximum communication rate to transmit. Then the weight $w(S)$ of the face S is given by

$$w(S) = (P_{tx} + kP_{rx}) \frac{L}{\max_{j=1,\ldots,q} \{r_{Sj}\}},$$

where P_{tx} and P_{rx} are the transmission power and the reception power, respectively, and L is a constant.

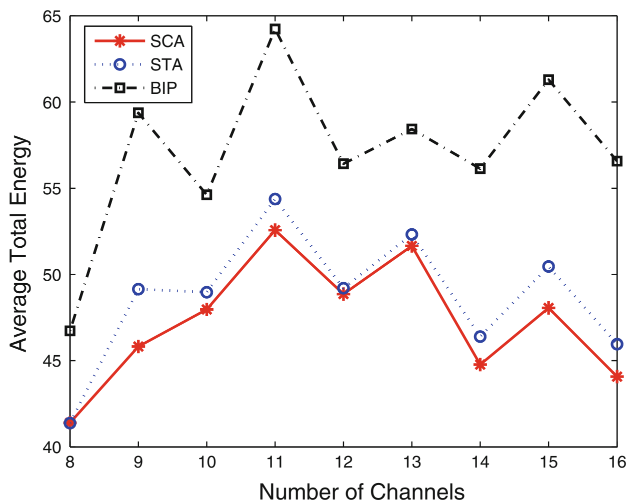

Fig. 8 Average total energy versus number of channels. Parameters: $n = 20, \quad R = 4, \quad P_{tx} = 1, \quad P_{rx} = 0.01, \quad L = 100, \quad r_i = i$ for $1 \le i \le C$

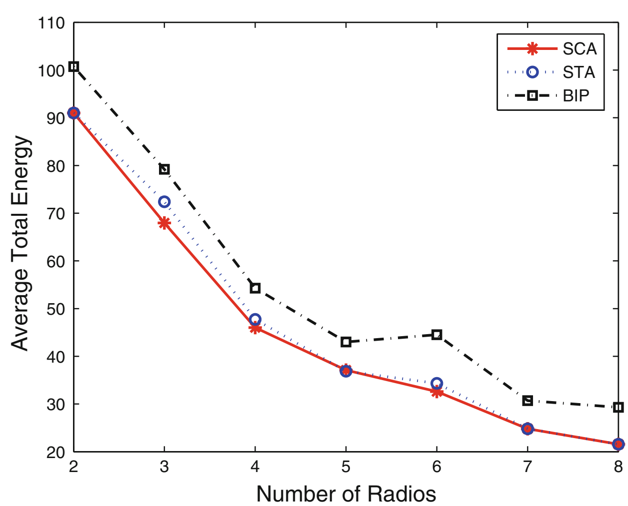

Fig. 9 Average total energy versus number of radios. Parameters: $n = 20, \quad C = 12, \quad P_{tx} = 1, \quad P_{rx} = 0.01, \quad L = 100, \quad r_i = i$ for $1 \le i \le 12$

The BIP algorithm for the broadcast problem in an SR-SC network can be easily extended to solve the broadcast problem in an MR-MC network. The major difference is that two nodes may be connected by more than one link if they share more than one common channel. At each iteration, after the optimal link on one channel is selected, the cost of other related links on the same channel is updated in the same manner. In Figs. 7, 8 and 9, the average total energy of the solutions returned by SCA and STA is compared with that of the BIP with respect to the underlying graph of the WSC. The average is taken over 10 random channel assignments. We see that the performances of SCA and STA are extremely close, and their performances are significantly better than that of BIP.

5 Conclusion and future work

In this paper, we study the minimum cost broadcast problem in multi-radio multi-channel ad hoc networks, where the total cost is the sum of the costs associated with the transmissions during the broadcast. We formulate it as the minimum spanning problem in simplicial complexes. We show that it is NP-complete. Hence we propose two approximation algorithms for this minimum spanning problem: one is to transform it into the connected set cover problem; the other is to transform it into the node-weighted Steiner tree problem and then apply the corresponding algorithm. Despite their distinct approaches, both approximation algorithms are shown to be order-optimal and offer a tradeoff in terms of performance versus complexity.

As a starting point, we have assumed that the channel assignment scheme is designed independent of the broadcast scheme. The joint optimization of the two schemes will further reduce the broadcast cost. Another future direction is to develop distributed versions of the approximation algorithms for the minimum cost broadcast problem.

Acknowledgments This work was supported by the Army Research Laboratory NS-CTA under Grant W911NF-09-2-0053.

References

1. Bazlamacci, C. F., & Hindi, K. S. (2001). Minimum-weight spanning tree algorithms: A survey and empirical study. *Computers & Operations Research, 28*(8), 767–785.
2. Čagalj, M., Hubaux, J., Enz, C. (2002). Minimum-energy broadcast in all-wireless networks: Np-completeness and distribution issues. In *Proceedings of the ACM MobiCom*, pp. 172–182.
3. Chiu, H. S., Wu, B., Yeung, K. L., & Lui, K. S. (2008). Widest spanning tree for multi-channel multi-interface wireless mesh networks. In *Proceedings of the IEEE WCNC*, pp. 2194–2199.

4. Chou, C. T., Misra, A., & Qadir, J. (2006). Low-latency broadcast in multirate wireless mesh networks. *IEEE Journal on Selected Areas in Communications (JSAC), 24*(11), 2081–2091.
5. Erdős, P., & Rényi, A. (1960). On the evolution of random graphs. In *Publication of the Mathematical Institute of the Hungarian Academy of Sciences*, pp. 17–61.
6. Feige, U. (1998). A threshold of ln *n* for approximating set cover. *Journal of the ACM, 45*(4), 634–652.
7. Guha, S., & Khuller, S. (1998). Approximation algorithms for connected dominating sets. *Algorithmica, 20*(4), 374–387.
8. Guha, S., & Khuller, S. (1999). Improved methods for approximating node weighted steiner trees and connected dominating sets. *Information and Computation, 150*(1), 57–74.
9. Klein, P. N. & Ravi, R. (1995). A nearly best-possible approximation algorithm for node-weighted steiner trees. *Journal of Algorithms, 19*(1), 104–115.
10. Kleinberg, J. & Tardos, E. (2005). *Algorithm design*. Boston, MA: Addison Wesley.
11. Kyasanur, P., & Vaidya, N. H. (2006). Routing and link-layer protocols for multi-channel multi-interface ad hoc wireless networks. *ACM SIGMOBILE Mobile Computing and Communications Review, 10*(1), 31–43.
12. Kyasanur, P., So, J., Chereddi, C., Vaidya, N. H. (2006). Multichannel mesh networks: Challenges and protocols. *IEEE Wireless Communications, 13*(2), 30–36.
13. Li, L., Qin, B., Zhang, C., & Li, H. (2007). Efficient broadcasting in multi-radio multi-channel and multi-hop wireless networks based on self-pruning. In *Proceedings of the HPCC*, pp. 484–495.
14. Liang, W. F. (2006). Approximate minimum-energy multicasting in wireless ad hoc networks. *IEEE Transactions on Mobile Computing 5*(4), 377–387.
15. Mannucci, M. A., Sparks, L., & Struppa, D. C. (2006). *Simplicial models of social aggregation I*. URL: http://arxiv.org/abs/cs/0604090.
16. Munkres, J. R. (1984). *Elements of algebraic topology*. Menlo Park, CA: Addison-Wesley.
17. Nguyen, H. L., & Nguyen, U. T. (2009). Channel assignment for multicast in multi-channel multi-radio wireless mesh networks. *Wireless Communications and Mobile Computing, 9*(4), 557–571.
18. Ramanathan, R., Bar-Noy, A., Basu, P., Johnson, M., Ren, W., Swami, A., et al. (2011). Beyond graphs: Capturing groups in networks. In *Proceedings of the IEEE NetSciCom*.
19. Raniwala, A., & Chiueh, T. (2005). Architecture and algorithms for an IEEE 802.11-based multi-channel wireless mesh network. In *Proceedings of the IEEE INFOCOM*, pp. 2223–2234.
20. Redi, J., & Ramanathan, R. (2011). The DARPA WNaN network architecture. In *Proceedings of the IEEE MILCOM*, Baltimore.
21. Ren, W., & Zhao, Q. (2011). A note on: 'Algorithms for connected set cover problem and fault-tolerant connected set cover problem'. *Theoretical Computer Science, 412*(45), 6451–6454.
22. Tang, J., Xue, G., & Zhang, W. (2005). Interference-aware topology control and QoS routing in multi-channel wireless mesh networks. In *Proceedings of the ACM MobiHoc*, pp. 68–77.
23. Wieselthier, J. E., Nguyen, G. D., & Ephremides, A. (2002). Energy-efficient broadcast and multicast trees in wireless networks. *Mobile Networks and Applications, 7*(6), 481–492.
24. Zhang, Z., Gao, X. F., & Wu, W. L. (2009). Algorithms for connected set cover problem and fault-tolerant connected set cover problem. *Theoretical Computer Science, 410*(8–10), 812–817.

Author Biographies

Wei Ren received a Ph.D. degree in 2011 from University of California, Davis, and an M.Sc. degree in 2006 and a B.Sc. degree in 2003, both from Peking University, China. During the Ph.D. study, his research interests were in cognitive radio systems and wireless networks, and He was also interested in algorithmic theory and optimization techniques for communications. He is currently a Software Development Engineer at Microsoft Corporation.

Qing Zhao received the Ph.D. degree in Electrical Engineering in 2001 from Cornell University, Ithaca, NY, USA. In August 2004, she joined the Department of Electrical and Computer Engineering at University of California, Davis, where she is currently a Professor. Her research interests are in the general area of stochastic optimization, decision theory, and algorithmic theory in dynamic systems and communication and social networks. She received the 2010 IEEE Signal Processing Magazine Best Paper Award and the 2000 Young Author Best Paper Award from the IEEE Signal Processing Society. She holds the title of UC Davis Chancellors Fellow and received the 2008 Outstanding Junior Faculty Award from the UC Davis College of Engineering. She was a plenary speaker at the 11th IEEE Workshop on Signal Processing Advances in Wireless Communications (SPAWC), 2010. She is also a coauthor of two papers that received student paper awards at ICASSP 2006 and the IEEE Asilomar Conference 2006.

Ram Ramanathan is a Principal Scientist in the Network Research Department at Raytheon BBN Technologies, where he helps invent, develop and prototype new technologies, primarily in wireless networks. Ram's current research areas broadly include mobile multi-hop wireless networks, network analysis and control algorithms, and network science. Over the course of his career, he has made significant contributions to the areas of scalable routing, topology control and the use of directional antennas in multihop networks and led several projects for the DoD, in particular DARPA. Ram has published widely in international journals and conferences, including best paper award winning papers at IEEE MILCOM, IEEE INFOCOM, and ACM SIGCOMM. He has served on the program committees of several conferences including MobiCom, Mobihoc and Infocom, and was the TPC co-chair of Mobicom 2007. Ram has served on the editorial boards of IEEE Transactions on Mobile Computing (TMC) and Ad Hoc Networks journal, and is currently the Associate Editor-in-Chief of TMC. He is a named inventor on eight patents. Ram received his Bachelor of Technology in Electrical Engineering from the Indian Institute of Technology, Madras, India, and received his M.S. and Ph.D. degrees in Computer and Information Sciences from the University of Delaware. He is a Fellow of the IEEE.

Jianhang Gao received his B.S. degree from Tsinghua University, Beijing, People's Republic of China in 2010. He joined the signal processing and adaptive networking group under instruction of Prof. Qing Zhao California, USA in 2010, researching in the area of algorithmic study of hypergraph and simplicial complex. Currently, he is studying for the Ph.D. degree in Electrical and Computer Engineering at University of California, Davis.

Ananthram Swami is with the U.S. Army Research Laboratory (ARL) as the Army's ST (Senior Research Scientist) for Network Science. He is an ARL Fellow and Fellow of the IEEE. He has held positions with Unocal Corporation, the University of Southern California (USC), CS-3 and Malgudi Systems. He was a Statistical Consultant to the California Lottery, developed a MATLAB-based toolbox for non-Gaussian signal processing, and has held visiting faculty positions at INP, Toulouse. He received the B.Tech. degree from IIT-Bombay; the M.S. degree from Rice University, and the Ph.D. degree from the University of Southern California (USC), all in Electrical Engineering. His research interests are in the broad area of network science: the study of interactions and co-evolution, prediction and control of inter-dependent networks, with applications in composite tactical networks.

Amotz Bar-Noy received the B.Sc. degree (1981) in Mathematics and Computer Science (summa cum laude) and the Ph.D. degree (1987) in Computer Science, both from the Hebrew University, Israel. He was a postdoc fellow in Stanford University, California (1987–1989); a Research Staff Member with IBM Research Center, New York (1989–1996); an associate Professor with the Electrical Engineering Department of Tel Aviv University, Israel (1996–2001); a Principal Technical Staff Member with AT&T research labs in New Jersey (1999-2001). Since 2002 he is a Professor with the Computer and Information Science Department of Brooklyn College—CUNY, Brooklyn New York and with the Computer Science Department of the Graduate Center of CUNY, Manhattan New York. He has published more than eighty refereed journal articles and more than 100 refereed conference and workshop articles. He served as a program committee member for many conference. He is an editor for the IEEE Transactions on Mobile Computing (TMC) journal. He served as an editor for the Wireless Networks (WINET) journal for about 10 years and served as a guest editor for two special issues one of Wireless Network and one of Mobile Networking and Applications. He served as a co-chair of the ALGOSENSORS 2012 Symposium. His field of expertise belongs to the Theoretical Computer Science community and to the Networking community. The scope of his research is to bridge the gap between these two communities.

Matthew P. Johnson studied philosophy and computer science at Columbia and Lawrence Universities and received his Ph.D. in computer science from the City University of New York in 2010. He is currently a postdoc at UCLA.

Prithwish Basu is a Senior Scientist at Raytheon BBN Technologies, where he has been playing leading roles in several DoD funded research programs. His research interests include network science, energy efficiency and routing issues in wireless ad hoc and sensor networks; and theoretical aspects of networking, in general. He is a technical lead of the ongoing Army funded Network Science Collaborative Technology Alliance (NS CTA) program and the Principal Investigator for BBN in the US/UK International Technology Alliance (ITA) program. Prithwish holds a Ph.D. (2003) and an M.S. (1999) in Computer Engineering from Boston University and a B.Tech. (1996) in Computer Science & Engineering from Indian Institute of Technology, Delhi. He has published over 65 papers in leading networking journals and conferences. Prithwish has served on the TPC of IEEE INFOCOM for the past few years, is an Associate Editor of IEEE Trans. Mobile Computing, and is currently serving as a Guest Editor for the IEEE J-SAC Special Issue on Network Science. In 2006, Prithwish received MIT Technology Reviews TR35 award given to top 35 young innovators under the age of 35.