Chapter 11
Minimum-Latency Scheduling

The key is not to prioritize what's
on your schedule, but to schedule your priorities.
STEPHEN R. COVEY

11.1 Motivation and Overview

Consider a multihop wireless network in which all network nodes V lie in plane and have a unit communication radius. Its communication topology G is the unit disk graph (UDG) of V. Under the protocol interference model, every node has a communication radius normalized to one, and an interference radius ρ for some parameter $\rho \geq 1$ (see Fig. 11.1). A node v can receive the message successfully from a transmitting node u if v is within the transmission range of u but is outside the interference range of any other node transmitting simultaneously.

In this chapter, we study minimum-latency schedulings for the following four group communications in the multihop wireless network:

- *Broadcast*. A distinguished source node sends a common packet to all other nodes.
- *Data Aggregation*. A distinguished sink node collects the data aggregated from all the packets at the nodes other than the sink node. In other words, every intermediate node combines all received packet with its own packet into a single packet of fixed size according to some aggregation function such as logical and/or, maximum, or minimum.
- *Data Gathering*. A distinguished sink node collects a packet from every other node.
- *Gossiping*. Every node broadcasts a common packet to all other nodes.

Suppose that all communications proceed in synchronous time slots and a node can transmit at most one packet of a fixed size in each time slot. A communication schedule for each of these four communication tasks not only specifies a

D.-Z. Du and P.-J. Wan, *Connected Dominating Set; Theory and Applications*, Springer Optimization and Its Applications 77, DOI 10.1007/978-1-4614-5242-3_11,

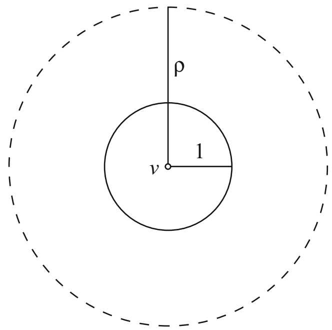

Fig. 11.1 Each node has a unit communication radius and an interference radius $\rho \geq 1$

communication routing but also assigns a time slot to every communication link in the routing subject to two constraints:

1. The link ordering given by the routing should be followed.
2. All communication links assigned in each time slot are interference-free.

The latency of a communication schedule is the number of time slots during which at least one transmission occurs. The minimum-latency schedulings for the above four group communications are all NP-hard. In this chapter, we present constant-approximation algorithms for them.

The CDS plays a critical role in the design of scheduling algorithms [109, 110]. Indeed, all relaying nodes of a message must form a CDS of G. However, in order to achieve a short latency, the CDS has to be short and sparse instead of being just small. Specifically, consider a node $s \in V$ and let R be the graph radius of G with respect to s. A CDS U of G with $s \in U$ should be "short" in the sense that the graph radius of $G[U]$ with respect to s is bounded by a constant factor of R and "sparse" in the sense that the maximum degree of $G[U]$ is bounded by a small constant. In Sect. 11.3, we will present a construction of such short and sparse CDS.

The following terms and notations will be adopted throughout this chapter. Let $n = |V|$. The connected domination number of G is denoted by γ_c. The unit disk centered at a node v is denoted by $D(v)$. The topological boundary of a planar set A is denoted by ∂A. The directed version of G, denoted by $\overrightarrow{G}$, is the digraph obtained from G by replacing every edge e in G with two oppositely oriented links between the two endpoints of e. A subset U of nodes is said to be distance-d independent for some $d > 0$ if and only if their mutual Euclidean distances are greater than d. Equivalently, a set of nodes are distance-d independent if and only if they form an independent set of the d-disk graph on V. For any $d > 0$, a distance-d coloring of a subset U of nodes is an assignment of colors to the nodes in U such that any pair of nodes of distance at most d receives distinct colors. Let X and Y be two disjoint subsets of V. Y is a *cover* of X if each node in X is adjacent to some node in Y, and a *minimal cover* of X if Y is a cover of X but no proper subset of Y is a cover of X. Any ordering $y_1, y_2, \cdots, y_m$ of Y induces a minimal cover $Z \subseteq Y$ of X by the following sequential pruning method: Initially, $Z = Y$. For each $i = m$ down to 1,

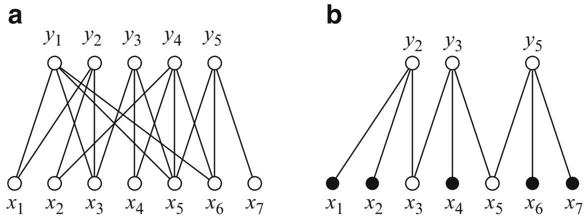

Fig. 11.2 (**a**) $X = \{x_i : 1 \leq i \leq 7\}$ is covered by $Y = \{y_i : 1 \leq i \leq 5\}$. (**b**) $\{y_2, y_3, y_4\}$ is a minimal cover of X. The *black* nodes are private neighbors

Table 11.1 Summary on the approximation bounds of the scheduling algorithms

Communication	Approximation bound
Broadcast	$2\beta_\rho$
Aggregation with $\rho > 1$	$(\alpha_\rho + 12)\,\beta_\rho$
Gathering	$2\beta_\rho$
Gossiping	$4\beta_\rho$

if $Z \setminus \{y_i\}$ is a cover of X, remove y_i from Z. Figure 11.2 is an illustration of such sequential pruning method. Suppose that Y is a cover of X. A node $x \in X$ is called a *private* neighbor of a node $y \in Y$ with respect to Y if y is the only node in Y which is adjacent to x. Clearly, if Y is a minimal cover of X, then each node in Y has at least one private neighbor with respect to Y.

Short and sparse CDS has important applications in the design of scheduling algorithms for group communications in wireless networks. In this chapter, we have constructed a short and sparse CDS and built a dominating tree on such CDS which is used as the routing for the group communications. By exploiting the rich structural properties of the dominating tree, we are able to design scheduling algorithms with constant approximation bounds for the group communications. Table 11.1 summarizes the approximation bounds of our scheduling algorithms described in this chapter for the four group communications.

11.2 Geometric Preliminaries

For any $\rho > 1$, let α_ρ denote the maximum number of points in a unit disk whose mutual distances are greater than $\rho - 1$. For any $\rho \geq 1$, let β_ρ denote the maximum number of points in a half disk of radius $\rho + 1$ whose mutual distances are greater than one. The upper bounds on α_ρ and β_ρ can be derived by the following classic result on disk packing.

Theorem 11.2.1 (Zassebhaus-Groemer-Oler Inequality). *Suppose that S is a compact convex set and U is a set of points with mutual distances at least one. Then*

$$|U \cap S| \leq \frac{\operatorname{area}(S)}{\sqrt{3}/2} + \frac{\operatorname{peri}(S)}{2} + 1,$$

where $\operatorname{area}(S)$ *and* $\operatorname{peri}(S)$ *are the area and perimeter of S, respectively.*

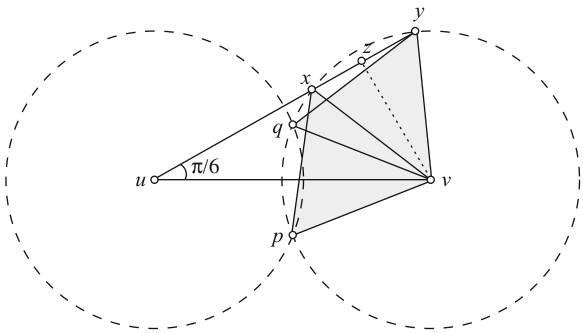

Fig. 11.3 The two circles have unit radius, and $1 \leq \|uv\| \leq 2$. Then, both $\triangle pvx$ and $\triangle qvy$ are equilateral

When the set S is a disk or a half disk, we have the following packing bound.

Corollary 11.2.2. *Suppose that S (respectively, S') is a disk (respectively, half disk) of radius r, and U is a set of points with mutual distances at least one. Then*

$$|U \cap S| \leq \frac{2\pi}{\sqrt{3}} r^2 + \pi r + 1,$$
$$|U \cap S'| \leq \frac{\pi}{\sqrt{3}} r^2 + \left(\frac{\pi}{2} + 1\right) r + 1.$$

By the above corollary,

$$\alpha_\rho \leq \left\lfloor \frac{2\pi/\sqrt{3}}{(\rho-1)^2} + \frac{\pi}{\rho-1} \right\rfloor + 1$$

and

$$\beta_\rho \leq \left\lfloor \frac{\pi}{\sqrt{3}} (\rho+1)^2 + \left(\frac{\pi}{2} + 1\right)(\rho+1) \right\rfloor + 1.$$

Now, we introduce an interesting "equilateral triangle property,"which will be used in the proof of Lemma 11.2.4.

Lemma 11.2.3. *Consider two nodes u and v with $1 \leq \|uv\| \leq 2$. Let p and q be their two intersection points of $\partial D(u)$ and $\partial D(v)$ (see Fig. 11.3). Suppose that x and y are the two intersection points between $\partial D(v)$ and the ray emanating from u which is apart from uv by 30° and is on the same side of uv as q, with x being between u and y. Then $\triangle pvx$ and $\triangle qvy$ are equilateral.*

Proof. Let z be the midpoint of xy. Then, vz is perpendicular to xy and

$$\widehat{xvz} = \arccos \|vz\| = \arccos \frac{\|uv\|}{2} = \widehat{pvu}.$$

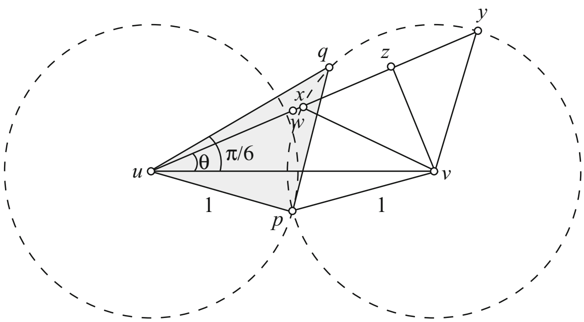

Fig. 11.4 If $\theta \leq 2\arcsin\frac{1}{4}$, then $\|uy\| \geq \|uv\|$, and hence $w \in ux \subset \triangle upq$

Hence,

$$\widehat{pvx} = \widehat{pvu} + \widehat{uvx} = \widehat{xvz} + \widehat{uvx} = \widehat{uvz} = \frac{\pi}{3}.$$

Similarly, $\widehat{yvz} = \widehat{uvq}$ and hence

$$\widehat{qvy} = \widehat{yvz} + \widehat{qvz} = \widehat{uvq} + \widehat{qvz} = \widehat{uvz} = \frac{\pi}{3}.$$

Thus, the lemma follows. □

The next lemma presents two sufficient conditions for the intersection of two unit disks being covered by a third unit-disk.

Lemma 11.2.4. *Consider three nodes u, v, and w satisfying that $1 < \|uv\| \leq 2$ and $\|vw\| > 1$. Then,*

$$D(u) \cap D(v) \subseteq D(w)$$

if one of the following two conditions holds

1. $\|uw\| \leq 1$ *and* $\widehat{vuw} \leq \frac{\pi}{6}$.
2. $1 < \|uw\| \leq \|uv\|$ *and* $\widehat{vuw} \leq 2\arcsin\frac{1}{4} \approx 28.955°$ *(see Fig. 11.4).*

Proof. Let p be the intersection point of $\partial D(u)$ and $\partial D(v)$ which lies on the different side of uv from w, and q be the point on $\partial D(v)$ satisfying that q is on the same side of uv, $\widehat{quv} = \frac{\pi}{6}$ and $\widehat{uqv} \geq \frac{\pi}{2}$ (see Fig. 11.4). By 11.2.3, $\|pq\| = 1$. We will show that w lies in $\triangle upq$. This would imply that $\|pw\| \leq 1$ and consequently

$$D(u) \cap D(v) \subseteq D(w).$$

Under the first condition in the lemma, w lies in $\triangle upq$ obviously. So we assume the second condition in the lemma holds. Let x and y be the intersection points of the ray uw and $\partial D(v)$ with x being closer to u than y (see Fig. 11.4). We claim that $\|uy\| \geq \|uv\|$. Note that

$$\widehat{uyv} = \widehat{vxy} = \widehat{xuv} + \widehat{xvu},$$

$$\widehat{uvy} = \widehat{xvu} + \widehat{xvy}.$$

It is sufficient to show that $\widehat{xvy} \geq \widehat{xuv}$. For the simplicity of presentation, we denote $\widehat{xuv}$ by θ. Let z be the midpoint of xy. Then, z is the perpendicular foot of v on uw, and

$$\begin{aligned}\|vz\| &= \|uv\| \sin\theta \leq 2\sin\theta \\ &= 4\sin\frac{\theta}{2}\cos\frac{\theta}{2} \\ &\leq 4\sin\left(\arcsin\frac{1}{4}\right)\cos\frac{\theta}{2} \\ &= \cos\frac{\theta}{2}.\end{aligned}$$

Hence,

$$\widehat{xvy} = 2\arccos\|vz\| \geq 2\arccos\left(\cos\frac{\theta}{2}\right) = \theta.$$

Thus, our claim holds. So,

$$\|uw\| \leq \|uv\| \leq \|uy\|,$$

which means w is on the line segment uy. As $w \notin D(v)$, w must be on the line segment ux. Consequently, w lies in $\triangle upq$ as well. □

11.3 Dominating Tree

In this section, we describe a rooted spanning tree T of G constructed from a connected dominating set (CDS). This tree will be used in the routings of all the four group communications. Depending on the type of the group communications, the root of T, denoted by s, is chosen as follows. For broadcast, s is the source of the broadcast; for aggregation or gathering, s is the sink node; for gossiping, s is a graph center of G. In either case, we use L to denote the graph radius of G with respect to s.

We begin with the construction of a small, short, and sparse CDS of G. We first select a maximal independent set (MIS) I of G in the first-fit manner in a breadth-first-search (BFS) ordering (with respect to s) of V. All nodes in I form a dominating set, and hence are referred to as dominators. Then, we select a set C of connectors to interconnect I as follows. Let G' be the graph on I in which there is edge between two dominators if and only if they have a common neighbor. The radius of G' with respect to s is denoted by L'. Clearly, $L' \leq L-1$. For each $0 \leq l \leq L'$, let I_l be the set of dominators of depth l in G'. Then, $I_0 = \{s\}$. For each $0 \leq l < L'$, let P_l be the set of nodes adjacent to at least one node in I_l and at least one node in I_{l+1}, and compute a minimal cover $C_l \subseteq P_l$ of I_{l+1} (see an illustration in Fig. 11.5). Set $C = \bigcup_{l=0}^{L'-1} C_l$. Then, $I \cup C$ is a CDS of G. We will prove the following lemma on the sparsity of $I \cup C$.

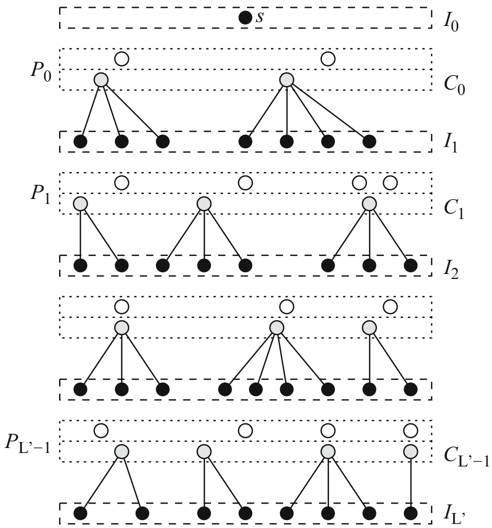

Fig. 11.5 The selection of connectors (marked by *gray*)

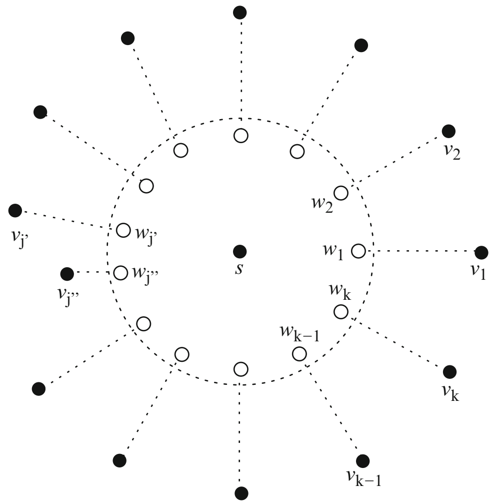

Fig. 11.6 $w_1, w_2, \ldots, w_k$ are the connectors in C_0. Each v_j is a private dominator neighbor of w_j in I_1 with respect to C_0

Lemma 11.3.1. *$|C_0| \leq 12$ and each dominator in I_l with $1 \leq l \leq L' - 1$, is adjacent to at most 11 connectors in C_l.*

Proof. We first prove that $|C_0| \leq 12$. Assume to the contrary that $C_0 = \{w_1, w_2, \ldots, w_k\}$ for some $k \geq 13$. By the minimality of C_0, for each $1 \leq j \leq k$, there is a node $v_j \in I_1$ such that v_j is adjacent to w_j but not to any other node in C_0 (see Fig. 11.6). Among the k nodes $v_1, \ldots, v_k$, there exist two, say $v_{j'}$ and $v_{j''}$, satisfying that $\angle v_{j'} s v_{j''} \leq \frac{2\pi}{13}$. Assume by symmetry that $v_{j''}$ is closer to s than $v_{j'}$. Since the distance between $v_{j'}$ and $v_{j''}$ is greater than one, $D(s) \cap D(v_{j'}) \subseteq D(v_{j''})$ by Lemma 11.2.4. Hence, $w_{j'} \in D(v_{j''})$, which is a contradiction.

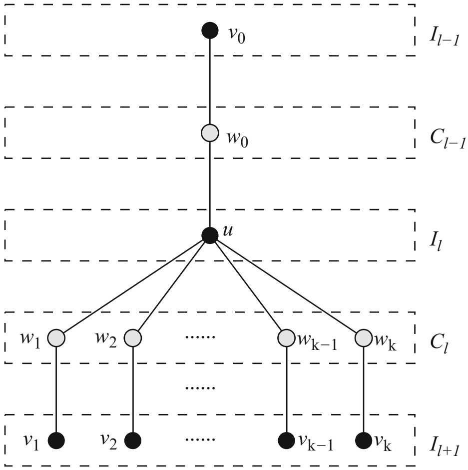

Fig. 11.7 $w_1, w_2, \ldots, w_k$ are the connectors in C_l adjacent to u. Each v_j with $1 \leq j \leq k$ is a private dominator neighbor of w_j in I_{l+1} with respect to C_l

Next, we show that each dominator u in I_l for some $1 \leq l \leq L' - 1$ is adjacent to at most 11 connectors in C_l. Suppose that $w_1, w_2, \ldots, w_k$ are the connectors in C_l which are adjacent to u. By the minimality of C_l, for each $1 \leq j \leq k$, there is a node $v_j \in I_{l+1}$ such that v_j is a private neighbor of w_j with respect to C_l (see Fig. 11.7). Let w_0 be a connector in C_{l-1} which is adjacent to u, and v_0 be a dominator in I_{l-1} which is adjacent to w_0. Then, for each $0 \leq j \leq k$, w_j is the only node in $\{w_0, w_1, \ldots, w_k\}$ which is adjacent to v_j. By the same argument above, we can show that $k + 1 \leq 12$, which implies $k \leq 11$. $\square$

Now, we construct T by specifying the parent of each node other than s. First, each dominator in I_l with $1 \leq l \leq L'$ chooses the neighboring connector of the smallest ID in C_{l-1} as its parent. Second, each connector in C_l with $0 \leq l \leq L' - 1$ chooses the neighboring dominator of the smallest ID in I_l as its parent. Third, each other node, referred to as dominate, chooses the neighboring dominator of the smallest ID as its parent. Clearly, T is a spanning tree and is called a *dominating tree*. Figure 11.8 is an illustration of the construction of T. By the property of the CDS $I \cup U$, the maximum depth of T is at most $2L' + 1 \leq 2L - 1$, s has at most 12 connector children, and each other dominator has at most 11 connector children.

In the remaining of this section, we present a first-fit distance-$(\rho + 1)$ coloring of an arbitrary subset U of dominators. In the lexicographic order of U, all nodes in U are sorted from the left to the right with ties broken by the ordering from the bottom to the top. Suppose that $\langle u_1, u_2, \ldots, u_k \rangle$ is the lexicographic order of U. The first-fit coloring in this order uses colors represented by natural numbers and runs as follows: Assign the color 1 to u_1. For $i = 2$ up to k, assign to u_i with the smallest color not used by any v_j with $j < i$ and $\left\| v_i v_j \right\| \leq (\rho + 1)$. We claim that at most β_ρ colors are used by this coloring. Indeed, consider an arbitrary node $u \in U$. All other nodes in U which precede u and are apart from u by a distance at most $\rho + 1$ lie in the left half disk of radius $\rho + 1$ centered at u. The number of these dominators is at most $\beta_\rho - 1$, where the -1 term is due to that u is also in this half disk. Hence, the color number received by u is at most $\left(\beta_\rho - 1\right) + 1 = \beta_\rho$. Thus, our claim holds.

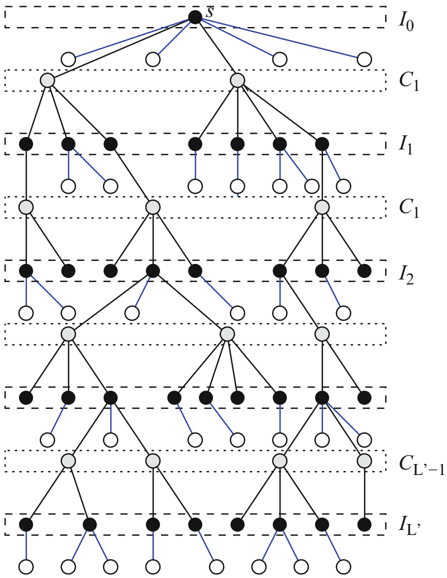

Fig. 11.8 An illustration of dominating tree

11.4 Broadcast Scheduling

Let s be the source of the broadcast. We first construct the dominating tree T rooted at s as in Sect. 11.3. The routing of the broadcast is the spanning s-aborescence oriented from T. The broadcast schedule is then partitioned in $2L' + 1$ rounds sequentially dedicated to the transmissions by

$$I_0, C_0, I_1, C_1, \ldots, I_{L'-1}, C_{L'-1}, I_{L'}$$

respectively. For each $1 \leq l \leq L'$, we compute a first-fit distance-$(\rho + 1)$ coloring of I_l in the lexicographic order. The individual rounds are then scheduled as follows:

- In the round for I_0, only the source node s transmits, and hence this round has only one time slot.
- In the round for C_0, all nodes in C_0 transmit one by one, and thus this round takes at most 12 time slots.
- In the round for I_l with $1 \leq l \leq L'$, a dominator of the ith color transmits in the ith time slot, and hence this round takes at most β_ρ time slots.
- In the round for C_l with $1 \leq l \leq L' - 1$, a connector with a child dominator of the ith color transmits in the ith time slot, and hence this round also takes at most β_ρ time-slots.

Thus, the latency of the entire broadcast schedule is at most

$$\begin{aligned} &1+12+\left(2L'-1\right)\beta_\rho \\ &\le 13+\beta_\rho\left(2L-3\right) \\ &= 2\beta_\rho L-\left(3\beta_\rho-13\right). \end{aligned}$$

Since L is a trivial lower bound on the minimum broadcast latency, the above broadcast schedule is a $2\beta_\rho$-approximation of the optimum.

11.5 Aggregation Scheduling

Let s be the sink of the aggregation. Let Δ denote the maximum degree of G, and L be the graph radius of G with respect to s. For the trivial case that $L=1$, we simply let all nodes other than s transmit one by one. Such trivial schedule has latency $n-1=\Delta$. Subsequently, we assume that $L>1$. We first construct the dominating tree T rooted as s as in Sect. 11.3. The routing of the aggregation schedule is the spanning inward s-aborescence oriented from T. Let W denote the set of dominates. The aggregation schedule is then partitioned in $2L'+1$ rounds sequentially dedicated to the transmissions by

$$W, I_{L'}, C_{L'-1}, I_{L'-1}, \ldots, C_1, I_1, C_0$$

respectively. We describe a procedure used by the scheduling in the round for W and the round for each C_l with $1\le l\le L'-1$.

Let B be a set of links whose receiving endpoints are all dominators. Suppose that ϕ is the maximum number of links with a common dominator endpoint. We first partition into at most ϕ subsets B_j with $1\le j\le \phi$ such that each dominator is incident to at most one link in each B_j. The schedule of B is then further partitioned into ϕ sub-rounds dedicated to $B_1, B_2, , \ldots, B_\phi$, respectively. In the sub-round for B_j, we compute a first-fit distance-$(\rho+1)$ coloring of the dominators incident to the links in B_j, and then all links in B_j whose dominator endpoints receive the ith color are scheduled in the ith time slot. Thus, each of the ϕ consists of at most β_ρ time slots. Hence, the total number of slots is at most $\phi\beta_\rho$.

Now, we are ready to describe the schedule in the individual rounds.

- In the round for W, we adopt the above procedure to produce a schedule in this round. Since each dominator is adjacent to at least one dominate, the maximum number of nodes in W adjacent to a dominator is at most $\Delta-1$. Hence, this round takes at most $(\Delta-1)\beta_\rho$ time slots.
- In the round for C_l with $1\le l\le L'-1$, we also adopt the above procedure to produce a schedule in this round. Since each dominator in I_{l-1} is adjacent to at most 11 connectors in C_l, this round takes at most $11\beta_\rho$ time slots.
- In the round for C_0, all nodes in C_0 transmit one by one, and thus this round takes at most 12 time slots.

- In the round for I_l with $1 \le l \le L'$, we compute a first-fit distance-$(\rho+1)$ coloring of I_l in the lexicographic order and let each dominator with the ith color transmit in the ith time slot. This round takes at most β_ρ time slots.

Thus, the latency of the entire aggregation schedule is at most

$$\begin{aligned}&(\Delta-1)\beta_\rho + 11\beta_\rho\left(L'-1\right) + 12 + L'\beta_\rho \\ &= \Delta\beta_\rho + 12\beta_\rho\left(L'-1\right) + 12 \\ &\le \Delta\beta_\rho + 12\beta_\rho\left(L-2\right) + 12 \\ &= \Delta\beta_\rho + 12\beta_\rho L - 12\left(2\beta_\rho - 1\right).\end{aligned}$$

Since the trivial case takes Δ time slots and $\beta_\rho > 1$, we have the following theorem.

Theorem 11.5.1. *The latency of the above aggregation schedule is at most* $\Delta\beta_\rho + 12\beta_\rho L - 12\left(2\beta_\rho - 1\right)$.

In the next, we present a lower bound on the minimum aggregation latency in terms of Δ.

Lemma 11.5.2. *For any* $\rho > 1$*, the minimum aggregation latency is at least* Δ/α_ρ.

Proof. Let u be a node with maximum degree in G, and S be the unit disk centered at u. Then, S contains $\Delta+1$ nodes. If s is not in S, then all these $\Delta+1$ nodes in S have to transmit; otherwise, exactly Δ nodes in S have to transmit. In either case, at least Δ nodes in S have to transmit. Since all nodes transmitting in the same time slot must be apart from each other by a distance greater than $\rho-1$, at most α_ρ nodes in C can transmit in a time slot. Hence, the Δ transmissions by the nodes in S take at least Δ/α_ρ time slots. □

Since L is also a trivial lower bound on the minimum aggregation latency, the approximation bound of the aggregation schedule is at most

$$\alpha_\rho\beta_\rho + 12\beta_\rho = \left(\alpha_\rho + 12\right)\beta_\rho.$$

11.6 Gathering Scheduling

Let s be the sink of the gathering. If $L=1$, then all other nodes transmit to s one by one, and this schedule is optimal. So, we assume subsequently that $L>1$. We first construct the dominating tree of G rooted at s. The routing of the gathering schedule is the spanning inward s-aborescence oriented from T. Our gather schedule utilizes a labelling of the edges of T, which is described below.

Let $\langle v_1, v_2, \ldots, v_{n-1}\rangle$ be an ordering of $V \setminus \{s\}$ in the descending order of depth in T with ties broken arbitrarily. For $1 \le i \le n$, we assign the jth edge in the tree path from s to v_j with a label $2(i-1)+j$ (see an example in Fig. 11.9). Clearly,

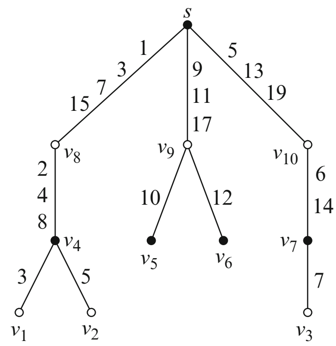

Fig. 11.9 A multi-labelling of the edges in the dominating tree

the number of labels received by an edge connecting v and its parent is equal to the number of descendents (including v itself) of v in T. If v is a connector (respectively, dominator), all labels received by the edge between v and its parent are odd (respectively, even). In addition, all edges across two consecutive layers of the dominating tree receive distinct labels. We further claim that the largest label is $2n-3$. Consider a node v_i and let h be the length of the path from s to v_i. The maximum label assigned to the edges in the path from s to v_i is $2(i-1)+h$. It is sufficient to show that

$$2(i-1)+h \leq 2n-3.$$

Since none of $v_1, v_2, \ldots, v_{i-1}$ belongs to the path from s to v_i, we have

$$h+i-1 \leq n-1.$$

and hence $i \leq n-h$. Therefore,

$$2(i-1)+h \leq 2(n-h-1)+h = 2n-h-2 \leq 2n-3.$$

So, the claim holds.

For each $1 \leq k \leq 2n-3$, let E_k denote the set of edges of T which has been assigned with a label k, and A_k denote the links in the inward s-arborescence oriented from the edges in E_k. Then, for odd (respectively, even) k, all the receiving (respectively, transmitting) endpoints of links in A_k are dominators. In addition, for each $1 \leq k \leq 2n-3$, every dominator is incident to at most one link in A_k.

Now, we are ready to describe the gathering schedule. The schedule is partitioned in $2n-3$ rounds sequentially dedicated to

$$A_{2n-3}, A_{2n-2}, \ldots, A_2, A_1$$

respectively. For each $1 \leq k \leq 2n-3$, the round for A_k is scheduled as follows. We first compute a first-fit distance-$(\rho+1)$ coloring of the dominator endpoints of the links A_k. Then each link whose dominator endpoint receives the ith color is

dominator subframe	connector subframe	dominator subframe	connector subframe	dominator subframe	connector subframe

Fig. 11.10 Framing of the time slots

scheduled in the ith time slot of the kth round. Thus, each round takes at most β_ρ time slots. Consequently, the latency of the gathering schedule is $\beta_\rho\,(2n-3)$. So, we have the following theorem.

Theorem 11.6.1. *The latency of the above gathering schedule is at most* $\beta_\rho\,(2n-3)$.

Since $n-1$ is a trivial lower bound on the minimum gathering latency, the approximation ratio of the gathering schedule presented in this section is at most $2\beta_\rho$.

11.7 Gossiping Scheduling

Let s be the graph center of G. If $L=1$, we adopt the following two-phased schedule. In the first phase, all nodes other than s transmit one by one. This phase takes $n-1$ time slots. In the second phase, the source node transmit all the received packets and its own packet one by one. This phase takes n time slots. So, the total latency is $2n-1$. Clearly, n is a trivial lower bound on the minimum gossiping latency, as every node has to transmits at least once and receive at least $n-1$ times. Thus, its approximation factor is at most 2.

From now on, we assume that $L>1$. Our gossiping schedule consists of two phases. In the first phase s collects all the packets from all other nodes, and in the second phase s broadcasts all the n packets to all other nodes. We adopt the gathering schedule presented in the previous section for the first phase. In the sequel, the node s disseminates all received packets and its own packet to all other nodes. We present a schedule for the second phase in the next.

We first construct the dominating tree T of G rooted at s. The routing of the second phase is the spanning s-aborescence oriented from T. Then, we compute the first-fit coloring distance-$(\rho+1)$ coloring of dominators. Let k be the number of colors used by this coloring. Then, $k\le\beta_\rho$. By proper renumbering of the colors, we assume that s has the first color. We group the time slots into $2k$-slot frames (see Fig. 11.10). In each frame, the first k slots form a dominator subframe, and the remaining k slots form a connector subframe. Only dominators (respectively, connectors) are allowed to transmit in the dominator (respectively, connector) subframe in each frame. Each dominator with color i is only allowed to transmit in the ith slot of a dominator (respectively, connector) subframe. Each connector is only allowed to transmit in the subsets of time slots corresponding to the colors of its child dominators. The source node s transmits one packet in each frame.

Each connector receiving a packet in a dominator subframe transmits the received packet in all the time slots corresponding to the colors of its child dominators of the connector subframe of the same frame. Each dominator with color i receiving a packet in a connector subframe transmits the received packet in the ith time slot of the dominator subframe of the subsequent frame.

The correctness of the above schedule is obvious. Next, we bound the latency of the second phase. After $n-1$ frames, s transmits the last packet. After another L' frames, the last packet reaches all nodes in $I_{L'}$. Finally, after another half frame, the last packet reaches all nodes. So, the total number of time slots taken by the second phase is at most

$$\begin{aligned}
&2k\left(n-1+L'\right)+k \\
&\quad \leq 2k\left(n+L-2\right)+k \\
&\quad = 2k\left(n+L-1.5\right) \\
&\quad \leq 2\beta_\rho\left(n+L-1.5\right).
\end{aligned}$$

By Theorem 11.6.1, the first phase takes at most $\beta_\rho\left(2n-3\right)$ time slots. Hence, the total number of time slots taken by the two phases is at most

$$\begin{aligned}
&\beta_\rho\left(2n-3\right)+2\beta_\rho\left(n+L-1.5\right) \\
&\quad = \beta_\rho\left(4n-6+2L\right).
\end{aligned}$$

Therefore, we have the following theorem.

Theorem 11.7.1. *The latency of the two-phased gossiping schedule is at most* $\beta_\rho\left(4n-6+2L\right)$.

In the next, we present a lower bound on the minimum gossiping latency.

Lemma 11.7.2. *The minimum gossiping latency of G is at least $n-1+L$.*

Proof. The broadcasting of each message requires at least L transmissions. So, the total number of transmissions in any gossiping schedule is at least nL. This implies that some node must take at least L transmissions. On the other hand, every node must take $n-1$ receptions. Therefore, some node takes at least $n-1+L$ transmissions and receptions. This implies that $n-1+L$ is a lower bound on the minimum gossiping latency. □

Since

$$4n-6+2L = 4\left(n-1+L\right)-2\left(L+1\right) < 4\left(n-1+L\right).$$

Therefore, the approximation factor of our gossiping schedule is at most $4\beta_\rho$.