
Department of Electrical Engineering and Computer Science
Case Western Reserve University

Circular Arrangements

V. Liberatore

Technical Report EECS-02-11, 2002

Abstract

Motivated by a scheduling problem in multicast environments, we consider the problem of arranging a weighted graph around a circle so as to minimize the total weighted arc length. We describe the first polynomial-time approximation algorithms for this problem, and specifically an $O(\log n)$ -approximation algorithm for undirected circular arrangements and a $\tilde{O}(\sqrt{n})$ -approximation algorithm for directed circular arrangements. We will show that a simplification of the latter algorithm has better performance than previous heuristics on graphs obtained from a busy Web server log.

1 Introduction

A fundamental issue in server design is to ensure the server ability to support an arbitrarily large number of requests (*scalability*) [PD00]. A scalable server has the tremendous advantage of providing constant service time to clients even during request bursts and peak times. Server scalability can be ensured by *multicast*, whereby a single data unit is duplicated several times within the network infrastructure [Ste98]. Consequently, a single multicast server needs only to send a data unit once to reach an arbitrarily large number of clients, which guarantees the server scalability. Multicast methods have originated research (e.g., [KL00, KSY00, BLMR98]) and commercial companies [DF, HN, PAS]. Applications range from heavily loaded Web servers [AAF98] to high-throughput database systems [HGLW87], and we believe that multicast applications will revolutionize the way users download most Web contents [CLP01].

A fundamental question is the order in which the server should multicast data (*scheduling*). The scheduling problem must take into account that client access patterns often show dependencies between consecutive requests, so that the request for a data unit will make it more likely or less likely that certain other data will be requested next. In this scenario, we have modeled the server data set as a weighted directed graph where nodes represent server data units and arc weights represent the strength of the dependency. The scheduling problem becomes the following question in combinatorial optimization: given a weighted directed graph $G = (N, A)$, arrange the nodes N around a circle so as to minimize the total weighted arc length [Lib02]. We call such question the *directed circular arrangement problem* and show that it is NP-hard. The main objectives of this paper are to

- Present the first polynomial-time approximation algorithms for the circular arrangement problem, and
- Measure algorithm performance on graphs obtained from the workload of a heavily loaded Web site.

Arrangements: Circular and Linear. A problem related to circular arrangements is that of finding *linear arrangements*, where the graph is to be arranged along a line (rather than a circle) so as to minimize the total weighted length of the arrangement. If the graph

is directed, it is further required that the graph be acyclic and the linear arrangement be a topological ordering. The linear arrangement problem is NP-hard [GJ79] and can be approximated to within an $O(\log n)$ factor [RR98]. Furthermore, the minimum linear arrangement problem admits a polynomial-time approximation scheme if the graph is dense and weights are uniform [AFK96]. The linear arrangement problem naturally suggests its analogous on a circle, but, in spite of superficial similarities, the two problems are unrelated as far as the approximation ratio is concerned. Specifically, an optimum linear arrangement can cost $\Omega(n)$ times as much as a circular arrangement of the same directed graph, while an $O(n)$ -approximation algorithm for circular arrangements is trivial [Lib02]. Consequently, the approximation of the two problems is in general unrelated. Indeed, the circular arrangement problem poses a specific technical issue that we discuss next.

Technical Issues. The circular arrangement problem presents two main technical obstacles. The first obstacle is the weakness of the linear relaxation of the problem. Such hurdle is common to other sequencing problems such as linear arrangements [Bak74]. The second difficulty is specific to circular arrangements and is due to the weakness of the divide-and-conquer approach (e.g., [ENRS95, RR98]). Indeed, we will argue that the “divide” cost does not depend only on the problem partition, but also on how each of the individual subproblems is solved. In particular, an arc will “cross” to the other graph component depending not only on how the top level separator is chosen but also on how the nodes are arranged within one of the two components. Due to this issue, we will not be able to use the known divide-and-conquer paradigm directly and we will provide a different approach to partition the original instance into subinstances.

Our Results. In order to attack the directed circular arrangement problem, we will first consider a version of the problem on undirected graphs. As it turns out, the undirected version of the problem is simpler in that it can be reduced to linear arrangements. More to the point, the reduction highlights certain structural properties that will be used for the directed version of the problem. Our result on undirected graphs is that

Theorem 1.1. *There is a polynomial-time $O(\log n)$ -approximation algorithm for the optimum undirected circular arrangement problem.*

The directed version of the problem is the one that is directly applicable to our multicast application. The directed problem will not be reduced to linear arrangements, and, actually, the directed circular arrangement problem does not use in any way results for the linear arrangement problem. However, the algorithm will use certain structural properties that we originally proved for the undirected case. We obtain that

Theorem 1.2. *There is a polynomial-time $\tilde{O}(\sqrt{n})$ -approximation algorithm for the optimum directed circular arrangement problem.*

Finally, we will show that a simplification of the algorithm in Theorem 1.2 has better performance than previous heuristics on graphs obtained from a busy Web server log.




Figure 1: Example of edges that cross vertices. The canonical orientation of e is $(f^{-1}(0), f^{-1}(6))$ ($h(e) = 6$) and the canonical orientation of g is $(f^{-1}(10), f^{-1}(0))$ ($h(g) = 4$). Edge e crosses $f^{-1}(0), f^{-1}(1), \dots, f^{-1}(5)$ and edge g crosses $f^{-1}(10), f^{-1}(11), \dots, f^{-1}(13)$.

Contents The paper is organized as follows. Section 2 contains our results on undirected graphs and section 3 contains our results on directed graphs. Section 4 reports on our algorithm engineering and on our experimental results in the context of multicast-based data dissemination. Several proofs are omitted from this abstract but are available on-line.

2 Undirected Graphs

We first consider the decision version of the optimum undirected circular arrangement problem, which we call the *Undirected Circular Arrangement Problem* (CA).

Instance: An undirected graph $G = (V, E)$, positive arc weights $w(e) \in \mathbb{N}$ for each $e \in E$, and a positive integer K .

Question: Is there a one-to-one function $f : V \rightarrow \{0, 1, \dots, n-1\}$ such that $\sum_{e \in E} w(e)h(e) \leq K$, where $n = |V|$ and $h(\{u, v\}) = \min\{(f(v) - f(u)) \bmod n, (f(u) - f(v)) \bmod n\}$?

Proposition 2.1. *The Undirected Circular Arrangement Problem is NP-complete.*

Proof Sketch. The proof is a reduction from the optimal linear arrangement problem (GT42) [GJ79]. \square

We consider the minimum circular arrangement problem where we seek a solution that minimizes K . We begin with some definitions. Henceforth, $m = |E|$. If $X \subseteq E$ is a set of edges, then we define $w(X) = \sum_{e \in X} w(e)$. We will fix the *canonical orientation* of $\{u, v\} \in E$ in the circular arrangement f to be (u, v) if $(f(v) - f(u)) \bmod n < (f(u) - f(v)) \bmod n$. If $(f(v) - f(u)) \bmod n = (f(u) - f(v)) \bmod n$, we fix arc orientation arbitrarily. By definition, $h((u, v)) = (f(v) - f(u)) \bmod n$. We will say that an edge (u, v) *crosses* $x \in V$ in an arrangement f if $(f(x) - f(u)) \bmod n < h(e)$. Hence, $e = (u, v)$ crosses exactly the $h(e)$ vertices in $C((u, v)) = \{x : (f(x) - f(u)) \bmod n < h(e)\}$. An example of these definition is shown in Fig. 1.

We will now turn to a series of lemmata on properties of edges crossing vertices, and we postpone their proofs to the full paper.

Lemma 2.2. For any $e \in E$, $h(e) \leq n/2$.

Lemma 2.3. Suppose that e crosses v , and let u be the vertex that e does not cross and that minimizes the quantity $(f(u) - f(v)) \bmod n$. Then, $u \in e$.

Lemma 2.4. The cost of a circular arrangement f is equal to $\sum_{v \in V} \rho_v$, where $\rho_v = \sum_{e: v \in C(e)} w(e)$ is the total weight of edges that cross v .

Definition 2.1. The nodes v_1, v_2, \dots, v_l are placed next to each other in a circular arrangement f if $f(v_{i+1}) = (f(v_i) + 1) \bmod n$ ($1 \leq i < l$).

Lemma 2.5. Let $U = \{v_1, v_2, \dots, v_l\} \subseteq V$ with $f(v_{i+1}) = (f(v_i) + 1) \bmod n$ ($1 \leq i < l$), so that vertices in U are placed next to each other. If $e \cap U, e - U \neq \emptyset$, then e crosses either v_l or $f^{-1}((f(v_1) - 1) \bmod n)$.

Definition 2.2. Let $G = (V, E)$ be a graph, $U \subseteq V$, and f a circular arrangement of G . A vertex $u \in U$ is the U -successor of a vertex v in the arrangement f if and only if $(f(u) - f(v)) \bmod n \leq (f(x) - f(v)) \bmod n$ for all $x \in U$.

The U -successor of node v in f will be denoted by $s_{f,U}(v)$, or simply as $s(v)$ when f and U are clear from the context. It is immediate to see that $s(v)$ is unique and belongs to U .

Definition 2.3. Let $G = (V, E)$ be a graph, $U \subseteq V$, and f a circular arrangement of G . A vertex $u \in U$ is the U -predecessor of a vertex v in the arrangement f if and only if $v = s_{f,U}^{-1}(s_{f,U}(v))$.

The following set $I(v)$ basically represents all vertices between v and its successor $s(v)$.

Definition 2.4. Let $G = (V, E)$ and $v \in V$. Let $d_{f,U}(v) = (f(s(v)) - f(v)) \bmod n$ and $I_{f,U}(v) = \{u : 0 < (f(u) - f(v)) \bmod n < d(v)\}$.

Subscripts will be omitted when f and U are clear from the context. The following easy properties are easily proven:

Lemma 2.6. The following properties hold on $d(v)$ and $I(v)$ for any f and U :

1. $d(v) = |I(v)| + 1$
2. $I(v) \cap U = \emptyset$
3. For any $U \subseteq V$, $\sum_{v \in U} d(v) = n$.
4. $I(v) \cap I(u) = \emptyset$ for all $u, v \in U$, $u \neq v$.

Definition 2.5. A graph $G = (V, E)$ is said to be the *unbiased union* of graphs $G_1 = (V_1, E_1), G_2 = (V_2, E_2), \dots, G_k = (V_k, E_k)$ if and only if G is the union of G_1, G_2, \dots, G_k , $V_i \cap V_j = \emptyset$, and $|V_i| \leq |V|/2$ ($1 \leq i, j \leq k$, $i \neq j$).

Lemma 2.7. *Let $G = (V, E)$ be the unbiased union of $G_1 = (V_1, E_1), G_2 = (V_2, E_2), \dots, G_k = (V_k, E_k)$ and f any circular arrangement of G . If an edge $e \in E_j$ crosses in f a vertex $v \in V_j$, then e crosses all vertices in $I_{f, V_j}(v)$.*

Proof. Let $n_j = |V_j|$ for $1 \leq j \leq k$. Since $v \in V_j$, we have that $n_j > 0$. We assume that $I(v) \neq \emptyset$ or else there is nothing to prove. The proof proceeds by contradiction, and so we also assume that there is a vertex in $I(v)$ that e does not cross. Let $u \in I(v)$ be the vertex that e does not cross and that minimizes the quantity $(f(u) - f(v)) \bmod n$. Hence, Lemma 2.3 gives $u \in e \subseteq V_j$, which contradicts the claim $I(v) \cap V_j = \emptyset$ in Lemma 2.6. We conclude that e crosses all vertices in $I(v)$. \square

Lemma 2.8. *Let $G = (V, E)$ be the unbiased union of $G_1 = (V_1, E_1), G_2 = (V_2, E_2), \dots, G_k = (V_k, E_k)$. Let $C_j(e) = C(e) \cap V_j$. Let $\rho_v^j = \sum_{e \in E_j: v \in C_j(e)} w(e)$ be the total weight of edges in E_j that cross $v \in V_j$. Then, the cost of the circular arrangement is at least $\sum_{j=1}^k \sum_{v \in V_j} \rho_v^j d(v)$.*

Proof Sketch. The proof applies Lemma 2.6 and 2.7. \square

The following lemma is critical in order to compare algorithm and optimum costs.

Lemma 2.9. *Let $G = (V, E)$ be the unbiased union of $G_1 = (V_1, E_1), G_2 = (V_2, E_2), \dots, G_k = (V_k, E_k)$. Then, any minimum circular arrangement places the vertices of V_j next to each other ($1 \leq j \leq k$).*

Proof. The proof is by contradiction: if a circular arrangement f does not place the vertices in V_j next to each other, then we will show an arrangement g that costs less than f . The intuition is that g will preserve the same ordering as f among vertices in V_j , but it will also place them next to each other. The point where the V_j are squeezed around has to be chosen so as to minimize the value of ρ_v^j .

Fix vertices $v_j = \arg \min_{u \in V_j} \rho_u^j$ ($1 \leq j \leq k$). Consider the following circular arrangement:

$$g(v) = \begin{cases} \sum_{i=1}^j n_i - 1 & \text{if } v = v_j \\ g(s_f(v)) - 1 & \text{otherwise} \end{cases}.$$

It is immediate to see that g establishes a one-to-one mapping $g : V \rightarrow \{0, 1, \dots, n-1\}$ and that g places the vertices of V_j 's next to each other. Moreover, $s_f(v) = s_g(v)$ for all $v \in V$.

We claim that in g an edge $e = (x, y) \in E_j$ crosses only vertices in $V_j - \{v_j\}$. Let $\tilde{V}_j = V - (V_j - \{v_j\})$. By definition of unbiased union, $|\tilde{V}_j| > n/2$. By contradiction, $C(e) \cap \tilde{V}_j \neq \emptyset$. First, suppose that $C(e) \supseteq \tilde{V}_j$. Then, $h(e) = |C(e)| \geq |\tilde{V}_j| > n/2$, which contradicts Lemma 2.2. Next, suppose that $\tilde{V}_j - C(e) \neq \emptyset$ (and again $C(e) \cap \tilde{V}_j \neq \emptyset$). By definition of crossing, e crosses $x \in V_j$. Let u be the vertex in $\tilde{V}_j - C(e)$ that minimizes $(g(u) - g(x)) \bmod n$. By Lemma 2.3 and $x \in C(e)$, we have that $u = y$, and since $e \subseteq V_j$, we also have $u = y = v_j$. However, by definition of crossing, e does not cross y , which is a contradiction. Thus, e crosses in g only vertices that belong to $V_j - \{v_j\}$.

We now introduce a charging scheme that will assign a cost to each node so as to upper bound the actual ρ_v^j relative to the ordering g . The scheme has two cases, depending on

whether e originally crosses v_j in f . Assume first that e does not cross v_j in f . Since e does not cross v_j in g , and, $s_f(v) = s_g(v)$ for all $v \in V$, e crosses exactly the same vertices $U \subseteq V_j$ in g as in f . In the charging scheme, the quantity ρ_v^j is charged $w(e)$ for all nodes $v \in U$. Assume on the other hand that $e = \{u, v\}$ crosses in f a set of vertices U with $v_j \in U \subseteq V_j$. The edge e does not cross v_j in g , and so it crosses $V_j - U$ at a cost of $w(e)|V_j - U|$. The charging scheme is that if edge e crosses in f a set U with $v_j \in U \subseteq V_j$, we charge $w(e)$ to U and $n_j - 1$ to v_j , so that the total charge is $w(e)(n_j - 1 + |U|) \geq w(e)|V_j| > w(e)|V_j - U|$. The charging scheme allows us to assume that each ρ_v^j remains unchanged from f to g for all $v \neq v_j$, and v_j is charged $(n_j - 1)\rho_{v_j}^j$. In this scheme, the cost change within G_j is no more than

$$\begin{aligned}
\sum_{v \in V_j - \{v_j\}} \rho_v^j + (n_j - 1)\rho_{v_j}^j - \sum_{v \in V_j} \rho_v^j d(v) &= (n_j - 2)\rho_{v_j}^j - \sum_{v \in V_j} \rho_v^j (d(v) - 1) \\
&\leq (n_j - 2)\rho_{v_j}^j - \rho_{v_j}^j \sum_{v \in V_j} (d(v) - 1) \\
&= (n_j - 2)\rho_{v_j}^j - \rho_{v_j}^j (n - n_j) \\
&= \rho_{v_j}^j (2(n_j - 1) - n) \\
&< 0.
\end{aligned}$$

because $\rho_{v_j}^j \leq \rho_v^j$, $\sum_{v \in V_j} d(v) = n$, and $n_j \leq n/2$. Hence, g costs less than f , a contradiction is reached, and the lemma is proven. \square

We define a β -separator of $G = (V, E)$ as a partition of V into V_1 and V_2 with the property that $\min\{|V_1|, |V_2|\} \geq \beta n$. A merely technical assumption that we will use later on is that since $\lfloor n/2 \rfloor \geq \min\{|V_1|, |V_2|\} \geq \lceil \beta n \rceil$, we can assume that $\lfloor n/2 \rfloor \geq \lceil \beta n \rceil$. For example, if $\beta = 2/5$, then $n \geq 4$. The cost of the separator is defined as the cost of the edges that cross from V_1 to V_2 . An optimum β -separator has minimum cost among all β -separators. The problem of finding a β -separator of small cost is NP-hard even when all edges have the same weight [LR99].

Theorem 2.10 ([LR99]). *There exists a polynomial-time algorithm that, given a weighted, undirected graph G , finds a $(1/3)$ -separator whose cost is $\alpha = O(\log n)$ times the cost of an optimum $(2/5)$ -separator.*

Lemma 2.11. *Let G be a graph with more than $n = 4$ vertices. Let s^* be the cost of an optimum $(2/5)$ -separator and a^* the minimum cost of a circular arrangement. Then, $s^* \leq (2a^*)/n$.*

Proof Sketch. Our proof is reminiscent of Hansen's for the case of linear arrangements [Han89]. However, we also need to apply the structural properties expressed by Lemma 2.3 and 2.4. \square

A component of our algorithm will be an approximate solution for the *minimum linear arrangement* problem: given an undirected graph $G = (V, E)$ with positive integer edge

weights $w(e)$ for all $e \in E$, find a one-to-one correspondence $f : V \rightarrow \{1, 2, \dots, |V|\}$ that minimizes $\sum_{e=\{u,v\} \in E} w(e)|f(v) - f(u)|$.

We can now state our approximation algorithm. First, we find a $(1/3)$ -separator as in Theorem 2.10, and denote by H_1 the smallest of the two sets of vertices. We remove H_1 and all edges incident on H_1 , and find a $(1/3)$ -separator for the rest of the graph. We call the smallest vertex set of this second separator H_2 and let $H_3 = V - H_1 - H_2$. The decomposition is illustrated in Fig. 2. Observe that $|H_3|/n \leq (2/3)^2 < 1/2$. We now run the $O(\log n)$ -




Figure 2: Decomposition of a graph into the components H_1 , H_2 , and H_3 .

approximation algorithm for minimum linear arrangement independently on H_1, H_2, H_3 to obtain linear orderings f_1, f_2, f_3 . We finally patch the three linear orderings together:

$$f(v) = \begin{cases} f_1(v) - 1 & \text{if } v \in H_1 \\ f_2(v) + |H_1| - 1 & \text{if } v \in H_2 \\ f_3(v) + |H_1| + |H_2| - 1 & \text{otherwise} \end{cases}.$$

The analysis of the algorithm is based on the following definitions and general considerations.

Definition 2.6. If $X \subseteq A$ is a set of arcs and f a circular arrangement, then the *contribution of X to the cost of f* is $c_f(X) = \sum_{e \in X} w(e)h(e)$.

The subscript f will be omitted when the circular arrangement is clear from the context. The main idea of the analysis is as follows. Consider a subgraph $G' = (V, E')$ of $G = (V, E)$ on the same vertex set V . Then, the minimum circular arrangement of G' does not cost more than the cost a^* of the minimum circular arrangement of G . Our algorithm considers a sequence of graphs $G_0, G_1 = (V, E_1), \dots, G_3 = (V, E_3)$. We define $G_0 = G$, $G_1 = (V, E_1)$ as the graph obtained from G by removing edges that are incident on both H_1 and $V - H_1$, $G_2 = (V, E_2)$ as the graph obtained by removing from E_1 edges that are incident on both H_2 and H_3 , and $G_3 = (V, \emptyset)$ (see Fig. 2). The decomposition has the properties that $E_i \subset E_{i+1}$ ($0 \leq i < k$). We will show that $c_f(E_i - E_{i+1}) \leq \alpha a^*$ for some $\alpha = O(\log n)$. As a result, the algorithm achieves an $O(\log n)$ approximation factor.

Lemma 2.12. *The cost a^* of the optimum circular arrangement is at least equal to the sum of the costs of the optimum linear arrangements of H_1, H_2, H_3 (i.e., $c(E_2 - E_3) \leq a^*$).*

Proof. If we remove the edges that are not contained within a single H_i ($1 \leq i \leq 3$), we do not increase the optimum cost. By Lemma 2.9, the optimum algorithm on the resulting graph places vertices in H_i next to each other. Since $|H_i| \leq n/2$, an edge $e \subseteq H_i$ crosses only vertices in H_i . Therefore, the cost of the arrangement is equal to the sum of the costs of the three independent linear arrangements, which proves the claim. \square

Lemma 2.13. *For the edge sets E_0 and E_1 defined above, $c(E_0 - E_1) \leq \alpha a^*$ for some $\alpha = O(\log n)$.*

Proof Sketch. Apply Theorem 2.10 and Lemma 2.11. \square

Lemma 2.14. *For the edge sets E_1 and E_2 defined above, $c(E_1 - E_2) \leq \alpha a^*$ for some $\alpha = O(\log n)$.*

We can now prove the main result of this section:

Theorem 1.1. By summing up the identities in the three previous lemmata, we obtain that

$$c(E) = c(E - E_1) + c(E_1 - E_2) + c(E_2) = \alpha a^* ,$$

for some $\alpha = O(\log n)$, which proves the theorem. \square

3 Directed Graphs

We consider the following decision version of the optimum circular arrangement problem, which we call the *Directed Circular Arrangement Problem* (DCA):

Instance: A directed graph $G = (N, A)$, non-negative arc weights $w(e) \in \mathbb{N}$ for each $e \in A$, and a positive integer K .

Question: Is there a one-to-one function $f : N \rightarrow \{0, 1, \dots, |N| - 1\}$ such that, for $n = |N|$,

$$\sum_{e=(u,v) \in A} w(e) h(e) \leq K ,$$

where $n = |N|$ and $h((u, v)) = (f(v) - f(u)) \bmod n$?

Proposition 3.1 ([Lib02]). *The Directed Circular Arrangement Problem is NP-complete.*

Proof Sketch. The proof is a reduction from the directed optimal linear arrangement problem (GT43 in [GJ79]). \square

We consider the optimum circular arrangement problem where we seek a solution that minimizes K . In order to simplify our notation, we make the following assumptions without loss of generality. We assume that G contains no loops as they do not contribute to the solution cost, and that $m = |A| = n(n-1)$. If $X \subseteq A$ is a set of arcs, then we define $w(X) = \sum_{e \in X} w(e)$. We will say that an arc (u, v) crosses node $x \in N$ if $(f(x) - f(u)) \bmod n < (f(v) - f(u)) \bmod n$. Hence, arc $e = (u, v)$ crosses exactly the $h(e)$ vertices in $C((u, v)) = \{x : (f(x) - f(u)) \bmod n < (f(v) - f(u)) \bmod n\}$.

As discussed in the introduction, the divide-and-conquer approach (e.g., [ENRS95, RR98]) fails in that the “divide” cost depends on how the individual subproblems are solved. Specifically, suppose that the directed circular arrangement problem on a graph G is decomposed to the subgraphs H_1, H_2, \dots, H_k . In the divide-and-conquer approach, we would ordinarily divide the cost into arcs that cross between two H_i ’s, and arcs that have both endpoints within a single H_i . Consider an arc (u, v) with u, v belonging to the same subgraph H_i . If we arrange $f(v) = (f(u) + 1) \bmod n$, then indeed arc (u, v) crosses only nodes in H_i . However, if $f(u) = (f(v) + 1) \bmod n$, then (u, v) can cross nodes that are not in H_i . As a result, the contribution of an arc depends not only on the node set partition, but also on how nodes are arranged within a single partition. Thus, this fact is a major obstacle to a divide-and-conquer approach. Our solution will balance partitions so as to sidestep this problem.

Definition 3.1 ([AMO93]). Let $G = (N, A)$ be a directed graph. The *head* of an arc (u, v) is node v and the *tail* is node u .

Lemma 3.2. *Suppose that e crosses u , and let u be the node that e does not cross and that minimizes the quantity $(f(u) - f(v)) \bmod n$. Then, v is e ’s head.*

Lemma 3.3. *The cost of a circular arrangement f is equal to $\sum_{v \in V} \rho_v$, where $\rho_v = \sum_{e: v \in C(e)} w(e)$ is the total weight of arcs crossing v .*

Lemma 3.4. *Let $G = (N, A)$ be a directed graph, and assume that the node set N can be partitioned into N_1, N_2, \dots, N_k with the property that for $e \in A$ there is a unique N_j ($1 \leq j \leq k$) on which e is incident. Then, there is a minimum circular arrangement that places the vertices of N_j next to each other ($1 \leq j \leq k$).*

The major differences between Lemma 2.9 and Lemma 3.4 are that $|N_j|$ can be larger than $n/2$, and that a property of one optimum arrangement is described, rather than a property of all optimum arrangements.

Proof Sketch. The proof is similar to that of Lemma 2.9, with the following difference. The arcs that cross v_j do not cross only nodes of N_j in g , but have to circumnavigate around the nodes in $N - N_j$. \square

A first component of our algorithm is a polynomial-time approximation algorithm for the *minimum feedback arc set* problem: given a directed graph $G = (N, A)$ with non-negative integer arc weights $w(e)$ for all $e \in A$, find an arc set of minimum total weight that

intersects every directed cycle in the graph G . The minimum feedback arc set problem is APX-complete [Kan92] and has a polynomial-time algorithm that approximates the optimum solution to within an $O(\log n \log \log n)$ factor [ENSS98]. Another component of our algorithm is expressed by the following lemma, which will give the termination condition of a sequence of separator computations.

Lemma 3.5. *Let $G = (N, A)$ be a directed acyclic graph and assume that the node set N can be partitioned into N_1, N_2, \dots, N_k with the property that for $e \in A$ there is a unique N_j ($1 \leq j \leq k$) on which e is incident. Let A_j be the set of arcs incident on N_j and consider the cost ℓ_j of any linear arrangement of the graph (N_j, A_j) . Then, $\ell_j \leq (|N_j| - 1)c_{f^*}(A_j)$.*

Proof. Let $W = \sum_{e \in A_j} w(e)$ and observe that $c_{f^*}(A_j) \geq W$ because A_j contains no loop. Meanwhile, a linear arrangement of N_j pays at most $W(|N_j| - 1)$, thus proving the lemma. \square

The main ideas of the proof are as follows. Consider a subgraph $G' = (N, A')$ of $G = (N, A)$ on the same node set N . Then, the minimum circular arrangement of G' costs no more than the minimum circular arrangement of G . Our algorithm will then consider a sequence of graphs $G_0, G_1 = (N, A_1), \dots, G_k = (N, A_k)$ with the properties that $G_0 = G$, that $A_{i+1} \subset A_i$ ($0 \leq i < k$), that $A_k = \emptyset$, and that $c_f(A_i - A_{i+1}) \leq \beta a^*$ for some $\beta \in \tilde{O}(\sqrt{n})$ and where a^* is the minimum cost of a circular arrangement of G . As a result, the algorithm achieves an $\tilde{O}(k\sqrt{n})$ approximation factor. Unlike the undirected case, we cannot show a decomposition with $k = O(1)$. Instead, we will remove a feedback arc set to make G acyclic, then we will keep partitioning the graph until the size of each connected components is so small that we can arrange those components with any topological ordering. Such decomposition guarantees a value of $k = O(\log n)$, so that the approximation factor will be $\tilde{O}(\sqrt{n})$. The graph $G_i = (N, A_i)$ will be partitioned into components $X_1, X_2, \dots, X_p \subseteq N$ with the property that no arc is incident on more than one X_j . In other words, the X_j 's contain weakly connected components of G_i . The cost of a circular arrangement of G_i is the sum of the costs due to the X_j 's: $c_f(A_i) = \sum_{j=1}^p \sum_{e=(x,y): x,y \in X_j} w(e)h(e)$. For the purpose of the analysis, we will sometimes consider a subgraph G'' obtained by removing nodes from a G_i along with all incident arcs and observe that a minimum circular arrangement of G'' costs no more than that of G_i . We will also need the following definition.

Definition 3.2. Given a complete weighted directed graph $G = (N, A)$, the *graph obtained by removing arc orientation from G* is the complete weighted undirected graph $G' = (N, E)$ with weights $w(\{u, v\}) = w((u, v)) + w((v, u))$ ($u, v \in N, u \neq v$).

We can now state our approximation algorithm. In the first step, we remove arc orientation from G . Then, we compute H_1, H_2, H_3 exactly as in the algorithm for undirected graph. At this point, $|H_1|, |H_2|, |H_3| < n/2$. Let G_1 be the directed graph obtained by removing from G arcs incident on both H_1 and $N - H_1$ and G_2 be the directed graph obtained by removing from G_1 arc incident on both H_2 and $N - H_2$. Observe that H_1, H_2, H_3 have the property that they contain weakly connected components of G_2 . Our algorithm then computes a feedback arc set for G_2 with an $O(\log n \log \log n)$ -approximation algorithm

[ENSS98], removes the feedback arcs, and obtains a graph G_3 . At this point, G_3 is acyclic and we start a procedure to be described below that finds $(1/3)$ -separators for the H_j 's while component size is bigger than a parameter $\theta = \sqrt{n \log n}$. At any step of this recursive procedure, the graph $G_i = (N, A_i)$ has the properties that the node set N can be partitioned into X_1, X_2, \dots, X_p and that arcs are incident on a unique X_j . In other words, the X_j 's have the property that they contain the weakly connected components of G_i . Initially, G_3 has $X_j = H_j$ ($j = 1, 2, 3$). To obtain G_{i+1} from G_i , we remove arc orientation from G_i , and for each X_j , if $|X_j| \geq \theta$, we find a $(1/3)$ -separator of the undirected subgraph induced by X_j , and correspondingly break X_j into X_{j1} and X_{j2} . Thus, the previous X_j 's contain the new X_j 's. The procedure continues until $|X_j| < \theta$ ($1 \leq j \leq p$).

Lemma 3.6. *Let $G = (N, A)$ be a directed acyclic graph and assume that the node set N can be partitioned into N_1, N_2, \dots, N_k with the property that for $e \in A$ there is a unique N_j ($1 \leq j \leq k$) on which e is incident. Let A_j be the set of arcs incident of N_j . Then, $c_f(A) = \sum_{j=1}^k c_f(A_j)$.*

Lemma 3.7. *The total contribution to the circular arrangement cost of the arcs that have endpoints in different H_j 's is no more than an $O(\log n)$ factor away from the cost a^* of an optimum circular arrangement (i.e., $c_f(A - A_2) \leq \alpha a^*$ for some $\alpha = O(\log n)$).*

Proof Sketch. Similar to the proof of Theorem 1.1, except that $h(e)$ can be as high as $n - 1$. However, this modification only doubles the constant factor. \square

Lemma 3.8. *The total contributions to the circular arrangement cost of the arcs in the feedback arc set is no more than an $O(\log n \log \log n)$ -factor away from the cost a^* of an optimum circular arrangement (i.e., $c_f(A_2 - A_3) \leq \gamma a^*$ for some $\gamma = O(\log n \log \log n)$).*

Proof. The minimum cost of a circular arrangement of G_2 is not larger than the minimum cost a^* of a circular arrangement of G . By Lemma 3.4, we can assume without loss of generality that the optimum arrangement f^* of G_2 places nodes in H_j next to each other. Furthermore, we can assume without loss of generality that the optimum circular arrangement f^* places the nodes of H_1 in $1, 2, \dots, |H_1|$. Let A' be the set of arcs with both endpoints in H_1 , A^a be the feedback arc set that the approximation algorithm finds for H_1 , and A^c be the set of arcs $(x, y) \in A'$ that cross nodes outside H_1 in f^* . Finally, let $G' = (H_1, A')$ and A^* be the optimum feedback arc set on G' . Arcs in $A' - A^c$ link nodes (x, y) with $f^*(y) > f^*(x)$, and so f^* is a topological ordering of $(H_1, A' - A^c)$. Therefore, A^c contains a feedback arc set of G' , so that $w(A^c) \geq w(A^*)$. Moreover, $h(e) \geq |H_2 \cup H_3| \geq n/2$ for all $e \in A^c$, so that $a^* \geq c_{f^*}(A^c) \geq nw(A^*)/2$. We now have that $w(A^a) \leq \gamma w(A^*)$, for some $\gamma = O(\log n \log \log n)$. Let f be the algorithm's circular arrangement, and $c_f(A^a) \leq nw(A^a) \leq n\alpha w(A^*) \leq 2\alpha a^*$. Similar arguments for H_2 and H_3 conclude the proof. \square

We now start with the DAG G_2 and, for each of components, H_1, H_2, H_3 , we compute an approximate $(1/3)$ -separator. We continue breaking the components with approximate $(1/3)$ -separators until component size is no more than a threshold $\theta = \sqrt{n \log n}$, at which point we arrange nodes within each component in topological order.

Lemma 3.9. *Let G be a directed graph with more than $n = 4$ vertices and G' be undirected graph obtained from G by removing arc orientation. Let s^* be the cost of an optimum $(2/5)$ -separator of G' and a^* the minimum cost of a circular arrangement of G . Then, $s^* \leq (2a^*)/n$.*

Corollary 3.10. *Let $G = (N, A)$ be a weighted directed graph with a node subset $H \subseteq N$ such that $|H| \geq 4$ and there is no arc that is incident on both H and $N - H$. Let A' be the set of arcs incident on H . Let $G' = (H, E')$ be the undirected subgraph obtained by removing arc orientation from the subgraph of G induced by H . Let s^* be the cost of an optimum $(2/5)$ -separator of G' . Then, $s^* \leq 2c_{f^*}(A')/|H|$.*

Proof. Consider $G'' = (H, A'')$, the weighted directed subgraph of G induced by G . We have that $c_{f^*}(A')$ is at least equal to the minimum cost of a circular arrangement of (N, A') , which in turn is at least equal to the minimum cost of a circular arrangement of (H, A') . By the previous lemma, s^* is at most $|H|$ times the minimum cost of a circular arrangement of (H, A') , which concludes the proof. \square

Lemma 3.11. *The arcs eliminated during each stage of the procedure above contribute to the arrangement cost no more than $O(\sqrt{n \log n})$ times the optimum cost a^* of a circular arrangement (i.e., $c_f(A_i - A_{i+1}) \leq \sqrt{n \log n}$).*

Proof. Consider a graph $G_i = (N, A_i)$ that has been decomposed into X_1, X_2, \dots, X_p . Let B_j be the set of arcs in $A_i - A_{i+1}$ that are incident on X_j . Thus, $A_i - A_{i+1} = \cup_{j=1}^p B_j$. We will show that $c_f(B_j) \leq \theta c_{f^*}(B_j)$ where $\theta = \sqrt{n \log n}$, and so

$$c_f(A_i - A_{i+1}) = \sum_{j=1}^p c_f(B_j) \leq \theta \sum_{j=1}^p c_{f^*}(B_j) = \theta c_{f^*}(A_i - A_{i+1}),$$

where the first and last equality follows from Lemma 3.6. It remains to show that $c_f(B_j) \leq \theta c_{f^*}(B_j)$. Consider a subset X_j of nodes found at some stage by the separator procedure above. If $|X_j| < \theta$, then the cost of arranging the arcs incident on X_j is less than θ times their contribution to the optimum cost by Lemma 3.5. If $|X_j| \geq \theta$, then we find a separator that costs at most αs^* , where $\alpha = O(\log n)$ and s^* is the cost of an optimum $(2/5)$ -separator of the undirected subgraph induced by X_j . By Corollary 3.10, $c_f(B_j) \leq n\alpha s^* \leq 2\alpha n c_{f^*}(B_j)/|X_j|$ which is no more than θ times $c_{f^*}(B_j)$, and so the lemma is proven. \square

Theorem 1.2. The discussion above and Lemmata 3.7, 3.8, and 3.11 yield a $\tilde{O}(k\sqrt{n})$ -approximation algorithm. Since $(1/3)$ -separators are guaranteed to split node subsets into subset containing at least $1/3$ of the original number of nodes, the splitting procedure terminates after $k = O(\log \theta) = O(\log n)$ steps, thus yielding the theorem. \square

4 Algorithm Engineering

We have implemented a fast version of our algorithm for directed graphs as well as three previous algorithms for the same problem. The approximation algorithm in section 3 requires the computation of feedback arc sets and balanced separators, which in turn require

the solution of multicommodity flow problems and the computation of spreading metrics. Furthermore, spreading metrics require the solution of several linear programs with an exponential number of constraints by means of the ellipsoid method [ENSS98, Hač79]. Thus, we do not believe that the algorithm can be practical in its original form and we made the following simplifications. First, we do not directly compute balanced separators. Instead, we find all the minimum $s - t$ cuts in the graph, which results in a separator tree [CH91]. Then, we start from the largest edge in the separator tree and we repeatedly add the largest incident edge to the current component until we isolate a set of vertices of the desired size. By construction, the resulting cut is a balanced separator and the cut value should not be too large as the separator edges have greedily joined the chosen component. Thus, we believe that the proposed heuristic is a reasonable way to compute a separator. The second simplification is in the calculation of the feedback arc set, where we replace an approximation algorithm with a greedy heuristic [CSS99]. The resulting algorithm is denoted as *Bsep*. In addition to the Bsep algorithm, we experimented with the following three algorithms. The first algorithm (*Rnd*) is a random ordering of the graph nodes. The second algorithm is based on maximum spanning trees (*MST*) to cluster nodes close to each other [Lib02]. The final algorithm originates from the observation that the heuristic for feedback arc set computes a linear ordering of the underlying graph and then discards all arcs that violate the topological ordering, that is, all arcs that link higher-numbered vertices to lower-numbered vertices [CSS99]. Thus, such heuristic can also be used to obtain directly a circular arrangement of the original graph, and it is the last algorithm that we consider in this paper (*Fesh*). We also considered the application of integer linear programming to solve the circular arrangement problem exactly. However, after two weeks of computation on a Sun Ultra60, Cplex had not found any integer solution other than the one that we manually inserted as a starting point, even when clique cuts and ordered sets were added. We explain such behavior by noticing that the linear relaxation of the integer formulation places an equal fraction of every node in each position of the circular arrangement, and that branch-and-cut algorithms suffer when the relaxation does not effectively direct the selection process.

We have tested the algorithms in the context of multicast-based data dissemination on directed graphs obtained from the access pattern to the Web server of the World Cup 98. The World Cup trace includes more than one billion requests over a period of 1 1/2 month and is one of the largest trace analyzed to date [AJ99]. Furthermore, the World Cup servers received up to 10 million requests per hour. As a result, the World Cup site is one of the most busy recorded so far, which makes it an ideal testbed for multicast data dissemination. Additional information on this workload can be found in [AJ99, Lib02].

The four algorithms are compared in table 1. The *Rnd* cost is roughly $n/2$, and in most cases the other algorithms improve over it. The *MST* and *Fesh* algorithm had comparable cost on graph 3 and 4, but *MST* is better than *Fesh* on graph 1 and *Fesh* is better than *MST* on graph 2. Thus, we believe that in general the performance of *MST* and *Fesh* is roughly comparable. Finally, the *Bsep* algorithm had the best performance in three graphs out of four, and on the other graph, its performance trails the best algorithm by 3%. Thus, we believe that *Bsep* is overall the best algorithm in terms of circular arrangement cost.

	Graph 1	Graph 2	Graph 3	Graph 4
n	225	174	206	200
Rnd	112.2294 ± 0.7440	87.0159 ± 1.2416	102.7072 ± 0.5072	100.0331 ± 1.5135
MST	107.26	79.6681	97.6602	103.693
Fesh	99.0344	88.6276	98.052	103.619
Bsep	97.1283	82.2139	89.5094	95.4161

Table 1: Cost of the circular arrangement returned by the four algorithms. Additionally, the worst-case cost on these graphs is $n - 1$, where n is the number of nodes in the graph. Thirty random arrangements were tried, and we report their cost in the form of average \pm standard deviation.

However, Bsep is slower than the other algorithms even with all the simplifications above. For example, Bsep took 17.3s on an unloaded Ultra60 for graph 1 while Fesh took only 1s on the same machine and with the same compiler (g++ -O).

Acknowledgments

This work has been supported in part under NSF grant ANI-0123929.

References

- [AAF98] K. C. Almeroth, M. H. Ammar, and Z. Fei. Scalable delivery of Web pages using cyclic best-effort (UDP) multicast. In *Proceedings of the Seventeenth Annual Joint Conference of the IEEE Computer and Communications Societies (INFO-COM 1998)*, 1998.
- [AFK96] Sanjeev Arora, Alan Frieze, and Haim Kaplan. A new rounding procedure for the assignment problem with applications to dense graph arrangement problems. In *37th Annual Symposium on Foundations of Computer Science (Burlington, VT, 1996)*, pages 21–30. IEEE Comput. Soc. Press, Los Alamitos, CA, 1996.
- [AJ99] Martin Arlitt and Tai Jin. Workload characterization of the 1998 World Cup web site. Technical Report HPL-1999-35R1, HP Labs, 1999.
- [AMO93] Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network flows*. Prentice Hall Inc., Englewood Cliffs, NJ, 1993. Theory, algorithms, and applications.
- [Bak74] Kenneth R. Baker. *Introduction to sequencing and scheduling*. Wiley, New York, 1974.

- [BLMR98] John W. Byers, Michael Luby, Michael Mitzenmacher, and Ashutosh Rege. A digital fountain approach to reliable distribution of bulk data. In *Proc. Sigcomm*, 1998.
- [CH91] C. K. Cheng and T. C. Hu. Ancestor tree for arbitrary multi-terminal cut functions. *Ann. Oper. Res.*, 33(1-4):199–213, 1991. Topological network design (Copenhagen, 1989).
- [CLP01] Panos K. Chrysanthis, Vincenzo Liberatore, and Kirk Pruhs. Middleware support for multicast-based data dissemination: A working reality. White paper, 2001.
- [CSS99] William W. Cohen, Robert E. Schapire, and Yoram Singer. Learning to order things. *Journal of Artificial Intelligence Research*, 10:243–270, 1999.
- [DF] <http://www.digitalfountain.com/>.
- [ENRS95] Guy Even, Joseph (Seffi) Naor, Satish Rao, and Baruch Schieber. Divide-and-conquer approximation algorithms via spreading metrics. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science*, pages 62–71, October 1995.
- [ENSS98] G. Even, J. Naor, B. Schieber, and M. Sudan. Approximating minimum feedback sets and multicuts in directed graphs. *Algorithmica*, 20(2):151–174, 1998.
- [GJ79] Michael R. Garey and David S. Johnson. *Computers and intractability*. W. H. Freeman and Co., San Francisco, Calif., 1979. A guide to the theory of NP-completeness, A Series of Books in the Mathematical Sciences.
- [Hač79] L. G. Hačijan. A polynomial algorithm in linear programming. *Dokl. Akad. Nauk SSSR*, 244(5):1093–1096, 1979.
- [Han89] Mark D. Hansen. Approximation algorithms for geometric embeddings in the plane with applications to parallel processing problems. In *Proceedings of the 30th Annual Symposium on Foundations of Computer Science*, pages 604–609, 1989.
- [HGLW87] Gary Herman, Gita Gopal, K. C. Lee, and Abel Weinrib. The datacycle architecture for very high throughput database systems. In *Proceedings of the 1987 ACM SIGMOD Conference International Conference on Management of Data*, pages 97–103, 1987.
- [HN] <http://www.hns.com/>.
- [Kan92] V. Kann. *On the Approximability of NP-complete Optimization Problems*. PhD thesis, Royal Institute of Technology, Stockholm, 1992.

- [KL00] Sanjeev Khanna and Vincenzo Liberatore. On broadcast disk paging. *SIAM Journal on Computing*, 29(5):1683–1702, 2000.
- [KSY00] Claire Kenyon, Nicolas Schabanel, and Neal Young. Polynomial-time approximation scheme for data broadcast. In *Proceedings of the Thirtisecond ACM Symposium on the Theory of Computing*, 2000.
- [Lib02] Vincenzo Liberatore. Multicast scheduling for list requests. In *21st Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2002)*, 2002. To appear.
- [LR99] Tom Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *J. ACM*, 46(6):787–832, 1999.
- [PAS] <http://www.panamsat.com/>.
- [PD00] Larry L. Peterson and Bruce S. Davie. *Computer Networks*. Morgan Kaufmann, 2000.
- [RR98] Satish Rao and Andréa W. Richa. New approximation techniques for some ordering problems. In *Proceedings of the Ninth Annual ACM-SIAM Symposium on Discrete Algorithms (San Francisco, CA, 1998)*, pages 211–218, New York, 1998. ACM.
- [Ste98] W. Richard Stevens. *Unix Network Programming*. PTR PH, 1998.

A Omitted Proofs

The proof of Proposition 2.1 is a reduction from the Optimal Linear Arrangement Problem (OLA) [GJ79]:

Instance: Graph $G = (V, E)$, positive integer K .

Question: Is there a one-to-one function $f : V \rightarrow \{1, 2, \dots, |V|\}$ such that $\sum_{\{u,v\} \in E} |f(u) - f(v)| \leq K$?

Proof of Proposition 2.1. The Undirected Circular Arrangement problem is obviously in NP. We now reduce OLA to CA. First, notice that OLA's optimal solution costs at most n^3 , where $n = |V|$. Hence, we can consider without loss of generality only values of $K \leq n^3$. Let $\gamma = n^3 + 1$. We construct an instance of CA as follows. Let $G' = (V', E')$ where $V' = V \cup \{x_1, x_2, \dots, x_\gamma\}$, $x_1, x_2, \dots, x_\gamma$ are newly introduced nodes, and $E' = E \cup \{\{x_i, x_{i+1}\} : 1 \leq i < \gamma\}$. We associate a unit weight to all edges $e \in E$ and $w(e) = \gamma$ for $e \in E' - E$. Finally, we set $K' = K + \gamma(\gamma - 1)$ and ask if there is a CA solution for G' that costs no more than K' . Observe that $K' < \gamma^2$, a fact that we will use later on. If there is a solution f to OLA that costs at most K , then there is a solution to CA that costs at most K' : simply assign the index $f(u)$ to all nodes $u \in V$ and assign $f(x_i) = (n + i) \bmod n'$ to the x_i 's, where $n' = |V'| = n + \gamma$. Conversely, suppose that there is a CA solution that costs no more than K' . First, $f(x_{i+1}) = (f(x_i) + 1) \bmod n'$, otherwise $\sum_{i=1}^{\gamma-1} w(\{x_i, x_{i+1}\})h(\{x_i, x_{i+1}\}) \geq \gamma(\gamma - 1) + \gamma = \gamma^2 > K'$. Hence, $f(x_{i+1}) = (f(x_i) + 1) \bmod n'$ and $\sum_{i=1}^{\gamma-1} w(\{x_i, x_{i+1}\})h(\{x_i, x_{i+1}\}) = \gamma(\gamma - 1)$. Moreover, vertices in V are assigned consecutive indices, and we can renumber so that those indices are in the range $\{1, 2, \dots, |N|\}$. We then have that such index assignment costs no more than K , which concludes the proof. \square

Proof of Lemma 2.2. Write $e = \{u, v\}$ and observe that $(f(v) - f(u)) \bmod n + (f(u) - f(v)) \bmod n = n$ and $h(e) = \min\{(f(v) - f(u)) \bmod n, (f(u) - f(v)) \bmod n\}$. \square

Lemma A.1. *If $e = (x, y)$ crosses u , then $(f(u) - f(x)) \bmod n < n/2$.*

Proof of Lemma A.1. Since $e = (x, y)$ crosses u , $(f(u) - f(x)) \bmod n < h(e) \leq n/2$. \square

Lemma A.2. *If e crosses u , e does not cross v , and $f(v) = (f(u) + 1) \bmod n$, then $v \in e$.*

Proof. Write $e = (x, y)$ and assume by contradiction $y \neq v$. Since e crosses u and does not cross v , $(f(v) - f(x)) \bmod n \geq (f(y) - f(x)) \bmod n > (f(u) - f(x)) \bmod n$. Since $y \neq v$, we also have $(f(v) - f(x)) \bmod n > (f(y) - f(x)) \bmod n$. Hence, $(f(v) - f(x)) \bmod n \geq (f(u) - f(x)) \bmod n + 2$, which contradicts the hypothesis $(f(v) - f(u)) \bmod n = 1$. \square

Proof of Lemma 2.3. Since u minimizes $(f(u) - f(v)) \bmod n$ among all non-crossed vertices, then e crosses $x = f^{-1}((f(u) - 1) \bmod n)$. Apply the previous lemma to x and u to obtain $u \in e$. \square

Proof of Lemma 2.4. An edge $e = \{u, v\}$ crosses $h(e)$ vertices and contributes $w(e)h(e)$ to the arrangement cost:

$$\sum_{e \in E} w(e)h(e) = \sum_{e \in E} w(e) \sum_{v \in C(e)} 1 = \sum_{v \in V} \sum_{e: v \in C(e)} w(e) = \sum_{v \in V} \rho_v .$$

Thus, the lemma follows. \square

Proof of Lemma 2.5. Let $x \in V$ and $u \in U$. Observe that if $(f(x) - f(u)) \bmod n \leq (f(v_l) - f(u)) \bmod n$ then $x \in U$. Analogously, if $(f(u) - f(x)) \bmod n < (f(v_l) - f(x)) \bmod n$, then $x \in U$. Let $e = \{u, v\}$, $u \in U$, and $v \in V - U$. If the canonical orientation of e is (u, v) , then $v \notin U$ implies $(f(v) - f(u)) \bmod n > (f(v_l) - f(u)) \bmod n$, and so e crosses v_l . If the canonical orientation is (v, u) , then $v \notin U$ implies $(f(u) - f(v)) \bmod n \geq (f(v_l) - f(v)) \bmod n$, and so e crosses $f^{-1}((f(v_l) - 1) \bmod n)$. \square

Proof of Lemma 2.8. By Lemma 2.7, we have that $C(e) \supseteq C_j(e) \cup \bigcup_{v \in C_j(e)} I(v) = \bigcup_{v \in C_j(e)} (I(v) \cup \{v\})$. By Lemma 2.6, the sets $I(v) \cup \{v\}$ are disjoint, so that $h(e) = |C(e)| \geq \sum_{v \in C_j(e)} d(v)$. We can express the cost of f as

$$\sum_{e \in E} w(e)h(e) = \sum_{j=1}^k \sum_{e \in E_j} w(e)h(e) \geq \sum_{j=1}^k \sum_{e \in E_j} \sum_{v \in C_j(e)} w(e)d(v) = \sum_{j=1}^k \sum_{v \in V_j} \rho_v^j d(v) .$$

\square

Proof of Lemma 2.11. Let $n' = \lceil 2n/5 \rceil$. Consider an optimum circular arrangement f . Let $U_i = f^{-1}(\{i, (i+1) \bmod n, \dots, (i+n'-1) \bmod n\})$ ($0 \leq i < n$). If an edge e crosses from U_i to $V - U_i$, that is $e \cap U_i, e - U_i \neq \emptyset$, then, by Lemma 2.5, e crosses either $u = f^{-1}((i-1) \bmod n)$ or $v = f^{-1}((i+n'-1) \bmod n)$. Let ξ_i be the sum of the weights of the edges that cross u or v , that is, $\xi_i = \rho_u + \rho_v$. By optimality of s^* , $\xi_i \geq s^*$. By Lemma 2.4, $a^* = \sum_{v \in V} \rho_v = \sum_{i=0}^{n-1} \xi_i/2 \geq ns^*/2$, which proves the claim. \square

Proof of Lemma 2.13. The cost due to an $e \in E - E_1$ is $w(e)h(e) \leq nw(e)/2$. Moreover, edge e belongs to an approximate $(1/3)$ -separator of G . Hence, the total weight of edges in $E - E_1$ is no more than αs_0^* , where s_0^* is the minimum cost of a $(2/5)$ -separator of G and $\alpha = O(\log n)$. Hence, $c(E - E_1) \leq n\alpha s_0^*/2 \leq \alpha a^*$ by Lemma 2.11, which concludes the proof. \square

Proof of Lemma 2.14. Let $G' = (V - H_1, E')$, where E' is the set of edges contained in E_1 that are not incident in H_1 , and let s_1^* be the minimum cost of a $(2/5)$ -separator of G' . The optimum circular arrangement of G costs at least as much as the optimum circular arrangement of G' . Hence, $s_1^* \leq (2a^*)/|V - H_1| \leq (4a^*)/n$, so that the cost for the edges in $E_1 - E_2$ is at most $n\alpha s_1^*/2 \leq 2\alpha a^*$, which proves the claim. \square

The proof of Proposition 3.1 is a reduction from the Directed Optimal Linear Arrangement Problem (DOLA) [GJ79]:

Instance: A directed graph $G = (N, A)$ and a positive integer K .

Question: Is there a topological ordering $f : N \rightarrow \{1, 2, \dots, |N|\}$ of N such that $\sum_{(u,v) \in A} (f(v) - f(u)) \leq K$?

Proof of Proposition 3.1. The Circular Arrangement problem is obviously in NP. We now reduce DOLA to CA. First, notice that DOLA's optimal solution costs at most n^3 , where $n = |N|$. Hence, we can consider without loss of generality only values of $K \leq n^3$. Let $\gamma = n^3 + 1$. We construct an instance of CA as follows. Let $G' = (N', A')$ where $N' = N \cup \{x_1, x_2, \dots, x_\gamma\}$, $x_1, x_2, \dots, x_\gamma$ are newly introduced nodes, and $A' = A \cup \{(x_i, x_{i+1}) : 1 \leq i < \gamma\}$. We associate a unit weight to all arcs $e \in A$ and $w(e) = \gamma$ for $e \in A' - A$. Finally, we set $K' = K + \gamma(\gamma - 1)$ and ask if there is a CA solution for G' that costs no more than K' . Observe that $K' < \gamma^2$, a fact that we will often use later on. If there is a solution to DOLA that costs at most K , then there is a solution to CA that costs at most K' : simply assign the index $f(u)$ to all nodes $u \in N$ and assign $f(x_i) = (n + i) \bmod n'$ to the x_i 's, where $n' = |N'| = n + \gamma$. Conversely, suppose that there is a CA solution that costs no more than K' . First, $f(x_{i+1}) = (f(x_i) + 1) \bmod n'$, otherwise $\sum_{i=1}^{\gamma-1} w((x_i, x_{i+1}))((f(x_{i+1}) - f(x_i)) \bmod n') \geq \gamma^2 > K'$. Hence, $f(x_{i+1}) = (f(x_i) + 1) \bmod n'$ and $\sum_{i=1}^{\gamma-1} w((x_i, x_{i+1}))((f(x_{i+1}) - f(x_i)) \bmod n') = \gamma(\gamma - 1)$. Moreover, nodes in N are assigned consecutive indices, and we can renumber so that those indices are in the range $\{1, 2, \dots, |N|\}$. We next claim that such index assignment is a topological ordering. Suppose by contradiction that there is an arc $e = (u, v) \in A$ with $f(v) < f(u)$. Then, $(f(v) - f(u)) \bmod n' \geq \gamma$ and so $(f(v) - f(u)) \bmod n' + \sum_{i=1}^{\gamma-1} w((x_i, x_{i+1}))((f(x_{i+1}) - f(x_i)) \bmod n') \geq \gamma^2 > K'$. We then have that the index assignment f induces a topological ordering of N and f 's cost is at most K , which concludes the proof. \square

Lemma A.3. *If e crosses u , e does not cross v , and $f(v) = (f(u) + 1) \bmod n$, then v is e 's head.*

Proof. Write $e = (x, y)$ and assume by contradiction $y \neq v$. Since e crosses u and does not cross v , $(f(v) - f(x)) \bmod n \geq (f(y) - f(x)) \bmod n > (f(u) - f(x)) \bmod n$. Since $y \neq v$, we also have $(f(v) - f(x)) \bmod n > (f(y) - f(x)) \bmod n$. Hence, $(f(v) - f(x)) \bmod n \geq (f(u) - f(x)) \bmod n + 2$, which contradicts the hypothesis $(f(v) - f(u)) \bmod n = 1$. \square

Proof of Lemma 3.2. Since u minimizes $(f(u) - f(v)) \bmod n$ among all non-crossed vertices, then e crosses $x = f^{-1}((f(u) - 1) \bmod n)$. Apply the previous lemma to x and u to obtain that u is e 's head. \square

Proof of Lemma 3.3. An arc $e = (u, v)$ crosses $h(e)$ vertices and contributes $w(e)h(e)$ to the arrangement cost:

$$\sum_{e \in A} w(e)h(e) = \sum_{e \in A} w(e) \sum_{v \in C(e)} 1 = \sum_{v \in N} \sum_{e: v \in C(e)} w(e) = \sum_{v \in N} \rho_v.$$

Thus, the lemma follows. \square

Proof of Lemma 3.4. The proof is by contradiction. Let $n_j = |N_j|$, let A_j be the set of arcs that have both endpoints in N_j , and let $G_j = (N_j, A_j)$ ($1 \leq j \leq k$). We will say that a vertex $u \in N_j$ is the *successor* of a vertex $v \in N_j$ in the arrangement f if and only if $(f(u) - f(v)) \bmod n \leq (f(x) - f(v)) \bmod n$ for all $x \in N_j$. The successor of $v \in N_j$ in f is unique, belongs to N_j and will be denoted as $s_f(v)$. We define $d(v) = (f(s_f(v)) - f(v)) \bmod n$ and observe that $\sum_{v \in N_j} d(v) = n$ for all j 's. Define $I(v) = \{u : 0 < (f(u) - f(v)) \bmod n < d(v)\}$. We have that $|I(v)| = d(v) - 1$. Moreover, by definition of $s_f(v)$, we have $I(v) \cap N_j = \emptyset$ when $v \in V_j$. It is easy to see that $I(v) \cap I(u) = \emptyset$ for all $u, v \in N_j$. We now claim that if an arc e crosses v and $I(v) \neq \emptyset$, edge e also crosses all nodes in $I(v)$. Suppose that this is not the case and let $u \in I(v)$ be the node that e does not cross and that minimizes $(f(u) - f(v)) \bmod n$. Hence, arc e crosses the node $x = f^{-1}((f(u) - 1) \bmod n) \in I(v) \cup \{v\}$. Lemma 3.2 now gives that u is e 's tail, which contradicts $I(v) \cap V_j = \emptyset$. We conclude that e crosses all nodes in $I(v)$. Let $C_j(e) = C(e) \cap E_j$. We have that $C(e) \supseteq C_j(e) \cup \bigcup_{v \in C_j(e)} I(v) = \bigcup_{v \in C_j(e)} (I(v) \cup \{v\})$. Since the sets $I(v) \cup \{v\}$ are disjoint for $v \in N_j$, we have that $h(e) = |C(e)| \geq \sum_{v \in C(e)} d(v)$. Let ρ_v^j be the total weight of arcs in E_j that cross v , that is, $\rho_v^j = \sum_{e \in V_j: v \in C(e)} w(e)$. We can express the cost of f as

$$\sum_{e \in A} w(e)h(e) = \sum_{j=1}^k \sum_{e \in A_j} w(e)h(e) \geq \sum_{j=1}^k \sum_{e \in A_j} \sum_{v \in C(e)} w(e)d(v) = \sum_{j=1}^k \sum_{v \in N_j} \rho_v^j d(v).$$

Fix nodes v_j with $\rho_{v_j}^j = \min_{u \in N_j} \rho_u^j$ ($1 \leq j \leq k$). Consider the following circular arrangement:

$$g(v) = \begin{cases} n_j - 1 & \text{if } v = v_j \\ g(s_f(v)) - 1 & \text{otherwise} \end{cases}.$$

It is immediate to see that g establishes a one-to-one mapping $g : N \rightarrow \{0, 1, \dots, n-1\}$ and that g places the vertices in the N_j 's next to each other. It remains to show that g costs less than f does. Since g preserves the successor $s_f(v)$ within nodes in N_j , we have that arcs cross exactly the same nodes of N_j in g as in f . The cost change due to G_j is no more than

$$\begin{aligned} \sum_{v \in N_j - \{v_j\}} \rho_v^j + (n - n_j + 1)\rho_{v_j}^j - \sum_{v \in N_j} \rho_v^j d(v) &= \rho_{v_j}^j (n - n_j) - \sum_{v \in N_j - \{v_j\}} \rho_v^j (d(v) - 1) \\ &\leq \rho_{v_j}^j (n - n_j) - \rho_{v_j}^j \sum_{v \in N_j} (d(v) - 1) = 0 \end{aligned}$$

because $d(v) \geq 1$, $\rho_{v_j}^j \leq \rho_v^j$, $\sum_{v \in V_j} d(v) = n$. Hence, g costs less than f , a contradiction is reached, and the lemma is proven. \square

Proof of Lemma 3.9. The proof is similar to that of Lemma 2.11. Let $n' = \lceil 2n/5 \rceil$. Consider an optimum circular arrangement f and let $U_i = f^{-1}(\{i, (i+1) \bmod n, \dots, (i+n'-1) \bmod n\})$ ($0 \leq i < n$). If an arc e has one endpoint in U_i , then, by Lemma 3.2, e crosses either $f^{-1}((i-1) \bmod n)$ or $f^{-1}((i+n'-1) \bmod n)$. By optimality of s^* , $\rho_i \geq s^*$. By Lemma 3.3, $a^* \geq \sum_{i=0}^{n-1} \rho_i/2 \geq ns^*/2$, which proves the claim. \square

Proof of Lemma 3.7. Let $G_1 = (N, A_1)$ be the subgraph obtained from G by removing arcs that are incident on both H_1 and $N - H_1$. Then, $w(A - A_1) = \alpha s_0^*$, where s_0^* is the minimum cost of a $(2/5)$ -separator on the undirected graph induced by G and $\alpha = O(\log n)$. The total contribution of $A - A_1$ to the arrangement cost is at most $n\alpha s_0^* \leq 2\alpha a^*$ by Lemma 3.9. Consider next $G' = (N - H_1, A')$ obtained by removing from G_1 the node set H_1 and all incident arcs. Let s_1^* be the minimum cost of a $(2/5)$ -separator of the undirected version of G' . The optimum circular arrangement of G costs at least as much as the optimum circular arrangement of G' . Hence, $s_1^* \leq (2a^*)/|N - H_1| \leq (4a^*)/n$, so that the cost for the arcs between H_2 and H_3 is at most $n\alpha s_1^* \leq 8\alpha a^*$. Since $\alpha = O(\log n)$, the proof is complete. \square