

Constructing Incremental Sequences in Graphs

Ralf Klasing^a Christian Laforest^b Joseph Peters^c Nicolas Thibault^b

^aLaBRI - Université Bordeaux 1 - CNRS, 351 cours de la Libération, 33405 Talence cedex (France)

^bIBISC, Tour Evry 2, Université d'Evry, 523 place des terrasses, 91000 EVRY (France)

^cSchool of Computing Science, Simon Fraser University, Burnaby, British Columbia, V5A 1S6 (Canada)

Abstract

Given a weighted graph $G = (V, E, w)$, we investigate the problem of constructing a sequence of $n = |V|$ subsets of vertices M_1, \dots, M_n (called groups) with small diameters, where the diameter of a group is calculated using distances in G . The constraint on these n groups is that they must be incremental: $M_1 \subset M_2 \subset \dots \subset M_n = V$. The cost of a sequence is the maximum ratio between the diameter of each group M_i and the diameter of a group N_i^* with i vertices and minimum diameter: $\max_{2 \leq i \leq n} \left\{ \frac{D(M_i)}{D(N_i^*)} \right\}$. This quantity captures the impact of the incremental constraint on the diameters of the groups in a sequence. We give general bounds on the value of this ratio and we prove that the problem of constructing an optimal incremental sequence cannot be solved approximately in polynomial time with an approximation ratio less than 2 unless $P = NP$. Finally, we give a 4-approximation algorithm and we show that the analysis of our algorithm is tight.

Key words: incremental sequence, graph, approximation algorithms

1. Introduction

We are given a weighted undirected graph $G = (V, E, w)$ where w is a function that assigns positive weights to the edges. We use $d_G(u, v)$ to denote the distance between u and v in G , that is, the weight of a minimum weight path between u and v in G . The *diameter of a group* $M \subseteq V$ is $D(M) = \max\{d_G(u, v) : u, v \in M\}$. Let $n = |V|$. A *group of size* i , $1 \leq i \leq n$, of *minimum diameter* is a group $N_i^* \subseteq V$ with $|N_i^*| = i$, and $D(N_i^*) = \min\{D(M) : M \subseteq V, |M| = i\}$. Our goal in this paper is to construct a sequence of groups M_1, M_2, \dots, M_n such that $M_1 \subset M_2 \subset \dots \subset M_n$ and each M_i has a diameter that is close to the optimal diameter (the diameter of N_i^*). We measure the quality of an incremental sequence M_1, M_2, \dots, M_n by the maximum ratio between the diameter of each M_i and the diameter of the corresponding N_i^* .

Definition 1. An *incremental sequence of groups* is a sequence M_1, M_2, \dots, M_n such that $M_1 \subset M_2 \subset \dots \subset M_n = V$ and $|M_i| = i$ for all i , $1 \leq i \leq n$. The *cost of*

an *incremental sequence of groups* M_1, M_2, \dots, M_n is

$$\text{cost}(M_1, \dots, M_n) = \max_{2 \leq i \leq n} \left\{ \frac{D(M_i)}{D(N_i^*)} \right\}.$$

Since we compare the diameter of successive incremental groups to groups of minimum diameter that are not constrained to be incremental, this cost measures the impact on the diameter of the constraint that the sequence of groups must be incremental.

Definition 2. An *optimal incremental sequence* is an incremental sequence $N_1^{\text{opt}}, N_2^{\text{opt}}, \dots, N_n^{\text{opt}}$ of minimum cost:

$$\text{cost}(N_1^{\text{opt}}, \dots, N_n^{\text{opt}}) = \min\{\text{cost}(M_1, \dots, M_n) : M_1 \subset \dots \subset M_n = V, |M_i| = i, 1 \leq i \leq n\}.$$

Our main contribution in this paper is a new cost measure, the *cost of an incremental sequence*, which allows the study of the impact of an incremental constraint on the quality of approximate solutions to NP -hard optimization problems. In this paper, we use the measure to study a diameter problem and a related eccentricity problem, but the approach is general and can be used to study other problems.

Our cost measure and our approach differ in several important ways from common previous approaches to studying approximation algorithms. Perhaps the most

Email: Ralf Klasing [ralf.klasing@labri.fr], Christian Laforest [laforest@ibisc.fr], Joseph Peters [peters@cs.sfu.ca], Nicolas Thibault [nthibaul@ibisc.fr].

common approach assumes that an entire problem instance is given in advance and the performance of the algorithm is measured in terms of the quality of the final solution. There are no constraints on the intermediate solutions produced by the algorithm and the cost measure does not take them into account. The *approximation ratio* is used to measure the intrinsic difficulty of constructing a solution in polynomial time compared to the best (non-polynomial time) solution. See [1,5,6,8] for comprehensive treatments of approximation algorithms. The major differences of our approach are that it requires that the final solution be built incrementally and the quality of the intermediate solutions is taken into account by our cost measure.

Another popular approach is to assume that a problem instance is revealed one element at a time. The quality of algorithms for these *on-line* versions of problems is measured using the *competitive ratio* which compares the final solution to the best that can be achieved by an *off-line* algorithm that knows an entire instance in advance. In one variant, changes to the existing partial solution are allowed when a new element is revealed; in another variant, changes are disallowed. See [2,3] for further references on on-line problems. The most important differences of our approach are that an entire instance is known in advance, and the order that elements are added is chosen by the algorithm.

The approach taken in [7] is to construct a sequence of incremental *trees* to cover successive groups. The main difference from our approach is that the successive groups are not chosen by the algorithm in [7]; they are given in advance.

Beyond theoretical interest in the incremental cost measure, a sequence of incremental groups could be used in applied situations such as the following. Suppose that the graph models a point to point network interconnecting a cluster of computers that is shared among several applications. Each application is allocated a subset of the computers that are available when it starts. An application starts with one active computer. As the need for computational power increases, computers are added, one by one, giving an incremental sequence of groups of computers. The computers need to communicate to exchange data and partial results, so the performance also depends on the communication latencies among the computers in the current group. The maximum latency in a group is the diameter of the group, so an optimal incremental sequence will give the

best performance.

In the next section, we derive matching upper and lower bounds on the cost of an optimal incremental sequence. In Section 3, we prove that the problem of constructing an optimal incremental sequence cannot be solved approximately with an approximation ratio less than 2 unless $P = NP$. In Section 4, we develop an optimal polynomial-time algorithm for the related problem of finding an incremental sequence of groups with small *eccentricities*. We then use this algorithm to develop a polynomial-time 4-approximation algorithm for the problem of constructing an optimal incremental sequence for a graph, and we show that our analysis of the algorithm is tight.

2. General bounds on the cost of an optimal incremental sequence

In this section, we derive matching upper and lower bounds on the cost of an optimal incremental sequence.

Theorem 1 $cost(N_1^{opt}, \dots, N_n^{opt}) \leq \sqrt{D(V)}$ for every weighted graph $G = (V, E, w)$ with $w(e) \geq 1$ for all $e \in E$.

PROOF. Let $G = (V, E, w)$ be a weighted graph with $w(e) \geq 1$ for all $e \in E$. For every i , $1 \leq i \leq n$, let N_i^* be a group of size i of minimum diameter. Let i_0 be the largest integer such that $D(N_{i_0}^*) \leq \sqrt{D(V)}$. Since $G = (V, E, w)$ is a weighted graph with $w(e) \geq 1$ for all $e \in E$, we have

$$1 \leq D(N_2^*) \leq \dots \leq D(N_{i_0}^*) \leq \sqrt{D(V)} < D(N_{i_0+1}^*) \leq \dots \leq D(N_n^*). \quad (1)$$

Let M_1, M_2, \dots, M_n be any incremental sequence such that $M_{i_0} = N_{i_0}^*$. Thus,

$$1 \leq D(M_2) \leq \dots \leq D(M_{i_0}) = D(N_{i_0}^*) \leq \sqrt{D(V)}. \quad (2)$$

As the diameter of $G = (V, E, w)$ is $D(V)$, we have

$$D(M_{i_0+1}) \leq \dots \leq D(M_n) \leq D(V). \quad (3)$$

By (1) and (2) we obtain

$$\max_{2 \leq i \leq i_0} \left\{ \frac{D(M_i)}{D(N_i^*)} \right\} \leq \sqrt{D(V)},$$




Fig. 1. The graph G_0

and by (1) and (3) we obtain $\max_{i_0+1 \leq i \leq n} \left\{ \frac{D(M_i)}{D(N_i^*)} \right\} \leq \frac{D(V)}{\sqrt{D(V)}} = \sqrt{D(V)}$. It follows that

$$\begin{aligned} \text{cost}(N_1^{opt}, \dots, N_n^{opt}) &\leq \text{cost}(M_1, \dots, M_n) \\ &\leq \sqrt{D(V)}. \quad \square \end{aligned}$$

The next theorem provides a lower bound that matches the upper bound of Theorem 2. Together, Theorems 1 and 2 give a tight bound on the worst case cost of an optimal incremental sequence for the class of graphs with all edge weights at least 1.

Theorem 2 $\text{cost}(N_1^{opt}, \dots, N_n^{opt}) \geq \sqrt{D(V_0)}$ for infinitely many weighted graphs with all edge weights at least 1.

PROOF. Let $G_0 = (V_0, E_0, w_0)$ be the weighted graph in Figure 1, where $K > 1$ is an arbitrary constant. The diameter of G_0 is $D(V_0) = K^2$. For every $i, 1 \leq i \leq 5$, let N_i^* be a group of size i of minimum diameter. Let M_1, M_2, \dots, M_5 be any incremental sequence for G_0 . If $M_2 \neq \{a, b\}$, then $\frac{D(M_2)}{D(N_2^*)} \geq \frac{K}{1} = K$. Otherwise, $M_2 = \{a, b\}$, and for all M_3 such that $M_2 \subset M_3$, $\frac{D(M_3)}{D(N_3^*)} = \frac{K^2}{K} = K$. Thus, $\text{cost}(M_1, \dots, M_5) \geq K = \sqrt{D(V_0)}$. The proof is easily generalized to any complete graph with all edge weights K^2 except a triangle with edge weights K and a pair of vertices that is disjoint from the triangle and connected by an edge with weight 1. \square

3. Non-approximability of constructing optimal incremental sequences

In this section, we investigate the complexity of constructing optimal incremental sequences. We first state the problems more formally.

UNWEIGHTED INCREMENTAL SEQUENCE

INSTANCE: A graph $G = (V, E)$.

SOLUTION: An incremental sequence of groups M_1, M_2, \dots, M_n in G , i.e., $M_1 \subset \dots \subset M_n = V$ with $|M_i| = i$ for all $i, 1 \leq i \leq n$.

MEASURE: $\text{cost}(M_1, \dots, M_n)$.

WEIGHTED INCREMENTAL SEQUENCE

INSTANCE: A weighted graph $G = (V, E, w)$ with $w(e) > 0$ for all $e \in E$.

SOLUTION: An incremental sequence of groups M_1, M_2, \dots, M_n in G , i.e., $M_1 \subset \dots \subset M_n = V$ with $|M_i| = i$ for all $i, 1 \leq i \leq n$.

MEASURE: $\text{cost}(M_1, \dots, M_n)$.

We now show that there is no polynomial-time approximation algorithm with an approximation ratio less than 2 for the problem of finding $N_1^{opt}, \dots, N_n^{opt}$ unless $P = NP$.

Theorem 3 *There is no polynomial time approximation algorithm with an approximation ratio less than 2 for UNWEIGHTED INCREMENTAL SEQUENCE unless $P = NP$.*

PROOF. Let $G = (V, E)$. Let $r' \notin V$, and let $G' = (V', E')$ be the graph such that $V' = V \cup \{r'\}$ and $E' = E \cup \{(u, r') : u \in V\}$. Constructing G' from G can be done in polynomial time. For all $S \subseteq V'$, let $D'(S)$ denote the diameter of S in G' . Let $M_1, \dots, M_{n'}$ be an incremental sequence for G' . For all $i, 1 \leq i \leq n'$, let N_i^* be a group of size i of minimum diameter in G' .

Suppose, by contradiction, that there is a polynomial time approximation algorithm for UNWEIGHTED INCREMENTAL SEQUENCE that guarantees an approximation ratio strictly less than 2. Let $V' = \{v'_1, v'_2, \dots, v'_{n'}\}$ such that $\{v'_1, \dots, v'_{i_0}\}$ is a maximum clique in G' . Consider the incremental sequence $N_1^{opt}, \dots, N_{n'}^{opt}$ in G' obtained by adding the vertices of G' in the order $v'_1, v'_2, \dots, v'_{n'}$. As G' is an unweighted graph of diameter at most 2, any subset $S \subseteq V'$ with $|S| \geq 2$ is such that $D'(S) = 1$ or $D'(S) = 2$. It follows that

- (1) $D'(N_i^{opt}) = D'(N_i^*) = 1$ for all $2 \leq i \leq i_0$,
- (2) $D'(N_i^{opt}) = D'(N_i^*) = 2$ for all $i_0+1 \leq i \leq n'$.

In particular, the incremental sequence $N_1^{opt}, \dots, N_{n'}^{opt}$ satisfies $D'(N_i^{opt}) = D'(N_i^*)$ for all $2 \leq i \leq n'$. Hence,

$$\frac{\max_{2 \leq i \leq n'} \left\{ \frac{D'(M_i)}{D'(N_i^*)} \right\}}{\max_{2 \leq i \leq n'} \left\{ \frac{D'(N_i^{opt})}{D'(N_i^*)} \right\}} = \max_{2 \leq i \leq n'} \left\{ \frac{D'(M_i)}{D'(N_i^*)} \right\}.$$

In G' , any subset $S \subseteq V'$ with $|S| \geq 2$ is such that $D'(S) = 1$ or $D'(S) = 2$, so the only two possible values of $\max_{2 \leq i \leq n'} \left\{ \frac{D'(M_i)}{D'(N_i^*)} \right\}$ are 1 and 2. This means that if $M_1, \dots, M_{n'}$ is an incremental sequence constructed in polynomial time by an algorithm with approximation ratio strictly less than 2, then $\max_{2 \leq i \leq n'} \left\{ \frac{D'(M_i)}{D'(N_i^*)} \right\} = 1$, and $D'(M_i) = D'(N_i^*)$ for all i , $2 \leq i \leq n'$. Thus, by choosing the largest integer i_0 such that $D'(M_{i_0}) = D'(N_{i_0}^*) = 1$, one can construct in polynomial time a maximum clique in G' (namely M_{i_0}), and therefore a maximum clique in G (namely $M_{i_0} \setminus \{r'\}$). This contradicts the fact that finding a clique of maximum size in G is NP -hard (see [4]). \square

Corollary 4 *There is no polynomial time approximation algorithm with an approximation ratio less than 2 for WEIGHTED INCREMENTAL SEQUENCE unless $P = NP$.*

4. A 4-approximation algorithm for constructing an optimal incremental sequence

In this section, we develop an optimal polynomial-time algorithm to find an incremental sequence of groups with small *eccentricities*. We then prove that our algorithm is a 4-approximation algorithm for the problem of finding an optimal incremental sequence for the diameter.

Definition 3. The *eccentricity* of a group $M \subseteq V$ with root $r \in M$ is $E(M, r) = \max\{d_G(u, r) : u \in M\}$. A group $M_i^* \subseteq V$ with $|M_i^*| = i$, $1 \leq i \leq n$, is a *group of size i of minimum eccentricity* if there exists a vertex $r_i^* \in M_i^*$ (called its *associated root*) such that $E(M_i^*, r_i^*) = \min\{E(M, r) : M \subseteq V, |M| = i, r \in M\}$. An *optimal incremental sequence for the eccentricity* is an incremental sequence of groups $M_1^{opt} =$

$\{r^{opt}\}, M_2^{opt}, \dots, M_n^{opt} = V$ with $|M_i^{opt}| = i$ for all i , $1 \leq i \leq n$, such that

$$\max_{2 \leq i \leq n} \left\{ \frac{E(M_i^{opt}, r^{opt})}{E(M_i^*, r_i^*)} \right\} = \min \left\{ \max_{2 \leq i \leq n} \left\{ \frac{E(M'_i, r')}{E(M_i^*, r_i^*)} \right\} : M'_1 \subset \dots \subset M'_n = V, \right. \\ \left. |M'_i| = i, M'_1 = \{r'\} \right\}.$$

Definition 4. Let $r \in V$ and let S be the sequence containing the values $\{d_G(r, u) : u \in V\}$ sorted in increasing order (note that $|S| \leq n = |V|$). Consider the partition $F_1(r), \dots, F_n(r)$ of V such that $F_j(r) = \{u : d_G(r, u) \text{ is the } j^{\text{th}} \text{ value in } S\}$, $1 \leq j \leq n$. A group $M \subseteq V$ is a *breadth-first subset from root $r \in M$* if it satisfies:

- If $|M| = 1$, then $M = \{r\}$.
- If $|M| \geq 2$, then there exists a $k \geq 2$ such that
 - $\forall j, 1 \leq j \leq k-1, F_j(r) \cap M = F_j(r)$,
 - $F_k(r) \cap M \neq \emptyset$,
 - $\forall l > k, F_l(r) \cap M = \emptyset$.

The following algorithm BE_i (for Best Eccentricity) finds a group of size i of minimum eccentricity for any i , $1 \leq i \leq n$.

Algorithm 1 (BE_i)

- (1) For each $r \in V$, construct a *breadth-first subset* $M_i(r) \subseteq V$ from root r with $|M_i(r)| = i$.
- (2) Choose r_i and its associated group $M_i(r_i)$ such that $E(M_i(r_i), r_i) = \min\{E(M_i(r), r) : r \in V\}$.

Note that for all $r \in V$, the partition $F_1(r), \dots, F_n(r)$ and the associated group $M_i(r)$ can be constructed in polynomial time using Dijkstra's algorithm. Thus, $M_i(r_i)$ can be constructed in polynomial time.

The following lemma shows that Algorithm BE_i constructs a group of size i of minimum eccentricity. The idea of the proof is to show that for a given root $r \in V$, the group of size i of minimum eccentricity associated with r is a *breadth-first subset from root r* . As algorithm BE_i checks each root $r \in V$, it necessarily finds the right subset.

Lemma 5 *Algorithm BE_i constructs a group of size i of minimum eccentricity for any i , $1 \leq i \leq n$.*

PROOF. Let $1 \leq i \leq n$. For all $r \in V$, let $M'_i(r) \subseteq V$ be any group of size i with $r \in M'_i(r)$ and let $M''_i(r) \subseteq V$ be a *breadth-first subset from root r* of size i . Thus, for any $r \in V$, we have

$$E(M_i''(r), r) = \max \{d_G(u, r) : u \in M_i''(r)\} \leq \max \{d_G(v, r) : v \in M_i'(r)\} = E(M_i'(r), r). \text{ Hence,}$$

$$\min \{E(M_i''(r), r) : r \in V\} \leq \min \{E(M_i'(r), r) : r \in V\}. \quad (4)$$

Let $M_i(r_i)$ be a group of size i constructed by Algorithm BE $_i$, $1 \leq i \leq n$, and let M_i^* be a group of size i with minimum eccentricity and associated root $r_i^* \in M_i^*$. By the definition of Algorithm BE $_i$, we have $E(M_i(r_i), r_i) = \min \{E(M_i''(r), r) : r \in V\}$ and by the definition of M_i^* , we have $E(M_i^*, r_i^*) = \min \{E(M_i'(r), r) : r \in V\}$. Thus, by (4), $E(M_i(r_i), r_i) \leq E(M_i^*, r_i^*)$. As M_i^* is a group of size i with the smallest eccentricity, $E(M_i^*, r_i^*) = E(M_i(r_i), r_i)$. \square

The next algorithm IBE (for Incremental Best Eccentricity) constructs an optimal incremental sequence of groups for the eccentricity.

Algorithm 2 (IBE)

- (1) For each $r \in V$:
 - Start with $M_1(r) = \{r\}$.
 - For each i , $1 \leq i \leq n$:
 - (a) Construct a breadth-first subset $M_i(r)$ from root r with $|M_i(r)| = i$.
 - (b) Compute the ratio $\frac{E(M_i(r), r)}{E(M_i^*, r_i^*)}$.
- (2) Choose $r_0 \in V$ and its associated sequence $M_1(r_0), \dots, M_n(r_0)$ such that

$$\max_{2 \leq i \leq n} \left\{ \frac{E(M_i(r_0), r_0)}{E(M_i^*, r_i^*)} \right\} = \min \left\{ \max_{2 \leq i \leq n} \left\{ \frac{E(M_i(r), r)}{E(M_i^*, r_i^*)} \right\} : r \in V \right\}.$$

Note that for all $r \in V$, the associated sequence $M_1(r), \dots, M_n(r)$ can be constructed in polynomial time using Dijkstra’s algorithm and that for all $r \in V$, and all i , $2 \leq i \leq n$, the ratio $\frac{E(M_i(r), r)}{E(M_i^*, r_i^*)}$ can be computed in polynomial time by using Algorithm BE $_i$ to compute $E(M_i^*, r_i^*)$. Thus, $M_1(r_0), \dots, M_n(r_0)$ can be constructed in polynomial time.

Lemma 6 Algorithm IBE finds an optimal incremental sequence for the eccentricity.

PROOF. Let $M_1(r_0) = \{r_0\}, M_2(r_0), \dots, M_n(r_0)$ be the incremental sequence constructed by IBE,

let $M_1^{opt} = \{r^{opt}\}, M_2^{opt}, \dots, M_n^{opt}$ be an optimal incremental sequence for the eccentricity, and let M_i^* be a group of size i of minimum eccentricity and associated root $r_i^* \in M_i^*$, $1 \leq i \leq n$. Algorithm IBE constructs an incremental sequence starting with each possible root, including the sequence $M_1(r^{opt}), \dots, M_n(r^{opt})$ starting with $M_1(r^{opt}) = \{r^{opt}\}$. Moreover, by the definition of Algorithm IBE, the groups $M_1(r^{opt}), \dots, M_n(r^{opt})$ are breadth-first subsets from root r^{opt} . Thus, we have $E(M_i(r^{opt}), r^{opt}) \leq E(M_i^{opt}, r^{opt})$, $1 \leq i \leq n$, and we obtain

$$\max_{2 \leq i \leq n} \left\{ \frac{E(M_i(r^{opt}), r^{opt})}{E(M_i^*, r_i^*)} \right\} \leq \max_{2 \leq i \leq n} \left\{ \frac{E(M_i^{opt}, r^{opt})}{E(M_i^*, r_i^*)} \right\}.$$

By the definition of Algorithm IBE (see the second part of the algorithm), and the fact that $M_1^{opt}, \dots, M_n^{opt}$ is an optimal incremental sequence for the eccentricity, we obtain

$$\max_{2 \leq i \leq n} \left\{ \frac{E(M_i(r_0), r_0)}{E(M_i^*, r_i^*)} \right\} = \max_{2 \leq i \leq n} \left\{ \frac{E(M_i^{opt}, r^{opt})}{E(M_i^*, r_i^*)} \right\}.$$

\square

We show that Algorithm IBE is a 4-approximation algorithm for the problem of finding an optimal incremental sequence for the diameter.

Theorem 7 Let M_1, \dots, M_n be the incremental sequence constructed by Algorithm IBE and let $N_1^{opt}, \dots, N_n^{opt}$ be an optimal incremental sequence. Then

$$\frac{\text{cost}(M_1, \dots, M_n)}{\text{cost}(N_1^{opt}, \dots, N_n^{opt})} \leq 4.$$

PROOF. For every $1 \leq i \leq n$, let N_i^* be a group of size i of minimum diameter, and let M_i^* be a group of size i of minimum eccentricity and associated root $r_i^* \in M_i^*$. Let $M_1^{opt}, M_2^{opt}, \dots, M_n^{opt}$ be an optimal incremental sequence for the eccentricity.

$$\max_{2 \leq i \leq n} \left\{ \frac{D(M_i)}{D(N_i^*)} \right\} \leq 2 \max_{2 \leq i \leq n} \left\{ \frac{E(M_i, r)}{D(N_i^*)} \right\}$$

(with $M_1 = \{r\}$ and because $D(M_i) \leq 2E(M_i, r)$)

$$\begin{aligned}
&\leq 2 \max_{2 \leq i \leq n} \left\{ \frac{E(M_i, r)}{E(M_i^*, r_i^*)} \right\} \\
&\quad \text{(by Definition 3, } E(M_i^*, r_i^*) \leq \\
&\quad E(N_i^*, c_i^*) \leq D(N_i^*), \text{ with } c_i^* \in N_i^*) \\
&= 2 \max_{2 \leq i \leq n} \left\{ \frac{E(M_i^{opt}, r^{opt})}{E(M_i^*, r_i^*)} \right\} \\
&\quad \text{(by Lemma 6, with } \\
&\quad M_1^{opt} = \{r^{opt}\}) \\
&\leq 2 \max_{2 \leq i \leq n} \left\{ \frac{E(N_i^{opt}, c^{opt})}{E(M_i^*, r_i^*)} \right\} \\
&\quad \text{(because } M_1^{opt}, \dots, M_n^{opt} \text{ is an} \\
&\quad \text{optimal incremental sequence} \\
&\quad \text{for the eccentricity, with} \\
&\quad N_1^{opt} = \{c^{opt}\}) \\
&\leq 2 \max_{2 \leq i \leq n} \left\{ \frac{D(N_i^{opt})}{E(M_i^*, r_i^*)} \right\} \\
&\quad \text{(since } c^{opt} \in N_i^{opt}, \text{ we have} \\
&\quad E(N_i^{opt}, c^{opt}) \leq D(N_i^{opt})) \\
&\leq 4 \max_{2 \leq i \leq n} \left\{ \frac{D(N_i^{opt})}{D(N_i^*)} \right\} \\
&\quad \text{(because } D(N_i^*) \leq D(M_i^*) \\
&\quad \leq 2E(M_i^*, r_i^*)) \quad \square
\end{aligned}$$

Note that we cannot obtain an approximation ratio less than 2 for this problem by Theorem 3. The next theorem shows that the approximation ratio of 4 for Algorithm IBE cannot be improved.

Theorem 8 *For every $0 < \epsilon < 1$, there exists a weighted graph such that the incremental sequence M_1, \dots, M_n constructed by Algorithm IBE gives*

$$\frac{\text{cost}(M_1, \dots, M_n)}{\text{cost}(N_1^{opt}, \dots, N_n^{opt})} = \frac{4}{1 + \epsilon}.$$

PROOF. Let $G_0(\epsilon)$ be the weighted graph in Figure 2.

For each i , $1 \leq i \leq 7$, let M_i^* be a group of size i of minimum eccentricity, let $r_i^* \in M_i^*$ be its associated root, and let N_i^* be a group of size i of minimum diameter. Given $G_0(\epsilon)$, Algorithm IBE will construct an incremental sequence starting with each of the vertices and will then choose the best incremental sequence for the eccentricity among these. The incremental sequence returned by Algorithm IBE is the sequence $M_1(a), \dots, M_7(a)$ obtained by adding the vertices of $G_0(\epsilon)$ in the order a, b, c, d, e, f, g . Indeed,




Fig. 2. The graph $G_0(\epsilon)$

this sequence leads to $\max_{2 \leq i \leq 7} \left\{ \frac{E(M_i(a), a)}{E(M_i^*, r_i^*)} \right\} = \frac{2}{1 + \epsilon}$, which is the minimum possible value for any incremental sequence in $G_0(\epsilon)$. The optimal incremental sequence for the diameter, $N_1^{opt}, \dots, N_7^{opt}$, is obtained by adding vertices in the order e, f, g, a, b, c, d . Using the sequence $M_1(a), \dots, M_7(a)$ for the diameter problem leads to $\max_{2 \leq i \leq 7} \left\{ \frac{D(M_i(a))}{D(N_i^*)} \right\} = \frac{4}{1 + \epsilon}$, whereas $\max_{2 \leq i \leq 7} \left\{ \frac{D(N_i^{opt})}{D(N_i^*)} \right\} = 1$. Thus, we have

$$\max_{2 \leq i \leq 7} \left\{ \frac{D(M_i(a))}{D(N_i^*)} \right\} / \max_{2 \leq i \leq 7} \left\{ \frac{D(N_i^{opt})}{D(N_i^*)} \right\} = \frac{4}{1 + \epsilon}. \square$$

5. Conclusions

In this paper, we have introduced a new measure to capture the impact of the incremental constraint on the quality of the solutions. We have used the approach to study a diameter problem, but the approach is general and can be used to study other optimization problems.

Our main complexity result is that the problem of constructing an optimal incremental sequence cannot be solved approximately in polynomial time with an approximation ratio less than 2 unless $P = NP$. In the process of developing a 4-approximation algorithm for this problem, we proved the somewhat surprising result that the related eccentricity problem can be solved optimally in polynomial time. The analysis of our 4-approximation algorithm is tight, so reducing the gap between the upper bound of 4 and the lower bound of 2 will require either a new algorithm or a stronger lower bound.

Acknowledgements: Research supported in part by the Equipe Associée RESEAUXCOM of INRIA, and

the European project IST FET CRESCCO (contract no. IST-2001-33135). Part of this work was done while J. Peters was visiting the MASCOTTE project at INRIA Sophia Antipolis, and while J. Peters and R. Klasing were visiting the IBISC (formerly LaMI) at the Université d'Evry.

References

- [1] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, and M. Protasi, *Complexity and approximation*, Springer, 1999.
- [2] A. Borodin and R. El-Yaniv, *Online computation and competitive analysis*, Cambridge University Press, 1998.
- [3] A. Fiat and G. Woeginger, *Online algorithms: The state of the art*, LNCS, no. 1442, Springer, 1998.
- [4] M. Garey and D. Johnson, *Computers and intractability*, Freeman and Company, 1979.
- [5] D. S. Hochbaum, *Approximation algorithms for NP-hard problems*, PWS Publishing Company, 1997.
- [6] J. Hromkovic, *Algorithmics for Hard Problems—Introduction to Combinatorial Optimization, Randomization, Approximation, and Heuristics*, Texts in Theoretical Computer Science, An EATCS Series, 2nd edition, Corr. 2nd printing, Springer, 2004.
- [7] N. Thibault and C. Laforest, *An optimal online strategy to increment connection trees*, in AWIN 2004, IEEE Workshop on Adaptive Wireless Networks, (in conjunction with GLOBECOM 2004), Dallas, USA, Dec 2004.
- [8] V. Vazirani, *Approximation algorithms*, 1st edition, Corr. 2nd printing, Springer, 2002.

Received 17 January 2006; accepted 28 March 2006