Online Maintenance of k-Medians and k-Covers on a Line*

Rudolf Fleischer, Mordecai J. Golin, and Zhang Yan

Dept. of Computer Science, Hong Kong University of Science and Technology,
Clear Water Bay Road, Kowloon, Hong Kong.
`rudolf,golin,cszy@cs.ust.hk`

Abstract. The standard dynamic programming solution to finding *k-medians on a line* with n nodes requires $O(kn^2)$ time. Dynamic programming speed-up techniques, e.g., use of the quadrangle inequality or properties of totally monotone matrices, can reduce this to $O(kn)$ time but these techniques are inherently static. The major result of this paper is to show that we can *maintain the dynamic programming speedup* in an online setting where points are added from left to right on a line. Computing the new k-medians after adding a new point takes only $O(k)$ amortized time and $O(k \log n)$ worst case time (simultaneously). Using similar techniques, we can also solve the *online k-coverage with uniform coverage on a line* problem with the same time bounds.

1 Introduction

In the *k-median* problem we are given a graph $G = (V, E)$ with nonnegative edge costs. We want to choose k nodes (the *medians*) from V so as to minimize the sum of the distances between each node and its closest median. As motivation, the nodes can be thought of as *customers*, the medians as *service centers*, and the distance between a customer and a service center as the cost of servicing the customer from that center. In this view, the k-median problem is about choosing a set of k service centers that minimizes the total cost of servicing all customers.

The k-median problem is often extended so that each customer (node) has a weight, corresponding to the *amount* of service requested. The distance between a customer and its closest service center (median) then becomes the cost of providing *one unit* of service, i.e., the cost of servicing a customer will then be the weight of the customer node times its distance from the closest service center. Another extension of the problem is to assign a *start-up cost* to each node representing the cost of building a service center at that node. The total cost we wish to minimize is then the sum of the start-up costs of the chosen medians plus the cost of servicing each of the customer requests. This is known as the *facility location problem.*

* This work partially supported by Hong Kong RGC grants HKUST6010/01E, HKUST6162/00E, HKUST6082/01E and HKUST6206/02. The authors would like to thank Gerhard Trippen for his help in proofreading and latexing the figures.

T. Hagerup and J. Katajainen (Eds.): SWAT 2004, LNCS 3111, pp. 102–113, 2004.

The k-Median on a Line Problem (kML)

Let $k \geq 0$. Let $x_1 < x_2 < \cdots < x_n$ be points on the real line. With each point x_j there are associated a *weight* $w_j \geq 0$ and a *start-up cost* $c_j \geq 0$. A *k-placement* is a subset $S \subseteq V_m = \{x_1, \ldots, x_m\}$ of size $|S|$ at most k. We define the *distance* of point x_j to S by $d_j(S) = \min_{y \in S} |x_j - y|$. The *cost* of S is (i) the cost of creating the service centers in S plus (ii) the cost of servicing all of the requests from S:

$$cost(S) = \sum_{x_i \in S} c_i + \sum_{j=1}^{n} w_j d_j(S) \,.$$

The *k-median on a line problem (kML)* is to find a k-placement S minimizing $cost(S)$. In *online* kML, the points are given to us in the order $x_1, x_2, \ldots$, and we have to compute optimal solutions for the known points at any time.

Fig. 1. The k-median on a line problem.

Lin and Vitter [7] proved that, in general, even finding an approximate solution to the k-median problem is NP-hard. They were able to show, though, that it is possible in polynomial time to achieve a cost within $O(1+\epsilon)$ of optimal if one is allowed to use $(1 + 1/\epsilon)(\ln n + 1)k$ medians. The problem remains hard if restricted to metric spaces. Guha and Khuller [5] proved that this problem is still MAX-SNP hard. Charikar, Guha, Tardos and Shmoys [4] showed that constant-factor approximations can be computed for any metric space. In the specific case of points in Euclidean space, Arora, Raghavan, and Rao [2] developed a PTAS.

There are some special graph topologies for which fast polynomial time algorithms exist, though. In particular, this is true for trees [8,10] and lines [6]. In this paper we will concentrate on the line case, in which all of the nodes lie on the real line and the distance between any two nodes is the Euclidean distance. See Fig. 1 for the exact definition of the k-median on a line problem (kML).

There is a straightforward $O(kn^2)$ dynamic programming (DP) algorithm for solving kML. It fills in $\Theta(kn)$ entries in a dynamic programming table[1] where calculating each entry requires minimizing over $O(n)$ values, so the entire algorithm needs $O(kn^2)$ time. Hassin and Tamir [6] showed that this DP formulation possesses a quadrangle or concavity property. Thus, the time to calculate the table entries can be reduced by an order of magnitude to $O(kn)$ using known DP speed-up techniques, such as those found in [9].

In this paper we study online kML. Since static kML can be solved in $O(kn)$ time our hope would be to be able to add new points in $O(k)$ time. The difficulty here is that Hassin and Tamir's approach cannot be made online because most DP speed-up techniques such as in [9] are inherently static. The best that can

[1] We do not give the details here because the DP formulation is very similar to the one shown in Lemma 1.

The k-Coverage on a Line Problem (kCL)

In addition to the requirements of kML, each node x_j is also given a *coverage radius* r_j. It is *covered* by a k-placement S if $d_j(S) \le r_j$. In that case, the service cost for x_j is zero. Otherwise, the service cost is w_j. The cost of S is then

$$cost(S) = \sum_{x_i \in S} c_i + \sum_{j=1}^{n} w_j I_j(S) ,$$

where $I_j(S) = 0$ if $d_j(S) \le r_j$ and $I_j(S) = 1$ if $d_j(S) > r_j$. The *k-coverage on a line problem (kCL)* is to find a k-placement S minimizing $cost(S)$. *Online* kCL is defined similarly to online kML.

Fig. 2. The k-coverage on a line problem.

be done using their approach is to totally recompute the dynamic programming matrix entries from scratch at each step using $O(kn)$ time per step[2].

Later, Auletta, Parente and Persiano [3] studied kML in the special case of *unit lengths*, i.e., $x_{i+1} = x_i + 1$ for all i, and *no start up costs*, i.e., $c_i = 0$ for all i. Being unaware of Hassin and Tamir's results they developed a new online technique for solving the problem which enabled them to add a new point in *amortized* $O(k)$ time, leading to an $O(kn)$ time algorithm for the static problem.

The major contribution of this paper is to bootstrap off of Auletta, Parente and Persiano's result to solve online kML when (i) the points can have arbitrary distances between them and (ii) start up costs are allowed. In Section 2 we prove the following theorem.

Theorem 1. *We can solve the online k-median on a line problem in $O(k)$ amortized and $O(k \log n)$ worst case time per update. These time bounds hold simultaneously.* □

A variant of kML is the *k-coverage problem (kCL)* where the cost of servicing customer x_j is zero if it is closer than r_j to a service center, or w_j otherwise. See Fig. 2 for the exact definition of kCL.

Hassin and Tamir [6] showed how to solve static kCL in $O(n^2)$ time (independent of k), again using the quadrangle inequality/concavity property. In Section 3 we restrict ourselves to the special case of uniform coverage, i.e., there is some $r > 0$ such that $r_j = r$ for all j. In this situation we can use a similar (albeit much simpler) approach as in Section 2 to maintain optimal partial solutions S as points are added to the right of the line. In Section 3 we will develop the following theorem.

[2] Although not stated in [6] it is also possible to reformulate their DP formulation in terms of finding row-minima in k $n \times n$ *totally monotone matrices* and then use the SMAWK algorithm [1] — which finds the row-minima of an $n \times n$ totally monotone matrix in $O(n)$ time — to find another $O(kn)$ solution. This was done explicitly in [11]. Unfortunately, the SMAWK algorithm is also inherently static, so this approach also can not be extended to solve the online problem.

Theorem 2. *We can solve the online k-coverage on a line problem with uniform coverage in $O(k)$ amortized and $O(k \log n)$ worst case time per update. These time bounds hold simultaneously.* □

2 The k-Median Problem

2.1 Notations and Preliminary Facts

In the online k-median problem, we start with an empty line and, at each step, append a new node to the right of all of the previous nodes. So, at step m we will have m points $x_1 < x_2 < \cdots < x_m$ and when adding the $(m+1)$st point we have $x_m < x_{m+1}$. Each node x_j will have a weight w_j, and a start-up cost c_j associated with it. At step m, the task is to pick a set S of at most k nodes from $x_1, x_2, \ldots, x_m$ that minimizes $cost(S) = \sum_{x_i \in S} c_i + \sum_{j=1}^{m} w_j d_j(S)$.

Our algorithm actually keeps track of $2k$ median placements for every step. The first k placements will be optimal placements for *exactly* $i = 1, \ldots, k$ resources, i.e., let

$$OPT_i(m) = \min_{S \subseteq V_m,\, |S|=i} \left(\sum_{x_i \in S} c_i + \sum_{j=1}^{m} w_j d_j(S) \right) .$$

The remaining k placements are *pseudo-optimal* placements with the additional constraint that x_m must be one of the chosen resources. That is, for $i = 1, \ldots, k$

$$POPT_i(m) = \min_{S \subseteq V_m,\, |S|=i, x_m \in S} \left(\sum_{x_i \in S} c_i + \sum_{j=1}^{m} w_j d_j(S) \right) .$$

In particular, if $i = 1$, then $S = \{x_m\}$ and $POPT_1(m) = c_m + \sum_{j=1}^{m-1} w_j (x_m - x_j)$. Optimal and pseudo-optimal placements are related by the following straightforward equations.

Lemma 1.

$$OPT_i(m) = \min_{1 \le j \le m} \left(POPT_i(j) + \sum_{l=j+1}^{m} w_l \cdot d(j,l) \right) \quad \text{and} \tag{1}$$

$$POPT_i(m) = \min_{1 \le j \le m-1} \left(OPT_{i-1}(j) + \sum_{l=j+1}^{m-1} w_l \cdot d(l,m) \right) + c_m , \tag{2}$$

where $d(j,l) = x_l - x_j$ is the distance between x_j and x_l. □

Denote by $MIN_i(m)$ the index j at which the "min" operation in Eq. (1) achieves its minimum value and by $PMIN_i(m)$ the index j at which the "min" operation in Eq. (2) achieves its minimum value. When computing the $OPT_i(m)$

and $POPT_i(m)$ values the algorithm will also compute and keep the $MIN_i(m)$ and $PMIN_i(m)$ indices.

The optimum cost we want to find is $OPT = \min_{1 \le i \le k}(OPT_i(n))$. It is not difficult to see that, knowing all values of $OPT_i(m)$, $MIN_i(m)$, $POPT_i(m)$ and $PMIN_i(m)$ for $1 \le i \le k$, $1 \le m \le n$, we can unroll the equations in Lemma 1 in $O(k)$ time to find the optimal set S of at most k medians that yields OPT. So, maintaining these $4nk$ variables suffices to solve the problem.

A straightforward calculation of the minimizations in Lemma 1 permits calculating the value of $POPT_i(m)$ from those of $OPT_{i-1}(j)$ in $O(m)$ time and the value of $OPT_i(m)$ from those of $POPT_i(j)$ in $O(m)$ time. This permits a dynamic programming algorithm that calculates *all* of the $OPT_i(m)$ and $POPT_i(m)$ values in $O\left(k \sum_{m=1}^{n} m\right) = O(kn^2)$ time, solving the problem.

As discussed in the previous section, this is very slow. The rest of this section is devoted to improving this by an order of magnitude; developing an algorithm that, at step m for each i, will calculate the value of $POPT_i(m)$ from those of $OPT_{i-1}(m)$ and the value of $OPT_i(m)$ from those of $POPT_i(m)$ in $O(1)$ amortized time and $O(\log n)$ worst case time.

2.2 The Functions $V_i(j, m, x)$ and $V'_i(j, m, x)$

As mentioned, our algorithm is actually an extension of the algorithm in [3]. In that paper, the authors defined two sets of functions which played important roles. We start by rewriting those functions using a slightly different notation which makes it easier to generalize their use. For all $1 \le i \le k$ and $1 \le j \le m$ define

$$V_i(j, m, x) = POPT_i(j) + \sum_{l=j+1}^{m} w_l \cdot d(j, l) + x \cdot d(j, m) . \tag{3}$$

For all $1 \le i \le k$ and $1 \le j \le m-1$ define

$$V'_i(j, m, x) = OPT_{i-1}(j) + \sum_{l=j+1}^{m-1} w_l \cdot d(l, m) + x \cdot \sum_{l=j+1}^{m-1} w_l . \tag{4}$$

Then Lemma 1 can be written as $OPT_i(m) = \min_{1 \le j \le m} V_i(j, m, 0)$ and $POPT_i(m) = \min_{1 \le j \le m-1} V'_i(j, m, 0) + c_m$.

The major first point of departure between this section and [3] is the following lemma, which basically says that $V_i(j, m, x)$ and $V'_i(j, m, x)$ can be computed in constant time when needed.

Lemma 2. *Suppose we are given* $W(m) = \sum_{l=1}^{m} w_l$ *and* $M(m) = \sum_{l=1}^{m} w_l \cdot d(1, l)$. *Then, given the values of* $POPT_i(j)$*, the function* $V_i(j, m, x)$ *can be evaluated at any* x *in constant time. Similarly, given the values of* $OPT_{i-1}(j)$*, the function* $V'_i(j, m, x)$ *can be evaluated at any* x *in constant time.*

Proof. We first examine $V_i(j, m, x)$. We already know $POPT_i(j)$ so we only need to compute the terms $\sum_{l=j+1}^{m} w_l \cdot d(j, l) + x \cdot d(j, m)$. We can compute

$\sum_{l=j+1}^{m} w_l \cdot d(j,l) = [M(m) - M(j)] - [W(m) - W(j)] \cdot d(1,j)$ in constant time. For $V_i'(j,m,x)$, we also only need to compute $\sum_{l=j+1}^{m-1} w_l \cdot d(l,m) + x \cdot \sum_{l=j+1}^{m-1} w_l$. But we can compute $\sum_{l=j+1}^{m-1} w_l \cdot d(l,m) = [W(m-1) - W(j)] \cdot d(1,m) - [M(m-1) - M(j)]$ and $\sum_{l=j+1}^{m-1} w_l = W(m-1) - W(j)$ in constant time. □

In the next two subsections we will see how to use this lemma to efficiently calculate $POPT_i(j)$ and $OPT_i(j)$.

2.3 Computing $OPT_i(m)$

We start by explaining how to maintain the values of $OPT_i(m)$. Our algorithm uses k similar data structures to keep track of the k sets of $OPT_i(m)$ values, for $1 \leq i \leq k$. Since these k structures are essentially the same we will fix i and consider how the i^{th} data structure permits the computation of the values of $OPT_i(m)$ as m increases.

The Data Structures. Recall Eq. (3). Consider the m functions $V_i(j,m,x)$ for $1 \leq j \leq m$. They are all linear functions in x so the lower envelope of these functions is a piecewise linear function to which each $V_i(j,m,x)$ contributes at most one segment.

We are only interested in $OPT_i(m) = \min_{1 \leq j \leq m} V_i(j,m,0)$ which is equivalent to evaluating this lower envelope at $x = 0$. In order to update the data structure efficiently, though, we will see that we will need to store the *entire* lower envelope for $x \geq 0$. We store the envelope by storing the *changes* in the envelope. More specifically, our data structures for computing the values of $OPT_i(m)$ consist of two arrays

$$\Delta_i(m) = (\delta_0, \delta_1, \ldots, \delta_s) \quad \text{and} \quad Z_i(m) = (z_1, \ldots, z_s), \tag{5}$$

such that

$$\text{if } \delta_{h-1} < x + W(m) < \delta_h, \text{ then } V_i(z_h, m, x) = \min_{j \leq m} V_i(j,m,x)\,. \tag{6}$$

The reasons for the shift term $W(m) = \sum_{l=1}^{m} w_l$ will become clear later. Since we only keep the lower envelope for $x \geq 0$, we have $\delta_0 < W(m) < \delta_1$.

An important observation is that the slope of $V(j,m,x)$ is $d(j,m)$ which decreases as j increases, so we have $z_1 < \cdots < z_s$ and $z_s = m$ at step m. In particular, note that $V(m,m,x)$, which is the rightmost part of the lower envelope, has slope $0 = d(m,m)$ and is a horizontal line.

Given such a data structure, computing the value of $OPT_i(m)$ becomes trivial. We simply have $MIN_i(m) = z_1$ and $OPT_i(m) = V_i(z_1, m, 0)$.

Updating the Data Structures. Assume that the data structure given by Eq. (5) and (6) is storing the lower envelope after step m and, in step $m+1$, point

x_{m+1} is added. We now need to recompute the lower envelope of $V_i(j, m+1, x)$, for $1 \le j \le m+1$ and $x \ge 0$. Note that in step m we have m functions

$$\{V_i(j, m, x) \,:\, 1 \le j \le m\}$$

but we now have $m+1$ functions

$$\{V_i(j, m+1, x) \,:\, 1 \le j \le m+1\}\,.$$

If we only consider the lower envelope of the first m functions $V_i(j, m+1, x)$ for $1 \le j \le m$, then the following lemma guarantees that the two arrays $\Delta_i(m)$ and $Z_i(m)$ do not change.

Lemma 3. *Assume $V_i(z_h, m, x)$ minimizes $V_i(j, m, x)$ for $1 \le j \le m$ when $\delta_{h-1} < x + W(m) < \delta_h$. Then $V_i(z_h, m+1, x)$ minimizes $V_i(j, m+1, x)$ for $1 \le j \le m$ when $\delta_{h-1} < x + W(m+1) < \delta_h$.*

Proof. It is easy to verify that for $1 \le j \le m$

$$V_i(j, m+1, x) = V_i(j, m, x + w_{m+1}) + (x + w_{m+1}) \cdot d(m, m+1)\,.$$

Since $\delta_{h-1} < x + W(m+1) < \delta_h$ iff $\delta_{h-1} < (x + w_{m+1}) + W(m) < \delta_h$, the above formula is minimized when $j = z_h$. □

This lemma is the reason for defining Eq. (5)and (6) as we did with the shift term instead of simply keeping the breakpoints of the lower envelope in $\Delta_i(m)$.

Note that the lemma *does not* say that the lower envelope of the functions remains the same (this could not be true since all of the functions have been changed). What the lemma does say is that the *structure* of the breakpoints of the lower envelope is the same after the given shift.

Now, we consider $V_i(m+1, m+1, x)$. As discussed in the previous subsection, $V_i(m+1, m+1, x)$ is the rightmost segment of the lower envelope *and* is a horizontal line. So, we only need to find the intersection point between the lower envelope of $V_i(j, m+1, x)$ for $1 \le j \le m$ and the horizontal line $y = V_i(m+1, m+1, x)$. Assume they intersect at the segment $V_i(z_{max}, m+1, x)$. Then, $Z_i(m+1)$ becomes $(z_1, \ldots, z_{max}, m+1)$, and $\Delta_i(m+1)$ changes correspondingly.

We can find this point of intersection either by using a binary search or a sequential search. The binary search would require $O(\log m)$ worst case comparisons between $y = V_i(m+1, m+1, x)$ and the lower envelope. The sequential search would scan the array $Z_i(m)$ from right to left, i.e. from z_s to z_1, discarding segments from the lower envelope until we find the intersection point of $y = V_i(m+1, m+1, x)$ with points on the lower envelope. The sequential search might take $\Theta(m)$ time in the worst case but only uses $O(1)$ in the amortized case since lines thrown off the lower envelope will never be considered again in a later step.

In both methods a comparison operation requires being able to compare the constant $V_i(m+1, m+1, x)$ to $V_i(j, m+1, x)$ for some j and some arbitrary value m. Recall from Lemma 2 that we can evaluate $V_i(j, m+1, x)$ at any particular

x in constant time. Thus, the total time required to update the lower envelope is $O(\log m)$ worst case and $O(1)$ amortized.

To combine the two bounds we perform the sequential and binary search alternately, i.e. , we use sequential search in odd numbered comparisons and binary search in even numbered comparisons. The combined search finishes when the intersection value is first found. Thus, the running time is proportional to the one that finishes first and we achieve both the $O(1)$ amortized time and the $O(\log m)$ worst case time.

Since we only keep the lower envelope for $x \geq 0$, we also need to remove from $Z_i(m+1)$ and $\Delta_i(m+1)$ the segments corresponding to negative x values. Set $z_{min} = \max\{z_h : \delta_{h-1} < W(m+1) < \delta_h\}$. Then $Z_i(m+1)$ should be $(z_{min}, \ldots, z_{max}, m+1)$, and $\Delta_i(m+1)$ should change correspondingly.

To find z_{min}, we also use the technique of combining sequential search and binary search. In the sequential search, we scan from left to right, i.e., from z_1 to z_s. The combined search also requires $O(1)$ amortized time and $O(\log m)$ worst case time.

2.4 Computing $POPT_i(m)$

In the previous section we showed how to update the values of $OPT_i(m)$ by maintaining a data structure that stores the lower envelope of $V_i(j, m, x)$ and evaluating the lower envelope at $x = 0$, i.e., $OPT_i(m) = \min_{1 \leq j \leq m} V_i(j, m, 0)$. In this section we will show how, in a very similar fashion, we can update the values of $POPT_i(m)$ by maintaining a data structure that stores the lower envelope of $V_i'(j, m, x)$. Note that

$$POPT_i(m) = c_m + \min_{1 \leq j \leq m-1} V_i'(j, m, 0),$$

i.e., evaluating the lower envelope at $x = 0$ and adding c_m.

As before we will be able to maintain the lower envelope of $V_i'(j, m, x)$, $1 \leq j \leq m-1$, in $O(1)$ amortized time and $O(\log m)$ worst case time. The data structure is almost the same as the one for maintaining $V_i(j, m, x)$ in the previous section so we only quickly sketch the ideas.

As before the algorithm uses k similar data structures to keep track of the k lower envelopes; for our analysis we fix i and consider the data structures for maintaining the lower envelope of $V_i'(j, m, x)$ (and thus $POPT_i(m)$) as m increases.

The Data Structures. By their definitions the $m-1$ functions $V_i'(j, m, x)$, for $1 \leq j \leq m-1$, are all linear functions, so their lower envelope is a piecewise linear function to which each $V_i(j, m, x)$ contributes at most one segment.

As before, in order to compute the values of $POPT_i(m)$, we only need to know the value of the lower envelope at $x = 0$ but, in order to update the structure efficiently, we will need to store the entire lower envelope.

Our data structures for computing the values of $POPT_i(m)$ consist of two arrays

$$\Delta_i'(m) = (\delta_0', \delta_1', \ldots, \delta_s') \quad \text{and} \quad Z_i'(m) = (z_1', \ldots, z_s'), \tag{7}$$

such that

$$\text{if } \delta_{h-1}' < x + d(1,m) < \delta_h', \text{ then } V_i'(z_h', m, x) = \min_{j \le m-1} V_i'(j, m, x)\,. \tag{8}$$

Since we only keep the lower envelope for $x \geq 0$, we have $\delta_0' < d(1,m) < \delta_1'$. Since the slopes $\left(\sum_{l=j+1}^{m-1} w_l\right)$ of $V_i'(j,m,x)$ decrease when j increases, we have $z_1' < \cdots < z_s'$ and $z_s' = m-1$ at step m. In particular, note that $V'(m-1,m,x)$, the rightmost part of the lower envelope, has slope 0 and is a horizontal line.

Given such data structures, computing the value of $POPT_i(m)$ becomes trivial. We simply have $PMIN_i(m) = z_1'$ and $POPT_i(m) = c_m + V_i'(z_1', m, 0)$.

Updating the Data Structures. Given the lower envelope of $V_i'(j,m,x)$, for $1 \leq j \leq m-1$ at step m we need to be able to recompute the lower envelope of $V_i'(j, m+1, x)$, for $1 \leq j \leq m$ after x_{m+1} is added.

As before, we will first deal with the functions $V_i'(j, m+1, x)$ for $1 \leq j \leq m-1$, and then later add the function $V_i'(m, m+1, x)$.

If we only consider the functions $V_i'(j, m+1, x)$ for $1 \leq j \leq m-1$, we have an analogue of Lemma 3 for this case that guarantees that the two arrays $\Delta_i'(m)$ and $Z_i'(m)$ do not change.

Lemma 4. *Assume* $V_i'(z_h', m, x)$ *minimizes* $V_i'(j,m,x)$ *for* $1 \leq j \leq m-1$ *when* $\delta_{h-1}' < x + d(1,m) < \delta_h'$. *Then* $V_i'(z_h', m+1, x)$ *minimizes* $V_i'(j, m+1, x)$ *for* $1 \leq j \leq m-1$ *when* $\delta_{h-1}' < x + d(1, m+1) < \delta_h'$.

Since the proof is almost exactly the same as that of Lemma 3 we do not, in this extended abstract, provide further details.

We note that, using *exactly* the same techniques as in the comments following Lemma 3, we can update the lower envelope of $V_i'(j,m,x)$ for $1 \leq j \leq m-1$ to the lower envelope of $V_i'(j, m+1, x)$ for $1 \leq j \leq m$ using a combined binary/sequential search that takes both $O(1)$ amortized and $O(\log m)$ worst case time per step (simultaneously).

2.5 The Algorithm

Given the data structures developed in the previous section the algorithm is very straightforward. After nodes $x_1 < x_2 < \cdots < x_m$ have been processed in step m the algorithm maintains

- $W(m) = \sum_{l=1}^{m} w_l$ and $M(m) = \sum_{l=1}^{m} w_l \cdot d(1,l)$.
- For $1 \leq i \leq k$, the data structures described in Sections 2.3 and 2.4 for storing the lower envelopes $\min_{j \le m} V_i(j,m,x)$ and $\min_{j \le m-1} V_i'(j,m,x)$.
- For $1 \leq i \leq k$ and $1 \leq j \leq m$, all of the values $OPT_i(j)$, $POPT_i(j)$ and corresponding indices $MIN_i(j)$, $PMIN_i(j)$.

Table 1. The values of $OPT_i(m)/MIN_i(m)$ in the upper rows and $POPT_i(m)/PMIN_i(m)$ in the lower rows.

		m								
		1	2	3	4	5	6	7	8	9
i	1	5400/1	2170/2	2264/2	691/4	761/4	785/4	1955/4	5241/4	6337/5
		5400/-	2170/-	3322/-	691/-	939/-	11048/-	18362/-	22093/-	32721/-
	2	-/-	7500/2	5270/3	2364/4	691/5	699/5	1817/5	4997/5	6089/5
		-/-	7500/1	5270/2	2364/3	691/4	10626/4	8885/6	8927/6	15649/6
	3	-/-	-/-	10600/3	5370/4	2364/5	2372/5	3490/5	6670/5	7762/5
		-/-	-/-	10600/2	5370/3	2364/4	10591/5	8799/6	8841/6	15563/6

After adding x_{m+1} with associated values w_{m+1} and c_{m+1} the algorithm updates its data structures by

- Calculating $W(m+1) = W(m) + w_{m+1}$ and $M(m+1) = M(m) + w_{m+1}d(1, m+1)$ in $O(1)$ time.
- Updating the $2k$ lower envelopes as described in Sections 2.3 and 2.4 in $O(\log m)$ worst case and $O(1)$ amortized time (simultaneously) per envelope.
- For $1 \leq i \leq k$, calculating $OPT_i(m+1) = \min_{j \leq m+1} V_i(j, m+1, 0)$ and $POPT_i(m+1) = c_m + \min_{j \leq m} V_i'(j, m+1, 0)$ in $O(1)$ time each.

Thus, in each step, the algorithm uses, as claimed, only a total of $O(k \log n)$ worst case and $O(k)$ amortized time (simultaneously).

The algorithm above only fills in the dynamic programming table. But, given the values $OPT_i(j)$, $POPT_i(j)$ and the corresponding indices $MIN_i(j)$, $PMIN_i(j)$ one can construct the optimal set of medians in $O(k)$ time so this fully solves the problem and finishes the proof of Theorem 1.

2.6 An Example

In this example, let $n = 9$ be the total number of nodes, and $k = 3$ the maximum number of resources. The nodes have x-coordinates 0, 5, 7, 10, 12, 13, 55, 72, 90, start-up costs c_j 5400, 2100, 3100, 100, 0, 9900, 8100, 7700, 13000, and weights w_j 14, 62, 47, 51, 35, 8, 26, 53, 14. Table 1 shows the values of OPT, MIN, $POPT$ and $PMIN$, respectively. From these tables, we can see that for $m = 9$ the optimal placement is to have two resources at x_4 and x_5.

Figure 3 shows the functions $V_2(j, 8, x)$ and $V_2(j, 9, x)$. The two arrays for the lower envelope when $m = 8$ are $Z_2(8) = (5, 8)$ and $\Delta_2(8) = (296, 361.5, +\infty)$. The two arrays for the lower envelope when $m = 9$ are $Z_2(9) = (5, 8, 9)$ and $\Delta_2(9) = (310, 361.5, 669.4, +\infty)$. As we can see, the intersection point of line 5 and line 8 in the left part of Figure 3 shifts to the left by w_9 when we add x_9 in the next step (right half of the figure), i.e., from 65.6 to 51.5. Actually, all intersection points will shift the same amount when a new node is added. That is why the partitioning value 361.5 does not change in the arrays $\Delta_2(8)$ and $\Delta_2(9)$ $(361.5 = 65.5 + W(8) = 51.5 + W(9))$.

3 The k-Coverage Problem

In this section we sketch how to solve online kCL with uniform coverage, i.e., to maintain a k-placement S minimizing

$$cost(S) = \sum_{x_i \in S} c_i + \sum_{j=1}^{n} w_j I_j(S)\,,$$

as m grows, where r is some fixed constant and $I_j(S) = 0$ if $d_j(S) \leq r_j$ and $I_j(S) = 1$ if $d_j(S) > r_j$. This problem has a simpler DP solution than the k-median problem, albeit one with a similar flavor.

We say that x_j is *covered* by a point in S if $d_j(S) \leq r$. For a point x_j, let cov_j denote the index of the smallest of the points $x_1, \ldots, x_j$ covered by x_j, and unc_j the index of largest of the points $x_1, \ldots, x_j$ not covered by x_j:

$$cov_j = \min\{i \,:\, i \leq j \text{ and } r + x_i \geq x_j\}, \qquad unc_j = \max\{i \,:\, i < j \text{ and } r + x_i < x_j\}.$$

Note that x_{unc_j} is the point to the left of x_{cov_j}, i.e., $unc_j = cov_j - 1$ if this point exists. The points that can cover x_j are exactly the points in $[x_{cov_j}, x_j]$. Similar to the k-median problem, let $OPT_i(m)$ denote the minimum cost of an i-cover for the first m points $x_1, \ldots, x_m$, for $i = 1, \ldots, k$ and $POPT_i(m)$ be the minimum cost of covering $x_1, \ldots, x_m$ if x_m is one of the resources. Then

$$OPT_i(m) = \min\left\{ w_m + OPT_i(m-1)\,,\ \min_{cov_m \leq j \leq m} POPT_i(j) \right\} \tag{9}$$

$$POPT_i(m) = c_m + \min_{unc_m \leq j \leq m-1} OPT_{i-1}(j). \tag{10}$$

The first term in the minimum of Eq. (9) corresponds to the possibility that x_m is not covered; the second term to the possibility that x_m is covered. It ranges over all possible covers.

In order to solve the problem in an online fashion we will need to be able to calculate the values of $OPT_i(m)$ and $POPT_i(m)$ efficiently at step m when processing x_m. This can be done in a fashion similar to that employed for kML

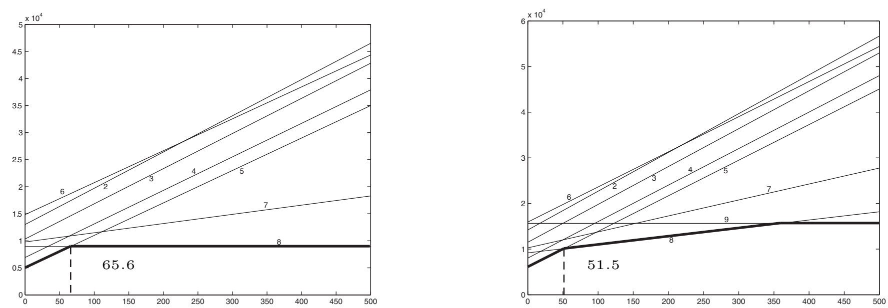

Fig. 3. Functions $V_2(j, 8, x)$, for $j = 2, \ldots, 8$, and $V_2(j, 9, x)$, for $j = 2, \ldots, 9$. The lines are labelled by j. The thick lines are the lower envelopes.

resulting in a similar result, i.e., computing all of the $OPT_i(m)$ and $POPT_i(m)$ values can be in $O(k)$ amortized time and $O(k \log m)$ worse case time per update. The details can be found in the full version of this paper.

4 Conclusion and Open Problems

In this paper we discussed how to solve the online k-median on a line problem in $O(k)$ amortized time and $O(k \log n)$ worst case time per point addition. This algorithm maintains in the online model the dynamic programming speed-up for the problem that was first demonstrated for the static version of the problem in [6]. The technique used is a generalization of one introduced in [3]. We also showed how a simpler form of our approach can solve the online k-coverage on a line problem with uniform coverage radius in the same time bounds. It is not clear how to extend our ideas to the non-uniform coverage radius case.

A major open question is how to solve the *dynamic* k-median and k-coverage on a line problem. That is, points will now be allowed to be inserted (or deleted!) anywhere on the line and not just on the right hand side. In this case would it be possible to maintain the k-medians or k-covers any quicker than recalculating them from scratch each time?

References

1. A. Aggarwal, M. Klawe, S. Moran, P. Shor, and R. Wilber. "Geometric applications of a matrix-searching algorithm," *Algorithmica,* **2**(2), 1987, pp. 195-208.
2. Sanjeev Arora, Prabhakar Raghavan, and Satish Rao. "Approximation schemes for Euclidean k-medians and related problems," *Proceedings of the 30th Annual ACM Symposium on the Theory of Computing*, 1998, pp. 106-113.
3. V. Auletta, D. Parente and G. Persiano. "Placing resources on a growing line," *Journal of Algorithms*, **26**, 1998, pp. 87-100.
4. M. Charikar, S. Guha, E. Tardos, and D.B. Shmoys. "A constant-factor approximation algorithm for the k-median problem," *Journal Computer System Sciences*, **65**, 2002, pp. 129-149.
5. Sudipto Guha, Samir Khuller. "Greedy strikes back: improved facility location algorithms," *Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA'98)*, 1998, pp. 649-657.
6. R. Hassin and A. Tamir. "Improved complexity bounds for location problems on the real line." *Operations Research Letters* **10**, 1991, pp. 395-402.
7. J-H. Lin and J.S. Vitter. "ϵ-approximations with minimum packing constraint violation," *Proceedings of the 24th Annual ACM Symposium on the Theory of Computing*, 1992, pp. 771-782.
8. A. Tamir. "An $O(pn^2)$ algorithm for the p-median and related problems on tree graphs," *Operations Research Letters*, **19** 1996,pp. 59-64.
9. R. Wilber. "The concave least-weight subsequence problem revisited," *Journal of Algorithms,* **9**, 1988, pp. 418-425.
10. A. Vigneron, L. Gao, M. Golin, G. Italiano and B. Li. "An algorithm for finding a k-median in a directed tree," *Information Processing Letters,* **74**, 2000, pp. 81-88.
11. G. Woeginger. "Monge strikes again: optimal placement of web proxies in the internet." *Operations Research Letters,* **27**, 2000 , pp. 93-96.