Tight Analysis of a Multiple-Swap Heurstic for Budgeted Red-Blue Median*

Zachary Friggstad†[1] and Yifeng Zhang[2]

1 **Department of Computing Science, University of Alberta, Edmonton, Canada**
`zacharyf@ualberta.ca`

2 **Department of Computing Science, University of Alberta, Edmonton, Canada**
`yifeng2@ualberta.ca`

Abstract

Budgeted Red-Blue Median is a generalization of classic k-Median in that there are two sets of facilities, say $\mathcal{R}$ and $\mathcal{B}$, that can be used to serve clients located in some metric space. The goal is to open k_r facilities in $\mathcal{R}$ and k_b facilities in $\mathcal{B}$ for some given bounds k_r, k_b and connect each client to their nearest open facility in a way that minimizes the total connection cost.

We extend work by Hajiaghayi, Khandekar, and Kortsarz [2012] and show that a *multiple-swap* local search heuristic can be used to obtain a $(5+\epsilon)$-approximation for Budgeted Red-Blue Median for any constant $\epsilon > 0$. This is an improvement over their single swap analysis and beats the previous best approximation guarantee of 8 by Swamy [2014].

We also present a matching lower bound showing that for every $p \geq 1$, there are instances of Budgeted Red-Blue Median with local optimum solutions for the p-swap heuristic whose cost is $5 + \Omega\left(\frac{1}{p}\right)$ times the optimum solution cost. Thus, our analysis is tight up to the lower order terms. In particular, for any $\epsilon > 0$ we show the single-swap heuristic admits local optima whose cost can be as bad as $7 - \epsilon$ times the optimum solution cost.

1998 ACM Subject Classification F.2.2 Computations on Discrete Structures, G.2.1 Combinatorial Algorithms, I.5.3 Clustering

Keywords and phrases Approximation Algorithms, Local search, Red-Blue Meidan

Digital Object Identifier 10.4230/LIPIcs.ICALP.2016.75

1 Introduction

Facility location problems crop up in many areas of computing science and operations research. A typical problem involves a set of clients and possible facility locations located in a metric space. The goal is to open some facilities and connect each client to some open facility as cheaply as possible. These problems become difficult when there are costs associated with opening facilities or additional constraints that ensure we cannot open too many facilities.

We study Budgeted Red-Blue Median, one particular instance of this type of problem. Here we are given a set of clients C, a set of *red* facilities $\mathcal{R}$, and a set of *blue* facilities $\mathcal{B}$. These are located in some metric space with metric distances $d(i,j) \geq 0$ for any two $i, j \in C \cup \mathcal{R} \cup \mathcal{B}$. Additionally, we are given two integer bounds $k_r \leq |\mathcal{R}|$ and $k_b \leq |\mathcal{B}|$. The

* The full version of this work is available at `http://arxiv.org/abs/1603.00973`.

† This research was undertaken, in part, thanks to funding from the Canada Research Chairs program and an NSERC Discovery Grant.

43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016).
Editors: Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi;
Article No. 75; pp. 75:1–75:13
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

goal is to select/open k_r red facilities R and k_b blue facilities B to minimize

$$\text{cost}(R \cup B) := \sum_{j \in C} \min_{i \in R \cup B} d(i, j) \,.$$

The classic NP-hard k-Median problem appears as a special case when, say, $\mathcal{R} = \emptyset$. Thus, Budgeted Red-Blue Median is NP-hard. In this paper, we focus on approximation algorithms for Budgeted Red-Blue Median, in particular on local search techniques.

1.1 Previous Work

The study of Budgeted Red-Blue Median from the perspective of approximation algorithms was initiated by Hajiaghayi, Khandekar, and Kortsarz [9], where they obtain a constant-factor approximation by a local search algorithm that iteratively tries to swap one red and/or one blue facility in the given solution. The do not specify the constant in their analysis, but it looks to be greater than 8. Citing [9] as inspiration, Krishnaswamy et al. studied a generalization of Budgeted Red-Blue Median known as Matroid Median [10]. Here, a matroid structure is given over the set of facilities and we can only open a set of facilities if they form an independent set in the matroid. They obtain a constant-factor approximation for Matroid Median through rounding an LP relaxation. This was later refined to an 8-approximation by Swamy [15].

The special case of k-Median is a classic optimization problem and has received a lot of attention from both theoretical and practical communities. The best approximation guarantee known so far is 2.675 by Byrka et al. [5], who build heavily on the breakthrough work of Li and Svensson for the problem [11].

While local search techniques have been used somewhat infrequently in the design of approximation algorithms in general, it may be fair to say that they have seen the most success in facility location problems. For almost 10 years, the best approximation for k-Median was based on a local search algorithm. Arya et al. [3] show that a multiple-swap heuristic leads to a $(3+\epsilon)$-approximation for k-Median for any constant $\epsilon > 0$. This analysis was simplified in [8], which inspires much of our analysis.

Another textbook application of local search is a $(1+\sqrt{2})$-approximation for Uncapacitated Facility Location [3, 6]. Local search has been very helpful in approximating Capacitated Facility Location, the first constant-factor approximation was by Pál, Tardos, and Wexler [13] and the current best approximation is a $(5+\epsilon)$-approximation by Bansal, Garg, and Gupta [4]. In the special case when all capacities are uniform, Aggarwal et al. [1] obtain a 3-approximation. Even more examples of local search applied to other facility location variants can be found in [2, 7, 8, 12, 14].

1.2 Our Results and Techniques

We show that a multiswap generalization of the local search algorithm considered in [9] is a $(5+\epsilon)$-approximation for Budgeted Red-Blue Median. That is, for a value p say the p-swap heuristic is the algorithm that, upon given an initial feasible solution, tries to swap up to p facilities of each colour. If no such swap produces a cheaper solution, it terminates. Otherwise, it iterates with the now cheaper solution. Algorithm 1 in Section 2 gives the formal description of our algorithm.

Say that a solution is locally optimum for the p-swap heuristic if no cheaper solution can be found by swapping up to p facilities of each colour. Let OPT denote the cost of an optimum solution. Our main result is the following.

▶ **Theorem 1.** *Any locally optimum solution for the p-swap heuristic has cost at most* $(5 + O(1/\sqrt{p})) \cdot OPT$.

Using standard techniques (briefly mentioned in Section 2), this readily leads to a polynomoial-time approximation algorithm. By choosing $p = \theta(1/\epsilon^2)$ we have the following.

▶ **Theorem 2.** *For any constant* $\epsilon > 0$, Budgeted Red-Blue Median *admits a polynomial-time* $(5 + \epsilon)$*-approximation.*

This improves over the 8-approximation for Budgeted Red-Blue Median in [15].

We emphasize the approximation guarantee from Theorem 1 result is for Budgeted Red-Blue Median only, the 8-approximation in [15] is still the best approximation for the general Matroid Median problem. Indeed, [10] show that Matroid Median cannot be approximated within any constant factor using any constant number of swaps even in the generalization of Budgeted Red-Blue Median where there can be a super-constant number facility colours.

We also present a lower bound that matches our analysis up the lower order terms.

▶ **Theorem 3.** *For any integers* p, ℓ *with* $1 \le p \le \ell/2$*, there is an instance of* Budgeted Red-Blue Median *that has a locally-optimum solution for the p-swap heuristic with cost at least* $\left(5 + \frac{2}{p} - \frac{10p}{\ell+1}\right) \cdot OPT$.

By letting $\ell \to \infty$ but keeping p fixed, we see that the p-swap heuristic cannot guarantee a ratio better than $5 + \frac{2}{p}$. So, Theorem 1 is tight up to lower order terms. Also, for $p = 1$ we see that the single-swap heuristic analyzed in [9] is not better than a 7-approximation.

Local search techniques are typically analyzed by constructing a set of candidate *test swaps* where some facilities in the optimum solution are swapped in and some from the local optimum are swapped out in order to generate a useful inequality. One of the main features of the k-Median analysis in [3] and [8] is that such swaps can be considered that ensure each facility in the global optimum is swapped in once and, by averaging some swaps, each facility in the local optimum is swapped out to the extent of at most $1 + O(\epsilon)$ times. Each time a facility in the local optimum is swapped out, they pay an additional 2 times the global optimum cost for some clients to reassign them.

We obtain only a $5 + \epsilon$ approximation because we end up swapping out some facilities in the local optimum solution to the extent of $2 + O(\epsilon)$, thereby paying an additional $2 + O(\epsilon)$ more than in the k-Median analysis. Ultimately, this is because some of our initial swaps generate inequalities that depend *positively* on client assignment costs in the local optimum. So we consider additional swaps that do not introduce any more positive dependence on the local optimum to cancel them out.

This issue was also encountered in the analysis in [9]. In some sense, we are showing that this is the only added difficulty over the standard k-Median analysis. However, the averaging arguments we use are a bit more sophisticated than the analysis for k-Median.

1.3 Organization

Section 2 presents the algorithm and describes some useful notation. In particular, it presents a way to decompose the global and local optimum solution into structured groups that are examined in the analysis. Section 3 analyzes the quality of locally optimum solutions to prove Theorem 1. Section 4 proves Theorem 3 with an explicit construction of a bad example. We conclude with some remarks in Section 5.

2 Notation and Preliminaries

Say that a *feasible solution* is a pair (R, B) of subsets $R \subseteq \mathcal{R}$ and $B \subseteq \mathcal{B}$ with $|R| = k_r$ and $|B| = k_b$. Algorithm 1 describes the local search algorithm.

Algorithm 1 The p-Swap Heuristic for BUDGETED RED-BLUE MEDIAN

```
Let (R, B) be an arbitrary feasible solution.
while there is some feasible solution (R', B') with |R − R'| ≤ p
        and |B − B'| ≤ p and cost(R' ∪ B') < cost(R ∪ B) do
    (R, B) ← (R', B')
end while
return (R, B)
```

While a single iteration of Algorithm 1 can be executed in $n^{O(p)}$ (where n is the total number of locations in the problem), it may be that the number of iterations is not polynomially bounded. We can employ a well-known trick to ensure it does terminate in a polynomial number of steps while losing only another ϵ in our analysis. The idea is to perform the update only if $\text{cost}(R' \cup B') \leq (1 - \epsilon/\Delta) \cdot \text{cost}(R \cup B)$ where Δ is some quantity that is polynomial in the input size. Our analysis is compatible with this approach; one can check that the total weight of all inequalities we consider is polynomially bounded. For example, see [3] for details. We do not focus any further on this issue, and instead work toward analyzing the cost of the solutions produced by Algorithm 1 as it is stated.

From now on, let $S = R \cup B$ with $R \subseteq \mathcal{R}, B \subseteq \mathcal{B}$ denote an arbitrary local optimum solution. That is, there is no cheaper solution (R', B') with $|R - R'| \leq p$ and $|B - B'| \leq p$. Also fix a global optimum solution $O = R^* \cup B^*$ where $R^* \subseteq \mathcal{B}$ and $B^* \subseteq \mathcal{B}$. We assume that $S \cap O = \emptyset$. This is without loss of generality, as we can duplicate each facility location in the input and say that S use the first copies and O use the second copies. It is easy to verify that S is still a local optimum solution.

To help analyze the cost, we will introduce some notation. For any client $j \in \mathcal{C}$, let $s_j \in S$ denote the local optimum facility is closest to j and $o_j \in O$ denote the global optimum facility that is closest to j, breaking ties arbitrarily. For brevity, let $c_j = d(j, s_j)$ be the cost of assigning j in the local optimum and $c_j^* = d(j, o_j)$ the cost of assigning j in the global optimum. Thus, $\text{cost}(S) = \sum_{j \in C} c_j$ and $\text{cost}(O) = \sum_{j \in C} c_j^*$. For any facility $i^* \in O$ we let $N^*(i^*) = \{j \in C : o_j = i^*\}$ and for any $i \in S$ we let $N(i) = \{j \in C : s_j = i\}$.

Let $\phi : O \to S$ map each facility in O to its nearest facility in S, breaking ties arbitrarily. For $i \in S$, let $\deg(i) = |\phi^{-1}(i)|$. If $\deg(i) \neq 0$, let $\text{cent}(i)$ be the facility in $\phi^{-1}(i)$ that is closest to i, again breaking ties arbitrarily.

We also borrow some additional notation from [9].

▶ **Definition 4** (very good, good, bad facility)**.** A facility $i \in S$ is *very good* if $\deg(i) = 0$, *good* if no $i^* \in \phi^{-1}(i)$ has the same colour as i, and *bad* otherwise.

The analysis in [9] divides $S \cup O$ into *blocks* that satisfy certain properties. We require slightly stronger properties than their blocks guarantee. We also use a slightly different notion of what it means for some $i \in S$ to be a *leader*. The required properties are summarized in the following lemma, which also serves as our definition of a block. The proof can be found in the full version of our work.

▶ **Lemma 5.** *We can partition $S \cup O$ into* blocks T *satisfying the following properties.*

- $|T \cap R| = |T \cap R^*|$ *and* $|T \cap B| = |T \cap B^*|$.

- *For every $i \in S \cap T$, we also have $\phi^{-1}(i) \subseteq T$. For every $i^* \in O \cap T$, we have $\phi(i^*) \in T$.*
- *There is some facility $\hat{i} \in T \cap S$ with $\deg(\hat{i}) > 0$ designated as the* leader *that has the following properties. Every other $i \in T \cap S - \{\hat{i}\}$ is either good or very good and all good $i \in T \cap S - \{\hat{i}\}$ have the same colour.*

We will focus on analyzing one block at a time to prove the approximation guarantee. This provides us with a cleaner way to describe the test swaps and the additional structure will help handle the inevitable cases where we have to swap out some $i \in S$ but cannot swap in all of $\phi^{-1}(i)$. For example, this can happen if all blue facilities $i \in B$ have $\deg(i)$ being very large (so all $\deg(i') = 0$ facilities are red). We will still need to close some of them in order to open facilities in B^* when generating bounds via test swaps.

Before delving into the analysis we note the following two bounds. The first has been used extensively in local search analysis and was first proven in [3] and the second was proven in [9].

▶ **Lemma 6.** *For any $j \in C$, $d(j, \phi(o_j)) - c_j \le 2c^*_j$.*

▶ **Lemma 7.** *For any $j \in C$, $d(j, \mathrm{cent}(\phi(o_j))) - c_j \le 3c^*_j + c_j$.*

Finally, we often consider operations that add or remove a single item from a set (usually to exclude the leader $\hat{i}$ of a block from some parts of the analysis). To keep the notation cleaner we let $A + i$ and $A - i$ refer to $A \cup \{i\}$ and $A - \{i\}$, respectively, for sets of facilities A and a single facility i.

3 Multiswap Analysis

Recall that we are assuming $S = R \cup B$ is a locally optimum solution with respect to the heuristic that swaps at most p facilities of each colour and that $O = R^* \cup B^*$ is some globally optimum solution. We assume $p = t^2 + 1$ for some sufficiently large integer t.

Focus on a single block T. For brevity, let $T^*_R = T \cap R^*$ and $T^*_B = T \cap B^*$ denote the red and blue facilities from the optimum solution in T. Similarly let $T_R = T \cap R$ and $T_B = T \cap B$ denote the red and blue facilities from the local optimum solution in T. The main goal of this section is to demonstrate the following inequality for group T.

▶ **Theorem 8.** *For some absolute constant γ that is independent of t, we have*

$$0 \le \sum_{j \in N^*(T^*_R \cup T^*_B)} \left[\left(1 + \frac{\gamma}{t}\right) c^*_j - c_j\right] + \sum_{j \in N(T_R \cup T_B)} \left[\left(4 + \frac{\gamma}{t}\right) \cdot c^*_j + \frac{\gamma}{t} \cdot c_j\right].$$

Theorem 1 follows by summing over all associated inequalities for the various blocks.

The analysis breaks into a number of cases based on whether T^*_R and/or T^*_B are large. In each of the cases, we use the following notation and assumptions. Let $\hat{i}$ denote the leader in T. Without loss of generality, assume all other $i \in T_B \cup T_R$ with $\deg(i) > 0$ are blue facilities. Let $\overline{B} = \{i \in T_B - \hat{i} : \deg(i) > 0\}$, so $\mathrm{cent}(\overline{B})$ denotes $\{i^* \in T^*_B \cup T^*_R : \mathrm{cent}(\phi(i^*)) = i^*\}$. Figure 1 illustrates this notation.

The swaps we consider in these cases are quite varied, but we always ensure we swap in $\mathrm{cent}(i)$ whenever some $i \in S \cap T$ with $\deg(i) > 0$ is swapped out. This way, we can always bound the reassignment cost of each client j by using either Lemma 6 or Lemma 7.

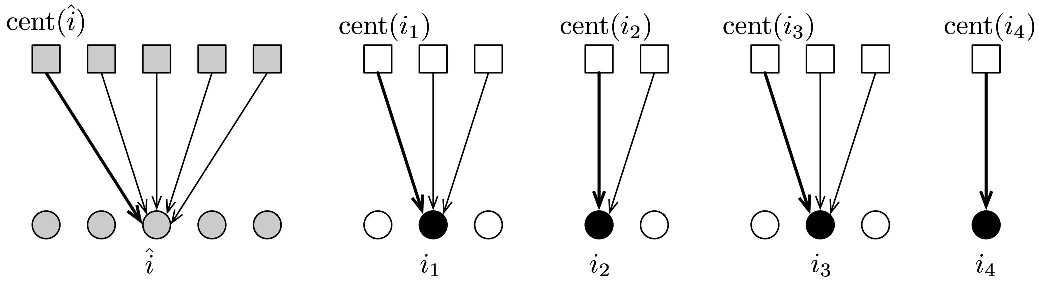

Figure 1 Illustration of a block T. The facilities on the top are in $T \cap O$ and the facilities on the bottom are in $T \cap S$. The directed edges depict ϕ, and the thick edges connect $\text{cent}(i)$ to i. The facilities coloured black lie in $\mathcal{B}$, the facilities coloured white lie in $\mathcal{R}$, and the facilities coloured grey could either lie in $\mathcal{B}$ or $\mathcal{R}$. Note that $\overline{B} = \{i_1, i_2, i_3, i_4\}$ and $\text{cent}(\overline{B}) = \{\text{cent}(i_1), \text{cent}(i_2), \text{cent}(i_3), \text{cent}(i_4)\}$. The layout of the figure is suggestive of how the block was constructed by adding "good" groups to the initial bad group in the procedure of generating blocks. The details of this procedure can be found in the full version of our work.

3.1 Case $|T^*_R| \leq t^2, |T^*_B| \leq t$

In this case, we simply swap out all of $T_R \cup T_B$ and swap in all of $T^*_R \cup T^*_B$. Because $R \cup B$ is a locally optimum solution and because this swaps at most t^2 facilities of each colour, we have:

$$0 \leq \text{cost}(S \cup (T^*_R \cup T^*_B) - (T_R \cup T_B)) - \text{cost}(S)\,.$$

Of course, after the swap each client will move to its nearest open facility. As is typical in local search analysis, we explicitly describe a (possibly suboptimal) reassignment of clients to facilities to upper bound this cost change.

Each $j \in N^*(T^*_R \cup T^*_B)$ is moved from s_j to o_j which incurs an assignment cost change of exactly $c^*_j - c_j$. Each $j \in N(T_R \cup T_B) - N^*(T^*_R \cup T^*_B)$ is moved to $\phi(o_j)$. Note that $\phi(o_j) \notin T$ so it remains open after the swap. By Lemma 6, the assignment cost change is bounded by $2c^*_j$. Every other client j that has not already been reassigned remains at s_j and incurs no assignment cost change. Thus,

$$0 \leq \sum_{j \in N^*(T^*_R \cup T^*_B)} (c^*_j - c_j) + \sum_{j \in N(T_R \cup T_B)} 2c^*_j$$

which is even better than what we are required to show for Theorem 8.

We note that the analysis Section 3.4 could be extended to subsume this analysis (with a worse constant), but we have included it here anyway to provide a gentle introduction to some of the simpler aspects of our approach.

3.2 Case $|T^*_R| \geq t^2 + 1, |T^*_B| \geq t + 1$

We start by briefly discussing some challenges in this case. In the worst case, all of the $i_b \in T_B$ have $\deg(i)$ being very large. The issue here is that we need to swap in each $i^*_b \in T^*_B$ in order to generate terms of the form $c^*_j - c_j$ for j with $o_j = i^*_b$. But this requires us to swap out some i_b. Since we cannot swap in all of $\phi^{-1}(i_b)$, we resort to only swapping in $\text{cent}(i_b)$.

Any client j with s_j being closed and $o_j \in \phi^{-1}(i_b) - \text{cent}(i_b)$ cannot be reassigned to $\phi(o_j)$, so we send it to $\text{cent}(\phi(o_j))$ and use Lemma 7 to bound the reassignment cost. This leaves a term of the form $+c_j$, so we have to consider additional swaps involving $-c_j$ in their bound to cancel this out. These additional swaps cause us to lose a factor of roughly 5 instead of 3.

Another smaller challenge is that we do not want to swap out the leader $\hat{i} \in T \cap S$ for a variety of technical reasons. However, since $|T^*_R|$ and $|T^*_B|$ are both big, this is not much of a problem. When we swap in some $i^* \in T \cap O$, we will just swap out a randomly chosen facility in $T \cap S - \hat{i}$ of the same colour. The probability any particular facility is swapped in this way is very small. Ultimately, each facility in $T \cap S - \hat{i}$ will be swapped out at most $2 + O(1/t)$ times in expectation.

To be precise, we partition the set of clients in $N(T_R \cup T_B)$ into two groups:

$$C_{bad} := N(\overline{B}) \cap N^*(T^*_R - \mathrm{cent}(\overline{B})) \quad \text{and} \quad C_{ok} := N(T_R \cup T_B - \hat{i}) - C_{bad}.$$

We have omitted $N(\hat{i})$ from C_{ok} because we will not close $\hat{i}$.

The first group is dubbed *bad* because there may be a swap where both s_j and $\phi(o_j)$ are closed yet o_j is not opened so we can only use Lemma 7 to bound their reassignment cost. In fact, some clients $j \in C_{good}$ may also be involved in such a swap, but we are able to use an averaging argument for these clients to show that the resulting $+s_j$ term from using Lemma 7 appears with negligible weight and does not need to be cancelled.

We consider the following two types of swaps to generate our initial inequality.

- For each $i^*_b \in T^*_B$, choose a random $i_b \in T_B - \hat{i}$. If $i_b \notin \overline{B}$ (i.e. $\deg(i_b) = 0$) then simply swap out i_b and swap in i^*_b. If $i_b \in \overline{B}$ then swap out i_b and a random $i_r \in T_R - \hat{i}$ and swap in i^*_b and $\mathrm{cent}(i_b)$.
- For each $i^*_r \in T^*_R - \mathrm{cent}(\overline{B})$, swap in i^*_r and swap out a randomly chosen $i_r \in T_R - \hat{i}$.

By choosing facilities at "random", we mean uniformly at random from the given set and this should be done independently for each invokation of the swap.

▶ **Lemma 9.**

$$0 \leq \sum_{j \in N^*(T^*_B \cup T^*_R)} \left(\frac{t+1}{t} \cdot c^*_j - c_j\right) + \sum_{j \in C_{ok}} \left[\left(2 + \frac{5}{t}\right) c^*_j + \frac{1}{t} c_j\right] + \frac{t+1}{t} \sum_{j \in C_{bad}} (3c^*_j + c_j)\,.$$

Proof. For brevity, we will let $\beta_R = \frac{|T_R|}{|T_R - \hat{i}|}$ and $\beta_B = \frac{|T_B|}{|T_B - \hat{i}|}$. Note that $\beta_R, \beta_B \leq \frac{t+1}{t}$ and that either $\beta_R = 1$ or $\beta_B = 1$.

First consider a swap of the first type that swaps in $\{i^*_b, \mathrm{cent}(i_b)\}$ and swaps out $\{i_b, i_r\}$ for some i_b with $\deg(i_b) > 0$. Because $R \cup B$ is a local optimum the cost of the solution does not decrease after performing this swap. We provide an upper bound on the reassignment cost.

Each $j \in N^*(\{i^*_b, \mathrm{cent}(i_b)\})$ is reassigned from s_j to o_j and incurs an assignment cost change of $c^*_j - c_j$. Every client $j \in N(\{i_b, i_r\})$ that has not yet been reassigned is first moved to $\phi(o_j)$. If this $\phi(o_j)$ remains open, assign j to it. By Lemma 6, the assignment cost for j increases by at most $2c^*_j$. If $\phi(o_j)$ is not open then $\phi(o_j) = i_b$ (because $\deg(i_r) = 0$) so we instead move j to $\mathrm{cent}(\phi(o_j)) = \mathrm{cent}(i_b)$. Lemma 7 shows the assignment cost increases by at most $3c^*_j + c_j$. This can only happen if $s_j \in \{i_r, i_b\}$ and $\phi(o_j) = i_b$.

Combining these observations and using slight overestimates, we see

$$0 \leq \sum_{j \in N^*(\{i^*_b, \mathrm{cent}(i_b)\})} (c^*_j - c_j) + \sum_{\substack{j \in N(\{i_b, i_r\}) \\ \phi(o_j) \neq i_b}} 2c^*_j + \sum_{\substack{j \in N(\{i_b, i_r\}) \\ \phi(o_j) = i_b}} (3c^*_j + c_j). \tag{1}$$

Now, if the random choice for i_b in the swap has $\deg(i_b) = 0$, then swapping $\{i_b\}$ out and $\{i^*_b\}$ in generates an even simpler inequality:

$$0 \leq \sum_{j \in N^*(i^*_b)} (c^*_j - c_j) + \sum_{j \in N(i_b)} 2c^*_j. \tag{2}$$

To see this, just reassign each $j \in N^*(i^*_b)$ from s_j to o_j and reassign the remaining $j \in N(i_b)$ from s_j to $\phi(o_j)$ (which remains open because $\deg(i_b) = 0$) and use Lemma 6.

Consider the expected inequality that is generated for this fixed i^*_b. We start with some useful facts that follow straight from the definitions and the swap we just performed.

- Any $j \in N^*(\mathrm{cent}(\overline{B}))$ has o_j being opened with probability $\frac{1}{|T_B - \hat{i}|}$.
- Any $j \in C_{bad}$ has s_j being closed with probability $\frac{1}{|T_B - \hat{i}|}$.
- Any $j \in C_{ok} - N(T_R)$ has s_j being closed with probability $\frac{1}{|T_B - \hat{i}|}$. When this happens, if o_j is not opened then $\phi(o_j)$ must be open. That is, $s_j \in C_{ok}$ means $o_j \in T^*_B \cup \mathrm{cent}(\overline{B})$.

  Furthermore, if $o_j \in T^*_B$ then $\phi(o_j) = \hat{i}$ (by the structure of block T) which remains open. If $o_j \in \mathrm{cent}(\overline{B})$ then either $\phi(o_j)$ was not closed, or else $\mathrm{cent}(\phi(o_j)) = o_j$ was opened.
- Any $j \in C_{ok} \cap N(T_R)$ has s_j being closed with probability $\frac{|\overline{B}|}{|T_B - \hat{i}|} \cdot \frac{1}{|T_R - \hat{i}|}$. If o_j and $\phi(o_j)$ are closed, then we move j to $\mathrm{cent}(\phi(o_j))$. However, this can only happen with probability $\frac{1}{|T_B - \hat{i}|} \cdot \frac{1}{|T_R - \hat{i}|}$ since it must be that $\phi(o_j)$ is the blue facility that was randomly chosen to be closed.

Averaging (1) over all random choices and using some slight overestimates we see

$$\begin{aligned} 0 \quad \leq \quad & \sum_{j \in N^*(i^*_b)} (c^*_j - c_j) + \frac{1}{|T_B - \hat{i}|} \cdot \sum_{j \in N^*(\mathrm{cent}(\overline{B}))} (c^*_j - c_j) \\ & + \frac{1}{|T_B - \hat{i}|} \left[\sum_{j \in C_{bad}} (3c^*_j + c_j) + \sum_{j \in C_{ok} - N(T_R)} 2c^*_j \right] \\ & + \frac{1}{|T_B - \hat{i}|} \cdot \frac{1}{|T_R - \hat{i}|} \sum_{j \in C_{ok} \cap N(T_R)} (|\overline{B}| 2c^*_j + 3c^*_j + c_j). \end{aligned}$$

Summing over all $i^*_b \in T^*_B$ (i.e. over all swaps of the first type) shows

$$\begin{aligned} 0 \quad \leq \quad & \sum_{j \in N^*(T^*_B)} (c^*_j - c_j) + \beta_B \cdot \sum_{j \in N^*(\mathrm{cent}(\overline{B}))} (c^*_j - c_j) \\ & + \beta_B \cdot \left[\sum_{j \in C_{bad}} (3c^*_j + c_j) + \sum_{j \in C_{ok} - N(T_R)} 2c^*_j \right] \\ & + \frac{\beta_B}{|T_R - \hat{i}|} \cdot \sum_{j \in C_{ok} \cap N(T_R)} ((2|\overline{B}| + 3)c^*_j + c_j). \end{aligned} \tag{3}$$

Next, consider the second type of swap that swaps in some $i^*_r \in T^*_R - \mathrm{cent}(\overline{B})$ and swaps out some randomly chosen $i_r \in T_R - \hat{i}$. Over all such swaps, the expected number of times each $i_r \in T_R - \hat{i}$ is swapped out is $\frac{|T^*_R| - |\overline{B}|}{|T_R - \hat{i}|} = \beta_R - \frac{|\overline{B}|}{|T_R - \hat{i}|}$. In each such swap, we reassign $j \in N^*(i^*_r)$ from s_j to o_j and every other $j \in N(i_r)$ from j to $\phi(o_j)$ which is still open because $\deg(i_r) = 0$. Thus,

$$0 \leq \sum_{j \in N^*(T^*_R - \mathrm{cent}(\overline{B}))} (c^*_j - c_j) + \left(\beta_R - \frac{|\overline{B}|}{|T_R - \hat{i}|} \right) \cdot \sum_{j \in C_{ok} \cap N(T_R)} 2c^*_j \,.$$

Scaling this bound by β_B, adding it to (3), and recalling $|T_R| \geq t^2$ shows

$$\begin{aligned} 0 \quad \leq \quad & \sum_{j \in N^*(T^*_B)} (c^*_j - c_j) + \beta_B \cdot \sum_{j \in N^*(T^*_R)} (c^*_j - c_j) \\ & + \beta_B \cdot \left[\sum_{j \in C_{bad}} (3c^*_j + c_j) + \sum_{j \in C_{ok} - N(T_R)} 2c^*_j + \right] \\ & + \beta_B \cdot \sum_{j \in C_{ok} \cap N(T_R)} \left[\left(2\beta_R + \frac{3}{t^2} \right) \cdot c^*_j + \frac{1}{t^2} \cdot c_j \right]. \end{aligned}$$

Recall that $\beta_B, \beta_R \leq \frac{t+1}{t}$ and also $\beta_B \cdot \beta_R \leq \frac{t+1}{t}$ to complete the proof of Lemma 9. ◀

Our next step is to cancel terms of the form $+c_j$ in the bound from Lemma 9 for $j \in C_{bad}$. To do this, we again perform the second type of swap for each $i \in T^*_R - \mathrm{cent}(\overline{B})$ but reassign clients a bit differently in the analysis.

▶ **Lemma 10.**

$$0 \leq \sum_{j \in C_{bad}} (c^*_j - c_j) + \frac{t+1}{t} \cdot \sum_{j \in C_{ok} \cap N(T_R)} 2c^*_j \,.$$

Proof. For each $i^*_r \in T^*_R - \mathrm{cent}(\overline{B})$, swap i^*_r in and a randomly chosen $i_r \in T_r - \hat{i}$. Rather than reassigning all $j \in N^*(i^*_r)$ to i^*_r, we only reassign those in $C_{bad} \cap N^*(i^*_r)$. Since $\deg(i_r) = 0$ then any other $j \in N(i_r)$ can be reassigned to $\phi(o_j)$ and which increases the cost by at most $2c^*_j$.

Summing over all i^*_r, observing that $C_{bad} \subseteq T^*_R - \mathrm{cent}(\overline{B})$, and also observing that each $j \in C_{ok}$ has s_j closed at most $\beta_R \leq \frac{t+1}{t}$ times in expectation, we derive the inequality stated in Lemma 10. ◀

Adding the bounds stated in Lemmas 9 and 10 shows that Theorem 8 holds in this case.

3.3 Case $|T^*_R| \geq t^2 + 1, |T^*_B| \leq t$

In this case, we start by swapping in all of T^*_B and swapping out all of T_B (including, perhaps, $\hat{i}$ if it is blue). In the same swap, we also swap in $\mathrm{cent}(T_B)$ and swap out a random subset of the appropriate number of facilities in $T_R - \hat{i}$. This is possible as $|T_R - \hat{i}| \geq t \geq |\mathrm{cent}(T_B)|$. By random subset, we mean among all subsets of $T_r - \hat{i}$ of the necessary size, choose one uniformly at random.

As with Section 3.2, we begin with a definition of bad clients that is specific to this case:

$$C_{bad} := N(T_B) \cap N^*(T^*_R - \mathrm{cent}(T_B))\,.$$

Clients $j \in C_{bad}$ have both s_j and $\phi(o_j)$ being closed yet o_j is not opened and we cannot make this negligible with an averaging argument.

▶ **Lemma 11.**

$$0 \quad \leq \quad \sum_{j \in N^*(T^*_B \cup \mathrm{cent}(T_B))} (c^*_j - c_j) + \frac{1}{t} \sum_{j \in N(T_R)} (3c^*_j + c_j) + \sum_{j \in C_{bad}} (3c^*_j + c_j)$$

Proof. After the swap, reassign every $j \in N^*(T^*_B \cup \mathrm{cent}(T_B))$ from s_j to o_j, for a cost change of $c^*_j - c_j$. Every other j that has s_j being closed is first reassigned to $\phi(o_j)$. If this is not open, then further move j to $\mathrm{cent}(o_j)$ which must be open because the only facilities $i \in T_R \cup T_B$ with $\deg(i) > 0$ that were closed lie in T_B and we opened $\mathrm{cent}(T_B)$.

If $j \in N(T_B) - C_{bad}$ then $o_j \in T^*_B \cup \mathrm{cent}(T_B)$ and we have already assigned j to o_j. If $j \in C_{bad}$ then we have moved j to $\mathrm{cent}(\phi(o_j))$ and the cost change is $3c^*_j + c_j$ by Lemma 7.

Finally, if $j \in N(T_R)$ then we either move j to $\phi(o_j)$ or to $\mathrm{cent}(\phi(o_j))$ if $\phi(o_j)$ is not open. The worst-case bound on the reassignment cost is $3c^*_j + c_j$ by Lemmas 6 and 7. However, note that $s_j \in T_R$ is closed with probability at most $1/t$ because we closed a random subset of $T_r - \hat{i}$ of size at most t and $|T_r - \hat{i}| \geq t^2$. ◀

We still need to swap in $T^*_R - \mathrm{cent}(T_B)$. For each such facility i^*_r, swap in i^*_r and swap out a randomly chosen $i_r \in T_R - \hat{i}$. The analysis of these swaps is nearly identical to analysis of the second type of swaps in Section 3.2, so we omit it and merely summarize what we get by combining the resulting inequalities with the inequality from Lemma 11.

▶ **Lemma 12.**

$$0 \leq \sum_{j \in N^*(T^*_R \cup T^*_B)} (c^*_j - c_j) + \sum_{j \in N(T_R)} \left(\frac{t^2+1}{t^2} \cdot 2c^*_j + \frac{1}{t} \cdot (3c^*_j + c_j) \right) + \sum_{j \in C_{bad}} (3c^*_j + c_j) \,.$$

We cancel the $+c_j$ terms for $j \in C_{bad}$ with one further collection of swaps. For each $i^*_r \in T^*_R - \mathrm{cent}(T_B)$ we swap in i^*_r and a randomly chosen $i_r \in T_R - \hat{i}$. The following lemma summarizes a bound we can obtain from these swaps. It is proven in essentially the same way as Lemma 10.

▶ **Lemma 13.**

$$0 \leq \sum_{j \in C_{bad}} (c^*_j - c_j) + \frac{t^2+1}{t^2} \cdot \sum_{j \in N(T_R)} 2c^*_j \,.$$

Adding this to the bound from Lemma 12 and recalling $C_{bad} \subseteq N(T_B)$ shows

$$0 \leq \sum_{j \in N^*(T^*_R \cup T^*_B)} (c^*_j - c_j) + \sum_{j \in N(T_R \cup T_B)} \left(\frac{t^2+3t+1}{t^2} \cdot 4c^*_j + \frac{1}{t} \cdot c_j \right) \,.$$

3.4 Case $|T^*_R| \leq t^2$, $|T^*_B| \geq t+1$

Because $\phi^{-1}(i) \subseteq T^*_R$ and $\deg(i) > 0$ for each $i \in \overline{B}$, then $|\overline{B}| \leq t^2$ as well. We will swap all of T^*_R for all of T_R, but we will also swap some blue facilities at the same time. Let $B' = \overline{B}$ and let $\overline{B}'$ be an arbitrary subset of T^*_B of size $|\overline{B}|$.

If $\hat{i} \notin T_R \cup B'$ then add $\hat{i}$ to B'. If $\mathrm{cent}(\hat{i}) \notin T^*_R \cup \overline{B}'$ then add $\mathrm{cent}(\hat{i})$ to $\overline{B}'$. At this point, $\left| |\overline{B}'| - |B'| \right| \leq 1$ Add an arbitrary $i^*_b \in T^*_B - \overline{B}'$ to $\overline{B}'$ or $i_b \in T_B - B'$ to B' to ensure $|\overline{B}'| = |B'|$.

We begin by swapping out $T_R \cup B'$ and swapping in $T^*_R \cup \overline{B}'$. The following list summarizes the important properties of this selection, the first point emphasizes that this swap will not improve the objective function since S is a locally optimum solution for the p-swap heuristic where $p = t^2 + 1$.

- $|B'| = |\overline{B}'| \leq t^2 + 1$ and $|T^*_R| \leq t^2$.
- T^*_R was swapped in and T_R was swapped out.
- For each $i \in T_R \cup T_B$ with $\deg(i) > 0$, i was swapped out and $\mathrm{cent}(i)$ was swapped in.

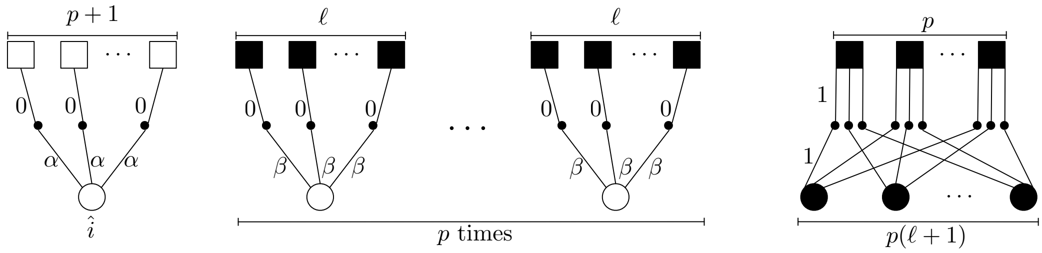

Figure 2 Illustration of the bad locality gap. Blue facilities are depicted with black and red facilities are depicted with white. The top facilities are the global optimum and the bottom are the local optimum (all of $\mathcal{R}$ and $\mathcal{B}$ is depicted in the picture). Each client is represented by a small black dot. There are $p^2 \cdot (\ell+1)$ clients in the right group, one for each pair of local and global optimum facilities in the group. The metric is the shortest path metric of the presented graph, if two locations are not connected in the picture then their distance is a very large value. Every edge in the right-most group with $p^2(\ell+1)$ clients has length 1. Recall $\beta = 2p$ and $\alpha = (\ell - p)2p$.

The following describes precisely the clients j that will be moved to $\mathrm{cent}(\phi(o_j))$ in our analysis.

$$C_{bad} := [N(T_R \cup B') - N^*(T_R^* \cup \overline{B}')] \cap \{j : \phi(o_j) \in T_R \cup B'\}\,.$$

The following bound is generated from swapping out $T_R \cup B'$ and swapping in $T_R^* \cup \overline{B}'$ and follows from the same arguments we have been using throughout the paper.

▶ **Lemma 14.**

$$0 \leq \sum_{j \in N^*(T_R^* \cup \overline{B}')} (c_j^* - c_j) + \sum_{j \in N(T_R \cup B') - C_{bad}} 2c_j^* + \sum_{j \in C_{bad}} (3c_j^* + c_j)\,.$$

Next, let $\kappa_B : (T_B^* - \overline{B}') \to (T_B - B')$ be an arbitrary bijection of the remaining blue facilities that were not swapped. For every $i_b^* \in T_B^* - \overline{B}'$, consider the effect of swapping in i_b^* and swapping out $\kappa_B(i_b^*)$. Note that every facility i_b swapped out in this way has $\deg(i_b) = 0$. So we can derive two possible inequalities from such swaps.

$$0 \leq \sum_{j \in N^*(i_b^*)} (c_j^* - c_j) + \sum_{j \in N(\kappa_B(i_b^*))} 2c_j^* \quad \text{and} \quad 0 \leq \sum_{j \in N^*(i_b^*) \cap C_{bad}} (c_j^* - c_j) + \sum_{j \in N(\kappa_B(i_b^*))} 2c_j^*. \quad (4)$$

The second inequality follows from only reassigning clients $j \in N^*(i_b^*) \cap C_{bad}$ from s_j to o_j.

Adding the bound in Lemma 14 to the sum of both inequalities over all $i_b^* \in T_B^* - \overline{B}'$ and noting that $\kappa_B(T_B^* - \overline{B}) \cap (T_R \cup B') = \emptyset$, we see

$$0 \leq \sum_{j \in N^*(T_R^* \cup T_B^*)} (c_j^* - c_j) + \sum_{j \in N(T_R \cup T_B)} 4c_j^*\,.$$

4 Locality Gaps

Here we prove Theorem 3. Let p, ℓ be integers satisfying $p \geq 1$ and $\ell \geq 2p$. Consider the instance with $k_r = p+1$ and $k_b = p(\ell+1)$ depicted in Figure 2. Here, $\beta = 2p$ and $\alpha = \beta \cdot (\ell - p)$.

The cost of the local optimum solution is $\alpha \cdot (p+1) + \beta \cdot p \cdot \ell + p^2(\ell+1)$ and the cost of the global optimum solution is simply $p^2(\ell+1)$. Through some careful simplification, we see the local optimum solution has cost at least $5 + \frac{2}{p} - \frac{10p}{\ell+1}$ times the global optimum solution.

To complete the proof of Theorem 3, we must verify that the presented local optimum solution indeed cannot be improved by swapping up to p facilities of each colour. The straightforward details appear in the full version of this paper.

5 Conclusion

We have demonstrated that a natural p-swap local search procedure for Budgeted Red-Blue Median is a $(5 + O(1/\sqrt{p}))$-approximation. This guarantees a better approximation ratio than the single-swap heuristic from [9], which we showed may find solutions whose cost is $(7 - \epsilon) \cdot OPT$ for arbitrarily small ϵ. Our analysis is essentially tight in that the p-swap heuristic may find solutions whose cost is $(5 + \frac{2}{p} - \epsilon) \cdot OPT$.

More generally, one can ask about the p-swap heuristic for the generalization where there are many different facility colours. If the number of colours is part of the input then any local search procedure that swaps only a constant number of facilities in total cannot provide good approximation guarantees [10]. However, if the number of different colours is bounded by a constant, then perhaps one can get better approximations through multiple-swap heuristics.

However, generalizing the approaches taken with Budgeted Red-Blue Median to this setting seems more difficult; one challenge is that it is not possible to get such nicely structured blocks. It would also be interesting to see what other special cases of Matroid Median admit good local-search based approximations. For example, the Mobile Facility Location problem studied in [2] is another special case of Matroid Median that admits a $(3 + \epsilon)$-approximation through local search.

Finally, the locality gap of the p-swap heuristic for k-Median is known to be $3 + \frac{2}{p}$ [3] and we have just shown it is at least $5 + \frac{2}{p}$ for Budgeted Red-Blue Median. Even if the multiple-swap heuristic for the generalization to a constant number of colours can provide a constant-factor approximation, this constant may be worse than the alternative 8-approximation obtained through Swamy's general Matroid Median approximation [15].

References

1. Ankit Aggarwal, Anand Louis, Manisha Bansal, Naveen Garg, Neelima Gupta, Shubham Gupta, and Surabhi Jain. A 3-approximation for facility location with uniform capacities. In *Proc. of IPCO*, pages 149–162, 2010.
2. Sara Ahmadian, Zachary Friggstad, and Chaitanya Swamy. Local-search based approximation algorithms for mobile facility location problems. In *Proc. of SODA*, pages 1607–1621, 2013.
3. Vijay Arya, Naveen Garg, Rohit Khandekar, Adam Meyerson, Kamesh Munagala, and Vinayaka Pandit. Local search heuristics for k-median and facility location problems. *SIAM J. Comput.*, 33(3):544–562, 2004.
4. Manisha Bansal, Naveen Garg, and Neelima Gupta. A 5-approximation for capacitated facility location. In *Proc. of ESA*, pages 133–144, 2012.
5. Jaroslaw Byrka, Thomas Pensyl, Bartosz Rybicki, Aravind Srinivasan, and Khoa Trinh. An improved approximation for k-median, and positive correlation in budgeted optimization. In *Proc. of SODA*, pages 737–756, 2015.
6. Moses Charikar and Sudipto Guha. Improved combinatorial algorithms for facility location problems. *SIAM J. Comput.*, 34(4):803–824, 2005.

7 Inge Li Gørtz and Viswanath Nagarajan. Locating depots for capacitated vehicle routing. In *Proc. of APPROX*, pages 230–241, 2011.
8 Anupam Gupta and Kanat Tangwongsan. Simpler analyses of local search algorithms for facility location. *CoRR*, abs/0809.2554, 2008.
9 MohammadTaghi Hajiaghayi, Rohit Khandekar, and Guy Kortsarz. Local search algorithms for the red-blue median problem. *Algorithmica*, 63(4):795–814, 2012.
10 Ravishankar Krishnaswamy, Amit Kumar, Viswanath Nagarajan, Yogish Sabharwal, and Barna Saha. The matroid median problem. In *Proc. of SODA*, pages 1117–1130, 2011.
11 Shi Li and Ola Svensson. Approximating k-median via pseudo-approximation. In *Proc. of STOC*, pages 901–910, 2013.
12 Mohammad Mahdian and Martin Pál. Universal facility location. In *Proc. of ESA*, pages 409–421, 2003.
13 Martin Pál, Éva Tardos, and Tom Wexler. Facility location with nonuniform hard capacities. In *Proc. of FOCS*, pages 329–338, 2001.
14 Zoya Svitkina and Éva Tardos. Facility location with hierarchical facility costs. *ACM Trans. Algorithms*, 6(2), 2010.
15 Chaitanya Swamy. Improved approximation algorithms for matroid and knapsack median problems and applications. In *Proc. of APPROX*, pages 403–418, 2014.