## Diplomarbeit

# Algorithmic Approaches to the String Barcoding Problem 

ausgeführt am<br>Institut für Computergraphik und Algorithmen<br>der Technischen Universität Wien

unter Anleitung von
Univ.-Prof. Dipl.-Ing. Dr. techn. Günther Raidl
durch
Philipp Neuner
Pfeilgasse 14/2/5
1080 Wien

## Abstract

This thesis deals with a heuristic approach based on Lagrangian relaxation to the string barcoding (SB) problem, a close cousin to the well-known combinatorial set cover (SC) problem. It has recently been proven to be NP-hard and has many real-world applications, particularly in the fields of medicine and biology.

Given a set of sequences over some alphabet, DNA for instance, we aim at finding a set of short sequences, so-called probes, such that we are able to identify an unknown sample sequence as one of the input sequences by determining which probes are subsequences of the sample, and which are not. The problem is twofold: the determination of all possible probes and the selection of a suitable subset of minimum cardinality.

The problem has been dealt with under various other names and has in this form been introduced by Rash and Gusfield in 2002. They proposed an exact approach based on integer linear programming and the use of suffix trees to generate a complete, nonredundant set of candidate probes.

We evaluated several approaches for the SB as well as the SC problem. One of the leading heuristics for the SC problem, based on Lagrangian relaxation, has been proposed by Caprara et al. in 1999. We adapted the algorithm to see if it works equally well when applied to the structurally very similar SB problem. Though the results we obtained are somewhat mixed, the heuristic shows its strength with very complex instances and delivers much better results compared to simpler heuristics.

## Contents

1 Introduction ..... 1
2 Preliminaries ..... 5
2.1 Linear Programming ..... 5
2.2 Lagrangian Relaxation ..... 6
2.3 Subgradient Optimization ..... 9
3 Formal Problem Definition ..... 11
3.1 String Barcoding ..... 11
3.2 Set Cover ..... 13
3.3 Minimum Test Collection ..... 13
3.4 Complexity and Inapproximability ..... 14
4 Probe Generation ..... 15
4.1 Suffix Trees ..... 15
4.2 Example ..... 16
4.3 Pitfalls ..... 18
5 Previous Work ..... 21
5.1 String Barcoding ..... 21
5.2 Set Cover Approaches ..... 25
6 A Lagrangian Heuristic for the String Barcoding Problem ..... 31
6.1 Outline ..... 31
6.2 Core Algorithm ..... 36
6.2.1 Lagrangian Multiplier Optimization ..... 36
6.2.2 Greedy Heuristic ..... 41
6.2.3 Probe Fixing ..... 42
6.2.4 Pricing ..... 43
6.2.5 Solution Refinement ..... 45
7 Implementation ..... 47
7.1 Workflow ..... 47
7.2 Data Formats ..... 49
7.3 Implementation Details ..... 50
7.4 Class overview ..... 51
8 Computational Results ..... 55
8.1 Instances ..... 55
8.2 Results ..... 56
8.3 Discussion ..... 57
9 Conclusions ..... 67
10 Bibliography ..... 69

## 1 Introduction

A common task in medical and biological applications is the identification of biological agents like bacteria, viruses, proteins, and the like. Basically, the problem is the same regardless of the exact kind of agent we are dealing with. To keep things simple we will consider only nucleotide sequences hereinafter. Nucleotides are the structural units of deoxyribonucleic and ribonucleic acids, commonly known as DNA and RNA, which carry the genetic code of all living organisms and viruses. They consist of a sugar group (ribose or deoxyribose), one or more phosphate groups, and one base. The four bases you may find in a DNA molecule are adenine (A), cytosine (C), guanine (G), and thymine (T). The same goes for RNA molecules with the only exception being that uracil ( U ) takes the place of thymine. Each base has a complementary counterpart to which it bonds exclusively. Accordingly, the bases can be divided into the complementary pairs A-T and C-G. The reaction that takes place when two complementary nucleotide sequences bond to each other is called hybridization.

Let us assume we are given a sample containing the DNA or RNA of some unknown virus and a list of viruses that we think may be contained in the sample. The most common techniques to match a genomic sample against a list of candidates rely on sequence alignments to find a perfect or near-perfect match and often incorporate heuristic elements for the sake of efficiency. A very widespread representative of this class of procedures is BLAST [29], which is short for basic local alignment search tool. These methods, however, can be very time consuming and require the sequence we want to identify to be available explicitly, i. e. we must perform a comprehensive sequence analysis to determine the exact succession of comprising nucleotides. Unfortunately, this is most often technically complex and expensive.

A similar problem arises when one strives to determine the gene expression of a cell, i. e. the genes that are active in the cell and the degree of activity. A gene is basically just a part of the DNA that encodes a particular protein. If a gene is active, it is transcribed
into messenger RNA (mRNA) molecules which are complementary the the respective gene. The mRNA molecules are transferred to the cell's ribosome where the protein is actually build. Thus, the kinds and amount of mRNA molecules found in a cell gives insight to its gene expression. When it comes to the identification of a particular mRNA molecule we face the exact same problem as stated above.

Instead of comparing the complete sample sequence with possible candidates, we may alternatively check the sample for the presence or absence of short nucleotide sequences, so-called oligonucleotides, which are typically between five and 50 nucleotides long. Modern microarray technology enables us to efficiently perform many such tests in parallel. Microarrays are small glass microscope slides, sometimes also silicon chips or nylon membranes, onto which many thousands of nucleotide sequences are attached at fixed locations. Before we apply our sample to the microarray, it has to be tagged


Figure 1.1: A DNA microarray with a fluorescent marker. This provides for a visible feedback on the hybridization process so that we can easily determine to which oligonucleotides on the array our sample hybridizes. Figure 1.1 exemplifies the outcome of this procedure, where the different colors are due to the different markers used. Special scanner devices are used for computer aided evaluation of the strength of the fluorescent reaction at each spot on the array.

From a more formal perspective the outcome of these tests provides some kind of binary signature, each binary digit of which tells us if a particular subsequence occurs as part of the sample sequence or not. The goal is to find a set of such subsequences, which we will call probes further on, so that the corresponding signature associated with each candidate sequence is unique over all sequences. The problem of finding a set of minimal cardinality that satisfies this condition has been introduced by Rash and Gusfield [32] and is referred to as the string barcoding (SB) problem. They provide a straight-forward integer linear programming formulation of the problem and propose a sophisticated method for generating a complete set of probes to choose from.

The same problem has been addressed by Borneman et al. [8]. They did, however,
factor out the problem of probe generation and assume a given set of candidate probes as part of the input. Moreover, the algorithm they propose for solving the problem is a heuristic based on Lagrangian relaxation. It should be noted, that the problem has been considered well before theses publications, in one form or another under the names probe selection, oligonucleotide fingerprinting, and the like. However, the vast majority of these publications puts the main focus on the biological and application-side aspects of the problem. This thesis focuses on the underlying combinatorial optimization problem and algorithmic approaches to solve it.

SB can be seen as a special case of both the well-known set covering (SC) problem and the minimum test collection (MTC) problem, which have been shown to be NP-hard combinatorial optimization problems. The question whether SB itself is NP-hard or not was left open in [32]. The authors did only state that the problem was NP-hard if an additional maximum length constraint was imposed on the probes. However, in the meantime it has been proven to be NP-hard [6, 25].

Scope and outline. Chapter 2 introduces some basic theoretical concepts like linear programming and Lagrangian relaxation, which are important for the understanding of later parts. Chapter 3 defines the SB problem and its relatives formally and presents complexity and inapproximability results. Chapter 4 describes in detail how complete probe sets can be generated using suffix trees as has been proposed in [32]. Chapter 5 summarizes previous work related to SB and some state-of-the-art SC approaches, which may be applied to SB instances as well. We decided to adapt one of the most cited and best performing heuristics for the SC problem, hoping that it would be equally successful in solving SB instances. A comprehensive description of the heuristic, including the adaptations we applied to accommodate for the peculiarities of the SB problem, and details on the implementation are given in Chapter 6. Computational results on selected real-world benchmark instances can be found in Chapter 8. Chapter 9 summarizes the results of this work.

1 Introduction

## 2 Preliminaries

The following sections introduce fundamental theoretical concepts, which are essential for the understanding of forthcoming parts. Section 2.1 defines linear programs and the notion of relaxation, Section 2.2 introduces the concept of Lagrangian relaxation, and Section 2.3 presents the subgradient optimization technique.

### 2.1 Linear Programming

This section defines linear programs and lists some of their properties. A comprehensive discourse on the topic can be found in [28]. Linear programs (LP) are optimization problems with a linear objective function that can be expressed as follows

$$
\begin{array}{ll}
\min _{x} & c^{T} x \\
\text { s. t. } & A x \geq b  \tag{2.2}\\
& x, b, c \in \mathbb{R}^{n}, A \in \mathbb{R}^{m \times n}
\end{array}
$$

Each linear inequality (2.2) defines a half-space of $\mathbb{R}^{n}$, the intersection of which forms a convex polytope. There are several efficient algorithms to solve problems of this kind to optimality. One of them is the simplex algorithm, which has been introduced by George Dantzig in 1947. The key idea of this method is to move along the polytope's facets from one vertex to the next. If we ensure that the objective function does not decrease when moving to an adjacent vertex, we will eventually reach the optimum. While very efficient in practice, the worst-case complexity is exponential. On the other hand, the more recent interior point methods like the predictor-corrector algorithm [27]
have polynomial worst-case complexity and are also in practice competitive with the simplex method.

When dealing with combinatorial optimization problems like those presented in this thesis, we concern ourselves with discrete entities like the probes in an instance of the SB problem. Many of these problems can be formulated as LPs as well, but in this case some or all decision variables are are required to be integer. Accordingly, these problems are called mixed integer linear programs (MIP), or integer linear programs (ILP) if all variables are required to be integer. In contrast to LPs, finding the optimal solution of a MIP is in general a NP-hard problem and therefore no polynomial-time algorithm exists for this task.

### 2.2 Lagrangian Relaxation

Given some ILP

$$
\begin{array}{ll}
\min & c^{T} x \\
\text { s. t. } & A x \geq b \\
& x \in \mathbb{Z}^{n} \tag{2.5}
\end{array}
$$

we may drop (relax) some of its constraints, thereby enlarging the feasible region of the problem. If we, for instance, relax the integrality constraints (2.5), we get what is called the LP relaxation (LPR) of an ILP. As every feasible solution to the ILP is also a feasible solution to its LPR, the optimal solution of the LPR is guaranteed to be better or exactly as good as the optimal solution of the ILP itself. Thus, the LPR if often used to derive a lower bound for a given ILP.

In practice, the LPR might nevertheless be impractical, either because its optimal solution is too resource consuming to compute, or because the computed solution provides an inferior lower bound to the underlying ILP. In these cases, there are other relaxations which can be used to compute a valid, and potentially better, lower bound, one of them being the Lagrangian relaxation (LGR). This method is based on splitting the
constraints into two disjoint sets:

$$
\begin{array}{ll}
\min & c^{T} x \\
\text { s. t. } & A x \geq b \\
& C x \geq d \\
& x \in \mathbb{Z}^{n} \tag{2.9}
\end{array}
$$

Matrix $A$ together with vector $b$ represent the constraints we want to keep, matrix $C$ in companion with vector $d$ those that make the problem hard to solve and that we want to relax. However, when performing Lagrangian relaxation, we do not just drop the constraints but move them into the objective function. We do this by associating a Lagrangian multiplier $0 \leq u_{i} \in \mathbb{R}$ with each relaxed constraint and adding a penalty term to the objective function:

$$
\begin{array}{ll}
\text { min } & \operatorname{LGR}(u):=c^{T} x+u^{T}(d-C x) \\
\text { s. t. } & A x \geq b \\
& x \in \mathbb{Z}^{n} \tag{2.12}
\end{array}
$$

That the optimal solution to this ILP constitutes a valid lower bound for the originating ILP stems from the fact that, given a valid solution $x$ of the ILP defined by (2.6) - (2.9) with objective value $v_{\text {ILP }}$, we can state two things:

- $x$ is also a feasible solution of (2.10) - (2.12), because only a subset of constraints has to be satisfied
- Let $v_{\text {LGR }}$ be the objective value of $x$ with respect to (2.10). Since $x$ satisfies constraint (2.8), $u^{T}(d-C x) \leq 0$ holds for any $0 \leq u \in \mathbb{R}^{m}$, which in turn implies that $v_{\text {LGR }} \leq v_{\text {ILP }}$

Given a set of multipliers, the LGR should be easy to solve to be useful. However, now we need to find a set of multipliers $u$ that yields the highest possible objective value.

## 2 Preliminaries

This problem is called the Lagrangian dual and is formally defined as

$$
\begin{equation*}
\max _{0 \leq u \in \mathbb{R}} \min \operatorname{LGR}(u) \tag{2.13}
\end{equation*}
$$

The optimal objective value $v_{\mathrm{LGR}}^{*}$ of the LG relaxation is guaranteed to be greater than or equal to the optimal objective value $v_{\text {LPR }}^{*}$ of the LP relaxation. This can be proven as follows:

$$
\begin{align*}
v_{\mathrm{LGR}}^{*} & =\max _{u \in \mathbb{R}^{m}} \min _{x \in \mathbb{Z}^{n}}\left\{c^{T} x+u^{T}(d-C x): A x \geq b\right\}  \tag{2.14}\\
& \geq \max _{u \in \mathbb{R}_{+}^{m}}^{\min _{x \in \mathbb{R}^{n}}\left\{c^{T} x+u^{T}(d-C x): A x \geq b\right\}}  \tag{2.15}\\
& =\max _{u \in \mathbb{R}_{+}^{m}} \max _{y \in \mathbb{R}_{+}^{m}}\left\{b^{T} y+u^{T} d: A^{T} y=c-u^{T} C\right\}  \tag{2.16}\\
& =\max _{u, y \in \mathbb{R}_{+}^{m}}\left\{b^{T} y+d^{T} u: A^{T} y=c-C^{T} u\right\}  \tag{2.17}\\
& =\min _{z \in \mathbb{R}^{n}}\left\{c^{T} z: A x \geq b, C x \geq d\right\}  \tag{2.18}\\
& =v_{\mathrm{LPR}}^{*} \tag{2.19}
\end{align*}
$$

where $\mathbb{R}_{+}$denotes the set $\{x \in \mathbb{R}: x \geq 0\}$. In (2.15) we relax the integrality constraints (2.12), which results in a solution with a possibly lower objective value. Equalities (2.16) and (2.18) are valid due to the the strong duality theorem, which is beyond the scope of this work - for an introduction to the topic see Nemhauser and Wolsey [28].

If the relaxation of the integrality constraints does not change the objective value of the solution, the Lagrangian relaxation is said to exhibit the integrality property. In this case, we can infer from (2.14) through (2.18) that we cannot derive a lower bound superior to that provided by the LP relaxation.

There are several techniques to solve the Lagrangian dual. One of the most prominent ones is called subgradient optimization and will be described in further detail in Section 2.3. A different and more recent approach is the volume algorithm introduced by Barahona and Anbil [2].

### 2.3 Subgradient Optimization

Subgradient optimization is an iterative procedure to find Lagrangian multipliers that yield a high lower bound. It is motivated by the fact that the Lagrangian dual is a piecewise linear, convex problem, from which we can infer that the optimum is located at a non-differentiable point. The technique was first introduced by Held and Karp [19, 20], who applied it to the traveling salesman problem.

In each iteration, we try to update the current set of multipliers $u$ so as to reach a higher lower bound than before. To decide on the direction of our move, we use a subgradient. Given some function $f: \mathbb{R}^{n} \mapsto \mathbb{R}$, a vector $v \in \mathbb{R}^{n}$ is called a subgradient at point $x_{0} \in \mathbb{R}^{n}$ if

$$
\begin{equation*}
f(x)-f\left(x_{0}\right) \geq v \cdot\left(x-x_{0}\right) \tag{2.20}
\end{equation*}
$$

holds for any $x \in \mathbb{R}^{n}$. Note that a subgradient is defined at non-differentiable points, too.

Consider the generic LGR given in (2.10) through (2.12). Let $u$ be a given multiplier vector and $x^{*}$ the corresponding optimal solution. Then, we can define the following subgradient

$$
\begin{equation*}
s(u):=b-A x^{*} \tag{2.21}
\end{equation*}
$$

The basic procedure is depicted in Algorithm 1. The step size parameter $\lambda$ controls how far we move along the subgradient and it will usually be adapted throughout the procedure to avoid moving past the optimum. Moreover, we have to decide on initial values of the Lagrangian multipliers, define the conditions that need to be satisfied to terminate the process, and provide a heuristic that produces feasible solutions to derive valid upper bounds.

```
Algorithm 1 Subgradient optimization
    Initialize step size parameter \(\lambda \in(0,2]\)
    Initialize the set of multipliers \(u\)
    Find some upper bound \(U\)
    repeat
        \(L \leftarrow \min \operatorname{LGR}(u)\)
        Let \(x^{*}\) be the corresponding solution
        \(T \leftarrow \lambda(U-L) /\|s(u)\|^{2}\)
        \(u_{i} \leftarrow \max \left\{0, u_{i}+T \cdot s_{i}(u)\right\}\)
    until termination criterion met
```


## 3 Formal Problem Definition

This chapter defines SB and related optimization problems more formally. Section 3.1 provides an ILP formulation of SB and proposes possible extensions, which may arise in real-world applications. Sections 3.2 and 3.3 define the well-known set covering (SC) and minimum test collection (MTC) problems, and state how they relate to SB. Section 3.4 deals with the complexity aspects of SB.

### 3.1 String Barcoding

Given sequences $s_{1}, \ldots, s_{n}$ over some finite alphabet $\Sigma$ and a set of probes $p_{1}, \ldots, p_{m}$, each probe being a subsequence of at least one sequence $s_{i}$. A probe $p_{k}$ is said to distinguish sequences $s_{i}$ and $s_{j}$ if and only if $p_{k}$ is a subsequence of $s_{i}$ or $s_{j}$, but not both. Accordingly, we define

$$
\delta_{i, j}(k)= \begin{cases}1 & \text { if } p_{k} \text { distinguishes } s_{i} \text { and } s_{j} \\ 0 & \text { otherwise }\end{cases}
$$

and denote the set of probes distinguishing sequences $s_{i}$ and $s_{j}$ by $\Delta_{i, j}$.

We aim at finding a minimum cardinality subset of probes such that every pair of sequences is distinguished by at least $r$ probes in this subset. This can be formulated as the following ILP:

$$
\begin{array}{ll}
\text { min } & \sum_{i=0}^{m} x_{i} \\
\text { s. t. } & \sum_{i=0}^{m} \delta_{i, j}(k) \cdot x_{k} \geq r \\
& x_{i} \in\{0,1\} \tag{3.3}
\end{array} \quad \forall(i, j): 1 \leq i<j \leq n
$$

The value of decision variable $x_{i}$ determines if the corresponding probe $p_{i}$ is to be part of the solution $\left(x_{i}=1\right)$ or not $\left(x_{i}=0\right)$. The redundancy parameter $r$ defines by how many probes each sequence pair has to be distinguished. In real-world applications a higher redundancy may be suitable, because we cannot expect the outcome of each single subsequence test to be $100 \%$ accurate. Setting $r>1$ compensates, at least to some degree, for false negatives or positives, thereby enlarging our safety margin and making the resulting probe set more robust. One very fortunate effect of requiring a higher redundancy is that the ILP might become easier to solve. This is because, as $r$ increases, the number of possibilities to satisfy constraints (3.2) usually decreases.

Another possibility to make the generated probe set more robust, particularly with regard to small mutations, is to require the minimum pairwise distance between probes in the solution to be above a particular threshold. One possible metric for measuring the distance between two strings or probes is the edit or Levenshtein distance. It defines the distance between two strings $s$ and $t$ as the minimum number of unit operations character insertion, deletion, or substitution - required to transform $s$ into $t$ or vice versa. Note that it is a symmetric metric and that, depending on the particular application, the unit operations may be weighed differently. To enforce a minimum pairwise edit distance, we add one additional constraint to ILP (3.1) - (3.3) for each pair of candidate probes $p_{i}$ and $p_{j}$ whose edit distance from one another is below the required threshold:

$$
\begin{equation*}
p_{i}+p_{j} \leq 1 \tag{3.4}
\end{equation*}
$$

Thus we assure that at most one probe of each such pair will show up in the solution.

### 3.2 Set Cover

Set cover (SC) is one of Karp's 21 NP-complete problems [22]. The problem is defined by a set of items $S=\left\{s_{1}, \ldots, s_{n}\right\}$ and a collection $\mathcal{C}=\left\{C_{1}, \ldots, C_{m}: C_{i} \subseteq S\right\}$ of subsets of $S$ with each subset $C_{i}$ having associated cost $0<c_{i} \in \mathbb{R}$. The goal is to determine a collection $\mathcal{C}^{*} \subseteq \mathcal{C}$ of minimum total cost that covers all items in $S$ :

$$
\begin{array}{ll}
\min _{\mathcal{C}^{\prime} \subseteq \mathcal{C}} & \sum_{C_{i} \in \mathcal{C}^{\prime}} c_{i} \\
\text { s. t. } & \bigcup_{C \in \mathcal{C}^{\prime}} C=S
\end{array}
$$

We restrict our discussion to the unicost case where $c_{i}=1$ for all $1 \leq i \leq m$. It is easy to see that SB is a special case of SC and can therefore be stated in these terms. In this case, set $S$ comprises all sequence pairs and collection $\mathcal{C}$ the sets of sequence pairs distinguished by the candidate probes. This relationship enables us to apply algorithms originally designed to solve SC instances to SB instances as well, which will be demonstrated in Chapter 6 .

### 3.3 Minimum Test Collection

The minimum test collection (MTC) problem is very similar to the SB problem and can even be regarded as a higher-level formulation of the latter. It is defined by a set of items $D$ and a collection $\mathcal{T}=\left\{T_{1}, \ldots, T_{n}: T_{i} \subseteq D\right\}$ of subsets of $D$, called tests. We aim at finding a collection $\mathcal{T}^{*}$ of minimum size such that for each pair of elements $\left(d_{i}, d_{j}\right) \in D \times D$ there is at least one test $T_{i} \in \mathcal{T}^{*}$ with $\left|\left\{d_{i}, d_{j}\right\} \cap T_{i}\right|=1$. Garey and Johnson [17] have shown that the problem is NP-complete by reduction from 3 -dimensional matching. As is the case with SC, any SB instance can easily be transformed into an MTC instance in a straight-forward manner.

### 3.4 Complexity and Inapproximability

Rash and Gusfield [32] state that the max-length string barcoding problem, which adds the additional constraint that no candidate probe exceeds a given maximum length, is NP-complete. However, they did not answer the question for the unconstrained case. In the meantime, Lancia and Rizzi [25] and Berman et al. [6] have proven this. Their findings rely on previous work that has shown that SC cannot be approximated in polynomial time within a factor of $(1-\varepsilon) \log n$, with $\varepsilon>0$ and $n$ being the number of items to be covered. More precisely, this means that no polynomial approximation algorithm exists that produces a solution of size $z \leq z^{*} \cdot(1-\varepsilon) \log n$, where $z^{*}$ is the size of the optimal solution. Lancia and Rizzi [25] provide polynomial time transformations, which preserve this property, from SB to a special subclass of MTC and from there to SC.

Borneman et al. [8] claim, though without proof, that NP-completeness can be shown by reduction from vertex cover.

## 4 Probe Generation

As shown in Section 3.4, SB is a NP-hard optimization problem for which no polynomial time algorithm is likely to exist (unless $\mathrm{P}=\mathrm{NP}$ ). For this reason, it is strongly advisable to keep the number of candidate probes as low as possible. On the other hand, it is important not to leave out any essential probe as this would change the solution space and we may end up with a suboptimal solution. A rather naive approach would be to enumerate all possible substrings of all given input sequences. This would, however, result in a large number of redundant probes that distinguish the same set of sequence pairs. Fortunately, there is a much more sophisticated way to generate a complete set of probes in a very efficient manner by means of a suffix tree. Section 4.1 will outline this method and Section 4.2 presents a clarifying example.

### 4.1 Suffix Trees

A suffix tree is a data structure which encodes all suffixes of a number of strings; Gusfield [18] provides an in-depth coverage of the topic. The suffix tree for a number of strings $S=\left\{s_{1}, \ldots, s_{n}\right\}$ has the following properties:

- It is a directed tree with a designated root node
- Each edge is labeled with a non-empty string that must be a substring of at least one $s_{i}$
- Each leaf node is associated with one input string $s_{i}$. The concatenation of the edge labels on the path from the root node to this node is a suffix of $s_{i}$
- Each internal node has at least 2 children


## 4 Probe Generation

- The labels of all outgoing edges of a node start with different characters

As any possible substring is the prefix of some suffix, it can be found on a unique path, starting at the root node and ending either at some node or inside an edge (label). According to the properties of a suffix tree, the substring occurs in every string that is associated with one of the leaf nodes that are located below this point in the tree.

The defining property of a probe $p_{k}$ is the set of strings $S_{k} \subseteq S$ in which it occurs as substring and, derived from that, the set of sequence pairs it distinguishes. Thus, two probes $p_{i}$ and $p_{j}$ can be considered equivalent if either $S_{i}=S_{j}$ or $S_{i}=S \backslash S_{j}$. As a consequence, each suffix tree node $v$ represents a whole set of substrings that are equivalent in this regard. This set comprises all substrings that we get by concatenating the labels of the edges on the path from the root node to the parent node of $v$ plus a non-empty suffix of the label of the edge from the parent node to $v$.

To exploit this fact, we assign a unique binary signature $b_{0} b_{1} \ldots b_{n-1}, b_{i} \in\{0,1\}$, to each node, with $b_{i}=1$ if the corresponding set of strings occurs in input string $s_{i}$, and $b_{i}=0$ otherwise. Note that two nodes with different signatures might still be equivalent to each other, namely if one signature is the binary complement of the other. To get rid of this special case, we require the first binary digit of every signature to be 1 and flip signatures starting with 0 . Apart from this, we can ignore nodes with a signature $11 \ldots 1$ because probes that occur in every sequence do not distinguish any sequence pairs and are of no use.

The probe generation procedure can be summarized as follows: we build the suffix tree for the given set of sequences; this can be done in linear time. Then, we traverse the tree bottom-up and determine the signature of each node as defined above. For each distinct signature we choose as many probes as necessary. Compared to more naive approaches like the complete enumeration of all distinct substrings this method yields $75 \%$ less probes on average, in some cases the reduction exceeds $90 \%$.

### 4.2 Example

To make things clearer, we will present a simple example. Consider the strings $s_{0}=$ CAGTGC, $s_{1}=$ CAGTTC, and $s_{2}=$ CATGGA. The corresponding suffix tree is depicted in

Fig. 4.1 with the black node being the root. Equivalent nodes, which distinguish the same pairs of strings, are labeled with identical Greek letters. Our example gives rise to three different classes of nodes with the following signatures

| Class | Signature |
| :--- | ---: |
| $\alpha$ | 101 |
| $\beta$ | 100 |
| $\gamma$ | 110 |

If we choose one substring per node class, we get three probes. On the other hand there are 41 distinct substrings of $s_{0}, s_{2}$, and $s_{2}$, so we save about $93 \%$.


Figure 4.1: Suffix tree for the strings CAGTGC, CAGTTC, and CATGGA

## 4 Probe Generation

### 4.3 Pitfalls

There are some issues as to what subsequences to associate with a particular signature. Especially if a redundancy $r \geq 1$ or a minimum edit distance $d \geq 1$ is required, we have to take special care to not artificially restrict the solution space. Otherwise we risk to cut off the optimal solution. If we aim at a solution with a higher redundancy $r \geq 1$, it is not anymore sufficient to choose one probe for each distinct signature as described above. In this case we need to select $r$ or, in case the total number of possible probes for this signature is smaller than $r$, as many probes as possible. If we do not do this, the problem may still be feasible but incomplete in the sense that the optimal solution may not be reachable anymore. For the sake of completeness, Figure 4.2 displays a counter example comprising six input sequences and the ILP it translates to when we require a redundancy of $r=2$.

The $x_{i}$ are binary variables because we selected one only probe per signature. One optimal solution to this ILP with objective value 6 is given by $x_{0}=x_{1}=x_{2}=x_{3}=$ $x_{6}=x_{8}=1$ and $x_{i}=0$ for all other variables. However, if we do not restrict ourselves to one probe per signature but generate $r=2$ probes per signature instead, the binary constraint changes to $x_{i} \in\{0,1,2\}$ and we are able to achieve a better solution with objective value 5 , given by $x_{1}=2, x_{2}=x_{3}=x_{7}=1$ and $x_{i}=0$ otherwise.
$s_{1}=$ GGTAATAATACTACTAGTAGTACTCATCATAGTCCTCGTGATGCTGGTACTAAATAACTACC
$s_{2}=$ GGTAATAATACTACTAGTAGTACTCATCATAGTCCTCGTGATGCTGGTACTAAATAAGTACG
$s_{3}=$ GGTAATAATACTACTAGTAGTACTCATCATAGTCCTCGTGATGCTGGTACTAAGTAGA
$s_{4}=$ GGTAATAATACTACTAGTAGTACTCATCATAGTCCTCGTGATGCTGGTACTAAATAATACATAGC
$s_{5}=$ GGTAATAATACTACTAGTAGTACTCATCATAGTCCTCGTGATGCTGGTACTAACTAGG
$s_{6}=$ GGTAATAATACTACTAGTAGTACTCATCATAGTCCTCGTGATGCTGGTACTACATCAA

$$
\begin{aligned}
& \text { minimize } \quad x_{0}+x_{1}+x_{2}+x_{3}+x_{4}+x_{5}+x_{6}+x_{7}+x_{8} \\
& \text { subject to } x_{0}+x_{2} \quad+x_{7} \geq 2 \\
& x_{0}+x_{1}+x_{2} \quad+x_{7}+x_{8} \geq 2 \\
& x_{0}+x_{2}+x_{3} \quad+x_{6} \quad \geq 2 \\
& x_{0}+x_{1} \quad+x_{5} \quad \geq 2 \\
& x_{0}+x_{1}+x_{2}+x_{3}+x_{4} \quad \geq 2 \\
& x_{1} \quad+x_{8} \geq 2 \\
& x_{3} \quad+x_{6}+x_{7} \quad \geq 2 \\
& x_{1}+x_{2}+x_{5}+x_{7} \geq 2 \\
& x_{1}+x_{3}+x_{4} \quad+x_{7} \quad \geq 2 \\
& x_{1} \quad+x_{3} \quad+x_{6}+x_{7}+x_{8} \geq 2 \\
& x_{2} \quad+x_{5}+x_{7}+x_{8} \geq 2 \\
& x_{3}+x_{4} \quad+x_{7}+x_{8} \geq 2 \\
& x_{1}+x_{2}+x_{3} \quad+x_{5}+x_{6} \quad \geq 2 \\
& x_{1} \quad+x_{4} \quad+x_{6} \quad \geq 2 \\
& x_{2}+x_{3}+x_{4}+x_{5} \quad \geq 2 \\
& x_{i} \in\{0,1\}
\end{aligned}
$$

Figure 4.2: SB instance and corresponding ILP, showing effect of redundancy on probe generation

## 4 Probe Generation

## 5 Previous Work

This chapter presents a selection of approaches to solve the SB and SC problems, which have been defined in Sections 3.1 and 3.2. For each algorithm, we will give a short summary and try to convey the underlying key idea. Being one of the classic NPhard combinatorial optimization problems, the extent of scientific coverage of the SC problem is enormous. For this reason, we picked out those approaches that seemed most promising, with respect to both efficiency and solution quality, and suitable for our purpose. Section 5.1 addresses approaches specific to the SB problem, whereas Section 5.2 summarizes several algorithms targeted at the more general SC problem. With the exception of the ILP approach of Rash and Gusfield [32], all algorithms are of heuristic nature because none of the exact approaches we came across appeared to be powerful enough to solve large-scale instances as those presented in Chapter 8.

### 5.1 String Barcoding

Rash and Gusfield [32] introduced the problem as described in Section 3.1. They describe how a complete set of candidate probes can be generated using suffix trees as a means to ensure that no two probes distinguish the same set of sequence pairs; Chapter 4 explains this procedure in full detail. Applied to real-world instances, the reduction in the number of generated probes, compared to the number of all distinct subsequences, exceeds $75 \%$ on average and reaches over $90 \%$ in many cases.

Given a set of sequences $s_{1}, \ldots, s_{n}$ and a valid solution, comprising probes $p_{i_{1}}, \ldots, p_{i_{k}}$, we can assign a unique binary signature $\sigma_{i}=b_{1}^{i} b_{2}^{i} \ldots b_{k}^{i}$ to each sequence $s_{i}$ with $b_{j}^{i}=1$ if probe $p_{i_{j}}$ is a subsequence of $s_{i}$, and $b_{j}^{i}=0$ otherwise. Obviously, the number of possible signatures of length $l$ is $2^{l}$, so the number of digits, which in turn corresponds to the number of probes, must be greater than or equal to $\log _{2}(n)$. This can be used to

## 5 Previous Work

make the ILP formulation slightly stronger by adding the constraint

$$
\begin{equation*}
\sum_{i=1}^{m} x_{i} \geq\left\lceil\log _{2} n\right\rceil \tag{5.1}
\end{equation*}
$$

Doing so does not affect the problem structure but may help to accelerate the branch-and-bound procedure employed by ILP solvers to find the optimal solution. During our computational experiments (Chapter 8), however, the observed effect on the running time was negligible.

The computational results presented by Rash and Gusfield were generated from realworld instances by solving the ILP given in (3.1) through (3.3) and (5.1). All benchmark instances were derived from biological data available through the GenBank [31] database. They considered only probes between 15 and 40 nucleotides long and required the minimum edit distance and the redundancy $r$ to be greater than 1 for most benchmark instances.

As described in Chapter 4, it is sufficient to take $r$ probes per node class in the suffix tree. Unfortunately, this is no longer true if we want to produce a probe set with a minimum pairwise edit distance greater than 1. In this case there is no way but to include all possible probes at each suffix tree node, which in most cases makes the resulting ILP so large that it is not solvable anymore - at least not with the computational resources available to us. Most probably, Rash and Gusfield made some kind of preselection instead of taking all probes per node. If this is true, they did not state on which criteria this selection was based.

Another issue is the handling of degenerate bases. Since sequence analysis techniques are not perfectly accurate, the resulting sequence data usually contains some positions for which we cannot state the nucleotide with certainty. For this reason, real-word data often contains degenerate bases, which represent the non-empty subsets of nucleotides that are possible for a particular sequence position (see Table 5.1 for a complete list). They are, however, not treated differently from simple nucleotides in [32].

Borneman et al. [8] present a heuristic approach based on Lagrangian relaxation. They factor out the problem of probe generation and assume a given set of predeter-

| Degenerate base | Set of possible nucleotides |
| :---: | :--- |
| R | $\{A, G\}$ |
| Y | $\{C, T\}$ |
| M | $\{A, C\}$ |
| K | $\{G, T\}$ |
| W | $\{C, G\}$ |
| B | $\{A, T\}$ |
| D | $\{C, G, T\}$ |
| H | $\{A, G, T\}$ |
| V | $\{A, C, T\}$ |
| N | $\{A, C, G\}$ |

Table 5.1: List of degenerate bases
mined probes like all distinct subsequences of a particular length. Furthermore, the computational results cannot be compared to those published in [32] because different benchmark instances have been used. Two of the four data sets used are considerably larger with respect to the number of contained sequences than any benchmark instance used in [32].

The heuristic uses Lagrangian relaxation of all constraints (3.2) and subgradient optimization to find Lagrangian multipliers that yield a high lower bound. Due to the integrality property, the solution of the relaxation is integral, albeit not necessarily feasible, so at each step of the subgradient optimization procedure a new partial solution is generated. It is greedily extended by iteratively adding the probe with maximum score to the solution until all sequence pairs are covered. The score of a probe is derived from its Lagrangian cost and the number of undistinguished sequence pairs it distinguishes the same measure has been used by Caprara et al. [10].

To speed up the procedure, a clustering technique is used. Similar sequences build a cluster, out of which we select one representative sequence. The set of representatives defines a subproblem that can be solved faster and whose solution can be used as an initial solution for the original instance.

They additionally propose a simulated annealing approach for a problem variant where the size of the probe set is fixed and the number of distinguished sequence pairs is to be maximized.

Berman et al. [6] have shown that no polynomial time approximation algorithm for SB can guarantee an approximation factor of $(1-\varepsilon) \log n$ for any $\varepsilon>0$ unless NP $=$ DTIME $\left(n^{\log \log n}\right)$ and present an information content heuristic with an approximation factor of $1+\log n$.

Each partial solution defines a partition of the set of sequences into disjoint classes. Two sequences belong to the same class if and only if they cannot be distinguished by any probe in the solution. Let $n$ be the number of sequences and $C_{1}, \ldots, C_{k}$ the classes defined by the current solution. The entropy of the solution is defined as

$$
H:=\frac{1}{n} \log _{2}\left(\prod_{i=1}^{k}\left|C_{i}\right|!\right)
$$

Starting with an empty solution, the algorithm iteratively adds the probe that maximizes the decrease in entropy, until the solution is feasible.

The problem of probe generation is not addressed.

DasGupta et al. [13] describe another greedy algorithm with an approximation factor of $2 \log n$ which, compared to [6], allegedly produces results of virtually identical quality. The algorithm has two phases: a probe generation phase and a probe selection phase. To keep the number of generated probes low, those that appear in all sequences are discarded. On the other hand, for each sequence only one probe is generated that appears exclusively in that sequence. While being faster than the suffix tree approach, this procedure may produce redundant candidate probes, which appear in the same set of sequences.

The selection phase is a very simple iterative procedure, which is executed until all sequence pairs are distinguished. At each iteration we select the probe that distinguishes the highest number of yet undistinguished sequences. The source code as well as an online server is publicly available at http://dna.engr.uconn.edu/~software/DNA-BAR.

### 5.2 Set Cover Approaches

As has been noted in Chapter 3, SB is a special case of SC. This gives us the opportunity to use any of the numerous SC approaches to solve SB instances as well. Note, however, that the structure of SB instances, at least of those used by Rash and Gusfield, differs from those commonly used for benchmarking SC algorithms. The instances in Beasley's OR-Library [5], which are often used for benchmarking purposes, range in size from 500 to 5,000 items and from 10,000 to $1,000,000$ sets per instance, the density of the ILP matrices ranges from below $1 \%$ to up to $20 \%$.

On the other hand, the number of columns of the ILP representation of an SB instance grows quadratically with the number of input sequences. The benchmark instances used by Rash and Gusfield [32] give rise to constraint matrices with 5,000-50,000 rows, $500-2,000$ columns, and a density of about $20 \%$. Note too, that the number of columns depends very much on the length constraints, which may be imposed on the probes. The numbers given result from an allowed probe length of 15 to 40 nucleotides.

Many of the approaches outlined below have been designed to handle weighted SC instances, where each set incurs some cost and the goal is to minimize the total cost instead of the number of sets contained in the solution. The following summary, however, considers the unicost case only, because the SB problem as discussed in this thesis is unweighted, too.

When we refer to an SC instance further on in this chapter, we will denote the set of items to be covered by $S:=\left\{s_{1}, \ldots, s_{n}\right\}$ and the collection of subsets of $S$ as $\mathcal{C}:=$ $\left\{C_{1}, \ldots, C_{m}\right\}$

Caprara et al. [10] propose one of the most cited and seemingly most successful heuristic approaches to date to tackle large-scale SC instances. Like many others, it is based on Lagrangian relaxation and uses several well-known techniques to improve its performance.

The algorithm has three parts, all of which rely on information derived from Lagrangian relaxation. The first part is the most important and produces actual heuristic solutions. The second part examines the best solution produced by the first part and tries to

## 5 Previous Work

determine sets in the solution that are highly likely to be part of an optimal solution. These will be fixed, which means we will make them a mandatory part of the solution and discard all items that are covered by them. As a result, we get a new subinstance of smaller size to which we reapply the whole procedure. The third part is responsible for keeping the size of an instance manageable. This is achieved by making only a subset of sets visible to the algorithm. This subset is updated regularly.

The complete procedure is used repeatedly until we either exceed a predefined time limit or produce a solution of sufficient quality.

The algorithm's main part uses subgradient optimization to find a near-optimal set of Lagrangian multipliers. Then, a set of multiplier vectors of comparable quality is generated, and for each of them we build a solution by means of a simple greedy heuristic, which is guided by the Lagrangian costs of the sets. Finally, we select some promising sets, fix them as part of the solution, and remove from the instance all items covered by these sets. We repeat this procedure until either all items are covered by fixed sets, or until we can conclude that no further improvement is possible.

This thesis focuses on the adaptation of this algorithm for the SB problem, so a more detailed description will be given in Chapter 6.

Beasley and Chu [3] propose a steady-state genetic algorithm (GA). Inspired by the natural evolutionary process, a genetic algorithm operates on a pool of solutions (population) whose average quality improves over time. In each iteration, solutions are selected based on their quality (fitness) and pairwise recombined into new solutions (parents and children), to which, with a certain probability, slight random modifications are applied (mutation). There are many variants of genetic algorithms which differ from each other in the number of solutions selected, the criteria on which the selection is based, and which solutions go into the population of the next iteration. In case of the steady-state GA at hand, we select only two solutions per iteration, recombine them, apply the mutation operator, and substitute the result for the worst solution in the current pool.

Beasley and Chu represent solutions as binary strings $x_{1} x_{2} \ldots x_{m}$, with $x_{i}=1$ if and only if set $C_{i}$ is part of the solution. They use binary tournament selection, which means that they randomly choose two solutions from the population and select the one with better fitness. The recombination operator randomly takes the value for each $x_{i}$ from either
parent, with a selection probability proportional to the fitness of the respective parent solution. The mutation probability is variable and increases as the algorithm converges. Mutation as well as recombination operators may yield infeasible solutions, so a repair operation is applied if necessary. This is done by greedily extending the partial solution until all items are covered and heuristically removing redundant sets afterward.

Eremeev [14] proposes another steady-state GA, which uses a non-binary representation. A valid solution $\mathcal{C}^{*}=\left\{C_{i_{1}}, \ldots, C_{i_{k}}\right\}$ is represented as a vector $\left\langle g_{1}, g_{2}, \ldots, g_{n}\right\rangle$, where $g_{i}=j, j \in\left\{i_{1}, \ldots, i_{k}\right\}$, means that item $s_{i}$ is covered by set $C_{j}$, and maybe some others too. This mode of representation is not unique in that a particular solution may have more than one representation. On the other hand, it has the advantage that both the recombination and mutation operators always yield feasible solutions, so a repair operator is not necessary.

The selection operator picks solutions randomly from the population with each solution having a probability to be chosen proportional to its fitness. Accordingly, this is often called fitness-proportional selection.

The recombination operator constructs a subproblem, which depends on the two solutions on which it is applied, and tries to solve the corresponding LP relaxation. If the number of constraints is above a certain threshold, solving the LP relaxation takes too long, or if the solution is not integral, the first parent solution will be taken without further modification.

Given two solutions $\mathcal{C}_{1} \subseteq \mathcal{C}$ and $\mathcal{C}_{2} \subseteq \mathcal{C}$, the recombination operator performs the following steps to construct the subproblem:

1. Let $\mathcal{C}^{\prime}:=\mathcal{C}_{1} \cup \mathcal{C}_{2}$ and $S^{\prime} \subseteq S$ be the set of items that are covered by exactly one set from $\mathcal{C}^{\prime}$
2. Determine the set $\mathcal{C}^{\prime \prime}:=\left\{C \in \mathcal{C}^{\prime}: C \cap S^{\prime} \neq \emptyset\right\}$
3. Remove from the instance all sets in $\mathcal{C}^{\prime \prime}$ as well as all items covered by $\mathcal{C}^{\prime \prime}$

The mutation operator takes as input a single element $g_{i}$ of a solution's representation vector and merely assigns a new value to that element. From all sets covering

## 5 Previous Work

the corresponding item $s_{i}$ one is chosen randomly, with more expensive sets having a lower probability to be picked. The mutation operator is applied to each element in turn, although only with a fixed probability. Note, however, that the solution does not necessarily change when its representation does.

The solution produced by the successive application of the recombination and mutation operators is always feasible but may contain redundant sets. Eremeev proposes two different greedy heuristics to eliminate redundant elements from a solution. Both of them are applied independently, and the better result goes into the next population.

Marchiori and Steenbeek [26] introduce an iterative heuristic. The heuristic operates on a dynamic subproblem only, which is recomputed every 100 iterations. Each iteration starts with a partial solution, which is empty in the beginning and is derived non-deterministically from the best solution found so far after each iteration. At each iteration, we alternately add and remove some elements to or from the solution until it is feasible and does not contain any redundant sets.

When adding elements, we do so in a greedy manner by repeatedly adding a set that covers a maximum number of yet uncovered items. The removal of elements is done unconditionally if the solution contains redundant sets, by removing one of them arbitrarily, or otherwise, only with a particular probability, by removing a random set.

Finally, a local optimization procedure identifies and adds sets that make two or more sets of the current solution redundant, which are removed in consequence.

Yagiura et al. [37] describe a local search approach combined with a variable fixing scheme, which is particularly useful when solving large scale instances. Both local search and variable fixing are guided by information derived from Lagrangian relaxation.

The neighborhood of a solution $S$ is defined as the set of solutions that differ from $S$ in at most three sets. The neighborhood has size $O\left(n^{3}\right)$, so its complete exploration is very expensive. Accordingly, the proposed implementation searches only part of the neighborhood, yet without sacrificing solution quality.

As another key point, they allow infeasible solutions by use of a penalized objective
function. A penalty weight is assigned to each item to be covered, which determines the extent by which the objective value of a solution increases if the respective item is uncovered.

The procedure repeatedly performs local search, until a local optimum is reached, followed by an update to the penalty weights. This update allows us to control the degree of infeasibility allowed and is used to intensify the search around the border of the feasible region. The rationale behind this is that any optimal solution $S^{*}$ must be located in this region as removing a single set from $S^{*}$ would render it infeasible.

Brusco et al. [9] propose a simulated annealing algorithm, supplemented with a morphing procedure. They target SC instances that exhibit a high degree of coverage correlation, which measures the extent to which the sets of an instance cover the same items.

Simulated annealing [23, 35] is inspired by a conceptually similar technique used in materials science, which involves heating followed by controlled cooling of a material like metal or glass to improve its quality. It is an iterative procedure which repeatedly determines the best neighbor of the incumbent solution, substituting it for the current solution if it is either superior or otherwise with a certain probability. This probability depends on the current temperature which decreases over time. The probabilistic acceptance of inferior solutions provides a means of escaping local optima.

Brusco et al. define the neighborhood of a solution by means of a non-deterministic algorithm, which first removes a set from the current solution, then adds sets to make the solution feasible again, and finally removes redundant sets. The set to be removed in the first step is either a random one or the set that would leave the least number of items uncovered if removed. During the repair step we greedily add sets that have a low cost and cover many items that are not covered yet.

The proposed morphing procedure involves maintaining for each set in the solution a list of sets that cover similar items. In each iteration we replace a set of the solution with one from the respective morphing list if this leads to an improvement.

## 5 Previous Work

Caprara et al. [11] give an extensive survey of further SC approaches, which covers several exact algorithms, too.

## 6 A Lagrangian Heuristic for the String Barcoding Problem

### 6.1 Outline

In essence, the heuristic presented in this chapter conforms to the CFT algorithm published by Caprara et al. [10]. Apart from some minor extensions to improve overall performance, we merely made changes that enabled the algorithm to deal with instances of the SB problem as defined in Section 3.1. Note that the original algorithm was designed to handle weighted SC instances, where each subset incurs some fixed cost and the goal is to find a cover of minimum total cost. Although the same concept could be applied to the SB problem as well, we will restrict the discussion to the unicost variant as presented in [32]. At some sites the original algorithm uses set costs as a sorting or selection criterion. In these cases, we use random perturbation or selection instead.

The algorithm, which we will refer to as the CFTSB heuristic, relies heavily on the concept of Lagrangian relaxation, which has been introduced in Section 2.2. Caprara et al. propose the use of subgradient optimization (Section 2.3) to compute good Lagrangian multipliers. However, when starting from scratch, finding satisfactory multipliers can be a very time-consuming process. We have found that an initial set of optimal multipliers can be computed in reasonable time by solving the LP relaxation instead. In general, this approach may not scale very well because the required effort to solve a particular LP is hard to estimate. With regard to our benchmark instances, however, this modification of the original algorithm leads to considerably lower run times without sacrificing solution quality.

The heuristic pursues a two-level, iterated approach. At the top level, we repeatedly generate feasible solutions for the given instance, each time starting from a different

## 6 A Lagrangian Heuristic for the String Barcoding Problem

partial solution. This partial solution is empty at the beginning of the first iteration and is subsequently derived from the solution generated in the previous iteration. We repeat this procedure until we either reach a predefined time limit or find a solution that is good enough for our purposes.

The actual solutions are generated at the lower level, which is conceptually similar to the top level procedure. We perform a series of iterations, each one starting from a partial solution, which is empty at the very beginning. As the first step of each iteration, we try to find good Lagrangian multipliers for the subproblem defined by the given partial solution. These are used to produce a set of heuristic solutions, each one extending the current partial solution. At the end of each iteration, some further probes are chosen to extend the partial solution, which serves as a starting point for the next iteration. We repeat this procedure until either the partial solution itself has become feasible, or until it is evident that we cannot find a solution better than the best one found so far.

To improve overall performance, especially with regard to large-scale instances, a pricing scheme is applied to the whole procedure. This means that, at any time, the algorithm takes into account only a subset of all available candidate probes; this subset will be dynamically updated at regular intervals. The use of this technique may reduce the run time considerably, even for medium-sized instances. We will discuss this topic in more detail in Section 6.2.4.

Before we delve further into the details, we introduce some notational conventions, which will be used throughout the chapter, and present specifics on the application of Lagrangian relaxation and subgradient optimization to the SB problem.

Each SB instance $\mathcal{I}=(S, P, r, F)$ is defined by a set $S$ of sequences, for which we seek a probe set of minimum size, a set $P$ of possible probes to choose from, the number $r$ of probes by which each sequence pair must be distinguished, and an optional partial solution $F \subseteq P$. Obviously, the set $F$ affects the number probes by which a particular sequence pair must be distinguished; for any sequence pair $\left(s_{i}, s_{j}\right) \in S \times S, i \neq j$, this number is given by:

$$
\begin{equation*}
r_{i, j}(F)=\max \left\{r-\sum_{p_{k} \in F} \delta_{i, j}(k), 0\right\} \tag{6.1}
\end{equation*}
$$

Any sequence pair $\left(s_{i}, s_{j}\right)$ with $r_{i, j}(F)=0$ is distinguished by the required number of probes and therefore needs not to be taken into consideration anymore. Thus, any partial solution $F$ defines a subinstance $\mathcal{I}_{F}$ of the original instance. For convenience, we will denote the set of sequence pairs that need further probes to distinguish them by

$$
\begin{equation*}
S^{2}(F):=\left\{\left(s_{i}, s_{j}\right) \in S \times S: i \neq j, r_{i, j}(F)>0\right\} \tag{6.2}
\end{equation*}
$$

We apply Lagrangian relaxation to the ILP formulation of the SB problem given in (3.1) through (3.3) and additionally relax the integrality constraints. This leads to the following minimization problem:

$$
\begin{align*}
\min L(u) & :=\sum_{k=1}^{m} x_{k}+\sum_{1 \leq i<j \leq n} u_{i, j}\left(\sum_{k=1}^{m} \delta_{i, j}(k) \cdot x_{k}\right)  \tag{6.3}\\
& =\sum_{k=1}^{m} x_{k} \underbrace{\left(1-\sum_{1 \leq i<j \leq n} \delta_{i, j}(k) \cdot u_{i, j}\right)}_{\text {Lagrangian cost } c_{k}(u)}+r \sum_{1 \leq i<j \leq n} u_{i, j} \\
& =\sum_{k=1}^{m} c_{k}(u) \cdot x_{k}+r \sum_{1 \leq i<j \leq n} u_{i, j}
\end{align*}
$$

$$
\begin{equation*}
\text { subject to } \quad 0 \leq x_{k} \leq 1, k=1, \ldots, m \tag{6.4}
\end{equation*}
$$

Each variable $x_{i}$ corresponds to probe $p_{i}$. Recall that the term $c_{i}$ denotes the Lagrangian cost of variable $x_{i}$ or probe $p_{i}$, respectively.

Given some Lagrangian multiplier vector $u$, an optimal solution $x^{*}$ to (6.3) - (6.4) can be derived easily as follows:

$$
x_{i}^{*}(u)= \begin{cases}1 & \text { if } c_{i}(u)<0 \\ 0 & \text { if } c_{i}(u) \geq 0\end{cases}
$$

Note that the Lagrangian relaxation defined by (6.3) through (6.4) has the integrality

## 6 A Lagrangian Heuristic for the String Barcoding Problem

property (see Section 2.2). As a consequence, any optimal solution is integral, which is the reason why we relax the integrality constraints too. Unfortunately, this does also imply that we will not be able to find a lower bound higher than that provided by the optimal solution of the LP relaxation.

When using subgradient optimization as a means to find good Lagrangian multipliers, we define the subgradient $s$ as

$$
\begin{equation*}
s_{i, j}(u):=r-\sum_{k=1}^{m} \delta_{i, j}(k) \cdot x_{k}^{*}(u) \tag{6.5}
\end{equation*}
$$

Accordingly, the update formula for the Lagrangian multipliers is given by

$$
\begin{equation*}
u_{i, j}^{k+1}:=\max \left\{u_{i, j}^{k}+\lambda \frac{U-L\left(u^{k}\right)}{\left\|s\left(u^{k}\right)\right\|^{2}} s_{i, j}\left(u^{k}\right), 0\right\} \tag{6.6}
\end{equation*}
$$

where $U$ is a valid upper bound, derived from any feasible solution, and $\lambda$ denotes the step size parameter, which will be updated periodically.

Algorithm 2 depicts the top level procedure in pseudocode. In line 7, we produce a valid solution for subinstance $\mathcal{I}_{F}$ defined by partial solution $F$. This is the most important part of the complete heuristic and will be addressed in Section 6.2. In addition to a feasible probe set, function CoreOpt returns the best set of Lagrangian multipliers it has found. Observe that the multipliers returned are associated with subinstance $\mathcal{I}_{F}$ and do not necessarily yield a good lower bound for the original problem. Therefore, we use the returned multiplier vector only if $F$ is empty, in which case $\mathcal{I}=\mathcal{I}_{F}$ holds, and if it results in a better lower bound for $\mathcal{I}$ than the best multiplier vector found so far (lines 12-14). Parameter $\pi$ controls the size of subinstance $\mathcal{I}_{F}$ derived from $F$ and is updated in lines $15-19$. Section 6.2.5 is dedicated to the inner workings of function Refine that computes partial solution $F$ at the end of each iteration.

```
Algorithm 2 CFTSB heuristic - top level perspective
    function \(\operatorname{CFTSB}(S, P, r)\)
        \(Q \leftarrow P\)
        \(F \leftarrow \emptyset\)
        \(u \leftarrow 0\)
        \(\pi \leftarrow \pi_{\text {min }}\)
        repeat
            \(\left(Q^{\prime}, u^{\prime}\right) \leftarrow \operatorname{CoreOpt}(S, P, r, F, u)\)
            \(Q^{\prime} \leftarrow Q^{\prime} \cup F\)
            if \(\left|Q^{\prime}\right|<|Q|\) then
                \(Q \leftarrow Q^{\prime}\)
            end if
            if \(F=\emptyset\) then
                \(u \leftarrow u^{\prime}\)
            end if
            if \(\left|Q^{\prime}\right|=|Q|\) or \(F=\emptyset\) then
                \(\pi \leftarrow \pi_{\text {min }}\)
            else
                \(\pi \leftarrow 1.1 \cdot \pi\)
            end if
            \(F \leftarrow \operatorname{Refine}\left(Q^{\prime}, \pi\right)\)
        until \(L(u)=|Q|\) or time limit exceeded
        return \(Q\)
    end function
```


### 6.2 Core Algorithm

Actual solutions for a given instance, or a subinstance thereof, are generated at the lower level of the CFTSB heuristic. The complete process is illustrated in Algorithm 3 and can be broken down into three steps:

Step 1 In lines 4-7, we attempt to find a near-optimal set of Lagrangian multipliers. In the first iteration, we achieve this by solving the LP relaxation to optimality. Subsequently, a slightly modified subgradient optimization procedure is employed for this purpose. Details on this step will be presented in Section 6.2.1.

Step 2 In the next step (line 8) we attempt to generate feasible solutions for the given problem instance using a simple greedy heuristic that iteratively adds probes to an initially empty set of probes until it constitutes a feasible solution. The decision which probe to add next is guided by the number of still insufficiently distinguished sequence pairs a probe covers in relation to its Lagrangian cost. A detailed rundown will be given in Section 6.2.2.

Step 3 During the third step (line 16), we select probes which have a presumably high likelihood of appearing in an optimal solution and add them to partial solution $F$. Details on how the selection is performed can be found in Section 6.2.3. Before we loop back to step one, we perturb the Lagrangian multipliers randomly so as to facilitate the generation of diverse solutions (line 17).

Function CoreOpt terminates when either $F$ itself constitutes a feasible solution observe that we never remove a probe from it - or when we can infer that we will not be able to construct a solution that is better than the best one found so far; this is the case if the number of probes in $F$ plus the lower bound for the remaining subproblem is equal to or greater than the current best solution.

### 6.2.1 Lagrangian Multiplier Optimization

In this section, we describe two possible methods that can be used to find good or even optimal Lagrangian multipliers for a given problem instance $\mathcal{I}=(S, P, r, F)$. An initial set of optimal multipliers is computed by solving the LP corresponding to instance $\mathcal{I}$.

```
Algorithm 3 Core component of CFTSB heuristic
    function \(\operatorname{CoreOpt}(S, P, r, F, u)\)
        \(Q \leftarrow P\)
        repeat
            \(u^{\prime} \leftarrow \operatorname{Mult} \operatorname{Opt}(u) \quad\) Find (near-)optimal multipliers
            if \(L\left(u^{\prime}\right)<L(u)\) then
                \(u \leftarrow u^{\prime}\)
            end if
            \(\left(Q^{\prime}, u^{\prime}\right) \leftarrow \operatorname{HeurOpt}(S, P, r, F, u) \quad\) Generate heuristic solution
            \(Q^{\prime} \leftarrow Q^{\prime} \cup F\)
            if \(\left|Q^{\prime}\right|<|Q|\) then Update solution
                    \(Q \leftarrow Q^{\prime}\)
            end if
            if \(L\left(u^{\prime}\right)<L(u)\) then
                \(u \leftarrow u^{\prime}\)
            end if
            Extend \(F^{\prime}\) by fixing additional probes
            Choose \(r \in[-0.1,0.1]\) randomly Perturb multipliers
            \(u_{i, j} \leftarrow(1+r) \cdot u_{i, j}\)
        until \(F=P\) or \(|F|+L(u)>|Q|\)
        return \((Q \backslash F, u)\)
    end function
```


## 6 A Lagrangian Heuristic for the String Barcoding Problem

Subgradient optimization is used to compute ensuing multipliers, starting from the most recent set of multipliers.

## LP Relaxation Method

This method computes optimal Lagrangian multipliers. We have shown in Section 2.2 that, if the integrality property holds, the dual values associated with an optimal solution to the LP relaxation pose an optimal solution to the Lagrangian relaxation as well. Given some optimal solution to the LP relaxation, let $y^{*}$ be the corresponding dual solution. If we assign the value of $y_{i, j}^{*}$ to Lagrangian multiplier $u_{i, j}$, both relaxations yield the same objective value. This is optimal because the objective value of the Lagrangian dual cannot exceed that of the LP relaxation due to the integrality property.

The main disadvantage of this approach is that it is hard to estimate the computational effort required to solve the LP relaxation to optimality. State-of-the-art LP solvers, however, may be able to solve the LP relaxation faster than the subgradient optimization process takes. At least this is the case with the benchmark instances we use.

## Subgradient Method

Caprara et al. propose the use of subgradient optimization, which has been introduced in Section 2.3, to find high quality Lagrangian multipliers. It is particularly useful when tackling very large instances but, on the other hand, there is no guarantee as to the quality of the multipliers it produces. For the most part, the procedure conforms to the pseudocode given in Algorithm 1.

As has been noted above, we do not use this technique to compute Lagrangian multipliers from scratch. Instead, we take the best multipliers found during the previous iteration, perturb them randomly, and use the result as a starting point.

In each iteration the current multiplier vector $u$ will be updated according to (6.6). As for the upper bound $U$, we use the size of the best known solution for the current subproblem. Prior to the first iteration, we compute an initial solution by means of the greedy heuristic described in Section 6.2.2. Caprara et al. propose to use $u=0$ as
input when computing the initial solution but we decided to initialize the Lagrangian multipliers beforehand, as this seems more reasonable to us.

The step size parameter $\lambda$ controls how far we move along the subgradient. If we are still far from the optimum, its value should be high. If we come closer to the optimum, its value should get smaller because otherwise we risk to overshoot. We use an initial value of $\lambda=0.1$. In the classical approach, the step size is halved if the lower bound cannot be improved for a particular number of iterations. We use a slightly different method which, according to Caprara et al., makes for faster convergence. Every $p=20$ iterations we compare the best and worst lower bounds, $L_{\max }$ and $L_{\min }$, achieved in this period. If these two values are within $0.1 \%$ of other, we assume that we either are very close to an optimum or the low improvement is due to the step size being to small. Either way, we will enlarge the step size by a factor of 1.5 . Even if we were already very close to the optimum, we do not loose much, because near-optimal multipliers are good enough for our purposes. If, on the other hand, $L_{\min }$ and $L_{\max }$ differ by more than $1 \%$, we decrease the step size by a factor of 2 to prevent moving past the optimum. In summary, this amounts to the following updating scheme:

$$
\lambda= \begin{cases}1.5 \cdot \lambda & \text { if } L_{\max } \leq 1.001 \cdot L_{\min } \\ 0.5 \cdot \lambda & \text { if } L_{\max } \geq 1.01 \cdot L_{\min } \\ \lambda & \text { otherwise }\end{cases}
$$

In practice, a large number of probes will have Lagrangian costs very close to zero, which usually implies the existence of many almost optimal solutions each of which can be obtained by including different subsets of those probes into the solution. Correspondingly, there are just as many subgradients (6.5) to choose from. Although it is theoretically possible to determine the subgradient providing the steepest ascent, it is computationally onerous to do so.

Instead, Caprara et al. [10] propose the use of a heuristically selected small-norm subgradient direction, which can be computed quickly and, compared to the subgradient defined in (6.5), leads to faster convergence of the subgradient optimization procedure in most cases.

We start with a partial solution $G:=\left\{p_{i} \in P: c_{i}(u) \leq 0.001\right\}$. Let $Q_{G}$ be the set of

## 6 A Lagrangian Heuristic for the String Barcoding Problem



Figure 6.1: Comparison between alternative subgradient definitions
sequence pairs that are distinguished by at least one probe in $G$. Let $R$ denote the set of redundant probes in $G$, i. e. those which can be removed safely from $G$ without affecting set $Q_{G}$. We walk through set $R$ in arbitrary order, removing a probe from $G$ if this does not affect set $Q_{G}$. Thus, we need to check for each probe in turn if its removal would leave any sequence pair in $Q_{G}$ undistinguished. However, for reasons of efficiency we do not update set $R$ after having deleted a probe from $G$. The small-norm subgradient is defined with respect to set $G$, after removing redundant elements as described above, as follows:

$$
\begin{equation*}
\dot{s}_{i, j}(u):=r-\sum_{p_{k} \in G} \delta_{i, j}(k) \tag{6.7}
\end{equation*}
$$

Note that, strictly speaking, $\dot{s}(u)$ may not qualify as a subgradient anymore. Figure 6.1 contrasts the standard subgradient as defined in (6.5) with the reduced-norm variant (6.7). It exemplifies the evolution of the lower bound for both subgradient variants, when applied to one of our benchmark instances. As one can see, the use of the smallnorm subgradient causes the subgradient procedure to converge to a higher lower bound. When applied to other benchmark instances, the results look similar.

We terminate the subgradient optimization when the lower bound does not improve
anymore, or only to a very small extent. This is the case if the absolute and relative lower bound improvement over the the last 300 iterations is below 1.0 and $0.1 \%$, respectively. Furthermore, to be on the safe side we impose a hard limit of $10 \cdot\left|Q_{G}\right|$ on the number of iterations.

### 6.2.2 Greedy Heuristic

Given a set of Lagrangian multipliers, the greedy heuristic presented in this section produces a feasible solution for a given SB instance. More precisely, we compute a series of Lagrangian multipliers, starting with the best known multiplier vector. For each multiplier vector in this series, we generate a feasible solution and return the best solution found.

We perform a fixed number of iterations of the subgradient optimization procedure given in Algorithm 1 to generate a series of Lagrangian multiplier sets. Apart from the fixed number of iterations, we also use the standard subgradient definition given in (6.5). Note that, as we are basically continuing the subgradient optimization procedure, we may as well find some better Lagrangian multipliers. See Algorithm 4 for the corresponding pseudocode.

```
Algorithm 4 Iterative heuristic
    function \(\operatorname{HeurOpt}(S, P, r, F, u)\)
        \(Q \leftarrow P \backslash F\)
        \(u^{*} \leftarrow u\)
        for \(i=1, \ldots, 250\) do
            \(Q^{\prime} \leftarrow \operatorname{Greedy}(S, P, r, F, u)\)
            if \(\left|Q^{\prime}\right|<|Q|\) then
                \(Q \leftarrow Q^{\prime}\)
            end if
            Update \(u\) according to (6.6)
            if \(L(u)>L\left(u^{*}\right)\) then
                \(u^{*} \leftarrow u\)
            end if
        end for
        return \(\left(Q, u^{*}\right)\)
    end function
```

The greedy procedure we use to compute a feasible solution for a given set of Lagrangian multipliers is depicted in Algorithm 5. We start with a given, possibly empty, partial

## 6 A Lagrangian Heuristic for the String Barcoding Problem

solution $Q$, which we extend probe by probe until we reach a feasible solution. The key point is the way we select the probes to add to the partial solution. For this, we assign to each candidate probe $p_{i}$ a score $\sigma_{i}$, which is based on the number sequence pairs not yet completely distinguished by the probes in $Q$ that $p_{i}$ distinguishes and its Lagrangian cost. More formally:

$$
\begin{aligned}
& \mu_{i}:=\sum_{\left(s_{j}, s_{k}\right) \in S^{2}(Q)} \delta_{j, k}(i) \\
& \gamma_{i}:=1-\sum_{\left(s_{j}, s_{k}\right) \in S^{2}(Q)} \delta_{j, k}(i) \cdot u_{j, k} \\
& \sigma_{i}:= \begin{cases}\frac{\gamma_{i}}{\mu_{i}} & \text { if } \gamma_{i}>0 \\
\gamma_{i} \mu_{i} & \text { if } \gamma_{i} \leq 0\end{cases}
\end{aligned}
$$

In each iteration we add the probe $p_{i}$ with minimum score $\sigma_{i}$ to $Q$.

As soon as $Q$ has become feasible, we try to remove as many redundant probes from it as possible. A probe is considered redundant if $Q$ is still feasible after its removal. As long as there are more than 10 of them, we choose one of them randomly, remove it from the solution and update the set of redundant probes. As soon as there are only 10 or less left we perform an exhaustive search to find a subset of redundant probes with maximum cardinality that can be removed safely from our solution without affecting its feasibility.

### 6.2.3 Probe Fixing

Especially when tackling large instances, the solution available at the end of an iteration of function CoreOpt can be further improved by fixing some probes, that is we require some probes to be part of the solution, and re-applying the function to the resulting subinstance.

Our goal is to fix probes which are highly likely to be member of an optimal solution. To find probes that match this criterion, we employ a simple greedy procedure.

```
Algorithm 5 Greedy heuristic
    function \(\operatorname{Greedy}(S, P, r, F, u)\)
        \(Q \leftarrow F\)
        while \(S^{2}(Q) \neq \emptyset\) do
            Let \(p^{*} \in P \backslash Q\) be the probe with minimum score \(\sigma\)
            \(Q \leftarrow Q \cup\left\{p^{*}\right\}\)
        end while
        Let \(R\) be the set of redundant probes in \(Q\)
        while \(|R|>10\) do
            Remove random probe \(r \in R\) from \(Q\)
            Update \(R\)
        end while
        Remove the maximum number of redundant probes from \(Q\)
        return \(Q \backslash F\)
    end function
```

First, we compute the Lagrangian costs of all probes with respect to the best Lagrangian multipliers found for the current subproblem. Let $F^{\prime}$ be to set of probes with Lagrangian costs below -0.001 and $Q$ the set of sequence pairs distinguished by at least one probe in $F^{\prime}$. From set $F^{\prime}$ we remove those probes that do not distinguish at least one sequence pair that is not distinguished by any other probe in $F^{\prime}$. In other words, we remove a probe from $F^{\prime}$ if and only if this does not affect set $Q$.

We apply the greedy heuristic presented in Section 6.2 .2 to the subinstance defined by $F^{\prime}$. Let $F^{\prime \prime}$ be the generated solution. In addition to the probes in $F^{\prime}$, we fix at most $\lfloor|Q| / 200\rfloor$ probes of $F^{\prime \prime}$; if $F^{\prime \prime}$ contains more probes, we select a random subset.

### 6.2.4 Pricing

The computationally most expensive parts of the overall algorithm are the multiplier optimization (Section 6.2.1) and the greedy heuristic (Section 6.2.2). Especially in case of very large instances, these parts should be implemented as efficient as possible to keep the run time low. Inspired by similar techniques used to solve very large LPs, we use a pricing technique to achieve this goal. According to Caprara et al., this has not been done before in combination with Lagrangian relaxation.

Basically, pricing makes only a subset of the available probes visible to the two aforementioned parts of the algorithm. This subset, which we will refer to further on as

## 6 A Lagrangian Heuristic for the String Barcoding Problem

the core, will be updated regularly. The dynamic nature of the core is meant to avoid missing very good or even optimal solutions due to this artificial restriction.

The choice of the probes that make up the core is once more based on their Lagrangian costs. In the beginning, we randomly choose five distinguishing probes for each sequence pair. When updating the core subsequently, the selection process works like this: Let $\mathcal{I}=$ $(S, P, r, F)$ be the instance we are operating on. We start by computing the Lagrangian costs based on the current multiplier vector $u^{k}$ and initialize the new core with set $C_{1}:=\left\{p_{i} \in P: c_{i}(u)<0.1\right\}$. Additionally, we select for each sequence pair five distinguishing probes at random, just like we did during initialization, and denote this set by $C_{2}$. If $C_{1}$ contains more than $l:=5 \cdot\left|S^{2}(F)\right|$ elements, we keep only the $l$ probes with the lowest Lagrangian costs and remove all others. The updated core is defined as the union of $C_{1}$ and $C_{2}$.

We also have to decide on the pricing frequency $T$, which tells us after how many subgradient iterations to update the core. We initially set $T:=10$. After each core update, we compute the lower bound, once with respect to all available probes and once with respect to the probes that are part of the current core $C^{k}$, subtract the first from the second, and divide the result by the best known upper bound $U$, i. e. the size of the best known solution for instance $\mathcal{I}$ :

$$
\Theta:=-\frac{\sum_{p_{i} \in P \backslash C^{k}} \min \left\{c_{i}\left(u^{k}\right), 0\right\}}{U}
$$

A small $\Theta$ indicates that the core has settled, and that changing it will most probably not lead to any improvements. Thus, increasing $T$ is the right thing to do. On the other hand, if $\Theta$ is large, it is advisable to make the update interval smaller. This leads to the following update formula:

$$
T:= \begin{cases}10 \cdot T & \text { if } \Theta \leq 10^{-6} \\ 5 \cdot T & \text { if } \Theta \leq 0.02 \\ 2 \cdot T & \text { if } \Theta \leq 0.2 \\ 10 & \text { otherwise }\end{cases}
$$

### 6.2.5 Solution Refinement

Although the solutions produced at the lower level procedure as implemented by function CoreOpt are mostly of high quality, we can often find an even better solution by fixing part of the returned solution and rerun the complete algorithm on the resulting subinstance. Similar to the probe fixing procedure described in Section 6.2.3, we intent to determine those probes of the solution that are most likely to be part of an optimal solution.

The procedure has one parameter $\pi$ that affects the size of the subinstance that results from fixing the probes we are going to select. Let $u^{*}$ be the best Lagrangian multiplier vector, $Q^{*}$ the best solution found for the original input instance $\mathcal{I}=(S, P, r, \emptyset)$, without any fixed probes, and let $\zeta_{i, j}, 1 \leq i<j \leq n$, denote the number of probes in $Q^{*}$ that distinguish sequence pair ( $s_{i}, s_{j}$ ). The gap $\Delta$ between the corresponding lower and upper bound, that is the size of the best overall solution found so far, is given by:

$$
\begin{aligned}
\Delta & =\left|Q^{*}\right|-\left(\sum_{1 \leq i<j \leq n} u_{i, j}^{*}+\sum_{i=1}^{m} \min \left\{c_{i}\left(u^{*}\right), 0\right\}\right) \\
& =\sum_{p_{i} \in Q^{*}}(\underbrace{c_{i}\left(u^{*}\right)+\sum_{1 \leq j<k \leq n} \delta_{j, k}(i) \cdot u_{i, j}^{*}}_{=1})-\sum_{1 \leq i<j \leq n} u_{i, j}^{*}-\sum_{i=1}^{m} \min \left\{c_{i}\left(u^{*}\right), 0\right\} \\
& =\sum_{p_{i} \in Q^{*}: c_{i}\left(u^{*}\right)>0} c_{i}\left(u^{*}\right)-\sum_{p_{i} \notin Q^{*}: c_{i}\left(u^{*}\right)<0} c_{i}\left(u^{*}\right)+\sum_{1 \leq i<j \leq n} u_{i, j}^{*}\left(\zeta_{i, j}-1\right)
\end{aligned}
$$

We estimate the extent to which any single probe $p_{i} \in Q^{*}$ contributes to this gap, denoted by $\eta_{i}$, as follows:

$$
\begin{equation*}
\eta_{i}:=\max \left\{c_{i}\left(u^{*}\right), 0\right\}+\sum_{1 \leq j<k \leq n} \delta_{j, k}(i) \cdot u_{j, k}^{*}\left(1-\frac{1}{\zeta_{j, k}}\right) \tag{6.8}
\end{equation*}
$$

Based on the assumption that probes with a small $\eta$ are more likely to be part of an optimal solution, we consider the elements of solution $Q^{*}=\left\{p_{q_{1}}, \ldots, p_{q_{l}}\right\}$ in ascending

## 6 A Lagrangian Heuristic for the String Barcoding Problem

order with respect to their gap estimate, i. e. $\eta_{q_{1}} \leq \ldots \leq \eta_{q_{l}}$. Depending on parameter $\pi$, the set of probes to fix is the smallest subset $F:=\left\{p_{q_{1}}, \ldots, p_{q_{i}}\right\}, i \leq l$, such that the percentage of sequence pairs fully distinguished by $F$ is greater than or equal to parameter $\pi$ :

$$
\begin{equation*}
1-\frac{2 \cdot\left|S^{2}(F)\right|}{|S| \cdot(|S|-1)} \geq \pi \tag{6.9}
\end{equation*}
$$

## 7 Implementation

This chapter provides details on our implementation of the CFTSB heuristic. We implemented the complete algorithm in C++ under the Linux operating system, with notable use of the standard template (STL) and the Boost [1] libraries. Section 7.1 summarizes the overall workflow of computing a probe set for a given set of sequences, Section 7.2 specifies the data formats of the input files, Section 7.3 addresses some implementation details, and Section 7.4 gives an overview of the classes of our design.

### 7.1 Workflow

Given a set of sequences, the workflow to compute a high-quality probe set is as follows. A separate program, which we will refer to as probe generator hereinafter, is used to compute an exhaustive set of candidate probes according to the description in Chapter 4. The probe generator also enforces any length constraints. The generated probes, along with the information in which sequences each probe occurs, will be handed as input to the actual implementation of the CFTSB heuristic. The separation of these two stages has the advantage that we do not need to repeat the probe generation phase when running the CFTSB heuristic multiple times on the same instance. However, when we change the length constraints or the redundancy parameter, we have to call the probe generator again.

Figures 7.1 and 7.2 summarize the command line options accepted by the probe generator and the main heuristic, respectively.

| --help | Display list of supported command line options along with a short <br> description. |
| :--- | :--- |
| --ifile file | Read instance from file. <br> --ofile file <br> --redundancy $n$ |
| Write output to file. |  |
| Set the number of probes in the solution by which each sequence <br> pair has to be distinguished to $n . ~ A s ~ p o i n t e d ~ o u t ~ i n ~ S e c t i o n ~$ 4, |  |
| we have to generate at least this number of candidate probes per |  |
| signature (or as many as possible). |  |

Figure 7.1: Command line options of the probe generator

| --help | Display list of supported command line options along with a short <br> description. |
| :--- | :--- |
| --seed $n$ | Seed the pseudo-random number generator with $n$ |
| --ifile file | Read instance from file. |
| --redundancy $n$ | Set the number of probes in the solution by which each sequence <br> pair has to be distinguished to $n$. |
| --maxrt $n$ | Set the maximum run time to $n$ seconds. <br> --optimum $n$ |
| Set the size of the optimal solution, if known, to $n . ~ I f ~ n o t ~ o m i t t e d, ~$ <br> the algorithm terminates as soon as it finds a solution of this size. <br> Chiefly useful for benchmarking. |  |

Figure 7.2: Command line options of the CFTSB implementation

# >gi|2769620|emb|Y16151.1|HIVY16151 HIV-1 pol gene, isolate RO-BCI9 ATGGGTGCGAGAGCGTCAGTATTAAGTGGGGGAAAATTAGATGCATGGGAAAAAATTCGGTTACGGCCAG GGGGAAAGAAAAAATATAGGATAAAACATTTAGTATGGGCAAGCAGAGAGTTAGAAAGATTCGCTCTTAA CCCTGGCCTTTTAGAAACAGCAGAAGGATGTCAACAAATACTAGAACAGTTACAGTCAACTCTCAAGACA GGATCAGAAGAACTTAAATCATTATTTAATACAATAGCAACCCTCTGGTGCGTACACCAAAGGATAGATG TAAAAGACACCAAGGAAGCTTTAGATAAAATAGAGGAAGTACAAAAGAAAAGCCAGCAAAAGACACAGCA GGCAGCAGCTGGCCCAGGAAGCAGCAGCAAAGTCAGCCAAAATTATCCTATAGTGCAAAATGCACAACCA CAAATGGTACAC 

Figure 7.3: Single DNA sequence in FASTA format

### 7.2 Data Formats

The probe generation application operates on plain text files that comply with the FASTA format [30]. Each file contains a set of sequences, and each sequence starts with a single line description, followed by lines of sequence data. Description lines always start with a > symbol. Figure 7.3 provides an example.

The probe generator produces a plain text file with the following structure. The first line contains the number of input sequences found in the original input file. As described in Section 4.1, each node in the suffix tree is associated with a binary signature that tells us which sequences contain the probes that correspond to this node. For each such set of probes the output file contains a sequence of lines, each containing a probe and its minimum length, followed by the indices of the sequences, which these probes occur in. An example can be found in Figure 7.4. The first line represents 15 distinct probes, namely all prefixes of length at least 22:

- AGCCTGGGAGCTCTCTGGCTAG
- AGCCTGGGAGCTCTCTGGCTAGC
- AGCCTGGGAGCTCTCTGGCTAGCT
- ...
- AGCCTGGGAGCTCTCTGGCTAGCTAGGGAACCCACT

The same goes for the second line. The third line tells us, that all these probes are contained exclusively in sequences 0,8 , and 19 .

### 7.3 Implementation Details

One task that we have to accomplish over and over again during the course of the algorithm is to iterate over all sequence pairs distinguished by a particular probe or, vice versa, over all probes distinguishing a particular sequence pair. To make this as efficient as possible, we have to keep this information in memory. Unfortunately, this information changes whenever we are fixing or hiding some probes. To solve this problem, we represent each probe as a dynamic vector containing references to those sequence pairs it distinguishes. Correspondingly, each sequence pair is represented as a dynamic vector containing references to those probes by which it can be distinguished. The set of all candidate probes and sequence pairs are held in globally available dynamic vectors, which are implemented as singleton objects [36].

A dynamic vector offers the possibility to hide and unhide elements at any time. It is implemented as a continuous array with all hidden elements located at the end; we maintain the index $I_{\text {hidden }}$ of the first hidden element. When hiding an element, we actually exchange it with the last active element, i. e. the predecessor of the first hidden element, and decrement $I_{\text {hidden }}$. Reactivation of a previously hidden element works accordingly. While these operations can be accomplished in constant time, the overall process takes linear time. This is due to the fact that we know the containing vector of an element but have to traverse all elements to find its actual position therein. The main disadvantage of this approach is that hiding or reactivating an element changes the order of elements in the data structure. In contrast to the implementation of the original CFT heuristic, the possibility to efficiently hide probes and sequence pairs temporarily frees us from the necessity to construct subinstances explicitly in memory. Since memory usage is an important issue, especially as instances become larger, this is a great advantage.

At any point in time, we operate on exactly one subinstance, defined by the set of hidden probes and sequence pairs. This subinstance changes

```
AGCCTGGGAGCTCTCTGGCTAGCTAGGGAACCCACT 22
AGCTAGGGAACCCACT 15
O 19
```

Figure 7.4: Sample probe specification

- during probe fixing (Section 6.2.3) and solution refinement (Section 6.2.5), where we choose some probes to be fixed as part of the forthcoming solution. In this case, we need to hide the respective probes as well as all sequence pairs that are fully distinguished by these probes. After doing so, we need to recompute the Lagrangian costs and the lower bound.
- when we initialize or update the core during the pricing procedure (Section 6.2.4). In this case, we hide some probes to make the algorithm faster. However, as we do not hide any sequence pairs, only the lower bound is invalidated and needs to be recomputed.

Contrary to the original CFT algorithm, some of the decisions inside the CFTSB heuristic are random. In these cases, we use a SIMD-oriented Fast Mersenne Twister [33] as pseudo-random number generator (PRNG). This kind of PRNG is known for its speed and provides a very high degree of equidistribution.

### 7.4 Class overview

This section gives an overview of the most important classes of our implementation. An UML class diagram depicting the inheritance relations between these classes is shown in Figure 7.5. Below, we describe briefly the capabilities of each class. Note that, out of all listed classes, class Instance is missing from Figure 7.5. Although it bears no inheritance relation to any of the other classes it is used by the majority of them, so its inclusion would have cluttered the diagram unnecessarily.

DynamicElement Abstract superclass representing objects that can be active or hidden. Objects that are to be stored in a DynamicList must be instances of a subclass of this class. Accordingly, both Probe and SequencePair derive from it.

DynamicList A vector of instances of DynamicElement or a subclass thereof. When the state of any contained element changes, either from hidden to active or vice versa, we need to call updateActive if we hid elements, updateHidden if we reactivated elements, or update if we did both. It is not until the appropriate update method has been called,that the actual masking takes effect. This class offers the possibility to iterate over either all, only the hidden, or only the active elements.


Figure 7.5: UML class diagram

Probe Derives from DynamicElement and DynamicList. Each instance of this class represents a particular candidate probe and holds a list of all sequence pairs that this probe distinguishes. Additionally provides a method to query its Lagrangian cost.

SequencePair Derives from DynamicElement and DynamicList. Each instance of this class represents a particular sequence pair. Provides methods to query the currently assigned Lagrangian multiplier and the number of probes that are required to distinguish the sequence pair. Note that this number changes whenever probes are fixed as part of the solution.

ProbeManager Derives from DynamicList. There is exactly one instance that holds a list of all candidate probes of an instance. This instance is globally available as singleton object.

SequencePairManager Derives from DynamicList. There is exactly one instance that holds a list of all sequence pairs of an instance. This instance is globally available as singleton object.

Instance This class maintains important data pertaining to the instance as a whole. This includes the required number of probes by which each sequence pair must be distinguished, the minimum and maximum probe length, and the number of sequences. Moreover, it provides methods to recompute Lagrangian costs and the lower bound. Moreover, this classes handles the parsing of input instances. There is exactly one instance, globally available as singleton object.

Sublnstance Represents a subinstance. Provides methods fixProbes, restoreFixed, hideProbes, and restoreHidden to fix, hide, or reactivate a subset of probes.

Core This class implements all aspects, that is initialization and update, of the pricing technique addressed in Section 6.2.4.

GreedySolver Implements the greedy heuristic described in Section 6.2.2.

RedProbeCleaner Implements the removal of redundant probes from a given solution as described in Section 6.2.2. A recursive method enumerates all subsets of redundant probes and eventually removes one with maximum cardinality.

LPSolver Computes an optimal solution to the current subinstance, we are operating
on. This is accomplished by constructing the corresponding ILP and solving its LP relaxation using the ILOG CPLEX LP solver.

SubgradientOptimizer This class encapsulates the implementation of the subgradient optimization procedure using the standard subgradient definition (6.5) as it is used by the greedy heuristic, which has been presented in Section 6.2.2. Method run launches the subgradient optimization procedure, whereas method step performs a single iteration only. The latter will be used when we are computing heuristic solutions (Section 6.2.2). The constructor accepts an optional Core object that will be used throughout its lifetime.

ImprovedSubgradientOptimizer Subclass of SubgradientOptimizer that employs the smallnorm subgradient definition given in Section 6.2.1.

TriPhase Encompasses the complete lower level part of the CFTSB heuristic according to its description in Section 6.2.

CFTHeuristic Realization of the top level part of the CFTSB heuristic as outlined in Algorithm 2.

## 8 Computational Results

This chapter presents the computational results obtained from a selection of benchmark instances and compares the quality of the solutions computed by the CFTSB heuristic to those of the simpler DNA-BAR heuristic [13] as well as to optimal solutions. Section 8.1 describes the benchmark instances we used, where we obtained them from or how they have been generated. The actual results are presented in Section 8.2 with a detailed discussion following in Section 8.3.

### 8.1 Instances

For our benchmark runs, we selected some of the instances on which the results presented by Rash and Gusfield [32] were based. These are derived from real-world genomic data of 4,548 different HIV strains publicly available through the GenBank [31]. Each instance class (hiv0, hiv1, hiv4) contains 25 instances. To assess the quality of the solutions produced by the CFTSB heuristic, we compared our results to the optimal solutions, which we computed by solving the corresponding ILPs. We found that the computational effort to solve these instances to optimality tends to grow as the required level of redundancy $r$ decreases. For $r=1$ we were not able to compute optimal solutions for any of the given instances within reasonable time or had to abort the runs due to excessive memory consumption. Therefore, we additionally constructed a set of smaller instances from those in instance class hiv1 - we chose this instance class arbitrarily. From each original instance, we selected a random subset of sequences, containing between 30 and $75 \%$ of the sequences of the original instance. For each instance we tried different percentages until the resultant instance could be solved to optimality in reasonable time. The resulting 24 instances constitute instance class hiv1a. All instances are available for download at http://www.ads.tuwien.ac.at/phil/mthesis/instances.tar.bz2.

8 Computational Results

| $\overline{\text { CPU }}$ | $\begin{aligned} & 2 \times \text { Intel Xeon } 5150 \\ & \text { Clock speed ... } 2.66 G H z \\ & \text { L2 cache ..... } 4 \mathrm{MBB} \\ & \text { Cores ........ } 2 \end{aligned}$ |
| :---: | :---: |
| Memory | 8GB |
| Operating System | Debian Linux 4.0 amd64 |
| Compiler | GCC 4.1.2 20061115 (prerelease) (Debian 4.1.1-21) |
| Optimization Flags | -02 -march=nocona -funroll-loops -ftree-vectorize -msse2 -mfpmath=sse -fomit-frame-pointer |
| LP Solver | ILOG CPLEX 10.0.1 |

Table 8.1: Hard- and software environment

### 8.2 Results

In addition to the results provided by the CFTSB heuristic, we computed an optimal solution for each instance using an ILP solver and compared the result to that of the simpler DNA-BAR heuristic [13]. We chose the latter primarily because out of the SB approaches summarized in Section 5.1, this is the only one for which the source code is publicly available. Due to the non-deterministic nature of the CFTSB heuristic, we performed ten runs per instance, each time with a different PRNG seed between 0 and 9. Accordingly, the reported values below are averages over these runs.

Like Rash and Gusfield, we required the generated probes to be at least 15 and at most 40 nucleotides long. This corresponds to restrictions as they are applied in practice and results in instances that can still be solved optimally. As noted in Section 5.1, requiring a minimum pairwise edit distance greater than one increases the number of candidate probes to an extent that makes computing an optimal solution impossible. For this reason, we refrained from doing so.

For details on the used hard- and software, we refer to Table 8.1. Note that, though the benchmark runs have been performed on a multi-core machine, our implementation uses only one core at a time. The computational results are presented in Tables 8.2 through 8.6. The column labels and their meaning are:

| seqs | number of sequences |
| :--- | :--- |
| len | average sequence length |

```
pgen CPU time in seconds spent for probe generation
opt size of the optimal solution
dbar size of the solution produced by the DNA-BAR heuristic
cft }\mp@subsup{}{}{0}\quad\mathrm{ size of the first solution encountered by the CFTSB heuristic
cft* size of the best solution found by the CFTSB heuristic
```

We restricted the maximum run time for each CFTSB run to one hour. For each solution, we also provide the CPU time spent to compute it (hours, minutes, and second, separated by colons) and, in case of the heuristic solutions, the relative distance to the optimal solution in percentages (gap). Given the size $\omega^{*}$ of an optimal solution, the gap for a heuristic solution of size $\omega$ is given by $\left(\omega-\omega^{*}\right) / \omega$.

We also provide the standard deviation (column sdev) for the CFTSB results. The last row of each table shows the mean, or median in case of run times, over all instances of a class.

Figures 8.1 through 8.5 provide a more visual comparison of the quality of the solutions computed by each algorithm. Each bar represents a single instance, and its total height corresponds to the size of the solution produced by the DNA-BAR heuristic. Each algorithm is associated with a different color, black for the optimal solution and different shades of gray for the other algorithms. The size of a solution produced by an algorithm is given by the total height of the part of the bar with the color that corresponds to the respective algorithm and all the parts below that. Due to this layout, it is easy to see that the the following relation holds for all results: opt $\leq \mathrm{cft}^{*} \leq \mathrm{cft}^{0} \leq \mathrm{dbar}$.

### 8.3 Discussion

It is noticeable that the solutions generated by the CFTSB heuristic are less than satisfactory if we require a redundancy greater than one. For many instances, finding the best solution takes longer than computing the optimal solution by solving the ILP. The average gap for benchmark runs with a redundancy greater than one is about $2 \%$. With the exception of instance class hiv4, redundancy 2 , the median run times for computing optimal solutions lie between 44.99 and 187.97 seconds. For the same instances, the median run times of the CFTSB heuristic lie between about 10 and 20 minutes, which is nearly 4 to 26 times as much. Moreover, the algorithm reaches an optimal solution
Table 8．2：Results for instance class hiv0，probe length $\in[15,40]$ ，redundancy 4

| モ1： $27: 07: 0$ | $\mathrm{LO}^{2} \mathrm{z}$ |  | も「66 | 86：00：00：0 | 02＇t |  | 9．8LI | 87＊00：00：0 | 89： 29 | 8¢9］ | 66 柆：00：0 | も゙26 | $87^{\circ} 0$ | \＆ 696 | ち「16 | $\varnothing$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 06：90：87：0 | $00^{\circ}$ | ¢9 0 | 8． 16 | モ6．10：00：0 | 00＇61 | $8 L^{\prime} 7$ | I．20］ | 79：00：00：0 | 68.89 | z¢L | モ9：97：00：0 | 06 | モE\％ | ¢ ¢08 | 06 | $\ddagger 2$ |
| 80＇29：0T：0 | $87^{\prime}$ I | 7ャ．0 | \％$¢ 6$ | LF＇00：00：0 | L6．91 | 88.7 | 6.605 | LZ：00：00：0 | $61^{\prime} \mathrm{E}$ ¢ | もも | てФ＇ $70: 00: 0$ | モ6 | LI＇0 | $0 \cdot$ L99 | 12 | ¢z |
| LL $68: 6 \mathrm{E}: 0$ | 76．I | で0 | 8＇t6 | \＆8＇ $10: 00: 0$ | 97．6I | $99^{\prime} 7$ | 6001 | 87＊00：00：0 | L\＆： 79 | LGL | 68＇も！ $70: 0$ | 86 | ¢L．0 | \＆919 | 76 | ъ\％ |
| モ6 20：08：0 | 86． 1 | もL．0 | 6.26 | 87＊00：00：0 | 20．81 | 96.7 | \＆ELI | 67：00：00：0 | 2F09 | モ¢L | 80＇LZ：00：0 | 96 | 87\％ | L．999 | 08 | LZ |
| 70．ce：ct：0 | $67^{\prime}$ I | 7ャ0 | Z＇モ6 | 88：00：00：0 | 98．2I | 89＇8 | 960 L | 88：00：00：0 | $69^{\circ} 9$ | モ¢L | 9000：90：0 | 86 | て，${ }^{\text {co }}$ | L＇879 | 88 | 07 |
| \＆L．0モ：ti：0 | 78.7 | $8 \mathrm{~F}^{\circ}$ | \＆ 10 L | 07＇10：00：0 | 98．9］ | $00 \cdot 8$ | I＇GzI | 7I＇10：00：0 | 99－99 | ¢91 | 78＇IL：07：0 | 66 | 97.67 | 7•8908 | 86 | 6I |
| 80＊8L：LI：0 | 20． | 7ャ．0 | z＇eLI | 㕵 $10: 00: 0$ | 78．61 | 91＇¢ | て＇も¢L | 89．70：00：0 | 70．89 | LLI | L2＇0才： $\mathrm{ZI}: 0$ | ZLI | \％．0 | \＆：209 | 001 | 8I |
| L98¢：6］：0 | GL＇ | 9ti | 6．201 | 87＇ T ：00：0 | 96．0z | $0 L^{\prime} 7$ | 8．GGI | 8600：00：0 | L7：89 | GLI | 26．98：8г：0 | п0I | 9＊8 | $6 . \mathrm{ZLLZ}$ | 66 | LI |
| 06 㸬：0¢：0 | $80^{\prime} 7$ | 20． | 9 20I | 01＇70：00：0 | L861 | 99.8 | 8．GZI | 90＇ L ： 000 ：0 | ¢069 | 291 | 06．9z：LZ： | 901 | ¢\％ 0 | 7．708 | 801 | 9I |
| ¢6： $28: 7 \%: 0$ | 82＇I | $85^{\circ}$ | 8＇tL | 62：00：00：0 | 98．61 | 09＇8 | 9.28 | 09：00：00：0 | 09\％ 72 | 97I | 8L＇¢t：0：0 | \＆ 2 | ¢9．9 | ［•820］ | 82 | GI |
| 8L＇90： LZ ：0 | ¢\＆\％ | 76.0 | 7：96 | 88：00：00：0 | 28．7\％ | 18＇8 | ¢ ¢LI | 67：00：00：0 | 97＇も¢ | 9 CL | 66．97：TG：0 | モ6 | 81．0 | ¢9\％9 | 76 | もI |
| 70．09：tito | モ\＆\％ | 76.0 | 7．96 | 86．00：00：0 | 08：8\％ | 29.7 | 9．07I | 99：00：00：0 | 62．82 | 891 | \＆\＆＇Lも：00：0 | モ6 | 88.0 | 6092 | 76 | ¢I |
| 81．07：07：0 | $6 \square^{\circ} \mathrm{I}$ | LI＇I | も「96 | LL＇00：00：0 | L6．LZ | $9^{9}{ }^{\text { }}$ | 9 9tL | 理00：00：0 | L969 | $0 ¢ \mathrm{~L}$ | 8¢゙ 7 L ：00：0 | モ6 | 71．0 | ［＇099 | 98 | 2I |
| \＆L：LG：LZ：0 | $8{ }^{\circ} \mathrm{T}$ | LL． 0 | G．got | 0＇10：00：0 | $09^{\prime} ¢ 8$ | $27^{\circ} \mathrm{C}$ | Z＇LZI | L9：00：00：0 | 70．99 | ILI | モ0 7 \％：00：0 | 801 | 80 | LELL | ¢6 | LI |
| 68．LS：6T：0 | Lq．t | 02.0 | も＇も6 | 91． T ：00：0 | ¢f：0z | 91＇$¢$ | 0 0LI | 98：00：00：0 | L8＇88 | LLI | で「0さ：も0：0 | ¢6 | $85^{\circ} 0$ | 7 ＇189 | モ6 | 0I |
| 29＇70：97：0 | $88^{\prime} 7$ | 26.0 | も．LZI | 92：00：00：0 | 68.8 | 92.8 | 9 9\％t | 79：00：00：0 | 99.64 | 2LV | £8．68：00：0 | 8LI | 61．0 | 6．8L9 | 26 | 60 |
| L2：6I：97：0 | $69^{\circ}$ | $67^{\circ} \mathrm{L}$ | 600I | 㛡 $70: 00: 0$ | 07：97 | 07＇$¢$ | \＆98L | 08． $10: 00: 0$ | \＆889 | LLI | L2．0才：67：0 | 801 | 82．81 | $0 \cdot \mathrm{8c} \mathrm{\%}$ | 0LI | 80 |
| 98＇LE：7\％：0 | $60^{\circ} \mathrm{I}$ | 㕵 0 | I＇70］ | 76：00：00：0 | ［9．81 | $89 \cdot 8$ | 8．6LI | ¢7．00：00：0 | 67＇TL | ELI | 67＇LI：00：0 | L0I | $7 \times 0$ | ¢＇も¢9 | 88 | 20 |
| 89：0才： $20: 0$ | GZ ${ }^{\text {I }}$ | 76.0 | 7： 26 | ［6．00：00：0 | \＆2：0］ | ¢9． 7 | $6.9 L I$ | 98：00：00：0 | 81．8L | LLI | 72：9\％：00：0 | 96 | LI＇0 | ［＇t99 | 76 | 90 |
| 99＇ti：67：0 | 96 ${ }^{\text {I }}$ | $88^{\circ}$ | 6.86 | LI＇L0：00：0 | 87．65 | L\＆： | $6 \varrho[\mathrm{~L}$ | 67＊00：00：0 | 01．02 | 991 | 軞とて：00：0 | 26 | $87^{\circ} 0$ | 8．889 | 96 | 90 |
| ठL＇L0：¢L：0 | $00^{\circ}$ | 76.0 | 8．16 | 08：00：00：0 | 68.78 | L゙ 7 | 90LI | 07．00：00：0 | 7\％： 72 | g¢L |  | 06 | 980 | c．et8 | 88 | ¢0 |
|  | $L Z^{\prime}$ I | 96.0 | \＆¢0I | 88：00：00：0 | 76.88 | L\＆： | ¢．97I | 67＇00：00：0 | ¢9．89 | ZLI | 997\％I：00：0 | z01 | 970 | ［＇969 | 88 | ¢0 |
| L6 $2 \mathrm{~L}: 8 \mathrm{~T}: 0$ | 98＇1 | $¢^{\circ} 0$ | 9．78 | 玫00：00：0 | Lz＇\＆z | ¢0＇z | 8.66 | 珧00：00：0 | 99：28 | Z¢T | 69：00：10：0 | L8 | $90^{\circ}$ I | ［＇88¢L | 92 | z0 |
| ¢1：2\％：07：0 | 76．I | 78.0 | 0 0 901 | 70＇10：00：0 | LI＇もを | $20^{\circ} \mathrm{Z}$ | 6.2 LI | 7ヶ＊00：00：0 | 28．02 | 92I | 89\％ $\mathrm{cq}: 00: 0$ | ¢01 | 61＇0 | 8．929 | 96 | L0 |
| \＆LLz：¢I：0 | 78．7 | 66.0 | 6． 001 | 78． $10: 00: 0$ | ¢9 \＆ | 76.7 | ［＇97I | 7¢00：00：0 | \＆L＇¢9 | 291 | 现 $77: 00: 0$ | 701 | $77^{\circ} 0$ | 72¢9 | 96 | 00 |

[^0]

| モ¢ 99：7I：0 | もI＇z |  | 0 zc | 80＇ $10: 00: 0$ | $7 \% \cdot 7 \%$ |  | 7．79 | 9800：00：0 | 69． 19 | \＆ 78 | L6． $20: 80: 0$ | 6．09 | 7．0 | 7＇889 | \＆ 68 | $\varnothing$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| L97才：It：0 | 72＇I | 02．0 | 961 | －T0：00：0 | L＇93 | I6T | 79 | 7600：00：0 | ç＇29 | 78 | 8： 2 ： $10: 0$ | 67 | \＆0 | 0.924 | L6 | ஏて |
| GL̇®：gi：0 | LI＇Z | $8 \square^{\circ} 0$ | \＆＇L9 | 96：00：00：0 | 88＇61 | 81＇\％ | 6．12 | 0屯 $00: 00: 0$ | 29＇LL | 801 | L゙9 GI：00：0 | 09 | ¢¢ 0 | 0＇708 | 76 | 86 |
|  | 09.7 | $29^{\circ}$ | \＆¢ $¢$ | 98：00：00：0 | 00＇9\％ | LZ＇\％ | $0 \cdot 99$ | L800：00：0 | 97＇¢9 | 98 | ¢1．97：00：0 | \％9 | $61^{\circ} 0$ | †¢¢¢9 | 98 | ヶ7 |
| 7920：It：0 | Ə0＇ 8 | $85^{\circ} 0$ | 2．29 | 98． $10: 00: 0$ | L¢ ¢ \％ | 78． | 769 | 68：00：00：0 | 98＇29 | モ6 | ธ6．$¢ 8: 00: \%$ | 99 | $87^{\circ} 0$ | も． 769 | 86 | Lz |
| 87＇87：6I：0 | \＆6\％ | 衡I | 7．89 | LL• $10: 00: 0$ | ちt．26 | 01＇\％ | Z＇tL | 79：00：00：0 | 8689 | 68 | 29．87：I0：0 | 99 | $88^{\circ}$ | c． 972 | モ0 | 07 |
| 0才¢80：8г：0 | LG＇L | 62.0 | 8．89 | 7ヶ＇10：00：0 | 98： 2 z | $89^{\circ}$ | 9.29 | 9900：00：0 | 91＇t9 | 28 | L6：20：80：0 | 89 | $\downarrow^{\circ} 0$ | ［．0¢8 | 66 | 6I |
| モ6 ¢z：¢г：0 | 69.7 | L2．0 | 969 | 99． $10: 00: 0$ | ¢6． 76 | $9 L^{\prime} \mathrm{Z}$ | \＆ 12 | LS．00：00：0 | 8t＇ゅ8 | 20I | 67＇90：80：0 | 89 | $28^{\circ} 0$ | \＆．878 | 96 | 8 I |
| 81＇そI：9¢：0 | 0才＇L | 78．0 | LOS | 0¢：00：00：0 | 00＇61 | $89^{\circ} \mathrm{L}$ | ¢ 69 | LZ：00：00：0 | 00．$\ddagger$ g | LL | ¢9．68：00：0 | 09 | $80^{\circ} \mathrm{I}$ | ［＇086 | ¢8 | LI |
| モ¢ 99：zi：0 | 9 St | 9［＇I | \＆＇ts | 玌：00：00：0 | 02．LZ | 99 ${ }^{\text {I }}$ | Q＇t9 | 0ヵ00：00：0 | 09：99 | ¢8 | 98．00：［I： 0 | ¢9 | $6 \mathrm{I}^{\circ} 0$ | L＇8z9 | 68 | 9 I |
| 89－98：9さ：0 | 08＇z | $58^{\circ}$ | も．ts | 81•10：00：0 | 0才゙もを | 78． | \％＇79 | 07：00：00：0 | 00＇z9 | L8 | 68．8¢：07：$¢$ | 09 | LI．0 | \＆゙tis | L6 | ¢I |
| 78： $28:$ zI：0 | 70＇L | $8^{\circ} 0$ | 967 | Lz＇ $10: 00: 0$ | 71．9\％ | 比 ${ }^{\text {c }}$ | 8＇L9 | 8®00：00：0 | も゙29 | 12 | 08 80：68： | 67 | 67＊ 0 | 6．806 | 86 | ¢I |
| LI＇もの：あt：0 | 08． I | EL＇0 | 6．09 | 86：00：00：0 | 0才＇93 | $8{ }^{\prime}$ I | 7＇¢9 | 8800：00：0 | 00TS | 12 | 97＇87：00：0 | 0 S | ¢9 $9^{\text {I }}$ | İ29LI | 68 | \＆1 |
| 29＇もて：0さ：0 | 99 ${ }^{\text {z }}$ | $85^{\circ} 0$ | \＆09 |  | 71．9\％ | $89^{\circ} 7$ | 8＇t9 | L0＇T0：00：0 | 81．69 | 82 | 90．89：モた：0 | 67 | 80 | I＇m92 | ¢6 | ZI |
| 6999：c0：0 | 80 ¢ | $7 \square^{\circ} 0$ | 7．0才 | 86：00：00：0 | 形して | 96.7 | 2．67 | 7\％：00：00：0 | 76．92 | 69 | $68 \cdot 18: 90: 0$ | 68 | $7 \%$ | 68.8 | 92 | II |
| 69 69：90：0 | ঢz＇z | L9．0 | ［09 | 79：00：00：0 | ¢t CZ | LL＇ Z | $0 \cdot 09$ | 97：00：00：0 | 70＇ts | UL | L8． $8 ¢: 00: 0$ | 67 | 70 | 8．729 | 18 | 0I |
| LI＇90：85：0 | $60^{\circ}$ I | 02．0 | 9．99 | 站 $70: 00: 0$ | 78．21 | 28． | 8＇79 | モ8：00：00：0 | ¢9：¢7 | 62 | LF＇gI：gg：z | g9 | $2 \mathrm{~L}^{\circ}$ | 0 ¢ 81 | 26 | 60 |
| LEGG：8L：0 | 92． L | 780 | 6．ts | モz＇00：00：0 | 50．81 | $9 L^{\prime} \mathrm{Z}$ | 7．09 | ZI＇00：00：0 | 68．08 | 76 | 86：20：00：0 | ts | $80^{\circ} 0$ | 6.827 | \＆2 | 80 |
| ［9887：90：0 | 70＇8 | 69．0 | 9． ¢ $^{\text {c }}$ | 97＇10：00：0 | 61．0z | $68^{\prime}$ I | L＇89 | 97：00：00：0 | L869 | 06 | 67．90：61：0 | ¢ 9 | Z $\mathrm{I}^{\circ} 0$ | L＇tzs | 76 | 20 |
| 86．67：90：0 | $86^{\circ} 0$ | 78.0 | ［＇87 | 28：00：00：0 | 8691 | 02． 1 | 809 | 87：00：00：0 | もじ89 | 89 | でったで0さ：0 | ¢1 | 2I＇0 | 0.989 | 18 | 90 |
| 7\％＇18：90：0 | $00 \cdot 8$ | $\mathrm{Ec}^{\circ} 0$ | G．LG | 80＇ 0 0：00：0 | 00＇ 77 | モ6． | 0＇19 | 07：00：00：0 | 00TS | LL | 98： $27: 00: 0$ | 09 | 77\％0 | $8 \cdot 869$ | 68 | 90 |
| くも゙も0：9さ：0 | $00 \cdot 8$ | L 20 | G．LS | 07＇10：00：0 | 09＇9z | $0 I^{\prime} \%$ | 8＇79 | 7¢ 00：00：0 | 00\％09 | 92 | 9\％＇te：90：L | 09 | $67^{\circ} 0$ | 9．87L | 06 | †0 |
| เ\＆゙てた：もた：0 | ¢9＇z | z9．0 | も现 | 89．70：00：0 | 80＇ 76 | LL＇Z | L＇79 | モ8：00：00：0 | ¢T＇t9 | 28 | 9］．9c：$¢ 8: \ddagger$ | ¢S | ［＇0 | 9．87t | 801 | 80 |
| 08¢を¢ | $68^{\circ} 0$ | 79．0 |  | 69．00：00：0 | 29．91 | 82＇L | 9.69 | 9800：00：0 | LI＇ts | 89 | 09＇ts：70：0 | Gt | モ゙．0 | ¢ 929 | 92 | 70 |
| L\＆GL：Lu：0 | $89^{\circ} 0$ | $29^{\circ} 0$ | \＆ 7 c | LC：00：00：0 | 8891 | 02． 1 | $0 \cdot 09$ | 07：00：00：0 | 97＊$¢ 9$ | 98 | 10 68：00：0 | 6s | L＇0 | ［．687 | \＆8 | L0 |
| 92：07：c0：0 | $67 \cdot 7$ | 78\％0 | ［67 | 68：00：00：0 | L8＇もI | ［6．1 | $6 . ⿰ ㇒ ⿻ 土 ㇒ 𠃋$ | 97：00：00：0 | 80．79 | \＆ 2 | 89： ES ：tI：0 | 87 | ［＇0 | \＆＇767 | 28 | 00 |
| эш！7 | des | ләр | サフ | әu！ | des | ләр | ${ }^{\text {¢¢ }}$ | әш！ | des | тес | әи！ | 7 ¢ | иәภ¢ | นəІ | sbəs | 7 Su |



Figure 8.1: Results for instance class hiv0, probe length $\in[15,40]$, redundancy 4


Figure 8.2: Results for instance class hiv1, probe length $\in[15,40]$, redundancy 2


| 89：98：7I：0 | 96．${ }^{\text {I }}$ |  | G．LG | 71＇10：00：0 | ¢ $¢$ ¢ L |  | \＆ 19 | 78：00：00：0 | 切＇L9 | 9•18 | 7\％＇LG：LI：0 | 9．09 | LZ $\%$ | 6.15 | 8．06 | $\varnothing$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | 07＇0 | 78\％0 | ［1． | 6．00：00：0 | ¢¢¢ | 㖪 | ¢9 | －00：00：0 | 29.99 | ¢8 | ¢9＇28：80：0 | IS | 78.0 | ［＇802 | 86 | ஏて |
| LI＇9\％：70：0 | $00 \cdot 0$ | $00^{\circ} 0$ | $0 \cdot 87$ | Le．00：00：0 | LI＇ti | $66^{\text {I }}$ | － IS | 91．00：00：0 | 97．99 | GL | 90 61：70： | 81 | 20.0 | c．tLt | 82 | 86 |
| モ¢ 0i：80：0 | $00 \cdot 0$ | $00^{\circ} 0$ | $0 \cdot$ 䛔 | 88：00：00：0 | EL2I | L8＇I | 8．19 | 07：00：00：0 | L6．c9 | \＆2 | 07＊07：00：0 | 形 | $9{ }^{\circ} \mathrm{O}$ | \＆－0¢9 | LL | ヶ7 |
| 2669：90：0 | モ0＇z | LT0 | 009 | 20＇10：00：0 | L9．81 | L8． | ［．89 | 1800：00：0 | もし29 | LL | 90 6\％：L0：0 | 67 | LT＊ | 0．8L6 | 18 | Lz |
| 01．0才：85：0 | L8．7 | LI＇I | 9.89 | 轨 $70: 00: 0$ | 98＇8z | 2I＇\％ | 9．02 | 72：00：00：0 | て「＇6 | 98 |  | 29 | ¢0＇I | 9676 | 601 | 07 |
| L\＆ $7 \mathrm{I}: 87: 0$ | \＆9．$¢$ | で0 | 8．7s | 98：00：00：0 | ちで8\％ | $9{ }^{\prime}$ | †＇99 | 0才＇00：00：0 | GL＇79 | ¢8 | 61＇98：I0：0 | IS | LI＇0 | \＆8G9 | 28 | 6I |
| 9L：67：97：0 | 98.0 | z9．0 | もくも | 99：00：00：0 | 99 L 1 | Lz． 7 | \＆ 99 | 96：00：00：0 | Sti29 | UL | 1095：00：0 | 27 | $9 \mathrm{~T}^{\circ} 0$ | $9 \cdot 979$ | 82 | 8I |
| \＆г ¢¢：zг：0 | 79．t | で0 | 8＇t9 | 69＇10：00：0 | ¢7＇8z | LI＇Z | ¢．92 | 0600：00：0 | 7下゙LL | 0LI | 70： $29: 2 ¢: 0$ | 79 | $89^{\circ} \mathrm{G}$ | G．ct9 | $\mathrm{c}_{0}$ | LI |
| 98：08：80：0 | 78．${ }^{\text {I }}$ | 80＇${ }^{\text {I }}$ | 8＇玵 | 89＇ 0 0：00：0 | LZ＇LI | L2．L | 9．t9 | L800：00：0 | 6069 | 02 | 7L：CZ：00：0 | 比 | ¢7\％0 | L＇t29 | 06 | 9 I |
| 6L＇もて：¢0：0 | 28＇ | 780 | 6.87 | 08： $10: 00: 0$ | \＆8＇SI | 96.1 | 9．gs | 97：00：00：0 | 762\％ | LL |  | 87 | 81．0 | 7．929 | 88 | ¢I |
| LLG¢：もt：0 | $00^{\circ}$ | 290 | 0．ts |  | 0才＇9\％ | LI＇z | L＇79 | L800：00：0 | 00TS | 12 | 09．モ¢：0才：0 | 09 | $67^{\circ} 0$ | 8.902 | 96 | ¢I |
| 7¢ 9¢：z\％：0 | $67^{\prime}$ I | $29^{\circ} 0$ | 2： 27 | 7I•10：00：0 | \＆G＇96 | 02＇I | 069 | 87：00：00：0 | 09.92 | ¢8 | ET－88：0¢：もI | $L$ | 61．0 | ［＇989 | 98 | \＆1 |
| LS＇LZ：LI：0 | 81 ${ }^{\circ}$ | 6ャ0 | 8．99 | 97＇70：00：0 | 96． L \％ | \＆$\square^{\circ} \mathrm{E}$ | \＆ 89 | L9：00：00：0 | もじ29 | 88 | 6868：Lも：0 | 99 | $88^{\circ}$ | モ．092 | 901 | ZI |
| 80 ¢t： $\mathrm{tu}=0$ | 08： | 20＇I | 9．97 | 86：00：00：0 | \＆8＇ъ7 | 82＇ | ¢ 99 | 97：00：00：0 | 02＊89 | \＆ 2 | ET＇Lt：I0：0 | 97 | $87^{\circ} 0$ | C．782 | \＆8 | LI |
| 顽 $78: 0 \mathrm{c}: 0$ | gq． 0 | $87^{\circ} 0$ | \＆Gg | LL $2 \cdot 70: 00: 0$ | LZ＇LZ | $89^{\prime}$ | L．99 | 0才•00：00：0 | St＇99 | L6 | 20＇98：Lも： 77 | g 9 | ¢T．0 | $0 \cdot 999$ | I0I | 0I |
| 098を： 0 0：0 | \＆z＇0 | 780 | ［＇Et | ［9．00：00：0 | 02．0z | $67^{\prime}$ I | 6． LS $^{\text {c }}$ | 87：00：00：0 | 乙T＇c9 | L2 | モ0：0币：00：0 | ¢t | E\％\％ | \＆ 862 | GL | 60 |
| 27：08：20：0 | L2： 7 | $00^{\circ}$ | $0 \cdot 9 \pm$ | LL＇00：00：0 | cq．tz | GL＇I | 8．ts | 97：00：00：0 | 6069 | 02 | モでもた：80：0 | 切 | $61^{\circ} 0$ | \＆ 689 | 62 | 80 |
| 67＇9\％：\％I：0 | 68\％ | 780 | ［．27 | L0＇L0：00：0 | 97＊81 | 98． | も゙も¢ | 08：00：00：0 | $28^{\circ} 09$ | UL | 97－88：6¢：0 | 97 | \＆1．0 | 7＇999 | 28 | 20 |
| 08 Lて：ちで0 | Et＇I | $87^{\prime}$ | 8.99 | 28：00：00：0 | $6 z^{\prime} \downarrow$ \％ | $29^{7}$ | 9.69 | 6800：00：0 | 6でゅ9 | 76 | 77＇TG：LI：0 | 99 | $87^{\circ} 0$ | 9．899 | 26 | 90 |
| 09－00：07：0 | $79^{\circ}$ | $87^{\circ} 0$ | \＆ 81 | 76：00：00：0 | Zだ¢ | $89^{\circ} \mathrm{L}$ | － 99 | 7\％：00：00：0 | LI＇tg | 焐 | 61 ¢ \％$\% 7 \%: 0$ | 87 | $9{ }^{\circ} \mathrm{O}$ | L689 | 98 | 90 |
| 7L＇77：85：0 | 89＇z | L20 | 9.89 | 71． 0 0：00：0 | 02．0z | 80＇I | 8.89 | 9800：00：0 | 89\％\＆ | 66 | LI＇88：80：0 | L9 | モ\＆ 0 | 9．28L | 26 | †0 |
| 87゙LもことI：0 | $67^{\circ} \mathrm{E}$ | 6200 | \％$¢ 9$ | 61＇70：00：0 | \＆1＇もz | $9^{\prime}{ }^{\prime}$ | 7.84 | 09：00：00：0 | L6EG | 26 | モ0：9\％：89：96 | ¢9 | $9 \mathrm{c}^{\circ} 0$ | †．98¢ | 9LI | 80 |
| ¢\＆20：モ0：0 | 91＇z | 780 | I＇zs |  | $69^{\circ} 0 z^{\circ}$ | T9．${ }^{\text {L }}$ | $\mathrm{C}^{\text {¢ }}$ ¢ 9 | 99．00：00：0 | GL．79 | ¢8 | 96．67：20：0 | LS | $85^{\circ} 0$ | L．976 | $\mathrm{c}^{6}$ | 70 |
| 962L：EL：0 | てギも | 90＇ I | \＆${ }^{\text {¢ }}$ | LS． $10: 00: 0$ | t¢ $7 \%$ | LI＇8 | $9 \cdot 89$ | 9800：00：0 | LLGg | L8 | 67－9さ：97：も | \％ 9 | $7^{\circ} 0$ | 8：619 | モ6 | L0 |
| 89 98：zI：0 | \＆1＇$¢$ | $80^{\circ} \mathrm{I}$ | 967 | 72＇80：00：0 | 79．07 | L6．${ }^{\text {I }}$ | 6.29 | 7800：00：0 | 76．72 | ¢8 | 67＇\＆z：97：0 | $8 \pm$ | LZ 0 | 6.799 | 76 | 00 |
| әu！7 | des | ләр | サフ | әu！ | des | әр | ${ }^{\ddagger \ddagger}$ | әu！ | des | тес | әш！ | 7 do | นә．ภ¢ | นə | sbas | 7 Su |



| 98：9t：6 | 81＇ |  | $0 \cdot 66$ | 61＇L0：00：0 | 99｀ L \％ |  | 62LI | 67＇00：00：0 | ¢9 99 | ¢09I | 90＇10：10：0 | 6.96 | $77 \%$ | 6.75 | 8．06 | $\varnothing$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| ， | ¢9 | $26^{\circ}$ | L0L | 86．00：00：0 | ＇12 | $6 L^{\prime} 7$ | 07L | モ9：00：00：0 | LLOL | 691 | 00：0 | 66 | 780 | 802 | ¢6 | ， |
| 9629：80：0 | 92＇I | 78．0 | $9^{\prime} 76$ | LF＇00：00：0 | 88＊0 | LL\％ | 000 L | 97：00：00 | ¢¢ $9^{\circ}$ | LSL | 09 \＆L：00 | L6 | $80^{\circ}$ | G＇tL | 82 | ¢\％ |
| 69．28：6I：0 | モz＇0 | 270 | 7．98 | 8700：00：0 | $88 \cdot 81$ | $86^{\circ}$ | 8.96 | L800：00：0 | 88＇99 |  | 97＇10：00：0 | 98 | LI＇0 | 8．079 | IL | z7 |
| 18：99：cz：0 | 09＊ | L2．0 | q．96 | 1800：00：0 | 88．81 | $90^{\circ}$ | L．til | 9900：00：0 | $68 \cdot 79$ | get | L9．80：00：0 | モ6 | $85^{\circ} 0$ | 0．8L6 | t8 | Lz |
| 98：80：07：0 | Gt ${ }^{\text {d }}$ | $85^{\circ}$ | L．zil | 68：70：00：0 | c9．7\％ | 78.8 | 8＇もEI | 60＇L0：00：0 | \＆L＇z¢ | 891 | 90＇LZ：I0：0 | 0LI | L0＇ | 9.676 | 601 | 07 |
| 10＇70：8¢：0 | 90＇t | 01．${ }^{\text {I }}$ | 6.66 | 82：00：00：0 | 09．7\％ | 79.8 | L2II | 79：00：00：0 | で゚09 | ¢¢L | －6：07：00：0 | 96 | LI．0 | 8899 | 28 | 6 I |
| 7991：0工：0 | II＇I | $00^{\circ}$ | 0＇L6 | 6¢00：00：0 | 牫して | 0才＇$¢$ | \＆60L | L゙00：00：0 | 68.89 | \＆もI | 0T＇9\％：00：0 | 06 | 2.0 | 9.979 | 82 | 81 |
| 29：80：78：0 | $9 L^{\prime}$＇ | 90＇I | \＆\＆6I | 82＇L0：00：0 | 80＇27 | 90＇t | Grgi | 97＇ L O：00：0 | LI＇t9 | L6I | 90＇88：70：0 | 07I | L9 9 | c．ctic | 901 | LI |
| LE2も：60：0 | モ\＆：0 | $85^{\circ} 0$ | \＆ 28 | L8＇10：00：0 | LT．gI | $90 \cdot 8$ | 7．00I | 97＊00：00：0 | 7\％¢ 8 | てロ1 | 90＇t0： 0 0：0 | 28 | ঢで0 | L＇L29 | 06 | 9 I |
| 8667： $20: 0$ | Li＇0 | 78\％ | I＇96 | Et $50: 00: 0$ | 68＇LI | $68^{\text { }}$ | \＆90I | 0才•00：00：0 | てギ8t | Lも | ¢7：67：00：0 | ¢6 | LI＇0 | 7．929 | 88 | S |
| 89－8\％：もち：0 | 29． | 02.0 | 9：26 | 89＇ $10: 00: 0$ | \＆1．¢z | $77 \cdot 8$ | 7．8LI | 67＊00：00：0 | LL29 | ［9］ | 1869：zI：0 | 96 | $67^{\circ} 0$ | 8.902 | 96 | もI |
| 69．2¢：97：0 | ¢9．${ }^{\circ}$ | $62^{\circ}$ | \％＇68 | 80＇ 0 0：00：0 | モ9＇7\％ | LZ＇\％ | L．90I | 佗00：00：0 | LZ＇98 | 791 | ¢モ．$¢ 7: 67: 0$ | 28 | 7．0 | ［＇989 | 98 | EL |
| 002も：0¢：0 | $00^{\circ} \mathrm{S}$ | $78^{\circ} 0$ | \＆Lil | 66． $10: 00: 0$ | Lも゚0E | $07 \cdot 8$ | \＆88I | 9600：00：0 | 7629 | 8LI | ¢\％＇LS：70：0 | 901 | $28^{\circ} 0$ | ¢ 092 | 90 I | ZI |
| 00＇t0：g¢：0 | 81＇z | 520 | 6．88 | 6I＇t0：00：0 | モ6．もz | 86.7 | L801 | 75：00：00：0 | L゙でてL | OSL | 7¢ ¢¢：80：0 | 28 | $87^{\circ} 0$ | G．78L | ¢8 | II |
| LZ：65：98：0 | LZ：\％ | $85^{\circ} 0$ | 8．901 | 28＇70：00：0 | 09 77 | LE\％ | G．LZI | 89：00：00：0 | L\＆ 29 | TLI | 98：$て: 00: 0$ | 01 | $9{ }^{\circ} 0$ | 0.999 | L0I | 0I |
| 989\％：6I：0 | 98.0 | 2900 | L．78 | L9：00：00：0 | 97．91 | $98^{\circ} \mathrm{I}$ | $9 \cdot \mathrm{G} 6$ | 穏00：00：0 | LTEL | 7．1 | LZ＇60：00：0 | 78 | $97^{\circ}$ | \＆ 88 | GL | 60 |
| も¢9¢：gi：0 | Et＇I | ¢9＊0 | \％＇98 | 08：00：00：0 | 埇も\％ | 96. | 9701 | 68：00：00：0 | 8才＇ 99 | 681 | 90．97：00：0 | ¢8 | $6{ }^{\circ} 0$ | \＆ 689 | 62 | 80 |
| もで6も：¢卬：0 | 68.7 | $29^{\circ}$ | ［＇06 | 20＇10：00：0 | St ${ }^{\circ} 07$ | $89^{\prime}$ | 0．90工 | 7¢00：00：0 | 09＇79 | ¢tI | 70：07：9¢：6 | 88 | 2I．0 | $7 \times 99$ | 28 | 20 |
| LI＇ṫ：67：0 | Et＇ 8 | 96.0 | L．til | 8600：00：0 | 82： 27 | $89^{\circ} 7$ | 088E | 99：00：00：0 | 68.89 | LLI | で0 07：00：0 | 801 | $97 \%$ | 9．899 | 26 | 90 |
| ¢5：87：80：0 | $60^{\text { }}$ | $00^{\circ}$ | $0 \cdot 86$ | 7600：00：0 | 81＇tI | ¢ $\mathrm{C}^{\circ} \mathrm{E}$ | 0＇coi | 88：00：00：0 | 96． 59 | 6 tI | £0．cg： $0: 0$ | 76 | $9{ }^{\circ} \mathrm{O}$ | L689 | 98 | 90 |
| 98＇78： 2 ［ 0 | 98.7 | LI＇I | 972IL | IS L0：00：0 | ¢9\％0z | $89^{\prime}$ | L＇78I | 09：00：00：0 | L6．08 | 661 | 98＇¢L：00：0 | 0LI | ¢80 | 9.282 | 26 | Ђ0 |
| 89＇LI：¢Z：0 | 96.8 | E\＆${ }^{\prime}$ | L．8GI | 89＇70：00：0 | $87^{\circ} 27$ | 02＇I | L．LSI | 88：00：00：0 | 97．99 | 981 | 08＇LE：90：0 | 6IL | 9T\％ | †98¢ | ¢IL | 80 |
| 乙I＇zг：\％\％：0 | 71＇z | LS＇0 | L＇t0I | 89． $10: 00: 0$ | 86： 27 | 7ヵ． | L．97I | 68：00：00：0 | モ9＇¢9 | 791 | ¢6． $98: 10: 0$ | 66 | LTO | L．976 | 96 | 70 |
| て1－8て：もち：0 | 06.7 | 㛡0 | 6701 | 88＇ $10: 00: 0$ | 07．7\％ | L9＇¢ | 7．7\％I | 99．00：00：0 | 00－99 | 991 | L6 Le：90：0 | 00L | \％．0 | \＆\％69 | $\pm 6$ | L0 |
| 9L＇IZ：8Z：0 | 0才＇${ }^{\text {I }}$ | $87^{\circ} 0$ | \＆＇76 | 67＇70：00：0 | 0L｀g | $06^{\circ} \mathrm{I}$ | 9＊20 | 87＊00：00：0 | L602 | 69L | ¢0＇L0：9\％：0 | ¢6 | $77^{\circ} 0$ | 6799 | 76 | 00 |
|  | de．${ }^{\text {d }}$ |  |  |  | ． 8 |  |  |  | de． 8 |  |  |  |  |  |  |  |



Figure 8.3: Results for instance class hiv4, probe length $\in[15,40]$, redundancy 2


Figure 8.4: Results for instance class hiv4, probe length $\in[15,40]$, redundancy 4


| 01：00：00：0 | 98：0 |  | も．8\％ | 70：00：00：0 | もI＇6 |  | Ə＇9\％ | 80：00：00：0 | 76．9L | 27 | 97：87： $77: 7$ | \＆¢ \％ | 20.0 | て＇キてん | $97 \%$ | $\varnothing$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 90：00：00：0 | $00 \cdot 0$ | $00^{\circ} 0$ | 0.88 | 90＇00：00：0 | 00 CL | $87^{\circ} \mathrm{I}$ | 7．78 | 90：00：00：0 | 00．gz | 98 | 28＇79：99：I | 88 | 71．0 | 8.802 | 81 | ¢\％ |
| 80．00：00：0 | $00 \cdot 0$ | $00^{\circ}$ | 0＇もち | 80：00：00：0 | 88：8 | $78^{\circ} 0$ | 097 | 80：00：00：0 | L9＇91 | 85 | LI＇ $7 \%$ ： $27: 0$ | $\dagger \%$ | ［＇0 | L＇069 | 88 | ఒ7 |
| 90：00：00：0 | $00 \cdot 0$ | 00.0 | 0.78 | 10：00：00：0 | 8L： | $29^{\circ}$ | L＇¢z | L0：00：00：0 | $60 \%$ | も\％ | 99．89：LZ：0 | 77 | $70^{\circ} 0$ | 8668 | 07 | 17 |
| 7ヵ＊00：00：0 | $00 \cdot 0$ | $00 \cdot 0$ | $0 \cdot 97$ | 60：00：00：0 | 0才， 0 I | 26.0 | 9.27 | 90：00：00：0 | 00．07 | 08 | 78：0才：ctic | 97 | O0．0 | 8－867 | 09 | 07 |
| 9800：00：0 | $00 \cdot 0$ | 00.0 | 0＇IZ | 80：00：00：0 | 㕵 2 | \＆¢ ${ }^{\circ}$ | 9.76 | 70：00：00：0 | 79．6 | \＆6 | 66．60：60：62 | L\％ | 75：0 | c． 806 | 98 | 6I |
| 20：00：00：0 | $00 \cdot 0$ | $00 \cdot 0$ | 0．12 | 70：00：00：0 | 06.1 | 79．0 | も「LZ | L0：00：00：0 | z9＇6 | 88 | 80＇29：co：0 | L\％ | $80^{\circ} 0$ | 0 ［t\％ | 08 | 8I |
| 80：00：00：0 | 00＇0 | $00 \cdot 0$ | 0＇LZ | 70：00：00：0 | Et＇LI | 02．0 | ¢ ¢ ¢ | 70：00：00：0 | 906L | 96 | 86．67：90：0 | 17 | 8.0 | ［＇089］ | 98 | LI |
| 70：00：00：0 | 00\％ | $00 \cdot 0$ | 0．2I | 10：00：00：0 | \＆9．$¢$ | z9．0 | 9．2I | 10：00：00：0 | 92．LI | 61 | 09＊89：6z：も | LI | $80^{\circ} 0$ | Le\＆ | $L 7$ | 9I |
| L゙・00：00：0 | $00 \cdot 0$ | $00 \cdot 0$ | 0.97 | 90：00：00：0 | 98.8 | $29^{\circ}$ | \＆8\％ | 70：00：00：0 | 酸LI | 67 | L6．97：89：0］ | 97 | $90^{\circ} 0$ | 8769 | 形 | ¢I |
| LI＇00：00：0 | $00 \cdot 0$ | $00 \cdot 0$ | $0 \cdot 8 \%$ | 01＇00：00：0 | 00．01 | $78^{\circ} 0$ | \＆：9\％ | 90：00：00：0 |  | 88 | 88：9ォ：67：9才 | ¢\％ | 81．0 | L．998 | 09 | 㕵 |
| ¢0．00：00：0 | $00 \cdot 0$ | $00 \cdot 0$ | $0.7 \%$ | 80：00：00：0 | L2： 2 | 79．0 | $9 \cdot 8 \%$ | 70：00：00：0 | 79\％ 8 | 96 | L0＇Et：Lg：0 | 77 | $\chi^{\text {¢ }}$ L | 9．Lz¢L | 07 | \＆1 |
| 90：00：00：0 | $00 \cdot 0$ | 00.0 | 0＇LZ | モ0：00：00：0 | 92． t | $78^{\circ}$ | $0.7 \%$ | 70：00：00：0 | 6でゅI | も\％ |  | L\％ | $90^{\circ} 0$ | $0 \cdot 689$ | 88 | ZI |
| 90：00：00：0 | $00 \cdot 0$ | $00 \cdot 0$ | 0．81 | 80：00：00：0 | 99．0］ | $88^{\circ}$ | 6．6I | 70：00：00：0 | 82： 27 | 88 | 28：87：90：8 | 81 | \＆1．0 | ¢ ¢ L0L | \＆8 | II |
| 88：00：00：0 | $00 \cdot 0$ | $00 \cdot 0$ | $0 \cdot 07$ | 80：00：00：0 | 00＇tI | ¢9 0 | $8.7 \%$ | 70：00：00：0 | 00＇gz | 96 | 76：80：00：0 | 07 | ¢0．0 | \＆67¢ | 98 | 0 I |
| 61．00：00：0 | $00 \cdot 0$ | $00 \cdot 0$ | $0 \cdot 8 \%$ | 85：00：00：0 | 00．0］ | $29^{\circ}$ | \＆ 96 | 90：00：00：0 | モ0¢ | 97 | 98 $28: 6 \mathrm{6}: 8$ | \＆\％ | 90.0 | $0 \cdot 985$ | 87 | 60 |
| L6 80：00：0 | L9．0 | て下＇0 | 788 | 20：00：00：0 | ［900 | $2 Z^{\circ} \mathrm{I}$ | ¢．98 | 90：00：00：0 | 9T．¢L | 88 | 7\％゙も0：87：0 | \＆¢ | 90.0 | 0 ¢¢¢ | g9 | 80 |
| 20：00：00：0 | $00 \cdot 0$ | $00 \cdot 0$ | $0 \cdot 97$ | 90：00：00：0 | 0才＇も | W2．0 | ［＇97 | 80：00：00：0 | 00．8 | $L 7$ | モ1．0て：0z：0 | 96 | 90.0 | $0 \cdot T L S$ | 87 | 20 |
| 67＇ $70: 00: 0$ | 00\％ | $00 \cdot 0$ | 065 | 72：00：00：0 | 96.8 | 29.0 | L：07 | 10：00：00：0 | 6L＇9L | 77 | 82＇tg：zg：L | 6I | ¢0．0 | $8 \cdot 895$ | 81 | 90 |
| 90：00：00：0 | $00 \cdot 0$ | $00 \cdot 0$ | 065 | 70：00：00：0 | L8：2 | 79．0 | キ＊0を | 10：00：00：0 | ¢G：0I | L\％ | 91．㘪：6¢：9 | 61 | 71．0 | 6.968 | ¢¢ | 90 |
| LT＇00：00：0 | 0＇8 | $87^{\circ} 0$ | L＇8z | 01：00：00：0 | 82＇も | 20． | †．97 | LI＇00：00：0 | DL＇LZ | 88 |  | $8 \%$ | LZ．0 | L＇926 | 87 | モ0 |
| モL＇00：00：0 | 00\％ | $00^{\circ}$ | $0 \cdot$ \％ | 60：00：00：0 | 98．9 | 79．0 | － 8 ¢ | 8000：00：0 | $60 \%$ | ¢\％ | 09．89：z¢：¢ | $7 \%$ | ¢0．0 | $8 \cdot 96$ | ¢t | ¢0 |
| 9600：00：0 | 9．＇t | $87^{\circ}$ | \＆ 97 | 60：00：00：0 | ct．li | 66.0 | 6.87 | 9000：00：0 | 87\％ 61 | LE | 70002：8L：99 | 97 | 75．0 | \＆¢п9 | 09 | 70 |
| 87：08：00：0 | 77： 7 | 7\％0 | 8.98 | EL＇00：00：0 | 99．0］ | 62.0 | 868 | 60：00：00：0 | 顽6L | ¢ | 19．7\％：¢¢： | 98 | $90^{\circ}$ | 9．89才 | ¢9 | L0 |
| 90：00：00：0 | $00 \cdot 0$ | $00 \cdot 0$ | 0＇もz |  | 89.6 | $87^{\circ}$ | 897 | 70：00：00：0 | 09＇ZI | 27 | モ0 $\llcorner$ ¢：9\＆：9 | ¢\％ | $70^{\circ}$ | ¢． $70 \pm$ | $8 \pm$ | 00 |
| u！ | e． 8 | ләрs | J | u！ 7 | ． 8 | эps |  | әu！ | e． 8 | req | эш！7 | 7 d | นəภฺ |  |  |  |



Figure 8.5: Results for instance class hiv1a, probe length $\in[15,40]$, redundancy 1
in few cases only. On the other hand, the run times for the CFTSB heuristic differ to a much lesser extent. Whereas the run times for computing optimal solutions vary wildly, even within one instance class, the run times of the CFTSB heuristic are all within the same order of magnitude. The low variance of the results over several runs is another advantage. Thus, for complex instances, where solving the ILP takes very long, the CFTSB heuristic provides a viable alternative.

If we look at the runs for instance class hiv1a, where we lower the redundancy to one, the picture changes. The CFTSB heuristic is able to compute an optimal solution for the majority of the instances in a very short time. Computing optimal solutions takes more than one hour for more than half of the instances. For some instances (2, 14-16, and 19) we experienced run times between 10 and 80 hours. In these cases, the CFTSB shines as it is able to produce optimal results (with the exception of instance 2) in under one second.

We assume that the heuristic works best when used with a redundancy parameter of one. A reason for this might be that the original approach has been designed for the standard set cover problem, which does not address the aspect of solution redundancy. We extended the original algorithm in a very straight-forward manner to handle this addi-

## 8 Computational Results

tional requirement, but as it seems, the performance seems to deteriorate with increasing redundancy values.

The DNA-BAR heuristic stands out by its speed, which seems to be independent of the particular instances it operates on. On the other hand, the solutions it produces are considerably worse than the first ones found by the CFTSB heuristic. Since in most cases it takes seldom more than a few seconds to find the first solution, it seems to be a better alternative wherever the focus does not lie on speed alone. Note, however, that you have to take into account the time spent for probe generation, which increases with the number and length of the input sequences.

## 9 Conclusions

The string barcoding problem is a variant of the well-known set cover problem. Like the latter, it is provably NP-hard and has many practical applications, especially in biology and medicine. In general, we want to identify an unknown sample sequence as one of a set of known sequences without inspecting the sample as a whole. Instead, we are seeking a set of short subsequences (probes) whose presence or absence in the sample we can check efficiently. We are seeking a set of probes of minimum cardinality such that the outcome of these tests allows us to identify the sample sequence uniquely. For each pair of input sequences, a feasible solution contains a predefined number (redundancy) of probes that is part of exactly one of the two sequences.

The problem can be divided into two parts. First, we need to identify the set of possible candidate probes. Each probe must be subsequence of at least one input sequence and must satisfy optional minimum or maximum length constraints. The set of all subsequences is in most cases much too large and contains lots of redundant elements, which appear in the same set of input sequences. A suffix tree can be used to identify a complete set of probes without redundant elements. It can be constructed in linear time and also has linear space complexity.

Few exact and many heuristic approaches address the problem. We adapted a very successful algorithm from literature that targets the set cover problem. It grounds on Lagrangian relaxation and proposes subgradient optimization for finding good Lagrangian multipliers. It is an iterative procedure which uses a Lagrangian heuristic to repeatedly solve different subproblems. Each subproblem is the result of setting up a partial solution of probes with a presumably high probability to appear in an optimal solution. Contrary to the original design, we solve the LP relaxation to find optimal initial Lagrangian multipliers for each subproblem.

We tested our implementation on a set of real-world instances derived from genomic data.

## 9 Conclusions

Unfortunately, the heuristic did not deliver results as good as the original approach did for the set cover problem. Especially when a higher solution redundancy was required, the solution quality deteriorated with optimal solutions being found in rare cases only.

On the other hand, our algorithm scales well with regard to instance size and offers a much greater degree of run time predictability. Though the computation of an optimal solution was sometimes even faster than the CFTSB heuristic, the run time over all instances showed an enormous variance. In contrast, the CFTSB heuristic produced good, though in most cases suboptimal, solutions in reasonable time. It was able to compute high quality solutions for very hard instances for which the computation of an optimal solution was excessively resource consuming, with respect to time as well as memory usage. Compared to the much simpler DNA-BAR heuristic proposed by DasGupta et al., our adaptation consistently produced results of superior quality within the same time frame.

## 10 Bibliography

[1] Boost C ++ libraries, 2007. http://www.boost.org. [Online; accessed 09/08/2007].
[2] Francisco Barahona and Ranga Anbil. The volume algorithm: producing primal solutions with a subgradient method. Mathematical Programming, 87(3):385-399, 2000.
[3] J.E. Beasley and P.C. Chu. A genetic algorithm for the set covering problem. European Journal of Operational Research, 94:392-404, 1996.
[4] John E. Beasley. Modern heuristic techniques for combinatorial problems, chapter Lagrangian Relaxation, pages 243-303. John Wiley \& Sons, Inc., New York, NY, USA, 1993.
[5] John E. Beasley. OR-library, 2005. http://people.brunel.ac.uk/~mastjjb/ jeb/info.html. [Online; accessed 10/03/2007].
[6] Piotr Berman, Bhaskar DasGupta, and Ming-Yang Kao. Tight approximability results for test set problems in bioinformatics. In Torben Hagerup and Jyrki Katajainen, editors, Algorithm Theory - SWAT 2004: 9th Scandinavian Workshop on Algorithm Theory, Humlebaek, Denmark, July 8-10, 2004. Proceedings, volume 3111 of Lecture Notes in Computer Science, pages 39-50. Springer, 2004.
[7] Piotr Berman, Bhaskar DasGupta, and Eduardo Sontag. Randomized approximation algorithms for set multicover problems with applications to reverse engineering of protein and gene networks. In Klaus Jansen, Sanjeev Khanna, José D. P. Rolim, and Dana Ron, editors, Approximation, Randomization, and Combinatorial Optimization: Algorithms and Techniques, volume 3122 of Lecture Notes in Computer Science, pages 39-50. Springer, 2004.

## 10 Bibliography

[8] James Borneman, Marek Chrobak, Gianluca Della Vedova, Andres Figueroa, and Tao Jiang. Probe selection algorithms with applications in the analysis of microbial communities. Bioinformatics, 17(90001):39-48, 2001.
[9] Michael J. Brusco, Larry W. Jacobs, and Garry M. Thompson. A morphing procedure to supplement a simulated annealing heuristic for cost- and coverage-corrleated set covering problems. Annals of Operations Research, 86:611-627, 1999.
[10] Alberto Caprara, Matteo Fischetti, and Paolo Toth. A heuristic method for the set covering problem. Operations Research, 47(5):730-743, 1999.
[11] Alberto Caprara, Paolo Toth, and Matteo Fischetti. Algorithms for the set covering problem. Annals of Operations Research, 98:353-371, 2000.
[12] Sebastián Ceria, Paolo Nobili, and Antonio Sassano. A lagrangian-based heuristic for large-scale set covering problems. Mathematical Programmming, 81(2):215-228, 1998. doi: http://dx.doi.org/10.1007/BF01581106.
[13] Bhaskar DasGupta, Kishori M. Konwar, Ion I. Mandoiu, and Alex A. Shvartsman. Highly scalable algorithms for robust string barcoding. International Journal of Bioinformatics Research and Applications, 1(2):145-161, 2005.
[14] Anton Eremeev. A genetic algorithm with a non-binary representation for the set covering problem. In Operations Research Proceedings 1998, pages 175-181. Springer, 1999.
[15] Markus Finger, Thomas Stützle, and Helena Lourenço. Exploiting fitness distance correlation of set covering problems. In Proceedings of the Applications of Evolutionary Computing on EvoWorkshops 2002, pages 61-71. Springer, 2002.
[16] Marshall L. Fisher. The Lagrangian relaxation method for solving integer programming problems. Management science, 27(1):1-18, 1981.
[17] Michael R. Garey and David S. Johnson. Computers and Intractability; A Guide to the Theory of NP-Completeness. W. H. Freeman \& Co., New York, NY, USA, 1990.
[18] Dan Gusfield. Algorithms on Strings, Trees, and Sequences. Cambridge University Press, New York, NY, USA, 1997.
[19] Michael Held and Richard M. Karp. The traveling-salesman problem and minimum spanning trees. Operations Research, 18(6):1138-1162, 1970.
[20] Michael Held and Richard M. Karp. The traveling-salesman problem and minimum spanning trees: Part II. Mathematical Programming, 1(1):6-25, 1971.
[21] Karla L. Hoffman and Manfred Padberg. Solving airline crew scheduling problems by branch-and-cut. Management Science, 39(6):657-682, 1993.
[22] Richard M. Karp. Reducibility among combinatorial problems. In Raymond E. Miller and James W. Thatcher, editors, Complexity of Computer Computations, pages 85-103. Plenum Press, New York, NY, USA, 1972.
[23] Scott Kirkpatrick, C. D. Gelatt Jr., and Mario P. Vecchi. Optimization by simulated annealing. Science, 220(4598):671-680, 1983.
[24] Gunnar W. Klau, Sven Rahmann, Alexander Schliep, Martin Vingron, and Knut Reinert. Optimal robust non-unique probe selection using integer linear programming. Bioinformatics, 20 (Suppl. 1):i186-i193, 2004.
[25] Giuseppe Lancia and Romeo Rizzi. The approximability of the string barcoding problem. Algorithms for Molecular Biology, 1, 2006.
[26] Elena Marchiori and Adri Steenbeek. An evolutionary algorithm for large scale set covering problems with application to airline crew scheduling. In Real-World Applications of Evolutionary Computing, EvoWorkshops 2000: EvoIASP, EvoSCONDI, EvoTel, EvoSTIM, EvoROB, and EvoFlight, pages 367-381. Springer, 2000.
[27] Sanjay Mehrotra. On the implementation of a primal-dual interior point method. SIAM Journal on Optimization, 2(4):575-601, 1992. doi: 10.1137/0802028.
[28] George L. Nemhauser and Laurence A. Wolsey. Integer and Combinatorial Optimization. Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley, 1999.

## 10 Bibliography

[29] National Center of Biotechnology Information. BLAST - basic local alignment search tool. http://www.ncbi.nlm.nih.gov/blast. [Online; accessed 09/26/2007].
[30] National Center of Biotechnology Information. Fasta format description. http: //www.ncbi.nlm.nih.gov/blast/fasta.shtml. [Online; accessed 09/02/2007].
[31] National Center of Biotechnology Information. GenBank overview, 2006. http: //www.ncbi.nlm.nih.gov/Genbank. [Online; accessed 09/03/2007].
[32] Sam Rash and Dan Gusfield. String barcoding: Uncovering optimal virus signatures. In RECOMB '02: Proceedings of the sixth annual international conference on Computational biology, pages 254-261, New York, NY, USA, 2002. ACM Press.
[33] Mutsuo Saito. An application of finite field: Design and implementation of 128-bit instruction-based fast pseudorandom number generator. Master's thesis, Hiroshima University, February 2007.
[34] Kenneth Steiglitz and Christos H. Papadimitriou. Combinatorial Optimization: Algorithms and Complexity. Dover, 1998.
[35] V. Černý. Thermodynamical approach to the traveling salesman problem: An efficient simulation algorithm. Journal of Optimization Theory and Applications, 45(1):41-51, 1985.
[36] Wikipedia. Singleton pattern - Wikipedia, the free encyclopedia, 2007. http: //en.wikipedia.org/wiki/Singleton_pattern. [Online; accessed 08/08/2007].
[37] Mutsunori Yagiura, Masahiro Kishida, and Toshihide Ibaraki. A 3-flip neighborhood local search for the set covering problem. European Journal of Operational Research, 2001.


[^0]:    

