



HAL
open science

Efficient algorithms for de novo assembly of alternative splicing events from RNA-seq data

Gustavo Akio Tominaga Sacomoto

► **To cite this version:**

Gustavo Akio Tominaga Sacomoto. Efficient algorithms for de novo assembly of alternative splicing events from RNA-seq data. Quantitative Methods [q-bio.QM]. Université Claude Bernard - Lyon I, 2014. English. NNT : 2014LYO10043 . tel-01095280v2

HAL Id: tel-01095280

<https://inria.hal.science/tel-01095280v2>

Submitted on 20 Feb 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

N° d'ordre: 43-2014

Année 2014

THÈSE

Présentée

devant L'UNIVERSITÉ CLAUDE BERNARD - LYON 1

pour l'obtention

du DIPLÔME DE DOCTORAT

(arrêté du 7 août 2006)

et soutenue publiquement le

6 Mars 2014

par

Gustavo Akio TOMINAGA SACOMOTO

**Efficient Algorithms for de novo Assembly of Alternative
Splicing Events from RNA-seq Data**

Directeur de thèse: Marie-France SAGOT

Co-Directeur: Pierluigi CRESCENZI

Co-Encadrant: Vincent LACROIX

JURY: Céline BROCHIER-ARMANET, Examineur
Michael BRUDNO, Rapporteur
Rodéric GUIGO, Rapporteur
Thierry LECROQ, Examineur
Peter WIDMAYER, Rapporteur

UNIVERSITÉ CLAUDE BERNARD-LYON 1

Président de l'Université

Vice-Président du Conseil d'Administration

Vice-Président du Conseil des Etudes et
de la Vie Universitaire

Vice-Président du Conseil Scientifique

Secrétaire Général

M. le Professeur M. A. BONMARTIN

M. le Professeur G. ANNAT

M. le Professeur D. SIMON

M. le Professeur J-F. MORNEX

M. G. GAY

SECTEUR SANTÉ

Composantes

Faculté de Médecin Lyon-Est - Claude
Bernard

Faculté de Médecine et de Maïeutique Lyon
Lyon Sud – Charles Mérieux

UFR d'Ontologie

Institut des Sciences Pharmaceutiques
et Biologiques

Institut Techniques de Réadaptation

Département de Formation et Centre de
Recherche en Biologie Humaine

Directeur: M. le Professeur J. ETIENNE

Directeur: M. le Professeur F-N. GILLY

Directeur: D. BOURGEOIS

Directeur: M. le Professeur F. LOCHER

Directeur: M. le Professeur MATILLON

Directeur: M. le Professeur P. FARGE

SECTEUR SCIENCES

Composantes

Faculté des Sciences et Technologies

Département Biologie

Département Chimie Biochimie

Département Génie Electrique et
des Procédés

Département Informatique

Département Mathématiques

Département Mécanique

Département Physique

Département Sciences de la Terre
des Activités Physiques et Sportives

UFR Sciences et Techniques

Observatoire de Lyon

Ecole Polytechnique Universitaire de Lyon 1

Ecole Supérieure de Chimie

Physique Electronique

Institut Universitaire de Technologie de
Lyon 1

Institut de Science Financière
et d'Assurances

Directeur: M. le Professeur S. De MARCHI

Directeur: M. le Professeur F. FLEURY

Directeur: Mme. le Professeur H. PARROT

Directeur: M. N. SIAUVE

Directeur: M. le Professeur S. AKKOUCHE

Directeur: M. le Professeur A. GOLDMAN

Directeur: M. le Professeur H. BEN HADID

Directeur: M. le Professeur S. FLECK

Directeur: M. le Professeur I. DANIEL

Directeur: M. C. COLLIGNON

Directeur: M. B. GUIDERDONI

Directeur: M. P. FOURNIER

Directeur: M. G. PIGNAULT

Directeur: M. le Professeur C. COULET

Directeur: M. le Professeur J. C. AUGROS

Acknowledgments

First and foremost, I would like to thank my advisors, Marie, Pilu and Vincent. *This thesis would not been possible without your guidance and support.* A special thanks to Vincent, who was a constant presence during the last three years, particularly in the beginning when I was trying to figure out my thesis topic.

My gratitude to my co-authors, Pavlos Antoniou, Étienne Birmelé, Rayan Chikhi, Pierluigi Crescenzi, Rui Ferreira, Roberto Grossi, Vincent Lacroix, Alice Julien-Laferrrière, Andrea Marino, Vincent Miele, Nadia Pisanti, Janice Kielbassa, Gregory Kucherov, Pierre Peterlongo, Marie-France Sagot, Kamil Salikhov, Romeo Rizzi and Raluca Uricaru. It has been a great pleasure to work with all of you. I learned a lot from you, not only new techniques or algorithms, but different ways to do research.

I would like to thank all KisSplice team, Alice Julien-Laferrrière, Camille Marchet, Vincent Miele and Vincent Lacroix. Thanks for all the bug reports, patches and discussions. I would also like to thank all Bamboo team, who created a great work environment. You made me feel home in a distant country.

Last but not least, I would like to write a few words in Portuguese dedicated to my family: *“Essa tese com certeza não seria possível sem o suporte incondicional da minha família, minha mãe, Miltes, meu pai, João, e minha irmã, Natália. Vocês que desde cedo me ensinaram que toda conquista é uma mistura de talento e muito suor. Vocês que sempre me encorajaram a seguir os meu sonhos. Obrigado por isso e muito mais!”*

Abstract

In this thesis, we address the problem of identifying and quantifying variants (alternative splicing and genomic polymorphism) in RNA-seq data when no reference genome is available, without assembling the full transcripts. Based on the fundamental idea that each variant corresponds to a recognizable pattern, a bubble, in a de Bruijn graph constructed from the RNA-seq reads, we propose a general model for all variants in such graphs. We then introduce an exact method, called KisSplice, to extract alternative splicing events. Finally, we show that it enables to identify more correct events than general purpose transcriptome assemblers (Grabherr et al. (2011)).

In order to deal with ever-increasing volumes of NGS data, an extra effort was put to make our method as scalable as possible. The main time bottleneck in the KisSplice is the bubble enumeration step. Thus, in order to improve the running time of KisSplice, we propose a new algorithm to enumerate bubbles. We show both theoretically and experimentally that our algorithm is several orders of magnitude faster than the heuristics based on cycle enumeration. The main memory bottleneck in KisSplice is the construction and representation of the de Bruijn graph. Thus, in order to reduce the memory consumption of KisSplice, we propose a new compact way to build and represent a de Bruijn graph improving over the state of the art (Chikhi and Rizk (2012)). We show both theoretically and experimentally that our approach uses 30% to 40% less memory than such state of the art, with an insignificant impact on the construction time.

Additionally, we show that the same techniques used to list bubbles can be applied in two classical enumeration problems: cycle listing and the K-shortest paths problem. In the first case, we give the first optimal algorithm to list cycles in undirected graphs, improving over Johnson's algorithm, the long-standing state of the art. This is the first improvement to this problem in almost 40 years. We also give the first optimal algorithm to list *st*-paths in undirected graphs. In the second case, we consider a different parameterization of the classical K-shortest simple (loopless) paths problem: instead of bounding the number of *st*-paths, we bound the weight of the *st*-paths. We present new algorithms with the same time complexities but using exponentially less memory than previous approaches.

Contents

Introduction	xi
1 Background	1
1.1 Biological Concepts	2
1.1.1 DNA, RNA and Protein	2
1.1.2 Alternative Splicing	4
1.1.3 Next Generation Sequencing (NGS)	5
1.2 Mathematical Concepts	8
1.2.1 Sets, Sequences and Strings	8
1.2.2 Graphs	8
1.2.3 Modeling and assembling NGS data	12
1.2.4 Enumeration Algorithms	19
2 Kissplice: calling alternative splicing from RNA-seq	23
2.1 Introduction	23
2.2 Methods	25
2.2.1 De Bruijn graph models	25
2.2.2 The KISSPLICE algorithm	31
2.3 Results	33
2.3.1 Simulated data	33
2.3.2 Real data	34
2.3.3 Characterization of novel AS events	35
2.3.4 Comparison with TRINITY	38
2.4 Discussion and conclusions	39
3 Listing in unweighted graphs	41
3.1 Efficient bubble enumeration in directed graphs	42
3.1.1 Introduction	42
3.1.2 De Bruijn graphs and bubbles	43
3.1.3 Turning bubbles into cycles	44
3.1.4 The algorithm	45
3.1.5 Proof of correctness and complexity analysis	49
3.1.6 Practical speedup	51
3.2 Optimal listing of cycles and <i>st</i> -paths in undirected graphs	52
3.2.1 Introduction	52
3.2.2 Preliminaries	53
3.2.3 Overview and main ideas	54
3.2.4 Amortization strategy	61

3.2.5	Certificate implementation and maintenance	63
3.2.6	Extended analysis of operations	66
3.3	Discussion and conclusions	71
4	Listing in weighted graphs	73
4.1	Listing bounded length bubbles in weighted directed graphs	74
4.1.1	Introduction	74
4.1.2	De Bruijn graphs and bounded length bubbles	74
4.1.3	An $O(n(m + n \log n))$ delay algorithm	76
4.1.4	Implementation and experimental results	78
4.1.5	A natural generalization	81
4.2	Listing bounded length paths	85
4.2.1	Introduction	85
4.2.2	Preliminaries	86
4.2.3	A simple polynomial delay algorithm	86
4.2.4	An improved algorithm for undirected graphs	88
4.2.5	Listing paths in increasing order of their lengths	92
4.3	Discussion and conclusions	94
5	Memory efficient de Bruijn graph representation	95
5.1	Introduction	96
5.2	Preliminaries	96
5.3	Cascading Bloom filter	97
5.4	Memory and time usage	98
5.4.1	Using different values of r for different filters	100
5.4.2	Query distribution among filters	100
5.5	Experimental results	102
5.5.1	Construction algorithm	102
5.5.2	Implementation and experimental setup	102
5.5.3	<i>E. coli</i> dataset, varying k	103
5.5.4	<i>E. coli</i> dataset, varying coverage	105
5.5.5	<i>E. coli</i> dataset, query statistics	105
5.5.6	Human dataset	106
5.6	Discussion and conclusions	106
	Conclusion and perspectives	109
	Bibliography	113

Introduction

The general question addressed in this thesis is how to extract biologically meaningful information from next generation sequencing (NGS) data without using a reference genome. The NGS technology allows to read, although in a fragmented way, the full content of the genetic material (DNA) in a given organism. The main difficulty lies in the fragmented nature of the NGS information. The NGS data forms a huge “jigsaw puzzle” that needs to be, at least partially solved (or assembled) in order to retrieve some biologically meaningful information. The challenge is made harder by the no reference genome assumption, which means that the puzzle needs to be solved relying only on the intrinsic information that two “pieces” (reads) are compatible, and thus should be together; there is no prior information about the “full picture”.

The usual route to solve this kind of problem is to first assemble the NGS reads and then analyze the result to draw conclusions about a given biological question. The difficulty with this approach is that even the simplest formalization of genome assembly as an optimization problem is NP-hard (shortest superstring problem). In addition, there is no guarantee that a solution to the optimization problem is unique or really corresponds to the original sequence. Actually, since the genome may contain repeats much larger than the read size, we may not have enough information to completely solve this problem regardless of the formulation. In practice, several heuristics are applied to this problem. The main goal of these heuristics is to produce long contigs, i.e. contiguous *consensus* sequences of overlapping reads. Of course, being heuristics, there are no strong guarantees about the results. Additionally, since they try to maximize the length of the contigs, genomic polymorphism (SNPs, indels and CNVs), corresponding to local variations in a (diploid) genome, are not modeled explicitly and thus systematically overlooked, for each variation only a consensus sequence is produced. In this thesis, we propose an alternative strategy avoiding the use of heuristics. We argue that for certain biological questions, it is not necessary to first solve the hard, often ill-posed, problem of completely assembling the NGS reads; it is instead sufficient, and sometimes even better (as in the case of searching for genomic polymorphism), to only *locally* assemble the data.

The classical view of the information flow inside the cell, or central dogma of molecular biology, can be summarized as: genes (DNA sequence) are transcribed into messenger molecules (mRNA) which are then translated into proteins. The NGS technology is not restricted to the sequence in the first step of this flow (DNA), it can as well be applied to the entire set of mRNAs of a cell (that is, to the *transcriptome*) through what is called an RNA-seq experiment. In this case, the assembly problem is not anymore to solve a single “jigsaw puzzle”, but several puzzles where the pieces are mixed together. Intuitively, this is a generalization of the genome assembly problem, and thus certainly no easier than it.

The specific problem we address here is the identification of variations (including alternative splicing) in RNA-seq data. The way the central dogma was stated may induce us to think that there is a one-to-one correspondence between genes and proteins. In general that

is not true, a single gene can produce several distinct proteins. Alternative splicing is one of the main factors responsible for this variability. It is a mechanism where several distinct mRNAs are produced from a same gene through an RNA sequence editing process. The local assembly strategy is specially suited to identify alternative splicing events because these are intrinsically local: an alternative splicing event usually generates two similar mRNAs molecules (isoforms) sharing the majority of their sequence (constitutive exons). The heuristics used in transcriptome assemblers tend to overlook such similar, but not identical, sequences.

Since the first automatic sequencing instruments (Sanger) were introduced in 1998, the cost per base sequenced has decreased dramatically: from \$2,400 to currently \$0.07 per million of bases sequenced (Illumina). This is due to an exponential increase in throughput; while the early Sanger machines produced 10 Kb (10^4 base pairs) per run, the current Illumina HiSeq 2000 produces 600 Gb/run. This represents an impressive 10^7 -fold increase. During the same time period, the processing capacity of an off-the-shelf computer had only a 10-fold increase, and the memory a 10^2 -fold increase. For that reason, any algorithm dealing with NGS data has to be highly efficient both in terms of memory usage and time consumption, and not rely on hardware improvements to compensate for ever-increasing volumes of data. In this thesis, we focus on time and memory efficient algorithms, from both the theoretical and the practical point of view.

The first step towards a solution to our variation identification problem is to have a suitable representation for the set of RNA-seq reads. A natural way to represent NGS data is to consider a directed graph (overlap graph), where each vertex corresponds to a read and the “compatibility” information, i.e. suffix-prefix overlaps, is stored in the arcs. This representation, however, does not scale well to large volumes of NGS data, since to compute the arcs, in principle, a quadratic number of read comparisons is necessary. A more suitable one, proposed by Pevzner et al. (2001), is to use de Bruijn graphs. This is the representation used here, and also in the majority of the recent NGS assemblers. A more in-depth comparison between several possible representations of NGS data, along with the necessary biological and mathematical background to follow this thesis, is given in Chapter 1.

We then show that variations in RNA-seq correspond to certain subgraphs in the de Bruijn graph built from the set of RNA-seq reads. More specifically, a variation creates a *bubble*, that is a pair of vertex-disjoint paths, in a de Bruijn graph. Hence, the problem of finding variations can be reduced to the problem of listing bubbles satisfying certain properties in the de Bruijn graph built from the set of reads. In Chapter 2, based on our paper Sacomoto et al. (2012), we describe a method, called KISSPLICE, implementing this strategy, along with a complete description of the relationship between variations and bubbles. We then show that, for the specific case of alternative splicing identification, our method is more sensitive than general purpose transcriptome assemblers and, although using relatively simple algorithms to build the graph and list the bubbles, uses roughly the same amount of memory and time.

The main time bottleneck in the KISSPLICE algorithm is the bubble enumeration step. Thus, in an effort to make our method as scalable as possible, in the first part of Chapter 3, which is based on our paper Birmelé et al. (2012), we modified Johnson’s cycle listing algorithm (Johnson (1975)) to enumerate bubbles in general directed graphs, while maintaining the same time complexity. For a directed graph with n vertices and m arcs containing η bubbles, the method we propose lists all bubbles with a given source in $O((n + m)(\eta + 1))$ total time and $O(m + n)$ delay (time elapsed between the output of two consecutive solutions). For the general problem of listing bubbles, this algorithm is exponentially faster than the listing algorithm of KISSPLICE. However, in the particular case of listing bubbles corresponding to alternative splicing events, this algorithm is outperformed by KISSPLICE. This issue is

addressed in the first part of Chapter 4, which is based on our paper [Sacomoto et al. \(2013\)](#). Using a different enumeration technique, we propose an algorithm to list bubbles with path length constraints in weighted directed graphs. The method we propose lists all bubbles with a given source in $O(n(m + n \log n))$ delay. Moreover, we experimentally show that this algorithm is several orders of magnitude faster than the listing algorithm of KISSPLICE to identify bubbles corresponding to alternative splicing events.

The main memory bottleneck in KISSPLICE is the construction and representation of the de Bruijn graph. Thus, again with the goal to make our method as scalable as possible, in Chapter 5, which is based on our paper [Salikhov et al. \(2013\)](#), we propose a new compact way to build and represent a de Bruijn graph improving over the state of the art [Chikhi and Rizk \(2012\)](#). We show both theoretically and experimentally that our approach uses 30% to 40% less memory than such state of the art, with an insignificant impact on the construction time. Our de Bruijn graph representation is general, in other words it is not restricted to the variation finding or RNA-seq context, and can be used as part of any algorithm that represents NGS data with de Bruijn graphs.

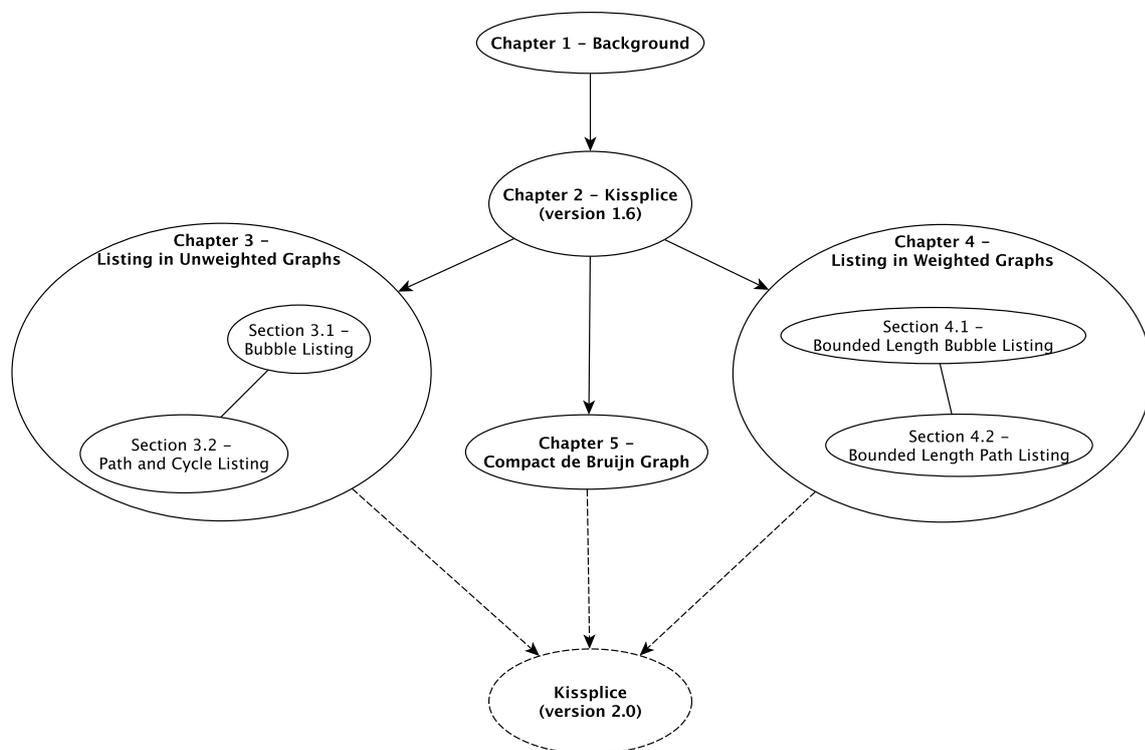
The central result of this thesis can be summarized as KISSPLICE version 2.0, the current version of KISSPLICE that includes both improvements, time and memory, discussed above. This version is able to treat medium-sized datasets (up to 100M Illumina reads) in a desktop computer (8GB of RAM) and large datasets (already tested with 1G Illumina reads) in a standard high-memory server (100GB of RAM). This is however not the only result of this thesis. As another example of the role of serendipity in scientific research, the techniques we developed while designing the two bubble listing algorithms turned out to be useful in the context of two classical enumeration problems: listing simple cycles in undirected graphs and listing the K -shortest paths.

The problem of efficiently listing all the simple cycles in a graph has been studied since the early 70s. For a graph with n vertices and m edges, containing η cycles, the most efficient solution was presented by [Johnson \(1975\)](#) and takes $O((\eta + 1)(m + n))$ time. This solution is not optimal for undirected graphs. Nevertheless, no theoretical improvements have been proposed in the past decades. In the second part of Chapter 3, which is based on our paper [Birmelé et al. \(2013\)](#), we present the first optimal solution to list all the simple cycles in an undirected graph G . Specifically, let $\mathcal{C}(G)$ denote the set of all these cycles ($|\mathcal{C}(G)| = \eta$). For a cycle $c \in \mathcal{C}(G)$, let $|c|$ denote the number of edges in c . Our algorithm requires $O(m + \sum_{c \in \mathcal{C}(G)} |c|)$ time and is asymptotically optimal: $\Omega(m)$ time is necessarily required to read G as input, and $\Omega(\sum_{c \in \mathcal{C}(G)} |c|)$ time is required to list the output. We also present the first optimal solution to list all the simple paths from s to t (shortly, st -paths) in an undirected graph G . Let $\mathcal{P}_{st}(G)$ denote the set of st -paths in G and, for an st -path $\pi \in \mathcal{P}_{st}(G)$, let $|\pi|$ be the number of edges in π . Our algorithm lists all the st -paths in G optimally in $O(m + \sum_{\pi \in \mathcal{P}_{st}(G)} |\pi|)$ time.

The K -shortest paths problem, that is returning the first K distinct shortest simple st -paths, has also been studied for more than 30 years, since the early 60s. For a weighted graph with n vertices and m edges, the most efficient solution is an $O(K(mn + n^2 \log n))$ time algorithm for directed graphs ([Yen \(1971\)](#); [Lawler \(1972\)](#)), and an $O(K(m + n \log n))$ time algorithm for undirected graphs ([Kato et al. \(1982\)](#)), both algorithms using $O(Kn + m)$ memory. In the second part of Chapter 4, which is based on our paper [Grossi et al. \(2014\)](#) (in preparation), we consider a different parameterization for this problem: instead of bounding the number of st -paths, we bound the weight of the st -paths. In other words, we consider the problem of listing st -paths with a weight bounded by α in a weighted graph. We present a general scheme to list bounded length st -paths in weighted graphs that takes $O(nt(n, m))$

time, where $t(n, m)$ is the time for a single source shortest path computation and η is the number of paths. This algorithm uses memory linear in the size of the graphs, independent of the number of paths output. For undirected non-negatively weighted graphs, we also show an improved algorithm that lists all st -paths with length bounded by α in $O((m + t(n, m))\eta)$ total time. In particular, this is $O(m\eta)$ for unit weights and $O((m + n \log n)\eta)$ for general non-negative weights. The time spent per path by our algorithms in both directed and undirected graphs matches the complexity of the best algorithms for the K -shortest path problem, while only using memory linear in the size of the graph. Moreover, we also show how to modify the general scheme to output the paths in increasing order of their lengths, providing an alternative solution to the K -shortest paths problem.

A summary of the thesis organization, and the relation between the chapters, is given in the figure below. As suggested by the figure, Chapters 3, 4 and 5 can be read more or less independently of each other.



Thesis organization and the corresponding versions of KISSPLICE.

Chapter 1

Background

Contents

1.1 Biological Concepts	2
1.1.1 DNA, RNA and Protein	2
1.1.2 Alternative Splicing	4
1.1.3 Next Generation Sequencing (NGS)	5
1.2 Mathematical Concepts	8
1.2.1 Sets, Sequences and Strings	8
1.2.2 Graphs	8
1.2.3 Modeling and assembling NGS data	12
1.2.4 Enumeration Algorithms	19

In this chapter, we cover the background and introduce the main notations necessary to follow the rest of this thesis. Certainly it would be infeasible to cover all the material with sufficient detail to make a thesis self-contained; we therefore do not attempt to be exhaustive or even complete. For the majority of the topics presented in this chapter, we give only a brief introduction, whereas we spend more time on a few others, which we consider to be important to the thesis and less familiar to the reader. Whenever possible, we give the main intuition behind the concepts presented and their inter-relationship. The chapter is divided in two sections: *biological concepts* (Section 1.1) and *mathematical concepts* (Section 1.2).

The purpose of Section 1.1 is to present three main topics. The first is the *central dogma of molecular biology*, which intuitively gives a roadmap for the information flow inside the cell, from the “blueprints” (DNA) to the “workers” (proteins). The next topic is *alternative splicing*, one of the steps of the information flow in “complex organisms” (eukaryotes), where the “message” can be *modulated*, that is from the same input message, several distinct output messages (mRNAs) can be produced. Finally, the third topic is *next-generation sequencing (NGS)* with emphasis on *RNA-seq*, a technology that allows to read, although in a fragmented way, the mRNAs inside the cell. A central question in this thesis is to find, from RNA-seq data, all the alternative ways in which the mRNAs are *spliced* (i.e. to find alternative splicing events).

The main goal of Section 1.2 is to introduce, along with the main definitions, notations and some properties of the mathematical structures used in this thesis, two seemingly unrelated topics: modeling and assembling NGS data, and enumeration algorithms. The relationship, although not immediately apparent, stands at the very core of this thesis. By modeling NGS

data, more specifically RNA-seq, using a special kind of directed graphs, namely de Bruijn graphs, the problem of identifying biologically interesting structures (e.g. alternative splicing events) can be seen as an enumeration problem of special structures in those graphs. This relationship is further detailed and explored in Chapter 2.

It should be noted that several standard computer science topics used throughout this thesis are not covered in Section 1.2. For the analysis of algorithms, the asymptotic big O notation and basic data structures (e.g. stacks, queues and heaps), we refer to Cormen et al. (2001). For the computational complexity theory and a compendium of NP-hard problems, we refer to Ausiello et al. (1999) and Garey and Johnson (1979). For basic graph algorithms, e.g. depth-first search (DFS), breadth-first search (BFS) and Dijkstra's algorithm, we refer to Cormen et al. (2001) and Sedgewick (2001). Finally, for further information about graphs and digraphs, we refer to Diestel (2005) and Bang-Jensen and Gutin (2008), respectively.

1.1 Biological Concepts

1.1.1 DNA, RNA and Protein

Deoxyribonucleic acid (DNA) is a long biopolymer that carries the hereditary information in almost all known organisms. They contain the “blueprint” for a complete organism. In other words, it is a long molecule composed by a huge number of repeated subunits, called *nucleotides*, chemically bonded together. There are four different nucleotides, namely: *adenine* (A), *cytosine* (C), *thymine* (T) and *guanine* (G). A DNA molecule has a *double-stranded* structure, where each molecule is composed by two strands, i.e. two chains of nucleotides, of the same length, running in opposite directions and respecting a fixed pairing rule between the corresponding nucleotides in each strand. The rules for nucleotide pairing, or *hybridization*, are: A with T (and vice-versa), and C with G (and vice-versa). We say that A (resp. C) is *complementary* to T (resp. G). That way, each strand is the *reverse complement* of the other, i.e. the sequence of nucleotides in one strand is equal to the reverse sequence of the other strand after substituting each nucleotide by its complementary. A *genome* is the set of all the genetic material, in the form of DNA (except for some viruses), of a given organism.

Ribonucleic acid (RNA) is, as DNA, a long biopolymer composed by adenine, cytosine, guanine and *uracil* (U) instead of thymine. However, unlike DNA, an RNA is not double-stranded: it contains only a single chain of nucleotides, and the sequence is usually much shorter. Despite being single-stranded, the nucleotide hybridization rules still hold for RNA molecules, with U substituting T. The hybridization can occur inside the same molecule of RNA, with complementary stretches of RNA folding and binding together; between two molecules of RNA; or, under certain circumstances, between an RNA and partially single-stranded DNA molecules. In terms of function, except for some viruses, the main role of RNA is not to carry hereditary information, but to transport genetic information from the DNA to other parts of the cell. An RNA is the main “messenger” within the cell. The *transcriptome* is the set of all RNAs present in the organism. Unlike the genome, it is not the same in all cells at all times.

Proteins are another kind of biopolymers where, unlike nucleic acids (DNA and RNA), the subunits are called *amino acids* and, instead of 4, there are 20 different types. A protein may contain several linear chains of amino acids, called *polypeptides*, but there is no strict pairing rules for amino acids like for nucleic acids. Proteins perform a large number of

functions within an organism, including: catalyzing certain reactions, replicating DNA, and transporting molecules from one location to another. The proteins are the main “workers” of the cell.

The *central dogma of molecular biology* states that the “coded genetic information hard-wired into DNA is transcribed into individual transportable cassettes, composed of messenger RNA (mRNA)¹; each mRNA¹ cassette contains the program for synthesis of a particular protein (or small number of proteins)” (Lodish et al. (2000)). In other words, the “blueprint” for each protein encoded in the DNA sequence is transcribed to RNA which is then translated to proteins. Note that, for the passage from DNA to RNA, we use the term transcription whereas from RNA to protein we use translation. That is because RNA and DNA use the same “language”, they are encoded using the same set of letters (nucleotides), while proteins use a different set, the amino acids, so when passing from an RNA to a protein, there is a translation from one language (nucleotides) to another (amino acids). A diagram of the central dogma is shown in Fig. 1.1(a).

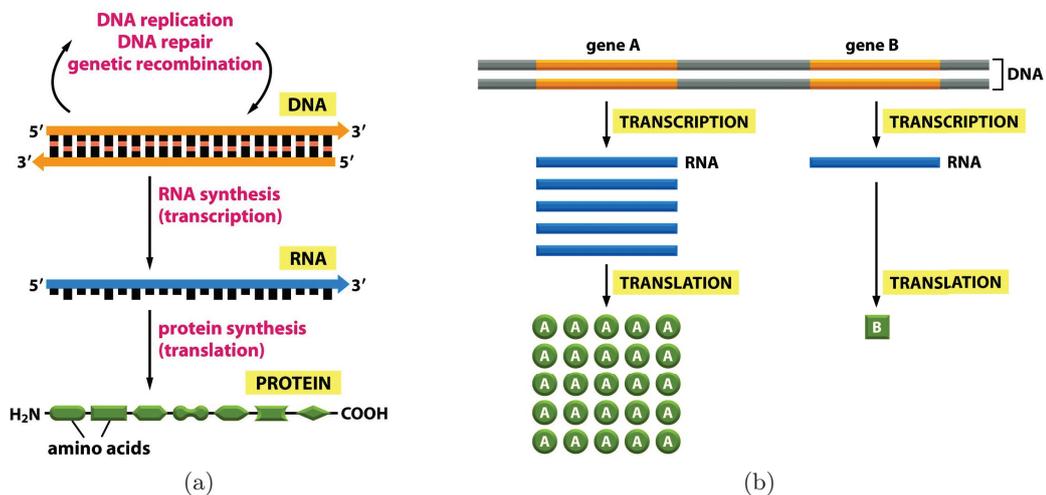


Figure 1.1: The central dogma of molecular biology and gene expression. (a) The central dogma states that genetic information flows from DNA to RNA to proteins. In the diagram, DNA is transcribed to RNA, which is then translated to proteins. (b) The genes are transcribed and translated at different rates to respond to different demands of the proteins corresponding to each gene. Two genes are shown in the left, a highly expressed gene and, in the right, a lowly expressed gene. Reproduced from Alberts et al. (2003).

Classically, a *gene* is a region of a DNA molecule that encodes for a protein (actually, a polypeptide chain). As stated before, translation and transcription are two main ways in which the cell reads out, or *expresses*, their genetic information, or genes. In a simplified view of the genome, we can assume that each gene is present in only one copy, and in this case, to respond to different demands of each protein, the cell has to translate and transcribe each gene with different efficiencies. A gene that is transcribed with higher rates is called *highly expressed*; on the other hand, a gene transcribed at lower rates is called *lowly expressed*. A diagram of the variable ranges of gene expression is shown in Fig. 1.1(b).

¹This is the classical view of molecular biology, which we present for simplicity. However, it is known that not every transcribed RNA is later translated into protein (Birney et al. (2007); van Bakel et al. (2010)). These molecules are known as non-coding RNA (ncRNA).

1.1.2 Alternative Splicing

In eukaryotic organisms (that comprises all living organisms except for bacteria and archaea); the cells have separate compartments for the nucleus and other structures (organelles). The genetic material (DNA) is stored in the nucleus which is separated from the cytoplasm by a membrane. In these organisms, transcription is done inside the nucleus, the RNA is then processed into an mRNA and exported to the cytoplasm to be translated into proteins. Another relevant particularity of eukaryotes is that their genes contain two different types of regions: *exons* and *introns*. The gene is then composed by alternating sequences of exons and introns. In the RNA processing step to produce an mRNA, the majority of the introns are removed, *spliced*, and a long chain of A's, the *poly-A tail*, is added to one of the ends. See Fig. 1.2(a) for a diagram of gene expression in eukaryotes.

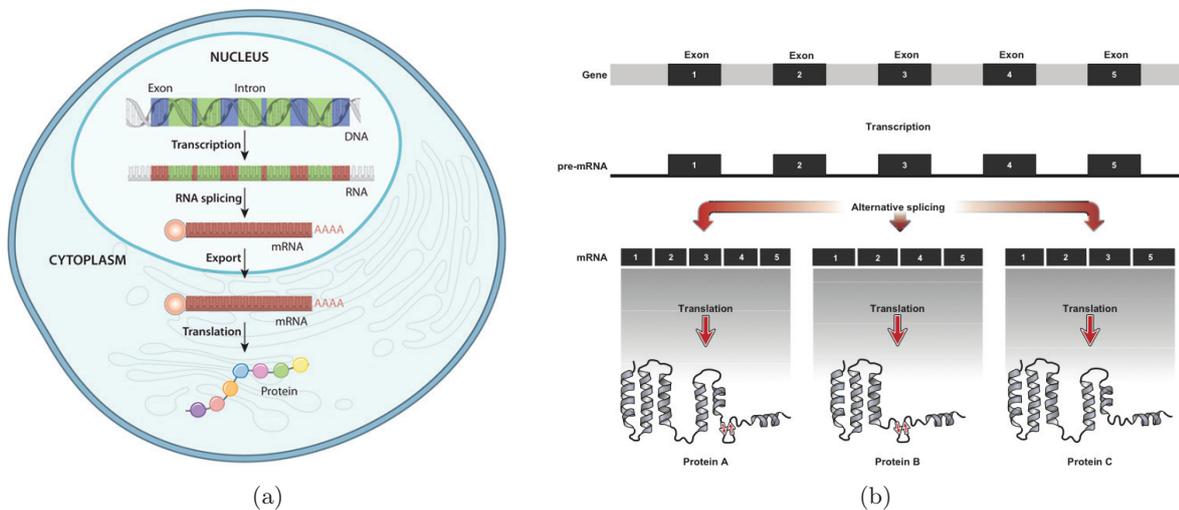


Figure 1.2: From DNA to RNA to protein in eukaryotes. (a) An eukaryotic gene transcription and translation is shown. The DNA inside the nucleus is transcribed into RNA, which is then processed into mRNA still inside the nucleus. Finally, the mature mRNA is transported to the cytoplasm to be translated into a polypeptide chain. (b) An eukaryotic gene with the introns and exons highlighted is shown. The same gene can produce different proteins through alternative splicing. In the example, three possible mRNAs are shown: one containing all the exons and two others with one of the exons skipped. Reproduced from [Schulz \(2010\)](#).

The discussion about the central dogma and gene expression may lead us to think that there is a one-to-one correspondence between genes and mRNAs (and proteins). In general, that is not the case, one gene can give rise to several distinct mRNAs (and proteins). Actually, it is estimated that 95% of all human genes give rise to more than one mRNA ([Pan et al. \(2008\)](#)). There are three main mechanism responsible for this variability: *alternative splicing*, *alternative promoters* and *alternative polyadenylation*. Alternative promoter and polyadenylation sites change the start and the end of the RNA transcription, respectively. Alternative splicing, our main interest here, is a post-transcription modification of the transcribed RNA (pre-mRNA). A diagram is shown in Fig. 1.2(b). Different mRNAs originating from the same gene are called *alternative isoforms* or simply *isoforms*. An exon is *constitutive* if it is present in all isoforms, and *alternative* otherwise.

Alternative splicing takes place when the transcribed RNA (pre-mRNA) is spliced to

produce the mature mRNA, instead of only remove all introns, some of exons may be skipped, included, shortened, or extended, and some introns may be retained; in each case a single pre-mRNA produces different mRNA variants (isoforms). An overview of the five types of alternative splicing events are shown in Fig. 1.3. A *splice site* is a sequence marking the border of a spliced region, usually the beginning or end of an intron, but it can also occur inside an exon in the case of alternative 3' or 5' splice sites; in this case, when performing alternative splicing, the exon is shortened or extended, respectively.

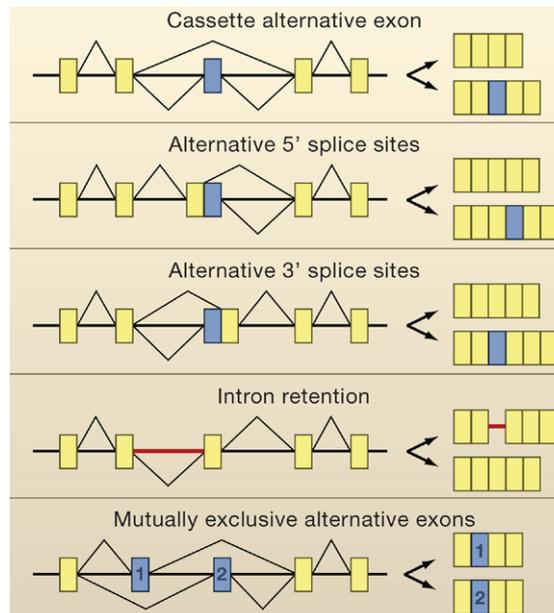


Figure 1.3: The five types of alternative splicing events are shown. In the left hand side, alternating sequences of exons (rectangles) and introns (straight lines) are represented; the zig-zag lines represent the removal (splicing) of a region. In the right hand side, in each case we have the two alternative isoforms after the splicing indicated by the zig-zag lines. Constitutive exons are shown in yellow, alternative exons in blue and retained introns in red. Reproduced and modified from [Blencowe \(2006\)](#).

1.1.3 Next Generation Sequencing (NGS)

The *shotgun sequencing* process was proposed by [Staden \(1979\)](#) to overcome the limitations of the semi-automated chain termination sequencing method introduced by [Sanger et al. \(1977\)](#). The main difficulty of Sanger's method when applied to whole genome sequencing was the length of the DNA strands that could be sequenced. Typically, only DNA strands with at most 1000 base pairs could be *read*, while even the smallest eukaryotic genome is several orders of magnitude larger. In order to overcome this, Staden's key idea was to randomly shear the whole genome of an organism into small fragments, and to independently sequence each fragment using Sanger's method. The resulting DNA *reads* would then be combined together, *in silico*, to reconstruct the original genome.

In the past few years, several sequencing technologies have been developed to replace Sanger sequencing in the shotgun sequencing context. These new methods, although producing shorter reads, are fully-automated and massively parallel and therefore can sequence a huge number of fragments in a same *run*, resulting in a huge number of reads in comparable

time and for a fraction of the cost of a Sanger sequencing. These high-throughput approaches are collectively known as *next generation sequencing (NGS)* methods. These technologies have been released as commercial products by several companies, e.g., the Solexa Genome Analyzer (Illumina, San Diego), the SOLiD platform (Applied Biosystems; USA) and 454 Genome Sequencers (Roche Applied Science; Basel). Although the specific details vary from one method to another, they can be seen as implementations of the cyclic-array sequencing (Shendure and Ji (2008)), which can be summarized as “the sequencing of a dense array of DNA features by iterative cycles of enzymatic manipulation and imaging-based data collection” (Mitra and Church (1999)). See Fig. 1.4 for a comparison between Sanger and cyclic-array methods for the shotgun sequencing process.

The DNA fragmentation process in the shotgun sequencing approach can be seen as a random sampling process, where each fragment comes from a random position in the genome. The reads are then obtained from both ends of each fragment, and usually do not cover the full DNA fragment. Assuming an uniform sampling, it can be theoretically shown (Lander and Waterman (1988)) that in order for all bases of the genome to be included in at least one read with high probability, a given amount of over-sampling is needed. In other words, for the whole genome to be sampled at least once, the average number of reads including a given genomic position, that is the *coverage*, should be higher than one.

RNA-seq

RNA-seq or *transcriptome sequencing* is the process of sequencing transcribed RNA using NGS technologies. Although the term RNA-seq is usually applied to both mRNA and non-coding RNA (e.g. micro RNAs) sequencing (Wang et al. (2009)), in this work we use it exclusively for mRNA sequencing. The basic RNA-seq protocol is very similar to genomic NGS, differing only in the first two phases. The first step is the extraction of mRNA from the cell, which is possible due to a distinctive structural property of mRNAs, namely its poly-A tail². The second step is the conversion to complementary DNA (cDNA), or reverse transcription. Then, the cDNA library is sequenced using the same NGS methods as for whole genome sequencing. The second step implies that the result of an RNA-seq is usually not strand specific, i.e. the reads obtained are a mixture of both the original strand of the mRNA and its reverse complement. There are alternative RNA-seq protocols (Levin et al. (2010)) where the strand information is not lost, however in this work we only consider the non strand-specific RNA-seq.

It is important to understand what new information can RNA-seq provide with regard to genomic sequencing. In addition, of course, to the information of which region is transcribed in the genome. As stated in Section 1.1.2, there is no one-to-one correspondence between genes and proteins, and in particular mRNAs, which implies that using only the genomic sequence, it is not possible to infer which set of mRNAs is expressed. In addition, as stated in Section 1.1.2, the expression level may vary from gene to gene, and this is reflected in the number of mRNA molecules transcribed for each gene. In that way, the sampling process in the fragmentation step of the NGS protocol is now a biased sampling towards the more transcribed mRNAs, i.e. the coverage of the mRNA is a proxy for its expression level. Therefore, the RNA-seq data brings at least two new dimensions in comparison to NGS genomic data: the variability of the mRNAs of a same gene and the expression level of each mRNA. Since 2008, several works (Cloonan and Grimmond (2008); Cloonan et al. (2008); Lister et al. (2008); Nagalakshmi et al.

²Actually, there are mRNAs *without* the poly-A tail (Yang et al. (2011); Djebali et al. (2012)). They are not selected, and thus not sequenced, by the standard RNA-seq protocol.

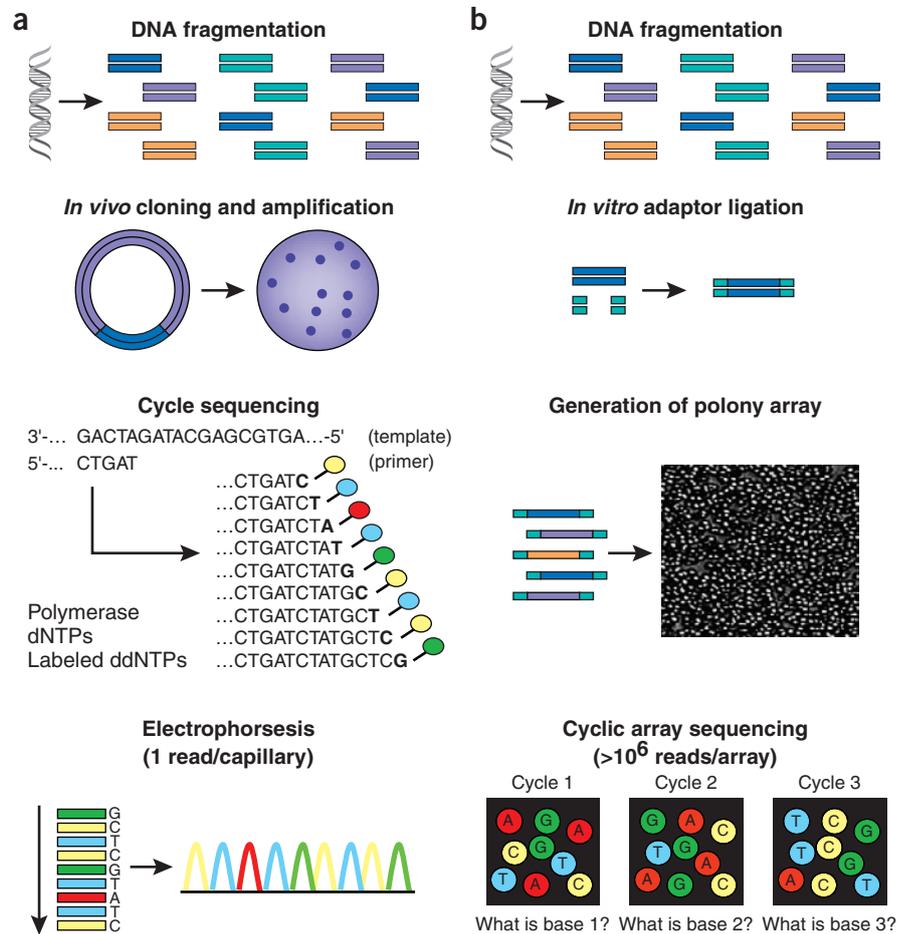


Figure 1.4: Workflow of shotgun Sanger sequencing versus next-generation sequencing. (a) In the high-throughput shotgun Sanger sequencing, DNA is first fragmented and subsequently integrated into a plasmid vector (a circular bacterial DNA that can replicate) that is later inserted into *Escherichia coli* to be amplified (copied several times). A single bacterial colony is selected for each sequencing reaction and the DNA is isolated. Each cycle sequencing reaction creates a ladder of dye-labeled products, which are subjected to electrophoretic separation in one run of a sequencing instrument. A detector for fluorescently labeled fragments of discrete sizes in the four-channel emission spectrum facilitates the sequencing trace. (b) In the next-generation shotgun sequencing, common adaptors are ligated to fragmented genomic DNA. The DNA is treated to create millions of immobilized PCR (polymerase chain reaction) colonies, called polonies, each containing copies of a single shotgun library fragment. In cyclic reactions, sequencing and detection of fluorescence labels determine a contiguous sequencing read for each polony. Reproduced from Shendure and Ji (2008).

(2008); Mortazavi et al. (2008); Sultan et al. (2008); Trapnell et al. (2010a); Pickrell et al. (2010); Peng et al. (2012)) have used RNA-seq data to shed a new light into the dynamics of gene expression in eukaryotic cells; see Cloonan and Grimmond (2008) and Wang et al. (2009) for comprehensive reviews.

1.2 Mathematical Concepts

1.2.1 Sets, Sequences and Strings

Given a set $X = \{x_1, \dots, x_n\}$, the cardinality of X is denoted by $|X|$. The power set 2^X is the set of all subsets of X , including the empty set. A *sequence* S is an ordered multi-set and is denoted by (s_1, \dots, s_n) . A *subsequence* of S is a sequence obtained from S by removing some elements, without changing the order of the others. A *prefix* (*suffix*) is the subsequence of S obtained after removing (s_i, \dots, s_n) ((s_1, \dots, s_i)), where $1 \leq i \leq n$. The *length* of the sequence, denoted by $|S|$, is the cardinality of the multi-set. The *concatenation* of S with an element x is the sequence $(s_1, \dots, s_n, s_{n+1})$ with $s_{n+1} = x$ and is denoted by Sx . Analogously, the concatenation of two sequences S_1, S_2 is denoted by S_1S_2 .

A *string*³ is a sequence where each element belongs to a set Σ , the *alphabet*. The prefix, suffix, concatenation and notation for the length are defined the same way as for sequences. The set of all strings over Σ is denoted by Σ^* . An element or *letter* of a string $w \in \Sigma^*$ at the position i , where $1 \leq i \leq |w|$, is denoted by $w[i]$. A *substring* of w is a contiguous subsequence of w denoted by $w[i, j]$, where $1 \leq i \leq j \leq |w|$. A *k-mer* is a substring of length k . Given two strings $x, y \in \Sigma^*$, the *suffix-prefix overlap* or simply *overlap*, is the longest suffix of x that is also a prefix of y . The *edit distance* for $x, y \in \Sigma^*$ is the minimum of number edit operations – substitutions, deletions or insertions – necessary to transform x in y (or vice-versa, it is symmetrical). The edit distance for $x, y \in \Sigma^*$ can be computed in $O(|x||y|)$ time using dynamic programming (Cormen et al. (2001)).

1.2.2 Graphs

A *directed graph* G is a pair of sets (V, E) such that $E \subseteq V^2$ is a set of ordered pairs. An element $v \in V$ is called a *vertex* of G , while an ordered pair $(u, v) \in E$ is called an *arc* of G . Given an arc $e = (u, v) \in E$, the *head* of e is vertex u and the *tail* is v . An *undirected graph* G is a pair of sets (V, E) such that $E \subseteq V^2$ is a set of unordered pairs, i.e. $(u, v) = (v, u)$. An unordered pair $(u, v) \in E$ is called an *edge*. Given a directed graph $G = (V, E)$, the *underlying undirected graph* is the undirected graph G' obtained from G disregarding the arc directions. Whenever it is clear from the context, or we are referring to both, we omit terms directed or undirected, saying simply “graph”. Finally, the graphs considered here, unless otherwise stated, are simple, that is, do not contain self-loops, i.e. $(v, v) \notin E$, nor multiple edges, i.e. E is a set not a multi-set.

Given a directed graph $G = (V, E)$ and a vertex $v \in V$, the in and out-neighborhoods of v are denoted by $N^+(v)$ and $N^-(v)$, respectively. For an undirected graph, $N^+(v) = N^-(v)$ and is denoted by $N(v)$. The *in* and *out-degree* of v are $d^-(v) = |N^-(v)|$ and $d^+(v) = |N^+(v)|$, respectively; for an undirected graph, the *degree* of v is $d(v) = |N(v)|$. A vertex $v \in V$ is a *source* of G if $N^-(v) = \emptyset$; symmetrically, it is a *sink* if $N^+(v) = \emptyset$. The reverse graph of $G = (V, E)$, denoted by $G^R = (V, E')$, is the directed graph obtained by reversing all arcs of G , i.e. $E' = \{(u, v) | (v, u) \in E\}$. The line graph of a directed graph G is the directed graph

³Although not strictly correct, we may use sequence when referring to strings.

$L(G)$ whose vertex set corresponds to the arc set of G and there is an arc directed from an arc e_1 to an arc e_2 if in G , the head of e_1 meets the tail of e_2 . See Fig. 1.5 for an example.

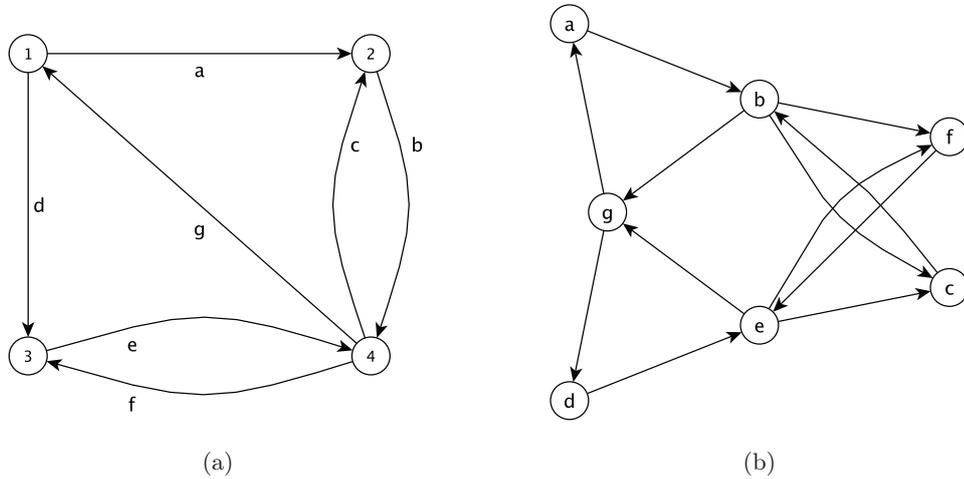


Figure 1.5: A line graph example. (a) A directed graph G with the vertices and arcs labeled. (b) The line graph $L(G)$, where the label of each vertex is the label of the corresponding arc in G . For example, arc a in G has arc g entering 1 and arc b leaving 2, in $L(G)$ the vertex a has g as an in-neighbor and b as an out-neighbor.

A *walk* in G is a sequence of arcs or vertices $p = (v_1, v_2) \dots (v_{n-1}, v_n) = (v_1, v_2, \dots, v_n)$, such that $(v_{i-1}, v_i) \in E$ for $1 \leq i \leq n$. The first (last) vertex of p is called *source* (*target*). A (simple)⁴ *path* is a walk in which all vertices are distinct: $v_i \neq v_j$ for all distinct $1 \leq i, j \leq n$. The path p from s to t is called *st-path* and is denoted by $s \rightsquigarrow t$, p_{st} or π_{st} . A *trail* is a walk in which all arcs are distinct, i.e. $(v_{i-1}, v_i) \neq (v_{j-1}, v_j)$ for all distinct $1 \leq i, j \leq n$. A subwalk of p is a subsequence $(v_{i-1}, v_i) \dots (v_{j-1}, v_j)$ of $p = (v_1, v_2) \dots (v_{n-1}, v_n)$. It is not hard to prove that every walk p , such that $v_1 \neq v_n$, contains a subwalk p' with the same source and target such that p' is a path. A (simple)⁵ *cycle* $c = (v_1, v_2) \dots (v_{n-1}, v_n)$ is a closed path, i.e. $v_1 = v_n$. A path or cycle p is *Hamiltonian* if it includes all vertices of G . A trail p is *Eulerian* if it includes all arcs of G .

A graph $H = (V_H, E_H)$ is a *subgraph* of $G = (V, E)$ if $V_H \subseteq V$ and $E_H \subseteq E$. The subgraph *induced* by a set of vertices $V' \subseteq V$ is the subgraph $G' = (V', E')$, where $E' = \{(u, v) : (u, v) \in E, u, v \in V'\}$ which is denoted by $G[V']$. The induced subgraph $G[V \setminus \{u\}]$ for $u \in V$ is denoted by $G - u$. Similarly for an edge $e \in E$, we adopt the notation $G - e = (V, E \setminus \{e\})$, and for any $F \subseteq E$, $G - F = (V, E \setminus F)$. Let p be a walk of G ; the induced subgraph $G[V \setminus p]$ is denoted by $G - p$.

A *weighted directed graph* $G = (V, E)$ is a directed graph with weights $w : E \rightarrow Q$ associated to the arcs. The weight of a walk $p = (v_1, v_2) \dots (v_{n-1}, v_n)$ is the sum of the weights of the arcs and is denoted by $w(p)$. The *distance* from s to t , denoted by $d_G(s, t)$ (we drop the subscript when the graph is clear from the context), is the weight of the shortest path from s to t . In an unweighted graph G , the distance between two vertices is equal to

⁴A path is, by definition, simple although we may say “simple path” to emphasize the fact that there are no duplicated vertex.

⁵A cycle is, by definition, simple although we may say “simple cycle” to emphasize the fact that there are no duplicated vertex.

the distance in the corresponding weighted graph with all the weights equal to one.

An undirected graph $G = (V, E)$ is *connected* if, for any two vertices $x, y \in V$, there exists a path from x to y . A *connected component* is a maximal connected subgraph. Any undirected graph can be uniquely partitioned into connected components. The connected components of $G = (V, E)$ can be found using any graph traversal algorithm; in particular, they can be computed in $O(|V| + |E|)$ time with a depth-first search (DFS) or breadth-first search (BFS) (Cormen et al. (2001)). A directed graph $G = (V, E)$ is *strongly connected* if for any $x, y \in V$ there exists the paths $x \rightsquigarrow y$ and $y \rightsquigarrow x$. It is *weakly connected* if its underlying undirected graph is connected. Of course, any strongly connected graph is also weakly connected.

Trees

A connected acyclic undirected graph G is called a (*unrooted*) *tree*. A *rooted tree* T is a tree with a special vertex r called *root*. The *parent* of a vertex v in T is the neighbor of v closer to the root. Every vertex, except the root, has a unique parent. The root has no parent. A *child* of v is a vertex of which v is the parent. Intuitively, a rooted tree is a tree where the edges are directed away from the root. The set of all children of v is denoted by $N^+(v)$. A vertex w is an *ancestor* of v if it belongs to the path $v \rightsquigarrow r$. Conversely, w is a *descendent* of v if v belongs to $w \rightsquigarrow r$. The descendent or ancestor is *proper* if it is different from v . A *subtree* of T rooted at v , denoted by T_v , is the subgraph of T induced by all descendents of v , which is also a tree, with root at v . A *leaf* is a vertex without any children. The *depth* of a vertex is the length of its unique path to the root. The *height* of a vertex is the length of the longest downward path to a leaf from that vertex.

Biconnected Graphs

An undirected graph $G = (V, E)$ is *biconnected* if it is connected and for any $x \in V$ the graph $G - x$ is still connected. Generalizing the definition of a connected graph, an undirected graph $G = (V, E)$ is *2-connected* (or *2-vertex-connected*) if for any $x, y \in V$ there exist two internally vertex-disjoint paths from x to y . By Menger's theorem (Diestel (2005)), the two definitions are equivalent, except when G is a single vertex⁶ or a single edge; in those cases the graphs are biconnected but not 2-connected. Before giving a simple characterization of the structure of 2-connected graphs (Lemma 1.1), we need another definition. Given an undirected graph H , a path p is an *H-path* or *ear* of H if p meets H exactly at its endpoints, i.e. the only vertices of p in common with H are its endpoints.

Lemma 1.1 (Diestel (2005)). *A graph is 2-connected if and only if it can be constructed from a cycle by successively adding H-paths to the graphs H already constructed.*

This process of adding H -paths to construct a 2-connected graph is also known as an *ear decomposition* (Bang-Jensen and Gutin (2008)). A similar characterization, using directed ears, can be stated for strongly connected graphs. Biconnected graphs have other interesting properties. For instance, given a biconnected graph $G = (V, E)$ and three distinct vertices $x, y, z \in V$, there is a xy -path passing through z . Indeed, let us construct G' by adding a new vertex w and the edges (x, w) , (y, w) to G ; the graph G' is also biconnected, so it contains two vertex-disjoint paths p_1, p_2 from w to z , one passing through x and the other through y . Thus, since p_1, p_2 are vertex-disjoint, the concatenation contains a path in G from x to

⁶We are using the convention that the *null* graph, i.e. the graph containing no vertices, is connected.

y passing through z . Using a similar argument, we can also prove the following. For any distinct $x, y \in V$ and an edge $e \in E$, there is a xy -path passing through e in G .

Similarly to connected components, for any undirected graph, a *biconnected component* (BCC) is a maximal biconnected subgraph. An *articulation point* or *cut vertex* is a vertex such that its removal increases the number of connected components. A biconnected component decomposition uniquely defines a partition on the edges, but not on the vertices. In other words, two distinct BCCs may share vertices but not edges. Actually, as stated in Lemma 1.2, a decomposition into BCCs can be defined as an equivalence relation for the edges, where each equivalence class is a BCC, and the common vertices are exactly the articulation points.

Lemma 1.2 (Tarjan (1972)). *Let $G = (V, E)$ be an undirected graph. We define an equivalence relation on the set of edges as follows: two edges are equivalent if and only if they belong to a common cycle. Let the distinct equivalence classes under this relation be E_i , $1 \leq i \leq l$, and let $B_i = (V_i, E_i)$, where V_i is the set of edges incident to E_i in G . Then:*

1. $\{B_1, B_2, \dots, B_l\}$ is the set of biconnected components of G ;
2. Each articulation point of G occurs more than once among the V_i , $1 \leq i \leq l$.
3. Each non-articulation point of G occurs exactly once among the V_i , $1 \leq i \leq l$;
4. The set $V_i \cap V_j$ contains at most one vertex, for any $1 \leq i, j \leq l$. Such vertex is an articulation point of G .

As a corollary, we have that a BCC decomposition is also a partition on the cycles of G , i.e. every cycle is contained in exactly one biconnected component. In addition, for any connected graph G , we have that the BCCs form a tree-like structure, the *block tree*, where two BCCs are adjacent if they share an articulation point. See Fig. 1.6 for an example. More precisely, let A be the set of articulation points of G and \mathcal{B} its set of biconnected components. Then, consider the graph \mathcal{T} whose vertices are $A \cup \mathcal{B}$ and there is an edge from $a \in A$ to $B \in \mathcal{B}$ if $a \in B$ (there are no edges between two vertices of A or \mathcal{B} , i.e. it is a bipartite graph). The graph \mathcal{T} is a tree.

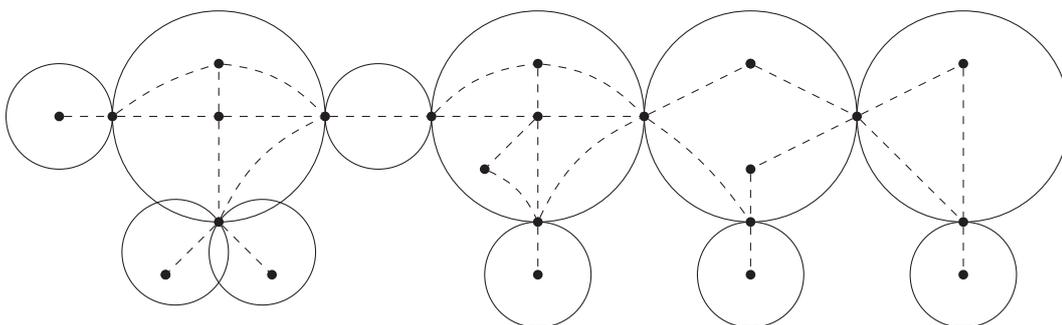


Figure 1.6: An example of connected graph G with its biconnected components highlighted. The articulation points are precisely the nodes in the intersection of the circle (BCCs). The circles plus the intersections form the block tree of G .

The biconnected components of $G = (V, E)$ can be computed in $O(|V| + |E|)$ using a modified DFS (Tarjan (1972); Cormen et al. (2001)). The same algorithm can also be used to find all the articulation points of G in linear time.

1.2.3 Modeling and assembling NGS data

As stated in Section 1.1.3, the NGS technology when applied to a genome (or a transcriptome, in the case of RNA-seq) produces a large number of fragments, or reads, from the original sequences. In this context, the most natural question is, given the set of reads, how to reconstruct the original sequence by combining the reads. This is called genome (transcriptome) assembly problem. At this point, it is important to make a clear distinction between *de novo* genome (transcriptome) assembly, which aims to reconstruct the genome (transcriptome) without using any other information but the reads, and comparative (re-sequencing) approaches that use knowledge on the genome of a closely related organism to guide the reconstruction. The *de novo* assembly problem is, as we show in this section, NP-hard under three common formalizations. On the other hand, comparative assembly is a considerably easier task admitting a polynomial algorithm, basically, it is sufficient to map the reads back to the reference genome (transcriptome) (Flicek and Birney (2009); Pop (2009)). Provided there exists a close enough reference, otherwise a mixture of both problems could be considered. Our main interest is in *de novo* assembly. Hereafter, we omit the term *de novo* when referring to it.

One basic assumption commonly made when modeling the assembly problem is that every read in the input must be present in the original genome (transcriptome). This neglects the fact that the reads may contain errors. Under this hypothesis, the genome (transcriptome) assembly problem can be formally stated as, given a set of strings $\mathcal{R} \subset \Sigma^* = \{A, C, T, G\}^*$, such that $r \in \mathcal{R}$ is a substring of an unknown string $S \in \Sigma^*$ (set of strings $S \subset \Sigma^*$), reconstruct the original string S (set of strings S). From now on, for the sake of a clear exposition, we consider only the genome assembly problem, and in the end of the section we highlight the differences with transcriptome assembly.

A simple way to formulate the assembly problem as an optimization problem, i.e. a problem of maximizing or minimizing a given objective function, is to require the reconstructed string to be of minimal length. Formally, given a set of strings $r \in \mathcal{R}$, find the minimum length string S^* such that every $r \in \mathcal{R}$ is a substring of S^* . This is precisely the *shortest superstring problem*, which is known to be NP-hard for $|\Sigma| \geq 2$ (Garey and Johnson (1979)). Despite that fact, some assemblers (Warren et al. (2007); Dohm et al. (2007)) employed this formulation. Of course, they do not solve the shortest superstring problem exactly; an exponential algorithm would indeed be impractical for all, but very small, instances. Instead, they employ variations of the following iterative greedy heuristic: at a given step the algorithm maintains a superstring S' for a subset $\mathcal{R}' \subseteq \mathcal{R}$, then extends S' with the read $r \in \mathcal{R} \setminus \mathcal{R}'$ such that the suffix-prefix overlap with S' is maximum and then adds r to \mathcal{R}' .

There are two main problems with the greedy strategy. The first issue is mainly due to the problem formulation: requiring the superstring to be of minimal length, although motivated by parsimony, is in the best case questionable. The reason is that the majority of the genomes have repeats, i.e. multiple identical, or nearly identical, substrings, while requiring a superstring of minimum length tends to over-collapse these substrings in the obtained solution. Consider for instance the example shown in Fig. 1.7. The second problem is due to the inherent local nature of the greedy heuristic: the choices are iteratively made without taking into account the global relationships between the reads. It is likely that in the true solution, the genome from which the reads were generated, several suffix-prefix read overlaps are not locally optimal. In order to address these issues, actually more the second one than the first, two high-level strategies were proposed (Pop (2009)): the overlap-layout-consensus (OLC) and the Eulerian path. In the core of each strategy is a representation of

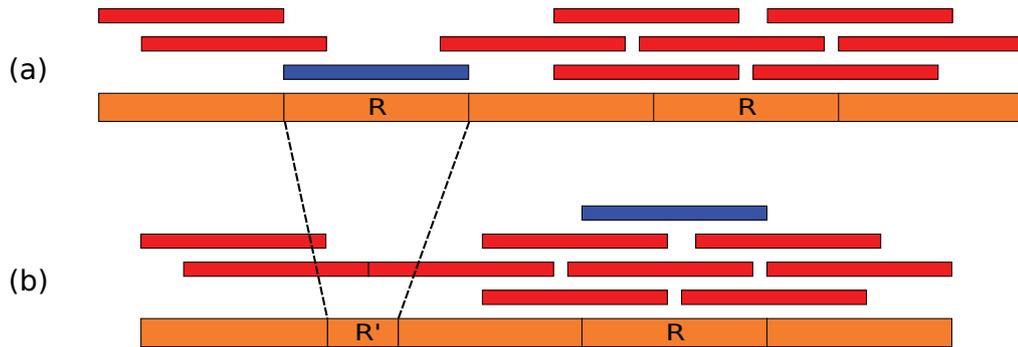


Figure 1.7: An example of a shortest superstring with an over-collapsed repeat. (a) The original string contains two exact copies of R ; the set of reads is shown above it, and the read in blue is entirely contained in one of the copies of R . (b) The shortest superstring for the same set of reads; the first copy R' is *over-collapsed*, and the blue read is now assigned to the second copy of R .

the reads set \mathcal{R} in terms of a (weighted) directed graph, the overlap graph for the OLC and the de Bruijn graph for the Eulerian path. The graph representation of \mathcal{R} allows for non-local analysis of the reads which is not possible with the greedy strategy.

Overlap-layout-consensus (OLC) Strategy

The overlap-layout-consensus strategy divides the assembly problem into three major stages. In the overlap stage, similarly to the approach used in the greedy strategy, for each pair of reads in \mathcal{R}^2 the maximal suffix-prefix overlaps are computed. In the layout stage, the overlap graph is constructed. That is, a complete weighted directed graph where each read of \mathcal{R} is a vertex, and there is a directed edge between any pair of reads (r_1, r_2) with weight equal to the length of the suffix-prefix overlap between r_1 and r_2 . The formal definition is given below (Definition 1.3). Next, still in the layout stage, the overlap graph is simplified. Finally, in the consensus stage, the genome sequence is obtained as consensus sequence, corresponding to a path or walk in the simplified overlap graph.

Definition 1.3 (Overlap Graph). *Given a set of reads $\mathcal{R} \subseteq \Sigma^*$, the overlap graph $G(\mathcal{R}) = (V, E)$, $w : E \rightarrow \mathbf{N}$ is a complete weighted directed graph such that:*

1. $V = \mathcal{R}$ and $E = \mathcal{R}^2$;
2. $w(u, v) = \text{length of the maximal suffix of } u \text{ that is equal}^7 \text{ to a prefix of } v$.

The main goal of the graph simplification in the layout stage is to reduce the complexity of the overlap graph. The first step is usually to remove all arcs that have weights below a given threshold. In practice, those edges are not even added to the original graph. A possible way to further reduce the complexity, proposed by Myers (2005), is to perform a transitive reduction in the graph, that is, to remove from the graph all the edges that are transitive inferable, i.e. consider the edges (x, y) , (y, z) and (x, z) , the last edge (x, z) is transitive

⁷For simplicity, we do not consider the more general definition where a small number of mismatches is allowed for the suffix-prefix overlap.

inferable since there is still a path from x to z after removing (x, z) . The subgraph of the overlap graph obtained by this process is called *string graph* (Myers (2005); Medvedev et al. (2007); Pop (2009)).

In the consensus stage, the problem of finding a walk in the graph corresponding to a consensus sequence can be formulated as an optimization problem by considering a constrained walk in the string graph and requiring it to be of minimum length (Medvedev et al. (2007)). The walk is constrained in the sense that some arcs of the string graph should be present at least once and others exactly once. Formally, we have a selection function s that classifies the arcs of string graph in the three categories: *optional* (no constraint), *required* (present at least once) and *exact* (present exactly once). The rationale for this classification is that some portions of the graph correspond to repeats in the genome, implying that they should be present more than once in the consensus sequence (walk), whereas other correspond to unique sequences, that should be present exactly once. This arc classification can be computed using the A-statistics (Myers et al. (2000)), as shown in Myers (2005). For a given selection function s and a string graph G , a walk of G respecting s is called an *s-walk*. Medvedev et al. (2007) showed, using a reduction from Hamiltonian path, that the problem of finding a minimum length *s-walk* is NP-hard.

Similarly to the shortest superstring problem, despite the fact that the minimum *s-walk* problem is NP-hard, several assemblers employed this formulation, using diverse heuristics to simplify the graph, i.e. make it as linear as possible, and find the consensus sequence. The OLC strategy was used by assemblers for various whole genome shotgun sequencing technologies, not only NGS technologies, for instance, the Celera assembler (Myers et al. (2000)), Arachne2 (Jaffe et al. (2003)) and Cap3 (Huang and Madan (1999)) for Sanger reads; Newbler (Margulies et al. (2005)) and Cabog (Miller et al. (2008)) for 454 reads; Edena (Hernandez et al. (2008)) and Shorty (Hossain et al. (2009)) for Illumina reads. For longer reads, i.e. Sanger and 454, the OLC seemed to be the more suitable approach (Pop (2009)). However, for shorter reads and much deeper coverages, the overlap computation step becomes a computational bottleneck. For that reason, most of the more recent assemblers use the Eulerian path strategy. With the notable exception of SGA (Simpson and Durbin (2012)) where they manage to overcome the overlap computation bottleneck using a FM-index (Ferragina and Manzini (2005)), which is a full-text compressed index based on the Burrows-Wheelers transformation allowing for fast substring queries.

Eulerian Path Strategy

The de Bruijn graph of order $k \in \mathbf{N}$ of a set of reads \mathcal{R} is a directed graph where each k -mer present in \mathcal{R} corresponds⁸ to a vertex and there is an arc (u, v) if the k -mers corresponding to u and v share a suffix-prefix overlap of size $k - 1$ and the corresponding $(k + 1)$ -mer, the k -mer u concatenated with the last character of v , is present in \mathcal{R} . The formal definition is given below (Definition 1.4). Actually, this is a *subgraph* of the de Bruijn graph under its classical combinatorial definition (Bang-Jensen and Gutin (2008)). However, following the terminology common to the bioinformatics literature, we still call it a de Bruijn graph. One of the most important aspects of de Bruijn graphs is that, unlike overlap graphs, they are not subjected to the overlap computation bottleneck. De Bruijn graphs can be efficiently computed using hashing or sorting. Indeed, given a read set \mathcal{R} , we can build a de Bruijn graph $G_k(\mathcal{R})$ using a hash table (Cormen et al. (2001)) to store all $(k + 1)$ -mers present in \mathcal{R} . As each insertion

⁸From now on, when considering de Bruijn graphs, we make no distinction between the k -mer corresponding to a vertex and the vertex itself.

and membership query in the hash table takes $O(1)$ (expected) time, the de Bruijn graph can be built in time linear in the size of \mathcal{R} , i.e. $O(\sum_{r \in \mathcal{R}} |r|)$.

Definition 1.4 (De Bruijn Graph). *Given a set of reads $\mathcal{R} \subseteq \Sigma^*$ and a parameter $k \in \mathbf{N}$, the de Bruijn graph $G_k(\mathcal{R}) = (V, E)$ is a directed graph such that:*

1. $V =$ the set of k -mers of \mathcal{R} ;
2. $E =$ the set of $(k+1)$ -mers of \mathcal{R} , in the sense that, given a $(k+1)$ -mer e of \mathcal{R} , we have that $e = (u, v)$, where $u = e[1, k]$ and $v = e[2, k+1]$.

Although not apparent from their definitions, intuitively, a de Bruijn can be seen as a special case of the overlap graph, where all the reads were further divided in k -mers and all the suffix-prefix overlaps have length exactly $k-1$. Indeed, the arcs in the overlap and de Bruijn graph represent the same structure, a suffix-prefix overlap between the strings corresponding to the vertices. In fact, in the particular case where all the reads of \mathcal{R} have length exactly $k+1$, the line graph of $G_k(\mathcal{R})$ is exactly the overlap graph of \mathcal{R} with the arcs of weight zero removed. Moreover, in a de Bruijn graph there is a loss of information with regard to the overlap graph: in de Bruijn graphs we do not have the information that two k -mers came from the same read. As a consequence there are walks in the de Bruijn graph that are not *read coherent*, i.e. are not entirely covered by an ordered set of reads where two adjacent reads have a non-empty suffix-prefix overlap (a tiling of the reads). An example of a de Bruijn and an overlap graph built from the same set of reads is shown in Fig. 1.8.

Interestingly, de Bruijn graphs were first used in computational biology in the context of sequencing by hybridization (SBH) (Pevzner (1989)). The outcome of a SBH experiment is the set of all distinct substrings of size k in the original sequence. Years later, it re-appeared in a pre-NGS context as an alternative to the OLC that could potentially lead to a polynomial algorithm for the genome assembly problem, although no such algorithm was provided (Pevzner et al. (2001)). The intuition was that differently from the OLC strategy that models the genome assembly problem as special case of the Hamiltonian path problem where the goal is to visit all vertices in the graph, genome assembly in a de Bruijn graph could be modelled as an Eulerian path (trail) problem, where the goal is to visit all *arcs* of the graph, for which there are polynomial algorithms (Cormen et al. (2001)). Unfortunately, there can be an exponential number of Eulerian trails in a graph (Diestel (2005)) and in order to select the one corresponding to the original sequence it is necessary to impose some constraints to the Eulerian trail, resulting in an NP-hard problem (Medvedev et al. (2007)).

As with the shortest superstring formulation for genome assembly, it is natural to require that all reads should be substrings of the solution of the genome assembly problem. In order to transpose this to the de Bruijn graph context, we observe that every read $r \in \mathcal{R}$ corresponds to a walk in the de Bruijn graph $G_k(\mathcal{R})$, possibly containing repeated vertices and arcs. This means that the solution should be a walk S in the de Bruijn graph $G_k(\mathcal{R})$, such that the each walk r_w corresponding to a read $r \in \mathcal{R}$ is a subwalk of S , i.e. S is a *superwalk* of $G_k(\mathcal{R})$. Now, motivated by parsimony, the optimization problem can be formulate as the problem of finding a minimum length superwalk in $G_k(\mathcal{R})$. Using a reduction from the shortest superstring problem, Medvedev et al. (2007) showed this problem is NP-hard.

The efficiency of a hash-based approach to construct a de Bruijn graph made it the ideal structure to represent NGS reads as the throughput of new technologies continued to increase. This is clear as the majority of the recent genome assemblers, although not trying to solve the minimum superwalk problem exactly (as with the OLC approaches, several heuristics to

instance, the minimum superwalk in $G_k(\mathcal{R})$ can be generalized to the problem of finding the set of walks \mathcal{S} , such that $\alpha \leq |\mathcal{S}| \leq \beta$, i.e. \mathcal{S} contains at least α *non-empty* walks and at most β ; each read $r \in \mathcal{R}$ is a subwalk of some $s \in \mathcal{S}$; and the total sum of the lengths of the walks in \mathcal{S} is minimum. The bounds α, β are part of the input of the problem and reflect the expected number of transcripts. Now, by choosing $\alpha = \beta = 1$ we obtain exactly the minimum superwalk problem, thus the generalization is also NP-hard. The same holds for the other two generalizations.

The majority of the transcriptome assemblers use de Bruijn graphs to represent the set of reads, for instance, in chronological order: Trans-ABYSS (Robertson et al. (2010)), Trinity (Grabherr et al. (2011)), Oases (Schulz et al. (2012)) and IDBA-tran (Peng et al. (2013)). As in the genome assembly case, in practice the transcriptome assemblers do not attempt to solve the generalized minimum superwalk problem exactly, employing instead various heuristics. There are, however, important differences in the heuristics used in both cases. Unlike the heuristics for genome assemblers where the main goal of the heuristics is to simplify the graph by linearizing it, the heuristics for transcriptome assembly have three main steps (Grabherr et al. (2011); Schulz et al. (2012)):

1. **Graph simplification.** This step is very similar to the de Bruijn graph simplification in genome assembly (see Fig. 1.9); the goal is the same, remove branching structures, that ideally correspond to sequencing errors, to transform the graph in a path. However, since alternative isoforms also produce branching structures, compared to genome assembly this simplification is done in a less aggressive way.
2. **Graph partition.** In the ideal case where two distinct genes do not share any k -mer (vertex), each connected component of the graph corresponds to the set of alternative isoforms of each gene. Unfortunately, genomes contain repeats, so two unrelated genes may share k -mers. The goal of this step is then to deal with these k -mers (vertices) linking two genes (connected components), in such a way that allows for the graph to be partition in subgraphs corresponding to genes. In the case of Schulz et al. (2012) this is, approximately, done by identifying the vertices corresponding to repeats and, for each vertex, duplicating it and dividing the arcs among the two copies.
3. **Path decomposition.** The goal of this step is, for each subgraph (gene) obtained in the last step: find the set of paths corresponding to the set of alternative isoforms of the gene. In the case of Schulz et al. (2012) this done by first applying an heuristic to remove cycles, and then iteratively applying a dynamic programming algorithm to find the path with largest read support (coverage).

We should stress that this is an *heuristic*, there are no guarantees that all transcripts are going to be (correctly) assembled. For instance, in the graph simplification step, even using a less aggressive approach, there are no guarantees that only sequencing errors are removed, actually it is quite likely that some alternative isoforms are also removed. Moreover, in the graph partition step, a gene can be split in two or more subgraphs, that way, in the path decomposition step, the resulting isoforms are necessarily fragmented. On the other hand, if two unrelated genes are in the same subgraph, in the path decomposition step, the resulting paths can be *chimeras*, containing parts of two unrelated transcripts. In general, transcripts from highly expressed genes are better assembled than lowly expressed ones; within the same gene, dominant isoforms are better assembled than minor ones.

1.2.4 Enumeration Algorithms

In this section, which is based on our paper [Marino et al. \(2014\)](#), we give a brief introduction to the area of an enumeration algorithms area. Naturally, the goal of enumeration is to list all feasible solutions of a given problem. For instance, given a graph $G = (V, E)$, listing the paths or shortest paths from a vertex $s \in V$ to a vertex $t \in V$, enumerating cycles, or enumerating all feasible solutions of a knapsack problem, are classical examples of *enumeration problems* or *listing problems*. An algorithm to solve an enumeration problem is called *enumeration algorithm* or *listing algorithm*.

While an optimization problem aims to find just *the best* solution according to an objective function, an enumeration problem aims to find *all* solutions satisfying a given set of constraints. This is particularly useful when the data is incomplete or the objective function is not clear: in these cases the best solution should be chosen among the results obtained by enumeration. Moreover, enumeration algorithms can be also applied to solve exactly NP-hard problems, by listing all feasible solutions and choosing the best one, as well as counting the number of feasible solutions in #P-hard problems ([Valiant \(1979\)](#)).

Complexity Classes

The classical complexity classes: P, NP, co-NP, NP-complete, #P, among others; are extremely useful but can only deal with problems with small (polynomial) outputs with regard to input size, e.g. the decision problems return 1 (true) or 0 (false). Quite often the number of solutions in a enumeration problem is exponential in the size of input, e.g. listing st -paths in a graphs, there is $G = (V, E)$ such that the number of st -paths is $\Omega(2^{|V|/2})$. To overcome this problem, the enumeration complexity classes are defined in an output-sensitive way; in other words, the time complexity takes into account the size of input and the *output*. In this way, if the number of solutions is small, an efficient algorithm has to terminate after a short (polynomial) time, otherwise it is allowed to spend more time. According to this idea, the following complexity classes were defined in [Johnson et al. \(1988\)](#).

Definition 1.5 (Polynomial Total Time). *An enumeration algorithm is polynomial total time if the time required to output all the configurations is bounded by a polynomial input size and the number of configurations.*

Definition 1.6 (Incremental Polynomial Time). *An enumeration algorithm is incremental polynomial time if it generates the configurations, one after the other in some order, in such a way that the time elapsed (delay) until the first is output, and thereafter the delay between any two consecutive solutions, is bounded by a polynomial in the input size and the number of configurations output so far.*

Definition 1.7 (Polynomial Delay). *An enumeration algorithm is polynomial delay if it generates the configurations, one after the other in some order, in such a way that the delay until the first is output, and thereafter the time elapsed (delay) between any two consecutive solutions, is bounded by a polynomial in the input size.*

Intuitively, the polynomial total time definition means that the delay between any two consecutive solutions is polynomial on *average*, while the polynomial delay definition implies that the maximum delay is polynomial. Hence, Definition 1.7 contains Definition 1.5 and Definition 1.6 is in between. It is important to stress that these complexity classes impose no restriction on the space complexity of the algorithms, e.g. a polynomial total time algorithm

can use memory exponential in the input size. However, in this thesis, we are mainly concerned about polynomial delay algorithm with space complexity polynomial on the *input size*. In some sense, these are efficient listing algorithm, or as defined in Fukuda et al. (1997) *strongly P-enumeration* algorithms.

The basic framework for efficient listing algorithms are: *backtracking* (unconstrained depth-first search with lexicographic ordering), *binary partition* (branch and bound-like recursive partition algorithm) and *reverse search* (traversal on the tree defined by the parent-child relation). In the remaining of this section, we give a brief introduction for the first two strategies; the backtracking method is used in Chapters 2 and 3, while the binary partition method is used in Chapter 3, and a variation of it in Chapter 4. An introduction of the reverse search method can be found in Avis and Fukuda (1993) and Marino et al. (2014).

The Backtracking Method

The backtracking method is a recursive¹⁰ listing technique based on the following simple idea: given a partial solution (i.e. a set that can be extended to a solution) recursively try all possible extensions leading to a solution. This technique has been successfully applied by several algorithms to list cycles in directed graphs (Tiernan (1970); Tarjan (1973); Johnson (1975); Szwarcfiter and Lauer (1976)); list bubbles in directed graphs (Birmelé et al. (2012)); list maximal cliques in undirected graphs (Bron and Kerbosch (1973); Koch (2001); Eppstein et al. (2010); Eppstein and Strash (2011)); and list maximal independent set in undirected graphs (Johnson et al. (1988)). The last two problems are particular cases of the more general problem of listing (maximal) sets in an *independence system*. This is not a coincidence, the backtracking method is particularly useful for listing problems that can be described as the enumeration of (maximal) sets in an independence system.

A collection of sets I is an independence system if for any set $X \in I$ all its subsets, $X' \subseteq X$, are also in I . More formally, a family of sets I over an universe U , i.e. $I \subseteq 2^U$, is an *independence system* if it satisfies the following properties: (i) the empty set belongs to I ; and (ii) every subset of some set in I also belongs to I , i.e. $X' \subseteq X$ and $X \in I$ implies that $X' \in I$. The sets in an independence system can be listed using the backtracking method: given a set X (initially the empty set) we recursively try to extend it, adding one new element $e \in U \setminus X$, and obtaining $X \cup \{e\}$. However, this is not enough to guarantee an efficient algorithm; for that, we also need to efficiently decide if the extension $X \cup \{e\}$ belongs to I , i.e. we need a polynomial membership oracle for I . An example of an efficient listing algorithm for the sets in an independence system is shown next. See Marino et al. (2014) for an example of an algorithm to list only the *maximal* sets in an independence system.

Enumerating all the subsets of a collection $U = \{a_1, \dots, a_n\} \subset \mathbb{Z}_{\geq 0}$ whose sum is less than b . The family of sets I over $U = \{a_1, \dots, a_n\}$ whose sum is less than b , form an independence system. Indeed, the empty set has sum 0, so it belongs to I ; since $a_i \geq 0$, the sum of $X' \subseteq X$ is not greater than X , so if $X \in I$ then $X' \in I$. Moreover, for a given subset X of U we can decide if it belongs to I in linear time, we just have to compute the sum of elements of X . We already have all requirements to design a backtracking-based algorithm for this problem: starting from the empty set $S = \emptyset$, we recursively try to add a new element $a_i \in U \setminus S$ to S , provided the resulting set $S \cup \{a_i\}$ belongs to I (the sum is smaller than b). The pseudocode is shown in Algorithm 1.1. Now, each iteration outputs a solution, and

¹⁰Of course, it can also be implemented in an iterative way, but it is not as natural as the recursive implementation.

takes $O(n)$ time, where $n = |U|$, thus the algorithm spend $O(n)$ time per solution. It is worth observing that by sorting the elements of U , each recursive call can generate a solution in $O(1)$ time, resulting in optimal $O(1)$ time per solution.

Algorithm 1.1: SUBSETSUM(S)

Input: S a set (initially empty) of integers belonging to the collection

$$U = \{a_1, \dots, a_n\} \subset \mathbb{Z}_{\geq 0}$$

Output: The subsets of U whose sum is less than b .

```

1 output  $S$ 
2 foreach  $a_i \in U \setminus S$  do
3   if  $a_i + \sum_{x \in S} x \leq b$  then
4     SUBSETSUM( $S \cup \{a_i\}$ )

```

The Binary Partition Method

The binary partition method, similarly to the backtracking method, is a recursive technique; based on the following simple idea: recursively divide the solution space into two *disjoint* parts until it becomes trivial, i.e. each part contains exactly one solution. More formally, let X be a subset of F , the set solutions, such that all elements of X satisfy a property \mathcal{P} ; recursively *partition* X into two subsets X_1 and X_2 (i.e. $X = X_1 \cup X_2$ and $X_1 \cap X_2 = \emptyset$), characterized by disjoint properties \mathcal{P}_1 and \mathcal{P}_2 , respectively. This procedure is repeated until the current set of solutions is a singleton. This technique has been successfully applied to many listing problems in graphs, including: st -paths in undirected graphs (Birmelé et al. (2013)), cycles in undirected graphs (Birmelé et al. (2013)), perfect matchings in bipartite graphs (Uno (2001)), and k -trees in undirected graphs (Ferreira et al. (2011)).

The recursion tree of any algorithm implementing the binary partition method is binary, since there are at most two recursive calls in the algorithm, one for each set X_1 and X_2 partitioning X . Moreover, unlike the backtracking method, the solutions are output only in the leaves of the tree, when the partition is a singleton. In order to design an efficient algorithm based on this technique, in any given call, when X is partitioned into X_1, X_2 , before proceeding with the recursion, we have to decide if X_1 and X_2 are non-empty, otherwise we would have many calls leading to no solution. Assuming that we have a polynomial (in the input size) oracle to decide if X_1 and X_2 are non-empty, and the height of the tree is bounded by a polynomial in the input size; the resulting algorithm has polynomial delay. Indeed, considering the recursion tree, the time elapsed between two solutions being output is bounded by the time spent in the nodes in any leaf-to-leaf path in the tree (recall that the solutions are only output in the leaves). As the height of the tree is polynomial in the input size, the number of nodes in any leaf-to-leaf path is also polynomial in the input size; and since the emptiness oracle is polynomial in the input size, the time spent in each node is also polynomial in the input size. An example of application of binary partition method is presented below.

Enumerating all the st -paths in an undirected graph $G = (V, E)$. The first requirement to apply the binary partition method is a property that allows to recursively partition the set of solutions, st -paths in this case. Let v be *any neighbor* of s , the set X of all st -paths in G can be partitioned in two sets: X_1 , the set of st -paths that do not include the edge (s, v) ;

and X_2 , the set of st -paths that include it. Actually, these sets can be described in a more “recursive” way: X_1 , the set of st -paths in $G - (s, v)$; and X_2 , the set of vt -paths in $G - s$ concatenated with the edge (s, v) ; in that way, both X_1 and X_2 are described in terms of the sets of xt -paths in a graph G' . Thus, a procedure $stPATHS(s, t, G)$, to list st -paths in G , can be implemented with recursive calls to $stPATHS(s, t, G - (s, v))$, corresponding to the st -paths in the partition X_1 ; and $stPATHS(v, t, G - s)$ with paths prepended with (s, v) , corresponding to the st -paths in the partition X_2 . In the base case, where $s = t$ and the current partition has only one solution, the corresponding st -path is output. The pseudocode is shown in Algorithm 1.2.

Recall that in order to have an efficient algorithm, before performing a recursive call we need to efficiently decide if the corresponding partition, X_1 or X_2 , is not empty, and only perform the call in that case. Clearly, X_1 is not empty if and only if there is at least one st -path in $G - (s, v)$, and X_2 is not empty if and only if there is at least one vt -path in $G - s$. In both cases, the test can be done in $O(|V| + |E|)$ time using a DFS traversal. Let us now analyze the delay of the algorithm. The height of the recursion is bounded by $|V| + |E|$, since at every call one vertex or one edge is removed from G , after $|V| + |E|$ calls the graph is empty. Hence, there are at most $2(|V| + |E|)$ nodes in any leaf-to-leaf path in the recursion tree. As the time spend in each node is $O(|V| + |E|)$, the delay is thus $O((|V| + |E|)^2)$.

Algorithm 1.2: $stPATHS(G, s, t, \pi)$

Input: An undirected graph G , vertices s and t , and a path π (initially empty).

Output: The paths from s to t in G .

```

1 if  $s = t$  then
2   |   output S
3   |   return
4 choose an edge  $e = (s, v)$ 
5 if there is a vt-path in  $G - s$  then
6   |    $stPATHS(G - s, v, t, \pi(s, v))$ 
7 if there is a st-path in  $G - e$  then
8   |    $stPATHS(G - e, s, t, \pi)$ 

```

Chapter 2

Kissplice: de novo calling alternative splicing events from RNA-seq data

Contents

2.1 Introduction	23
2.2 Methods	25
2.2.1 De Bruijn graph models	25
2.2.2 The KISSPLICE algorithm	31
2.3 Results	33
2.3.1 Simulated data	33
2.3.2 Real data	34
2.3.3 Characterization of novel AS events	35
2.3.4 Comparison with TRINITY	38
2.4 Discussion and conclusions	39

This chapter is strongly based on our paper [Sacomoto et al. \(2012\)](#). Here, we address the problem of identifying and quantifying variations (alternative splicing and genomic polymorphism) in RNA-seq data when no reference genome is available, without assembling the full transcripts. Based on the fundamental idea that each variation corresponds to a recognizable pattern in a de Bruijn graph constructed from the RNA-seq reads, we propose a general model for all variations in such graphs. We then introduce an exact algorithm, called KISSPLICE, to extract alternative splicing events. Finally, we show that it enables to identify more correct events than general purpose transcriptome assemblers. The algorithm presented in this chapter corresponds to KISSPLICE version 1.6. Further improvements in time and memory efficiency are shown in Chapters 4 and 5, respectively. The current implementation of KISSPLICE (version 2.0) already includes those improvements.

2.1 Introduction

Thanks to recent technological advances, sequencing is no longer restricted to genomes and can now be applied to many new areas, including the study of gene expression and splicing. As stated in Section 1.1.3, the so-called RNA-seq protocol consists in applying fragmentation

and reverse transcription to an RNA sample followed by sequencing the ends of the resulting cDNA fragments. The short sequencing reads then need to be reassembled to get back to the initial RNA molecules. As stated in Section 1.2.3, a lot of effort has been put on this assembly task, whether in the presence or in the absence of a reference genome but the general goal of identifying and quantifying all RNA molecules initially present in the sample remains hard to reach.

The main challenge is certainly that reads are short, and can therefore be ambiguously assigned to multiple transcripts. In particular, in the case of alternative splicing (AS for short), reads stemming from constitutive exons can be assigned to any alternative transcript containing this exon. Finding the correct transcript is often not possible given the data we have, and any choice will be arguable. As pointed out in [Martin and Wang \(2011\)](#), reference-based and de novo assemblers each have their own limitations. Reference-based assemblers ([Guttman et al. \(2010\)](#); [Trapnell et al. \(2010b\)](#); [Montgomery et al. \(2010\)](#); [Mezlini et al. \(2012\)](#); [Roberts and Pachter \(2013\)](#)) depend on the quality of the reference while only a small number of species currently have a high-quality reference genome available. De novo assemblers ([Robertson et al. \(2010\)](#); [Grabherr et al. \(2011\)](#); [Schulz et al. \(2012\)](#); [Peng et al. \(2013\)](#)), as stated in Section 1.2.3, implement reconstruction heuristics which may lead them to miss infrequent alternative transcripts while genes sharing repeats are likely to be assembled together and create chimeras.

We argue here that it is not always necessary to aim at the difficult goal of assembling full-length molecules. Instead, identifying the variable parts between molecules is already very valuable and does not require to solve the problem of assigning a read from a constitutive exon to the correct transcript. We therefore focus in this work on the simpler task of identifying variations in RNA-seq data. Three kinds of variations have to be considered: (i) alternative splicing (AS) that produces several alternative transcripts for a same gene, (ii) single nucleotide polymorphism (SNPs) that may also produce several transcripts for a same gene whenever they affect transcribed regions, and (iii) genomic insertions or deletions (indels). Our contribution in this chapter is double: we first give a general model which captures these three types of variations by linking them to characteristic structural patterns called “bubbles” in the de Bruijn graph (DBG for short) built from a set of RNA-seq reads, and second, we propose a method dedicated to the problem of identifying AS events in a DBG, including read-coverage quantification. We notice here that only splicing events but not transcriptional events, such as alternative start and polyadenylation sites, are covered by our method.

The identification of bubbles or bulges in DBG has been studied before in the context of genome assembly ([Pevzner et al. \(2004\)](#); [Zerbino and Birney \(2008\)](#); [Simpson et al. \(2009\)](#)), but the goal was not list them as variation-related structures, instead to simplify the de Bruijn graph. On the other hand, methods to identify variations as a restricted type of bubbles were proposed ([Peterlongo et al. \(2010\)](#); [Iqbal et al. \(2012\)](#); [Leggett et al. \(2013\)](#)), these works deal only with genomic NGS data and the variations considered are genomic polymorphisms, mainly SNPs and small indels. More recently, [Nijkamp et al. \(2013\)](#) presented a method to list bubble-like structures in the metagenomic context using the same graph decomposition previously proposed in [Sacomoto et al. \(2012\)](#).

When no reference genome is available, efforts have focused on assembling the full-length RNA molecules, not the variable parts which are our interest here. As stated in Section 1.2.3, most RNA-seq assemblers ([Robertson et al. \(2010\)](#); [Grabherr et al. \(2011\)](#); [Schulz et al. \(2012\)](#); [Peng et al. \(2013\)](#)) do rely on the use of a DBG, but, since the primary goal of an assembler is to produce the longest contigs, heuristics are applied, such as tip or bubble removal, in order to linearize the graph. The application of such heuristics results in a loss of information

which may in fact be crucial if the goal is to study expressed variations (alternative splicing and genomic polymorphism).

To our knowledge, this work is the first attempt to characterize variations in RNA-seq data without assembling full-length transcripts. We stress that it is not a general purpose transcriptome assembler and when we benchmark it against such methods, we only focus on the specific task of AS event calling. Finally, our method can be used in a comparative framework with two or more conditions and our quantification module outputs a coverage (number of reads mapped) for both the shorter and the longer isoform(s) of each AS event, in each experiment.

The chapter is organized as follows. We first present the model (Section 2.2.1) linking structures of the DBG for a set of RNA-seq reads to variations (AS, SNPs and indels), and then introduce a method, that we call KISSPLICE, for identifying DBG structures associated with AS events (Section 2.2.2). We show in Section 2.3 the results of using KISSPLICE compared with other methods on simulated and real data.

2.2 Methods

2.2.1 De Bruijn graph models

De Bruijn graph

In Section 1.2.3, we defined a de Bruijn graph for a read set $\mathcal{R} \subset \{A, C, T, G\}^*$ as a directed graph where each vertex corresponds to a k -mer and the arcs represent suffix-prefix overlaps of size $k - 1$ and correspond to a $(k + 1)$ -mer. See Fig. 2.1(a) for an example of *directed* de Bruijn graph. One problem with this definition is that it does not capture very well the double stranded nature of the DNA molecule, that is each k -mer present in the reads has a reverse complementary k -mer essentially representing the same information. Recall that, even though we are dealing with mRNAs sequencing data, and RNA is single stranded, one of the early steps of the RNA-seq protocol is reverse transcription, where the more stable double stranded cDNA is obtained from the mRNA extracted from the cell. Thus, RNA-seq data is also double¹ stranded.

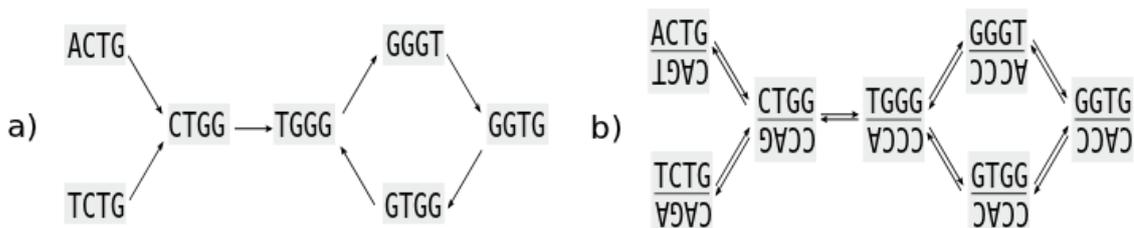


Figure 2.1: (a) The *directed* de Bruijn graph, with $k = 4$, for the set of reads $\mathcal{R} = \{ACTGG, TCTGGG, CTGGGTGGG\}$ is shown. (b) The *bidirected* de Bruijn graph, with $k = 4$, for the same set of reads is shown.

In order to better model the DNA double stranded nature, Medvedev et al. (2007), based on Kececioğlu (1992), modified the de Bruijn graph definition to associate to each vertex not only a k -mer $w \in \{A, C, T, G\}^k$ but its reverse complement $\bar{w} \in \{A, C, T, G\}^k$. In such a

¹As stated in Section 1.1.3, there are strand specific RNA-seq protocols.

context, a de Bruijn graph is a directed² multigraph $G = (V, E)$, where each vertex $v \in V$ is associated to a k -mer w and its reverse complement \bar{w} . The sequence w , denoted by $F(v)$, is the *forward sequence* of v , while \bar{w} , denoted by $R(v)$, is the reverse complement sequence of v . An arc exists from vertex v_1 to vertex v_2 if the suffix of length $k-1$ of $F(v_1)$ or $R(v_1)$ overlaps perfectly with the prefix of $F(v_2)$ or $R(v_2)$. See Fig. 2.1(b) for an example of *bidirected* de Bruijn graph. This is formally stated in the following definition.

Definition 2.1 (Bidirected de Bruijn Graph). *Given a set of reads $\mathcal{R} \subseteq \Sigma^*$ and a parameter $k \in \mathbb{N}$, the bidirected de Bruijn graph $B_k(\mathcal{R}) = (V, E)$ is a directed multigraph such that:*

1. $V = \{\{w, \bar{w}\} | w \text{ is a } k\text{-mer of } \mathcal{R}\}$,
2. $E = \{(x, y) \in V^2 | F(x) \text{ or } R(x) \text{ has a } k-1 \text{ suffix-prefix overlap with } F(y) \text{ or } R(y)\}$,

where \bar{w} is the reverse complement of w , and $F, R : V \rightarrow \Sigma^k$ are functions such that, for³ $v = \{w, \bar{w}\} \in V$, $F(v) = w$ and $R(v) = \bar{w}$.

It is convenient to augment this definition with arc labels in the set $\{F, R\}^2$. The first letter of the arc label indicates which of $F(v_1)$ or $R(v_1)$ has a suffix-prefix overlap with $F(v_2)$ or $R(v_2)$, this latter choice being indicated by the second letter. Moreover, because of the reverse complements, a suffix-prefix overlap from v_1 to v_2 induces a symmetrical suffix-prefix overlap from the corresponding complementary k -mers of v_2 to v_1 . As a result, there is an even number of arcs in the bidirected de Bruijn graph: if there is an arc from v_1 to v_2 then, necessarily, there is a *twin arc* from v_2 to v_1 with the corresponding label (i.e. if the first arc has label FF, RF, FR, RR then the second has label RR, RF, FR, FF , respectively). An example of a bidirected de Bruijn graph with the corresponding arc labels is shown in Fig. 2.2. The de Bruijn graphs considered in this chapter are all bidirected, so we omit this term, referring to them simply as de Bruijn graphs (or DBG for short).

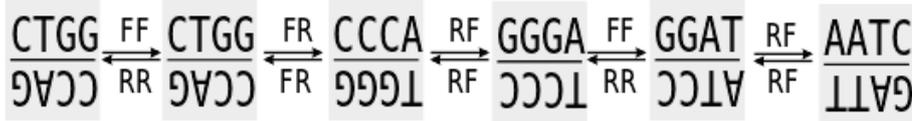


Figure 2.2: An example of a bidirected de Bruijn graph with arc labels.

Definition 2.2 (Valid path). *Given a bidirected de Bruijn graph $B_k(\mathcal{R}) = (V, E)$, a simple path $p = (v_1, v_2) \dots (v_{n-1}, v_n)$ is valid if for any two adjacent arcs (v_{i-1}, v_i) and (v_i, v_{i+1}) the labels are of the form L_1L_2 and L_2L_3 , respectively, where $L_1, L_2, L_3 \in \{R, F\}$.*

Consider the arc $e = (x, y)$ with label $L_1L_2 \in \{R, F\}^2$, we say that e enters y in the forward or reverse direction if $L_2 = F$ or $L_2 = R$, respectively, analogously for e leaving x . Basically, Definition 2.2 says that for a path to be valid all pairs of adjacent arcs should enter and leave a vertex in the same direction. For instance, for the graph shown in Fig. 2.2, the path from the leftmost vertex ($CTGG/CCAG$) going to the vertex $GGAT/ATCC$ is valid,

²The original definition of Medvedev et al. (2007) for the bidirected de Bruijn graph is based on *bidirected* graphs (Edmonds and Johnson (1970)). However, for the sake of a clearer exposition, Definition 2.1 is based on directed multigraph with arc labels instead. It can be shown that they are equivalent.

³Given a vertex $v \in V$ and its two corresponding k -mers, it is arbitrary, but fixed, which k -mer is the forward $w = F(v)$ and reverse $\bar{w} = R(v)$.

with (FF, FR, RF, FF) being the corresponding sequence of arc-labels. On the other hand, the path from the leftmost vertex ($CTGG/CCAG$) to the rightmost vertex ($AATC/GATT$) is not valid, since there is no arc leaving the *forward* part of $GGAT/ATCC$ and entering $AATC/GATT$. Finally, due the reverse complement relationship between the pair of labels, every valid path $p = s \rightsquigarrow t$ induces a complementary valid path $\bar{p} = t \rightsquigarrow s$, where each arc is substituted by its twin.

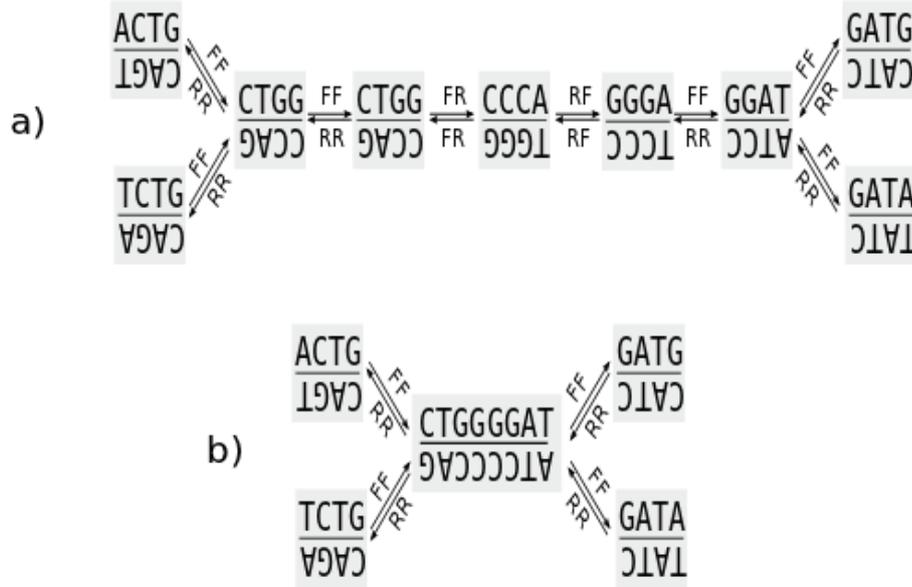


Figure 2.3: (a) An example of a de Bruijn graph, with $k = 4$. (b) The corresponding compressed de Bruijn graph. The compressed path in the DBG has 5 vertices, and the corresponding vertex in the cDBG has a pair of sequences each of length $k + (i - 1) = 4 + 4 = 8$.

A DBG can be compressed without loss of information by merging the vertices of non-branching valid paths. A non-branching valid path $p = s \rightsquigarrow t$ is a valid path such that for any internal vertex there is only one valid extension. See Fig. 2.3(a) for an example of a non-branching valid path from $CTGG/CCAG$ to $GGAT/ATCC$. Let $d_F^+(v)$ be the number of arcs leaving v in the forward direction, analogously for incoming arcs and the reverse direction. More explicitly, a path is non-branching if $d_F^+(s) = 1$ and $d_F^-(t) = 1$, and for every internal vertex v we have that $d_F^+(v) = d_F^-(v) = 1$ (the conditions can be stated considering the reverse direction instead, as one implies the other). Two adjacent vertices in a non-branching valid path are merged into one by removing the redundant information, that is keeping only one copy of the $k - 1$ suffix-prefix overlap. A valid path composed by $i > 1$ vertices is merged into one vertex containing as labels, not a pair of k -mers, but a pair of sequences of length $k + (i - 1)$ as each vertex in the path adds one new character to the first vertex. See Fig. 2.3 for an example of DBG and the corresponding compressed DBG. In the remaining of the chapter, we denote by cDBG a compressed DBG. Moreover, Definition 2.2 also applies to cDBG.

Bubble patterns in the cDBG

Variations (alternative splicing events and genomic polymorphisms) in a transcriptome, correspond to recognizable patterns in the cDBG, which we call a *bubble*. Intuitively, the variable

parts will correspond to alternative paths and the common parts will correspond to the beginning and end points of these paths. See Fig. 2.4 for an example of a bubble in a cDBG. We now formally define the notion of bubble, taking carefully into account the bidirected and arc labeled nature of the cDBG.

Definition 2.3 (Switching Vertex). *Given a path $p = (v_1, v_2) \dots (v_{n-1}, v_n)$ in a bidirected de Bruijn graph (or a cDBG), a vertex $v_i \in p$ is a switching vertex of p if the arc (v_i, v_{i+1}) leaves v_i in the complementary direction the arc (v_i, v_{i-1}) enters v_i .*

Definition 2.4 (Bubble). *Given a bidirected de Bruijn graph (or a cDBG) $B_k(\mathcal{R}) = (V, E)$, a bubble is a cycle with at least four distinct vertices such that there are exactly two switching vertices, denoted S_{left} and S_{right} .*

It follows directly from this definition, that for any bubble there are two valid paths, not sharing any internal vertex, from S_{left} to S_{right} . In the remaining of the chapter, we refer to these two paths as the paths of the bubble. If they differ in length, we refer to, respectively, the longer and the shorter path of the bubble. Where the length of a valid path in a cDBG is the length of the corresponding sequence of that path, not the number of vertices. In the example of Fig. 2.4 the switching vertices are encircled in blue and the longer path is shown above the shorter path.

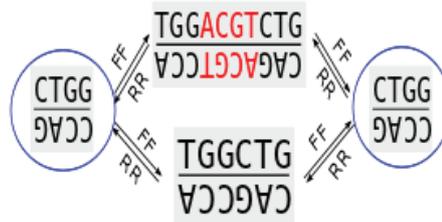


Figure 2.4: An example of a bubble in a bidirected de Bruijn graph. The bubble was generated by the sequences: CTGGACGTCTGG (asb) and CTGGCTGG (ab). The switching vertices are encircled in blue.

In general, any process generating patterns asb and $as'b$ in the sequences, with $a, b, s, s' \in \Sigma^*$, $|a| \geq k$, $|b| \geq k$ and s and s' not sharing any k -mer, creates a bubble in the cDBG. Indeed, all k -mers entirely contained in a (resp. b) compose the vertex S_{left} (resp. S_{right}). Since $|a| \geq k$ and $s \neq s'$, there is at least one pair of k -mers, one in as and the other in as' , sharing the $k - 1$ prefix and differing by the last letter, thus creating a branch in S_{left} from which the two paths in the bubble diverge. The same applies for sb , $s'b$ and S_{right} , where the paths merge again. All k -mers contained in s (resp. s') and in the junctions as and sb (resp. as' and $s'b$) compose the paths of the bubble. In the case where s is empty, the shorter path is composed of k -mers covering the junction ab . As we show later most AS events fall into this case.

We show next that this model is general as it captures SNPs, indels and AS events. However, the main focus of the algorithm we present in this work is the detection of bubbles generated by AS events.

Bubbles generated by AS events

As stated in Section 1.1.2, a single gene may give rise to multiple alternative spliceforms through the process of AS. Alternative spliceforms differ locally from each other by the in-

clusion or exclusion of subsequences. These subsequences may correspond to exons (exon skipping), exon fragments (alternative donor or acceptor sites) or introns (intron retention) as shown in Fig. 2.5(a). We should stress that we do not model mutually exclusive exons (another less frequent type of AS), since, as we show next, it does not correspond to the same pattern in terms of path lengths, and is therefore harder to treat. Additionally, alternative start and polyadenylation sites (Alberts et al. (2003)), which are not considered as AS events but as transcriptional events, are also not taken into account.

An alternative splicing event corresponds to a local variation between two alternative transcripts. It is characterized by two common sequences (a and b in the AS events given in Fig. 2.5(a)) and a single variable part (s in Fig. 2.5(a)). As stated in the last section, if $|a| \geq k$, $|b| \geq k$, then the patterns asb and ab generate a bubble in the cDBG. See Fig. 2.5(b) for an example of bubble generated by an AS event. In this example, the flanking sequences a and b correspond to the switching vertices, the variable part s to the longer path, and the shorter path corresponds to the k -mers covering the junction ab . Moreover, as there are $k - 1$ k -mers covering the junction between the two common sequences a and b , the shorter path is composed of *exactly* $k - 1$ k -mers. This however is not true in general. These two properties – correspondence between flanking sequences and switching vertices and exactly $k - 1$ k -mers in the shorter path – do not hold in general. They are not true when a and s share a suffix or b and s share a prefix. This case (which actually happens in more than 50% of the AS events, since it suffices that 1 out of 4 possible nucleotides are shared) is illustrated in Fig 2.5(c). In this example, the sequence of the switching vertex opening the bubble is a concatenated with the longest common prefix between b and s , and the shorter path contains $k - 3$ k -mers. In general, the length of the shorter path for a bubble generated by the pattern asb and ab is $k - 1 - lcp(s, b) - lcs(s, a)$, where $lcp(s, b)$ (resp. $lcs(s, a)$) is the length of the longest common prefix (resp. suffix) between s and b (resp. a). Overall, a bubble generated by an AS event always corresponds to a local variation between two RNA sequences. The shorter variant always has a length bounded by $2k - 2$. In human, 99% of the annotated exon skipping events yield a bubble with a shorter path length between $2k - 8$ and $2k - 2$ (Kuhn et al. (2009)).

Bubbles generated by SNPs, indels and repeats

Variations at the genomic level will necessarily also be present at the transcriptomic level whenever they affect transcribed regions. Two major types of variations can be observed at the genomic level: SNPs and indels. As shown in Fig. 2.6(a) and Fig. 2.6(b), they also generate bubbles in the cDBG.

However, these bubbles have characteristics which enable to differentiate them from bubbles generated by AS events. Indeed, bubbles generated by SNPs exhibit two paths of length exactly $2k - 1$, which is larger than $2k - 2$, the maximum size of the shorter path in a bubble generated by an AS event.

Genomic insertions or deletions (indels for short) may also generate bubbles with similar path lengths as bubbles generated by splicing events. In this case, the difference of length between the two paths is usually smaller, less than 3 nt for 85% of indels in human transcribed regions (Sherry et al. (2001)) whereas it is more than 3 nt for 99% of AS events. This suggests an initial criterion to separate between AS and indels: when the difference of path lengths is strictly below 3 we classify them as an indel; and AS event, otherwise. In Section 2.3.3, we refine the classification by considering that in an AS event in a coding region the difference of length is more likely to be a multiple of 3; since each codon is composed of 3 bases, an AS event with the length of the variable part not a multiple of 3 would cause a frame shift,

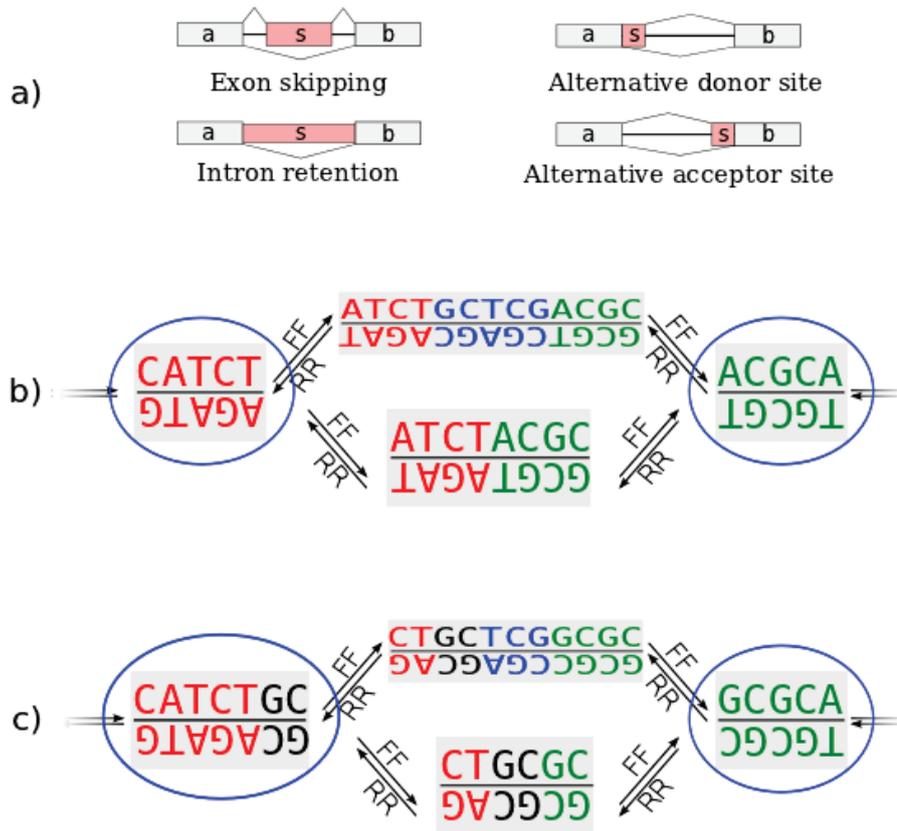


Figure 2.5: (a) AS events generating a bubble in the DBG. These events create a bubble in the DBG or cDBG, in which the shorter path is composed by k -mers covering the ab junction. This path, composed by $k - 1$ vertices in the DBG, is compressed into a sequence of length $2k - 2$ in the cDBG. (b) A bubble in a cDBG, with $k = 5$, due to the variable part $GCTCG$ (s). This bubble is generated by the sequences $CATCTACGCA$ (ab) and $CATCTGCTCGACGCA$ (asb). The shorter path has length $2k - 2 = 8$. (c) A bubble in a cDBG, with $k = 5$, due to the skipped exon $GCTCG$ (s) with the flanking sequences $CATCT$ (a) and $GCGCA$ (b). This bubble is generated by the sequences $CATCTGCGCA$ (ab) and $CATCTGCTCGGCGCA$ (asb). Observe that s and b share the prefix GC . As a result, the k -mers $ATCTG$ and $TCTGC$ are common to both paths and are represented only once; and the length of the shorter path is $2k - 2 - 2 = 6$.

potentially change completely the amino acid sequence.

Finally, inexact repeats may generate bubbles with a similar path length as bubbles generated by splicing events, but the sequences of the paths exhibit a clear pattern which can be easily identified: the longer path contains an inexact repeat. More precisely, as outlined in Fig 2.6(c), it is sufficient to compare the shorter path with one of the ends of the longer path. We treat this kind of event as false positive, bubbles that do not correspond to a true variation in the dataset. However, there is a type of true genomic polymorphism that may be included in this group: copy number variations (CNVs).

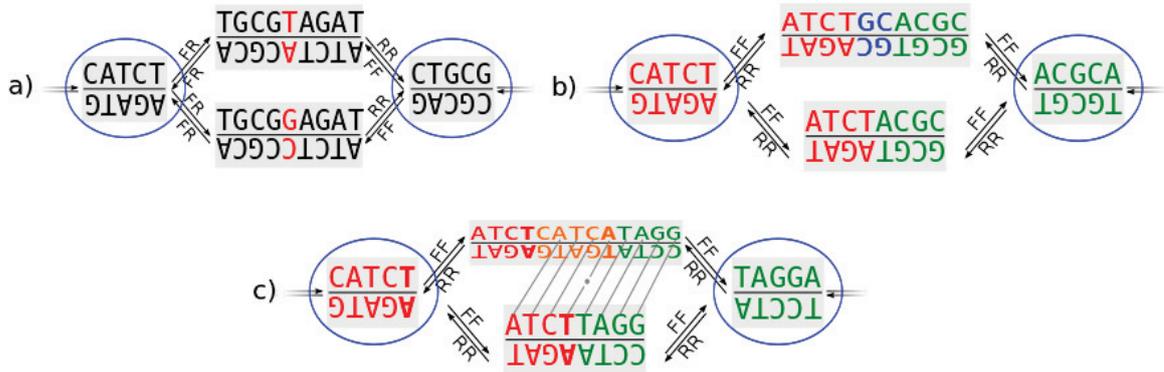


Figure 2.6: (a) Bubble due to a SNP (substitution is the red letter). Starting from the forward strand in the leftmost (switching) vertex would generate the sequences *CATCTACGCAG* (upper path) and *CATCTCCGCAG* (lower path). (b) Bubble due to the deletion *GC*. This bubble is generated by the sequences *CATCTACGCA* and *CATCTGCACGCA*. (c) Bubble due to an inexact repeat. This bubble is generated by the sequences *CATCTTAGGA* and *CATCTCATCATAGGA*, where *CATCTCATCA* is an inexact repeat.

In the following, we focus on bubbles generated by AS events. In the output of the method we present in the next section, we do provide as a collateral result three additional collections of bubbles: one corresponding to putative SNPs, one to short indels, and one to putative repeats associated bubbles. The post-treatment of these collections to discard false positives caused by sequencing errors, or recover the ones corresponding to CNVs, is beyond the scope of this work.

2.2.2 The KISSPLICE algorithm

The KISSPLICE algorithm detects in the cDBG all the bubble patterns generated by AS events, i.e. the bubbles having a shorter path of length at most $2k - 2$. Essentially, the algorithm lists all the *cycles* verifying the following criteria:

- i the cycle contains exactly two switching vertices, i.e. it corresponds to a bubble;
- ii the length of the shorter path linking the two switching vertices is smaller than $2k - 2$;
- iii both paths have length greater than $2k - 8$;
- iv the length of the longer path is smaller than α (a parameter, set to 1000 by default).

The last condition imposes an upper bound on the length of the exon or intron skipped in a AS event, and is necessary due to performance issues, a larger value considerably increases

the running times. Further criteria are applied to make the algorithm more efficient without loss of information, and to eliminate bubbles that do not correspond to AS.

Since the number of cycles in a graph may be exponential in the size of the graph, the naive approach of listing all cycles of the cDBG and verifying which of them satisfy our conditions is only viable for very small cases. Nonetheless, KISSPLICE is able to enumerate a potentially exponential number of bubbles for real-sized dataset in very reasonable time and memory consumption. This is in part due to the fact that, previous to cycle enumeration, the graph is pre-processed in a way that, along with the pruning criteria of Step 4 (see below), is responsible for a good performance in practice.

KISSPLICE is indeed composed of six main steps which are described next. The pre-processing just mentioned corresponds to Step 2, and the enumeration algorithm is described in Step 4. This description corresponds to KISSPLICE version 1.6. A memory efficient replacement for Step 1 is presented in Chapter 5. And a time efficient replacement for Step 4 is presented in Chapter 4. The current implementation of KISSPLICE (version 2.0) includes both improvements.

1. *cDBG construction*. Construction of the cDBG of the reads of one or several RNA-seq experiments. The first step is to obtain the list of unique k -mers, with the corresponding multiplicities (coverage), from the reads. This is done, using constant memory, by applying an algorithm similar to the *external merge-sort* (Knuth (1998)) to the multiset of k -mers. Basically, the method works by partitioning the multiset of k -mers, and performing several iterations where only a fixed amount of k -mers is loaded in memory, sorted, and re-written on the disk. As a result, we obtain a list of k -mers and its coverage. In order to get rid of most of the sequencing errors, k -mers with a minimal k -mer coverage of mkC (a parameter) are removed. The second step is to actually build the DBG, this is done in the naive way by reading the list of k -mers and adding the corresponding arcs. In the next step, using a greedy non-branching path extension algorithm all maximal non-branching valid paths are found. Then, we obtain the cDBG by merging each path into a single vertex.
2. *Biconnected component (BCC) decomposition*. As stated in Section 1.2.2, a connected undirected graph is biconnected if it remains connected after the removal of any vertex, and a BCC of an undirected graph is a maximal biconnected subgraph. Moreover, as stated in Lemma 1.2 the BCCs of an undirected graph form a partition of the edges with two important properties: every cycle is contained in exactly one BCC, and every edge not contained in a cycle forms a singleton BCC.

From Definition 2.4, it is clear that every bubble in a cDBG corresponds to a cycle in the underlying undirected graph. Thus, applying on the underlying undirected graph of the cDBG Tarjan's lowpoint method (Tarjan (1972)) which performs a modified depth-first search traversal of the graph, Step 2 detects all BCCs, and discards the ones with less than 4 vertices, they cannot contain any bubble. Without modifying the results, this considerably reduces the memory footprint and the computation time of the whole process. To give an idea of the effectiveness of this step, the cDBG of a 5M reads dataset had 1.7M vertices, but the largest BCC only 2961 vertices.

3. *Simple bubbles compression*. Single substitution events (SNPs, sequencing errors) generate a large number of cycles themselves included into bigger ones, creating a combinatorial explosion of the number of possible bubbles. This step of KISSPLICE detects and compresses all bubbles composed of just four vertices: two switching vertices and

two *non-branching internal vertices* each corresponding to sequences differing by just one position. Fig. 2.6(a) shows an example of a simple bubble. Simple bubbles are output as potential SNPs and then replaced by a single vertex in the graph. The two non-branching internal vertices are merged into one, storing a consensus sequence where the unique substitution is replaced by N.

4. *Bubble enumeration.* The cycles are detected in the cDBG using a simple backtracking procedure proposed by [Tiernan \(1970\)](#), which is an unconstrained DFS augmented with four pruning criteria. Indeed, from a path prefix $\pi = s \rightsquigarrow u$ the algorithm recursively explores the vertices of $N^+(u)$ minus the internal vertices of π . Every time a new vertex $v \in N^+(u)$ is added to π the algorithm checks whether: $\pi \cdot (u, v)$ contains more than two switching vertices, the length of the shorter path is greater than $2k - 2$, the length of the longer path is greater than α , or the length of one of the paths is smaller than $2k - 8$; if any of the conditions is satisfied the algorithm stops the recursion on that branch. On the other hand, if $\pi \cdot (u, v)$ is a cycle, i.e. $v = s$, and it satisfies the conditions (i) to (iv) the algorithm outputs a bubble.

This approach has the same theoretical time complexity as Tiernan’s algorithm for cycle listing, i.e. in the worst case the complexity is proportional to the number of paths in the graph, which might be exponential in the size of the graph and the number of bubbles. Tiernan’s algorithm is worse than Tarjan’s ([Tarjan \(1973\)](#)) or Johnson’s ([Johnson \(1975\)](#)) polynomial delay algorithms, but it appears to be not immediate how to use the pruning criteria with them while preserving their theoretical complexity. Moreover, the pruning criteria are very effective for the type of instances we are dealing with. In practice, Tiernan’s algorithm with prunings is faster than a complete cycle listing using Tarjan’s or Johnson’s with a post-processing step to check the four conditions.

5. *Results filtration and classification.* The two paths of each bubble are aligned. If the whole of the shorter path aligns with high similarity to the longer path, we decide that the bubble is due to inexact repeats (see Section 2.2.1). After this alignment, a bubble is classified either as an SNP, AS event, repeat associated bubble, or a small indel.
6. *Read coherence and coverage computation.* Reads from each input dataset are mapped to each path of the bubble. If at least one nucleotide of a path is covered by no read, the bubble is said to be not *read-coherent* and is discarded. The coverage of each position of the bubble corresponds to the number of reads overlapping this position.

2.3 Results

2.3.1 Simulated data

In order to assess the sensitivity and specificity of our approach, we simulated the sequencing of genes for which we are able to control the number of alternative transcripts. We show that the method is indeed able to recover AS events whenever the alternative transcripts are sufficiently expressed. For our sensitivity tests, we used simulated RNA-seq single end reads (75 bp) with sequencing errors. We first tested a pair of transcripts with a 200 nt skipped exon. Simulated reads were obtained with MetaSim ([Richter et al. \(2008\)](#)) which is a reference software for simulating sequencing experiments. As in real experiments, it produces heterogeneous coverage and authorizes to use realistic error models.

In order to find the minimum coverage for which we are able to work, we created datasets for several coverages (from 4X to 20X, which corresponds to 60 to 300 Reads Per Kilobase or RPK for short), with 3 repetitions for each coverage, and tested them with different values of k ($k = 13, \dots, 41$). The purpose of using 3 repetitions for each coverage was to obtain results which did not depend on irreproducible coverage biases. For coverages below 8X (120 RPK), KISSPLICE found the correct event in some but not all of the 3 tested samples. The failure to detect the event was due to the heterogeneous and thus locally very low coverage around the skipped exon, e.g. some nucleotides were not covered by any read or the overlap between the reads was smaller than $k-1$. Above 8X (120 RPK), KISSPLICE detected the correct exon skipping event in all samples.

For each successful test, there was a maximal value k_{max} for k above which the event was not found, and a minimal value k_{min} below which KISSPLICE also reported false positive events. Indeed, if k is too small, then the pattern $ab, as'b$, with $|a| \geq k, |b| \geq k$ is more likely to occur by chance in the transcripts, therefore generating a bubble in the DBG. Between these two thresholds, KISSPLICE found only one event: the correct one. The values of k_{min} and k_{max} are clearly dependent on the coverage of the gene. At 8X (120 RPK), the 200 nucleotides exon was found between $k_{min} = 17$ and $k_{max} = 29$. At 20X (300 RPK), it was found for $k_{min} = 17$ and $k_{max} = 39$. We performed similar tests on other datasets, varying the length of the skipped exon. As expected, if the skipped exon is shorter (longer), KISSPLICE needed a lower (higher) coverage to recover it.

Since KISSPLICE is, to our knowledge, the first method able to call AS events without a reference genome, it cannot be easily benchmarked against other programs. Here, we compare it to a general purpose transcriptome assembler, TRINITY (Grabherr et al. (2011)). Both methods are compared only on the specific task of AS event calling. The current version of TRINITY being restricted to a fixed value of $k = 25$, we systematically verified that this value was included in $[k_{min}, k_{max}]$.

We found out that TRINITY was able to recover the AS event in all 3 samples only when the coverage was above 18X (270 RPK), which clearly shows that KISSPLICE is more sensitive for this task. This can be explained by the fact that TRINITY uses heuristics which tend to over-simplify the cDBG.

All these results were obtained using a minimal k -mer coverage (mkC for short) of 1. We also tested with $mkC = 2$ (i.e. k -mers present only once in the dataset are discarded), leading to the same main behavior. We noticed however a loss in sensitivity for both methods, but a significant gain in the running time. KISSPLICE found the event in all 3 samples for a coverage of 12X (180 RPK) which remains better than the sensitivity of TRINITY for $mkC = 1$.

2.3.2 Real data

We further tested our method on RNA-seq data from human. Even though we do not use any reference genome in our method, we applied it to cases where an annotated reference genome is indeed available in order to be able to assess if our predictions are correct.

We ran KISSPLICE with $k = 25$ and $mkC = 2$ on a dataset which consists of 32M reads from human brain and 39M reads from liver from the Illumina Body Map 2.0 Project (downloaded from the Sequence Read Archive, accession number ERP000546). As in all DBG based assemblers, the most memory consuming step was the DBG construction which we performed on a cluster. The memory requirement is directly dependent on the number of unique k -mers in the dataset.

Despite the fact that we do not use any heuristic to discard k -mers (except for the minimum

coverage threshold) from our index, our memory performances are very similar to the ones of Inchworm, the first step of TRINITY, as indicated in Fig. 2.11a. In addition, for the specific task of calling AS events, KISSPLICE is faster than TRINITY as shown in Fig. 2.11b.

KISSPLICE identified 5923 biconnected components which contained at least one bubble, 664 of which consisted of bubbles generated by repeats associated events and 1160 which consisted of bubbles generated by short indels (less than 3 nt). Noticeably, the BCCs which generated most cycles and were most time consuming were associated to repeats. As these bubbles are not of interest for KISSPLICE, this observation prompted us to introduce an additional parameter in KISSPLICE to stop the computation in a BCC if the number of cycles being enumerated reaches a threshold. This enabled us to have a significant gain of time.

Out of the 4099 remaining BCCs, we found that 3657 were read-coherent (i.e. each nucleotide is covered by at least one read) and we next focused on this set. For each of the 3657 cases, we tried to align the two paths of each bubble to the reference genome using Blat (Kent (2002)). If the two paths align with the same initial and final coordinates, then we consider that the bubble is a real AS event. If they align with different initial and final coordinates, then we consider that it is a false positive. Out of the 3657 BCCs, 3497 (95%) corresponded to real AS events, while the remaining corresponded to false positives. A first inspection of these false positives led to the conclusion that the majority of them correspond to chimeric transcripts. Indeed, the shorter path and the longer path both map in two blocks within the same gene, but the second block is either upstream of the first block, or on the reverse strand, in both cases contradicting the annotations and therefore suggesting that the transcripts are chimeric and could have been generated by a genomic rearrangement or a trans-splicing mechanism.

For each of the 3497 real cases, we further tried to establish if they corresponded to annotated splicing events. We therefore first computed all annotated AS events using AStalavista (Sammeth et al. (2008)) and the UCSC Known Genes annotation (Kuhn et al. (2009)). Then, for each aligned bubble, we checked if the coordinates of the aligned blocks matched the splice sites of the annotated AS events. If the answer was positive, then we considered that the AS event we found was known, otherwise we considered it was novel. Out of a total of 3497 cases, we find that only 1538 are known while 1959 are novel. This clearly shows that current annotations largely underestimate the number of alternative transcripts per multi-exon genes as was also reported recently (Wang et al. (2008)).

Additionally, we noticed that 719 BCCs contained more than one AS event, which all mapped to the same gene. This corresponds to complex splicing events which involve more than 2 transcripts. Such events have been described in Sammeth (2009). Their existence suggests that more complex models could be established to characterize them as one single event, and not as a collection of simple pairwise events. An example of novel complex AS event is given in Fig. 2.7.

We also found the case where the same AS event maps to multiple locations on the reference genome (423 cases). We think these correspond to families of paralogous genes, which are “collectively” alternatively spliced. We were able to verify this hypothesis on all tested instances. In this case, we are unable to decide which of the genes of the family are producing the alternative transcripts, but we do detect an AS event.

2.3.3 Characterization of novel AS events

In order to further characterize the 1959 novel AS events we found, we compared them with annotated events considering their abundance, length of the variable region and use of splice

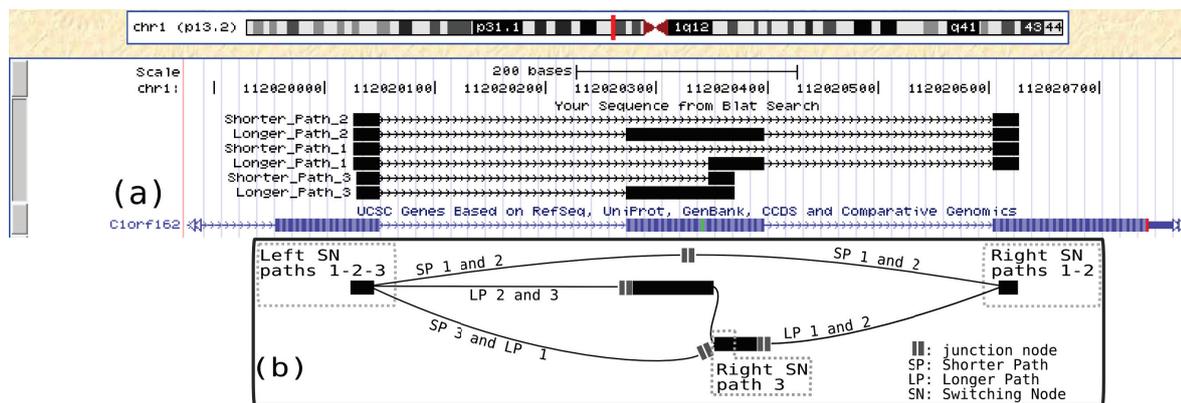


Figure 2.7: BCC corresponding to a novel complex AS event. The intermediate annotated exon is either present, partially present, or skipped. (a) The annotations (blue track) report only the version where it is present while black tracks report all events found by KISSPLICE. (b) The cDBG associated to this complex event where the junction vertices are composed by $2k - 2$ nucleotides.

sites. For each AS event, we have 4 abundances, one for each spliceform (i.e. path of the bubble), and one for each condition. We computed the abundance of an event as the abundance of the minor spliceform. As outlined in Fig. 2.8, we show that novel events are less abundant than annotated events. This in itself could be one of the reasons why they had not been annotated so far. Interestingly, we also found that while annotated events are clearly more expressed in brain than liver (median coverage, in reads per nucleotide, of 3.4 Vs 1.2), this trend was weaker for novel events (2.4 Vs 1.2). This may reflect the fact that, since tissue-specific splicing in brain has been intensely studied, annotations may be biased in their favor.

We then computed the length of each event as the difference of the length between the two paths of the bubble. We found that for annotated events, there is a clear preference (59%) for lengths that are a multiple of 3, which is expected if the event affects a coding region. However, although still very different from random, this preference is less strong for novel events (45%), which, in addition, are particularly enriched in short lengths as shown in Fig. 2.9.

Finally, we computed the splice sites of annotated and novel events, and we found that a vast majority (99.5%) of known events exhibit canonical splice sites, while this is again less strong for novel events (75.3%). Out of the non canonical cases, 13 correspond to U12 introns, but most correspond to short events.

Altogether, while we cannot discard that short non canonical events do occur and have been under-annotated so far, we think that the observations we make on the length and splice site features can be explained by the presence of genomic indels in our results. We had indeed already stated in Section 2.2.1 that while most annotated genomic indels are below 3nt, some may still be above. In order to assess the proportion of bubbles, with length below 10nt, corresponding to indels and AS events, we mapped them to the reference genome. The results are shown in Fig. 2.10. It is clear that bubbles with length smaller than 6nt and not a multiple of 3 are more likely to correspond to genomic indels than AS events. In KISSPLICE (version 2.0) we changed our criterion to classify events with lengths 1, 2, 4, and 5 nt as indels. Moreover, events larger than 10nt have canonical splice sites 92.5% of the cases. More

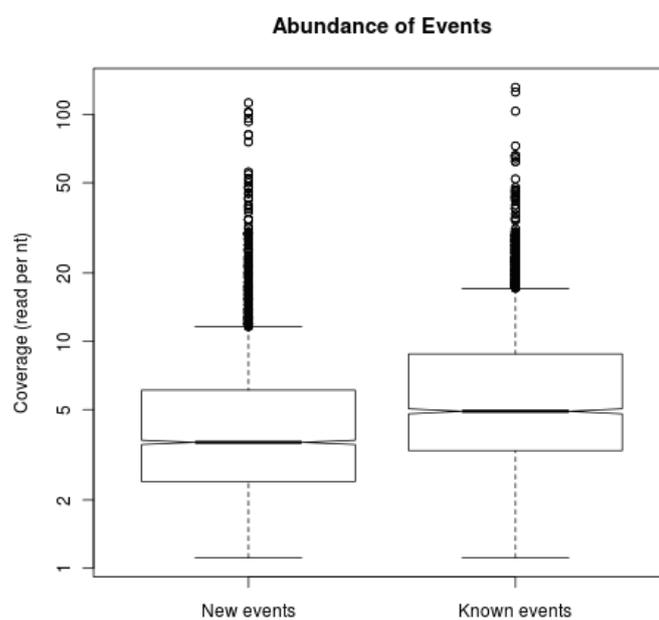


Figure 2.8: Abundance of known and novel events.

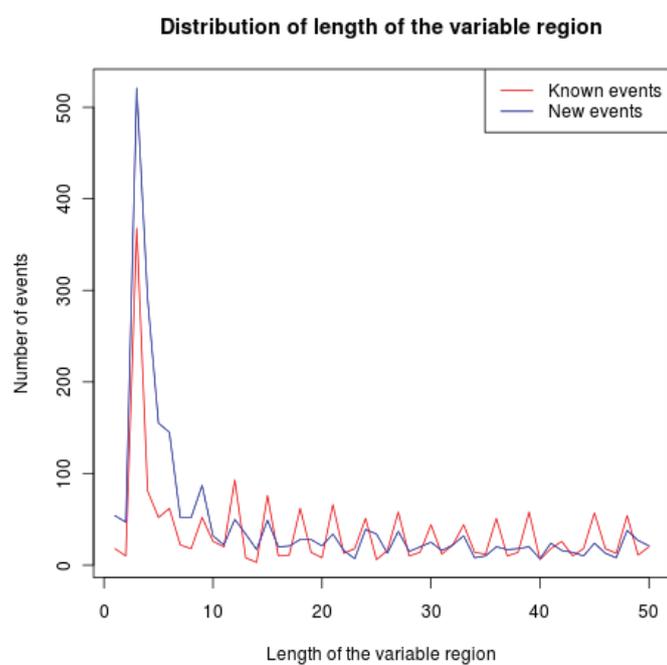


Figure 2.9: Distribution of lengths of the variable regions for known and novel events. Only the initial part of the distribution is given.

generally, we wish to stress that this confusion between genomic indels and AS events is currently being made by all transcriptome assemblers.

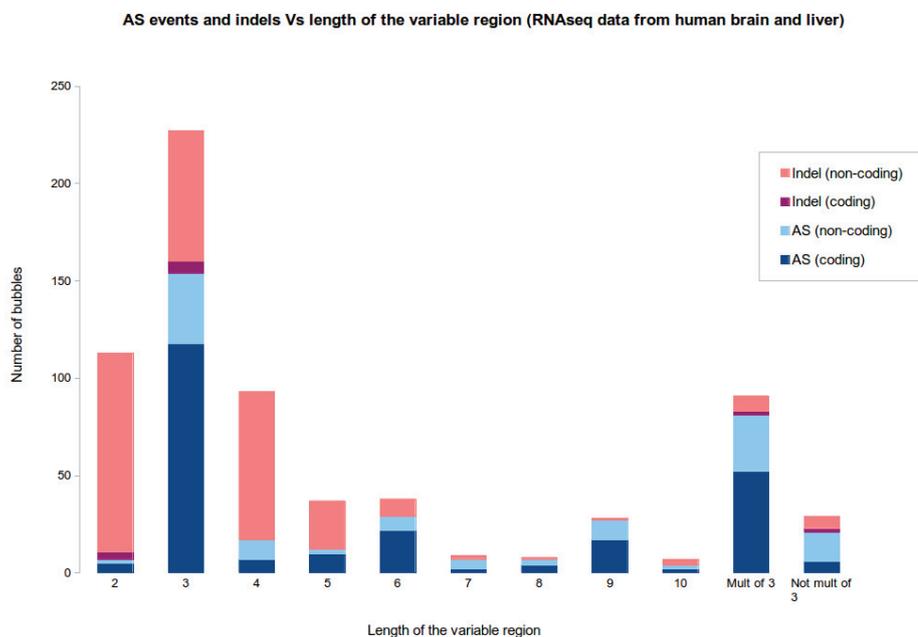


Figure 2.10: Distribution of bubbles corresponding to alternative splicing events and indels, according to the length of the variable region.

2.3.4 Comparison with TRINITY

Finally, in order to further discuss the sensitivity of our method on real data, we compared our results with TRINITY. Although TRINITY is not tailored to find AS events, we managed to retrieve this information from the output. Whenever TRINITY found several alternative transcripts for one gene, we selected this gene. We further focused on cases which contained a cycle in the splicing graph reconstructed from this gene and we compared them with the events found by KISSPLICE. Whenever we found that both the longer and the shorter path of a bubble were mapping to the transcripts of a TRINITY gene, we decided that both methods had found the same event. In total, KISSPLICE found 4099 cases, TRINITY found 1123 out of which 553 were common. While the sensitivity is overall larger for KISSPLICE, we see that 570 cases are found by Trinity and not by KISSPLICE. We then mapped these transcripts to the human genome using Blat. In many instances (348 cases), the transcripts did not align on their entire length, or to different chromosomes, indicating that they corresponded to chimeras. A first inspection of the remaining 222 cases revealed that they correspond to the complex BCCs we chose to neglect at an early stage of the computation, because they contain a very large number of repeat-associated bubbles. A first simple way to deal with this issue is to increase the value of k . The effect of this is to break the large BCCs into computable cases, enabling to recover a good proportion of the missed events. For instance, for $k = 35$, we found back 84 cases. More generally, this shows that more work on the model and on the

algorithms is still required to characterize better AS events which are intricate with inexact repeats. We think that TRINITY manages to identify some of them because it uses heuristics, which enables it to simplify these complex graph structures.

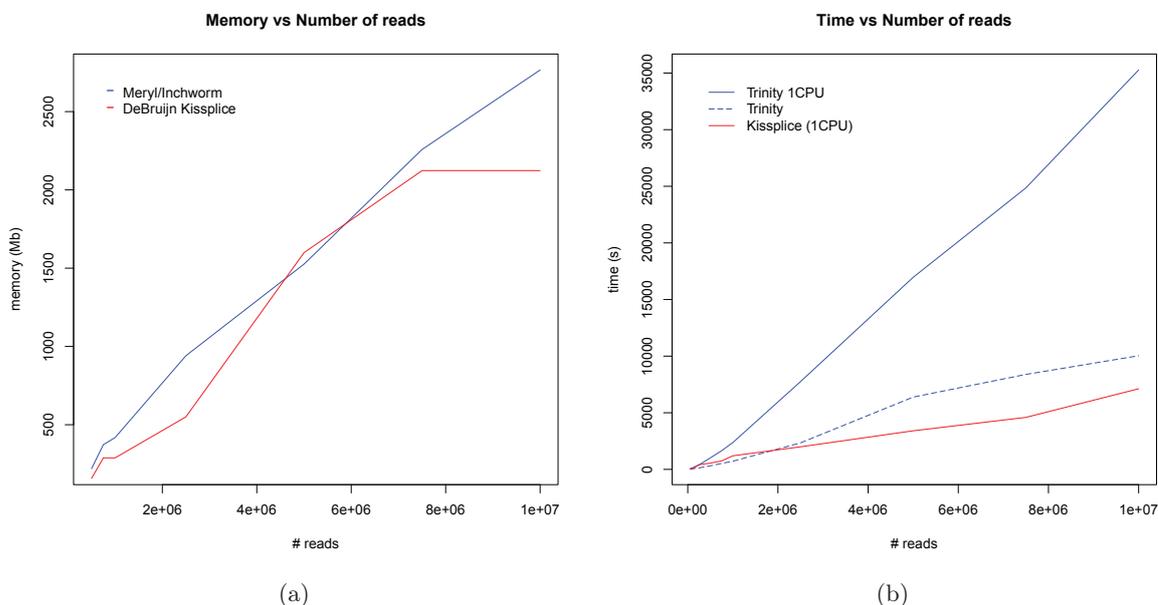


Figure 2.11: (a) Memory usage of KISSPLICE and Inchworm as a function of input size. (b) Time performances of KISSPLICE and TRINITY as a function of input size.

2.4 Discussion and conclusions

This chapter presents two main contributions. First, we introduced a general model for detecting variations in de Bruijn graphs, and second, we developed an algorithm, KISSPLICE, to detect AS events in such graphs. This approach enables to tackle the problem of finding AS events without assembling the full-length transcripts, which may be time consuming and uses heuristics that may lead to a loss of information. To our knowledge, this approach is new and should constitute a useful complement to general purpose transcriptome assemblers.

Results on human data show that this approach enables de novo calling of AS events with a higher sensitivity than obtained by the approaches based on a full assembly of the reads, while using similar memory requirements and less time. 5% of the extracted events correspond to false positives, while the 95% remaining can be separated into known (44%) and novel events (56%). Novel events exhibit similar sequence features as known events as long as we focus on events longer than 10 bp. Below this, novel events seem to be enriched in genomic indels.

KISSPLICE is an user-friendly tool under active development available for download at <http://kissplice.prabi.fr/>, which is mature enough to be used in real life projects to establish a more complete catalog of AS events in any species, whether it has a reference genome or not. Despite the fact that more and more genomes are now being sequenced, the new genome assemblies obtained usually do not reach the level of quality of the ones we have for model organisms. Hence, we think that methods which do not rely on a reference genome

are not going to be easily replaced in the near future.

There is of course room for further improvements. For instance, the current bubble listing algorithm, the core of the KISSPLICE pipeline, is not entirely satisfactory. In Chapter 3, we present a linear delay algorithm to list all cycles satisfying condition (i) of Section 2.2.2, that is to directly list all bubbles in a de Bruijn graph. In Chapter 4, we propose an improved, completely unrelated, polynomial delay algorithm to list all cycles satisfying conditions (i), (ii) and (iv), and experimentally show that this method outperforms the algorithm of Section 2.2.2.

Another point not satisfactory in the initial KISSPLICE (version 1.6) pipeline is the memory consumption. As stated in Section 2.3.2, the memory bottleneck is the de Bruijn graph construction. We address this issue in Chapter 5 where we propose a practical algorithm to build the de Bruijn improving over the state of the art.

In addition, the coverage could be used for distinguishing SNPs from sequencing errors, and the splicing site signature, i.e. canonical splicing sites (Burset et al. (2000)) GT-AG, could be used to distinguish between intron retention and the others AS events. Moreover, the sequences surrounding the bubbles could be locally assembled using a third party tool (Peterlongo and Chikhi (2012)). This would allow to output their context or the full contig they belong to.

Last, the complex structure of BCCs associated to repeats seems to indicate that more work on the model and on the algorithms is required to efficiently deal with the identification of repeat associated bubbles, which may be highly intertwined with other events.

Chapter 3

Listing in unweighted graphs

Contents

3.1 Efficient bubble enumeration in directed graphs	42
3.1.1 Introduction	42
3.1.2 De Bruijn graphs and bubbles	43
3.1.3 Turning bubbles into cycles	44
3.1.4 The algorithm	45
3.1.5 Proof of correctness and complexity analysis	49
3.1.6 Practical speedup	51
3.2 Optimal listing of cycles and <i>st</i>-paths in undirected graphs . . .	52
3.2.1 Introduction	52
3.2.2 Preliminaries	53
3.2.3 Overview and main ideas	54
3.2.4 Amortization strategy	61
3.2.5 Certificate implementation and maintenance	63
3.2.6 Extended analysis of operations	66
3.3 Discussion and conclusions	71

In this chapter, we are mainly concerned with listing problems in *unweighted* graphs. In directed graphs, we consider the problem of listing bubbles, defined as a pair of internally vertex-disjoint paths (Chapter 2). In undirected graphs, we consider the classical problem of listing *st*-path and cycles. The chapter is divided in two main parts.

The first part (Section 3.1) is strongly based on our paper [Birmelé et al. \(2012\)](#). The goal is to show a non-trivial adaptation of Johnson’s cycle¹ listing algorithm ([Johnson \(1975\)](#)) to identify all bubbles in a directed graph maintaining the same complexity. For a directed graph with n vertices and m arcs, containing η bubbles, the method we propose lists all bubbles with a given source in $O((n + m)(\eta + 1))$ total time and $O(m + n)$ delay. For the general problem of listing bubbles, this algorithm is exponentially faster than the algorithm based on Tiernan’s algorithm ([Tiernan \(1970\)](#)) presented in Chapter 2. However, it should be noted that, contrary to Chapter 2, the graph here is not a bidirected de Bruijn graph.

The second part (Section 3.2) is strongly based on our paper [Birmelé et al. \(2013\)](#). The goal is to show an algorithm to list cycles in undirected graphs improving over the state of

¹Johnson uses the term *elementary circuits*.

the art (Johnson’s algorithm). Indeed, we present the first optimal solution to list all the simple cycles in an undirected graph G . Specifically, let $\mathcal{C}(G)$ denote the set of all these cycles. For a cycle $c \in \mathcal{C}(G)$, let $|c|$ denote the number of edges in c . Our algorithm requires $O(m + \sum_{c \in \mathcal{C}(G)} |c|)$ time and is asymptotically optimal: $\Omega(m)$ time is necessarily required to read G as input, and $\Omega(\sum_{c \in \mathcal{C}(G)} |c|)$ time is required to list the output. We also present the first optimal solution to list all the simple paths from s to t in an undirected graph G .

3.1 Efficient bubble enumeration in directed graphs

3.1.1 Introduction

In the previous chapter, a method (KISSPLICE) to identify variants (alternative splicing events, SNPs, indels and inexact tandem repeats) in RNA-seq data without a reference genome was introduced. Each variant corresponds to a recognizable pattern in a (bidirected) de Bruijn graph built from the reads of the RNA-seq experiment. In each case, the pattern corresponds to a bubble defined as two vertex-disjoint paths between a pair of source and target vertices s and t . Properties on the lengths or sequence similarity of the paths then enable to differentiate between the different types of variants.

Bubbles have been studied before in the context of genome assembly (Peng et al. (2010); Li et al. (2010); Simpson et al. (2009); Zerbino and Birney (2008)) where they also have been called bulges (Pevzner et al. (2004)). However, the purpose in these works was not to list all bubbles, but “only” to remove them from the graph in order to provide longer contigs for a genome assembly. More recently, ad-hoc listing methods have been proposed but are restricted to (almost) non-branching bubbles (Peterlongo et al. (2010); Iqbal et al. (2012); Leggett et al. (2013)), i.e. each vertex from the bubble has in-degree and out-degree 1, except for s and t . Furthermore, in all these applications (Pevzner et al. (2004); Zerbino and Birney (2008); Simpson et al. (2009); Peng et al. (2010); Li et al. (2010); Iqbal et al. (2012); Leggett et al. (2013)), since the patterns correspond to SNPs or sequencing errors, the authors only considered paths of length smaller than a constant.

On the other hand, bubbles of arbitrary length have been considered in the context of splicing graphs (Sammeth (2009)). However, in this context, a notable difference is that the graph is a DAG. Additionally, in the case of Iqbal et al. (2012) the vertices are colored and only unicolor paths are then considered for forming bubbles. Finally, the concept of bubble also applies to the area of phylogenetic networks (Gusfield et al. (2004)), where it corresponds to the notion of a recombination cycle. Again for this application, the graph is a DAG. To our knowledge, no enumeration algorithm for recombination cycles has been proposed.

In this chapter, we consider the more general problem of listing all bubbles in an arbitrary directed graph. That is, our solution is not restricted to acyclic or de Bruijn graphs, neither imposes restrictions on the path length or the degrees of the internal nodes. This problem is quite general but it remained an open question whether a polynomial delay algorithm could be proposed for solving it. The algorithm briefly presented in Chapter 2 (also in Sacomoto et al. (2012)) was an adaptation of Tiernan’s algorithm for cycle listing (Tiernan (1970)) which is not polynomial delay. Actually, since in the worst case Tiernan’s algorithm can explore all the st -paths while the graph only contains a constant number of cycles, the algorithm of Chapter 2 is not even polynomial total time. The time spent by the algorithm is, in the worst case, exponential in the size of the input graph and the number of bubbles output.

The first part of this chapter is organized as follows. We start by discussing in Section 3.1.2 the correspondence between bubbles in bidirected de Bruijn graphs (Chapter 2) and directed de Bruijn graphs (Chapter 1). We then explain in Section 3.1.3 how to transform the directed graph where we want to list the bubbles into a new directed graph such that the bubbles correspond to cycles satisfying some extra properties. We present in Section 3.1.4 the algorithm to list all cycles corresponding to bubbles in the initial graph and prove in Section 3.1.5 that this algorithm has linear delay. Finally, we briefly describe, in Section 3.1.6, a slightly more complex version of the algorithm that could lead to a more space and time efficient implementation, but with the same overall complexity.

3.1.2 De Bruijn graphs and bubbles

In the previous chapter, we defined bubbles (Definition 2.4) as a pair of vertex-disjoint valid paths in a bidirected de Bruijn graph. Recall that, a bidirected de Bruijn graph is directed multigraph where each vertex is labeled by a k -mer and its reverse complement and the arcs represent a $k - 1$ suffix-prefix overlap and are labeled depending on which k -mer, forward or reverse, the overlap refers to, whereas a (directed) de Bruijn graph (Definition 1.4) is a directed graph where each vertex is labeled by a k -mer and the arcs correspond to $k - 1$ suffix-prefix overlaps. Here, we consider bubbles in a directed de Bruijn graph.

Definition 3.1 ((s, t) -bubble). *Given a directed graph $G = (V, E)$, an (s, t) -bubble is a pair of internally vertex-disjoint st -paths.*

Both de Bruijn graph definitions are roughly equivalent. Indeed, given a bidirected DBG we transform it into a regular DBG by splitting every vertex in two vertices, one corresponding to the forward k -mer and the other to the reverse k -mer, and maintaining the arcs accordingly. This transformation, however, does not induce a one-to-one correspondence between bubbles in the bidirected DBG (Definition 2.4) and (s, t) -bubbles in the corresponding DBG. Indeed, every valid path in the bidirected DBG corresponds to a simple path in the directed DBG, but the converse is not true, a simple path in the directed DBG containing a k -mer and its reverse complement is not a valid path in the bidirected DBG, implying that, every bubble in a bidirected DBG corresponds to a (s, t) -bubble in the directed DBG, but the converse is not true. We, however, disregard this nonequivalence, since no true bubble is lost by considering the directed DBG.

From now on, we consider the more general problem of listing bubbles in an arbitrary directed graph, not necessary a DBG.

Problem 3.2 (Listing bubbles). *Given a directed graph $G = (V, E)$, output all (s, t) -bubbles in G , for all pairs $s, t \in V$.*

In order to solve Problem 3.2, we consider the problem of listing all bubbles with a given source (Problem 3.3). Indeed, by trying all possible sources s we can list all (s, t) -bubbles.

Problem 3.3 (Listing $(s, *)$ -bubbles). *Given a directed graph $G = (V, E)$ and vertex s , output all (s, t) -bubbles in G , for all $t \in V$.*

The number of vertices and arcs of G is denoted by n and m , respectively.

3.1.3 Turning bubbles into cycles

Let $G = (V, E)$ be a directed graph, and let $s \in V$. We want to find all (s, t) -bubbles for all possible target vertices t . We transform G into a new graph $G'_s = (V'_s, E'_s)$ where $|V'_s| = 2|V|$ and $|E'_s| = O(|V| + |E|)$. Namely,

$$V'_s = \{v, \bar{v} \mid v \in V\}$$

$$E'_s = \{(u, v), (\bar{v}, \bar{u}) \mid (u, v) \in E \text{ and } v \neq s\} \cup \{(v, \bar{v}) \mid v \in V \text{ and } v \neq s\} \cup \{(\bar{s}, s)\}$$

Let us denote by \bar{V} the set of vertices of G'_s that were not already in G , that is $\bar{V} = V'_s \setminus V$. The two vertices $x \in V$ and $\bar{x} \in \bar{V}$ are said to be *twin vertices*. Observe that the graph G'_s is thus built by adding to G a reversed copy of itself, where the copy of each vertex is referred to as its *twin*. The arcs incoming to s (and outgoing from \bar{s}) are not included so that the only cycles in G'_s that contain s also contain \bar{s} . New arcs are also created between each pair of twins: the new arcs are the ones leading from a vertex u to its twin \bar{u} for all u except for s where the arc goes from \bar{s} to s . An example of a transformation is given in Figure 3.1.

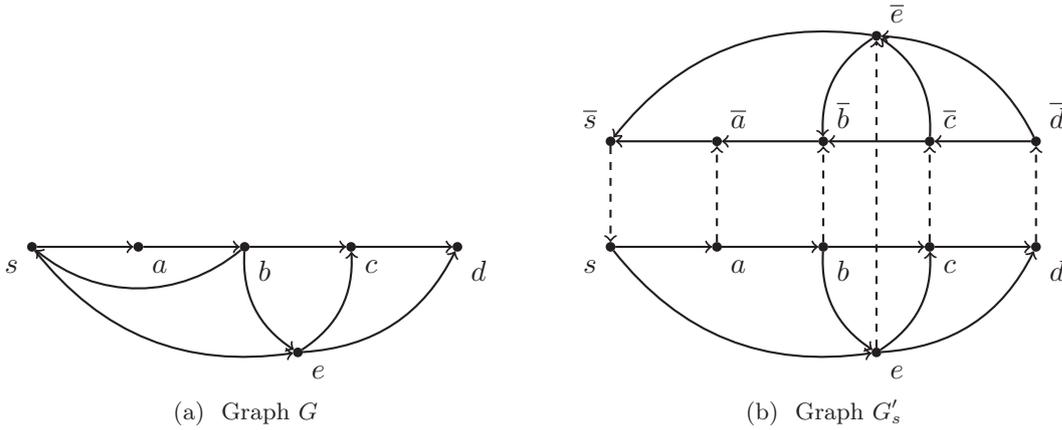


Figure 3.1: Graph G and its transformation G'_s . We have that $\langle s, e, \bar{e}, \bar{b}, \bar{a}, \bar{s}, s \rangle$ is a bubble-cycle with swap arc (e, \bar{e}) that has a correspondence to the (s, e) -bubble composed by the two vertex-disjoint paths $\langle s, e \rangle$ and $\langle s, a, b, e \rangle$.

We define a cycle of G'_s as being *bipolar* if it contains vertices of both V and \bar{V} . As the only arc from \bar{V} to V is (\bar{s}, s) , then every bipolar cycle C contains also only one arc from V to \bar{V} . This arc, which is the arc (t, \bar{t}) for some $t \in V$, is called the *swap arc* of C . Moreover, since (\bar{s}, s) is the only incoming arc of s , all the cycles containing s are bipolar. We say that C is *twin-free* if it contains no pair of twins except for (s, \bar{s}) and (t, \bar{t}) .

Definition 3.4 (Bubble-cycle). A bubble-cycle in G'_s is a twin-free cycle of size greater than four².

Proposition 3.5. Given a vertex s in G , there is a one-to-two correspondence between the set of (s, t) -bubbles in G for all $t \in V$, and the set of bubble-cycles of G'_s .

²The only twin-free cycles in of size four in G'_s are generated by the outgoing edges of s . There are $O(|V|)$ of such cycles.

Proof. Let us consider an (s, t) -bubble in G formed by two vertex-disjoint st -paths P and Q . Consider the cycle of G'_s obtained by concatenating P (resp. Q), the arc (t, \bar{t}) , the inverted copy of Q (resp. P), and the arc (\bar{s}, s) . Both cycles are bipolar, twin-free, and have (t, \bar{t}) as swap arc. Therefore both are bubble-cycles.

Conversely, consider any bubble-cycle C and let (t, \bar{t}) be its swap arc. C is composed by a first subpath P from s to t that traverses vertices of V and a second subpath \bar{Q} from \bar{t} to \bar{s} composed of vertices of \bar{V} only. By definition of G'_s , the arcs of the subpath P form a path from s to t in the original graph G ; given that the vertices in the subpath \bar{Q} from \bar{t} to \bar{s} are in \bar{V} and use arcs that are those of E inverted, then \bar{Q} corresponds to another path from s to t of the original graph G . As no internal vertex of \bar{Q} is a twin of a vertex in P , these two paths from s to t are vertex-disjoint, and hence they form an (s, t) -bubble.

Notice that there is a cycle s, v, \bar{v}, \bar{s} for each v in the out-neighborhood of s . Such cycles do not correspond to any bubble in G , and the condition on the size of C allows us to rule them out. \square

3.1.4 The algorithm

Johnson (1975) introduced a polynomial delay algorithm for the cycle enumeration problem in directed graphs. We propose to adapt the principle of this algorithm, the pruned backtracking, to enumerate bubble-cycles in G'_s . Indeed, we use a similar pruning strategy, modified to take into account the twin nodes. Proposition 3.5 then ensures that running our algorithm on G'_s for every $s \in V$ is equivalent to the enumeration of (twice) all the bubbles of G . To do so, we explore G'_s by recursively traversing it while maintaining the following three variables. We denote by $N^+(v)$ the set of out-neighbors and $N^-(v)$ as the set of in-neighbors of v .

1. A variable *stack* which contains the vertices of a path (with no repeated vertices) from s to the current vertex. Each time it is possible to reach \bar{s} from the current vertex by satisfying all the conditions to have a bubble-cycle, this stack is completed into a bubble-cycle and its content output.
2. A variable *status*(v) for each vertex v which can take three possible values:
 - free*: v should be explored during the traversal of G'_s ;
 - blocked*: v should not be explored because it is already in the stack or because it is not possible to complete the current stack into a cycle by going through v – notice that the key idea of the algorithm is that a vertex may be blocked without being on the stack, avoiding thus useless explorations;
 - twinned*: $v \in \bar{V}$ and its twin is already in the stack, so that v should not be explored.
3. A set $B(v)$ of in-neighbors of v where vertex v is blocked and for each vertex $w \in B(v)$ there exists an arc (w, v) in G'_s (that is, $w \in N^-(v)$). If a modification in the stack causes that v is unblocked and it is possible to go from v to \bar{s} using free vertices, then w should be unblocked if it is currently blocked.

Algorithm 3.1 enumerates all the bubble-cycles in G (Problem 3.2) by fixing the source s of the (s, t) -bubble, computing the transformed graph G'_s and then listing all bubble-cycles with source s in G'_s (Problem 3.3). This procedure is repeated for each vertex $s \in V$. To list the bubble-cycles with source s , procedure `cycle`(s) is called. As a general approach, Algorithm 3.3 uses classical backtracking with a pruned search tree. The root of the recursion

corresponds to the enumeration of all bubble-cycles in G'_s with starting point s . The algorithm then proceeds recursively: for each free out-neighbor w of v the algorithm enumerates all bubble-cycles that have the vertices in the current stack plus w as a prefix. If $v \in V$ and \bar{v} is twinned, the recursion is also applied to the current stack plus \bar{v} , (v, \bar{v}) becoming the current swap arc. A base case of the recursion happens when \bar{s} is reached and the call to `cycle(\bar{s})` completed. In this case, the path in *stack* is a twin-free cycle and, if this cycle has more than 4 vertices, it is a bubble-cycle to output.

The key idea that enables to make this pruned backtracking efficient is the block-unblock strategy. Observe that when `cycle(v)` is called, v is pushed in the stack and to ensure twin-free extensions, v is blocked and \bar{v} is twinned if $v \in V$. Later, when backtracking, v is popped from the stack but it is *not necessarily* marked as free. If there were no twin-free cycles with the vertices in the current stack as a prefix, the vertex v would remain blocked and its status would be set to free only at a later stage. The intuition is that either v is a dead-end or there remain vertices in the stack that block all twin-free paths from v to \bar{s} . In order to manage the status of the vertices, the sets $B(w)$ are used. When a vertex v remains blocked while backtracking, it implies that every out-neighbor w of v has been previously blocked or twinned. To indicate that each out-neighbor $w \in N^+(v)$ (also, $v \in N^-(w)$ is an *in-neighbor* of w) blocks vertex v , we add v to each $B(w)$. When, at a later point in the recursion, a vertex $w \in N^+(v)$ becomes unblocked, v must also be unblocked as possibly there are now bubble-cycles that include v . Algorithm 3.2 implements this recursive unblocking strategy.

Algorithm 3.1: Main algorithm

```

1 for  $s \in V$  do
2   stack =  $\emptyset$ 
3   for  $v \in G'_s$  do
4     status( $v$ ) = free
5      $B(v) = \emptyset$ 
6   cycle( $s$ )

```

Algorithm 3.2: Procedure `unblock(v)`

```

/* recursive unblocking of vertices for which popping  $v$  creates a path to
    $\bar{s}$  */
1 status( $v$ ) = free
2 for  $w \in B(v)$  do
3   delete  $w$  from  $B(v)$ 
4   if status( $w$ ) = blocked then
5     unblock( $w$ )

```

An important difference between the algorithm introduced here and Johnson's is that we now have three possible states for any vertex, *i.e.* free, blocked and twinned, instead of only the first two. The twinned state is necessary to ensure that the two paths of the bubble share no internal vertex. Whenever \bar{v} is twinned, it can only be explored from v . On the other hand, a blocked vertex should never be explored. A twin vertex \bar{v} can be already blocked when the algorithm is exploring v , since it could have been unsuccessfully explored by some other call. In this case, it is necessary to verify the status of \bar{v} , as it is shown in the graph

Algorithm 3.3: Procedure $\text{cycle}(v)$

```

1  $f = \text{false}$ 
2 push  $v$ 
3  $\text{status}(v) = \text{blocked}$ 
  /* Exploring forward the edges going out from  $v \in V$  */
4 if  $v \in V$  then
5   if  $\text{status}(\bar{v}) = \text{free}$  then
6      $\text{status}(\bar{v}) = \text{twinned}$ 
7   for  $w \in N^+(v) \cap V$  do
8     if  $\text{status}(w) = \text{free}$  then
9       if  $\text{cycle}(w)$  then
10         $f = \text{true}$ 
11   if  $\text{status}(\bar{v}) = \text{twinned}$  then
12     if  $\text{cycle}(\bar{v})$  then
13        $f = \text{true}$ 
  /* Exploring forward the edges going out from  $v \in \bar{V}$  */
14 else
15   for  $w \in N^+(v)$  do
16     if  $w = \bar{s}$  then
17       output the cycle composed by the stack followed by  $\bar{s}$  and  $s$ 
18        $f = \text{true}$ 
19     else if  $\text{status}(w) = \text{free}$  then
20       if  $\text{cycle}(w)$  then
21          $f = \text{true}$ 
22 if  $f$  then
23    $\text{unblock}(v)$ 
24 else
25   for  $w \in N^+(v)$  do
26     if  $v \notin B(w)$  then
27        $B(w) = B(w) \cup \{v\}$ 
28 pop  $v$ 
29 return  $f$ 

```

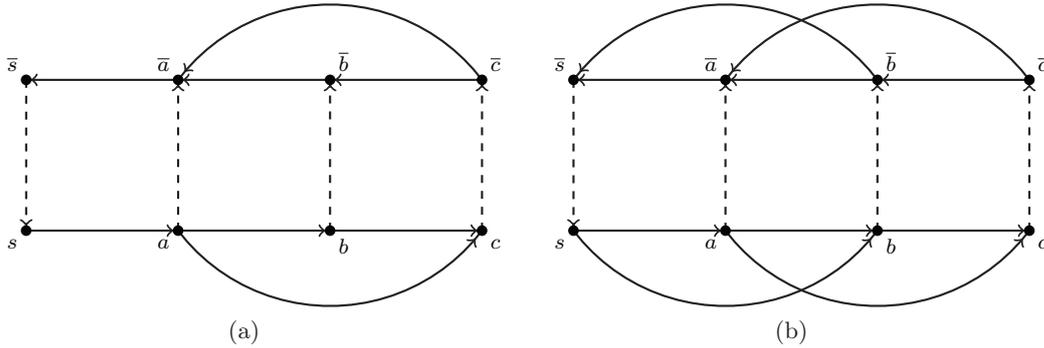


Figure 3.2: (a) Example where the twin \bar{v} is already blocked when the algorithm starts exploring v . By starting in s and visiting first (s, a) and (a, b) , the vertex \bar{c} is already blocked when the algorithm starts exploring c . (b) Counterexample for the variant of the algorithm visiting first the twin and then the regular neighbors. By starting in s and visiting first (s, a) and (a, b) , the algorithm misses the bubble-cycle $\langle s, a, c, \bar{c}, \bar{b}, \bar{s} \rangle$.

of Figure 3.2a. Indeed, consider the algorithm starting from s with (s, a) and (a, b) being the first two arcs visited in the lower part. Later, when the calls `cycle(\bar{c})` and `cycle(\bar{b})` are made, since \bar{a} is twinned, both \bar{b} and \bar{c} remain blocked. When the algorithm backtracks to a and explores (a, c) , the call `cycle(c)` is made and \bar{c} is already blocked.

Another important difference with respect to Johnson's algorithm is that there is a specific order in which the out-neighborhood of a vertex should be explored. In particular, notice that the order in which Algorithm 3.3 explores the neighbors of a vertex v is: first the vertices in $N^+(v) \setminus \{\bar{v}\}$ and then \bar{v} . A variant of the algorithm where this order would be reversed, visiting first \bar{v} and then the vertices in $N^+(v) \setminus \{\bar{v}\}$, would fail to enumerate all the bubbles. Indeed, intuitively a vertex can be blocked because the only way to reach \bar{s} is through a twinned vertex and when that vertex is untwinned the first one is not unblocked. Indeed, consider the graph in Figure 3.2b and the twin-first variant starting in s with (s, a) and (a, b) being the first two arcs explored in the lower part of the graph. When the algorithm starts exploring b the stack contains $\langle s, a, b \rangle$. After, the call `cycle(\bar{b})` returns *true* and `cycle(c)` returns *false* because \bar{a} and \bar{b} are twinned. After finishing exploring b , the blocked list $B(b)$ is empty. Thus, the only vertex unblocked is b, c (and \bar{c}) remaining blocked. Finally, the algorithm backtracks to a and explores the edge (a, c) , but c is blocked, and it fails to enumerate $\langle s, a, c, \bar{c}, \bar{b}, \bar{s} \rangle$.

One way to address the problem above would be to modify the algorithm so that every time a vertex \bar{v} is untwinned, a call to `unblock(\bar{v})` is made. All the bubble-cycles would be correctly enumerated. However, in this case, it is not hard to find an example where the delay would then no longer be linear. Intuitively, visiting first $N^+(v) \setminus \{\bar{v}\}$ and then \bar{v} , works because every vertex u that was blocked (during the exploration of $N^+(v) \setminus \{\bar{v}\}$) should remain blocked when the algorithm explores \bar{v} . Indeed, a bubble would be missed only if there existed a path starting from \bar{v} , going to \bar{s} through u and avoiding the twinned vertices. This is not possible if no path from $N^+(v) \setminus \{\bar{v}\}$ to u could be completed into a bubble-cycle by avoiding the twinned vertices, as we will show later on.

3.1.5 Proof of correctness and complexity analysis

Proof of correctness: Algorithm 3.3 enumerates all bubbles with source s

Lemma 3.6. *Let v be a vertex of G'_s such that $\text{status}(v) = \text{blocked}$, S the set of vertices currently in the stack, and T the set of vertices whose status is equal to twinned. Then $S \cup T$ is a (v, \bar{s}) separator, that is, each path, if any exists, from v to \bar{s} contains at least one vertex in $S \cup T$.*

Proof. The result is obvious for the vertices in $S \cup T$. Let v be a vertex of G'_s such that $\text{status}(v) = \text{blocked}$ and $v \notin S \cup T$. This means that when v was popped for the last time, $\text{cycle}(v)$ was equal to *false* since v remained blocked.

Let us prove by induction on k that each path to \bar{s} of length k from a blocked vertex not in $S \cup T$ contains at least one vertex in $S \cup T$.

We first consider the base case $k = 1$. Suppose that v is a counter-example for $k = 1$. This means that there is an arc from v to \bar{s} (\bar{s} is an out-neighbor of v). However, in that case the output of $\text{cycle}(v)$ is *true*, a contradiction because v would then be unblocked.

Suppose that the result is true for $k - 1$ and, by contradiction, that there exists a blocked vertex $v \notin S \cup T$ and a path (v, w, \dots, \bar{s}) of length k avoiding $S \cup T$. Since (w, \dots, \bar{s}) is a path of length $k - 1$, we can then assume that w is free. Otherwise, if w were blocked, by induction, the path (w, \dots, \bar{s}) would contain at least one vertex in $S \cup T$, and so would the path (v, w, \dots, \bar{s}) .

Since the call to $\text{cycle}(v)$ returned *false* (v remained blocked), either w was already blocked or twinned, or the call to $\text{cycle}(w)$ made inside $\text{cycle}(v)$ gave an output equal to *false*. In any case, after the call to $\text{cycle}(v)$, w was blocked or twinned and v put in $B(w)$.

The conditional at line 11 of the cycle procedure ensures that when untwinned, a vertex immediately becomes blocked. Thus, since w is now free, a call to $\text{unblock}(w)$ was made in any case, yielding a call to $\text{unblock}(v)$. This contradicts the fact that v is blocked. \square

Theorem 3.7. *The algorithm returns only bubble-cycles. Moreover, each of those cycles is returned exactly once.*

Proof. Let us first prove that only bubble-cycles are output. As any call to unblock (either inside the procedure cycle or inside the procedure unblock itself) is immediately followed by the popping of the considered vertex, no vertex can appear twice in the stack. Thus, the algorithm returns only cycles. They are trivially bipolar as they have to contain s and \bar{s} to be output.

Consider now a cycle C output by the algorithm with swap arc (t, \bar{t}) . Let (v, w) in C with $v \neq s$ and $v \neq t$. If \bar{v} is free when v is put on the stack, then \bar{v} is twinned before w is put on the stack and cannot be explored until w is popped. If \bar{v} is blocked when v is put on the stack, then by Lemma 3.6 it remains blocked at least until v is popped. Thus, \bar{v} cannot be in C , and consequently the output cycles are twin-free.

So far we have proven that the output produces bubble-cycles. Let us now show that all cycles $C = \{v_0 = s, v_1, \dots, v_{l-1}, v_l = \bar{s}, v_0\}$ satisfying those conditions are output by the algorithm, and each is output exactly once.

The fact that C is not returned twice is a direct consequence of the fact that the stack is different in all the leaves of a backtracking procedure. To show that C is output, let us prove by induction that the stack is equal to $\{v_0, \dots, v_i\}$ at some point of the algorithm, for every $0 \leq i \leq l - 1$. Indeed, it is true for $i = 0$. Moreover, suppose that at some point, the stack is $\{v_0, \dots, v_{i-1}\}$.

Suppose that v_{i-1} is different from t . As the cycle contains no pair of twins except for those composing the arcs (s, \bar{s}) and (t, \bar{t}) , the path $\{v_i, v_{i+1}, \dots, v_l\}$ contains no twin of $\{v_0, \dots, v_{i-1}\}$ and therefore no twinned vertex. Thus, it is a path from v_i to \bar{s} avoiding $S \cup T$. Lemma 3.6 then ensures that at this point v_i is not blocked. As it is also not twinned, its status is free. Therefore, it will be explored by the backtracking procedure and the stack at some point will be $\{v_0, \dots, v_i\}$. If $v_{i-1} = t$, $v_i = \bar{t}$ is not blocked using the same arguments. Thus it was twinned by the call to `cycle(t)` and is therefore explored at Line 12 of this procedure. Again, the stack at some point will be $\{v_0, \dots, v_i\}$. \square

Analysis of complexity: Algorithm 3.3 has linear delay

As in Johnson (1975), we show that Algorithm 3.3 has delay $O(|V| + |E|)$ by proving that a cycle has to be output between two successive unblockings of the same vertex and that with linear delay some vertex has to be unblocked again. To do so, let us first prove the following lemmas.

Lemma 3.8. *Let v be a vertex such that `cycle(v)` returns true. Then a cycle is output after that call and before any call to `unblock`.*

Proof. Let y be the first vertex such that `unblock(y)` is called inside `cycle(v)`. Since `cycle(v)` returns true, there is a call to `unblock(v)` before it returns, so that y exists. Certainly, `unblock(y)` was called before `unblock(v)` if $y \neq v$. Moreover, the call `unblock(y)` was done inside `cycle(y)`, from line 23, otherwise it would contradict the choice of y . So, the call to `cycle(y)` was done within the recursive calls inside the call to `cycle(v)`. `cycle(y)` must then return true as y was unblocked from it.

All the recursive calls `cycle(z)` made inside `cycle(y)` must return false, otherwise there would be a call to `unblock(z)` before `unblock(y)`, contradicting the choice of y . Since `cycle(y)` must return true and the calls to all the neighbors returned false, the only possibility is that $\bar{s} \in N^+(y)$. Therefore, a cycle is output before `unblock(y)`. \square

Lemma 3.9. *Let v be a vertex such that there is a $v\bar{s}$ -path P avoiding $S \cup T$ at the moment a call to `cycle(v)` is made. Then the return value of `cycle(v)` is true.*

Proof. First notice that if there is such a path P , then v belongs to a cycle in G'_s . This cycle may however not be a bubble-cycle in the sense that it may not be twin-free, that is, it may contain more than two pairs of twin vertices. Indeed, since the only constraint that we have on P is that it avoids all vertices that are in S and T when v is reached, then if $v \in V$, it could be that the path P from v to \bar{s} contains, besides s and \bar{s} , at least two more pairs of twin vertices. An example is given in Figure 3.1b. It is however always possible, by construction of G'_s from G , to find a vertex $y \in V$ such that y is the first vertex in P with \bar{y} also in P . Let P' be the path that is a concatenation of the subpath $s \rightsquigarrow y$ of P , the arc (y, \bar{y}) , and the subpath $\bar{y} \rightsquigarrow \bar{s}$ in P . This path is twin-free, and a call to `cycle(v)` will, by correctness of the algorithm, return true. \square

Theorem 3.10. *Algorithm 3.3 has linear delay.*

Proof. Let us first prove that between two successive unblockings of any vertex v , a cycle is output. Let w be the vertex such that a call to `unblock(w)` at line 23 of Algorithm 3.3 unblocks v for the first time. Let S and T be, respectively, the current sets of stack and twinned vertices after popping w . The recursive structure of the unblocking procedure then ensures that there exists a vw -path avoiding $S \cup T$. Moreover, as the call to `unblock(w)`

was made at line 23, the answer to `cycle(w)` is *true* so there exists also a $w\bar{s}$ -path avoiding $S \cup T$. The concatenation of both paths is again a $v\bar{s}$ -path avoiding $S \cup T$. Let x be the first vertex of this path to be visited again. Note that, if no vertex in this path is visited again there is nothing to prove, since v is free, `cycle(v)` needs to be called before any `unlock(v)` call. When `cycle(x)` is called, there is a $x\bar{s}$ -path avoiding the current $S \cup T$ of stack and twinned vertices. Thus, applying Lemma 3.9 and then Lemma 3.8, we know that a cycle is output before any call to `unlock`. As no call to `unlock(v)` can be made before the call to `cycle(x)`, a cycle is output before the second call to `unlock(v)`.

Let us now consider the delay of the algorithm. In both its exploration and unblocking phases, the algorithm follows the arcs of the graph and transforms the status or the B lists of their endpoints, which overall require constant time. Thus, the delay only depends on the number of arcs which are considered during two successive outputs. An arc (u, v) is considered once by the algorithm in the three following situations: the exploration part of a call to `cycle(u)`; an insertion of u in $B(v)$; a call to `unlock(v)`. As shown before, `unlock(v)` is called only once between two successive outputs. `cycle(u)` cannot be called more than twice. Thus the arc (u, v) is considered at most 5 times between two outputs. This ensures that the delay of the algorithm is $O(m + n)$. \square

3.1.6 Practical speedup

Speeding up preprocessing. In Section 3.1.3, the bubble enumeration problem was reduced to the enumeration of some particular cycles in the transformed graph G'_s for each s . It is worth observing that this does not imply building from scratch G'_s for each s . Indeed, notice that for any two vertices s_1 and s_2 , we can transform G'_{s_1} into G'_{s_2} by: (a) removing from G'_{s_1} the arcs (\bar{s}_1, s_1) , (s_2, \bar{s}_2) , (v, s_2) , and (\bar{s}_2, \bar{v}) for each $v \in N^-(s_2)$ in G ; (b) adding to G'_{s_1} the arcs (s_1, \bar{s}_1) , (\bar{s}_2, s_2) , (v, s_1) , and (\bar{s}_1, \bar{v}) for each $v \in N^-(s_1)$ in G .

Avoiding duplicate bubbles. The one-to-two correspondence between cycles in G'_s and bubbles starting from s in G , claimed by Proposition 3.5, can be reduced to a one-to-one correspondence in the following way. Consider an arbitrary order on the vertices of V , and assign to each vertex of \bar{V} the order of its twin. Let C be a cycle of G'_s that passes through s and contains exactly two pairs of twin vertices. Denote again by t the vertex such that (t, \bar{t}) is the arc through which C swaps from V to \bar{V} . Denote by *swap predecessor* the vertex before t in C and by *swap successor* the vertex after \bar{t} in C .

Proposition 3.11. *There is a one-to-one correspondence between the set of (s, t) -bubbles in G for all $t \in V$, and the set of cycles of G'_s that pass through s , contain exactly two pairs of twin vertices and such that the swap predecessor is greater than the swap successor.*

Proof. The proof follows the one of Proposition 3.5. The only difference is that, if we consider a bubble composed of the paths P_1 and P_2 , one of these two paths, say P_1 , has a next to last vertex greater than the next to last vertex of P_2 . Then the cycle of G'_s made of P_1 and \bar{P}_2 is still considered by the algorithm whereas the cycle made of P_2 and \bar{P}_1 is not. Moreover, the cycles of length four which are of the type $\{s, t, \bar{t}, \bar{s}\}$ are ruled out as \bar{s} is of the same order as s . \square

3.2 Optimal listing of cycles and st -paths in undirected graphs

3.2.1 Introduction

Listing all the simple cycles (hereafter just called cycles) in a graph is a classical problem whose efficient solutions date back to the early 70s. For a graph with n vertices and m edges containing η cycles, the best known solution in the literature is given by Johnson's algorithm (Johnson (1975)) and takes $O((\eta + 1)(m + n))$ time.

Previous work

The classical problem of listing all the cycles of a graph has been extensively studied for its many applications in several fields, ranging from the mechanical analysis of chemical structures Sussenguth (1965) to the design and analysis of reliable communication networks, and the graph isomorphism problem (Welch (1966)). In particular, at the turn of the seventies several algorithms for enumerating all cycles of an undirected graph have been proposed. There is a vast body of work, and the majority of the algorithms listing all the cycles can be divided into the following three classes (see Bezem and Leeuwen (1987) and Mateti and Deo (1976) for excellent surveys).

1. *Search space algorithms.* According to this approach, cycles are looked for in an appropriate search space. In the case of undirected graphs, the *cycle vector space* (Diestel (2005)) turned out to be the most promising choice: from a basis for this space, all vectors are computed and it is tested whether they are a cycle. Since the algorithm introduced in Welch (1966), many algorithms have been proposed: however, the complexity of these algorithms turns out to be exponential in the dimension of the vector space, and thus in n . For planar graphs, an algorithm listing cycles in $O((\eta + 1)n)$ time was presented in Syslo (1981).
2. *Backtrack algorithms.* By this approach, all paths are generated by backtrack and, for each path, it is tested whether it is a cycle. One of the first algorithms is the one proposed in Tiernan (1970), which is however exponential in η . By adding a simple pruning strategy, this algorithm has been successively modified in Tarjan (1973): it lists all the cycles in $O(nm(\eta + 1))$ time. Further improvements were proposed in (Johnson (1975); Szwarcfiter and Lauer (1976); Read and Tarjan (1975)), leading to $O((\eta + 1)(m + n))$ -time algorithms that work for both directed and undirected graphs. Apart from the algorithm in Tiernan (1970), all the algorithms based on this approach are *polynomial-time delay*, that is, the time elapsed between the outputting of two cycles is polynomial in the size of the graph (more precisely, $O(nm)$ in the case of the algorithm of Tarjan (1973) and $O(m)$ in the case of the other three algorithms).
3. *Using the powers of the adjacency matrix.* This approach uses the so-called *variable adjacency matrix*, that is, the formal sum of edges joining two vertices. A non-zero element of the p -th power of this matrix is the sum of all walks of length p : hence, to compute all cycles, we compute the n th power of the variable adjacency matrix. This approach is not very efficient because of the non-simple walks. Algorithms based on this approach (e.g. Ponstein (1966) and Yau (1967)) basically differ only on the way they avoid to consider walks that are neither paths nor cycles.

Almost 40 years after Johnson's algorithm Johnson (1975), the problem of efficiently listing all cycles of a graph is still an active area of research (e.g. Halford and Chugg (2004); Horváth

et al. (2004); Liu and Wang (2006); Sankar and Sarad (2007); Wild (2008); Schott and Staples (2011)). New application areas have emerged in the last decade, such as bioinformatics: for example, two algorithms for this problem have been proposed in Klamt and et al. (2006) and Klamt and von Kamp (2009) while studying biological interaction graphs. Nevertheless, no significant improvement has been obtained from the theory standpoint: in particular, Johnson's algorithm is still the theoretically most efficient. His $O((\eta + 1)(m + n))$ -time solution is surprisingly not optimal for undirected graphs as we show in this chapter.

Results

We present the first optimal solution to list all the cycles in an undirected graph G . Specifically, let $\mathcal{C}(G)$ denote the set of all these cycles ($|\mathcal{C}(G)| = \eta$). For a cycle $c \in \mathcal{C}(G)$, let $|c|$ denote the number of edges in c . Our algorithm requires $O(m + \sum_{c \in \mathcal{C}(G)} |c|)$ time and is asymptotically optimal: indeed, $\Omega(m)$ time is necessarily required to read G as input, and $\Omega(\sum_{c \in \mathcal{C}(G)} |c|)$ time is necessarily required to list the output. Since $|c| \leq n$, the cost of our algorithm never exceeds $O(m + (\eta + 1)n)$ time.

Along the same lines, we also present the first optimal solution to list all the simple paths from s to t (shortly, st -paths) in an undirected graph G . Let $\mathcal{P}_{st}(G)$ denote the set of st -paths in G and, for an st -path $\pi \in \mathcal{P}_{st}(G)$, let $|\pi|$ be the number of edges in π . Our algorithm lists all the st -paths in G optimally in $O(m + \sum_{\pi \in \mathcal{P}_{st}(G)} |\pi|)$ time, observing that $\Omega(\sum_{\pi \in \mathcal{P}_{st}(G)} |\pi|)$ time is necessarily required to list the output.

We prove the following reduction to relate $\mathcal{C}(G)$ and $\mathcal{P}_{st}(G)$ for some suitable choices of vertices s, t : if there exists an optimal algorithm to list the st -paths in G , then there exists an optimal algorithm to list the cycles in G . Hence, we can focus on listing st -paths.

Difficult graphs for Johnson's algorithm

It is worth observing that the analysis of the time complexity of Johnson's algorithm is not pessimistic and cannot match the one of our algorithm for listing cycles. For example, consider the sparse "diamond" graph $D_n = (V, E)$ in Fig. 3.3 with $n = 2k + 3$ vertices in $V = \{a, b, c, v_1, \dots, v_k, u_1, \dots, u_k\}$. There are $m = \Theta(n)$ edges in $E = \{(a, c), (a, v_i), (v_i, b), (b, u_i), (u_i, c), \text{ for } 1 \leq i \leq k\}$, and three kinds of (simple) cycles: (1) $(a, v_i), (v_i, b), (b, u_j), (u_j, c), (c, a)$ for $1 \leq i, j \leq k$; (2) $(a, v_i), (v_i, b), (b, v_j), (v_j, a)$ for $1 \leq i < j \leq k$; (3) $(b, u_i), (u_i, c), (c, u_j), (u_j, b)$ for $1 \leq i < j \leq k$, totalizing $\eta = \Theta(n^2)$ cycles. Our algorithm takes $\Theta(n + k^2) = \Theta(\eta) = \Theta(n^2)$ time to list these cycles. On the other hand, Johnson's algorithm takes $\Theta(n^3)$ time, and the discovery of the $\Theta(n^2)$ cycles in (1) costs $\Theta(k) = \Theta(n)$ time each: the backtracking procedure in Johnson's algorithm starting at a , and passing through v_i, b and u_j for some i, j , arrives at c : at that point, it explores all the vertices u_l ($l \neq i$) even if they do not lead to cycles when coupled with a, v_i, b, u_j , and c .

3.2.2 Preliminaries

Let $G = (V, E)$ be an undirected connected graph with $n = |V|$ vertices and $m = |E|$ edges, without self-loops or parallel edges. For a vertex $u \in V$, we denote by $N(u)$ the neighborhood of u and by $d(u) = |N(u)|$ its degree. $G[V']$ denotes the subgraph induced by $V' \subseteq V$, and $G - u$ is the induced subgraph $G[V \setminus \{u\}]$ for $u \in V$. Likewise for edge $e \in E$, we adopt the notation $G - e = (V, E \setminus \{e\})$. For a vertex $v \in V$, the *postorder* DFS number of v is the relative time in which v was *last* visited in a DFS traversal, i.e. the position of v in the vertex list ordered by the last visiting time of each vertex in the DFS.

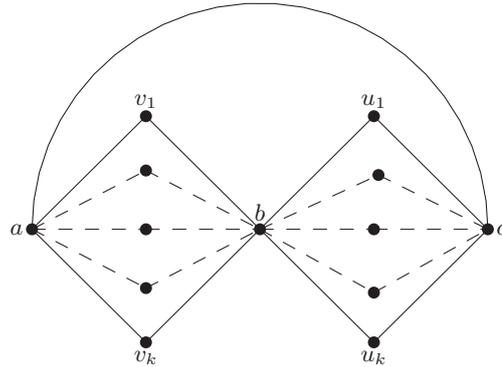


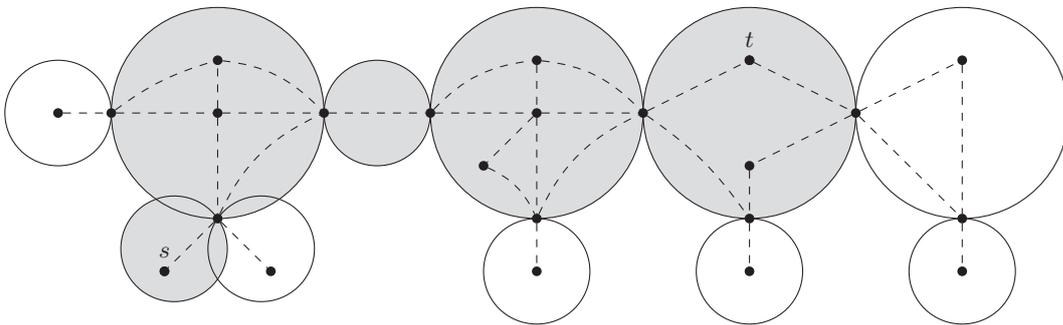
Figure 3.3: Diamond graph.

Paths are simple in G by definition: we refer to a path π by its natural sequence of vertices or edges. A path π from s to t , or st -path, is denoted by $\pi = s \rightsquigarrow t$. Additionally, $\mathcal{P}(G)$ is the set of all paths in G and $\mathcal{P}_{s,t}(G)$ is the set of all st -paths in G . When $s = t$ we have cycles, and $\mathcal{C}(G)$ denotes the set of all cycles in G . We denote the number of edges in a path π by $|\pi|$ and in a cycle c by $|c|$. In this section, we consider the following problems.

Problem 3.12 (Listing st -Paths). *Given an undirected graph $G = (V, E)$ and two distinct vertices $s, t \in V$, output all the paths $\pi \in \mathcal{P}_{s,t}(G)$.*

Problem 3.13 (Listing Cycles). *Given an undirected graph $G = (V, E)$, output all the cycles $c \in \mathcal{C}(G)$.*

Our algorithms assume without loss of generality that the input graph G is connected, hence $m \geq n - 1$, and use the decomposition of G into biconnected components. Recall that an *articulation point* (or cut-vertex) is a vertex $u \in V$ such that the number of connected components in G increases when u is removed. G is *biconnected* if it has no articulation points. Otherwise, G can always be decomposed into a tree of biconnected components, called the *block tree*, where each biconnected component is a maximal biconnected subgraph of G (see Fig. 3.4), and two biconnected components are adjacent if and only if they share an articulation point.

Figure 3.4: Block tree of G with bead string $B_{s,t}$ in gray.

3.2.3 Overview and main ideas

While the basic approach is simple (see the binary partition in point 3), we use a number of non-trivial ideas to obtain our optimal algorithm for an undirected (connected) graph G as

summarized in the steps below.

1. Prove the following reduction. If there exists an optimal algorithm to list the st -paths in G , there exists an optimal algorithm to list the cycles in G . This relates $\mathcal{C}(G)$ and $\mathcal{P}_{st}(G)$ for some choices s, t .
2. Focus on listing the st -paths. Consider the decomposition of the graph into biconnected components (BCCs), thus forming a tree T where two BCCs are adjacent in T iff they share an articulation point. Exploit (and prove) the property that if s and t belong to distinct BCCs, then (i) there is a unique *sequence* $B_{s,t}$ of adjacent BCCs in T through which each st -path must necessarily pass, and (ii) each st -path is the concatenation of paths connecting the articulation points of these BCCs in $B_{s,t}$.
3. Recursively list the st -paths in $B_{s,t}$ using the classical binary partition (i.e. given an edge e in G , list all the cycles containing e , and then all the cycles not containing e): now it suffices to work on the *first* BCC in $B_{s,t}$, and efficiently maintain it when deleting an edge e , as required by the binary partition.
4. Use a notion of *certificate* to avoid recursive calls (in the binary partition) that do not list new st -paths. This certificate is maintained dynamically as a data structure representing the first BCC in $B_{s,t}$, which guarantees that there exists at least one *new* solution in the current $B_{s,t}$.
5. Consider the binary recursion tree corresponding to the binary partition. Divide this tree into *spines*: a spine corresponds to the recursive calls generated by the edges e belonging to the same adjacency list in $B_{s,t}$. The amortized cost for each listed st -path π is $O(|\pi|)$ when there is a guarantee that the amortized cost in each spine S is $O(\mu)$, where μ is a lower bound on the number of st -paths that will be listed from the recursive calls belonging to S . The (unknown) parameter μ , which is different for each spine S , and the corresponding cost $O(\mu)$, will drive the design of the proposed algorithms.

Reduction to st -paths

We now show that listing cycles reduces to listing st -paths while preserving the optimal complexity.

Lemma 3.14. *Given an algorithm that solves Problem 3.12 in optimal $O(m + \sum_{\pi \in \mathcal{P}_{s,t}(G)} |\pi|)$ time, there exists an algorithm that solves Problem 3.13 in optimal $O(m + \sum_{c \in \mathcal{C}(G)} |c|)$ time.*

Proof. Compute the biconnected components of G and keep them in a list L . Each (simple) cycle is contained in one of the biconnected components and therefore we can treat each biconnected component individually as follows. While L is not empty, extract a biconnected component $B = (V_B, E_B)$ from L and repeat the following three steps: (i) compute a DFS traversal of B and take any back edge $b = (s, t)$ in B ; (ii) list all st -paths in $B - b$, i.e. the cycles in B that include edge b ; (iii) remove edge b from B , compute the new biconnected components thus created by removing edge b , and append them to L . When L becomes empty, all the cycles in G have been listed.

Creating L takes $O(m)$ time. For every $B \in L$, steps (i) and (iii) take $O(|E_B|)$ time. Note that step (ii) always outputs distinct cycles in B (i.e. st -paths in $B - b$) in $O(|E_B| + \sum_{\pi \in \mathcal{P}_{s,t}(B-b)} |\pi|)$ time. However, $B - b$ is then decomposed into biconnected components whose edges are traversed again. We can pay for the latter cost: for any edge $e \neq b$ in a

biconnected component B , there is always a cycle in B that contains both b and e (i.e. it is an st -path in $B - b$), hence $\sum_{\pi \in \mathcal{P}_{s,t}(B-b)} |\pi|$ dominates the term $|E_B|$, i.e. $\sum_{\pi \in \mathcal{P}_{s,t}(B-b)} |\pi| = \Omega(|E_B|)$. Therefore steps (i)–(iii) take $O(\sum_{\pi \in \mathcal{P}_{s,t}(B-b)} |\pi|)$ time. When L becomes empty, the whole task has taken $O(m + \sum_{c \in \mathcal{C}(G)} |c|)$ time. \square

Decomposition in biconnected components

We now focus on listing st -paths (Problem 3.12). We use the decomposition of G into a block tree of biconnected components. Given vertices s, t , define its *bead string*, denoted by $B_{s,t}$, as the unique sequence of one or more adjacent biconnected components (the *beads*) in the block tree, such that the first one contains s and the last one contains t (see Fig. 3.4): these biconnected components are connected through articulation points, which must belong to all the paths to be listed.

Lemma 3.15. *All the st -paths in $\mathcal{P}_{s,t}(G)$ are contained in the induced subgraph $G[B_{s,t}]$ for the bead string $B_{s,t}$. Moreover, all the articulation points in $G[B_{s,t}]$ are traversed by each of these paths.*

Proof. Consider an edge $e = (u, v)$ in G such that $u \in B_{s,t}$ and $v \notin B_{s,t}$. Since the biconnected components of a graph form a tree and the bead string $B_{s,t}$ is a path in this tree, there are no paths $v \rightsquigarrow w$ in $G - e$ for any $w \in B_{s,t}$ because the biconnected components in G are maximal and there would be a larger one (a contradiction). Moreover, let B_1, B_2, \dots, B_r be the biconnected components composing $B_{s,t}$, where $s \in B_1$ and $t \in B_r$. If there is only one biconnected component in the path (i.e. $r = 1$), there are no articulation points in $B_{s,t}$. Otherwise, all of the $r - 1$ articulation points in $B_{s,t}$ are traversed by each path $\pi \in \mathcal{P}_{s,t}(G)$: indeed, the articulation point between adjacent biconnected components B_i and B_{i+1} is their only vertex in common and there are no edges linking B_i and B_{i+1} . \square

We thus restrict the problem of listing the paths in $\mathcal{P}_{s,t}(G)$ to the induced subgraph $G[B_{s,t}]$, conceptually isolating it from the rest of G . For the sake of description, we will use interchangeably $B_{s,t}$ and $G[B_{s,t}]$ in the rest of the chapter.

Binary partition scheme

We list the set of st -paths in $B_{s,t}$, denoted by $\mathcal{P}_{s,t}(B_{s,t})$, by applying the binary partition method (where $\mathcal{P}_{s,t}(G) = \mathcal{P}_{s,t}(B_{s,t})$ by Lemma 3.15): we choose an edge $e = (s, v)$ incident to s and then list all the st -paths that include e and then all the st -paths that do not include e . Since we delete some vertices and some edges during the recursive calls, we proceed as follows.

Invariant: At a generic recursive step on vertex u (initially, $u := s$), let $\pi_s = s \rightsquigarrow u$ be the path discovered so far (initially, π_s is empty $\{\}$). Let $B_{u,t}$ be the current bead string (initially, $B_{u,t} := B_{s,t}$). More precisely, $B_{u,t}$ is defined as follows: (i) remove from $B_{s,t}$ all the vertices in π_s but u , and the edges incident to u and discarded so far; (ii) recompute the block tree on the resulting graph; (iii) $B_{u,t}$ is the unique bead string that connects u to t in the recomputed block tree.

Base case: When $u = t$, output the st -path π_s .

Recursive rule: Let $\mathcal{P}(\pi_s, u, B_{u,t})$ denote the set of st -paths to be listed by the current recursive call. Then, it is the union of the following two disjoint sets, for an edge $e = (u, v)$ incident to u :

- *Left branching*: the st -paths in $\mathcal{P}(\pi_s \cdot e, v, B_{v,t})$ that use e , where $B_{v,t}$ is the unique bead string connecting v to t in the block tree resulting from the deletion of vertex u from $B_{u,t}$.
- *Right branching*: the st -paths in $\mathcal{P}(\pi_s, u, B'_{u,t})$ that do *not* use e , where $B'_{u,t}$ is the unique bead string connecting u to t in the block tree resulting from the deletion of edge e from $B_{u,t}$.

Hence, $\mathcal{P}_{s,t}(B_{s,t})$ (and so $\mathcal{P}_{s,t}(G)$) can be computed by invoking $\mathcal{P}(\{\}, s, B_{s,t})$. The correctness and completeness of the above approach is discussed in Section 3.2.3.

At this point, it should be clear why we introduce the notion of bead strings in the binary partition. The existence of the partial path π_s and the bead string $B_{u,t}$ guarantees that there surely exists at least one st -path. But there are two sides of the coin when using $B_{u,t}$.

1. One advantage is that we can avoid useless recursive calls: If vertex u has only one incident edge e , we just perform the left branching; otherwise, we can safely perform both the left and right branching since the *first* bead in $B_{u,t}$ is always a biconnected component by definition (thus there exists both an st -path that traverses e and one that does not).
2. The other side of the coin is that we have to maintain the bead string $B_{u,t}$ as $B_{v,t}$ in the left branching and as $B'_{u,t}$ in the right branching by Lemma 3.15. Note that these bead strings are surely non-empty since $B_{u,t}$ is non-empty by induction (we only perform either left or left/right branching when there are solutions by item 1).

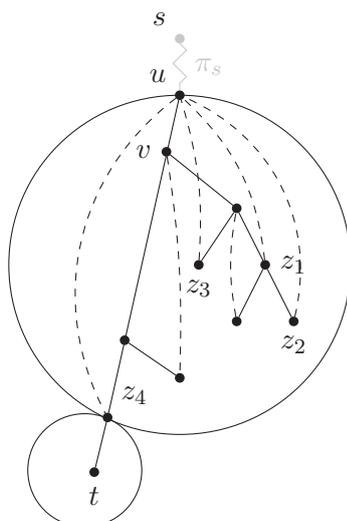
To efficiently address point 2, we need to introduce the notion of certificate as described next.

Introducing the certificate

Given the bead string $B_{u,t}$, we call the *head* of $B_{u,t}$, denoted by H_u , the first biconnected component in $B_{u,t}$, where $u \in H_u$. Consider a DFS tree of $B_{u,t}$ rooted at u that changes along with $B_{u,t}$, and classify the edges in $B_{u,t}$ as tree edges or back edges (there are no cross edges since the graph is undirected).

To maintain $B_{u,t}$ (and so H_u) during the recursive calls, we introduce a *certificate* C (see Fig. 3.5): It is a suitable data structure that uses the above classification of the edges in $B_{u,t}$, and supports the following operations, required by the binary partition scheme.

- **choose**(C, u): returns an edge $e = (u, v)$ with $v \in H_u$ such that $\pi_s \cdot (u, v) \cdot u \rightsquigarrow t$ is an st -path such that $u \rightsquigarrow t$ is inside $B_{u,t}$. Note that e always exists since H_u is biconnected. Also, the chosen v is the last one in DFS postorder among the neighbors of u : in this way, the (only) tree edge e is returned when there are no back edges leaving from u . (As it will be clear in Sections 3.2.4 and 3.2.5, this order facilitates the analysis and the implementation of the certificate.)
- **left_update**(C, e): for the given $e = (u, v)$, it obtains $B_{v,t}$ from $B_{u,t}$ as discussed in Section 3.2.3. This implies updating also H_u, C , and the block tree, since the recursion continues on v . It returns bookkeeping information I for what is updated, so that it is possible to revert to $B_{u,t}, H_u, C$, and the block tree, to their status before this operation.

Figure 3.5: Example certificate of $B_{u,t}$

- **right_update**(C, e): for the given $e = (u, v)$, it obtains $B'_{u,t}$ from $B_{u,t}$ as discussed in Section 3.2.3, which implies updating also H_u , C , and the block tree. It returns bookkeeping information I as in the case of **left_update**(C, e).
- **restore**(C, I): reverts the bead string to $B_{u,t}$, the head H_u , the certificate C , and the block tree, to their status before operation $I := \text{left_update}(C, e)$ or $I := \text{right_update}(C, e)$ was issued (in the same recursive call).

Note that a notion of certificate in listing problems has been introduced in [Ferreira et al. \(2011\)](#), but it cannot be directly applied to our case due to the different nature of the problems and our use of more complex structures such as biconnected components.

Using our certificate and its operations, we can now formalize the binary partition and its recursive calls $\mathcal{P}(\pi_s, u, B_{u,t})$ described in Section 3.2.3 as Algorithm 3.4, where $B_{u,t}$ is replaced by its certificate C .

Algorithm 3.4: $\text{list_paths}_{s,t}(\pi_s, u, C)$

```

1 if  $u = t$  then
2   |   output( $\pi_s$ )
3   |   return
4  $e = (u, v) := \text{choose}(C, u)$ 
5 if  $e$  is back edge then
6   |    $I := \text{right\_update}(C, e)$ 
7   |    $\text{list\_paths}_{s,t}(\pi_s, u, C)$ 
8   |    $\text{restore}(C, I)$ 
9  $I := \text{left\_update}(C, e)$ 
10  $\text{list\_paths}_{s,t}(\pi_s \cdot (u, v), v, C)$ 
11  $\text{restore}(C, I)$ 

```

The base case ($u = t$) corresponds to lines 1–4 of Algorithm 3.4. During recursion, the left branching corresponds to lines 5 and 11–13, while the right branching to lines 5–10. Note

that we perform only the left branching when there is only one incident edge in u , which is a tree edge by definition of **choose**. Also, lines 9 and 13 are needed to restore the parameters to their values when returning from the recursive calls.

Lemma 3.16. *Given a correct implementation of the certificate C and its supported operations, Algorithm 3.4 correctly lists all the st -paths in $\mathcal{P}_{s,t}(G)$.*

Proof. For a given vertex u the function **choose**(C, u) returns an edge e incident to u . We maintain the invariant that π_s is a path $s \rightsquigarrow u$, since at the point of the recursive call in line 10: (i) is connected as we append edge (u, v) to π_s and; (ii) it is simple as vertex u is removed from the graph G in the call to **left_update**(C, e) in line 9. In the case of recursive call in line 7 the invariant is trivially maintained as π_s does not change. The algorithm only outputs st -paths since π_s is a $s \rightsquigarrow u$ path and $u = t$ when the algorithm outputs, in line 2.

The paths with prefix π_s that do not use e are listed by the recursive call in line 7. This is done by removing e from the graph (line 6) and thus no path can include e . Paths that use e are listed in line 10 since in the recursive call e is added to π_s . Given that the tree edge incident to u is the last one to be returned by **choose**(C, u), there is no path that does not use this edge, therefore it is not necessary to call line 7 for this edge. \square

A natural question is what is the time complexity: we must account for the cost of maintaining C and for the cost of the recursive calls of Algorithm 3.4. Since we cannot always maintain the certificate in $O(1)$ time, the ideal situation for attaining an optimal cost is taking $O(\mu)$ time if at least μ st -paths are listed in the current call (and its nested calls). Unfortunately, we cannot estimate μ efficiently and cannot design Algorithm 3.4 so that it takes $O(\mu)$ adaptively. We circumvent this by using a different cost scheme in Section 11 that is based on the recursion tree induced by Algorithm 3.4. Section 3.2.5 is devoted to the efficient implementation of the above certificate operations according to the cost scheme that we discuss next.

Recursion tree and cost amortization

We now show how to distribute the costs among the several recursive calls of Algorithm 3.4 so that optimality is achieved. Consider a generic execution on the bead string $B_{u,t}$. We trace this execution by using a binary recursion tree R . The nodes of R are labeled by the arguments of Algorithm 3.4: specifically, we denote a node in R by the triple $x = \langle \pi_s, u, C \rangle$ iff it represents the call with arguments π_s , u , and C .³ The left branching is represented by the left child, and the right branching (if any) by the right child of the current node.

Lemma 3.17. *The binary recursion tree R for $B_{u,t}$ has the following properties:*

1. *There is a one-to-one correspondence between the paths in $\mathcal{P}_{s,t}(B_{u,t})$ and the leaves in the recursion tree rooted at node $\langle \pi_s, u, C \rangle$.*
2. *Consider any leaf and its corresponding st -path π : there are $|\pi|$ left branches in the corresponding root-to-leaf trace.*
3. *Consider the instruction $e := \mathbf{choose}(C, u)$ in Algorithm 3.4: unary (i.e. single-child) nodes correspond to left branches (e is a tree edge) while binary nodes correspond to left and right branches (e is a back edge).*
4. *The number of binary nodes is $|\mathcal{P}_{s,t}(B_{u,t})| - 1$.*

³For clarity, we use “nodes” when referring to R and “vertices” when referring to $B_{u,t}$.

Proof. We proceed in order as follows.

1. We only output a solution in a leaf and we only do recursive calls that lead us to a solution. Moreover every node partitions the set of solutions in the ones that use an edge and the ones that do not use it. This guarantees that the leaves in the left subtree of the node corresponding to the recursive call and the leaves in the right subtree do not intersect. This implies that different leaves correspond to different paths from s to t , and that for each path there is a corresponding leaf.
2. Each left branch corresponds to the inclusion of an edge in the path π .
3. Since we are in a biconnected component, there is always a left branch. There can be no unary node as a right branch: indeed for any edge of $B_{u,t}$ there exists always a path from s to t passing through that edge. Since the tree edge is always the last one to be chosen, unary nodes cannot correspond to back edges and binary nodes are always back edges.
4. It follows from point 1 and from the fact that the recursion tree is a binary tree. (In any binary tree, the number of binary nodes is equal to the number of leaves minus 1.)

□

We define a *spine* of R to be a subset of R 's nodes linked as follows: the first node is a node x that is either the left child of its parent or the root of R , and the other nodes are those reachable from x by right branching in R . Let $x = \langle \pi_s, u, C \rangle$ be the first node in a spine S . The nodes in S correspond to the edges that are incident to vertex u in $B_{u,t}$: hence their number equals the degree $d(u)$ of u in $B_{u,t}$, and the deepest (last) node in S is always a tree edge in $B_{u,t}$ while the others are back edges. Fig. 3.6 shows the spine corresponding to $B_{u,t}$ in Fig. 3.5. Summing up, R can be seen as composed by spines, unary nodes, and leaves where each spine has a unary node as deepest node. This gives a global picture of R that we now exploit for the analysis.

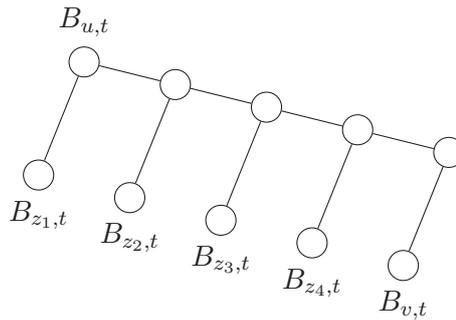


Figure 3.6: Spine of the recursion tree

We define the *compact head*, denoted by $H_X = (V_X, E_X)$, as the (multi)graph obtained by compacting the maximal chains of degree-2 vertices, except u, t , and the vertices that are the leaves of its DFS tree rooted at u .

The rationale behind the above definition is that the costs defined in terms of H_X amortize well, as the size of H_X and the number of st -paths in the subtree of R rooted at node $x = \langle \pi_s, u, C \rangle$ are intimately related (see Lemma 3.21 in Section 3.2.4) while this is not necessarily true for H_u .

Recall that each leaf corresponds to a path π and each spine corresponds to a compact head $H_X = (V_X, E_X)$. We now define the following abstract cost for spines, unary nodes, and leaves of R , for a sufficiently large constant $c_0 > 0$, that Algorithm 3.4 must fulfill:

$$T(r) = \begin{cases} c_0 & \text{if } r \text{ is unary} \\ c_0|\pi| & \text{if } r \text{ is a leaf} \\ c_0(|V_X| + |E_X|) & \text{if } r \text{ is a spine} \end{cases} \quad (3.1)$$

Lemma 3.18. *The sum of the costs in the nodes of the recursion tree $\sum_{r \in R} T(r) = O(\sum_{\pi \in \mathcal{P}_{s,t}(B_{u,t})} |\pi|)$.*

Section 3.2.4 contains the proof of Lemma 3.18 and related properties. Setting $u := s$, we obtain that the cost in Lemma 3.18 is optimal, by Lemma 3.15.

Theorem 3.19. *Algorithm 3.4 solves problem Problem 3.12 in optimal $O(m + \sum_{\pi \in \mathcal{P}_{s,t}(G)} |\pi|)$ time.*

By Lemma 3.14, we obtain an optimal result for listing cycles.

Theorem 3.20. *Problem 3.13 can be optimally solved in $O(m + \sum_{c \in \mathcal{C}(G)} |c|)$ time.*

3.2.4 Amortization strategy

We devote this section to prove Lemma 3.18. Let us split the sum in Eq. (3.1) in three parts, and bound each part individually, as

$$\sum_{r \in R} T(r) \leq \sum_{r: \text{unary}} T(r) + \sum_{r: \text{leaf}} T(r) + \sum_{r: \text{spine}} T(r). \quad (3.2)$$

We have that $\sum_{r: \text{unary}} T(r) = O(\sum_{\pi \in \mathcal{P}_{s,t}(G)} |\pi|)$, since there are $|\mathcal{P}_{s,t}(G)|$ leaves, and the root-to-leaf trace leading to the leaf for π contains at most $|\pi|$ unary nodes by Lemma 3.17, where each unary node has cost $O(1)$ by Eq. (3.1).

Also, $\sum_{r: \text{leaf}} T(r) = O(\sum_{\pi \in \mathcal{P}_{s,t}(G)} |\pi|)$, since the leaf r for π has cost $O(|\pi|)$ by Eq. (3.1).

It remains to bound $\sum_{r: \text{spine}} T(r)$. By Eq. (3.1), we can rewrite this cost as $\sum_{H_X} c_0(|V_X| + |E_X|)$, where the sum ranges over the compacted heads H_X associated with the spines r . We use the following lemma to provide a lower bound on the number of st -paths descending from r .

Lemma 3.21. *Given a spine r , and its bead string $B_{u,t}$ with head H_u , there are at least $|E_X| - |V_X| + 1$ st -paths in G that have prefix $\pi_s = s \rightsquigarrow u$ and suffix $u \rightsquigarrow t$ internal to $B_{u,t}$, where the compacted head is $H_X = (V_X, E_X)$.*

Proof. H_X is biconnected. In any biconnected graph $B = (V_B, E_B)$ there are at least $|E_B| - |V_B| + 1$ xy -paths for any $x, y \in V_B$. Find an ear decomposition (see Chapter 1, Lemma 1.1) of B and consider the process of forming B by adding ears one at the time, starting from a single cycle including x and y . Initially $|V_B| = |E_B|$ and there are 2 xy -paths. Each new ear forms a path connecting two vertices that are part of a xy -path, increasing the number of paths by at least 1. If the ear has k edges, its addition increases V by $k - 1$, E by k , and the number of xy -paths by at least 1. The result follows by induction. \square

The implication of Lemma 3.21 is that there are at least $|E_X| - |V_X| + 1$ leaves descending from the given spine r . Hence, we can charge to each of them a cost of $\frac{c_0(|V_X| + |E_X|)}{|E_X| - |V_X| + 1}$. Lemma 3.22 allows us to prove that the latter cost is $O(1)$ when H_u is different from a single edge or a cycle. (If H_u is a single edge or a cycle, H_X is a single or double edge, and the cost is trivially a constant.)

Lemma 3.22. *For a compacted head $H_X = (V_X, E_X)$, its density is $\frac{|E_X|}{|V_X|} \geq \frac{11}{10}$.*

Proof. Consider the following partition $V_X = \{r\} \cup V_2 \cup V_3$ where: r is the root; V_2 is the set of vertices with degree 2 and; V_3 , the vertices with degree ≥ 3 . Since H_X is compacted DFS tree of a biconnected graph, we have that V_2 is a *subset* of the leaves and V_3 contains the set of internal vertices (except r). There are no vertices with degree 1 and $d(r) \geq 2$. Let $x = \sum_{v \in V_3} d(v)$ and $y = \sum_{v \in V_2} d(v)$. We can write the density as a function of x and y , namely,

$$\frac{|E_X|}{|V_X|} = \frac{x + y + d(r)}{2(|V_3| + |V_2| + 1)}$$

Note that $|V_3| \leq \frac{x}{3}$ as the vertices in V_3 have at least degree 3, $|V_2| = \frac{y}{2}$ as vertices in V_2 have degree exactly 2. Since $d(r) \geq 2$, we derive the following bound

$$\frac{|E_X|}{|V_X|} \geq \frac{x + y + 2}{\frac{2}{3}x + y + 2}$$

Consider any graph with $|V_X| > 3$ and its DFS tree rooted at r . Note that: (i) there are no tree edges between any two leaves, (ii) every vertex in V_2 is a leaf and (iii) no leaf is a child of r . Therefore, every tree edge incident in a vertex of V_2 is also incident in a vertex of V_3 . Since exactly half the incident edges to V_2 are tree edges (the other half are back edges) we get that $y \leq 2x$.

With $|V_X| \geq 3$ there exists at least one internal vertex in the DFS tree and therefore $x \geq 3$.

$$\begin{aligned} & \text{minimize} && \frac{x + y + 2}{\frac{2}{3}x + y + 2} \\ & \text{subject to} && 0 \leq y \leq 2x, \\ & && x \geq 3. \end{aligned}$$

Since for any x the function is minimized by the maximum y s.t. $y \leq 2x$ and for any y by the minimum x , we get

$$\frac{|E_X|}{|V_X|} \geq \frac{9x + 6}{8x + 6} \geq \frac{11}{10}.$$

□

Specifically, let $\alpha = \frac{11}{10}$ and write $\alpha = 1 + 2/\beta$ for a constant β : we have that $|E_X| + |V_X| = (|E_X| - |V_X|) + 2|V_X| \leq (|E_X| - |V_X|) + \beta(|E_X| - |V_X|) = \frac{\alpha+1}{\alpha-1}(|E_X| - |V_X|)$. Thus, we can charge each leaf with a cost of $\frac{c_0(|V_X| + |E_X|)}{|E_X| - |V_X| + 1} \leq c_0 \frac{\alpha+1}{\alpha-1} = O(1)$. This motivates the definition of H_X , since Lemma 3.22 does not necessarily hold for the head H_u (due to the unary nodes in its DFS tree).

One last step to bound $\sum_{H_X} c_0(|V_X| + |E_X|)$: as noted before, a root-to-leaf trace for the string storing π has $|\pi|$ left branches by Lemma 3.17, and as many spines, each spine charging $c_0 \frac{\alpha+1}{\alpha-1} = O(1)$ to the leaf at hand. This means that each of the $|\mathcal{P}_{s,t}(G)|$ leaves is charged for a cost of $O(|\pi|)$, thus bounding the sum as $\sum_{r \text{ spine}} T(r) = \sum_{H_X} c_0(|V_X| + |E_X|) = O(\sum_{\pi \in \mathcal{P}_{s,t}(G)} |\pi|)$. This completes the proof of Lemma 3.18. As a corollary, we obtain the following result.

Lemma 3.23. *The recursion tree R with cost as in Eq. (3.1) induces an $O(|\pi|)$ amortized cost for each st -path π .*

3.2.5 Certificate implementation and maintenance

The certificate C associated with a node $\langle \pi_s, u, C \rangle$ in the recursion tree is a compacted and augmented DFS tree of bead string $B_{u,t}$, rooted at vertex u . The DFS tree changes over time along with $B_{u,t}$, and is maintained in such a way that t is in the leftmost path of the tree. We compact the DFS tree by contracting the vertices that have degree 2, except u , t , and the leaves (the latter surely have incident back edges). Maintaining this compacted representation is not a difficult data-structure problem. From now on we can assume w.l.o.g. that C is an augmented DFS tree rooted at u where internal nodes of the DFS tree have degree ≥ 3 , and each vertex v has associated the following information.

1. A doubly-linked list $lb(v)$ of back edges linking v to its descendants w sorted by postorder DFS numbering.
2. A doubly-linked list $ab(v)$ of back edges linking v to its ancestors w sorted by preorder DFS numbering.
3. An integer $\gamma(v)$, such that if v is an ancestor of w then $\gamma(v) < \gamma(w)$.
4. The smallest $\gamma(w)$ over all w , such that (h, w) is a back edge and h is in the subtree of v , denoted by $lowpoint(v)$.

Given three vertices $v, w, x \in C$ such that v is the parent of w and x is not in the subtree⁴ of w , we can efficiently test if v is an articulation point, i.e. $lowpoint(w) \leq \gamma(v)$. (Note that we adopt a variant of $lowpoint$ using $\gamma(v)$ in place of the preorder numbering Tarjan (1972): it has the same effect whereas using $\gamma(v)$ is preferable since it is easier to dynamically maintain.)

Lemma 3.24. *The certificate associated with the root of the recursion can be computed in $O(m)$ time.*

Proof. In order to set t to be in the leftmost path, we perform a DFS traversal of graph G starting from s and stop when we reach vertex t . We then compute the DFS tree, traversing the path $s \rightsquigarrow t$ first. When visiting vertex v , we set $\gamma(v)$ to depth of v in the DFS. Before going up on the traversal, we compute the lowpoints using the lowpoints of the children. Let z be the parent of v . If $lowpoint(v) \leq \gamma(z)$ and v is not in the leftmost path in the DFS, we cut the subtree of v as it does not belong to $B_{s,t}$. When first exploring the neighborhood of v , if w was already visited, i.e. $e = (u, w)$ is a back edge, and w is a descendant of v ; we add e to $ab(w)$. This maintains the DFS preordering in the ancestor back edge list. Now, after the first scan of $N(v)$ is over and all the recursive calls returned (all the children were explored), we re-scan the neighborhood of v . If $e = (v, w)$ is a back edge and w is an ancestor of v , we add e to $lb(w)$. This maintains the DFS postorder in the descendant back edge list. This procedure takes at most two DFS traversals in $O(m)$ time. This DFS tree can be compacted in the same time bound. \square

Lemma 3.25. *Operation $choose(C, u)$ can be implemented in $O(1)$ time.*

Proof. If the list $lb(v)$ is empty, return the tree edge $e = (u, v)$ linking u to its only child v (there are no other children). Else, return the last edge in $lb(v)$. \square

⁴The second condition is always satisfied when w is not in the leftmost path, since t is not in the subtree of w .

We analyze the cost of updating and restoring the certificate C . We can reuse parts of C , namely, those corresponding to the vertices that are not in the compacted head $H_X = (V_X, E_X)$ as defined in Section 11. We prove that, given a unary node u and its tree edge $e = (u, v)$, the subtree of v in C can be easily made a certificate for the left branch of the recursion.

Lemma 3.26. *On a unary node, `left_update`(C, e) takes $O(1)$ time.*

Proof. Take edge $e = (u, v)$. Remove edge e and set v as the root of the certificate. Since e is the only edge incident in v , the subtree v is still a DFS tree. Cut the list of children of v keeping only the first child. (The other children are no longer in the bead string and become part of I .) There is no need to update $\gamma(v)$. \square

We now devote the rest of this section to show how to efficiently maintain C on a spine. Consider removing a back edge e from u : the compacted head $H_X = (V_X, E_X)$ of the bead string can be divided into smaller biconnected components. Many of those can be excluded from the certificate (i.e. they are no longer in the new bead string, and so they are bookkept in I) and additionally we have to update the lowpoints that change. We prove that this operation can be performed in $O(|V_X|)$ total time on a spine of the recursion tree.

Lemma 3.27. *The total cost of all the operations `right_update`(C, e) in a spine is $O(|V_X|)$ time.*

Proof. In the right branches along a spine, we remove all back edges in $lb(u)$. This is done by starting from the last edge in $lb(u)$, i.e. proceeding in reverse DFS postorder. For back edge $b_i = (z_i, u)$, we traverse the vertices in the path from z_i towards the root u , as these are the only lowpoints that can change. While moving upwards on the tree, on each vertex w , we update $lowpoint(w)$. This is done by taking the endpoint y of the first edge in $ab(w)$ (the back edge that goes the topmost in the tree) and choosing the minimum between $\gamma(y)$ and the lowpoint of each child⁵ of w . We stop when the updated $lowpoint(w) = \gamma(u)$ since it implies that the lowpoint of the vertex can not be further reduced. Note that we stop before u , except when removing the last back edge in $lb(u)$.

To prune the branches of the DFS tree that are no longer in $B_{u,t}$, consider again each vertex w in the path from z_i towards the root u and its parent y . We check if the updated $lowpoint(w) \leq \gamma(y)$ and w is not in the leftmost path of the DFS. If both conditions are satisfied, we have that $w \notin B_{u,t}$, and therefore we cut the subtree of w and keep it in I to restore later. We use the same halting criterion as in the previous paragraph.

The cost of removing all back edges in the spine is $O(|V_X|)$: there are $O(|V_X|)$ tree edges and, in the paths from z_i to u , we do not traverse the same tree edge twice since the process described stops at the first common ancestor of endpoints of back edges b_i . Additionally, we take $O(1)$ time to cut a subtree of an articulation point in the DFS tree. \square

To compute `left_update`(C, e) in the binary nodes of a spine, we use the fact that in every left branching from that spine, the graph is the same (in a spine we only remove edges incident to u and on a left branch from the spine we remove the vertex u) and therefore its block tree is also the same. However, the certificates on these nodes are not the same, as they

⁵If $lowpoint(w)$ does not change we cannot pay to explore its children. For each vertex we dynamically maintain a list $l(w)$ of its children that have lowpoint equal to $\gamma(u)$. Then, we can test in constant time if $l(w) \neq \emptyset$ and y is not the root u . If both conditions are true $lowpoint(w)$ changes, otherwise it remains equal to $\gamma(u)$ and we stop.

are rooted at different vertices. Using the reverse DFS postorder of the edges, we are able to traverse each edge in H_X only a constant number of times in the spine.

Lemma 3.28. *The total cost of all operations $\text{left_update}(C, e)$ in a spine is amortized $O(|E_X|)$.*

Proof. Let t' be the last vertex in the path $u \rightsquigarrow t$ s.t. $t' \in V_X$. Since t' is an articulation point, the subtree of the DFS tree rooted in t' is maintained in the case of removal of vertex u . Therefore the only modifications of the DFS tree occur in the compacted head H_X of $B_{u,t}$. Let us compute the certificate C_i : this is the certificate of the left branch of the i th node of the spine where we augment the path with the back edge $b_i = (z_i, u)$ of $lb(u)$ in the order defined by $\text{choose}(C, u)$.

For the case of C_1 , we remove u and rebuild the certificate starting from z_1 (the last edge in $lb(u)$) using the algorithm from Lemma 3.24 restricted to H_X and using t' as target and $\gamma(t')$ as a baseline to γ (instead of the depth). This takes $O(|E_X|)$ time.

For the general case of C_i with $i > 1$ we also rebuild (part) of the certificate starting from z_i using the procedure from Lemma 3.24 but we use information gathered in C_{i-1} to avoid exploring useless branches of the DFS tree. The key point is that, when we reach the first bead in common to both $B_{z_i,t}$ and $B_{z_{i-1},t}$, we only explore edges internal to this bead. If an edge e leaving the bead leads to t , we can reuse a subtree of C_{i-1} . If e does not lead to t , then it has already been explored (and cut) in C_{i-1} and there is no need to explore it again since it will be discarded. Given the order we take b_i , each bead is not added more than once, and the total cost over the spine is $O(|E_X|)$.

Nevertheless, the internal edges E'_X of the first bead in common between $B_{z_i,t}$ and $B_{z_{i-1},t}$ can be explored several times during this procedure.⁶ We can charge the cost $O(|E'_X|)$ of exploring those edges to another node in the recursion tree, since this common bead is the head of at least one certificate in the recursion subtree of the left child of the i th node of the spine. Specifically, we charge the first node in the *leftmost* path of the i th node of the spine that has exactly the edges E'_X as head of its bead string: (i) if $|E'_X| \leq 1$ it corresponds to a unary node or a leaf in the recursion tree and therefore we can charge it with $O(1)$ cost; (ii) otherwise it corresponds to a first node of a spine and therefore we can also charge it with $O(|E'_X|)$. We use this charging scheme when $i \neq 1$ and the cost is always charged in the leftmost recursion path of i th node of the spine. Consequently, we never charge a node in the recursion tree more than once. \square

Lemma 3.29. *On each node of the recursion tree, $\text{restore}(C, I)$ takes time proportional to the size of the modifications kept in I .*

Proof. We use standard data structures (i.e. linked lists) for the representation of certificate C . Persistent versions of these data structures exist that maintain a stack of modifications applied to them and that can restore its contents to their previous states. Given the modifications in I , these data structures take $O(|I|)$ time to restore the previous version of C .

Let us consider the case of performing $\text{left_update}(C, e)$. We cut at most $O(|V_X|)$ edges from C . Note that, although we conceptually remove whole branches of the DFS tree, we only remove edges that attach those branches to the DFS tree. The other vertices and edges are left in the certificate but, as they no longer remain attached to $B_{u,t}$, they will never be

⁶Consider the case where z_i, \dots, z_j are all in the same bead after the removal of u . The bead strings are the same, but the roots z_i, \dots, z_j are different, so we have to compute the corresponding DFS of the first component $|j - i|$ times.

reached or explored. In the case of $\text{right_update}(C, e)$, we have a similar situation, with at most $O(|E_X|)$ edges being modified along the spine of the recursion tree. \square

From Lemmas 3.25 and 3.27–3.29, it follows that on a spine of the recursion tree we have the costs: $\text{choose}(u)$ on each node which is bounded by $O(|V_X|)$ time as there are at most $|V_X|$ back edges in u ; $\text{right_update}(C, e)$, $\text{restore}(C, I)$ take $O(|V_X|)$ time; $\text{left_update}(C, e)$ and $\text{restore}(C, I)$ are charged $O(|V_X| + |E_X|)$ time. We thus have the following result, completing the proof of Theorem 3.19.

Lemma 3.30. *Algorithm 3.4 can be implemented with a cost fulfilling Eq. (3.1), thus it takes total $O(m + \sum_{r \in R} T(r)) = O(m + \sum_{\pi \in \mathcal{P}_{s,t}(B_{u,t})} |\pi|)$ time.*

3.2.6 Extended analysis of operations

In this section, we present all details and illustrate with figures the operations $\text{right_update}(C, e)$ and $\text{left_update}(C, e)$ that are performed along a spine of the recursion tree. In order to better detail the procedures in Lemma 3.27 and Lemma 3.28, we divide them in smaller parts. We use bead string $B_{u,t}$ from Fig. 3.5 and the respective spine from Fig. 3.6 as the base for the examples. This spine contains four binary nodes corresponding to the back edges in $lb(u)$ and an unary node corresponding to the tree edge (u, v) . Note that edges are taken in order of the endpoints z_1, z_2, z_3, z_4, v as defined in operation $\text{choose}(C, u)$.

By Lemma 3.15, the impact of operations $\text{right_update}(C, e)$ and $\text{left_update}(C, e)$ in the certificate is restricted to the biconnected component of u . Thus we mainly focus on maintaining the compacted head $H_X = (V_X, E_X)$ of the bead string $B_{u,t}$.

Operation $\text{right_update}(C, e)$

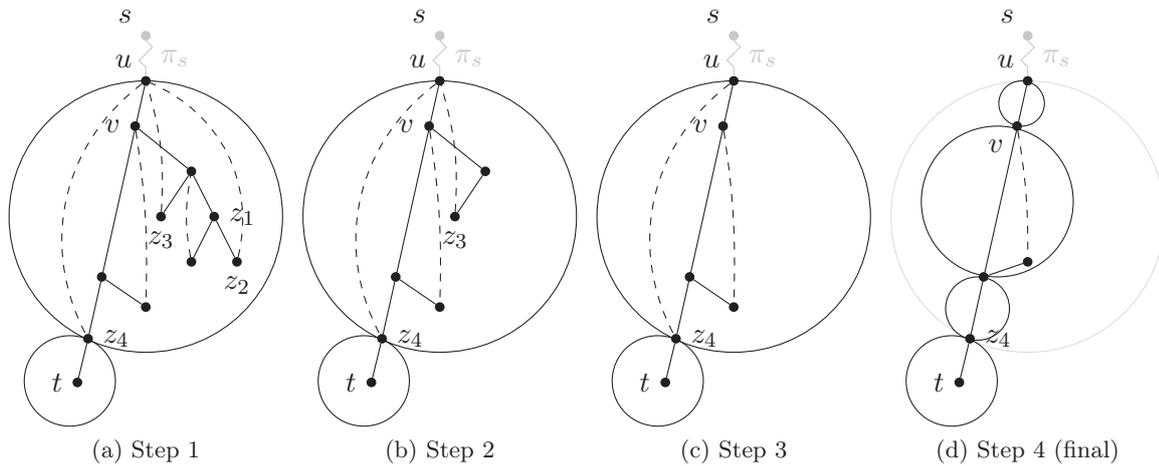


Figure 3.7: Example application of $\text{right_update}(C, e)$ on a spine of the recursion tree

Lemma 3.31. (Lemma 3.27 restated) *In a spine of the recursion tree, operations $\text{right_update}(C, e)$ can be implemented in $O(|V_X|)$ total time.*

In the right branches along a spine, we remove all back edges in $lb(u)$. This is done by starting from the last edge in $lb(u)$, i.e. proceeding in reverse DFS postorder. In the

example from Fig. 3.5, we remove the back edges $(z_1, u) \dots (z_4, u)$. To update the certificate corresponding to $B_{u,t}$, we have to (i) update the lowpoints in each vertex of H_X ; (ii) prune vertices that cease to be in $B_{u,t}$ after removing a back edge. For a vertex w in the tree, there is no need to update $\gamma(w)$.

Consider the update of lowpoints in the DFS tree. For a back edge $b_i = (z_i, u)$, we traverse the vertices in the path from z_i towards the root u . By definition of lowpoint, these are the only lowpoints that can change. Suppose that we remove back edge (z_4, u) in the example from Fig. 3.5, only the lowpoints of the vertices in the path from z_4 towards the root u change. Furthermore, consider a vertex w in the tree that is an ancestor of at least two endpoints z_i, z_j of back edges b_i, b_j . The lowpoint of w does not change when we remove b_i . These observations lead us to the following lemma.

Lemma 3.32. *In a spine of the recursion tree, the update of lowpoints in the certificate by operation `right_update(C, e)` can be done in $O(|V_X|)$ total time.*

Proof. Take each back edge $b_i = (z_i, u)$ in the order defined by `choose(C, u)`. Remove b_i from $lb(u)$ and $ab(z_i)$. Starting from z_i , consider each vertex w in the path from z_i towards the root u . On vertex w , we update $lowpoint(w)$ using the standard procedure: take the endpoint y of the first edge in $ab(w)$ (the back edge that goes the nearest to the root of the tree) and choosing the minimum between $\gamma(y)$ and the lowpoint of each child of w . When the updated $lowpoint(w) = \gamma(u)$, we stop examining the path from z_i to u since it implies that the lowpoint of the vertex can not be further reduced (i.e. w is both an ancestor to both z_i and z_{i+1}).

If $lowpoint(w)$ does not change we cannot pay to explore its children. In order to get around this, for each vertex we dynamically maintain, throughout the spine, a list $l(w)$ of its children that have lowpoint equal to $\gamma(u)$. Then, we can test in constant time if $l(w) \neq \emptyset$ and y (the endpoint of the first edge in $ab(w)$) is not the root u . If both conditions are satisfied $lowpoint(w)$ changes, otherwise it remains equal to $\gamma(u)$ and we stop. The total time to create the lists is $O(|V_X|)$ and the time to update is bounded by the number of tree edges traversed, shown to be $O(|V_X|)$ in the next paragraph.

The cost of updating the lowpoints when removing all back edges b_i is $O(|V_X|)$: there are $O(|V_X|)$ tree edges and we do not traverse the same tree edge twice since the process described stops at the first common ancestor of endpoints of back edges b_i and b_{i+1} . By contradiction: if a tree edge (x, y) would be traversed twice when removing back edges b_i and b_{i+1} , it would imply that both x and y are ancestors of z_i and z_{i+1} (as edge (x, y) is both in the path z_i to u and the path z_{i+1} to u) but we stop at the first ancestor of z_i and z_{i+1} . \square

Let us now consider the removal of vertices that are no longer in $B_{u,t}$ as consequence of operation `right_update(C, e)` in a spine of the recursion tree. By removing a back edge $b_i = (z_i, u)$, it is possible that a vertex w previously in H_X is no longer in the bead string $B_{u,t}$ (e.g. w is no longer biconnected to u and thus there is no simple path $u \rightsquigarrow w \rightsquigarrow t$).

Lemma 3.33. *In a spine of the recursion tree, the branches of the DFS that are no longer in $B_{u,t}$ due to operation `right_update(C, e)` can be removed from the certificate in $O(|V_X|)$ total time.*

Proof. To prune the branches of the DFS tree that are no longer in H_X , consider again each vertex w in the path from z_i towards the root u and the vertex y , parent of w . It is easy to check if y is an articulation point by verifying if the updated $lowpoint(w) \leq \gamma(y)$ and there exists x not in the subtree of w . If w is not in the leftmost path, then t is not in the subtree

of w . If that is the case, we have that $w \notin B_{u,t}$, and therefore we cut the subtree of w and bookkeep it in I to restore later. Like in the update the lowpoints, we stop examining the path z_i towards u in a vertex w when $\text{lowpoint}(w) = \gamma(u)$ (the lowpoints and biconnected components in the path from w to u do not change). When cutting the subtree of w , note that there are no back edges connecting it to $B_{u,t}$ (w is an articulation point) and therefore there are no updates to the lists lb and ab of the vertices in $B_{u,t}$. Like in the case of updating the lowpoints, we do not traverse the same tree edge twice (we use the same halting criterion). \square

With Lemma 3.32 and Lemma 3.33 we finalize the proof of Lemma 3.27. Fig. 3.7 shows the changes the bead string $B_{u,t}$ from Fig. 3.5 goes through in the corresponding spine of the recursion tree.

Operation `left_update(C, e)`

In the binary nodes of a spine, we use the fact that in every left branching from that spine the graph is the same (in a spine we only remove edges incident to u and on a left branch from the spine we remove the vertex u) and therefore its block tree is also the same. In Fig. 3.8, we show the resulting block tree of the graph from Fig. 3.5 after having removed vertex u . However, the certificates on these left branches are not the same, as they are rooted at different vertices. In the example we must compute the certificates $C_1 \dots C_4$ corresponding to bead strings $B_{z_1,t} \dots B_{z_4,t}$. We do not account for the cost of the left branch on the last node of spine (corresponding to $B_{v,t}$) as the node is unary and we have shown in Lemma 3.26 how to maintain the certificate in $O(1)$ time.

By using the reverse DFS postorder of the back edges, we are able to traverse each edge in H_X only an amortized constant number of times in the spine.

Lemma 3.34. (Lemma 3.28 restated) *The calls to operation `left_update(C, e)` in a spine of the recursion tree can be charged with a time cost of $O(|E_X|)$ to that spine.*

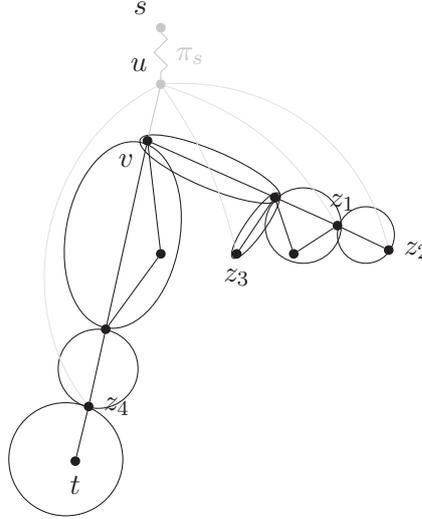
To achieve this time cost, for each back edge $b_i = (z_i, u)$, we compute the certificate corresponding to $B_{z_i,t}$ based on the certificate of $B_{z_{i-1},t}$. Consider the compacted head $H_X = (V_X, E_X)$ of the bead string $B_{u,t}$. We use $O(|E_X|)$ time to compute the first certificate C_1 corresponding to bead string $B_{z_1,t}$. Fig. 3.9 shows bead string $B_{z_1,t}$ from the example of Fig. 3.5.

Lemma 3.35. *The certificate C_1 , corresponding to bead string $B_{z_1,t}$, can be computed in $O(|E_X|)$ time.*

Proof. Let t' be the last vertex in the path $u \rightsquigarrow t$ s.t. $t' \in V_X$. Since t' is an articulation point, the subtree of the DFS tree rooted in t' is maintained in the case of removal vertex u . Therefore the only modifications of the DFS tree occur in head H_X of $B_{u,t}$.

To compute C_1 , we remove u and rebuild the certificate starting from z_1 using the algorithm from Lemma 3.24 restricted to H_X and using t' as target and $\gamma(t')$ as a baseline to γ (instead of the depth). In particular we do the following. To set t' to be in the leftmost path, we perform a DFS traversal of graph H_X starting from z_1 and stop when we reach vertex t' . Then compute the DFS tree, traversing the path $z_1 \rightsquigarrow t'$ first.

Update of γ . For each tree edge (v, w) in the $t' \rightsquigarrow z_1$ path, we set $\gamma(v) = \gamma(w) - 1$, using $\gamma(t')$ as a baseline. During the rest of the traversal, when visiting vertex v , let w be the parent of v in the DFS tree. We set $\gamma(v) = \gamma(w) + 1$. This maintains the property that $\gamma(v) > \gamma(w)$ for any w ancestor of v .

Figure 3.8: Block tree after removing vertex u

Lowpoints and pruning the tree. Bottom-up in the DFS-tree, compute the lowpoints using the lowpoints of the children. For z the parent of v , if $\text{lowpoint}(v) \leq \gamma(z)$ and v is not in the leftmost path in the DFS, cut the subtree of v as it does not belong to $B_{z_1,t}$.

Computing lb and ab . In the traversal, when finding a back edge $e = (v, w)$, if w is a descendant of v we append e to $ab(w)$. This maintains the DFS preorder in the ancestor back edge list. After the first scan of $N(v)$ is over and all the recursive calls returned, re-scan the neighborhood of v . If $e = (v, w)$ is a back edge and w is an ancestor of v , we add e to $lb(w)$. This maintains the DFS postorder in the descendant back edge list. This procedure takes $O(|E_X|)$ time. \square

To compute each certificate C_i , corresponding to bead string $B_{z_i,t}$, we are able to avoid visiting most of the edges that belong $B_{z_{i-1},t}$. Since we take z_i in reverse DFS postorder, on the spine of the recursion we visit $O(|E_X|)$ edges plus a term that can be amortized.

Lemma 3.36. *For each back edge $b_i = (z_i, u)$ with $i > 1$, let $E_X'_i$ be the edges in the first bead in common between $B_{z_i,t}$ and $B_{z_{i-1},t}$. The total cost of computing all certificates $B_{z_i,t}$ in a spine of the recursion tree is: $O(|E_X| + \sum_{i>1} |E_X'_i|)$.*

Proof. Let us compute the certificate C_i : the certificate of the left branch of the i th node of the spine where we augment the path with back edge $b_i = (z_i, u)$ of $lb(u)$.

For the general case of C_i with $i > 1$ we also rebuild (part) of the certificate starting from z_i using the procedure from Lemma 3.24 but we use information gathered in C_{i-1} to avoid exploring useless branches of the DFS tree. The key point is that, when we reach the first bead in common to both $B_{z_i,t}$ and $B_{z_{i-1},t}$, we only explore edges internal to this bead. If an edge e that leaves the bead leads to t , we can reuse a subtree of C_{i-1} . If e does not lead to t , then it has already been explored (and cut) in C_{i-1} and there is no need to explore it again since it is going to be discarded.

In detail, we start computing a DFS from z_i in $B_{u,t}$ until we reach a vertex $t' \in B_{z_{i-1},t}$. Note that the bead of t' has one entry point and one exit point in C_{i-1} . After reaching t' we proceed with the traversal using only edges already in C_{i-1} . When arriving at a vertex w that is not in the same bead of t' , we stop the traversal. If w is in a bead towards t , we reuse

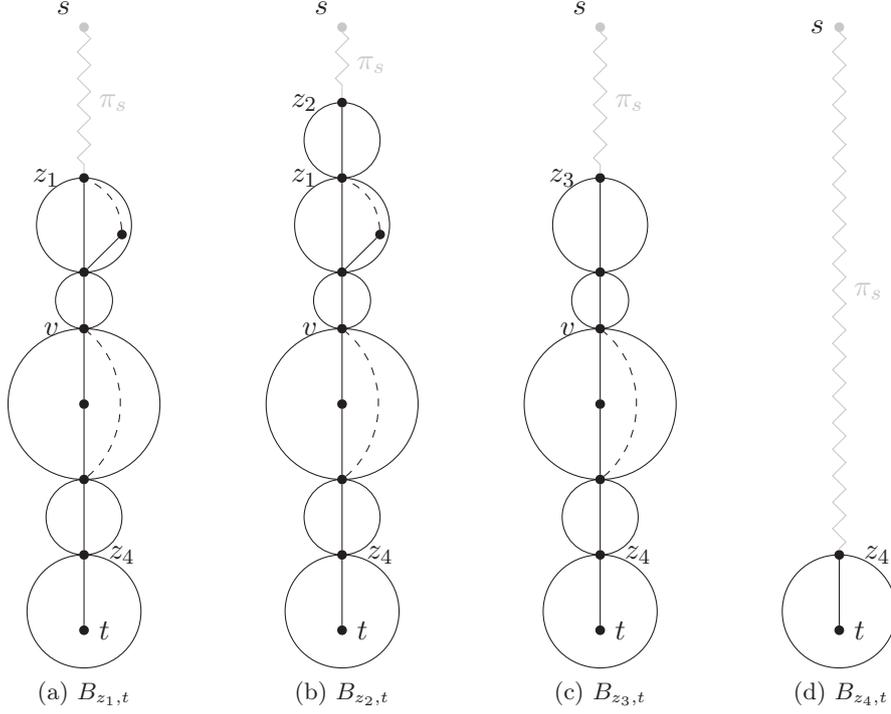


Figure 3.9: Certificates of the left branches of a spine

the subtree of w and use $\gamma(w)$ as a baseline of the numbering γ . Otherwise w is in a bead towards z_{i-1} and we cut this branch of the certificate. When all edges in the bead of t' are traversed, we proceed with visit in the standard way.

Given the order we take b_i , each bead is not added more than once to a certificate C_i , therefore the total cost over the spine is $O(|E_X|)$. Nevertheless, the internal edges $E_{X'_i}$ of the first bead in common between $B_{z_i,t}$ and $B_{z_{i-1},t}$ are explored for each back edge b_i . \square

Although the edges in $E_{X'_i}$ are in a common bead between $B_{z_i,t}$ and $B_{z_{i-1},t}$, these edges must be visited. The entry point in the common bead can be different for z_i and z_{i-1} , the DFS tree of that bead can also be different. For an example, consider the case where z_i, \dots, z_j are all in the same bead after the removal of u . The bead strings $B_{z_i,t} \dots B_{z_j,t}$ are the same, but the roots z_i, \dots, z_j of the certificate are different, so we have to compute the corresponding DFS of the first bead $|j - i|$ times. Note that this is not the case for the other beads in common: the entry point is always the same.

Lemma 3.37. *The cost $O(|E_X| + \sum_{i>1} |E_{X'_i}|)$ on a spine of the recursion tree can be amortized to $O(|E_X|)$.*

Proof. We can charge the cost $O(|E_{X'_i}|)$ of exploring the edges in the first bead in common between $B_{z_i,t}$ and $B_{z_{i-1},t}$ to another node in the recursion tree. Since this common bead is the head of at least one certificate in the recursion subtree of the left child of the i th node of the spine. Specifically, we charge the first and only node in the *leftmost* path of the i th child of the spine that has exactly the edges $E_{X'_i}$ as head of its bead string: (i) if $|E_{X'_i}| \leq 1$ it corresponds to a unary node or a leaf in the recursion tree and therefore we can charge it with $O(1)$ cost; (ii) otherwise it corresponds to a first node of a spine and therefore we can

also charge it with $O(|E_{X'_i}|)$. We use this charging scheme when $i \neq 1$ and the cost is always charged in the leftmost recursion path of i th node of the spine, consequently we never charge a node in the recursion tree more than once. \square

Lemmas 3.36 and 3.37 finalize the proof of Lemma 3.28. Fig. 3.9 shows the certificates of bead strings $B_{z_i,t}$ on the left branches of the spine from Figure 3.6.

3.3 Discussion and conclusions

In the first part of this chapter, we showed that it is possible (Algorithm 3.3) to list all bubbles with a given source in a directed graph with linear delay, thus solving Problem 3.3. Moreover, it is possible (Algorithm 3.1) to enumerate all bubbles, for all possible sources, thus solving Problem 3.2, in $O((m+n)(\eta+n))$ total time, where η is the number of bubbles.

Unfortunately, this algorithm is not a good replacement for KISSPLICE's listing algorithm (Section 2.2.2), since for the task listing bubbles corresponding to AS events, in practice, the latter performs better. Recall that KISSPLICE searches for *cycles* satisfying conditions (i) to (iv) of Section 2.2.2, the cycles satisfying condition (i) correspond to the (s,t) -bubbles, the remaining conditions, (ii) to (iv), are constraints for the length of the sequences corresponding to each path. In KISSPLICE's listing algorithm, several prunings based on these constraints are applied to avoid the enumeration of bubbles that are guaranteed not to satisfy the constraints. On the other hand, Algorithm 3.1 efficiently lists all bubbles directly, i.e. cycles satisfying condition (i), but it is not evident how to apply the same prunings for constraints (ii) to (iv). In the end, we have to list all bubbles and, in a post-processing step, filter out the ones not satisfying the constraints. Since in typical cases, the number of bubbles satisfying the constraints is small compared to the total⁷ number of bubbles, this approach is worse than KISSPLICE's listing algorithm. In Chapter 4, we present a *practical* polynomial delay algorithm that directly lists bubbles satisfying constraints (ii) and (iv).

Nonetheless, the problem of listing bubbles in a directed graph is interesting from a theoretical point of view, since (s,t) -bubbles are natural substructures in directed graphs⁸. Moreover, Algorithm 3.3 required a non-trivial adaptation of Johnson's algorithm (Johnson (1975)) for listing cycles in directed graphs, and is the first linear delay algorithm to list all bubbles with a given source in a directed graph.

In the second part of this chapter, we showed that Johnson's algorithm, the long-standing best known solution to list cycles, is surprisingly not optimal for undirected graphs. We then presented an $O(m + \sum_{c \in \mathcal{C}(G)} |c|)$ algorithm to list cycles in undirected graphs, where $\mathcal{C}(G)$ is the set of cycles and $|c|$ the length of cycle c . Clearly, $\Omega(m)$ time is necessary to read the graph and $\Omega(\sum_{c \in \mathcal{C}(G)} |c|)$ time to list the output. Thus, our algorithm is optimal. Actually, we presented an optimal algorithm to list st -paths in undirected graphs and used an optimality preserving reduction from cycle listing to st -path listing.

This chapter raises some interesting questions, for instance, whether it is possible to directly list the (s,t) -bubbles satisfying path length constraints. In Chapter 4, we give an affirmative answer to this question. Another natural question is whether it is possible to apply

⁷Bubbles not satisfying the constraints correspond to, among others, de Bruijn graph artifacts, other genomic polymorphisms (i.e. inversions), repeat related structures, and, more rarely, multiple exclusive exons.

⁸For instance, (s,t) -bubbles are related to 2-vertex-connected directed graphs (Bang-Jensen and Gutin (2008)) where every pair of vertices are extremities of at least one bubble.

techniques similar to the ones presented in Section 3.2 to improve Johnson's algorithm for directed graphs or Algorithm 3.3. An important invariant maintained by our optimal st -path listing algorithm is the following: in the beginning of every recursive call every edge in the graph is contained in some st -path. Intuitively, this means that at every step the graph is cleaned and only the necessary edges are kept. However, in directed graphs is NP-hard to decide if a given arc belongs to a st -path or (s, t) -bubble (Fortune et al. (1980)). Thus, it is unlikely that the same kind of cleaning can be done in directed graphs. This seems a hard barrier to overcome. The last question is, provided we are only interested in counting⁹, whether it is possible to improve our st -path listing algorithm to $O(|\mathcal{P}_{st}(G)|)$. In other words, is it possible to improve our algorithm to spend only a constant time per path if it is not required to output each st -path. This may seem impossible, but there are listing algorithms achieving this complexity, e.g. listing spanning trees (Marino et al. (2014)).

⁹Counting st -paths is #P-hard, so an algorithm polynomial in the size of the graph is very unlikely.

Chapter 4

Listing in weighted graphs

Contents

4.1 Listing bounded length bubbles in weighted directed graphs . . .	74
4.1.1 Introduction	74
4.1.2 De Bruijn graphs and bounded length bubbles	74
4.1.3 An $O(n(m + n \log n))$ delay algorithm	76
4.1.4 Implementation and experimental results	78
4.1.5 A natural generalization	81
4.2 Listing bounded length paths	85
4.2.1 Introduction	85
4.2.2 Preliminaries	86
4.2.3 A simple polynomial delay algorithm	86
4.2.4 An improved algorithm for undirected graphs	88
4.2.5 Listing paths in increasing order of their lengths	92
4.3 Discussion and conclusions	94

In this chapter, we present efficient algorithms to list paths and bubbles, satisfying path length constraints in *weighted* directed graphs. The chapter is divided in two main parts.

The first part (Section 4.1) is strongly based on our paper [Sacomoto et al. \(2013\)](#), and its goal is to present a polynomial delay algorithm to list all bubbles in weighted directed graphs, such that each path p_1, p_2 in the bubble has length bounded by α_1, α_2 respectively. For a directed graph with n vertices and m arcs, the method we propose lists all bubbles with a given source in $O(n(m + n \log))$ delay. Moreover, we experimentally show that this algorithm is significantly faster than the listing algorithm of KISSPLICE (version 1.6) to identify bubbles corresponding to alternative splicing events.

The second part (Section 4.2) is strongly based on our paper [Grossi et al. \(2014\)](#) (in preparation), and its goal is to present a general scheme to list bounded length st -paths in weighted directed or undirected graphs using memory linear in the size of the graph, independent of the number of paths output. For undirected non-negatively weighted graphs, we also show an improved algorithm that lists all st -paths with length bounded by α in $O((m + t(n, m))\gamma)$ total time, where γ is the number the st -paths with length bounded by α and $t(m, n)$ is the time to compute a shortest path tree. In particular, this is $O(m\gamma)$ for unit weights and $O((m + n \log n)\gamma)$ for general non-negative weights. Moreover, we show how to modify the general scheme to output the paths in increasing order of their lengths.

4.1 Listing bounded length bubbles in weighted directed graphs

4.1.1 Introduction

In the previous chapter, we proposed a linear delay algorithm to list all bubbles in a directed graph, in particular, applicable also to de Bruijn graphs. Although interesting from a theoretical point of view, the algorithm cannot replace the listing algorithm of KISSPLICE (Section 2.2.2), since for the task of listing bubbles corresponding to AS events, in practice, the latter performs better. Indeed, this is due to the fact that the bubbles corresponding to alternative splicing events (excluding mutually exclusive exons) satisfy some path length constraints. We can use this information in simple backtracking algorithm of KISSPLICE (version 1.6) to efficiently prune the branches of the search tree; we cannot, however, give any theoretical guarantees. In the worst case, the algorithm is still exponential in the number of bubbles and the size of the graph. On the other hand, it is not clear how to incorporate the same prunings in the linear delay algorithm of the last chapter. As a result, the algorithm lists a huge number of bubbles that have to be checked, in a post-processing step, for the path length constraints. In this chapter, we present a polynomial delay algorithm to directly list bubbles satisfying the path constraints. Moreover, we experimentally show that the algorithm is several orders of magnitude faster than KISSPLICE’s (version 1.6) listing algorithm.

As stated in Section 3.1, the problem of identifying bubbles with path length constraints was considered before in the genome assembly (Li et al. (2010); Peng et al. (2010); Zerbino and Birney (2008); Simpson et al. (2009)) and in the variant finding (Iqbal et al. (2012); Leggett et al. (2013)) contexts. However, in the first case the goal was not to list all bubbles. In general, assemblers perform a greedy search for bubbles in order to “linearize” a de Bruijn graph. Moreover, the path length constraints are symmetric, that is both paths should satisfy the same length constraint. In the second case, the goal is really to list bubbles, but in Iqbal et al. (2012) the search is restricted to non-branching bubbles, while in Leggett et al. (2013) this constraint is relaxed to a bounded (small) number branching internal vertices. Additionally, in Leggett et al. (2013) there is no strong theoretical guarantee for the time complexity; the algorithm is basically an unconstrained DFS, similar to the listing algorithm of KISSPLICE, where the search is truncated at a given depth.

In this chapter, we introduce the first polynomial delay algorithm to list all bubbles with length constraints in a weighted directed graph. Its complexity for general non-negatively weighted graphs is $O(n(m + n \log n))$ (Section 4.1.3) where n is the number of vertices in the graph, m the number of arcs. In the particular case of de Bruijn graphs, the complexity is $O(n(m + n \log \alpha))$ (Section 4.1.4) where α is a constant related to the length of the skipped part in an alternative splicing event. In practice, an algorithmic solution in $O(nm \log n)$ (Section 4.1.4) appears to work better on de Bruijn graphs built from such data. We implemented the latter, show that it is more efficient than previous approaches and outline that it allows to discover novel long alternative splicing events.

4.1.2 De Bruijn graphs and bounded length bubbles

As was shown in Chapter 2, polymorphisms (i.e. variable parts) in a transcriptome (including alternative splicing events) correspond to recognizable patterns in the DBG that are precisely the (s, t) -bubbles (Definition 3.1). Intuitively, the variable parts correspond to alternative paths and the common parts correspond to the beginning and end points of those paths. More formally, any process generating patterns awb and $aw'b$ in the sequences, with $a, b, w, w' \in \Sigma^*$, $|a| \geq k, |b| \geq k$ and w and w' not sharing any k -mer, creates a (s, t) -bubble in the DBG. In

the special case of AS events (excluding mutually exclusive exons), since w' is empty, one of the paths corresponds to the *junction* of ab , i.e. to k -mers that contain at least one letter of each sequence. Thus the number of vertices of this path in the DBG is predictable: it is at most¹ $k - 1$. An example is given in Fig. 4.1. In practice (see Section 2.2.1), an upper bound α to the other path and a lower bound β on both paths is also imposed. In other words, an AS event corresponds to a (s, t) -bubble with paths p_1 and p_2 such that p_1 has at most α vertices, p_2 at most $k - 1$ and both have at least β vertices.

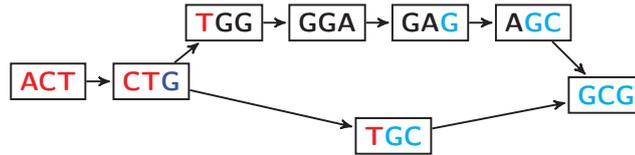


Figure 4.1: DBG with $k = 3$ for the sequences: **ACTGGAGCG** (awb) and **ACTGCG** (ab). The pattern in the sequence generates a (s, t) -bubble, from **CTG** to **GCG**. In this case, $b = \text{GCG}$ and $w = \text{GGA}$ have their first letter **G** in common, so the path corresponding to the junction ab has $k - 1 - 1 = 1$ vertex.

Given a directed graph G with non-negative arc weights $w : E \mapsto \mathbb{Q}_{\geq 0}$, we can extend Definition 3.1 to G by considering (s, t) -bubbles with length constraints in both paths.

Definition 4.1 ($(s, t, \alpha_1, \alpha_2)$ -bubble). *A $(s, t, \alpha_1, \alpha_2)$ -bubble in a weighted directed graph is a (s, t) -bubble with paths p_1, p_2 satisfying $w(p_1) \leq \alpha_1$ and $w(p_2) \leq \alpha_2$.*

As stated in Chapter 2, when dealing with DBGs built from RNA-seq data, in a lossless preprocessing step, all maximal non-branching linear paths of the graph (i.e. paths containing only vertices with in and out-degree 1) are compressed each into one single vertex, whose label corresponds to the label of the path (i.e. it is the concatenation of the labels of the vertices in the path without the overlapping part(s)). The resulting graph is the compressed de Bruijn graph (cDBG). In the cDBG, the vertices can have labels larger than k , but an arc still indicates a suffix-prefix overlap of size $k - 1$. Finally, since the only property of a bubble corresponding to an AS event is the constraint on the length of the path, we can disregard the labels from the cDBG and only keep for each vertex its label length². In this way, searching for bubbles corresponding to AS events in a cDBG can be seen as a particular case of looking for $(s, t, \alpha_1, \alpha_2)$ -bubbles satisfying the lower bound β in a non-negative weighted directed graph.

Actually, it is not hard to see that the enumeration of $(s, t, \alpha_1, \alpha_2)$ -bubbles, for all s and t , satisfying the lower bound β is NP-hard. Indeed, deciding the existence of at least one $(s, t, \alpha_1, \alpha_2)$ -bubble, for some s and t , with the lower bound β in a weighted directed graph where all the weights are 1 is NP-complete. It follows by a simple reduction from the Hamiltonian path problem (Garey and Johnson (1979)): given a directed graph $G = (V, E)$ and two vertices s and t , build the graph G' by adding to G the vertices s' and t' , the arcs (s, s') and (t, t') , and a new path from s' to t' with exactly $|V|$ nodes. There is a $(x, y, |V|+2, |V|+2)$ -bubble, for some x and y , satisfying the lower bound $\beta = |V| + 2$ in G' if and only if there is a Hamiltonian path from s to t in G .

¹The size is *exactly* $k - 1$ if w has no common prefix with b and no common suffix with a .

²Resulting in a graph with weights in the vertices. Here, however, we consider the weights in the arcs. Since this is more standard and, in our case, both alternatives are equivalent, we can transform one into another by splitting vertices or arcs.

From now on, we consider the more general problem of listing $(s, t, \alpha_1, \alpha_2)$ -bubbles (without the lower bound) for an arbitrary non-negative weighted directed graph G (not restricted to a cDBG).

Problem 4.2 (Listing bounded length bubbles). *Given a non-negatively weighted directed graph $G = (V, E)$, output all $(s, t, \alpha_1, \alpha_2)$ -bubbles, for all pairs $s, t \in V$.*

In order to solve Problem 3.2, we consider the problem of listing all bubbles with a given source (Problem 3.3). Indeed, by trying all possible sources s we can list all (s, t) -bubbles.

Problem 4.3 (Listing $(s, *, \alpha_1, \alpha_2)$ -bubbles). *Given a non-negatively weighted directed graph $G = (V, E)$ and vertex s , output all $(s, t, \alpha_1, \alpha_2)$ -bubbles, for all $t \in V$.*

The number of vertices and arcs of G is denoted by n and m , respectively.

4.1.3 An $O(n(m + n \log n))$ delay algorithm

In this section, we present an $O(n(m + n \log n))$ delay algorithm to enumerate, for a fixed source s , all $(s, t, \alpha_1, \alpha_2)$ -bubbles in a general directed graph G with non-negative weights. The pseudocode is shown in Algorithm 4.1. It is important to stress that this pseudocode uses high-level primitives, e.g. the tests in lines 5, 9 and 14. An efficient implementation for the test in line 9, along with its correctness and analysis, is implicitly given in Lemma 4.6. This is a central result in this section. For its proof we need Lemma 4.4.

Algorithm 4.1 uses a recursive strategy, inspired by the binary partition method, that successively divides the solution space at every call until the considered subspace is a singleton. In order to have a more symmetric structure for the subproblems, we define the notion of a *pair of compatible paths*, which is an object that generalizes the definition of a $(s, t, \alpha_1, \alpha_2)$ -bubble. Given two vertices $s_1, s_2 \in V$ and upper bounds $\alpha_1, \alpha_2 \in \mathbb{Q}_{\geq 0}$, the paths $p_1 = s_1 \rightsquigarrow t_1$ and $p_2 = s_2 \rightsquigarrow t_2$ are a *pair of compatible paths* for s_1 and s_2 if $t_1 = t_2$, $w(p_1) \leq \alpha_1$, $w(p_2) \leq \alpha_2$ and the paths are internally vertex-disjoint. Clearly, every $(s, t, \alpha_1, \alpha_2)$ -bubble is also a pair of compatible paths for $s_1 = s_2 = s$ and some t .

Given a vertex v , the set of out-neighbors of v is denoted by $N^+(v)$. Let now $\mathcal{P}_{\alpha_1, \alpha_2}(s_1, s_2, G)$ be the set of all pairs of compatible paths for s_1, s_2, α_1 and α_2 in G . We have³ that:

$$\mathcal{P}_{\alpha_1, \alpha_2}(s_1, s_2, G) = \mathcal{P}_{\alpha_1, \alpha_2}(s_1, s_2, G') \cup \bigcup_{v \in N^+(s_2)} (s_2, v) \mathcal{P}_{\alpha_1, \alpha'_2}(s_1, v, G - s_2), \quad (4.1)$$

where $\alpha'_2 = \alpha_2 - w(s_2, v)$ and $G' = G - \{(s_2, v) | v \in N^+(s_2)\}$. In other words, the set of pairs of compatible paths for s_1 and s_2 can be partitioned into: $\mathcal{P}_{\alpha_1, \alpha'_2}(s_1, v, G - s_2)$, the sets of pairs of paths containing the arc (s_2, v) , for each $v \in N^+(s_2)$; and $\mathcal{P}_{\alpha_1, \alpha_2}(s_1, s_2, G')$, the set of pairs of paths that do not contain any of them. Algorithm 4.1 implements this recursive partition strategy. The solutions are only output in the leaves of the recursion tree (line 3), where the partition is always a singleton. Moreover, in order to guarantee that every leaf in the recursion tree outputs at least one solution, we have to test if $\mathcal{P}_{\alpha_1, \alpha'_2}(s_1, v, G - s_2)$ (and $\mathcal{P}_{\alpha_1, \alpha_2}(s_1, s_2, G')$) is not empty before making the recursive call (lines 9 and 14).

The correctness of Algorithm 4.1 follows directly from the relation given in Eq. 4.1 and the correctness of the tests performed in lines 9 and 14. In the remaining of this section, we describe a possible implementation for the tests, prove correctness and analyze the time complexity. Finally, we prove that Algorithm 4.1 has an $O(n(m + n \log n))$ delay.

³The same relation is true using s_1 instead of s_2 .

Algorithm 4.1: `enumerate_bubbles`($s_1, \alpha_1, s_2, \alpha_2, B, G$)

```

1 if  $s_1 = s_2$  then
2   if  $B \neq \emptyset$  then
3     output( $B$ )
4     return
5   else if there is no  $(s, t, \alpha_1, \alpha_2)$ -bubble, where  $s = s_1 = s_2$  then
6     return
7 choose  $u \in \{s_1, s_2\}$ , such that  $N^+(u) \neq \emptyset$ 
8 for  $v \in N^+(u)$  do
9   if there is a pair of compatible paths using  $(u, v)$  in  $G$  then
10    if  $u = s_1$  then
11      enumerate_bubbles( $v, \alpha_1 - w(s_1, v), s_2, \alpha_2, B \cup (s_1, v), G - s_1$ )
12    else
13      enumerate_bubbles( $s_1, \alpha_1, v, \alpha_2 - w(s_2, v), B \cup (s_2, v), G - s_2$ )
14 if there is a pair of compatible paths in  $G - \{(u, v) | v \in N^+(u)\}$  then
15   enumerate_bubbles( $s_1, \alpha_1, s_2, \alpha_2, B, G - \{(u, v) | v \in N^+(u)\}$ )

```

Lemma 4.4. *There exists a pair of compatible paths for $s_1 \neq s_2$ in G if and only if there exists t such that $d(s_1, t) \leq \alpha_1$ and $d(s_2, t) \leq \alpha_2$.*

Proof. Clearly this is a necessary condition. Let us prove that it is also sufficient. Consider the paths $p_1 = s_1 \rightsquigarrow t$ and $p_2 = s_2 \rightsquigarrow t$, such that $w(p_1) \leq \alpha_1$ and $w(p_2) \leq \alpha_2$. Let t' be the first vertex in common between p_1 and p_2 . The sub-paths $p'_1 = s_1 \rightsquigarrow t'$ and $p'_2 = s_2 \rightsquigarrow t'$ are internally vertex-disjoint, and since the weights are non-negative, they also satisfy $w(p'_1) \leq w(p_1) \leq \alpha_1$ and $w(p'_2) \leq w(p_2) \leq \alpha_2$. \square

Using this lemma, we can test for the existence of a pair of compatible paths for $s_1 \neq s_2$ in $O(m + n \log n)$ time. Indeed, let T_1 be a shortest path tree of G rooted in s_1 and truncated at distance α_1 , the same for T_2 , meaning that, for any vertex w in T_1 (resp. T_2), the tree path between s_1 and w (resp. s_2 and w) is a shortest one. It is not difficult to prove that the intersection $T_1 \cap T_2$ is not empty if and only if there is a pair of compatible paths for s_1 and s_2 in G . Moreover, each shortest path tree can be computed in $O(m + n \log n)$ time, using Dijkstra's algorithm (Cormen et al. (2001)). Thus, in order to test for the existence of a $(s, t, \alpha_1, \alpha_2)$ -bubble for some t in G , we can test, for each arc (s, v) outgoing from s , the existence of a pair of compatible paths for $s \neq v$ and v in G . Since s has at most n out-neighbors, we obtain Lemma 4.5.

Lemma 4.5. *The test of line 5 can be performed in $O(n(m + n \log n))$.*

The test of line 9 could be implemented using the same idea. For each $v \in N^+(u)$, we test for the existence of a pair of compatible paths for, say, $u = s_2$ (the same would apply for s_1) and v in $G - u$, that is v is in the subgraph of G obtained by eliminating from G the vertex u and all the arcs incoming to or outgoing from u . This would lead to a total cost of $O(n(m + n \log n))$ for all tests of line 9 in each call. However, this is not enough to achieve an $O(n(m + n \log n))$ delay. In Lemma 4.6, we present an improved strategy to perform these tests in $O(m + n \log n)$ total time.

Lemma 4.6. *The test of line 9, for all $v \in N^+(u)$, can be performed in $O(m + n \log n)$ total time.*

Proof. Let us assume that $u = s_2$, the case $u = s_1$ is symmetric. From Lemma 4.4, for each $v \in N^+(u)$, we have that deciding if there exists a pair of compatible paths for s_1 and s_2 in G that uses (u, v) is equivalent to deciding if there exists t satisfying (i) $d(s_1, t) \leq \alpha_1$ and (ii) $d(v, t) \leq \alpha_2 - w(u, v)$ in $G - u$.

First, we compute a shortest path tree rooted in s_1 for $G - u$. Let V_{α_1} be the set of vertices at a distance at most α_1 from s_1 . We build a graph G' by adding a new vertex r to $G - u$, and for each $y \in V_{\alpha_1}$, we add the arcs (y, r) with weight $w(y, r) = 0$. We claim that there exists t in $G - u$ satisfying conditions (i) and (ii) if and only if $d(v, r) \leq \alpha_2 - w(u, v)$ in G' . Indeed, if t satisfies (i) we have that the arc (t, r) is in G' , so $d(t, r) = 0$. From the triangle inequality and (ii), $d(v, r) \leq d(v, t) + d(t, r) = d(v, t) \leq \alpha_2 - w(u, v)$. The other direction is trivial.

Finally, we compute a shortest path tree T_r rooted in r for the reverse graph G'^R , obtained by reversing the direction of the arcs of G' . With T_r , we have the distance from any vertex to r in G' , i.e. we can answer the query $d(v, r) \leq \alpha_2 - w(u, v)$ in constant time. Observe that the construction of T_r depends only on $G - u$, s_1 and α_1 , i.e. T_r is the same for all out-neighbors $v \in N^+(u)$. Therefore, we can build T_r only once in $O(m + n \log n)$ time, with two iterations of Dijkstra's algorithm, and use it to answer each test of line 9 in constant time. \square

Theorem 4.7. *Algorithm 4.1 has $O(n(m + n \log n))$ delay.*

Proof. The height of the recursion tree is bounded by $2n$ since at each call the size of the graph is reduced either by one vertex (lines 11 and 13) or all its out-neighborhood (line 15). After at most $2n$ recursive calls, the graph is empty. Since every leaf of the recursion tree outputs a solution and the distance between two leaves is bounded by $4n$, the delay is $O(n)$ multiplied by the cost per node (call) in the recursion tree. From Lemma 4.4, line 14 takes $O(m + n \log n)$ time, and from Lemma 4.6, line 9 takes $O(m + n \log n)$ total time. This leads to an $O(m + n \log n)$ time per call, excluding line 5. Lemma 4.5 states that the cost for the test in line 5 is $O(n(m + n \log n))$, but this line is executed only once, at the root of the recursion tree. Therefore, the delay is $O(n(m + n \log n))$. \square

4.1.4 Implementation and experimental results

We now discuss the details necessary for an efficient implementation of Algorithm 4.1 and the results on two sets of experimental tests. For the first set, our goal is to compare the running time of Dijkstra's algorithm (for typical cDBGs arising from applications) using several priority queue implementations. With the second set, our objective is to compare an implementation of Algorithm 4.1 to the KISSPLICE listing algorithm given in Section 2.2.2. For both cases, we retrieved from the *Short Read Archive* (accession code ERX141791) 14M Illumina 79bp single-ended reads of a *Drosophila melanogaster* RNA-seq experiment. We then built the de Bruijn graph for this dataset with $k = 31$. In order to remove likely sequencing errors, we discarded all k -mers that are present less than 3 times in the dataset. The resulting graph contained 22M k -mers, which after compressing all maximal linear paths, corresponded to 600k vertices.

In order to perform a fair comparison with KISSPLICE, we pre-processed the graph as described in Section 2.2.2. Namely, we decomposed the underlying undirected graph into biconnected components (BCCs) and compressed all non-branching bubbles with equal path

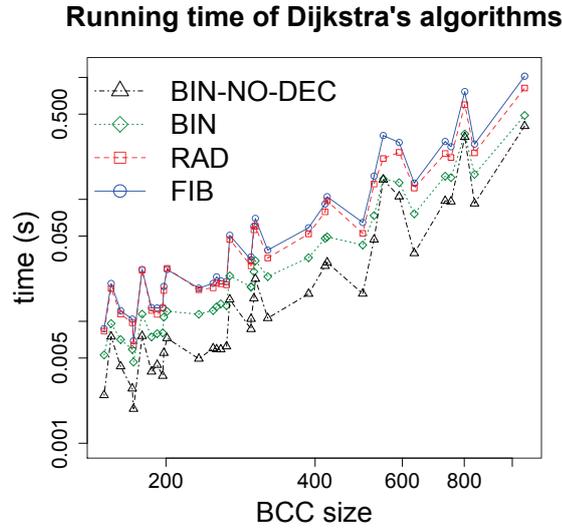


Figure 4.2: Running times for each version of Dijkstra’s algorithm: using Fibonacci heaps (FIB), using radix heaps (RAD), using binary heaps (BIN) and using binary heaps without the decrease-key operation (BIN-NO-DEC). The tests were done including all BCCs with more than 150 vertices. Both axes are in logarithmic scale.

lengths. In the end, after discarding all BCCs with less than 4 vertices (as they cannot contain a bubble), we obtained 7113 BCCs, the largest one containing 24977 vertices. This pre-processing is lossless, i.e. every bubble in the original graph is entirely contained in exactly one BCC. In KISSPLICE, the enumeration is then done in each BCC independently.

Dijkstra’s algorithm with different priority queues

Dijkstra’s algorithm is an important subroutine of Algorithm 4.1 that may have a big influence on its running time. Actually, the time complexity of Algorithm 4.1 can be written as $O(nt(n, m))$, where $t(n, m)$ is the complexity of Dijkstra’s algorithm. There are several variants of this algorithm (Cormen et al. (2001)), with different complexities depending on the priority queue used, including binary heaps ($O(m \log n)$) and Fibonacci heaps ($O(m+n \log n)$). In the particular case where all the weights are non-negative integers bounded by C , Dijkstra’s algorithm can be implemented using radix heaps ($O(m + n \log C)$) (Ahuja et al. (1990)). As stated in Section 4.1.2, the weights of the de Bruijn graphs considered here are integer, but not necessarily bounded. However, we can remove from the graph all arcs with weights greater than α_1 since these are not part of any $(s, t, \alpha_1, \alpha_2)$ -bubble. This results in a complexity of $O(m + n \log \alpha_1)$ for Dijkstra’s algorithm.

We implemented four versions of Lemma 4.5 (for deciding whether there exists a $(s, t, \alpha_1, \alpha_2)$ -bubble for a given s) each using a different version of Dijkstra’s algorithm: with Fibonacci heaps (FIB), with radix heaps (RAD), with binary heaps (BIN) and with binary heaps without decrease-key operation (BIN-NO-DEC). The last version is Dijkstra’s modified in order not to use the decrease-key operation so that we can use a simpler binary heap that does not support such operation (Chen et al. (2007)). We then ran the four versions, using $\alpha_1 = 1000$ and $\alpha_2 = 2k - 2 = 60$, for each vertex in all the BCCs with more than 150 vertices. The

results are shown⁴ in Fig. 4.2. Contrary to the theoretical predictions, the versions with the best complexities, FIB and RAD, have the worst results on this type of instances. It is clear that the best version is BIN-NO-DEC, which is at least 2.2 times and at most 4.3 times faster than FIB. One of the factors possibly contributing to a better performance of BIN and BIN-NO-DEC is the fact that cDBGs, as stated in Section 4.1.2, have bounded degree and are therefore sparse.

Comparison with the KISSPLICE algorithm

In this section, we compare Algorithm 4.1 to the KISSPLICE enumeration algorithm given in Section 2.2.2. To this purpose, we implemented Algorithm 4.1 using Dijkstra’s algorithm with binary heaps without the decrease-key operation for all shortest paths computation. In this way, the delay of Algorithm 4.1 becomes $O(nm \log n)$, which is worse than the one using Fibonacci or radix heaps, but is faster in practice. The goal of the KISSPLICE enumeration is to find all the potential alternative splicing events in a BCC, i.e. to find all $(s, t, \alpha_1, \alpha_2)$ -bubbles satisfying also the lower bound constraint (Section 4.1.2). In order to compare KISSPLICE (version 1.6) to Algorithm 4.1, we (naively) modified the latter so that, whenever a $(s, t, \alpha_1, \alpha_2)$ -bubble is found, we check whether it also satisfies the lower bound constraints and output it only if it does.

In KISSPLICE, the upper bound α_1 is an open parameter, $\alpha_2 = k - 1$ and the lower bound is $k - 7$. Moreover, there are two stop conditions: either when more than 10000 $(s, t, \alpha_1, \alpha_2)$ -bubbles satisfying the lower bound constraint have been enumerated or a 900s timeout has been reached. We ran both KISSPLICE (version 1.6) and the modified Algorithm 4.1, with the stop conditions, for all 7113 BCCs, using $\alpha_2 = 60$, a lower bound of 54 and $\alpha_1 = 250, 500, 750$ and 1000. The running times for all BCCs with more than 150 vertices (there are 37) is shown⁵ in Fig. 4.3. For the BCCs smaller than 150 vertices, both algorithms have comparable (very small) running times. For instance, with $\alpha_1 = 250$, KISSPLICE runs in 17.44s for *all* 7113 BCCs with less than 150 vertices, while Algorithm 4.1 runs in 15.26s.

The plots in Fig. 4.3 show a trend of increasing running times for larger BCCs, but the graphs are not very smooth, i.e. there are some sudden decreases and increases in the running times observed. This is in part due to the fact that the time complexity of Algorithm 4.1 is output sensitive. The delay of the algorithm is $O(nm \log n)$, but the total time complexity is $O(|\mathcal{B}|nm \log n)$, where $|\mathcal{B}|$ is the number of $(s, t, \alpha_1, \alpha_2)$ -bubbles in the graph. The number of bubbles in the graph depends on its internal structure. A large graph does not necessarily have a large number of bubbles, while a small graph may have an exponential number of bubbles. Therefore, the value of $|\mathcal{B}|nm \log n$ can decrease by increasing the size of the graph.

Concerning now the comparison between the algorithms, as we can see in Fig. 4.3, Algorithm 4.1 is usually several times faster (keep in mind that the axes are in logarithmic scale) than KISSPLICE, with larger differences when α_1 increases (10 to 1000 times faster when $\alpha_1 = 1000$). In some instances however, KISSPLICE is faster than Algorithm 4.1, but (with only one exception for $\alpha_1 = 250$ and $\alpha_1 = 500$) they correspond either to very small instances or to cases where only 10000 bubbles were enumerated and the stop condition was met. Finally, using Algorithm 4.1, the computation finished within 900s for all but 3 BCCs, whereas using KISSPLICE, 11 BCCs remained unfinished after 900s. The improvement in time

⁴The results for the largest BCC were omitted from the plot to improve the visualization. It took 942.15s for FIB and 419.84s for BIN-NO-DEC.

⁵The BCCs where *both* algorithms reach the timeout were omitted from the plots to improve the visualization. For $\alpha_1 = 250, 500, 750$ and 1000 there are 1, 2, 3 and 3 BCCs omitted, respectively.

therefore enables us to have access to bubbles that could not be enumerated with the previous approach.

On the usefulness of larger values of α_1

In KISSPLICE (version 1.6), the value of α_1 was experimentally set to 1000 due to performance issues, as indeed the algorithm quickly becomes impractical for larger values. On the other hand, the results of Section 4.1.4 suggest that Algorithm 4.1, that is faster than KISSPLICE, can deal with larger values of α_1 . From a biological point of view, it is a priori possible to argue that $\alpha_1 = 1000$ is a reasonable choice, because 87% of annotated exons in *Drosophila* indeed are shorter than 1000 bp (Pruitt et al. (2009)). However, missing the top 13% may have a big impact on downstream analyses of AS, not to mention the possibility that not yet annotated AS events could be enriched in long skipped exons. In this section, we outline that larger values of α_1 indeed produces more results that are biologically relevant. For this, we exploit another RNA-seq dataset, with deeper coverage.

To this purpose, we retrieved 32M RNA-seq reads from the human brain and 39M from the human liver from the Short Read Archive (accession number ERP000546). Next, we built the de Bruijn graph with $k = 31$ for both datasets, then merged and decomposed the DBG into 5692 BCCs (containing more than 10 vertices). We ran Algorithm 4.1 for each BCC with $\alpha_1 = 5000$. It took 4min25s for Algorithm 4.1 to run on all BCCs, whereas KISSPLICE, even using $\alpha_1 = 1000$, took 31min45s, almost 8 times more. There were 59 BCCs containing at least one bubble with the length of the longest path strictly larger than 1000bp potentially corresponding to alternative splicing events. In Fig. 4.4, we show one of those bubbles mapped to the reference genome. It corresponds to an exon skipping in the PRRC2B human gene, the skipped exon containing 2069 bp. While the transcript containing the exon is annotated, the variant with the exon skipped is not annotated.

Furthermore, we ran TRINITY on the same dataset and found that it was unable to report this novel variant. Our method therefore enables us to find new AS events, reported by no other method. This is, of course, just an indication of the usefulness of our approach when compared to a full-transcriptome assembler.

4.1.5 A natural generalization

An intractable case: Paths with length constraints

For the sake of theoretical completeness, in this section, we extend the definition of $(s, t, \alpha_1, \alpha_2)$ -bubble to the case where the length constraints concern d vertex-disjoint paths, for an arbitrary but fixed d . This situation also arises in real data, when more than 2 variants share the same flanking splice sites (for instance for single and double exon skipping), or when a SNP has 3 variants.

Definition 4.8 ((s, t, A) - d -bubble). *Let d be a natural number and $A = \{\alpha_1, \dots, \alpha_d\} \subset \mathbb{Q}_{\geq 0}$. Given a directed weighted graph G and two vertices s and t , an (s, t, A) - d -bubble is a set of d pairwise internally vertex-disjoint paths $\{p_1, \dots, p_d\}$, satisfying $p_i = s \rightsquigarrow t$ and $w(p_i) \leq \alpha_i$, for all $i \in [1, d]$.*

Analogously to $(s, t, \alpha_1, \alpha_2)$ -bubbles, we can define two variants of the enumeration problem: all bubbles with a given source (s fixed) and all bubbles with a given source and target (s and t fixed). In both cases, the first step is to decide the existence of at least one (s, t, A) - d -bubble in the graph.

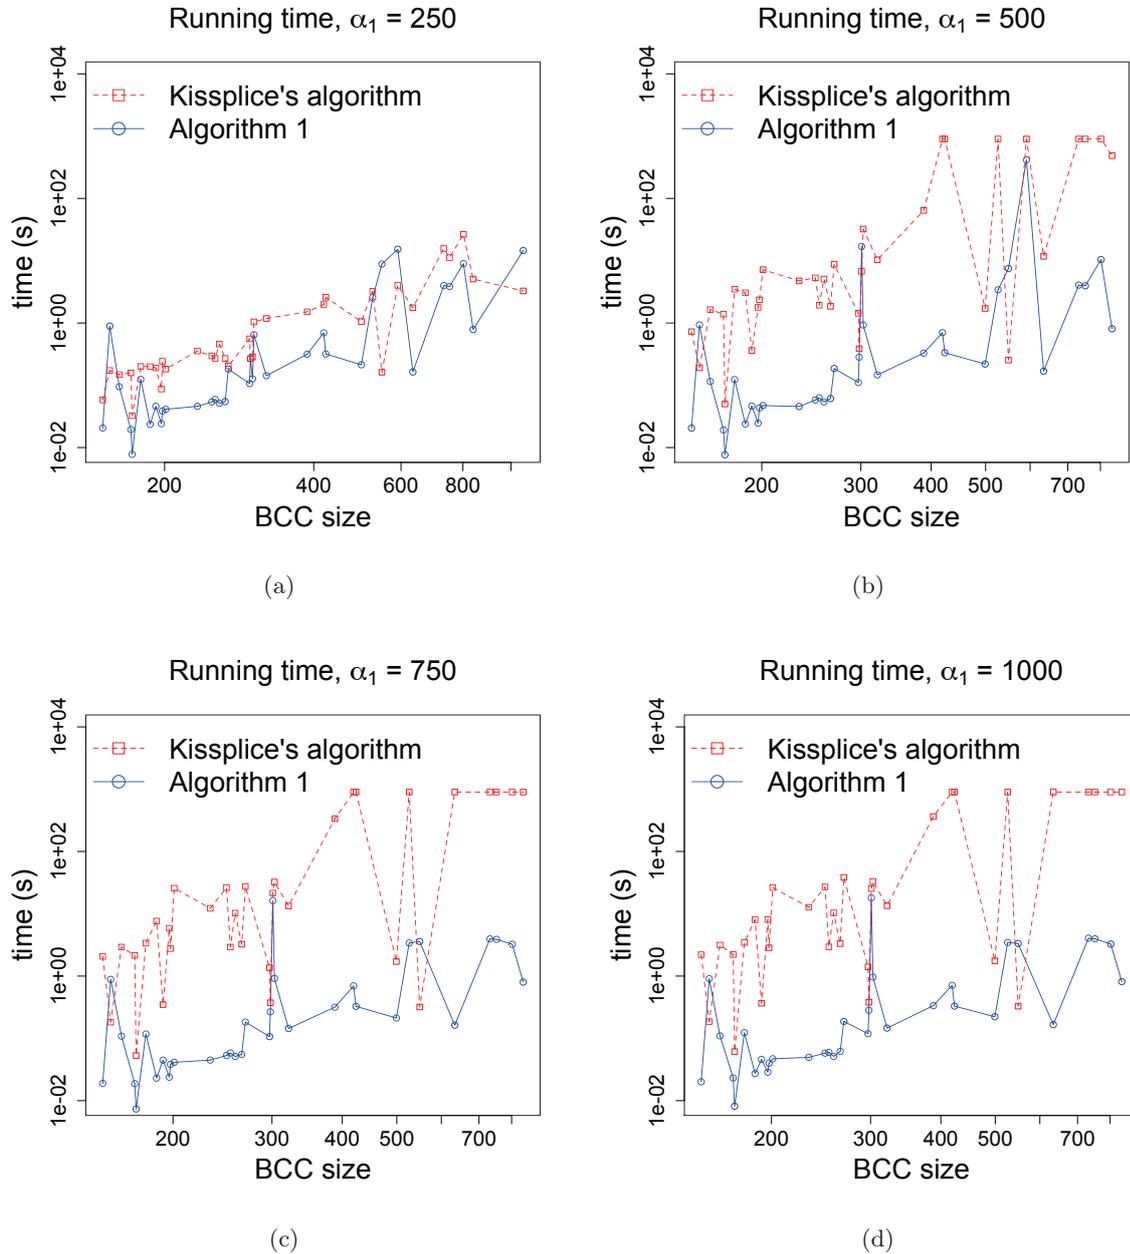


Figure 4.3: Running times of Algorithm 4.1 and of the KISSPLICE bubble listing algorithm for all the BCCs with more than 150 vertices. Each graph (a), (b), (c) and (d) shows the running time of both algorithms for $\alpha_1 = 250, 500, 750$ and 1000 , respectively.

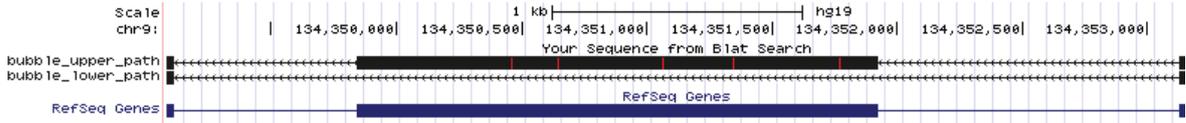


Figure 4.4: One of the bubbles with longest path larger than 1000 bp found by Algorithm 4.1 with the corresponding sequences mapped to the reference genome and visualized using the UCSC Genome Browser. The first two lines correspond to the sequences of, respectively, the shortest (exon exclusion variant) and longest paths of the bubble mapped to the genome. The blue lines are the UCSC human transcript annotations.

Problem 4.9 ($((s, t, A)$ - d -bubble decision problem). *Given a non-negatively weighted directed graph G , two vertices s, t , a set $A = \{\alpha_1, \dots, \alpha_d\} \subset \mathbb{Q}_{\geq 0}$ and $d \in \mathbb{N}$, decide if there exists a (s, t, A) - d -bubble.*

This problem is a generalization of the two-disjoint-paths problem with a min-max objective function, which is NP-complete (Li et al. (1990)). More formally, this problem can be stated as follows: given a directed graph G with non-negative weights, two vertices $s, t \in V$, and a maximum length M , decide if there exists a pair of vertex-disjoint paths such that the maximum of their lengths is less than M . The (s, t, A) - d -bubble decision problem, with $A = \{M, M\}$ and $d = 2$, is precisely this problem.

Problem 4.10 ($((s, *, A)$ - d -bubble decision problem). *Given a non-negatively weighted directed graph G , a vertex s , a set $A = \{\alpha_1, \dots, \alpha_d\} \subset \mathbb{Q}_{\geq 0}$ and $d \in \mathbb{N}$, decide if there exists a (s, t, A) - d -bubble, for some $t \in V$.*

The two-disjoint-path problem with a min-max objective function is NP-complete even for strictly positive weighted graphs. Let us reduce Problem 4.10 to it. Consider a graph G with strictly positive weights, two vertices $s, t \in V$, and a maximum length M . Construct the graph G' by adding an arc with weights 0 from s to t and use this as input for the $(s, *, \{M, M, 0\})$ -3-bubble decision problem. Since G has strictly positive weights, the only path with length 0 from s to t in G' is the added arc. Thus, there is a $(s, *, \{M, M, 0\})$ -3-bubble in G' if and only if there are two vertex-disjoint paths in G each with a length $\leq M$.

Therefore, the decision problem for fixed s (Problem 4.9) is NP-hard for $d \geq 2$, and for fixed s and t (Problem 4.10) is NP-hard for $d \geq 3$. In other words, the only tractable case is the enumeration of (s, t, A) -2-bubbles with fixed s , the one considered in Section 4.1.3.

A tractable case: Paths without length constraints

In the previous section, we showed that a natural generalization of $(s, t, \alpha_1, \alpha_2)$ -bubbles to contain more than two vertex-disjoint paths satisfying length constraints leads to an NP-hard enumeration problem. Indeed, even deciding the existence of at least one (s, t, \mathcal{A}) - d -bubble is NP-hard. In this section, we consider a similar generalization for (s, t) -bubbles instead of $(s, t, \alpha_1, \alpha_2)$ -bubbles, that is, we consider bubbles containing more than two vertex-disjoint paths without any path length constraints. The formal definition is given below.

Definition 4.11 ($((s, t)$ - d -bubble). *Let d be a natural number. Given a directed graph G and two vertices s and t , a (s, t) - d -bubble is a set of d pairwise internally vertex-disjoint paths $\{p_1, \dots, p_d\}$.*

Clearly, this definition is a special case of Definition 4.8: consider a weighted graph $G = (V, E)$ with unitary weights (i.e. an unweighted graph), the (s, t, \mathcal{A}) - d -bubbles with $\alpha_i = |V|$ for $i \in [1, d]$ are precisely the (s, t) - d -bubbles of G . As in Section 4.1.5, let us first consider the problem of deciding whether a graph contains a (s, t) - d -bubble for fixed s and t .

Problem 4.12 ((s, t) - d -bubble decision problem). *Given a directed graph G and two vertices s, t , decide whether there exists a (s, t) - d -bubble in G .*

Contrary to Problem 4.9, this problem can be decided in polynomial time. Indeed, given a directed graph $G = (V, A)$ and two vertices s and t , construct the graph $G' = (V', A')$ by splitting each vertex $v \in V$ in two vertices: an incoming part v_{in} with all the arcs entering v , and an outgoing part v_{out} with all the arcs leaving v ; and add the arc (v_{in}, v_{out}) . More formally, G' is defined as $V' = \{\{v_{in}, v_{out}\} | v \in V\}$ and $A' = \{(u_{out}, v_{in}) | (u, v) \in A\} \cup \{(v_{in}, v_{out}) | v \in V\}$. Now, it is not hard to prove that every set of *arc-disjoint* paths in G' corresponds to a set of vertex-disjoint paths in G . Thus, considering G' a network with unitary arc capacities (Cormen et al. (2001)), we have that G contains a (s, t) - d -bubble if and only if G' contains a (s, t) -flow f such that $|f| \geq d$. Therefore, using the augmenting path algorithm (Cormen et al. (2001)) for the max-flow problem, we can decide if there exists a (s, t) - d -bubble in G in $O(md)$ time. Actually, using an iterative decomposition of the (s, t) -flow f into (s, t) -paths, we can explicitly find a (s, t) - d -bubble in the time bound.

Lemma 4.13. *Given a directed graph $G = (V, A)$ and two vertices $s, t \in V$, a (s, t) - d -bubble in G can be found in $O(md)$ time.*

We now consider the problem of enumerating (s, t) - d -bubbles in G for fixed s and t . The reduction from (s, t) - d -bubbles to (s, t) -flows used in the last paragraph may induce us to think that we can enumerate (s, t) - d -bubbles in G by enumerating (s, t) -flows in G' , and since there is a polynomial delay algorithm for the latter (Bussieck and Lubbecke (1998)), we would be done. Unfortunately, there is no one-to-one correspondence between (s, t) -flows in G' and (s, t) - d -bubbles in G : we can always add a circulation c to a (s, t) -flow f to obtain a new (s, t) -flow f' , but f and f' correspond to the same (s, t) - d -bubble. In fact, there can be exponentially more (s, t) -flows in G' than (s, t) - d -bubbles in G . On the other hand, the strategy used in Algorithm 4.1 can be adapted to enumerate (s, t) - d -bubbles.

Similarly to Section 4.1.3, in order to have a more symmetric structure for the subproblems, we define the notion of a *set of compatible paths*, which is an object that generalizes the definition of a (s, t) - d -bubble. Given a set of sources $S = \{s_1, \dots, s_d\}$ and a target t , a set of paths $P_t = \{p_1, \dots, p_d\}$ is compatible if $p_i = s_i \rightsquigarrow t$ and they are internally vertex-disjoint. We then focus on the more general problem of enumerating sets of compatible paths. Let $\mathcal{P}(S, t, G)$ be the set of all compatible paths for S and t in G . The same partition given in Eq. 4.1 is also valid for $\mathcal{P}(S, t, G)$. Namely, for any $s \in S$ such that $\delta^+(s) \neq \emptyset$,

$$\mathcal{P}(S, t, G) = \mathcal{P}(S, t, G') \bigcup_{v \in \delta^+(s)} (s, v) \mathcal{P}(S \setminus \{s\} \cup \{v\}, t, G - s), \quad (4.2)$$

where $G' = G - \{(s, v) | v \in \delta^+(s)\}$. Now, adding a new source to G with an arc to each vertex in S , we can use an augmenting path algorithm to test whether $\mathcal{P}(S, t, G) \neq \emptyset$ in $O(md)$ time. That way, an algorithm implementing the partition scheme of Eq. 4.2 can enumerate (s, t) - d -bubbles in $O(n^2md)$ delay, where the bound on the delay holds since each node of the recursion tree costs $O(nmd)$ (at most n emptiness checks are performed) and the height of the tree is bounded by n .

Theorem 4.14. *Given a directed graph G and two vertices s, t , the (s, t) - d -bubbles in G can be enumerated in $O(n^2md)$ delay.*

4.2 Listing bounded length paths

4.2.1 Introduction

A natural generalization of the problem of listing st -paths in undirected graphs (Section 3.2) is obtained by imposing a length constraint for the paths, that is, listing only the st -paths such that the length is bounded by some constant. The problem of listing st -paths in a *weighted* directed graph with lengths bounded by a constant is a further generalization of that problem. In this section, we consider this more general problem along with restrictions to undirected and unweighted graphs.

The shortest path problem is probably one of the most studied ones in computer science with a huge number of applications; it would be infeasible to list any reasonable subset of them here. A natural generalization of it, falling into the enumeration context, is the K -shortest paths problem, that consists in returning the first K distinct shortest st -paths, where both the graph and the parameter K are part of the input. There are several applications for this problem ranging from finding suboptimal solutions in sequence alignment problems (Waterman (1983); Byers and Waterman (1984)), to heuristics to solve NP-hard multi-criteria path optimization problems (Abdel-Lateef (1988); El-Amin and Al-Ghamdi (1993)). See Eppstein (1999) for further references.

The K -shortest paths problem has been studied since the early 1960s (see the references in Dreyfus (1969)). However, the first efficient algorithm for this problem in directed graphs with non-negative weights only appeared 10 years later, in the early 1970s, by Yen (1971) and Lawler (1972). With Dijkstra's algorithm implemented with Fibonacci heaps (Cormen et al. (2001)), their algorithm runs in $O(K(mn + n^2 \log n))$ time, where m, n are the number of arcs and vertices, respectively. More recently, Eppstein (1999) showed that if the paths can have cycles, i.e. they are walks, then the problem can be solved in $O(K + m + n \log n)$ time. When the input graph is undirected, the K -shortest *simple* paths problem is solvable in $O(K(m + n \log n))$ time (Katoh et al. (1982)). For directed unweighted graphs, the best known algorithm for the problem is the $O(Km\sqrt{n})$ time⁶ randomized algorithm of Roditty and Zwick (2005). In a different direction, Roditty (2007) noticed that the K -shortest simple paths can be efficiently approximated. Building upon his work, Bernstein (2010) presented an $O(Km/\epsilon)$ time⁷ algorithm for a $(1 + \epsilon)$ -approximation. When the paths are to be computed exactly, however, the best running time is still the $O(K(mn + n^2 \log n))$ time of Yen and Lawler's algorithm for directed graphs and the $O(K(m + n \log n))$ time of Katoh's algorithm for undirected graphs. Both algorithms use $O(Kn + m)$ memory.

The problems of listing st -paths with a length bounded by α on one hand and of the K -shortest st -paths on the other are closely related, even though the first problem cannot be solved in polynomial time under the standard definition, i.e. there can be an exponential number of bounded length st -paths. Intuitively, they are both the same problem with different parameterizations; in the first case the enumeration is constrained by the maximum length of the path and in the second by the maximum number of paths. Although very similar,

⁶Polylog factors are omitted.

⁷Polylog factors are omitted.

the problem of listing bounded length st -paths has not, to the best of our knowledge, been explicitly considered before, except for Eppstein (1999) who mentions that his algorithm can be modified to the bounded length case maintaining the same time and space complexity. Actually, Yen and Lawler's algorithm can be modified to solve the bounded length st -path problem, but in this case the memory used by the algorithm is the same as in the original version, i.e. proportional to the number of bounded length st -paths output, which is potentially exponential in the size of the graph. We show here that it is possible to list bounded length st -paths using space that is only linear in the size of the graph.

In the remainder of the chapter, we consider the problem of listing all st -paths with length bounded by α in a graph G with n vertices and m edges/arcs. We give a general $O(nt(n, m))$ delay algorithm, where $t(n, m)$ is the cost for a single source shortest paths computation, to list them in weighted (including negative values) directed graphs (Section 4.2.3) using $O(m + n)$ space. Next, we improve the total complexity of this algorithm to $O((m + t(n, m))\gamma)$, where γ is the number of paths output, for undirected graphs with non-negative weights (Section 4.2.4) while maintaining the same memory complexity. Finally, we modify the general algorithm to output the paths in increasing order of their lengths (Section 4.2.3). This algorithm can be used to solve the K -shortest paths problem.

4.2.2 Preliminaries

Given a weighted (directed or undirected) graph G with weights $w : E \mapsto \mathbb{Q}$, we say that a path p is α -bounded if the weight, or *length*, of the path satisfies $w(p) \leq \alpha$ and $\alpha \in \mathbb{Q}$, in the particular case of unitary weights (i.e. unweighted graphs), we say that p is k -bounded if $w(p) \leq k$ and $k \in \mathbb{Z}_{\geq 0}$. The general problem, formally defined below, with which we are concerned in this section is listing α -bounded st -paths in G .

Problem 4.15 (Listing α -bounded st -paths). *Given a weighted directed graph $G = (V, E)$, two vertices $s, t \in V$, and an upper bound $\alpha \in \mathbb{Q}$, output all α -bounded st -paths.*

The general problem is stated in terms of *directed* weighted graph, because any solution for directed graphs also applies to the undirected graphs, and in fact in Section 4.2.3 we only provide a solution to the directed case. Moreover, whenever G contains negative weight arcs, we assume that G does not contain any negative cycle, otherwise the shortest paths cannot be efficiently computed (Cormen et al. (2001)). Finally, we assume that all directed graphs considered here are weakly connected and all undirected graphs are connected, that way $m \geq n$, where n is the number of vertices and m the number of arcs (edges).

4.2.3 A simple polynomial delay algorithm

In this section, we present a simple polynomial delay algorithm to list all st -paths with length bounded by α in a weighted directed graph G . This is the most general version of the problem. Consequently, the algorithm works for any version of the problem, weighted (including negative weights) or unweighted, directed or undirected. However, the complexity is different for each version of the problem. The algorithm, inspired by the binary partition method, recursively partitions the solution space at every call until the considered subspace is a singleton (contains only one solution) and in that case outputs the corresponding solution. It is important to stress that the order in which the solutions are output is fixed, but arbitrary. The pseudocode is given in Algorithm 4.2.

Let us describe the partition scheme. Let $\mathcal{P}_\alpha(s, t, G)$ be the set of all paths from s to t in G . Assuming $s \neq t$, we have that

$$\mathcal{P}_\alpha(s, t, G) = \bigcup_{v \in N^+(s)} (s, v) \mathcal{P}_{\alpha'}(v, t, G - s), \quad (4.3)$$

where $\alpha' = \alpha - w(s, v)$. In other words, the set of paths from s to t can be partitioned into the union of $(s, v) \mathcal{P}_\alpha(v, t, G - s)$, the sets of paths containing the edge (s, v) , for each $v \in N^+(s)$. Indeed, since $s \neq t$, every path in $\mathcal{P}_\alpha(s, t, G)$ necessarily contains an edge (s, v) , where $v \in N^+(s)$.

Algorithm 4.2 implements this recursive partition strategy. The solutions are only output in the leaves of the recursion tree (line 2), where the partition is always a singleton. Moreover, in order to guarantee that every leaf in the recursion tree outputs one solution, we have to test if $\mathcal{P}_{\alpha'}(v, t, G - u)$, where $\alpha' = \alpha - w(u, v)$, is not empty before the recursive call (line 6). This set is not empty if and only if the weight of the shortest path from v to t in $G - u$ is at most α' , i.e. $d_{G-u}(v, t) \leq \alpha' = \alpha - w(u, v)$. Hence, to perform this test it is enough to compute all the distances from t in the graph $G^R - u$, where G^R is the graph G with all arcs reversed.

Algorithm 4.2: `list_paths`($u, t, \alpha, \pi_{su}, G$)

```

1 if  $u = t$  then
2   output( $\pi_{su}$ )
3   return
4 compute the distances from  $t$  in  $G^R - u$ 
5 for  $v \in N^+(u)$  do
6   if  $d(v, t) \leq \alpha - w(u, v)$  then
7     list_paths( $v, t, \alpha - w(u, v), \pi_{su}(u, v), G - u$ )

```

The correctness of Algorithm 4.2 follows directly from the relation given in Eq. 4.3 and the correctness of the tests of line 6. We can perform those tests in $O(1)$ by pre-computing the distances from t to all vertices (single source shortest paths) in the reverse graph $G^R - u$, which can be computed in $O(t(n, m))$. The height of the recursion tree is bounded by n , since at every level of the recursion tree a new vertex is added to the current solution and any solution has at most n vertices. In that way, the path between any two leaves in the recursion tree has at most $2n$ nodes. Thus, the time elapsed between two solutions being output is $O(nt(n, m))$. Moreover, the space complexity of the algorithm is $O(m)$, since for each recursive call, we can store the difference with the previous graph.

Theorem 4.16. *Algorithm 4.2 has delay $O(nt(n, m))$, where $t(n, m)$ is the cost to compute a shortest path tree, and uses $O(m)$ space.*

For unweighted (directed and undirected) graphs, the single source shortest paths can be computed using breadth-first search (BFS) running in $O(m)$, so Theorem 4.16 guarantees an $O(km)$ delay to list all k -paths, since the height of the recursion tree is bounded by k instead of n . In the case of non-negative weights the single source shortest paths can be computed using Dijkstra's algorithm in $O(m + n \log n)$, resulting in an $O(nm + n^2 \log n)$ delay. Finally, for general weights, the single source shortest paths can be computed using the Bellman-Ford algorithm in $O(mn)$ time, resulting in an $O(mn^2)$ delay.

4.2.4 An improved algorithm for undirected graphs

In this section, we improve the total time complexity of Algorithm 4.2 from $O(nt(n, m)\gamma)$ to $O((m + t(n, m))\gamma)$ in the case of *non-negatively* weighted undirected graphs, where $\gamma = |\mathcal{P}_\alpha(s, t, G)|$ is the number of α -bounded st -paths. In other words, for undirected graphs we can list all α -bounded st -paths in $O((m + n \log n)\gamma)$ and all k -bounded st -paths in $O(m\gamma)$. However, the delay of the algorithm is still $O(nt(n, m))$ in the worst case, although the (worst case) average delay is $O(m + t(n, m))$. From now on, all the graphs considered are undirected unless otherwise stated.

The basis to improve the complexity of Algorithm 4.2 is to explore the structure of $\mathcal{P}_\alpha(s, t, G)$ to reduce the number of nodes in the recursion tree. More precisely, at every call, we identify the longest common prefix of $\mathcal{P}_\alpha(s, t, G)$, i.e. the longest (considering the number of edges) path $\pi_{ss'}$ such that $\mathcal{P}_\alpha(s, t, G) = \pi_{ss'}\mathcal{P}_\alpha(s', t, G)$, and append it to the current path prefix being considered in the recursive call. The pseudocode for this algorithm is very similar to Algorithm 4.2 and, for the sake of completeness, is given in Algorithm 4.3. We postpone the description of the $\text{lcp}(u, t, \alpha, G)$ function to the next section, along with a discussion about the difficulties to extend it to directed graphs or general weights graphs.

Algorithm 4.3: $\text{list_paths}(u, t, \alpha, \pi_{su}, G)$

```

1  $\pi_{uu'} = \text{lcp}(u, t, \alpha, G)$ 
2 if  $u' = t$  then
3   |   output( $\pi_{su}\pi_{uu'}$ )
4   |   return
5 else
6   |   compute a shortest path tree  $T'_t$  from  $t$  in  $G^R - \pi_{uu'}$ 
7   |   for  $v \in N(u')$  do
8   |   |   if  $d(v, t) + w(u, v) \leq \alpha$  then
9   |   |   |    $\text{list\_paths}(v, t, \alpha - w(\pi_{uu'}) - w(u', v), \pi_{su}\pi_{uu'}(u', v), G - \pi_{uu'})$ 

```

The correctness of Algorithm 4.3 follows directly from the correctness of Algorithm 4.2. The space used is the same of Algorithm 4.2, provided that $\text{lcp}(u, t, \alpha, G)$ uses linear space, which, as we show in the next section, is indeed the case (Theorem 4.20).

Let us now analyze the total complexity of Algorithm 4.3 as a function of the input size and of γ , the number of α -bounded (k -bounded) st -paths. Let R be the recursion tree of Algorithm 4.3 and $T(r)$ the cost of a given node $r \in R$. The total cost of the algorithm can be split in two parts, which we later bound individually, in the following way:

$$\sum_{r \in R} T(r) = \sum_{r: \text{internal}} T(r) + \sum_{r: \text{leaf}} T(r). \quad (4.4)$$

We have that $\sum_{r: \text{leaf}} T(r) = O((m + t(m, n))\gamma)$, since leaves and solutions are in one-to-one correspondence and the cost for each leaf is dominated by the cost of $\text{lcp}(u, t, \alpha, G)$, that is $O(m + t(m, n))$ (Theorem 4.20). Now, we have that every internal node of the recursion has at least two children, otherwise $\pi_{uu'}$ would not be the longest common prefix of $\mathcal{P}_\alpha(u, t, G)$. Thus, $\sum_{r: \text{internal}} T(r) = O((m + t(m, n))\gamma)$ since each internal node costs $O(m + t(m, n))$, the cost is also dominated by the cost of the longest prefix computation, and in any tree the number of branching nodes is at most the number of leaves. Therefore, the total complexity of Algorithm 4.3 is $O((m + t(n, m))\gamma)$. This completes the proof of Theorem 4.17.

Theorem 4.17. *Algorithm 4.3 outputs all α -bounded (or k -bounded) st -paths in $O((m + t(n, m))\gamma)$ using $O(m)$ space.*

This means that for unweighted graphs it is possible to list all k -bounded st -paths in $O(m)$ per path. Moreover, for non-negatively weighted graphs, it is possible to list all α -bounded st -paths in $O(m + n \log n)$ per path.

Computing the longest common prefix of $\mathcal{P}_\alpha(s, t, G)$

In this section, we present an efficient algorithm to compute the longest common prefix of the set of α -paths from s to t , completing the description of Algorithm 4.3. The naive algorithm for this problem runs in $O(nt(n, m))$, so that using it in Algorithm 4.3 would not improve the total complexity compared to Algorithm 4.2. Basically, the naive algorithm computes a shortest path π_{st} and then for each prefix in increasing order of length tests if there are at least two distinct extensions each with total weight less than α . In order to test the extensions, for each prefix π_{su} , we recompute the distances from t in the graph $G - \pi_{su}$, thus performing n shortest path tree computations (k computations in the unweighted case) in the worst case.

Algorithm 4.4 improves the naive algorithm by avoiding those recomputations. However, before entering the description of Algorithm 4.4, we need a better characterization of the structure of the longest common prefix of $\mathcal{P}_\alpha(s, t, G)$. Lemma 4.18 gives this. It does so by considering a shortest path tree rooted at s , denoted by T_s . Recall that T_s is a subgraph of G and induces a partition of the edges of G into tree edges and non-tree edges. In this tree, the longest common prefix of $\mathcal{P}_\alpha(s, t, G)$ is a prefix of the tree path from the root s to t . Additionally, any st -path in G , excluding the tree path, necessarily passes through at least one non-tree edge. The lemma characterizes the longest common prefix in terms of the non-tree edges from the subtrees rooted at siblings of the vertices in the tree path from s to t .

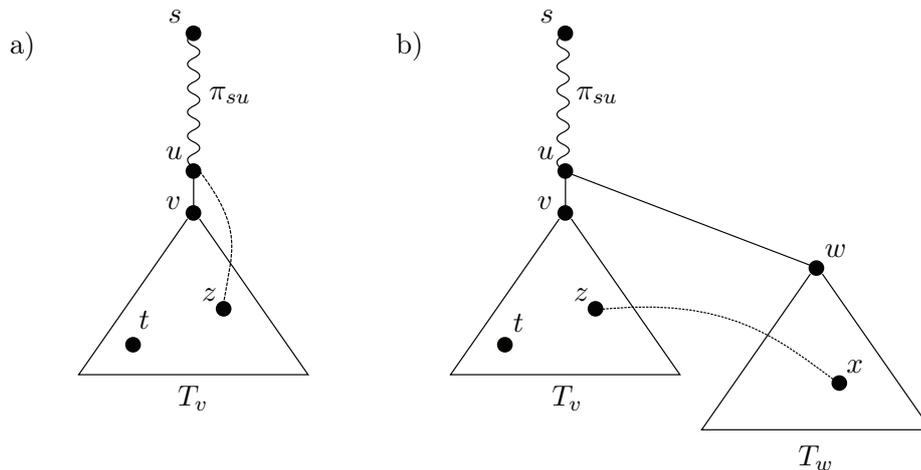


Figure 4.5: The common prefix π_{su} of $\mathcal{P}_\alpha(s, t, G)$ can always be extended into an st -path using the tree path of T_s from u to t . The path π_{su} is the longest common prefix if and only if it can also be extended with a path containing a non-tree edge (x, z) such that $z \in T_v$ and (a) $x = u$ or (b) $x \in T_w$ and w is sibling of v ; and $d_{G'}(s, x) + w(x, z) + d_{G'}(z, t) \leq \alpha$, where $G' = G - (u, v)$.

Lemma 4.18. *Let $\pi_{su} = (s = v_0, v_1), \dots, (v_{l-1}, v_l = u)$ be a common prefix of all paths in $\mathcal{P}_\alpha(s, t, G) \neq \emptyset$ and T_s a shortest path tree rooted at s . Then,*

1. *the path $\pi_{su}(u, v)$ is a common prefix of $\mathcal{P}_\alpha(s, t, G)$, if there is no edge (x, z) , with $z \in T_v$ and $x = u$ or $x \in T_w$ where w is a sibling of v in the tree T_s , such that $d_{G'}(s, x) + w(x, z) + d_{G'}(z, t) \leq \alpha$, where $G' = G - (u, v)$; (see Fig. 4.5)*
2. *π_{su} is the longest common prefix of $\mathcal{P}_\alpha(s, t, G)$, otherwise.*

Proof. Let us prove that if there exists a path π_{ut} , not containing the tree edge (u, v) , extending π_{su} such that $w(\pi_{su}\pi_{ut}) \leq \alpha$, then there is a non-tree edge (x, z) such that $d_{G'}(s, x) + w(x, z) + d_{G'}(z, t) \leq \alpha$ and $z \in T_v$. For the moment, we do not impose that $x \in T_w$ or $x = u$, we deal with this condition later. The paths of $\mathcal{P}_\alpha(s, t, G)$ that do not pass through (u, v) necessarily use some non-tree edge (x, z) , where $z \in T_v$, since t belongs to T_v . Now, consider the path $\pi_{su}\pi_{ut}$ and let (x, z) be the last non-tree edge of this path that enters T_v . This path can be rewritten as $\pi_{sx}(x, z)\pi_{zt}$. We have that $w(\pi_{sx}(x, z)\pi_{zt}) = w(\pi_{sx}) + w(x, z) + w(\pi_{zt})$, where the path π_{zt} is entirely contained in the induced subgraph of the vertices of T_v , because of our choice of (x, z) . Thus, $w(\pi_{zt}) \geq d_{G'}(z, t)$. Moreover, the path π_{sx} does not contain (u, v) , since $\pi_{su}\pi_{ut}$ is a simple path and (u, v) is not the first edge of π_{ut} . Thus, $w(\pi_{sx}) \geq d_{G'}(s, x)$. Therefore, combining the two inequalities, we have that $d_{G'}(s, x) + w(x, z) + d_{G'}(z, t) \leq w(\pi_{sx}) + w(x, z) + w(\pi_{zt}) \leq \alpha$ and (x, z) is a non-tree edge with $z \in T_v$.

It remains to prove that it is sufficient to consider only the non-tree edges (x, z) entering T_v , such that $x = u$ or $x \in T_w$ where w is sibling of v in the tree. Let π_{st} be a path in $\mathcal{P}_\alpha(s, t, G)$ including (x, z) that does not have $\pi_{su}(u, v)$ as a prefix. This path can be rewritten as $\pi_{st} = \pi_{sx}(x, z)\pi_{zt}$. Since π_{su} is a common prefix of $\mathcal{P}_\alpha(s, t, G)$, we have that π_{su} is a prefix of $\pi_{sx}(x, z)\pi_{zt}$. Thus, $\pi_{sx} = \pi_{su}\pi_{ux}$, and either π_{ux} is empty, so (x, z) is a non-tree edge from u , or π_{ux} enters a sibling subtree of v . This completes the proof of the first part of the lemma.

Let us prove the second part of the lemma. There is at least one non-tree edge (x, z) entering T_v from u or T_w , a sibling of v , such that $d_{G'}(s, x) + w(x, z) + d_{G'}(z, t) \leq \alpha$. Thus, the concatenation $\pi_{sx}(x, z)\pi_{zt}$ of the shortest paths contains $\pi_{su}(u, y)$ as prefix, where y is a neighbor of u . Moreover, there is a subpath π_{st}^* of $\pi_{sx}(x, z)\pi_{zt}$ that is simple and $w(\pi_{st}^*) \leq \alpha$, which also has $\pi_{su}(u, y)$ as prefix. Therefore, π_{su} has two possible extensions, using (u, y) or the tree edge (u, v) . \square

In order to use the characterization of Lemma 4.18 for the longest prefix of $\mathcal{P}_\alpha(s, t, G)$, we need to be able to efficiently test for the weight condition given in item 1, namely $d_{G'}(s, x) + w(x, z) + d_{G'}(z, t) \leq \alpha$, where $G' = G - (u, v)$ and (u, v) belongs to the tree path from s to t . We have that $d_{G'}(s, x) = d_G(s, x)$, since x does not belong to the subtree of v in the shortest path tree T_s . Indeed, only the distances of vertices in the subtree T_v can possibly change after the removal of the tree edge (u, v) . On the other hand, in principle we have no guarantee that $d_{G'}(z, t)$ also remains unchanged: recall that to maintain the distances from t we need a tree rooted at t not at s . Clearly, we cannot compute the shortest path tree from t for each G' , in the worst case, this would imply the computation of n shortest path trees. For this reason, we need Lemma 4.19. It states that, in the specific case of the vertices z we need to compute the distance to t in G' , we have that $d_{G'}(z, t) = d_G(z, t)$. A similar result was proved in [Hershberger and Suri \(2001\)](#).

Lemma 4.19. *Let T_s be a shortest path tree rooted at s and t a vertex of G . Then, for any edge (u, v) , with v closer to t , in the shortest path π_{st} in the tree T_s , we have that $d_G(z, t) = d_{G'}(z, t)$, where $z \in T_v$ and $G' = G - (u, v)$.*

Proof. Suppose that (u, v) belongs to the shortest path π'_{zt} in G . This path can be written as the concatenation $\pi'_{zv}\pi'_{vt}$ (assuming wlog v closer than u to t), where π'_{zv} and π'_{vt} are both simple paths. We also have that t and z belong to the subtree T_v which does not include the edge (u, v) , so the paths π_{vt} and π_{vz} in the tree T_v are shortest paths that do not include (u, v) . The concatenation of π_{vt} and π_{vz} contains a subpath π_{zt} from z to t such that $w(\pi_{zt}) \leq w(\pi_{vt}) + w(\pi_{vz})$. On the other hand, $w(\pi_{vt}) + w(\pi_{vz}) \leq w(\pi'_{zv}) + w(\pi'_{vt})$, since π_{vt} and π_{vz} are both shortest paths. Thus, $w(\pi_{zt}) \leq w(\pi'_{zv}) + w(\pi'_{vt})$. Therefore, the concatenation of π_{vt} and π_{vz} contains as a subpath a shortest path from z to t that does not include (u, v) . \square

It is not hard to verify that Lemma 4.18 is also valid for directed graphs. Indeed, in the proof above, the fact that G is undirected is not used. On the other hand, the non-negative hypothesis for the weights is necessary; more specifically, we need the monotonicity property for path weights which states that for any path the weight of any sub-path is not greater than the weight of the full path. Now, in Lemma 4.19 both the path monotonicity property and the fact that the graph is undirected are necessary. Since these two lemmas are the base for the efficiency of Algorithm 4.4, it seems difficult to extend it to general weights and/or directed graphs.

Algorithm 4.4 implements the strategy suggested by Lemma 4.18. Given a shortest path tree T_s of G rooted at s , the algorithm traverses each vertex v_i in the tree path $s = v_0, \dots, v_n = t$ from the root s to t , and at every step finds all non-tree edges (x, z) entering the subtree rooted at v_{i+1} from a sibling subtree, i.e. a subtree rooted at $w \in N^+(v_i) \setminus \{v_{i+1}\}$. For each non-tree (x, z) linking the sibling subtrees found, it checks if it satisfies the weight condition $d_{G'}(s, x) + w(x, z) + d_{G'}(z, t) \leq \alpha$, where $G' = G \setminus (v_i, v_{i+1})$, given in item 1 of Lemma 4.18. Item 2 of the same lemma implies that the first time an edge (x, z) satisfies the weight condition, the tree path traversed so far is the longest common prefix of $\mathcal{P}_\alpha(s, t, G)$. In order to test the weight conditions, as stated previously, we have that $d_{G'}(s, x) = d_G(s, x)$, since x does not belong to the subtree of v in T_s . In addition, Lemma 4.19 guarantees that $d_{G'}(z, t) = d_G(z, t)$. Thus, it is sufficient for the algorithm to compute only the shortest path trees from t and from s in G .

Algorithm 4.4: $\text{lcp}(s, t, \alpha, G)$

```

1 compute  $T_s$ , a shortest path tree from  $s$  in  $G$ 
2 compute  $T_t$ , a shortest path tree from  $t$  in  $G$ 
3 let  $\pi_{st} = (s = v_0, v_1) \dots (v_{n-1}, v_n = t)$  be the shortest path in  $T_s$ 
4 for  $v_i \in \{v_1, \dots, v_n\}$  do
5   for  $w \in N^+(v_i) \setminus \{v_{i+1}\}$  do
6     let  $T_w$  be the subtree of  $T_s$  rooted at  $w$ 
7     for  $(x, z) \in G$ , s.t.  $x \in T_w$  or  $x = v_i$  do
8       if  $z \in T_{v_{i+1}}$  and  $d_G(s, x) + w(x, z) + d_G(z, t) \leq \alpha$  then
9         break
10 return  $\pi_{sv_{i-1}}$ 

```

Theorem 4.20. *Algorithm 4.4 finds the longest common prefix of $\mathcal{P}_\alpha(s, t, G)$ in $O(m + t(n, m))$ using $O(m)$ space.*

Proof. The cost of the algorithm can be divided in two parts: the cost to compute the shortest path trees T_s and T_t , and the cost of the loop in line 4. The first part is bounded by $O(t(n, m))$. Let us now prove that the second part is bounded by $O(m + n)$. The cost of each execution of line 8 is $O(1)$, since we only need distances from s and t and the shortest path trees from s and t are already computed, and we pre-process the tree to decide in $O(1)$ if a vertex belongs to a subtree. In that way, the cost of the loop is bounded by the number of times line 8 is executed. Line 8 is executed at most m times, the neighborhood of each vertex is visited at most once, since the subtrees T_w are disjoint, they are rooted at vertices adjacent to some vertex in the tree path π_{su} but not included in it. \square

4.2.5 Listing paths in increasing order of their lengths

In this section, we modify Algorithm 4.2 to output the α -bounded st -paths in increasing order of their length, while maintaining (almost) the same time complexity but increasing the memory usage. As for Algorithm 4.2, this algorithm works for any version of the problem, directed or undirected graphs with general weights, and the complexity depends on the cost to compute a shortest path tree. The pseudocode is shown in Algorithm 4.5. This is a generic description of the algorithm, the container Q is not specified in the pseudocode, the only requirement is the support for two operations: *push*, to insert a new element in Q ; and *pop*, to remove and return an element of Q .

Algorithm 4.5: `list_paths_iterative($u, t, \alpha, \pi_{su}, G$)`

```

1 push  $\langle s, t, \emptyset, G \rangle$  in  $Q$ 
2 while  $Q$  is not empty do
3    $\langle u, t, \pi_{su}, G \rangle = Q.pop()$ 
4   if  $u = t$  then
5     output( $\pi_{su}$ )
6   else
7     compute a shortest path tree  $T_t$  from  $t$  in  $G^R - u$ 
8     for  $v \in N^+(u)$  do
9       if  $d(v, t) \leq \alpha - w(u, v)$  then
10        push  $\langle v, t, \alpha - w(u, v), \pi_{su}(u, v), G - u \rangle$  in  $Q$ 

```

Algorithm 4.5 is a non-recursive version of Algorithm 4.2, and uses the same strategy to partition the solution space (Eq. 4.3). However, the order in which the partitions are explored is not necessarily the same, depending on the type of container used for Q . We show that if Q is a stack then the solutions are output in the reverse order of Algorithm 4.2, and the maximum size of the stack is linear in the size of the input. If on the other hand, Q is a heap, using a suitable key, the solutions are output in increasing order of their lengths, but in this case the maximum size of the heap is linear in the number of solutions, which is not polynomial in the size of the input.

The recursive partition of $\mathcal{P}_\alpha(s, t, G)$, i.e. the set of α -bounded st -paths in $G = (V, E)$, according to Eq. 4.3 has a rooted tree structure. Indeed, the nodes are the sets $\mathcal{P}_{\alpha'}(v, t, G')$, where $G' = (V', E')$ is a subgraph of G , $\alpha' \in \mathbb{Q}$, and $v \in V'$; for a given node the children are

the sets in the partition of Eq. 4.3 satisfying the condition of line 9, i.e. the non-empty sets; the root is $\mathcal{P}_\alpha(s, t, G)$; and the leaves are the singletons $\mathcal{P}_{\alpha'}(t, t, G')$, which are in a one-to-one correspondence with the α -bounded st -paths. We denote this rooted tree by \mathcal{T} .

For any container Q supporting push and pop operations, Algorithm 4.5 visits each node of \mathcal{T} exactly once, since at every iteration a node from Q is deleted and its children are inserted in Q , and \mathcal{T} is a tree. In particular, this guarantees that every leaf of \mathcal{T} is visited exactly once, thus proving the following lemma.

Lemma 4.21. *Algorithm 4.5 outputs all α -bounded st -paths.*

Let us consider the case where Q is a stack. It is not hard to prove that Algorithm 4.2 is a DFS traversal of \mathcal{T} starting from the root, while Algorithm 4.5 is an *iterative* DFS (Sedgewick (2001)) traversal of \mathcal{T} also starting from the root. Basically, an iterative DFS keeps the vertices of the fringe of the non-visited subgraph in a stack, at each iteration the next vertex to be explored is popped from the stack, and recursive calls are replaced by pushing vertices in the stack. Now, for a fixed permutation of the children of each node in \mathcal{T} , the nodes visited in an iterative DFS traversal are in the reverse order of the nodes visited in a recursive DFS traversal (Sedgewick (2001)), thus proving Lemma 4.22.

Lemma 4.22. *If Q is a stack, then Algorithm 4.5 outputs the α -bounded st -path in the reverse order of Algorithm 4.2.*

For any rooted tree, at any moment during an iterative DFS traversal, the number of nodes in the stack is bounded by the sum of the degrees of the root-to-leaf path currently being explored. Recall that every leaf in \mathcal{T} corresponds to a path in $\mathcal{P}_\alpha(s, t, G)$. Actually, there is a one-to-one correspondence between the nodes of a root-to-leaf path P in \mathcal{T} and the vertices of the α -bounded st -path π associated to that leaf. Hence, the sum of the degrees of the nodes of P in \mathcal{T} is equal to the sum of the degrees of the vertices π in G , which is bounded by m , thus proving Lemma 4.23.

Lemma 4.23. *The maximum number of elements in the stack of Algorithm 4.5 over all iterations is bounded by m .*

Let us consider now the case where Q is a heap. There is a one-to-many correspondence between arcs in G and arcs in \mathcal{T} , i.e. if $\mathcal{P}_{\alpha'}(v, t, G'')$ is a child of $\mathcal{P}_{\alpha'}(u, t, G')$ in \mathcal{T} then (u, v) is an arc of G . For every arc of \mathcal{T} let us associate the weight of the corresponding arc in G . Intuitively, Algorithm 4.5 using a priority queue with $w(\pi_{su}) + d_G(u, t)$ as keys performs a Dijkstra-like traversal in a weighted version of \mathcal{T} starting from the root, where for a node $\langle u, t, \pi_{su}, G \rangle$ the distance from the root is $w(\pi_{su})$ and $d_G(u, t)$ is a (precise) estimation of the distance from $\langle u, t, \pi_{su}, G \rangle$ to the closest leaf of \mathcal{T} . In other words, it is an A^* -like traversal (Dechter and Pearl (1985)) in the weighted rooted tree \mathcal{T} , using the (optimal) heuristic $d_G(u, t)$. As such, Algorithm 4.5 explores first the nodes of \mathcal{T} leading to the cheapest non-visited leaf. This is formally stated in Lemma 4.24.

Lemma 4.24. *If Q is a priority queue with $w(\pi_{su}) + d_{G'}(u, t)$ as the priority key of $\langle u, t, \pi_{su}, G' \rangle$, then Algorithm 4.5 outputs the α -bounded st -paths in increasing order of their lengths.*

Proof. The priority of a node $N_u = \langle u, t, \pi_{su}, G' \rangle$ (i.e. $\mathcal{P}_{\alpha-w(\pi_{su})}(u, t, G')$) is the weight of the path π_{su} plus the weight of a shortest path π_{ut}^* from u to t in G' . Let v be an out-neighbor of u in G' , then the node $N_v = \langle v, t, \pi_{su}(u, v), G' - u \rangle$ is a child of N_u , and the priority of N_v is greater or equal to the priority of N_u . Indeed, suppose it is strictly smaller, then the path

$\pi_{su}(u, v)$ concatenated with the shortest path from v to t in $G' - u$ is shorter than $\pi_{su}\pi_{ut}^*$, contradicting the fact that π_{ut}^* is a shortest path of G' . Hence, the priorities of the nodes removed from Q are not decreasing, since for every node removed only nodes with greater or equal priorities are inserted. Moreover, the priority of a leaf $\langle t, t, \pi_{st}, G' \rangle$ is precisely $w(\pi_{st})$, the weight of a path in $\mathcal{P}_\alpha(s, t, G)$. Therefore, the leaves are visited in increasing order of the length of their corresponding st -path. \square

For any choice of the container Q , every node of \mathcal{T} is visited exactly once, that is, each node of \mathcal{T} is pushed at most once in Q . This proves Lemma 4.25.

Lemma 4.25. *The maximum number of elements in a priority queue of Algorithm 4.5 over all iterations is bounded by γ .*

Algorithm 4.5 uses $O(m\gamma)$ space, since for every node inserted in the heap, we also have to store the corresponding graph. Moreover, using a binary heap (Cormen et al. (2001)) as a priority queue, the push and pop operations can be performed in $O(\log \gamma)$ each, where by Lemma 4.25 γ is the maximum size of the heap. Therefore, combining it with Lemma 4.24 we obtain the following theorem.

Theorem 4.26. *Algorithm 4.5 using a heap outputs all α -bounded st -paths in increasing order of their lengths in $O((nt(n, m) + \log \gamma)\gamma)$ total time, using $O(m\gamma)$ space.*

4.3 Discussion and conclusions

In the first part of this chapter, we introduced a polynomial delay algorithm to list all bubbles with path length constraints in weighted directed graphs. This is a theoretically sound approach that in practice is considerably faster than the bubble listing algorithm of KISSPLICE (Section 2.2.2), and as a result enables us to enumerate more bubbles. Additionally, we gave an indication that these additional bubbles correspond to longer AS events, overseen previously but biologically very relevant. Moreover, as shown in Ahuja et al. (1990), by combining radix and Fibonacci heaps in Dijkstra, we can achieve a $O(n(m + n\sqrt{\log \alpha_1}))$ delay for Algorithm 4.1 in cDGBs. The current implementation of KISSPLICE (version 2.0) uses Algorithm 4.1 to list bubbles.

In the second part of this chapter, we introduced a general framework to list bounded length st -paths in weighted directed graphs. In the particular case of undirected graphs, we showed an improved algorithm to list bounded length st -paths in $O((m + n \log n)\gamma)$ time for non-negative weights and $O(m\gamma)$ time for unit weights, where γ is the number of bounded length paths. Moreover, we showed how to modify the general algorithm to output the paths in increasing order of their length, thus providing an alternative solution to the classical K -shortest paths problem, which does not improve the complexity but is simpler than previous approaches.

Actually, the general framework of Section 4.2 can be seen as a “simplification” of the bubble listing algorithm of Section 4.1 (extended to general weights). More precisely, listing bounded length st -paths can be reduced to listing bounded length bubbles with a given source s . Indeed, consider an instance of the first problem, a graph G and two vertices s, t , and build the graph G' by adding an arc (s, t) with weight α' , strictly smaller than the sum of all negative weight arcs (if any) of G ; listing st -paths with a length bounded by α in G is equivalent to listing (s, t, α', α) -bubbles in G' .

Chapter 5

Memory efficient de Bruijn graph representation

Contents

5.1	Introduction	96
5.2	Preliminaries	96
5.3	Cascading Bloom filter	97
5.4	Memory and time usage	98
5.4.1	Using different values of r for different filters	100
5.4.2	Query distribution among filters	100
5.5	Experimental results	102
5.5.1	Construction algorithm	102
5.5.2	Implementation and experimental setup	102
5.5.3	<i>E. coli</i> dataset, varying k	103
5.5.4	<i>E. coli</i> dataset, varying coverage	105
5.5.5	<i>E. coli</i> dataset, query statistics	105
5.5.6	Human dataset	106
5.6	Discussion and conclusions	106

This chapter is strongly based on our paper [Salikhov et al. \(2013\)](#). As shown in Chapter 2, the de Bruijn graph construction and representation are the memory bottleneck of KISSPLICE. In this chapter, we consider the problem of compactly representing a de Bruijn graph. We show how to reduce the memory required by the algorithm of [Chikhi and Rizk \(2012\)](#), that represents de Bruijn graphs using Bloom filters. Our method requires 30% to 40% less memory with respect to their method, with insignificant impact to construction time. At the same time, our experiments showed a better query time compared to their method. This is, to our knowledge, the best *practical* representation for de Bruijn graphs. The current implementation of KISSPLICE (version 2.0) uses the de Bruijn graph representation and construction presented in this chapter.

5.1 Introduction

As shown in Chapter 1, KISSPLICE is not the only NGS data analysis method using de Bruijn graphs. In fact, the majority of the more recent genome and transcriptome assemblers and some metagenome assemblers (Peng et al. (2011); Namiki et al. (2011)) use de Bruijn graphs. Due to the very large size of NGS datasets, it is essential to represent de Bruijn graphs as compactly as possible. This has been a very active line of research. Recently, several papers have been published that propose different approaches to compressing de Bruijn graphs Conway and Bromage (2011); Ye et al. (2012); Chikhi and Rizk (2012); Bowe et al. (2012); Pell et al. (2012).

Conway and Bromage (2011) proposed a method based on classical succinct data structures, i.e. bitmaps with efficient rank/select operations. On the same direction, Bowe et al. (2012) proposed a very interesting succinct representation that, assuming only one string (read) is present, uses only $4m$ bits, where m is the number of arcs in the graph. The more realistic case, where there are M reads, can be easily reduced to the one string case by concatenating all M reads using a special separator character. However, in this case the size of the structure is $4m + O(M \log m)$ bits (Bowe et al. (2012), Theorem 1). Since the multiplicative constant of the second term is hidden by the asymptotic notation, it is hard to know precisely what would be the size of this structure in practice.

Ye et al. (2012) proposed a different method based on a sparse representation of de Bruijn graphs, where only a subset of k -mers present in the dataset are stored. Pell et al. (2012) proposed a method to represent it approximately, the so called *probabilistic de Bruijn graph*. In their representation a vertex have a small probability to be a false positive, i.e. the k -mer is not present in the dataset. Finally, Chikhi and Rizk (2012) improved Pell's scheme in order to obtain an exact representation of the de Bruijn graph. This was, to our knowledge, the best *practical* representation of an exact de Bruijn graph.

In this chapter, we focus on the method proposed in Chikhi and Rizk (2012) which is based on Bloom filters. They were first used in Pell et al. (2012) to provide a very space-efficient representation of a subset of a given set (in our case, a subset of k -mers), at the price of allowing *one-sided errors*, namely *false positives*. The method of Chikhi and Rizk (2012) is based on the following idea: if all queried vertices (k -mers) are only those which are reachable from some vertex known to belong to the graph, then only a fraction of all false positives can actually occur. Storing these false positives explicitly leads to an exact (false positive free) and space-efficient representation of the de Bruijn graph.

Our contribution is an improvement of this scheme by changing the representation of the set of false positives. We achieve this by iteratively applying a Bloom filter to represent the set of false positives, then the set of "false false positives" etc. We show analytically that this cascade of Bloom filters allows for a considerable further economy of memory, improving the method of Chikhi and Rizk (2012). Depending on the value of k , our method requires 30% to 40% less memory with respect to the method of Chikhi and Rizk (2012). Moreover, with our method, the memory grows very little as k grows. Finally, we implemented our method and tested it against Chikhi and Rizk (2012) on real datasets. The tests confirm the theoretical predictions for the size of structure and show a 20% to 30% *improvement* in query times.

5.2 Preliminaries

A *Bloom filter* is a space-efficient data structure for representing a given subset of elements $T \subseteq U$, with support for efficient membership queries with one-sided error. That is, if a query

for an element $x \in U$ returns *no* then $x \notin T$, but if it returns *yes* then x may or not belong to T , i.e. with small probability $x \notin T$ (false positive). It consists of a bitmap (array of bits) B with size m and a set of p distinct hash functions $\{h_1, \dots, h_p\}$, where $h_i : U \mapsto \{0, \dots, m-1\}$. Initially, all bits of B are set to 0. An insertion of an element $x \in T$ is done by setting the elements of B with indices $h_1(x), \dots, h_p(x)$ to 1, i.e. $B[h_i(x)] = 1$ for all $i \in [1, p]$. The membership queries are done symmetrically, returning *yes* if all $B[h_i(x)]$ are equal 1 and *no* otherwise. As shown in Kirsch and Mitzenmacher (2008), when considering hash functions that yield equally likely positions in the bit array, and for large enough array size m and number of inserted elements n , the false positive rate \mathcal{F} is

$$\mathcal{F} \approx (1 - e^{-pn/m})^p = (1 - e^{-p/r})^p, \quad (5.1)$$

where $r = m/n$ is the number of bits (of the bitmap B) per element (of T represented). It is not hard to see that this expression is minimized when $p = r \ln 2$, giving a false positive rate of

$$\mathcal{F} \approx (1 - e^{-p/r})^p = (1/2)^p \approx 0.6185^r. \quad (5.2)$$

A de Bruijn graph, as defined in Chapter 1 (Definition 1.4), is entirely determined by the set of k -mers (vertices) and $(k+1)$ -mers (arcs) of the read set $\mathcal{R} \subseteq \Sigma^* = \{A, C, T, G\}^*$. For reasons that will be clear soon, we relax this definition, dropping the bijection between that $(k+1)$ -mers and arcs but keeping the $k-1$ suffix-prefix overlap requirement. That way, a de Bruijn graph, for a given parameter k , of a set of reads \mathcal{R} is entirely defined by the set $T \subseteq U = \Sigma^k$ of k -mers present in \mathcal{R} . Indeed, the vertices of the graph are precisely the k -mers of T and for any two vertices $u, v \in T$, there is an arc from u to v if and only if the suffix of u of size $k-1$ is equal to the prefix of v of the same size. Therefore, given a set $T \subseteq U$ of k -mers we can represent its de Bruijn graph using a Bloom filter B . This representation has the disadvantage of having false positive vertices, as direct consequence of the false positive queries in the Bloom filter, which can create false connections in the graph (see Pell et al. (2012) for the influence of false positive vertices on the topology of the graph). The naive way to remove those false positives vertices, by explicitly storing (e.g. using a hash table) the set of all false positives of B , is clearly inefficient, as the expected number of elements to be explicitly stored is $|U|\mathcal{F} = 4^k \mathcal{F}$.

The key idea of Chikhi and Rizk (2012) is to explicitly store only a subset of all false positives of B , the so-called *critical false positives*. This is possible because in order to perform an exact (without false positive vertices) graph traversal, only potential neighbors of vertices in T are queried. In other words, the set of critical false positives consists of the potential neighbors of T that are false positives of B , i.e. the k -mers from U that overlap the k -mers from T by $k-1$ letters and are false positives of B . Thus, the size of the set of critical false positives is bounded by $8|T|$, since each vertex of T has at most $2|\Sigma| = 8$ neighbors (for each vertex, there are $|\Sigma|$ k -mers overlapping the $k-1$ suffix and $|\Sigma|$ overlapping the $k-1$ prefix). Therefore, the expected number of critical false positives is bounded above by $8|T|\mathcal{F}$.

5.3 Cascading Bloom filter

Let \mathcal{R} be a set of reads and T_0 be the set of occurring k -mers (vertices of the de Bruijn graph) that we want to store. As stated in Section 5.2, the method of Chikhi and Rizk (2012) stores T_0 via a bitmap B_1 using a Bloom filter, together with the set T_1 of critical false positives. T_1 consists of those k -mers which have a $k-1$ overlap with k -mers from T_0 but which are stored

in B_1 “by mistake”, i.e. belong¹ to B_1 but not to T_0 . B_1 and T_1 are sufficient to represent the graph provided that the only queried k -mers are those which are potential neighbors of k -mers of T_0 .

The idea we introduce here is to use this structure recursively and represent the set T_1 by a new bitmap B_2 and a new set T_2 , then represent T_2 by B_3 and T_3 , and so on. More formally, starting from B_1 and T_1 defined as above, we define a series of bitmaps B_1, B_2, \dots and a series of sets T_1, T_2, \dots as follows. B_2 stores the set of false positives T_1 using another Bloom filter, and the set T_2 contains the critical false positives of B_2 , i.e. “true vertices” from T_0 that are stored in B_2 “by mistake” (we call them **false**² positives). B_3 and T_3 , and, generally, B_i and T_i are defined similarly: B_i stores k -mers of T_{i-1} using a Bloom filter, and T_i contains k -mers stored in B_i “by mistake”, i.e. those k -mers that do not belong to T_{i-1} but belong to T_{i-2} (we call them **false** ^{i} positives). Observe that $T_0 \cap T_1 = \emptyset$, $T_0 \supseteq T_2 \supseteq T_4 \dots$ and $T_1 \supseteq T_3 \supseteq T_5 \dots$

The following lemma shows that the construction is correct, that is it allows one to verify whether or not a given k -mer belongs to the set T_0 .

Lemma 5.1. *Given a k -mer (vertex) K , consider the smallest i such that $K \notin B_{i+1}$ (if $K \notin B_1$, we define $i = 0$). Then, if i is odd, then $K \in T_0$, and if i is even (including 0), then $K \notin T_0$.*

Proof. Observe that $K \notin B_{i+1}$ implies $K \notin T_i$ by the basic property of Bloom filters that membership queries have one-sided error, i.e. there are no false negatives. We first check the Lemma for $i = 0, 1$.

For $i = 0$, we have $K \notin B_1$, and then $K \notin T_0$.

For $i = 1$, we have $K \in B_1$ but $K \notin B_2$. The latter implies that $K \notin T_1$, and then K must be a false² positive, that is $K \in T_0$. Note that here we use the fact that the only queried k -mers K are either vertices of T_0 or their neighbors in the graph (see [Chikhi and Rizk \(2012\)](#)), and therefore if $K \in B_1$ and $K \notin T_0$ then $K \in T_1$.

For the general case $i \geq 2$, we show by induction that $K \in T_{i-1}$. Indeed, $K \in B_1 \cap \dots \cap B_i$ implies $K \in T_{i-1} \cup T_i$ (which, again, is easily seen by induction), and $K \notin B_{i+1}$ implies $K \notin T_i$.

Since $T_{i-1} \subseteq T_0$ for odd i , and $T_{i-1} \subseteq T_1$ for even i (for $T_0 \cap T_1 = \emptyset$), the lemma follows. \square

Naturally, the lemma provides an algorithm to check if a given k -mer K belongs to the graph: it suffices to check successively if it belongs to B_1, B_2, \dots until we encounter the first B_{i+1} which does not contain K . Then, the answer will simply depend on whether i is even or odd: K belongs to the graph if and only if i is odd.

In our reasoning so far, we assumed an infinite number of bitmaps B_i . Of course, in practice we cannot store infinitely many (and even simply many) bitmaps. Therefore, we “truncate” the construction at some step t and store a finite set of bitmaps B_1, B_2, \dots, B_t together with an explicit representation of T_t . The procedure of Lemma 5.1 is extended in the obvious way: if for all $1 \leq i \leq t$, $K \in B_i$, then the answer is determined by directly checking $K \in T_t$.

5.4 Memory and time usage

First, we estimate the memory needed by our data structure, under the assumption of an infinite number of bitmaps. Let N be the number of “true positives”, i.e. vertices of T_0 . As

¹By a slight abuse of language, we say that “an element belongs to B_j ” if it is accepted by the corresponding Bloom filter.

k	optimal r for $t = 4$	bits per k -mer for $t = 4$	optimal r for $t = 6$	bits per k -mer for $t = 6$	bits per k -mer for $t = 1$
16	5.777	8.556	5.506	8.459	12.078
32	6.049	8.664	5.556	8.47	13.518
64	6.399	8.824	5.641	8.49	14.958
128	6.819	9.045	5.772	8.524	16.398

Table 5.1: 1st column: k -mer size; 2nd and 4th columns: optimal value of r for Bloom filters (bitmap size per number of stored elements) for $t = 4$ and $t = 6$ respectively; 3rd and 5th columns: the resulting space per k -mer (for $t = 4$ and $t = 6$); 6th column: space per k -mer for the method of [Chikhi and Rizk \(2012\)](#) ($t = 1$)

stated in Section 5.2, if T_0 has to be stored via a bitmap B_1 of size rN , the false positive rate can be estimated as c^r , where $c = 0.6185$. And, the expected number of critical false positive vertices (set T_1) has been estimated in [Chikhi and Rizk \(2012\)](#) to be $8Nc^r$, as every vertex has eight extensions, i.e. potential neighbors in the graph. We slightly refine this estimation to $6Nc^r$ by noticing that for most of the graph vertices, two out of these eight extensions belong to T_0 (are real vertices) and thus only six are potential false positives. Furthermore, to store these $6Nc^r$ critical false positive vertices, we use a bitmap B_2 of size $6rNc^r$. Bitmap B_3 is used for storing vertices of T_0 which are stored in B_2 “by mistake” (set T_2). We estimate the number of these vertices as the fraction c^r (false positive rate of filter B_2) of N (size of T_0), that is Nc^r . Similarly, the number of vertices we need to put to B_4 is $6Nc^r$ multiplied by c^r , i.e. $6Nc^{2r}$. Continuing in this way, the memory needed for the whole structure is $rN + 6rNc^r + rNc^r + 6rNc^{2r} + rNc^{2r} + \dots$ bits. The number of bits per k -mer is then

$$r + 6rc^r + rc^r + 6rc^{2r} + \dots = (r + 6rc^r)(1 + c^r + c^{2r} + \dots) = (1 + 6c^r) \frac{r}{1 - c^r}. \quad (5.3)$$

A simple calculation shows that the minimum of this expression is achieved when $r = 5.464$, and then the minimum memory used per k -mer is 8.45 bits.

As mentioned earlier, in practice we store only a finite number of bitmaps B_1, \dots, B_t together with an explicit representation (such as array or hash table) of T_t . In this case, the memory taken by the bitmaps is a truncated sum $rN + 6rNc^r + rNc^r + \dots$, and a data structure storing T_t takes either $2k \cdot Nc^{\lceil \frac{t}{2} \rceil r}$ or $2k \cdot 6Nc^{\lceil \frac{t}{2} \rceil r}$ bits, depending on whether t is even or odd. The latter follows from the observations that we need to store $Nc^{\lceil \frac{t}{2} \rceil r}$ (or $6rNc^{\lceil \frac{t}{2} \rceil r}$) k -mers, each taking $2k$ bits of memory. Consequently, we have to adjust the optimal value of r minimizing the total space, and re-estimate the resulting space spent on one k -mer.

Table 5.1 shows estimations for optimal values of r and the corresponding space per k -mer for $t = 4$ and $t = 6$, and several values of k . The data demonstrates that even such small values of t lead to considerable memory savings. It appears that the space per k -mer is very close to the “optimal” space (8.45 bits) obtained for the infinite number of filters. Table 5.1 reveals another advantage of our improvement: the number of bits per stored k -mer remains almost constant for different values of k .

The last column of Table 5.1 shows the memory usage of the original method of [Chikhi and Rizk \(2012\)](#), obtained using the estimation $(1.44 \log_2(\frac{16k}{2.08}) + 2.08)$ the authors provided.

Note that according to that estimation, doubling the value of k results in a memory increment by 1.44 bits, whereas in our method the increment is of 0.11 to 0.22 bits.

Let us now estimate preprocessing and query times for our scheme. If the value of t is small (such as $t = 4$, as in Table 5.1), the preprocessing time grows insignificantly in comparison to the original method of Chikhi and Rizk (2012). To construct each B_i , we need to store T_{i-2} (possibly on disk, if we want to save on the internal memory used by the algorithm) in order to compute those k -mers which are stored in B_{i-1} “by mistake”. The preprocessing time increases little in comparison to the original method of Chikhi and Rizk (2012), as the size of B_i decreases exponentially and then the time spent to construct the whole structure is linear on the size of T_0 .

The query time can be split in two parts: the time spent on querying t Bloom filters and the time spent on querying T_t . Clearly, using t Bloom filters instead of a single one introduces a multiplicative factor of t to the first part of the query time. On the other hand, the set T_t is generally much smaller than T_1 , due to the above-mentioned exponential decrease. Depending on the data structure for storing T_t , the time saving in querying T_t vs. T_1 may even dominate the time loss in querying multiple Bloom filters. Our experimental results (Section 5.5.1 below) confirm that this situation does indeed occur in practice. Note that even in the case when querying T_t weakly depends on its size (e.g. when T_t is implemented by a hash table), the query time will not increase much, due to our choice of a small value for t , as discussed earlier.

5.4.1 Using different values of r for different filters

In the previous section, we assumed that each of our Bloom filters uses the same value of r , the ratio of bitmap size to the number of stored k -mers. However, formula (5.3) for the number of bits per k -mer shows a difference for odd and even filter indices. This suggests that using different parameters r for different filters, rather than the same for all filters, may reduce the space even further. If r_i denotes the corresponding ratio for filter B_i , then (5.3) should be rewritten to

$$r_1 + 6r_2c^{r_1} + r_3c^{r_2} + 6r_4c^{r_1+r_3} + \dots, \quad (5.4)$$

and the minimum value of this expression becomes 7.93 (this value is achieved with $r_1 = 4.41; r_i = 1.44, i > 1$).

In the same way, we can use different values of r_i in the truncated case. This leads to a small 2% to 4% improvement in comparison with case of unique value of r . Table 5.2 shows results for the case $t = 4$ for different values of k .

5.4.2 Query distribution among filters

The query algorithm of Lemma 5.1 simply queries Bloom filters B_1, \dots, B_t successively as long as the returned answer is positive. The query time then directly depends on the number of filters applied before getting a negative answer. Therefore, it is instructive to analyze how the query frequencies to different filters are distributed when performing a graph traversal. We provide such an analysis in this section.

We analyze query frequencies during an exhaustive traversal of the de Bruijn graph, when each true node is visited exactly once. We assume that each time a true node is visited, all its eight potential neighbors are queried, as there is no other way to tell which of those neighbors are real. Note however that this assumption does not take into account structural

k	r_1, r_2, r_3, r_4	bits per k -mer different values of r	bits per k -mer single value of r
16	5.254, 3.541, 4.981, 8.653	8.336	8.556
32	5.383, 3.899, 5.318, 9.108	8.404	8.664
64	5.572, 4.452, 5.681, 9.108	8.512	8.824
128	5.786, 5.108, 6.109, 9.109	8.669	9.045

Table 5.2: Estimated memory occupation for the case of different values of r vs. single value of r , for 4 Bloom filters ($t = 4$). Numbers in the second column represent values of r_i on which the minimum is achieved. For the case of single r , its value is shown in Table 5.1.

properties of the de Bruin graph, nor any additional statistical properties of the genome (such as genomic word frequencies).

For a filter B_i , we want to estimate the number of queried k -mers resolved by B_i during the traversal, that is queries on which B_i returns *no*. This number is the difference of the number of queries submitted to B_i and the number of queries for which B_i returns *yes*. Note that the queries submitted to B_i are precisely those on which the previous filter B_{i-1} returns *yes*.

If the input set T_0 contains N k -mers, then the number of queries in a graph traversal is $8N$, since for each true node each of its 8 potential neighbors are queried. Moreover, about $2N$ queries correspond to true k -mers, as we assume that most of the graph nodes have two true neighbors. Filter B_1 will return *yes* on $2N + 6c^r N$ queries, corresponding to the number of true and false positives respectively. For an arbitrary i , filter B_i returns *yes* precisely on the k -mers inserted to B_i (i.e. k -mers B_i is built on), and the k -mers which are inserted to B_{i+1} (which are the critical false positives for B_i). The counts then easily follow from the analysis of Section 5.4.

	B_1	B_2	B_3	B_4
nb of queries	$8N$	$(2 + 6c^r)N$	$(6c^r + 2c^r)N$	$(2c^r + 6c^{2r})N$
queries returning <i>yes</i>	$(2 + 6c^r)N$	$(6c^r + 2c^r)N$	$(2c^r + 6c^{2r})N$	$(6c^{2r} + 2c^{2r})N$
queries returning <i>no</i>	$(6 - 6c^r)N$	$(2 - 2c^r)N$	$(6c^r - 6c^{2r})N$	$(2c^r - 2c^{2r})N$
resolved queries	69.57%	23.19%	5.04%	1.68%

Table 5.3: Estimations of the number of queries made to filters B_1, B_2, B_3, B_4 in the case of infinite number of filters. Last row: fraction of queries resolved by each filter, estimated for the optimal value $r = 5.464$.

Table 5.3 provides counts for the first four filters, together with the estimated fraction of k -mers resolved by each filter (last row), for the case of infinite number of filters. The data shows that 99.48% of all k -mers are resolved by four filters. This suggests that a very small number of filters should be sufficient to cover a vast majority of k -mers. Furthermore, Table 5.4 shows data for 1-, 2- and 4-filter setups, this time with the optimal value of r for

each case. Even two filters are already sufficient to reduce the accesses to T_2 to 2.08%. In case of four filters, 99.7% of k -mers are resolved before accessing T_4 .

value of t	r	B_1	B_2	B_3	B_4	T_t
1	11.44	74.70%	0	0	0	25.3%
2	8.060	73.44%	24.48%	0	0	2.08%
4	6.049	70.90%	23.63%	3.88%	1.29%	0.3%

Table 5.4: Fractions of queries resolved by each filter for, 1, 2 and 4 filters. Estimations have been computed for $k = 32$ and optimal values of r shown in the second column. Last column shows the fraction of queries resolved at the last step, by testing against the explicitly stored set T_t .

5.5 Experimental results

5.5.1 Construction algorithm

In practice, constructing a cascading Bloom filter for a real-life read set is a computationally intensive step. To perform it on a commonly-used computer, the implementation makes an essential use of external memory. Here we give a short description of the construction algorithm for up to four Bloom filters. Extension for larger number of filters is straightforward.

We start from the input set T_0 of k -mers written on disk. We build the Bloom filter B_1 of appropriate size by inserting elements of T_0 successively. Next, all possible extensions of each k -mer in T_0 are queried against B_1 , and those which return true are written to the disk. Then, in this set only the k -mers absent from T_0 are kept, i.e. we perform a set difference from T_0 . We cannot afford to load T_0 entirely in memory, so we partition T_0 and perform the set difference in several iterations, loading only one partition of T_0 each time. This results in the set T_1 of critical false positives, which is also kept on disk. Up to this point, the procedure is identical to that of [Chikhi and Rizk \(2012\)](#).

Next, we insert all k -mers from T_1 into B_2 and to obtain T_2 , we check for each k -mer in T_0 if a query to B_2 returns true. This results in the set T_2 , which is directly stored on disk. Thus, at this point we have B_1 , B_2 and, by loading T_2 from the disk, a complete representation for $t = 2$. In order to build the data structure for $t = 4$, we continue this process, by inserting T_2 in B_3 and retrieving (and writing directly on disk) T_3 from T_1 (stored on disk). It should be noted that to obtain T_i we need T_{i-2} , and by always directly storing it on disk we guarantee not to use more memory than the size of the final structure. The set T_t (that is, T_1 , T_2 or T_4 in our experiments) is represented as a sorted array and is searched by a binary search. We found this implementation more efficient than a hash table.

5.5.2 Implementation and experimental setup

We implemented our method using MINIA software ([Chikhi and Rizk \(2012\)](#)) and ran comparative tests for 2 and 4 Bloom filters ($t = 2, 4$). Note that since the only modified part of MINIA was the construction step and the k -mer membership queries, this allows us to precisely evaluate our method against the one of [Chikhi and Rizk \(2012\)](#).

The first step of the implementation is to retrieve the list of k -mers that appear more than d times using DSK (Rizk et al. (2013)) – a constant memory streaming algorithm to count k -mers. Note, as a side remark, that performing counting allows us to perform off-line deletions of k -mers. That is, if at some point of the scan of the input set of k -mers (or reads) some of them should be deleted, it is done by a simple decrement of the counter.

Assessing the query time is done through the procedure of graph traversal, as it is implemented in Chikhi and Rizk (2012). Since the procedure is identical and independent on the data structure, the time spent on graph traversal is a faithful estimator of the query time.

We compare three versions: $t = 1$ (i.e. the version of Chikhi and Rizk (2012)), $t = 2$ and $t = 4$. For convenience, we define 1 Bloom, 2 Bloom and 4 Bloom as the versions with $t = 1, 2$ and 4, respectively.

5.5.3 *E. coli* dataset, varying k

In this set of tests, our main goal was to evaluate the influence of the k -mer size on principal parameters: size of the whole data structure, size of the set T_t , graph traversal time, and time of construction of the data structure. We retrieved 10M *E. coli* reads of 100bp from the *Short Read Archive* (ERX008638) without read pairing information and extracted all k -mers occurring at least two times. The total number of k -mers considered varied, depending on the value of k , from 6,967,781 ($k = 15$) to 5,923,501 ($k = 63$). We ran each version, 1 Bloom (Chikhi and Rizk (2012)), 2 Bloom and 4 Bloom, for values of k ranging from 16 to 64. The results are shown in Fig. 5.1.

The total size of the structures in bits per stored k -mer, i.e. the size of B_1 and T_1 (respectively, B_1, B_2, T_2 or B_1, B_2, B_3, B_4, T_4) is shown in Fig. 5.1a. As expected, the space for 4 Bloom filters is the smallest for all values of k considered, showing a considerable improvement, ranging from 32% to 39%, over the version of Chikhi and Rizk (2012). Even the version with just 2 Bloom filters shows an improvement of at least 20% over Chikhi and Rizk (2012), for all values of k . Regarding the influence of the k -mer size on the structure size, we observe that for 4 Bloom filters the structure size is almost constant, the minimum value is 8.60 and the largest is 8.89, an increase of only 3%. For 1 and 2 Bloom the same pattern is seen: a plateau from $k = 16$ to 32, a jump for $k = 33$ and another plateau from $k = 33$ to 64. The jump at $k = 32$ is due to switching from 64-bit to 128-bit representation of k -mers in the table T_t .

The traversal times for each version is shown in Fig. 5.1c. The fastest version is 4 Bloom, showing an improvement over Chikhi and Rizk (2012) of 18% to 30%, followed by 2 Bloom. This result is surprising and may seem counter-intuitive, as we have four filters to apply to the queried k -mer rather than a single filter as in Chikhi and Rizk (2012). However, the size of T_4 (or even T_2) is much smaller than T_1 , as the size of T_i 's decreases exponentially. As T_t is stored in an array, the time economy in searching T_4 (or T_2) compared to T_1 dominates the time lost on querying additional Bloom filters, which explains the overall gain in query time.

As far as the construction time is concerned (Fig. 5.1d), our versions yielded also a faster construction, with the 4 Bloom version being 5% to 22% faster than that of Chikhi and Rizk (2012). The gain is explained by the time required for sorting the array storing T_t , which is much higher for T_0 than for T_2 or T_4 . However, the gain is less significant here, and, on the other hand, was not observed for bigger datasets (see Section 5.5.6).

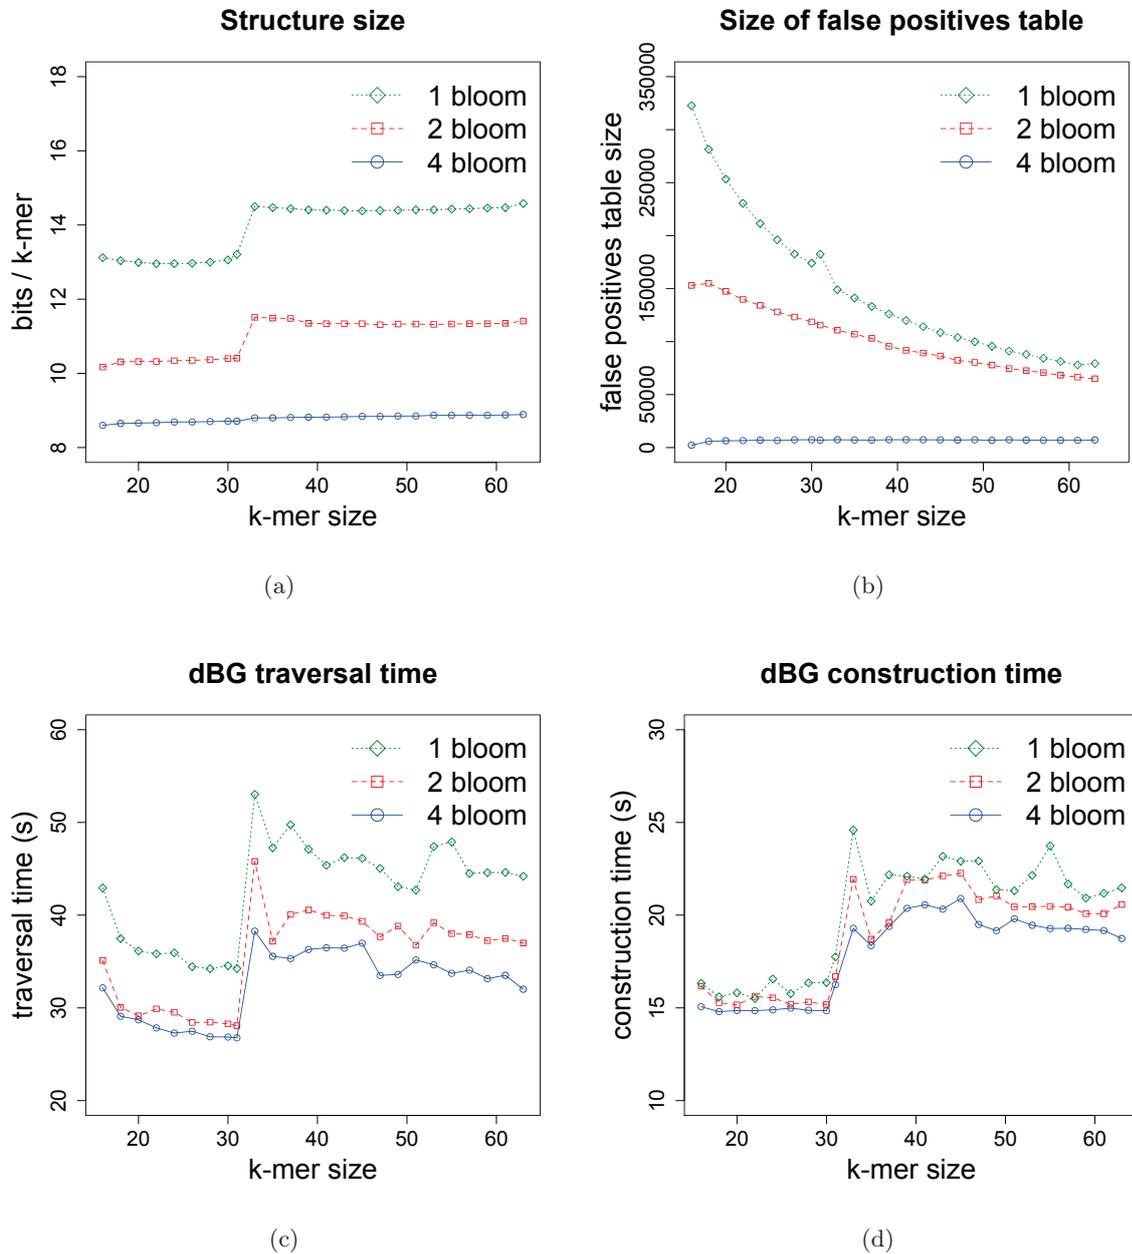


Figure 5.1: Results for 10M E.coli reads of 100bp using several values of k . The *1 Bloom* version corresponds to the one presented in [Chikhi and Rizk \(2012\)](#). (a) Size of the structure in bits used per k -mer stored. (b) Number of false positives stored in T_1 , T_2 or T_4 for 1, 2 or 4 Bloom filters, respectively. (c) De Bruijn graph construction time, excluding k -mer counting step. (d) De Bruijn graph traversal time, including branching k -mer indexing.

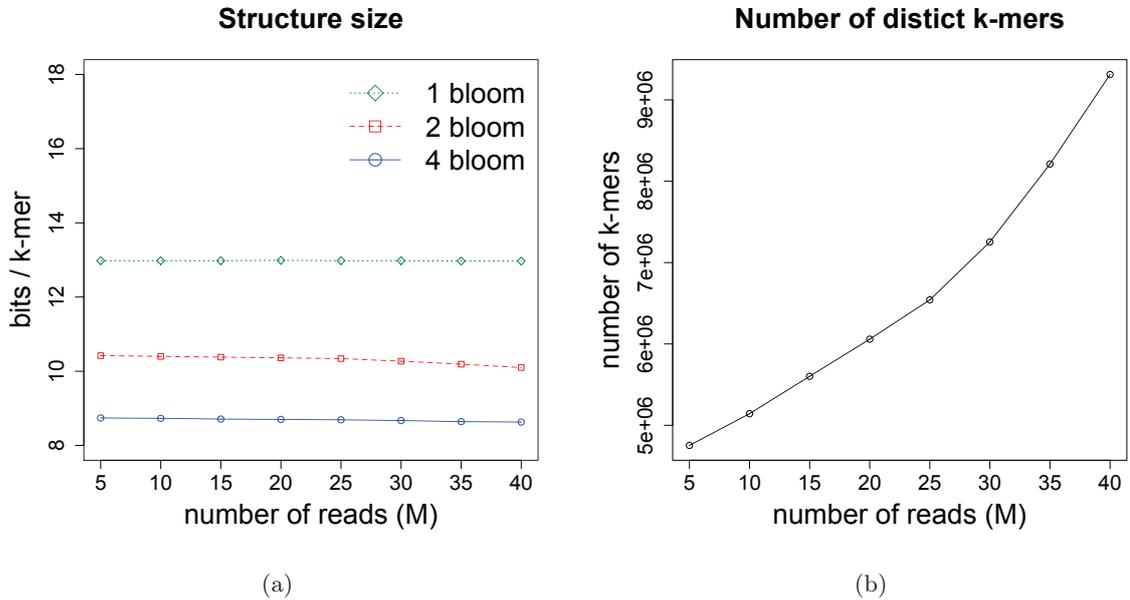


Figure 5.2: Results for *E. coli* reads of 100bp using $k = 27$. The 1 Bloom version corresponds to the one presented in Chikhi and Rizk (2012). (a) Size of the structure in bits used per k -mer stored. (b) Number of distinct k -mers.

5.5.4 *E. coli* dataset, varying coverage

From the complete *E. coli* dataset (≈ 44 M reads) from the previous section, we selected several samples ranging from 5M to 40M reads in order to assess the impact of the coverage on the size of the data structures. This strain *E. coli* (K-12 MG1655) is estimated to have a genome of 4.6M bp Blattner et al. (1997), implying that a sample of 5M reads (of 100bp) corresponds to ≈ 100 X coverage. We set $d = 3$ and $k = 27$. The results are shown in Fig. 5.2. As expected, the memory consumption per k -mer remains almost constant for increasing coverage, with a slight decrease for 2 and 4 Bloom. The best results are obtained with the 4 Bloom version, an improvement of 33% over the 1 Bloom version of Chikhi and Rizk (2012). On the other hand, the number of distinct k -mers increases markedly (around 10% for each 5M reads) with increasing coverage, see Fig. 5.2b. This is due to sequencing errors: an increase in coverage implies more errors with higher coverage, which are not removed by our cutoff $d = 3$. This suggests that the value of d should be chosen according to the coverage of the sample. Moreover, in the case where read qualities are available, a quality control pre-processing step may help to reduce the number of sequencing errors.

5.5.5 *E. coli* dataset, query statistics

In this set of tests we used the dataset of Section 5.5.3 to experimentally evaluate how the queries are distributed among the Bloom filters. We ran the graph traversal algorithm for each version, 1 Bloom (Chikhi and Rizk (2012)), 2 Bloom and 4 Bloom, using values of k ranging from 16 to 64 and retrieved the number of queries resolved in each Bloom filter and the table T_t . The results are shown in Fig. 5.3. The plots indicate that, for each version, the query distribution among the Bloom filters is approximately invariant to the value of k .

Indeed, on average 74%, 73% and 70% of the queries are resolved in B_1 for the 1, 2 and 4 Bloom version, respectively, and the variance is smaller than 0.01% in each case. For the 4 Bloom version, 70%, 24%, 4%, 1% and 0.2% of the queries are resolved in B_1 , B_2 , B_3 , B_4 and T_4 , respectively, showing that the values estimated theoretically in Section 5.4.2 (the last row of Table 5.4) are very precise. Furthermore, as a query to a Bloom filter is faster than to T_1 and the majority of the queries to 4 and 2 Bloom versions, 94% and 95% respectively, are resolved in the first two filters, it is natural that on average queries to 1 Bloom version are slower than to 2 and 4 Bloom versions, corroborating the results of Section 5.5.3.

5.5.6 Human dataset

We also compared 2 and 4 Bloom versions with the 1 Bloom version of Chikhi and Rizk (2012) on a large dataset. For that, we retrieved 564M Human reads of 100bp (SRA: SRX016231) without pairing information and discarded the reads occurring less than 3 times. The dataset corresponds to $\approx 17X$ coverage. A total of 2,455,753,508 k -mers were indexed. We ran each version, 1 Bloom (Chikhi and Rizk (2012)), 2 Bloom and 4 Bloom with $k = 23$. The results are shown in Table 5.5.

The results are in general consistent with the previous tests on *E.coli* datasets. There is an improvement of 34% (21%) for the 4 Bloom (2 Bloom) in the size of the structure. The graph traversal is also 26% faster in the 4 Bloom version. However, in contrast to the previous results, the graph construction time increased by 10% and 7% for 4 and 2 Bloom versions respectively, when compared to the 1 Bloom version. This is due to the fact that disk writing/reading operations now dominate the time for the graph construction, and 2 and 4 Bloom versions generate more disk accesses than 1 Bloom. As stated in Section 5.5.1, when constructing the 1 Bloom structure, the only part written on the disk is T_1 and it is read only once to fill an array in memory. For 4 Bloom, T_1 and T_2 are written to the disk, and T_0 and T_1 are read at least one time each to build B_2 and B_3 . Moreover, since the size coefficient of B_1 reduces, from $r = 11.10$ in 1 Bloom to $r = 5.97$ in 4 Bloom, the number of false positives in T_1 increases.

5.6 Discussion and conclusions

Using cascading Bloom filters for storing de Bruijn graphs has clear advantage over the single-filter method of Chikhi and Rizk (2012). In terms of memory consumption, which is the main parameter here, we obtained an improvement of around 30%-40% in all our experiments. Our data structure takes 8.5 to 9 bits per stored k -mer, compared to 13 to 15 bits by the method of Chikhi and Rizk (2012). This confirms our analytical estimations. The above results were obtained using only four filters and are very close to the estimated optimum (around 8.4 bits/ k -mer) produced by the infinite number of filters. An interesting characteristic of our method is that the memory grows insignificantly with the growth of k , even slower than with the method of Chikhi and Rizk (2012). Somewhat surprisingly, we also obtained a significant decrease, of order 20%-30%, of query time. The construction time of the data structure varied from being 10% slower (for the human dataset) to 22% faster (for the bacterial dataset).

As stated previously, another compact encoding of de Bruijn graphs has been proposed in Bowe et al. (2012), however no implementation of the method was made available. For this reason, we could not experimentally compare our method with the one of Bowe et al. (2012). We remark, however, that the space bound of Bowe et al. (2012) heavily depends

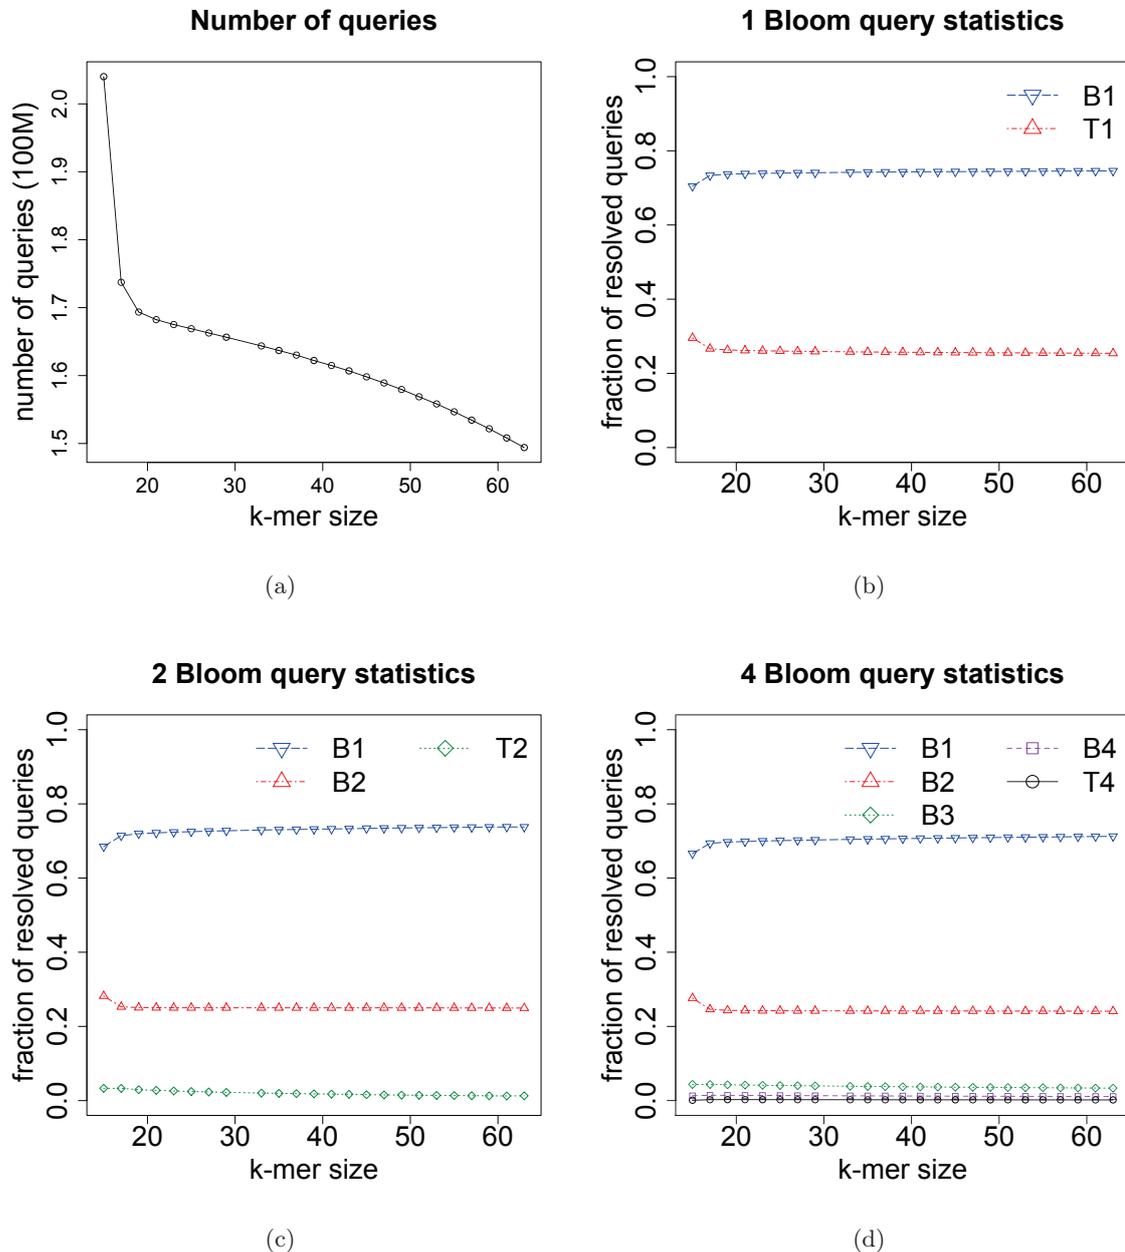


Figure 5.3: Query statistics results for 10M E.coli reads of 100bp using several values of k . The 1 Bloom version corresponds to the one presented in Chikhi and Rizk (2012). (a) Total number of queries performed, for each value of k , during a graph traversal. (b) Fraction of resolved queries in B_1 and T_1 (1 Bloom version) for each value of k . (c) Fraction of resolved queries in B_1, B_2 and T_2 (2 Bloom version) for each value of k . (d) Fraction of resolved queries in B_1, B_2, B_3, B_4 and T_4 for each value of k .

Method	1 Bloom	2 Bloom	4 Bloom
Construction time (s)	40160.7	43362.8	44300.7
Traversal time (s)	46596.5	35909.3	34177.2
r coefficient	11.10	7.80	5.97
Bloom filters size (MB)	$B_1 = 3250.95$	$B_1 = 2283.64$ $B_2 = 323.08$	$B_1 = 1749.04$ $B_2 = 591.57$ $B_3 = 100.56$ $B_4 = 34.01$
False positive table size (MB)	$T_1 = 545.94$	$T_2 = 425.74$	$T_4 = 36.62$
Total size (MB)	3796.89	3032.46	2511.8
Size (bits/k-mer)	12.96	10.35	8.58

Table 5.5: Results of 1, 2 and 4 Bloom filters version for 564M Human reads of 100bp using $k = 23$. The *1 Bloom* version corresponds to the one presented in [Chikhi and Rizk \(2012\)](#).

on the number of reads (i.e. coverage), while in our case, the data structure size is almost invariant with respect to the coverage (Section 5.5.4).

An interesting open question is whether the Bloom filter construction can be made online, so that new k -mers (reads) can be inserted without reconstructing the whole data structure from scratch. Note that the presented construction (Section 5.5.1) is inherently off-line, as all k -mers should be known before the data structure is built.

Another interesting prospect for further possible improvements of our method is offered by [Porat \(2009\)](#), where an efficient replacement to Bloom filter was introduced. The results of [Porat \(2009\)](#) suggest that we could hope to reduce the memory to about 5 bits per k -mer. However, there exist obstacles on this way: an implementation of such a structure would probably result in a significant construction and query time increase.

Conclusion and perspectives

In this thesis, we presented KISSPLICE, a time and memory efficient method to identify variations (alternative splicing and genomic polymorphisms) by locally assembling RNA-seq data without using a reference genome. The local nature of the KISSPLICE strategy allows to avoid some of the difficulties faced by standard full-length transcriptome assemblers, namely solving an ill-posed problem often formulated as a NP-hard optimization problem. As a result, we can avoid an extensive use of heuristics and, thus, obtain an overall more sensitive method with stronger theoretical guarantees. A lot of effort was put in order to make our method as scalable as possible in order to deal with ever-increasing volumes of NGS data. We improved the state-of-the-art de Bruijn graph construction and representation in an effort to reduce the memory footprint of KISSPLICE. We also developed a new time-efficient approach to list bubbles in de Bruijn graphs in order to reduce the running time of our method.

The techniques we developed while studying the bubble listing problem turned out to be useful in other enumeration contexts, namely cycle listing and the K -shortest paths problem. The classical problem of listing cycles in a graph has been studied since the early 70s, however, as shown here, the best algorithm for undirected graph is not optimal. In this thesis, we gave the first optimal algorithm to list cycles in undirected graphs, along with the first optimal algorithm to list st -paths. The classical K -shortest paths problem has been studied since the early 60s, and the best algorithm for solving it uses memory proportional to the number of path output, i.e. K . In this thesis, we gave an alternative parameterization of the problem. For this alternative version, we gave an algorithm that uses memory linear in the size of the graph, independent of the number of paths output.

In the past 3 years, KISSPLICE has evolved into not only a time and memory efficient method, but also into an user-friendly *software* for the bioinformaticians and biologists. We are now a 5-persons team actively developing KISSPLICE, which includes continuously: improving its robustness, correcting bugs, improving the usability and the documentation. It is important to highlight that, besides the algorithmic improvements already mentioned, a lot of effort was put on improving the *implementation* of KISSPLICE, including: parallelization of certain steps of the pipeline, careful implementation of the data structures, systematic removal of memory leaks, among others. This team effort produced a stable and user-friendly software.

KISSPLICE has been used in several projects, as evidenced by an average of 250 unique visitors per month to our website. A summary of some of the projects in which we are directly involved is shown in the figure below. We are aware that in order to further convince the biologists that the new alternative splicing events found by KISSPLICE are *real* events, it is desirable to experimentally validate some of them, and for that we are working in collaboration with D. Auboeuf's group to validate AS events found by KISSPLICE in K562 cell lines. Since the human genome is known and well annotated, our choice of human cell lines for these experiments may seem a bit odd. Actually, our goal is to show that KISSPLICE is useful even

when a good reference genome is available. For that, we need to validate AS events found by KISSPLICE that are not annotated and not found by mapping approaches (Trapnell et al. (2010b)) either, even when there is a good (annotated) reference genome. The preliminary results are promising: from the randomly selected events only found by KISSPLICE, around 40% of them were validated, while 30%, although not validated, are part of complex events (more than two isoforms) and the isoform amplified in the experiment matches the AS event found by KISSPLICE but not selected for validation. The remaining 30% corresponded to cases where the minor isoform had a relative abundance of less than 15%. Although for now we did not manage to validate these cases, it does not yet mean that they are not real. Indeed, since the experimental validation is based on RT-PCR, it may be that the early rounds of the PCR favor the major isoform, which causes the complete loss of the minor isoform in the final rounds. Finally, after clarifying how many of the novel events found by KISSPLICE are real, there still remains the central question whether these new isoforms are functional or just noise of the splicing machinery. Our point of view is that an exhaustive description of all isoforms present in the cell is a good prerequisite to help address this central question.

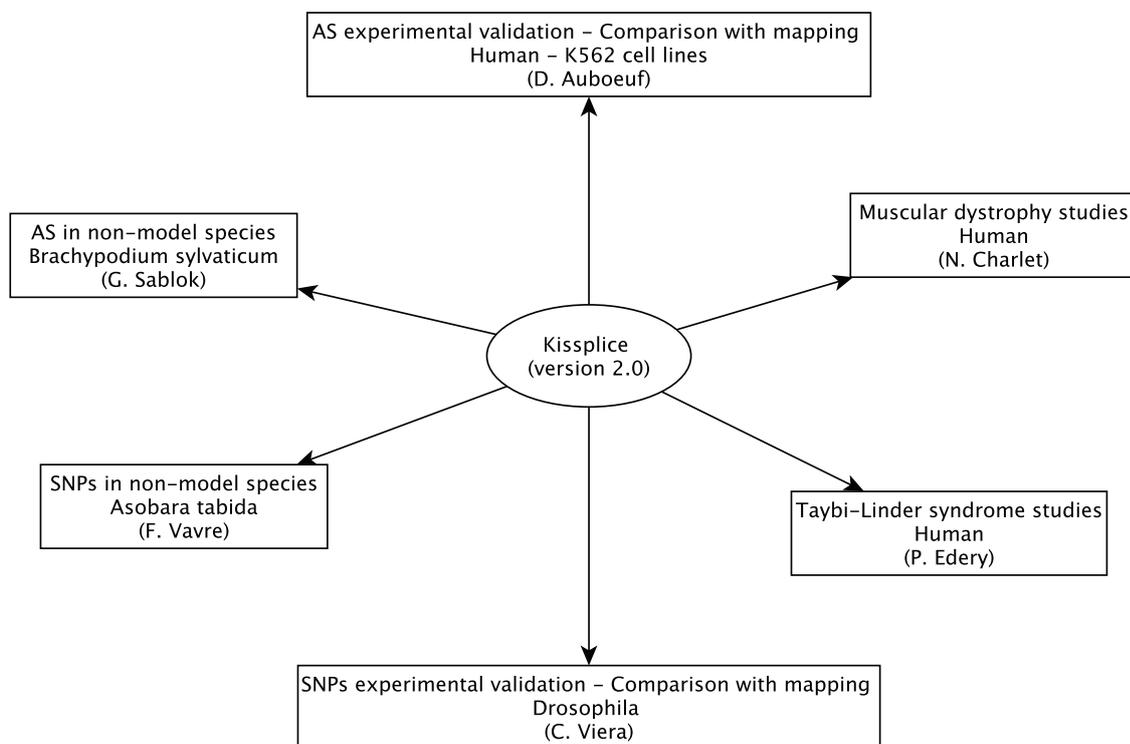


Figure 5.4: Projects using KISSPLICE.

At the end of each chapter, open problems and perspectives were discussed, providing the reader with an idea of possible extensions of the methods and techniques presented in this work. For that reason, in this chapter we focus on the main open problem concerning KISSPLICE: the complex biconnected components (BCCs).

The second step of the KISSPLICE algorithm is the biconnected decomposition. Since each bubble is entirely contained in one BCC, after the BCC decomposition, the bubbles in each BCC are enumerated independently. In the ideal case – a repeat-free genome – each connected component in the DBG corresponds to a single gene, as well as each BCC. In practice, the BCC

decomposition works well: the vast majority of the BCCs are relatively small and contain the sequence of a single gene (or a family of paralogous genes). However, there is a small number of large BCCs (often one or two) that contain the sequences of several unrelated genes, and it is infeasible to enumerate all the bubbles inside these complex BCCs. The problem is not the efficiency of the algorithm, but the number of bubbles satisfying our constraints. There are a huge number of bubbles in the complex BCCs (usually more than in all the other BCCs together), and most of them are repeat-associated bubbles. A manual exploration of a fraction of the bubbles contained in these complex BCCs led us to think that they are generated by transposable elements (and to a much lesser extent to other types of repeats). A transposable element (TE) is a DNA sequence that can change its position within the genome through a copy-and-paste or cut-and-paste process (Wicker et al. (2007)). Transposable elements are spread throughout the genome (including in many transcribed regions). We believe that old copies of TEs that invaded the genome a long time ago are responsible for the complex BCCs. Since they are old, these copies diverged, however, they still contain enough sequence similarity to merge several unrelated genes inside the same BCC. More importantly, they are in transcribed regions, mostly UTRs. In human, the transposable elements in the Alu family alone generate a BCC with millions of bubbles satisfying our constraints, so it is infeasible to enumerate all of them. The problem of simply ignoring the complex BCCs (that is what we have been doing so far), is that they potentially contain true events “trapped” inside.

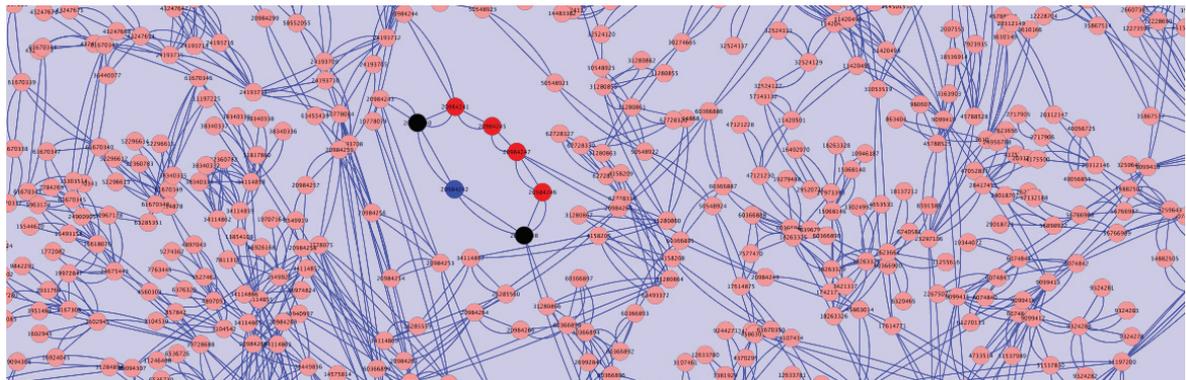


Figure 5.5: An alternative splicing event (intron retention) in the SCN5A gene (human) trapped inside a complex BCC. The switching vertices are shown in black.

Repeated elements are not a problem restricted to KISSPLICE or local assembly strategies; global transcriptome assemblers are possibly even more affected. Recall that the second step of the standard transcriptome heuristic (see Section 1.2.3) is to partition the graph into subgraphs corresponding to genes; in the presence of repeated elements this is a much harder task. Several genes are likely to be wrongly assigned to the same subgraph, and as a result the heuristic is going to produce chimeric transcripts. On the other hand, if to account for the presence of repeated elements, the heuristic adopts a more stringent graph partition strategy, more genes are likely to be split into several subgraphs, resulting in transcripts only partially assembled.

A first step towards a solution to this issue is to solve the following problem: given a DBG built from RNA-seq reads, identify the subgraph corresponding to the transposable elements (they are not the only repeated elements, but we believe they are the main source of problems). Observe that, unlike genomic NGS data where it is possible to use the coverage of a vertex as a proxy for uniqueness of that sequence in the genome, in RNA-seq data it is

not obvious how to determine the uniqueness of the sequence corresponding to a vertex. A solution to this problem would be useful in KISSPLICE as well as in full-length transcriptome assemblers.

Bibliography

- Abdel-Lateef, B. A.-H. (1988). *A Dual-Based Approach to a Multiobjective Location Problem*. PhD thesis, Univ. of Liverpool, Dept. of Mathematics.
- Ahuja, R. K., Mehlhorn, K., Orlin, J. B., and Tarjan, R. E. (1990). Faster algorithms for the Shortest Path Problem. *J. ACM*, 37:213–223.
- Alberts, B., Johnson, A., Lewis, J., Raff, M., Bray, D., Hopkin, K., Roberts, K., and Walter, P. (2003). *Essential Cell Biology, Second Edition*. Garland Science/Taylor & Francis Group.
- Ausiello, G., Protasi, M., Marchetti-Spaccamela, A., Gambosi, G., Crescenzi, P., and Kann, V. (1999). *Complexity and Approximation: Combinatorial Optimization Problems and Their Approximability Properties*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1st edition.
- Avis, D. and Fukuda, K. (1993). Reverse search for enumeration. *Discrete Applied Mathematics*, 65:21–46.
- Bang-Jensen, J. and Gutin, G. Z. (2008). *Digraphs: Theory, Algorithms and Applications*. Springer Publishing Company, Incorporated, 2nd edition.
- Bankevich, A., Nurk, S., Antipov, D., Gurevich, A. A., Dvorkin, M., Kulikov, A. S., Lesin, V. M., Nikolenko, S. I., Pham, S. K., Prjibelski, A. D., Pyshkin, A., Sirotkin, A., Vyahhi, N., Tesler, G., Alekseyev, M. A., and Pevzner, P. A. (2012). Spades: A new genome assembly algorithm and its applications to single-cell sequencing. *Journal of Computational Biology*, 19(5):455–477.
- Bernstein, A. (2010). A nearly optimal algorithm for approximating replacement paths and k shortest simple paths in general graphs. In Charikar, M., editor, *SODA*, pages 742–755. SIAM.
- Bezem, G. and Leeuwen, J. v. (1987). Enumeration in graphs. Technical Report RUU-CS-87-07, Utrecht University.
- Birmelé, E., Crescenzi, P., Ferreira, R. A., Grossi, R., Lacroix, V., Marino, A., Pisanti, N., Sacomoto, G. A. T., and Sagot, M.-F. (2012). Efficient bubble enumeration in directed graphs. In *SPIRE*, volume 7608 of *Lecture Notes in Computer Science*, pages 118–129. Springer.
- Birmelé, E., Ferreira, R. A., Grossi, R., Marino, A., Pisanti, N., Rizzi, R., and Sacomoto, G. (2013). Optimal listing of cycles and st -paths in undirected graphs. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 1884–1896. SIAM.

Birney, E., Stamatoyannopoulos, J., Dutta, A., Guigó, R., Gingeras, T., Margulies, E., Weng, Z., Snyder, M., Dermitzakis, E., Thurman, R., Kuehn, M., Taylor, C., Neph, S., Koch, C., Asthana, S., Malhotra, A., Adzhubei, I., Greenbaum, J., Andrews, R., Flicek, P., Boyle, P., Cao, H., Carter, N., Clelland, G., Davis, S., Day, N., Dhami, P., Dillon, S., Dorschner, M., Fiegler, H., Giresi, P., Goldy, J., Hawrylycz, M., Haydock, A., Humbert, R., James, K., Johnson, B., Johnson, E., Frum, T., Rosenzweig, E., Karnani, N., Lee, K., Lefebvre, G., Navas, P., Neri, F., Parker, S., Sabo, P., Sandstrom, R., Shafer, A., Vetric, D., Weaver, M., Wilcox, S., Yu, M., Collins, F., Dekker, J., Lieb, J., Tullius, T., Crawford, G., Sunyaev, S., Noble, W., Dunham, I., Dencoeud, F., Reymond, A., Kapranov, P., Rozowsky, J., Zheng, D., Castelo, R., Frankish, A., Harrow, J., Ghosh, S., Sandelin, A., Hofacker, I., Baertsch, R., Keefe, D., Dike, S., Cheng, J., Hirsch, H., Sekinger, E., Lagarde, J., Abril, J., Shahab, A., Flamm, C., Fried, C., Hackermüller, J., Hertel, J., Lindemeyer, M., Missal, K., Tanzer, A., Washietl, S., Korb, J., Emanuelsson, O., Pedersen, J., Holroyd, N., Taylor, R., Swarbreck, D., Matthews, N., Dickson, M., Thomas, D., Weirauch, M., Gilbert, J., Drenkow, J., Bell, I., Zhao, X., Srinivasan, K., Sung, W., Ooi, H., Chiu, K., Foissac, S., Alioto, T., Brent, M., Pachter, L., Tress, M., Valencia, A., Choo, S., Choo, C., Ucla, C., Manzano, C., Wyss, C., Cheung, E., Clark, T., Brown, J., Ganesh, M., Patel, S., Tamma, H., Chrast, J., Henrichsen, C., Kai, C., Kawai, J., Nagalakshmi, U., Wu, J., Lian, Z., Lian, J., Newburger, P., Zhang, X., Bickel, P., Mattick, J., Carninci, P., Hayashizaki, Y., Weissman, S., Hubbard, T., Myers, R., Rogers, J., Stadler, P., Lowe, T., Wei, C., Ruan, Y., Struhl, K., Gerstein, M., Antonarakis, S., Fu, Y., Green, E., Karaöz, U., Siepel, A., Taylor, J., Liefer, L., Wetterstrand, K., Good, P., Feingold, E., Guyer, M., Cooper, G., Asimenos, G., Dewey, C., Hou, M., Nikolaev, S., Montoya-Burgos, J., Löytynoja, A., Whelan, S., Pardi, F., Massingham, T., Huang, H., Zhang, N., Holmes, I., Mullikin, J., Ureta-Vidal, A., Paten, B., Srinivasan, M., Church, D., Rosenbloom, K., Kent, W., Stone, E., Program, N. C. S., of Medicine Human Genome Sequencing Center, B. C., Center, W. U. G. S., Institute, B., Institute, C. H. O. R., Batzoglou, S., Goldman, N., Hardison, R., Haussler, D., Miller, W., Sidow, A., Trinklein, N., Zhang, Z., Barrera, L., Stuart, R., King, D., Ameur, A., Enroth, S., Bieda, M., Kim, J., Bhinge, A., Jiang, N., Liu, J., Yao, F., Vega, V., Lee, C., Ng, P., Shahab, A., Yang, A., Moqtaderi, Z., Zhu, Z., Xu, X., Squazzo, S., Oberley, M., Inman, D., Singer, M., Richmond, T., Munn, K., Rada-Iglesias, A., Wallerman, O., Komorowski, J., Fowler, J., Couttet, P., Bruce, A., Dovey, O., Ellis, P., Langford, C., Nix, D., Euskirchen, G., Hartman, S., Urban, A., Kraus, P., Van Calcar, S., Heintzman, N., Kim, T., Wang, K., Qu, C., Hon, G., Luna, R., Glass, C., Rosenfeld, M., Aldred, S., Cooper, S., Halees, A., Lin, J., Shulha, H., Zhang, X., Xu, M., Haidar, J., Yu, Y., Ruan, Y., Iyer, V., Green, R., Wadelius, C., Farnham, P., Ren, B., Harte, R., Hinrichs, A., Trumbower, H., Clawson, H., Hillman-Jackson, J., Zweig, A., Smith, K., Thakkapallayil, A., Barber, G., Kuhn, R., Karolchik, D., Armengol, L., Bird, C., de Bakker, P., Kern, A., Lopez-Bigas, N., Martin, J., Stranger, B., Woodroffe, A., Davydov, E., Dimas, A., Eyas, E., Hallgrímsdóttir, I., Huppert, J., Zody, M., Abecasis, G., Estivill, X., Bouffard, G., Guan, X., Hansen, N., Idol, J., Maduro, V., Maskeri, B., McDowell, J., Park, M., Thomas, P., Young, A., Blakesley, R., Muzny, D., Sodergren, E., Wheeler, D., Worley, K., Jiang, H., Weinstock, G., Gibbs, R., Graves, T., Fulton, R., Mardis, E., Wilson, R., Clamp, M., Cuff, J., Gnerre, S., Jaffe, D., Chang, J., Lindblad-Toh, K., Lander, E., Koriabine, M., Nefedov, M., Osoegawa, K., Yoshinaga, Y., Zhu, B., and de Jong, P. (2007). Identification and analysis of functional elements in 1% of the human genome by the encode pilot project. *Nature*, 447(7146):799–816.

Blattner, F. R., Plunkett, G., and et al., C. B. (1997). The complete genome sequence of

- escherichia coli k-12. *Science*, 277(5331):1453–1462.
- Blencowe, B. J. (2006). Alternative Splicing: New Insights from Global Analyses. *Cell*, 126(1):37–47.
- Bowe, A., Onodera, T., Sadakane, K., and Shibuya, T. (2012). Succinct de Bruijn graphs. In Raphael, B. and J.Tang, editors, *Algorithms in Bioinformatics - 12th International Workshop, WABI 2012, Ljubljana, Slovenia, September 10-12, 2012. Proceedings*, volume 7534 of *Lecture Notes in Computer Science*, pages 225–235. Springer.
- Bron, C. and Kerbosch, J. (1973). Algorithm 457: Finding all cliques of an undirected graph. *Commun. ACM*, 16(9):575–577.
- Burset, M., Seledtsov, I. A., and Solovyev, V. V. (2000). Analysis of canonical and non-canonical splice sites in mammalian genomes. *Nucleic acids research*, 28(21):4364–4375.
- Bussieck, M. R. and Lubbecke, M. E. (1998). The vertex set of a 0/1-polytope is strongly p-enumerable. *Computational Geometry*, 11(2):103–109.
- Butler, J., MacCallum, I., Kleber, M., Shlyakhter, I. A., Belmonte, M. K., Lander, E. S., Nusbaum, C., and Jaffe, D. B. (2008). Allpaths: De novo assembly of whole-genome shotgun microreads. *Genome Research*.
- Byers, T. H. and Waterman, M. S. (1984). Determining all optimal and near-optimal solutions when solving shortest path problems by dynamic programming. *Operations Research*, 32:1381–1384.
- Chaisson, M. J. and Pevzner, P. A. (2008). Short read fragment assembly of bacterial genomes. *Genome Research*, 18(2):324–330.
- Chen, M., Chowdhury, R. A., Ramachandran, V., Roche, D. L., and Tong, L. (2007). Priority Queues and Dijkstra’s Algorithm. Technical Report TR-07-54, University of Texas, Austin.
- Chikhi, R. and Rizk, G. (2012). Space-efficient and exact de Bruijn graph representation based on a Bloom filter. In Raphael, B. J. and Tang, J., editors, *Algorithms in Bioinformatics - 12th International Workshop, WABI 2012, Ljubljana, Slovenia, September 10-12, 2012. Proceedings*, volume 7534 of *Lecture Notes in Computer Science*, pages 236–248. Springer.
- Cloonan, N., Forrest, A. R. R., Kolle, G., Gardiner, B. B. A., Faulkner, G. J., Brown, M. K., Taylor, D. F., Steptoe, A. L., Wani, S., Bethel, G., Robertson, A. J., Perkins, A. C., Bruce, S. J., Lee, C. C., Ranade, S. S., Peckham, H. E., Manning, J. M., McKernan, K. J., and Grimmond, S. M. (2008). Stem cell transcriptome profiling via massive-scale mRNA sequencing. *Nature Methods*, 5(7):613–619.
- Cloonan, N. and Grimmond, S. (2008). Transcriptome content and dynamics at single-nucleotide resolution. *Genome Biology*, 9(9):1–4.
- Conway, T. and Bromage, A. (2011). Succinct data structures for assembling large genomes. *Bioinformatics*, 27(4):479–486.
- Cormen, T. H., Leiserson, C. E., Rivest, R. L., and Stein, C. (2001). *Introduction to Algorithms*. The MIT Press, 2 edition.

- Dechter, R. and Pearl, J. (1985). Generalized best-first search strategies and the optimality of a^* . *J. ACM*, 32(3):505–536.
- Diestel, R. (2005). *Graph Theory (Graduate Texts in Mathematics)*. Springer.
- Djebali, S., Davis, C. A., Merkel, A., Dobin, A., Lassmann, T., Mortazavi, A., Tanzer, A., Lagarde, J., Lin, W., Schlesinger, F., Xue, C., Marinov, G. K., Khatun, J., Williams, B. A., Zaleski, C., Rozowsky, J., Roder, M., Kokocinski, F., Abdelhamid, R. F., Alioto, T., Antoshechkin, I., Baer, M. T., Bar, N. S., Batut, P., Bell, K., Bell, I., Chakraborty, S., Chen, X., Chrast, J., Curado, J., Derrien, T., Drenkow, J., Dumais, E., Dumais, J., Duttagupta, R., Falconnet, E., Fastuca, M., Fejes-Toth, K., Ferreira, P., Foissac, S., Fullwood, M. J., Gao, H., Gonzalez, D., Gordon, A., Gunawardena, H., Howald, C., Jha, S., Johnson, R., Kapranov, P., King, B., Kingswood, C., Luo, O. J., Park, E., Persaud, K., Preall, J. B., Ribeca, P., Risk, B., Robyr, D., Sammeth, M., Schaffer, L., See, L.-H., Shahab, A., Skancke, J., Suzuki, A. M., Takahashi, H., Tilgner, H., Trout, D., Walters, N., Wang, H., Wrobel, J., Yu, Y., Ruan, X., Hayashizaki, Y., Harrow, J., Gerstein, M., Hubbard, T., Reymond, A., Antonarakis, S. E., Hannon, G., Giddings, M. C., Ruan, Y., Wold, B., Carninci, P., Guigo, R., and Gingeras, T. R. (2012). Landscape of transcription in human cells. *Nature*, 489(7414):101–108.
- Dohm, J. C., Lottaz, C., Borodina, T., and Himmelbauer, H. (2007). Sharcgs, a fast and highly accurate short-read assembly algorithm for de novo genomic sequencing. *Genome Research*, 17(11):1697–1706.
- Dreyfus, S. E. (1969). An appraisal of some shortest-path algorithms. *Operations Research*, 17(3):395–412.
- Edmonds, J. and Johnson, E. L. (1970). Matching: a well-solved class of integer linear programs. In *Combinatorial structures and their applications*, pages 89–92.
- El-Amin and Al-Ghamdi (1993). An expert system for transmission line route selection. In *Int. Power Engineering Conf*, volume 2, pages 697–702. Nanyang Technol. Univ, Singapore.
- Eppstein, D. (1999). Finding the k shortest paths. *SIAM J. Comput.*, 28(2):652–673.
- Eppstein, D., Löffler, M., and Strash, D. (2010). Listing all maximal cliques in sparse graphs in near-optimal time. In Cheong, O., Chwa, K.-Y., and Park, K., editors, *ISAAC (1)*, volume 6506 of *Lecture Notes in Computer Science*, pages 403–414. Springer.
- Eppstein, D. and Strash, D. (2011). Listing all maximal cliques in large sparse real-world graphs. In Pardalos, P. M. and Rebennack, S., editors, *SEA*, volume 6630 of *Lecture Notes in Computer Science*, pages 364–375. Springer.
- Ferragina, P. and Manzini, G. (2005). Indexing compressed text. *J. ACM*, 52(4):552–581.
- Ferreira, R. A., Grossi, R., and Rizzi, R. (2011). Output-sensitive listing of bounded-size trees in undirected graphs. In *ESA*, pages 275–286.
- Flicek, P. and Birney, E. (2009). Sense from sequence reads: methods for alignment and assembly. *Nature Methods*, 6(11 Suppl):S6–S12.
- Fortune, S., Hopcroft, J. E., and Wyllie, J. (1980). The directed subgraph homeomorphism problem. *Theoretical Computer Science*, 10:111–121.

- Fukuda, K., Liebling, T. M., and Margot, F. (1997). Analysis of backtrack algorithms for listing all vertices and all faces of a convex polyhedron. *Comput. Geom.*, 8:1–12.
- Garey, M. R. and Johnson, D. S. (1979). *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, NY, USA.
- Grabherr, M. G., Haas, B. J., Yassour, M., Levin, J. Z., Thompson, D. A., Amit, I., Adiconis, X., Fan, L., Raychowdhury, R., Zeng, Q., Chen, Z., Mauceli, E., Hacohen, N., Gnirke, A., Rhind, N., di Palma, F., Birren, B. W., Nusbaum, C., Lindblad-Toh, K., Friedman, N., and Regev, A. (2011). Full-length transcriptome assembly from RNA-Seq data without a reference genome. *Nature Biotechnology*, 29(7):644–652.
- Grossi, R., Rizzi, R., Sacomoto, G., and Sagot, M.-F. (2014). Listing bounded length paths. In preparation.
- Gusfield, D., Eddhu, S., and Langley, C. (2004). Optimal, efficient reconstruction of phylogenetic networks with constrained recombination. *J. Bioinf. and Comput. Biol.*, 2(1):173–214.
- Guttman, M., Garber, M., Levin, J. Z., Donaghey, J., Robinson, J., Adiconis, X., Fan, L., Koziol, M. J., Gnirke, A., Nusbaum, C., Rinn, J. L., Lander, E. S., and Regev, A. (2010). Ab initio reconstruction of cell type-specific transcriptomes in mouse reveals the conserved multi-exonic structure of lincnas. *Nature Biotechnology*, 28(5):503–510.
- Halford, T. R. and Chugg, K. M. (2004). Enumerating and counting cycles in bipartite graphs. In *IEEE Communication Theory Workshop*.
- Hernandez, D., François, P., Farinelli, L., Østerås, M., and Schrenzel, J. (2008). De novo bacterial genome sequencing: Millions of very short reads assembled on a desktop computer. *Genome Research*, 18(5):802–809.
- Hershberger, J. and Suri, S. (2001). Vickrey prices and shortest paths: What is an edge worth? In *FOCS*, pages 252–259. IEEE Computer Society.
- Horváth, T., Gärtner, T., and Wrobel, S. (2004). Cyclic pattern kernels for predictive graph mining. In *Proc. of 10th ACM SIGKDD*, pages 158–167.
- Hossain, M. S., Azimi, N., and Skiena, S. (2009). Crystallizing short-read assemblies around seeds. *BMC Bioinformatics*, 10(S-1).
- Huang, X. and Madan, A. (1999). CAP3: A DNA Sequence Assembly Program. *Genome Research*, 9(9):868–877.
- Idury, R. M. and Waterman, M. S. (1995). A new algorithm for dna sequence assembly. *Journal of Computational Biology*, 2(2):291–306.
- Iqbal, Z., Caccamo, M., Turner, I., Flicek, P., and McVean, G. (2012). De novo assembly and genotyping of variants using colored de bruijn graphs. *Nature Genetics*.
- Jaffe, D. B., Butler, J., Gnerre, S., Mauceli, E., Lindblad-Toh, K., Mesirov, J. P., Zody, M. C., and Lander, E. S. (2003). Whole-genome sequence assembly for mammalian genomes: Arachne 2. *Genome Research*, 13(1):91–6.
- Johnson, D. B. (1975). Finding all the elementary circuits of a directed graph. *SIAM J. Comput.*, 4(1):77–84.

- Johnson, D. S., Papadimitriou, C. H., and Yannakakis, M. (1988). On generating all maximal independent sets. *Inf. Process. Lett.*, 27(3):119–123.
- Katoh, N., Ibaraki, T., and Mine, H. (1982). An efficient algorithm for K shortest simple paths. *Networks*, 12(4):411–427.
- Kececioglu, J. D. (1992). *Exact and approximation algorithms for DNA sequence reconstruction*. PhD thesis, University of Arizona, Tucson.
- Kent, W. J. (2002). BLAT—the BLAST-like alignment tool. *Genome Research*, 12(4):656–664.
- Kirsch, A. and Mitzenmacher, M. (2008). Less hashing, same performance: Building a better bloom filter. *Random Structures Algorithms*, 33(2):187–218.
- Klamt, S. and et al. (2006). A methodology for the structural and functional analysis of signaling and regulatory networks. *BMC Bioinformatics*, 7:56.
- Klamt, S. and von Kamp, A. (2009). Computing paths and cycles in biological interaction graphs. *BMC Bioinformatics*, 10:181.
- Knuth, D. E. (1998). *The Art of Computer Programming, Volume 3: Sorting and Searching*. Addison Wesley Longman Publishing Co., Inc., Redwood City, CA, USA.
- Koch, I. (2001). Enumerating all connected maximal common subgraphs in two graphs. *Theoretical Computer Science*, 250(1-2):1–30.
- Kuhn, R. M., Karolchik, D., Zweig, A. S., Wang, T., and et. al, K. E. S. (2009). The UCSC Genome Browser Database: update 2009. *Nucleic Acids Research*, 37(Database issue):D755–D761.
- Lander, E. S. and Waterman, M. S. (1988). Genomic mapping by fingerprinting random clones: a mathematical analysis. *Genomics*, 2(3):231–239.
- Lawler, E. L. (1972). A procedure for computing the K best solutions to discrete optimization problems and its application to the shortest path problem. *Management Science*, 18:401–405.
- Leggett, R. M., Ramirez-Gonzalez, R. H., Verweij, W., Kawashima, C. G., Iqbal, Z., Jones, J. D. G., Caccamo, M., and MacLean, D. (2013). Identifying and classifying trait linked polymorphisms in non-reference species by walking coloured de bruijn graphs. *PLoS ONE*, 8(3):e60058.
- Levin, J. Z., Yassour, M., Adiconis, X., Nusbaum, C., Thompson, D. A. A., Friedman, N., Gnirke, A., and Regev, A. (2010). Comprehensive comparative analysis of strand-specific RNA sequencing methods. *Nature methods*, 7(9):709–715.
- Li, C.-L., McCormick, S. T., and Simchi-Levi, D. (1990). The complexity of finding two disjoint paths with min-max objective function. *Discrete Applied Mathematics*.
- Li, R., Zhu, H., Ruan, J., Qian, W., Fang, X., Shi, Z., Li, Y., Li, S., Shan, G., Kristiansen, K., Li, S., Yang, H., Wang, J., and Wang, J. (2010). De novo assembly of human genomes with massively parallel short read sequencing. *Genome Research*, 20:265–272.

- Lister, R., O'Malley, R. C., Tonti-Filippini, J., Gregory, B. D., Berry, C. C., Millar, A. H., and Ecker, J. R. (2008). Highly integrated single-base resolution maps of the epigenome in *Arabidopsis*. *Cell*, 133(3):523–536.
- Liu, H. and Wang, J. (2006). A new way to enumerate cycles in graph. In *AICT and ICIW*, pages 57–59.
- Lodish, H., Berk, A., Zipursky, S. L., Matsudaira, P., Baltimore, D., and Darnell, J. (2000). *Molecular Cell Biology*. W. H. Freeman, 4 edition.
- Margulies, M., Egholm, M., Altman, W. E., Attiya, S., Bader, J. S., Bemben, L. A., Berka, J., Braverman, M. S., Chen, Y. J., Chen, Z., Dewell, S. B., Du, L., Fierro, J. M., Gomes, X. V., Godwin, B. C., He, W., Helgesen, S., Ho, C. H., Ho, C. H., Irzyk, G. P., Jando, S. C., Alenquer, M. L., Jarvie, T. P., Jirage, K. B., Kim, J. B., Knight, J. R., Lanza, J. R., Leamon, J. H., Lefkowitz, S. M., Lei, M., Li, J., Lohman, K. L., Lu, H., Makhijani, V. B., McDade, K. E., McKenna, M. P., Myers, E. W., Nickerson, E., Nobile, J. R., Plant, R., Puc, B. P., Ronan, M. T., Roth, G. T., Sarkis, G. J., Simons, J. F., Simpson, J. W., Srinivasan, M., Tartaro, K. R., Tomasz, A., Vogt, K. A., Volkmer, G. A., Wang, S. H., Wang, Y., Weiner, M. P., Yu, P., Begley, R. F., and Rothberg, J. M. (2005). Genome sequencing in microfabricated high-density picolitre reactors. *Nature*, 437(7057):376–380.
- Marino, A., Sacomoto, G., and Uno, T. (2014). Lecture notes on enumeration algorithms. In preparation.
- Martin, J. A. and Wang, Z. (2011). Next-generation transcriptome assembly. *Nature Reviews Genetics*, pages 1–12.
- Mateti, P. and Deo, N. (1976). On algorithms for enumerating all circuits of a graph. *SIAM J. Comput.*, 5(1):90–99.
- Medvedev, P., Georgiou, K., Myers, G., and Brudno, M. (2007). Computability of models for sequence assembly. In *WABI*, Lecture Notes in Computer Science, pages 289–301. Springer.
- Mezlini, A. M., Smith, E. J. M., Fiume, M., Buske, O., Savich, G. L., Shah, S., Aparicio, S., Chiang, D. Y., Goldenberg, A., and Brudno, M. (2012). iReckon: Simultaneous isoform discovery and abundance estimation from RNA-seq data. *Genome Research*, 23(3):519–529.
- Miller, J. R., Delcher, A. L., Koren, S., Venter, E., Walenz, B., Brownley, A., Johnson, J., Li, K., Mobarry, C. M., and Sutton, G. G. (2008). Aggressive assembly of pyrosequencing reads with mates. *Bioinformatics*, 24(24):2818–2824.
- Mitra, R. D. and Church, G. M. (1999). In situ localized amplification and contact replication of many individual dna molecules. *Nucleic Acids Research*, 27(24).
- Montgomery, S. B., Sammeth, M., Gutierrez-Arcelus, M., Lach, R. P., Ingle, C., Nisbett, J., Guigo, R., and Dermitzakis, E. T. (2010). Transcriptome genetics using second generation sequencing in a caucasian population. *Nature*, 464(7289):773–777.
- Mortazavi, A., Williams, B. A., McCue, K., Schaeffer, L., and Wold, B. (2008). Mapping and quantifying mammalian transcriptomes by rna-seq. *Nat Meth*, 5(7):621–628.
- Myers, E. W. (2005). The fragment assembly string graph. *Bioinformatics*, 21(2):79–85.

- Myers, E. W., Sutton, G. G., Delcher, A. L., Dew, I. M., Fasulo, D. P., Flanigan, M. J., Kravitz, S. A., Mobarry, C. M., Reinert, K. H., Remington, K. A., Anson, E. L., Bolanos, R. A., Chou, H. H., Jordan, C. M., Halpern, A. L., Lonardi, S., Beasley, E. M., Brandon, R. C., Chen, L., Dunn, P. J., Lai, Z., Liang, Y., Nusskern, D. R., Zhan, M., Zhang, Q., Zheng, X., Rubin, G. M., Adams, M. D., and Venter, J. C. (2000). A whole-genome assembly of drosophila. *Science*, 287(5461):2196–204.
- Nagalakshmi, U., Wang, Z., Waern, K., Shou, C., Raha, D., Gerstein, M., and Snyder, M. (2008). The transcriptional landscape of the yeast genome defined by rna sequencing. *Science*, 320:1344–1349.
- Namiki, T., Hachiya, T., Tanaka, H., and Sakakibara, Y. (2011). Metavelvet: An extension of velvet assembler to de novo metagenome assembly from short sequence reads. In *Proceedings of the 2Nd ACM Conference on Bioinformatics, Computational Biology and Biomedicine, BCB '11*, pages 116–124, New York, NY, USA. ACM.
- Nijkamp, J. F., Pop, M., Reinders, M. J., and de Ridder, D. (2013). Exploring variation-aware contig graphs for (comparative) metagenomics using marygold. *Bioinformatics*.
- Pan, Q., Shai, O., Lee, L. J., Frey, B. J., and Blencowe, B. J. (2008). Deep surveying of alternative splicing complexity in the human transcriptome by high-throughput sequencing. *Nature genetics*, 40(12):1413–1415.
- Pell, J., Hintze, A., Canino-Koning, R., Howe, A., Tiedje, J. M., and Brown, C. T. (2012). Scaling metagenome sequence assembly with probabilistic de Bruijn graphs. *Proc. Natl. Acad. Sci. U.S.A.*, 109(33):13272–13277.
- Peng, Y., Leung, H. C. M., Yiu, S.-M., and Chin, F. Y. L. (2010). Idba - a practical iterative de bruijn graph de novo assembler. In Berger, B., editor, *RECOMB*, volume 6044 of *Lecture Notes in Computer Science*, pages 426–440. Springer.
- Peng, Y., Leung, H. C. M., Yiu, S. M., and Chin, F. Y. L. (2011). Meta-IDBA: a de novo assembler for metagenomic data. *Bioinformatics*, 27(13):i94–i101.
- Peng, Y., Leung, H. C. M., Yiu, S.-M., Lv, M.-J., Zhu, X.-G., and Chin, F. Y. L. (2013). Idba-tran: a more robust de novo de bruijn graph assembler for transcriptomes with uneven expression levels. *Bioinformatics*, 29(13):326–334.
- Peng, Z., Cheng, Y., Tan, B. C., Kang, L., Tian, Z., Zhu, Y., Zhang, W., Liang, Y., Hu, X., Tan, X., Guo, J., Dong, Z., Liang, Y., Bao, L., and Wang, J. (2012). Comprehensive analysis of RNA-Seq data reveals extensive RNA editing in a human transcriptome. *Nature Biotechnology*, 30(3):253–260.
- Peterlongo, P. and Chikhi, R. (2012). Mapsembler, targeted and micro assembly of large ngs datasets on a desktop computer. *BMC Bioinformatics*, 13:48.
- Peterlongo, P., Schnel, N., Pisanti, N., Sagot, M.-F., and Lacroix, V. (2010). Identifying SNPs without a reference genome by comparing raw reads. In *SPIRE*, Springer LNCS 6393, pages 147–158.
- Pevzner, P. (1989). l -tuple DNA sequencing: computer analysis. *J. Biomol. Struct. Dyn.*, (7):63–73.

- Pevzner, P., Tang, H., and Tesler, G. (2004). *De novo* repeat classification and fragment assembly. In *RECOMB*, pages 213–222.
- Pevzner, P. A., Tang, H., and Waterman, M. S. (2001). An eulerian path approach to dna fragment assembly. *Proc Natl Acad Sci USA*, 98(17):9748–53.
- Pickrell, J. K., Marioni, J. C., Pai, A. A., Degner, J. F., Engelhardt, B. E., Nkadori, E., Veyrieras, J.-B., Stephens, M., Gilad, Y., and Pritchard, J. K. (2010). Understanding mechanisms underlying human gene expression variation with RNA sequencing. *Nature*, 464(7289):768–772.
- Ponstein, J. (1966). Self-avoiding paths and the adjacency matrix of a graph. *SIAM Journal on Applied Mathematics*, 14:600–609.
- Pop, M. (2009). Genome assembly reborn: recent computational challenges. *Briefings in Bioinformatics*, 10(4):354–366.
- Porat, E. (2009). An optimal Bloom filter replacement based on matrix solving. In *Computer Science - Theory and Applications, Fourth International Computer Science Symposium in Russia, CSR 2009, Novosibirsk, Russia, August 18-23, 2009. Proceedings*, volume 5675 of *Lecture Notes in Computer Science*, pages 263–273. Springer.
- Pruitt, K., Tatusova, T., Klimke, W., and Maglott, D. (2009). NCBI reference sequences: current status, policy and new initiatives. *Nucleic Acids Research*, 37(Database-Issue):32–36.
- Read, R. C. and Tarjan, R. E. (1975). Bounds on backtrack algorithms for listing cycles, paths, and spanning trees. *Networks*, 5(3):237–252.
- Richter, D. C., Ott, F., Auch, A. F., Schmid, R., and Huson, D. H. (2008). MetaSim: a sequencing simulator for genomics and metagenomics. *PLoS One*, 3(10):e3373.
- Rizk, G., Lavenier, D., and Chikhi, R. (2013). DSK: k-mer counting with very low memory usage. *Bioinformatics*.
- Roberts, A. and Pachter, L. (2013). Streaming fragment assignment for real-time analysis of sequencing experiments. *Nature Methods*, 10(1):71–73.
- Robertson, G., Schein, J., Chiu, R., Corbett, R., Field, M., Jackman, S. D., Mungall, K., Lee, S., Okada, H. M. M., Qian, J. Q., Griffith, M., Raymond, A., Thiessen, N., Cezard, T., Butterfield, Y. S., Newsome, R., Chan, S. K., She, R., Varhol, R., Kamoh, B., Prabhu, A.-L. L., Tam, A., Zhao, Y., Moore, R. A., Hirst, M., Marra, M. A., Jones, S. J., Hoodless, P. A., and Birol, I. (2010). De novo assembly and analysis of RNA-seq data. *Nature methods*, 7(11):909–912.
- Roditty, L. (2007). On the k-simple shortest paths problem in weighted directed graphs. In Bansal, N., Pruhs, K., and Stein, C., editors, *SODA*, pages 920–928. SIAM.
- Roditty, L. and Zwick, U. (2005). Replacement paths and k simple shortest paths in unweighted directed graphs. In Caires, L., Italiano, G., Monteiro, L., Palamidessi, C., and Yung, M., editors, *ICALP*, volume 3580 of *Lecture Notes in Computer Science*, pages 249–260. Springer Berlin Heidelberg.

- Sacomoto, G., Lacroix, V., and Sagot, M.-F. (2013). A polynomial delay algorithm for the enumeration of bubbles with length constraints in directed graphs and its application to the detection of alternative splicing in rna-seq data. In *WABI*, pages 99–111.
- Sacomoto, G. A. T., Kielbassa, J., Chikhi, R., Uricaru, R., Antoniou, P., Sagot, M.-F., Peterlongo, P., and Lacroix, V. (2012). Kissplice: de-novo calling alternative splicing events from rna-seq data. *BMC Bioinformatics*, 13(S-6):S5.
- Salikhov, K., Sacomoto, G., and Kucherov, G. (2013). Using cascading bloom filters to improve the memory usage for de bruijn graphs. In *WABI*, Lecture Notes in Computer Science, pages 364–376.
- Sammeth, M. (2009). Complete alternative splicing events are bubbles in splicing graphs. *Journal of Computational Biology*, 16(8):1117–1140.
- Sammeth, M., Foissac, S., and Guigó, R. (2008). A general definition and nomenclature for alternative splicing events. *PLoS Computational Biology*, 4(8):e1000147.
- Sanger, F., Nicklen, S., and Coulson, A. (1977). DNA sequencing with chain-terminating inhibitors. *Proc. Natl. Acad. Sci.*, 74.
- Sankar, K. and Sarad, A. (2007). A time and memory efficient way to enumerate cycles in a graph. In *Intelligent and Advanced Systems*, pages 498–500.
- Schott, R. and Staples, G. S. (2011). Complexity of counting cycles using Zeons. *Computers and Mathematics with Applications*, 62:1828–1837.
- Schulz, M. H. (2010). *Data Structures and Algorithms for Analysis of Alternative Splicing with RNA-Seq Data*. PhD thesis, Free University of Berlin.
- Schulz, M. H., Zerbino, D. R., Vingron, M., and Birney, E. (2012). Oases: robust de novo RNA-seq assembly across the dynamic range of expression levels. *Bioinformatics*, 28(8):1086–1092.
- Sedgewick, R. (2001). *Algorithms in c, part 5: graph algorithms, third edition*. Addison-Wesley Professional, third edition.
- Shendure, J. and Ji, H. (2008). Next-generation dna sequencing. *Nature Biotechnology*.
- Sherry, S. T., Ward, M. H., Kholodov, M., Baker, J., Phan, L., Smigielski, E. M., and Sirotkin, K. (2001). dbSNP: the NCBI database of genetic variation. *Nucleic Acids Research*, 29(1):308–311.
- Simpson, J., Wong, K., Jackman, S., Schein, J., Jones, S., and Birol, I. (2009). Abyss: a parallel assembler for short read sequence data. *Genome Research*, 19:1117.
- Simpson, J. T. and Durbin, R. (2012). Efficient de novo assembly of large genomes using compressed data structures. *Genome Research*, 22(3):549–556.
- Staden, R. (1979). A strategy of DNA sequencing employing computer programs. *Nucleic Acids Research*, 6:2601–2610.

- Sultan, M., Schulz, M. H., Richard, H., Magen, A., Klingenhoff, A., Scherf, M., Seifert, M., Borodina, T., Soldatov, A., Parkhomchuk, D., Schmidt, D., O’Keeffe, S., Haas, S., Vingron, M., Lehrach, H., and Yaspo, M.-L. (2008). A Global View of Gene Activity and Alternative Splicing by Deep Sequencing of the Human Transcriptome. *Science*, 321(5891):956–960.
- Sussenguth, E. (1965). A graph-theoretical algorithm for matching chemical structures. *J. Chem. Doc.*, 5:36–43.
- Syslo, M. M. (1981). An efficient cycle vector space algorithm for listing all cycles of a planar graph. *SIAM J. Comput.*, 10(4):797–808.
- Szwarcfiter, J. L. and Lauer, P. E. (1976). A search strategy for the elementary cycles of a directed graph. *BIT Numerical Mathematics*, 16.
- Tarjan, R. E. (1972). Depth-first search and linear graph algorithms. *SIAM Journal of Computing*, 1(2):146–160.
- Tarjan, R. E. (1973). Enumeration of the elementary circuits of a directed graph. *SIAM J. Comput.*, 2(3):211–216.
- Taylor, L. (2013). Phage assembly suite and tutorial (PHAST). <http://gcat.davidson.edu/phast/index.html>. Accessed: 2013-11-02.
- Tiernan, J. C. (1970). An efficient search algorithm to find the elementary circuits of a graph. *Communications ACM*, 13:722–726.
- Trapnell, C., Williams, B. A., Pertea, G., Mortazavi, A., Kwan, G., van Baren, M. J., Salzberg, S. L., Wold, B. J., and Pachter, L. (2010a). Transcript assembly and quantification by rna-seq reveals unannotated transcripts and isoform switching during cell differentiation. *Nature Biotechnology*, 28(5):511–515.
- Trapnell, C., Williams, B. A., Pertea, G., Mortazavi, A., Kwan, G., van Baren, M. J., Salzberg, S. L., Wold, B. J., and Pachter, L. (2010b). Transcript assembly and quantification by RNA-Seq reveals unannotated transcripts and isoform switching during cell differentiation. *Nature Biotechnology*, 28(5):511–515.
- Uno, T. (2001). A fast algorithm for enumerating bipartite perfect matchings. In *Proc. of the 12th International Symposium on Algorithms and Computation*, ISAAC’01, pages 367–379.
- Valiant, L. G. (1979). The complexity of computing the permanent. *Theoretical Computer Science*, 8:189–201.
- van Bakel, H., Nislow, C., Blencowe, B. J., and Hughes, T. R. (2010). Most “dark matter” transcripts are associated with known genes. *PLoS Biology*, 8(5).
- Wang, E., Sandberg, R., Luo, S., Khrebtkova, I., Zhang, L., Mayr, C., Kingsmore, S., Schroth, G. P., and Burge, C. (2008). Alternative isoform regulation in human tissue transcriptomes. *Nature*, 456(7221):470–476.
- Wang, Z., Gerstein, M., and Snyder, M. (2009). Rna-seq: a revolutionary tool for transcriptomics. *Nature Reviews Genetics*, 10(1).
- Warren, R. L., Sutton, G. G., Jones, S. J. M., and Holt, R. A. (2007). Assembling millions of short dna sequences using ssake. *Bioinformatics*, 23(4):500–501.

- Waterman, M. S. S. (1983). Sequence alignments in the neighborhood of the optimum with general application to dynamic programming. *Proceedings of National Academy of Science*, 80(10):3123–3124.
- Welch, Jr., J. T. (1966). A mechanical analysis of the cyclic structure of undirected linear graphs. *J. ACM*, 13:205–210.
- Wicker, T., Sabot, F., Hua-Van, A., Bennetzen, J. L., Capy, P., Chalhoub, B., Flavell, A., Leroy, P., Morgante, M., Panaud, O., Paux, E., SanMiguel, P., and Schulman, A. H. (2007). A unified classification system for eukaryotic transposable elements. *Nature Reviews Genetics*, 8(12):973–982.
- Wild, M. (2008). Generating all cycles, chordless cycles, and hamiltonian cycles with the principle of exclusion. *J. of Discrete Algorithms*, 6:93–102.
- Yang, L., Duff, M. O., Graveley, B. R., Carmichael, G. G., and Chen, L.-L. L. (2011). Genomewide characterization of non-polyadenylated RNAs. *Genome biology*, 12(2).
- Yau, S. (1967). Generation of all hamiltonian circuits, paths, and centers of a graph, and related problems. *IEEE Transactions on Circuit Theory*, 14:79–81.
- Ye, C., Ma, Z., Cannon, C., Pop, M., and Yu, D. (2012). Exploiting sparseness in de novo genome assembly. *BMC Bioinformatics*, 13(Suppl 6):S1.
- Yen, J. Y. (1971). Finding the K shortest loopless paths in a network. *Management Science*, 17:712–716.
- Zerbino, D. R. and Birney, E. (2008). Velvet: algorithms for de novo short read assembly using de Bruijn graphs. *Genome Research*, 18(5):821–829.

TITRE: Algorithmes efficaces pour l'assemblage de novo d'événements d'épissage alternatif dans des données de RNA-seq

RÉSUMÉ: Dans cette thèse, nous abordons le problème de l'identification et de la quantification de variants (épissage alternatif et polymorphisme génomique) dans des données de RNA-seq sans génome de référence, et sans faire un assemblage complet des transcripts. Basé sur l'idée que chaque variant correspond à une motif reconnaissable, qu'on appelle une bulle, dans un graphe de Bruijn construit à partir des lectures de RNA-seq, nous proposons un modèle pour les variants dans de tels graphes. Nous introduisons ensuite une méthode, appelé KisSplice, pour extraire les événements d'épissage alternatif, et nous montrons qu'il trouve plus d'événements corrects que les assembleurs de transcriptome traditionnels.

Afin d'améliorer son temps d'exécution, nous proposons un nouvel algorithme polynomial pour énumérer les bulles. On montre qu'il est plusieurs ordres de grandeur plus rapide que les approches précédentes. Afin de réduire sa consommation en mémoire, nous proposons une nouvelle façon de représenter un graphe de Bruijn. Nous montrons que notre approche utilise 30% à 40% moins de mémoire que l'état de l'art.

Nous appliquons les techniques développées pour énumérer les bulles à deux problèmes classiques. Nous donnons le premier algorithme optimal pour énumérer les cycles dans des graphes non orientés. Il s'agit de la première amélioration à ce problème en près de 40 ans. Nous considérons ensuite une variante du problème des K chemins plus courts: au lieu de limiter le nombre des chemins, nous limitons leurs poids. Nous présentons de nouveaux algorithmes qui utilisent exponentiellement moins mémoire que les approches précédentes.

MOTS-CLEFS: algorithme; complexité; énumération; génération; graphes; structure de données; exact exponentiel; biologie computationnelle; séquençage haut débit; NGS; RNA-seq; épissage alternatif; variante; SNP; insertion et délétion; assemblage; graphe de de Bruijn; analyse amortie; chaîne; cycle; bulle; filtre de Bloom

TITLE: Efficient Algorithms for de novo Assembly of Alternative Splicing Events from RNA-seq Data

ABSTRACT: In this thesis, we address the problem of identifying and quantifying variants (alternative splicing and genomic polymorphism) in RNA-seq data when no reference genome is available, without assembling the full transcripts. Based on the idea that each variant corresponds to a recognizable pattern, a bubble, in a de Bruijn graph constructed from the RNA-seq reads, we propose a general model for all variants in such graphs. We then introduce an exact method, called KisSplice, to extract alternative splicing events, and show that it outperforms general purpose transcriptome assemblers.

We put an extra effort to make KisSplice as scalable as possible. In order to improve its running time, we propose a new polynomial delay algorithm to enumerate bubbles. We show that it is several orders of magnitude faster than previous approaches. In order to reduce its memory consumption, we propose a new compact way to build and represent a de Bruijn graph. We show that our approach uses 30% to 40% less memory than the state of the art, with an insignificant impact on the construction time.

Additionally, we apply the techniques developed to list bubbles in two classical problems: cycle enumeration and the K-shortest paths problem. We give the first optimal algorithm to list cycles in undirected graphs, improving over Johnson's algorithm. This is the first improvement to this problem in almost 40 years. We then consider a different parameterization of the K-shortest (simple) paths problem: instead of bounding the number of st-paths, we bound the weight of the st-paths. We present new algorithms using exponentially less memory than previous approaches.

KEYWORDS: algorithm; complexity; enumeration; listing; graph; data structure; exact exponential; computational biology; high-throughput sequencing; NGS; RNA-seq; alternative splicing; variant; SNP; indel; assembly; de Bruijn graph; amortized analysis; path; cycle; bubble; Bloom filter