SpringerBriefs in Mathematics

Stasys Jukna

Tropical Circuit Complexity

Limits of Pure Dynamic Programming

Springer

SpringerBriefs in Mathematics

SpringerBriefs present concise summaries of cutting-edge research and practical applications across a wide spectrum of fields. Featuring compact volumes of 50 to 125 pages, the series covers a range of content from professional to academic. Briefs are characterized by fast, global electronic dissemination, standard publishing contracts, standardized manuscript preparation and formatting guidelines, and expedited production schedules.

Typical topics might include:

- A timely report of state-of-the art techniques
- A bridge between new research results, as published in journal articles, and a contextual literature review
- A snapshot of a hot or emerging topic
- An in-depth case study
- A presentation of core concepts that students must understand in order to make independent contributions

SpringerBriefs in Mathematics showcases expositions in all areas of mathematics and applied mathematics. Manuscripts presenting new results or a single new result in a classical field, new field, or an emerging topic, applications, or bridges between new results and already published works, are encouraged. The series is intended for mathematicians and applied mathematicians. All works are peer-reviewed to meet the highest standards of scientific literature.

Titles from this series are indexed by Scopus, Web of Science, Mathematical Reviews, and zbMATH.

Stasys Jukna

Tropical Circuit Complexity

Limits of Pure Dynamic Programming

Springer

Stasys Jukna
Faculty of Mathematics and Informatics
Vilnius University
Vilnius, Lithuania

ISSN 2191-8198 ISSN 2191-8201 (electronic)
SpringerBriefs in Mathematics
ISBN 978-3-031-42353-6 ISBN 978-3-031-42354-3 (eBook)
https://doi.org/10.1007/978-3-031-42354-3

Mathematics Subject Classification: 49L20; 68Q06; 68Q17; 90C39

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Paper in this product is recyclable.

Preface

Go to the roots of calculations! Group the operations. Classify them according to their complexities rather than their appearances! This, I believe, is the mission of future mathematicians.

–Evariste Galois

Understanding the power and weakness of algorithmic paradigms for solving decision or optimization problems in rigorous mathematical terms is an important long-term goal. Along with greedy and linear programming, dynamic programming (DP) is one of THE algorithmic paradigms for solving combinatorial optimization problems. Dynamic programming algorithms turned out to be quite powerful in many practical applications, so that we know what these algorithms *can* do. But what can DP algorithms *not* do (efficiently)? Answering this question is the subject of this book.

Roughly speaking, the idea of DP is to break up a given optimization problem into smaller subproblems in a divide-and-conquer manner and solve these subproblems recursively. Optimal solutions of smaller instances are found and retained for use in solving larger instances (smaller instances are never solved again). Many classical DP algorithms are *pure* in that they only apply the basic operations (min, $+$) or (max, $+$) in their recursion equations.

A rigorous mathematical model for pure DP algorithms is that of tropical circuits. These are conventional combinational circuits using (min, $+$) or (max, $+$) operations as gates. Pure DP algorithms are special (recursively constructed) tropical circuits. So, if one can prove that *any* tropical circuit solving a given optimization problem must use at least t gates, then we know that *no* pure DP algorithm can solve this problem by performing fewer than t (min, $+$) or (max, $+$) operations, be the designer of an algorithm even omnipotent. Thanks to the rigorous combinatorial nature of tropical circuits, ideas and arguments from the Boolean and arithmetic circuit complexity can be exploited to obtain lower bounds for topical circuits and, hence, also for pure DP algorithms.

For example, the classical Bellman–Held–Karp DP algorithm gives a tropical (min, $+$) circuit with about $n^2 2^n$ gates solving the travelling salesman problem on n-vertex graphs, while a trivial brute force algorithm results in about $n! \approx (n/\mathrm{e})^n$

gates. On the other hand, Jerrum and Snir in 1982 have shown that at least about $n^2 2^n$ gates are also necessary in any (min, +) circuit solving this problem. This shows that the Bellman–Held–Karp DP algorithm is *optimal* among all pure DP algorithms for this problem. The tropical (min, +) circuit corresponding to the (also classical) Floyd–Warshall–Roy pure DP algorithm for the all-pairs shortest paths problem on n-vertex graphs uses about n^3 gates. On the other hand, already in 1970, Kerr has shown that at least about n^3 gates are also necessary for this problem. So, the Floyd–Warshall–Roy pure DP algorithm is also optimal in the class of all pure DP algorithms.

After these and several other impressing lower bounds where obtained, a long break followed. Only in recent years, and mainly due to recognized connection with dynamic programming, tropical circuits have attracted growing attention again. The goal of this book is to survey the lower-bound ideas and methods that emerged during these last years.

We focus on presenting the lower-bound arguments themselves, rather than on quantitative bounds achieved using them. That is, the focus is on the proof arguments, on the ideas behind them. Because of a very pragmatic motivation of tropical circuits—their intimate relation to dynamic programming—the primary goal is to create as large as possible "toolbox" for proving lower bounds on the size of tropical circuits, not relying on unproven complexity assumptions like $\mathbf{P} \neq \mathbf{NP}$.

The difficulty in proving that a given optimization problem requires large tropical circuits lies in the nature of our adversary: the circuit. Small circuits may work in a counterintuitive fashion, using deep, devious, and fiendishly clever ideas. How can one prove that there is no clever way to quickly solve the problem? In this book, we will learn some tools to defeat this adversary.

Tropical algebra and geometry—where "adding" numbers means to take their minimum or maximum, and "multiplying" them means to add them—are now actively studied topics in mathematics. Tropical circuit complexity adds a computational complexity aspect to this topic.

The book is self-contained and is meant to be approachable already by graduate students in mathematics and computer science. The text assumes certain mathematical maturity (minor knowledge of basic concepts in graph theory, discrete probability, and linear algebra) but *no* special knowledge in the theory of computing or dynamic programming.

Supplementary material to the book can be found on my home page.

Vilnius, Lithuania/Frankfurt, Germany — Stasys Jukna
June 2023

Contents

Notation

We will use more or less standard concepts and notation. For ease of reference, let us collect some of most often used ones right now:

Nonnegative real numbers	$\mathbb{R}_+ = \{x \in \mathbb{R} \colon x \geqslant 0\}$
Nonnegative integers	$\mathbb{N} = \{0, 1, 2, \ldots\}$ and $[n] = \{1, \ldots, n\}$
K_n	The complete graph on $[n]$
$K_{n,n}$	A complete bipartite $n \times n$ graph
2^X for a set X	Family of all subsets of X
$\lvert X \rvert$ for a finite set X	Number of elements in X
Family $\mathcal{F} \subseteq 2^X$ is uniform	All sets in $\mathcal{F}$ have the same cardinality
Characteristic vector of $S \subseteq [n]$	Vector $a \in \{0, 1\}^n$ with $a_i = 1$ iff $i \in S$
Unit vector $\vec{e}_i$	$\vec{e}_i = (0, \ldots, 0, 1, 0, \ldots, 0)$ with 1 in the ith position
$a \leqslant b$ for $a, b \in \mathbb{R}^n$	$a_i \leqslant b_i$ for all $i = 1, \ldots, n$
$A \subseteq \mathbb{R}^n$ is an antichain	$a \not\leqslant b$ for all $a \neq b \in A$
Upward closure $A^{\uparrow}$ of $A \subseteq \mathbb{N}^n$	$A^{\uparrow} = \{b \in \mathbb{N}^n : b \geqslant a$ for some $a \in A\}$
Downward closure $A^{\downarrow}$ of $A \subseteq \mathbb{N}^n$	$A^{\downarrow} = \{b \in \mathbb{N}^n : b \leqslant a$ for some $a \in A\}$
$B \subseteq \mathbb{R}^n$ lies above $A \subseteq \mathbb{R}^n$	$B \subseteq A^{\uparrow}$, i.e., $\forall b \in B\ \exists a \in A :\ b \geqslant a$
$B \subseteq \mathbb{R}^n$ lies below $A \subseteq \mathbb{R}^n$	$B \subseteq A^{\downarrow}$, i.e., $\forall b \in B\ \exists a \in A :\ b \leqslant a$
Support of $a \in \mathbb{R}^n$	$\sup(a) = \{i : a_i \neq 0\}$
Degree of $a \in \mathbb{N}^n$	$\lvert a \rvert = a_1 + \cdots + a_n$
Lower envelope of $A \subseteq \mathbb{N}^n$	$\lfloor A \rfloor = \{a \in A \colon \lvert a \rvert$ is minimal$\}$
Higher envelope of $A \subseteq \mathbb{N}^n$	$\lceil A \rceil = \{a \in A \colon \lvert a \rvert$ is maximal$\}$
$A \subseteq \mathbb{N}^n$ is homogeneous	$\lceil A \rceil = \lfloor A \rfloor$
Sum of $a, b \in \mathbb{R}^n$	$a + b = (a_1 + b_1, \ldots, a_n + b_n)$
Minkowski sum of $A, B \subseteq \mathbb{R}^n$	$A + B = \{a + b \colon a \in A, b \in B\}$
Scalar product of $a, b \in \mathbb{R}^n$	$\langle a, b \rangle = a_1 b_1 + \cdots + a_n b_n$
Tropical (min, +) polynomial	$f(x) = \min_{a \in A}\{\langle a, x \rangle + c_a\};\ A \subseteq \mathbb{N}^n,\ c_a \in \mathbb{R}_+$

Chapter 1
Basics

Abstract In this chapter, we recall the models of arithmetic $(+, \times)$, Boolean $(\vee, \wedge)$, and tropical $(\min, +)$ and $(\max, +)$ circuits, introduce Minkowski $(\cup, +)$ circuits as a model taking all of them "under one hat," establish the basic structural properties of tropical polynomials, and relate the corresponding circuit complexity measures. The main message of this chapter is that: lower bounds on the tropical circuit complexity of optimization problems can be obtained by proving lower bounds on the monotone arithmetic circuit complexity of particular polynomials.

1.1 What Is This Book About?

We are interested in solving discrete optimization problems[1] $f : \mathbb{R}_+^n \to \mathbb{R}_+$ on given sets $A \subseteq \mathbb{N}^n$ of *feasible solutions*:

$$f(x) = \min_{a \in A} \sum_{i=1}^{n} a_i x_i \quad \text{or} \quad f(x) = \max_{a \in A} \sum_{i=1}^{n} a_i x_i \,. \tag{1.1}$$

If $A \subseteq \{0, 1\}^n$, then f is usually called a *combinatorial optimization* or $0/1$ *optimization* problem. The set $A \subseteq \mathbb{N}^n$ of feasible solutions can be described either explicitly (as a particular set of vectors), or as the set $A = \{a \in \mathbb{N}^n : Ma \leqslant b\}$ of nonnegative integer or $0/1$ solutions of a given system of linear inequalities (as in linear programming), or by other means. It is only important that the set A does not depend on the input weightings $x \in \mathbb{R}_+^n$.

For example, in the $0/1$ optimization problem, known as the shortest s-t path problem on a given graph G, the set A of feasible solutions consists of characteristic 0-1 vectors[2] of all paths in G between two vertices s and t, the paths being viewed as sets of their edges. In the minimum weight spanning tree problem, feasible solutions

[1] In what follows, $\mathbb{N} = \{0, 1, 2, \ldots\}$ stands for the set of all nonnegative integers, $[n] = \{1, \ldots, n\}$ for the set of the first n positive integers, and $\mathbb{R}_+$ for the set of all nonnegative real numbers.

[2] The *characteristic* 0-1 *vector* of a set $S \subseteq [n]$ is the vector $a \in \{0, 1\}^n$ such that $a_i = 1$ iff $i \in S$.

S. Jukna, *Tropical Circuit Complexity*, SpringerBriefs in Mathematics,
https://doi.org/10.1007/978-3-031-42354-3_1

are (characteristic 0-1 vectors of) spanning trees of a given graph, and the problem is to compute the minimum weight of such a tree. In the assignment problem, we deal with perfect matchings in a complete bipartite graph, etc.

Note that min, max, and $+$ are the only operations used to formulate the problems (1.1) themselves. It is easy to see that the optimization problem on *every* set $A \subseteq \{0,1\}^n$ of feasible solutions *can* be solved using at most $n|A|$ $(\min,+)$ or $(\max,+)$ operations: compute all $|A|$ sums $\sum_{i=1}^n a_i x_i$ for $a \in A$, and take their minimum or maximum using additional $|A|-1$ min or max operations. But, as Examples 1.6 to 1.8 in Sect. 1.4 show, this trivial *upper* bound can be very far from the truth. For example, if $A \subseteq \{0,1\}^n$ consists of all $|A| = \binom{n}{n/2}$ vectors with exactly $n/2$ ones, then the minimization problem on A is, given an input weighting $x \in \mathbb{R}_+^n$, to compute the sum of lightest $n/2$ weights. Although there are $|A| \geqslant 2^{n/2}$ feasible solutions, the problem can be solved by using at most $\mathcal{O}(n^2)$ $(\min,+)$ operations (Example 1.6). The main goal of this book is to learn how to prove *lower* bounds:

> $(*)$ *At least* how many $(\min,+)$ or $(\max,+)$ operations do we need to solve or to approximate a given discrete optimization problem?

For example, in the *lightest triangle problem*, inputs are assignments of nonnegative real weights to the edges of the complete graph K_n on $\{1,\dots,n\}$ and the goal is to compute the minimum weight of a triangle. How many $(\min,+)$ operations do we need to solve this problem (for all possible input weightings)? Since there are only $\binom{n}{3}$ triangles, the problem can be trivially solved using a cubic number $\mathcal{O}(n^3)$ of $(\min,+)$ operations. More interesting, however, is the question: does a cubic number $\Omega(n^3)$ of operations is also *necessary* to solve this problem? (Yes, Corollary 2.4.)

With a wish to make the question $(*)$ mathematically precise, we arrive to the classical model of *circuits* (also called *combinational circuits*). We will mainly be interested in tropical circuits, that is, in circuits over tropical[3] semirings $(\mathbb{R}_+,\min,+)$ and $(\mathbb{R}_+,\max,+)$. But, as we will see, the power of tropical circuits is related to that of circuits over the Boolean semiring $(\{0,1\},\vee,\wedge)$ as well as over the arithmetic semiring $(\mathbb{R}_+,+,\times)$. So, let us first recall what "circuits" over a semiring actually are.

[3] The adjective "tropical" is not to contrast with "polar geometry." It was coined by French mathematicians in honor of Imre Simon who lived in Sao Paulo (south tropic). Tropical algebra and tropical geometry are now intensively studied topics in mathematics.

1.2 Circuits

A (commutative) *semiring* $(R, \oplus, \odot)$ consists of a set R closed under two associative and commutative binary operations "addition" $x \oplus y$ and "multiplication" $x \odot y$, where multiplication distributes over addition: $x \odot (y \oplus z) = (x \odot y) \oplus (x \odot z)$. That is, in a semiring, we can "add" and "multiply" elements, but neither "subtraction" nor "division" is necessarily possible. A semiring is *additively idempotent* if $x \oplus x = x$ holds and is *multiplicatively idempotent* if $x \odot x = x$ holds for all elements $x \in R$. A semiring may (or may not) contain an additive neutral element $\mathbb{0} \in R$ satisfying $\mathbb{0} \oplus x = x \oplus \mathbb{0} = x$. We will only assume that the semiring contains a multiplicative neutral element $\mathbb{1} \in R$ such that $\mathbb{1} \odot x = x \odot \mathbb{1} = x$.

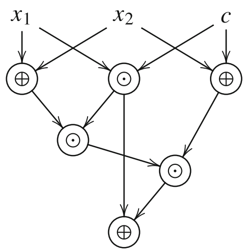

A *circuit* Φ (also known as a *combinational circuit*) over a semiring $(R, \oplus, \odot)$ is a directed acyclic graph; parallel edges joining the same pair of nodes are allowed. Each indegree-zero node (a *source* node) holds either one of the variables $x_1, \dots, x_n$ or a semiring element $c \in R$; if there are no semiring elements $c \in R$ other than $c = \mathbb{1}$ as inputs, then the circuit is called *constant-free*. Every other node, a *gate*, has indegree two and performs one of the semiring operations $\oplus$ or $\odot$ on the values computed at the two gates entering this gate. Usually (but not always), one of the gates is declared as the output gate. The *size* of a circuit Φ, denoted size(Φ), is the total number of gates in it. A circuit Φ *computes* a function $f : R^n \to R$ if $\Phi(x) = f(x)$ holds for all $x \in R^n$.

Proposition 1.1 *Over any semiring, there are at most* $2^s(2s + n + 1)^{2s}$ *distinct constant-free circuits* $\Phi(x_1, \dots, x_n)$ *with at most* s *gates.*

Proof Each gate in such a circuit is assigned a semiring operation (two choices) and acts on some two previous nodes. Each previous node can be either a previous gate (at most s choices) or an input variable (n choices) or the "constant" $\mathbb{1}$. Thus, each single gate has at most $N = 2(2s + n + 1)^2$ choices, and the number of choices for a circuit is at most N^s. □

In this book, we will consider circuits over the following four semirings $(R, \oplus, \odot)$: the arithmetic semiring $(\mathbb{R}_+, +, \times)$ with usual (arithmetic) addition and multiplication, the tropical $(\min, +)$ semiring $(\mathbb{R}_+, \min, +)$, the tropical $(\max, +)$ semiring $(\mathbb{R}_+, \max, +)$, and the Boolean $(\vee, \wedge)$ semiring $(\{0, 1\}, \vee, \wedge)$. That is, we will consider the following types of circuits: