Theoretical Computer Science 590 (2015) 17–26

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

Cancellation-free circuits in unbounded and bounded depth ☆

Joan Boyar [1], Magnus Gausdal Find *,[2],

Department of Mathematics and Computer Science, University of Southern Denmark, Denmark

ARTICLE INFO

Article history:
Received 7 November 2013
Received in revised form 23 June 2014
Accepted 1 October 2014
Available online 8 October 2014

Keywords:
Circuit complexity
Cancellation-free
Linear circuits

ABSTRACT

We study the notion of "cancellation-free" circuits. This is a restriction of XOR circuits, but can be considered as being equivalent to previously studied models of computation. The notion was coined by Boyar and Peralta in a study of heuristics for a particular circuit minimization problem. They asked how large a gap there can be between the smallest cancellation-free circuit and the smallest XOR circuit. We present a new proof showing that the difference can be a factor $\Omega(n/\log^2 n)$. Furthermore, our proof holds for circuits of constant depth. We also study the complexity of computing the Sierpinski matrix using cancellation-free circuits and give a tight $\Omega(n\log(n))$ lower bound.

1. Introduction

Let $\mathbb{F}_2$ be the field of order 2, and let $\mathbb{F}_2^n$ be the n-dimensional vector space over $\mathbb{F}_2$. For $n \in \mathbb{N}$, we let $[n] = \{1, \ldots, n\}$. A Boolean function $f : \mathbb{F}_2^n \to \mathbb{F}_2^m$ is said to be linear if there exists a Boolean $m \times n$ matrix A such that $f(\mathbf{x}) = A\mathbf{x}$ for every $\mathbf{x} \in \mathbb{F}_2^n$. This is equivalent to saying that f can be computed using only XOR gates.

An *XOR circuit* (or a *linear circuit*) C is a directed acyclic graph. There are n nodes with in-degree 0, called the *inputs*. All other nodes have in-degree 2 and are called *gates*. There are m nodes which are called the *outputs*; these are labeled $y_1, \ldots, y_m$. The value of a gate is the sum of its two children (addition in $\mathbb{F}_2$, denoted $\oplus$). The circuit C, with inputs $\mathbf{x} = (x_1, \ldots, x_n)$, *computes* the $m \times n$ matrix A if the output vector computed by C, $\mathbf{y} = (y_1, \ldots, y_m)$, satisfies $\mathbf{y} = A\mathbf{x}$. In other words, output y_i is defined by the ith row of the matrix. The *size* of a circuit C is the number of gates in C. The *depth* is the number of gates on a longest directed path from an input to an output. For simplicity, we will let $m = n$ unless otherwise explicitly stated. For a matrix A, let $|A|$ be the number of nonzero entries in A.

Our contributions In this paper we deal with a restriction of XOR circuits called *cancellation-free* circuits, coined in [2], where the authors noticed that many heuristics for finding small XOR circuits always produce cancellation-free XOR circuits. They asked the question of how large a separation there can be between these two models. Recently, Gashkov and Sergeev [3] showed that the work of Grinchuk and Sergeev [4] implied a separation of $\Omega(\frac{n}{\log^6 n \log\log n})$. An improved separation of $\Omega(\frac{n}{\log^2 n})$ follows from Lemma 4.1 and Lemma 4.2 in [5], although this implied separation was not published until

☆ A preliminary version of this paper appears in [1]. Both authors are partially supported by the Danish Council for Independent Research, Natural Sciences (grant number DFF-1323-00247).

* Corresponding author.
E-mail addresses: joan@imada.sdu.dk (J. Boyar), magnusgf@imada.sdu.dk (M.G. Find).

[1] Part of this work was done while visiting the University of Waterloo.

[2] Part of this work was done while visiting the University of Toronto.

http://dx.doi.org/10.1016/j.tcs.2014.10.014
0304-3975/

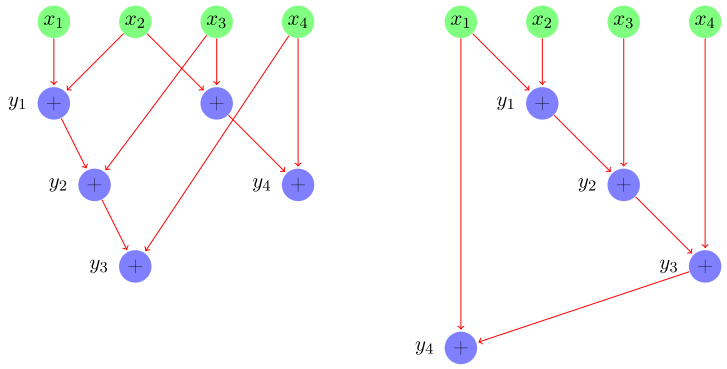

Fig. 1. Two circuits computing the matrix A. The circuit on the left is cancellation-free, and has size 5 – one more than the circuit on to the right.

recently [6]. We present an alternative proof of the same separation. Our proof is based on a different construction and uses communication complexity in a novel way that might have independent interest. Like the separation implied in the work [6], but unlike the separations demonstrated in [3,7], our separation holds even in the case of constant depth circuits. We conclude that many heuristics for finding XOR circuits do not approximate better than a factor of $\Theta(\frac{n}{\log^2 n})$ of the optimal. We also study the complexity of computing the Sierpinski matrix using cancellation-free circuits. We show that the complexity is exactly $\frac{1}{2}n\log n$. Furthermore, our proof holds for OR circuits. As a corollary to this we obtain an explicit matrix where the smallest OR circuit is a factor $\Theta(\log n)$ larger than the smallest OR circuit for its complement.

We also study the complexity of computing the *Sierpinski matrix* (described later), and show a tight $\frac{1}{2}n\log n$ lower bound for OR circuits and cancellation-free circuits. This results follows implicitly from the work of Kennes [8], however our proof is simpler and more direct. Also we hope that our proof can be strengthened to give an $\omega(n)$ lower bound for XOR circuits for the Sierpinski matrix. A similar lower bound was shown independently by Selezneva in [9,10].

2. Cancellation-free XOR circuits

For XOR circuits, the value computed by every gate is the parity of a subset of the n variables. That is, the output of every gate u can be considered as a vector $\kappa(u)$ in the vector space $\mathbb{F}_2^n$, where $\kappa(u)_i = 1$ if and only if x_i is a term in the parity function computed by the gate u. We call $\kappa(u)$ the *value vector* of u, and for input variables define $\kappa(x_i) = e^{(i)}$, the unit vector having the ith coordinate 1 and all others 0. It is clear by definition that if a gate u has the two children w, t, then $\kappa(u) = \kappa(w) \oplus \kappa(t)$, where $\oplus$ denotes coordinate-wise addition in $\mathbb{F}_2$. We say that an XOR circuit is *cancellation-free* if for every pair of gates u, w where u is an ancestor of w, then $\kappa(u) \geq \kappa(w)$, where $\geq$ denotes the usual coordinate-wise partial order. These are also called SUM circuits in [7,6].

If this is satisfied, the circuit never exploits the $\mathbb{F}_2$-identity, $a \oplus a = 0$, so things do not "cancel out" in the circuit.

Although it is not hard to see that cancellation-free circuits is equivalent to addition chains [11,12] and "ensemble computations" [13], we stick to the term "cancellation-free", since we will think of it as a special case of XOR circuits.

For a simple example demonstrating that cancellation-free circuits indeed are less powerful than general XOR circuits, consider the matrix

$$A = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 \end{pmatrix}.$$

In Fig. 1, two circuits computing the matrix A are shown, the circuit on the right uses cancellations, and the circuit on the left is cancellation-free, and has one gate more. For this particular matrix, any cancellation-free circuit must use at least 5 gates.

A different, but tightly related kind of circuits is OR circuits. The definition is exactly the same as for XOR circuits, but with $\vee$ (logical OR) instead of $\oplus$, see [14,6,13]. Cancellation-free circuits is a special case of OR circuits and every cancellation-free circuit can be interpreted as an OR circuit for the same matrix, as well as an XOR circuit.

For a matrix A, we will let $C_{\oplus}(A)$, $C_{CF}(A)$, $C_{\vee}(A)$ denote the smallest XOR circuit, the smallest cancellation-free circuit and the smallest OR circuit computing the matrix A.

By the discussion above, the following is immediate:

Proposition 1. *For every matrix, A, $C_{\vee}(A) \leq C_{CF}(A)$.*

This means in particular that any lower bound for OR circuits carries over to a lower bound for cancellation-free circuits. However, the converse does not hold in general [7]. A simple example showing this is the matrix