Discrete Applied Mathematics 179 (2014) 193–200

Contents lists available at ScienceDirect

Discrete Applied Mathematics

journal homepage: www.elsevier.com/locate/dam

On the maximum length of coil-in-the-box codes in dimension 8

Patric R.J. Östergård, Ville H. Pettersson *

Department of Communications and Networking, Aalto University School of Electrical Engineering, P.O. Box 13000, 00076 Aalto, Finland

ARTICLE INFO

Article history:
Received 6 February 2013
Received in revised form 27 June 2014
Accepted 20 July 2014
Available online 12 August 2014

Keywords:
Canonical augmentation
Induced cycle
Circuit code
Coil-in-the-box code
Snake-in-the-box code

ABSTRACT

The coil-in-the-box problem asks for a simple chordless cycle of maximum length in the n-cube. Such cycles are also known as n-dimensional spread 2 circuit codes, or n-coils. This problem has been solved earlier for $n \le 7$. An approach based on canonical augmentation is here used to solve the problem for $n = 8$ and show that the maximum length of a chordless cycle in the 8-cube is 96. Several new 8-coils of length 96 are presented.

1. Introduction

A path (resp. cycle) in a graph G is chordless if every two vertices in the path (cycle) that are not successive have a distance of at least 2 in G. An n-snake is a simple chordless path in an n-cube. An n-coil is a simple chordless cycle in an n-cube. An n-coil is also known as an n-dimensional spread 2 circuit code [12].

The problem of finding long coils and snakes in an n-cube was originally described by Kautz in 1958 [6]. Since then, the problem has attracted much attention in the research community, partly because long coils and snakes have some applications related to coding. For example, an n-coil corresponds to a cyclic gray code with certain error-detection capabilities, whose applications include analog to digital conversion [12]. One of the most recent examples is rank modulation [18].

So far the exact length of the longest n-snake and n-coil has been determined for $n \le 7$. For $n \ge 8$, only upper and lower bounds were previously known. Lower bounds are obtained constructively, and the techniques used for small n have included exhaustive and non-exhaustive search with pruning [4,15], nested Monte-Carlo search [7], evolutionary methods [1], reduction to SAT [20], gluing together smaller snakes [17], and methods based on symmetries and repetition [17,12]. We refer to [7] for an overview of these earlier methods. Several asymptotic upper and lower bounds have also been published [14,16,19], the best asymptotic upper bound for the length (number of edges) of a coil currently being [19]

$$2^{n-1}(1 - 1/(89\sqrt{n}) + O(1/n)).$$

This asymptotic bound is not very tight for small n.

For general graphs the problem of determining whether a given graph contains a chordless cycle of length at least k is known to be NP-complete [2,3].

* Corresponding author. Tel.: +358 504626499; fax: +358 947022474.
E-mail addresses: patric.ostergard@aalto.fi (P.R.J. Östergård), ville.pettersson@aalto.fi (V.H. Pettersson).

http://dx.doi.org/10.1016/j.dam.2014.07.010
0166-218X/

We refer to the length of the longest n-coil as $f(n)$. Previously the values of $f(n)$ have been determined for $n \leq 7$, and are 2, 4, 6, 8, 14, 26, and 48 (sequence A000937 in OEIS [13]). In the current work we extend this list and show that $f(8) = 96$. It was already known that 96 is a lower bound for $f(8)$ (see [12]). Earlier only one coil of length 96 has been published. In the current work we present several new coils of length 96.

We use an approach based on canonical augmentation to exhaustively generate (or show the nonexistence of) all n-coils of length greater than a given threshold. Our method relies on the observation that by the pigeonhole principle some $(n-1)$-dimensional part of a long n-coil obtained by fixing a coordinate will contain few disjoint snakes and many vertices. This observation can be applied recursively to the sets of disjoint snakes in smaller dimensions. We apply our method to the case $n = 8$ and show that there are no 8-coils with more than 96 edges.

The paper is organized as follows. In Section 2 we give some necessary background information on n-cubes and n-coils. In Section 3 we describe the canonical augmentation method and how it can be applied to the search of n-coils. In Section 4 we derive upper and lower bounds that help restrict the search space, and in Section 5 we give some details on how to implement different parts of the method. In Section 6 we show how the results of our search can be verified with double counting, and finally we present the results in Section 7. Some 8-coils of length 96 are listed in Table 5.

2. Preliminaries

The n-cube is a graph that has one vertex for each binary word of length n and an edge between a pair of vertices if the Hamming distance between the corresponding binary words is 1. A binary *code* of length n is a subset of binary words of length n. The Hamming distance between two codewords is equal to the distance between the corresponding vertices in the n-cube. Sometimes it is convenient to think of a code as a subset of vertices in the n-cube, and vice versa.

An automorphism in the automorphism group Γ_n of the n-cube consists of a permutation of coordinates followed by addition of a constant vector to the corresponding binary word, giving $|\Gamma_n| = 2^n n!$. Two n-coils are said to be *equivalent* if there is an automorphism in Γ_n that maps one to the other. The equivalence of n-coils can be studied in the framework of equivalence of codes [5, Section 2.3].

The binary words that correspond to two successive vertices in a coil differ in exactly one coordinate, so a coil can be represented as a list that shows coordinate changes between successive vertices. Such a list is called a *coordinate sequence* (or *transition sequence*). Assume that the vertices of the coil are $s_0 s_1 s_2 \ldots s_{k-1}$. In the coordinate sequence $t_0 t_1 \ldots t_{k-1}$, the value of t_i is the coordinate in which the binary words for vertices s_i and s_{i+1} differ (indices taken modulo k). An n-coil is said to be *degenerate*, if there is a dimension that does not appear in its coordinate sequence. Throughout this paper we focus on non-degenerate coils, since degenerate coils cannot be of maximum length.

Note that the length of an n-coil has to be an even number, since in order to leave a vertex and return back to it each coordinate has to occur an even number of times in the coordinate sequence.

We define an n-dimensional *snake collection* to be a generalized version of an n-snake that, instead of containing just one snake, can contain multiple snakes. In more detail, it is a subgraph of the n-cube that consists of chordless disjoint paths, such that each path is an n-snake, and the distance between any two vertices in different paths is at least 2. The motivation behind this definition is that we can create an $(n-1)$-dimensional snake collection or a *half* from an (non-degenerate) n-coil by considering the induced subgraph obtained by fixing some value v for some coordinate i. This subgraph consists of the nodes of the n-coil whose labels have value v in coordinate i and of the edges between these nodes. Note that if a snake collection is created in this way from a coil with more than 4 edges, it will only contain snakes with at least two edges, since a snake of length 1 would force a chord in the other $(n-1)$-dimensional half of the coil, and a snake of length 0 would have degree at most 1 in the original coil. These properties only hold for halves; in general, a snake collection can contain snakes with fewer than 2 edges.

3. Canonical augmentation

Canonical augmentation was originally introduced by McKay [10] as a method for recursively generating one representative from each equivalence class of a set of objects. In canonical augmentation it is required that objects be generated "in a canonical way". An overview of the method can be found in [5, Section 4.2.3].

Canonical augmentation works as follows. In the recursive search, each generated object X has a sequence of parents $p(X), p(p(X)), \ldots$. For each object X, we define a canonical parent $w(X)$ which fulfils $X \cong Y \Longrightarrow w(X) \cong w(Y)$. An object is rejected in the recursive search if it was not generated from its canonical parent, i.e., if $w(X) \ncong p(X)$. An object is also rejected if it is equivalent to some other object generated previously from the same parent. It can be shown that with this method exactly one representative from each equivalence class will be generated.

One benefit of using canonical augmentation is that no further isomorph rejection is needed after the algorithm finishes. Another benefit is a possibility for validating the results with double counting techniques, which are described in more detail in Section 6.

Canonical augmentation can be applied to (non-degenerate) coils by using an n-level recursive search that builds n-coils by extending partial coils, i.e., snake collections, in smaller dimensions. At levels $i < n$ of the recursive search i-dimensional snake collections are built by extending $(i-1)$-dimensional snake collections, and at level n complete n-coils are formed