Journal of Computer and System Sciences 85 (2017) 74–92

Contents lists available at ScienceDirect

Journal of Computer and System Sciences

www.elsevier.com/locate/jcss

Combinatorial filter reduction: Special cases, approximation, and fixed-parameter tractability

Fatemeh Zahra Saberifar [a], Ali Mohades [a,*], Mohammadreza Razzazi [b], Jason M. O'Kane [c]

[a] *Department of Mathematics and Computer Science, Amirkabir University of Technology, 424 Hafez Ave, Tehran, Iran*
[b] *Computer Engineering and IT Department, Amirkabir University of Technology, 424 Hafez Ave, Tehran, Iran*
[c] *Department of Computer Science and Engineering, University of South Carolina, 315 Main St, Columbia, SC, USA*

ARTICLE INFO

Article history:
Received 25 April 2016
Received in revised form 9 November 2016
Accepted 11 November 2016
Available online 22 November 2016

Keywords:
Combinatorial filters
Reduction
Automata
Approximation
Fixed-parameter

ABSTRACT

Recent research in algorithmic robotics considers *combinatorial filters*, which concisely capture the discrete structure underlying many reasoning problems for robots. An important recent result is that the *filter minimization problem*—Given a filter, find the smallest equivalent filter—is NP-hard. This paper extends that result along several dimensions, including hardness proofs for some natural special cases and for approximation, and new results analyzing the only known algorithm for this problem. We show that this problem is not fixed-parameter tractable for any of the obvious parameters, but it is fixed-parameter tractable for a certain combination of new parameters.

1. Introduction

This paper continues the thread of research initiated by LaValle [1,2] concerning *combinatorial filtering*. The idea is to model systems (such as robots or sensor networks) that must process and draw conclusions from a sequence of discrete sensor readings.

Fig. 1 shows a well-known example. In this environment, with two moving agents and three beams sensors, the task is to determine whether both agents are in the same region at each time. A naive solution requires considering $2^9 = 512$ possible states to solve the task, but Tovar, Cohen, Bobadilla, Czarnowski and LaValle [3] described a (hand-crafted) optimal combinatorial filter with just 4 states for this problem. See Fig. 1.

Part of the motivation for studying combinatorial filters can be understood by analogy to the probabilistic filters [5] commonly used by roboticists. Though such filters must, in the general case, be able to represent arbitrarily complicated probability densities, many systems behave in sufficiently structured ways that simpler representations can be used without loss of accuracy. For example, in a linear system with Gaussian noise, the resulting density is simply a Gaussian whose mean and covariance can be computed using the Kalman filter [6]. In the same way, optimized combinatorial filters can be viewed as a specialization of the general class of nondeterministic filters [1], which reason about possible states of the underlying system. If the number of states in such a filter can be reduced without altering the filter's behavior, then the result can be useful for understanding the elements of structure underlying the problem.

* Corresponding author.
E-mail addresses: fz.saberifar@aut.ac.ir (F.Z. Saberifar), mohades@aut.ac.ir (A. Mohades), razzazi@aut.ac.ir (M. Razzazi), jokane@cse.sc.edu (J.M. O'Kane).

http://dx.doi.org/10.1016/j.jcss.2016.11.007

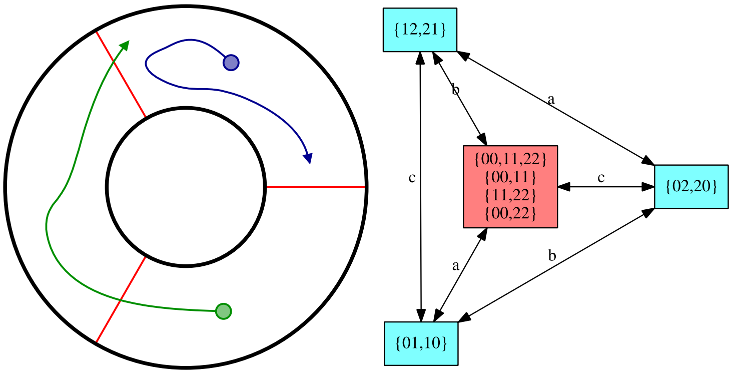

Fig. 1. (O'Kane and Shell [4]) [left] Two agents move through an environment having three beam sensors. [right] An optimal combinatorial filter reproduced by O'Kane and Shell's algorithm [4], for solving the task of whether the agents are in the same region at each time. The numbers 0, 1, and 2 denote the regions, and the letters a, b, and c denote observations of the three beam sensors. The shading of each state indicates the filter's output: red indicates that the agents are in the same region and blue indicates that the agents are not in the same region. (For interpretation of the references to color in this figure, the reader is referred to the web version of this article.)

A common feature of most of the early work on combinatorial filters is the need for human input to design and minimize the filters. Hence, O'Kane and Shell [4] worked on an automated method for filter reduction. Specifically, they consider the problem in which the input is a combinatorial filter expressed as a transition graph, and the goal is to produce the smallest reduced filter, measured by the number of states, equivalent to the given filter. They proved that this *filter minimization problem* is NP-hard and presented an efficient heuristic algorithm for it. This algorithm was able to reproduce the optimal filter for the problem in Fig. 1, automatically.

In spite of this progress, there remain some open problems in this area.

1. The filter minimization problem so far has been proved NP-hard only in the general case. However, this result seems very fragile, in the sense that there may exist some common special cases of it that are solvable efficiently.
2. In spite of the hardness of solving the problem optimally, there may still exist approximation algorithms, either for the general problem or for special cases. That is, perhaps we can guarantee to find, in polynomial time, solution that is within a factor α of optimal.
3. The only algorithm for this problem is presented by O'Kane and Shell. However, their analysis of the algorithm is incomplete. In particular, they conjecture, but do not prove, that if a certain graph coloring subroutine in that algorithm is performed optimally, then the resulting filter will also be optimally reduced.
4. Another way to attack this problem is to identify parameters that capture the difficulty of the problem, independent of problem size. This kind of analysis arises from the field of parameterized complexity theory. At present, nothing is known about which parameters govern the hardness of filter minimization.

The contribution of this paper is to provide some solutions for each of these four families of open problems.

1. We show that several natural special cases of the problem are also NP-hard. We also prove that a few special narrow cases belong to P. Table 1 shows a brief preview of these results.
2. We show that the filter minimization problem is NP-hard to approximate, even for the hard special cases mentioned above.
3. We show that O'Kane and Shell's conjecture—that their algorithm produces optimally reduced output filters if their graph coloring subroutine is optimal—is false, by providing a counterexample for which optimal coloring in their algorithm does not necessarily produce optimal filter reduction.
4. We show that for some natural kinds of parameters, the problem is not fixed-parameter tractable. We also introduce two new parameters and show, by describing an algorithm, that the problem is fixed-parameter tractable when it is parameterized by these two parameters.

Taken together, these results support a conclusion that filter minimization is a truly challenging problem, because it is not susceptible to any of the standard approaches for solving NP-hard problems. In practice, solving the filter minimization problem appears to require accepting at least one of three unpalatable options: (a) suboptimal heuristic solutions, (b) potentially exponential run time, or (c) restriction to a highly-constrained special case.

The balance of this paper is structured as follows. Section 2 reviews the related work and Section 3 introduces the basic definitions and the problem formulation. Section 4 presents a technique for converting the instances of the graph coloring problem to instances of the filter minimization problem, and proves several useful properties of this conversion that we reuse in some other sections. In Section 5, we prove hardness results for tree, bipartite, and planar filter reduction,