Computational Geometry 61 (2017) 1–23

Contents lists available at ScienceDirect

Computational Geometry: Theory and Applications

www.elsevier.com/locate/comgeo

Approximation algorithms for Max Morse Matching ☆

Abhishek Rathod *, Talha Bin Masood, Vijay Natarajan

Department of Computer Science and Automation, Indian Institute of Science, Bangalore, India

A R T I C L E I N F O

Article history:
Received 29 April 2016
Accepted 18 October 2016
Available online 24 October 2016

Keywords:
Discrete Morse theory
Computational topology
Approximation algorithms
Homology computation

A B S T R A C T

In this paper, we prove that the Max Morse Matching Problem is approximable, thus resolving an open problem posed by Joswig and Pfetsch [1]. For D-dimensional simplicial complexes, we obtain a $(D+1)/(D^2+D+1)$-factor approximation ratio using a simple edge reorientation algorithm that removes cycles. For $D \geq 5$, we describe a $2/D$-factor approximation algorithm for simplicial manifolds by processing the simplices in increasing order of dimension. This algorithm leads to $1/2$-factor approximation for 3-manifolds and $4/9$-factor approximation for 4-manifolds. This algorithm may also be applied to non-manifolds resulting in a $1/(D+1)$-factor approximation ratio. One application of these algorithms is towards efficient homology computation of simplicial complexes. Experiments using a prototype implementation on several datasets indicate that the algorithm computes near optimal results.

1. Introduction

Discrete Morse theory is a combinatorial analogue of Morse theory that is applicable to cell complexes [2]. It has become a popular tool in computational topology and visualization communities [3,4] and is actively studied in algebraic, geometric, and topological combinatorics [5,6].

The idea of using discrete Morse theory to speedup homology [7], persistent homology [8] and multidimensional persistence computations [9] hinges on the fact that discrete Morse theory helps reduce the problem of computing homology groups on an input simplicial complex to computing homology groups on a collapsed cell complex. Ideally, if one were to compute a discrete gradient vector field with minimum number of critical simplices (unmatched vertices in the Hasse graph) or maximum number of regular simplices (matched Hasse graph vertices), then the time required for computing homology over the collapsed cell complex would be very small. However, finding a vector field with maximum number of gradient pairs is an NP-hard problem as observed by Lewiner [10] and Joswig et al. [1] by showing a reduction from *the collapsibility problem* introduced by Eğecioğlu and Gonzalez in [11]. We study efficient approximations to the maximum number of gradient pairs in a discrete gradient vector field.

Computing the homology groups has several applications, particularly in material sciences, imaging, pattern classification and computer assisted proofs in dynamics [12]. More recently, homology and persistent homology have been appraised to be a more widely applicable computational invariant of topological spaces, arising from practical data sets of interest [13]. An approximate Morse matching designed using the algorithms described in this paper may be used to compute homology efficiently. One of the primary motivations for this work was that a previous study [7] involving discrete Morse theory

☆ The work is partially supported by the Department of Science and Technology, India under grant SR/S3/EECE/0086/2012.

* Corresponding author.
E-mail addresses: abhishek@jcrathod.in (A. Rathod), tbmasood@csa.iisc.ernet.in (T. Bin Masood), vijayn@csa.iisc.ernet.in (V. Natarajan).

http://dx.doi.org/10.1016/j.comgeo.2016.10.002
0925-7721/

in homology computation reported noteworthy speedup over existing methods. Their method used a modification of the coreduction heuristic [14] to construct discrete Morse functions. We start with a twin goal in mind — first to introduce rigor into the study by developing algorithms with approximation bounds and then to have a practical implementation that achieves nearly optimal solutions.

1.1. Max Morse Matching Problem

The Max Morse Matching Problem (MMMP) can be described as follows: Consider the Hasse graph $\mathcal{H}_{\mathcal{K}}$ of a simplicial complex $\mathcal{K}$ whose edges are all directed from a simplex to its lower dimensional facets. Associate a matching induced orientation to $\mathcal{H}_{\mathcal{K}}$ such that the resulting oriented graph $\overline{\mathcal{H}_{\mathcal{K}}}$ is acyclic. The goal is to maximize the cardinality of matched (regular) nodes. Equivalently, the goal is to maximize the number of gradient pairs. The approximate version of Max Morse Matching Problem seeks an algorithm that computes a Morse matching whose cardinality is within a factor α of the optimal solution for every instance of the problem.

1.2. Prior work

Joswig et al. [1] established the NP-completeness of Morse Matching Problem. They also posed the approximability of Max Morse Matching as an open problem pg. 6 Sec. 4 [1]. Several followup efforts seek optimality of Morse matchings either by restricting the problem to 2-manifolds or by applying heuristics [1,7,15–19]. Recently, Burton et al. [20] developed an FPT algorithm for designing optimal Morse functions.

1.3. Summary of results

We describe a $(D+1)/(D^2+D+1)$-factor approximation algorithm for Max Morse Matching Problem on D-dimensional simplicial complexes. This algorithm uses maximum-cardinality bipartite matching on the Hasse graph $\mathcal{H}_{\mathcal{K}}$ to orient it. We then use a BFS-like traversal of the oriented Hasse graph $\overline{\mathcal{H}_{\mathcal{K}}}$ to classify matching edges as either forward edges if they do not introduce cycles or backward edges if they do. We use a counting argument to prove an approximation bound that holds for manifold as well as non-manifold complexes.

For simplicial D-manifolds, we propose two approximation algorithms. The first approximation algorithm provides a ratio of $2/(D+1)$, for $D \geq 3$. The ratio is improved to $2/D$, for $D \geq 5$, via a refinement that specifies the order in which the graph is processed. It leads to $1/2$-factor and $4/9$-factor approximations for 3-manifolds and 4-manifolds, respectively. Both algorithms process simplices of lowest dimension first followed by higher dimensions in increasing order. Every d-dimensional simplex is first given the opportunity to match to a $(d-1)$-dimensional simplex. If unsuccessful, it is then given the option of matching to a $(d+1)$-dimensional simplex. Furthermore, both algorithms employ optimal algorithms for designing gradient fields for 0-dimensional and D-dimensional simplices (in the case of manifolds). The refinement processes subgraphs with small vertex degree with higher priority and hence achieves a better approximation ratio.

We provide evidence of practical utility of our algorithms through an extensive series of computational experiments.

2. Background

2.1. Discrete Morse theory

Our focus in this paper is limited to simplicial complexes and hence we restrict the discussion of Forman's Morse theory below to simplicial complexes. Please refer to [21] for a compelling expository introduction.

Definition 1. A simplicial complex $\mathcal{K}$ is a finite collection of simplices that satisfies the following conditions:

1. A face of a simplex in $\mathcal{K}$ also belongs to $\mathcal{K}$.
2. The intersection of two simplices $\sigma_1, \sigma_2 \in \mathcal{K}$ is either empty or a face of both σ_1 and σ_2.

Let $\mathcal{K}$ be a simplicial complex, and let σ^d, τ^{d-1} be simplices[1] of $\mathcal{K}$. The relation $\prec$ is defined as: $\tau \prec \sigma \Leftrightarrow \{\tau \subset \sigma \text{ and } \dim\tau = \dim\sigma - 1\}$. Alternatively, we say that τ is the *facet* of σ and σ is a *cofacet* of τ. The boundary $bd(\sigma)$ and the coboundary $cbd(\sigma)$ of a simplex are defined as: $bd(\sigma) = \{\tau | \tau \prec \sigma\}$ and $cbd(\sigma) = \{\rho | \sigma \prec \rho\}$. A function $f : \mathcal{K} \to \mathbb{R}$ is called a *discrete Morse function* if it assigns higher values to cofacets, with at most one exception at each simplex. Specifically, for a function $f : \mathcal{K} \to \mathbb{R}$ for every $\sigma \in \mathcal{K}$, let $\mathcal{N}_1(\sigma) = \#\{\rho \in cbd(\sigma) | f(\rho) \leq f(\sigma)\}$ and $\mathcal{N}_2(\sigma) = \#\{\tau \in bd(\sigma) | f(\tau) \geq f(\sigma)\}$. Such a function is called a discrete Morse function if for every $\sigma \in \mathcal{K}$, $\mathcal{N}_1(\sigma) + \mathcal{N}_2(\sigma) \leq 1$. If $\mathcal{N}_1(\sigma) = \mathcal{N}_2(\sigma) = 0$, then the simplex σ is *critical*, else it is *regular*.

[1] A d-dimensional simplex σ^d may be denoted either as σ or σ^d depending on whether we wish to emphasize its dimension.