Discrete Applied Mathematics 205 (2016) 27–34

Contents lists available at ScienceDirect

Discrete Applied Mathematics

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

Eccentricity, center and radius computations on the cover graphs of distributive lattices with applications to stable matchings ☆

Christine T. Cheng [a,*], Eric McDermid [b], Ichiro Suzuki [a]

[a] *Department of Computer Science, University of Wisconsin–Milwaukee, United States*

[b] *Apple, Austin, TX, United States*

ARTICLE INFO

Article history:
Received 20 November 2014
Received in revised form 16 November 2015
Accepted 18 November 2015
Available online 23 December 2015

Keywords:
Distributive lattices
Center
Radius
Eccentricity
Stable matchings

ABSTRACT

Birkhoff's fundamental theorem on distributive lattices states that for every distributive lattice $\mathcal{L}$ there is a poset $\mathcal{P}_{\mathcal{L}}$ whose lattice of down-sets is order-isomorphic to $\mathcal{L}$. Let $G(\mathcal{L})$ denote the cover graph of $\mathcal{L}$. In this paper, we consider the following problems: suppose we are simply given $\mathcal{P}_{\mathcal{L}}$. How do we compute the eccentricity of an element of $\mathcal{L}$ in $G(\mathcal{L})$? How about a center and the radius of $G(\mathcal{L})$? While eccentricity, center and radius computations have long been studied for various classes of graphs, our problems are different in that we are not given the graph explicitly; instead, we only have a structure that implicitly describes the graph. By making use of the comparability graph of $\mathcal{P}_{\mathcal{L}}$, we show that all the said problems can be solved efficiently. One of the implications of these results is that a center stable matching, a kind of fair stable matching, can be computed in polynomial time.

1. Introduction

A *finite distributive lattice* $\mathcal{L} = (L, \leq)$ is a partially ordered set (poset) where for any two elements $x, y \in L$, (i) their *meet* $x \wedge y$ or greatest lower bound exists, (ii) their *join* $x \vee y$ or least upper bound exists, and (iii) the meet and join operators distribute over each other. The *bottom element* of $\mathcal{L}$ is the unique element $\hat{0}$ such that $\hat{0} \leq x$ for all $x \in L$ while its *top element* is the unique element $\hat{1}$ such that $x \leq \hat{1}$ for all $x \in L$. As an example, consider the factors of a positive integer z ordered according to the divisibility relation. For any two factors f_1 and f_2, their meet is their greatest common factor, their join is their least common multiple while the bottom and top elements of the lattice are 1 and z respectively. Many more objects form a distributive lattice including the domino tilings of a polygon, the perfect matchings of a bipartite planar graph, alternating sign matrices, etc. (see [25,10] and references therein).

Let $\mathcal{P} = (P, \leq)$ be a poset. A subset P' of P is a *down-set* or *order ideal* of $\mathcal{P}$ if whenever $p \in P'$, all the predecessors of p in $\mathcal{P}$ are also in P'. Let $D(\mathcal{P})$ consist of the down-sets of $\mathcal{P}$. It is not difficult to see that $(D(\mathcal{P}), \subseteq)$ is a distributive lattice. An important characterization of distributive lattices states that the converse is true as well. (See Fig. 1 for an example.)

Theorem 1 (*Birkhoff [3]*). *For every distributive lattice $\mathcal{L}$, there is (up to isomorphism) a unique poset $\mathcal{P}_{\mathcal{L}}$ such that $(D(\mathcal{P}_{\mathcal{L}}), \subseteq)$ is order-isomorphic to $\mathcal{L}$.*

☆ A preliminary version of this paper appeared in the Proceedings of ICALP, 2011.

* Corresponding author.
E-mail addresses: ccheng@uwm.edu (C.T. Cheng), em4617@gmail.com (E. McDermid), suzuki@uwm.edu (I. Suzuki).

http://dx.doi.org/10.1016/j.dam.2015.11.016

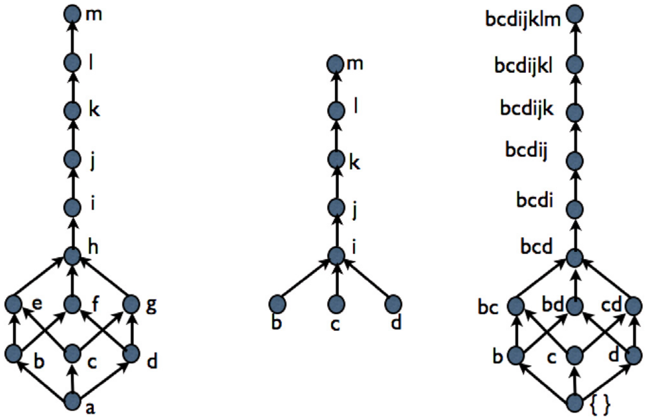

Fig. 1. Consider the distributive lattice $\mathcal{L}$ on the left. The poset $\mathcal{P}_{\mathcal{L}}$ that encodes $\mathcal{L}$ is in the middle. In Birkhoff's proof, $\mathcal{P}_{\mathcal{L}}$ is the subposet induced by the *join-irreducible* elements of $\mathcal{L}$ — i.e. the elements whose in-degree is 1 in the Hasse diagram of $\mathcal{L}$. The down-sets of $\mathcal{P}_{\mathcal{L}}$ (labeled without the curly braces) ordered according to the subset relation are shown on the right. Clearly, $\mathcal{L}$ and $(D(\mathcal{P}_{\mathcal{L}}), \subseteq)$ are isomorphic distributive lattices.

Given a graph $G = (V, E)$, let $d(v, u)$ denote the distance between v and u in G. The *eccentricity* of v, $\text{ecc}(v)$, is equal to $\max_u d(v, u)$. The *radius* of G, $\text{rad}(G)$, is equal to $\min_v \text{ecc}(v)$ while the *diameter* of G, $\text{diam}(G)$, is equal to $\max_v \text{ecc}(v)$. A node whose eccentricity is equal to the radius of G is referred to as a *center* of G. A node that has the smallest total (or average) distance from all other nodes of G is called a *median* of G. To illustrate some of these concepts, consider the popular game *Six Degrees of Kevin Bacon*. A person is given the name of an actor[1] a, and the goal is to identify a sequence of at most seven actors starting with a and ending with Kevin Bacon so that any two consecutive actors in the sequence have appeared in a movie together; i.e., actor a can "reach" Kevin Bacon by six steps or less. Graph theorists have long known that this is just a game on the *Actors' graph* where actors that have appeared in a movie are the vertices, and two actors are adjacent if and only if they have appeared in a movie together. An inherent assumption of the game is that the radius of the graph is at most six, and Kevin Bacon is one of its centers.

For a distributive lattice $\mathcal{L}$, let $G(\mathcal{L})$ denote the *cover graph* of $\mathcal{L}$, the undirected Hasse diagram of $\mathcal{L}$. In this paper, we consider the following problems: *given $\mathcal{P}_{\mathcal{L}}$, is there an efficient algorithm for finding a center of $G(\mathcal{L})$, expressed as a down-set of $\mathcal{P}_{\mathcal{L}}$? How about computing the radius of $G(\mathcal{L})$? Suppose we are additionally given an element of $\mathcal{L}$ expressed as a down-set of $\mathcal{P}_{\mathcal{L}}$, can we compute this element's eccentricity efficiently?*

The problems of computing eccentricities, radii, diameters and centers of graphs have a long and rich history starting with Jordan's theorem in 1869 [19] which states that a tree either has one center or two centers that are adjacent to each other. It is easy to compute the parameters in $O(mn)$ time, where n is the number of vertices of the graph and m is the number of edges, by running breadth-first search (BFS) from each node of the graph. Hence, most researchers' goal is to beat this brute force method. Seidel's [28] and Chan's [6] all-pairs shortest path algorithms for dense and sparse graphs respectively show that these parameters can be computed in $o(mn)$ time. Others like Roditty et al. [26] designed faster 3/2-approximation algorithms for computing the parameters in sparse graphs. They also provided evidence that it is unlikely that their diameter approximation result can be improved. For specific families of graphs, Corneil, Dragan and others [8,9] showed that Lex-BFS (a variant of BFS) can either compute the diameter of a graph exactly or within one if the graph is chordal, interval, AT-free, etc. On the other hand, Borassi et al. [5] developed heuristics that involve a small but related runs of BFS to determine the diameter of real-world networks. We note though that our problems are quite different from the ones considered in these papers because the graph under consideration is not given to us *explicitly*. Instead, we only have an auxiliary structure that encodes the graph so we have to rely on a different set of techniques to solve the problems.

Motivation. Our interest in finding a center of $G(\mathcal{L})$ given $\mathcal{P}_{\mathcal{L}}$ originated from our work on stable matchings. An instance I of the *stable marriage problem* (SM) has n men and n women each of whom has a preference list that ranks members of the opposite gender in a linear order. A *matching* is a set of n disjoint man–woman pairs; it is *stable* if there is no man–woman pair who prefer each other over their partners in the matching. The goal of the problem is to find a stable matching of I if one exists. A seminal result of Gale and Shapley in the 1960s [12] states that *every* SM instance has a stable matching that can be computed in $O(n^2)$ time. Today, centralized stable matching algorithms are used to match medical residents to hospitals [27] and students to schools [1,2].

An SM instance can have up to $2^{O(n)}$ stable matchings [14]. It turns out, however, that the Gale–Shapley algorithm outputs only two kinds: the *man-optimal/woman-pessimal* stable matching and the *woman-optimal/man-pessimal* stable matching. In the man-optimal stable matching, every man is matched to his best partner in all of the stable matchings while simultaneously every woman is matched to her worst partner in all of the stable matchings; the woman-optimal/man-pessimal stable matching has the opposite properties. Hence, in spite of the fact that the Gale–Shapley algorithm solves

[1] Although we use the word "actor", the person can be male or female.