Computers & Operations Research 40 (2013) 2387–2397

Contents lists available at SciVerse ScienceDirect

Computers & Operations Research

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

A parallel hybrid greedy branch and bound scheme for the maximum distance-2 matching problem

Ioannis T. Christou*, Spyridon Vassilaras

Athens Information Technology 19 km. Markopoulou Ave. PO Box 68, Paiania 19002, Greece

ARTICLE INFO

Available online 29 April 2013

Keywords:
Maximum distance-2 matching
Optimization methods
GRASP
Branch and bound
Parallel algorithms
Wireless ad-hoc networks
Capacity

ABSTRACT

We present a new highly parallel algorithm for fast determination of near-optimal solutions to the NP-hard problem of identifying a *maximum distance-2 matching* in arbitrary graphs. This problem, known as D2EMIS, has important applications such as determining the maximum capacity of the media access (MAC) layer in wireless ad-hoc networks [1]. It can also be seen as a *maximum 2-packing* problem [2] on the edge-to-vertex dual graph of the original graph. Our algorithm extends the GRASP [3] philosophy in that partial solutions are constructed by adding in a greedy adaptive manner the “best” nodes that can be found; however, when there are multiple alternatives that can be selected in an iteration, the algorithm branches into as many paths as there are (greedy) alternatives. The algorithm, using appropriate bounds to prune partial solutions that cannot be optimal, produces very fast near-optimal solutions that compare very well against other distributed algorithms and random greedy heuristics proposed before or variants thereof, or exact methods (Integer Programming or Maximum Satisfiability state-of-the-art solvers).

1. Introduction

The maximum distance-2 matching and its associated maximum 2-packing problems (also known as *largest* 2-packing) are well studied problems in graph theory with classical applications in optimal facility location, secure communications, VLSI design and network flow problems (see [4,5] for some examples). Given a graph $G(V,E)$, we define as the distance between two nodes $u,v \in V$ the number$d_v(u,v)$ of edges in any shortest path in $G(V,E)$ between the two nodes. We also define the distance $d_e(e_1,e_2)$ between two edges $e_1,e_2 \in E$ to be the number of nodes in any shortest path in G that joins the two edges e_1,e_2. A set $S_V \subseteq V$ is called a k-packing if $d_v(x,y) > k$ for all pairs of distinct nodes $x,y \in S_V$. A *maximum k-packing* is a k-packing with the maximum cardinality of all k-packings in $G(V,E)$. Similarly, a set $S_E \subseteq E$ is a distance-k matching if $d_e(e_1,e_2) > k$ for all pairs of distinct edges $e_1,e_2 \in S_E$. A maximum k-distance matching is a distance-k matching with the maximum cardinality of all distance-k matchings in $G(V,E)$. Note that any distance-k matching problem can be transformed to an equivalent k-packing problem by considering the edge-to-vertex dual graph[1] $D(V',E')$ where every node in V' corresponds to an edge in E, and an edge in E' connects two nodes in V' if and only if the corresponding edges in E have a common node in $G(V,E)$. However the reverse is not always possible, i.e., not every k-packing problem can be transformed to an equivalent distance-k matching problem. Indeed, given a graph G there always exist its line graph D, but given any graph D with maximum node degree >2 there is no graph G whose line graph is D because an edge has exactly two incident nodes while a node with node degree >2 has more than two incident edges.

These two problems have received renewed interest recently as they have found application in determining the maximum number of concurrent one hop transmissions in an arbitrary ad-hoc network under common MAC layer protocols (such as the MAC layer protocol in the IEEE 802.11 family of standards). More specifically, consider a set of wireless nodes with known positions on the plane and transmission ranges that form an ad hoc network subject to concurrent transmission constraints imposed by an 802.11 MAC layer protocol in the RTS/CTS mode. The problem of determining a set of links with maximum cardinality that can be active simultaneously without causing harmful interference to each other depends on the network topology itself and can be formulated as a maximum distance-2 matching problem (also known as the *maximum induced matching problem* or the *D2EMIS* problem) [1]. This number of maximum concurrent transmissions can be thought of as the network capacity under no traffic flow constraints. A random graph version of the same problem (i.e., determining the average number of maximum concurrent transmissions for a given probability distribution of nodes placement) is more relevant in mobile ad hoc networks (MANETs) in which nodes move around and the topology of the network changes in

* Corresponding author. Tel.: +30 210 668 2725.
E-mail addresses: ichr@ait.edu.gr (I.T. Christou), svas@ait.edu.gr (S. Vassilaras).

[1] also commonly known as the *line graph.*

0305-0548/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.cor.2013.04.009

a random way: In this case, the random graph average determines the MAC layer capacity of a MANET with the same marginal distribution of node positions, symmetric traffic flows (each node is a source for and destination to exactly one flow), a two-hop packet relaying scheme (first proposed in [6]) and unbounded end-to-end packet delivery delay, as explained in [7]. In this context, another interesting problem is to determine the transmission range which maximizes the average number of maximum concurrent transmissions ([1,7]).

The problem of determining a maximum distance-k matching is known to be NP-hard for $k \geq 2$ [1] (a polynomial time algorithm known as Edmonds' algorithm exists for $k=1$ in generic graphs). On the other hand the maximum k-packing problem is NP-hard for all k in the general case. For the D2EMIS problem, a simple random approximation heuristic by which "free" edges (i.e., edges that are not at distance ≤2 from any other already selected edge) are randomly chosen and added into the induced matching S_E until no more edges can be added, has a guaranteed performance approximation of $1/[2(\Delta-1)]$ where Δ is the maximum node degree of the graph [8] (i.e., the cardinality of the resulting set is larger or equal than $1/[2(\Delta-1)]$ times the cardinality of the maximum induced matching).

In this paper, we describe a highly parallel distributed algorithm for obtaining a near optimal solution to the maximum 2-packing problem and therefore for the D2EMIS problem that is especially suitable to multi-core processors and networks of multi-core computers. We show through extensive experimentation that the true optimal cardinality of S_E can be much higher than that produced by the random greedy heuristic algorithm. We compare the proposed algorithm against a weighted Maximum Satisfiability (max-sat) solver [9] by formulating the problem as a weighted max-sat problem and we show that it compares very well against such an approach. Our algorithm also compares very well against an implementation of a self-stabilizing distributed algorithm for finding maximal 2-packing solutions; a maximal 2-packing solution is one such that no more nodes can be added to it without violating any constraints. Comparisons with the state-of-the-art open-source MIP solver SCIP [9,10] reveal that the algorithm can find the true optimal solution or near-optimal solutions in a fraction of the time required by the solver to obtain the optimal solution. Experiments with the currently incumbent commercial MIP solver Gurobi © show that for large graph instances, our algorithm can quickly find near-optimal solutions whereas Gurobi fails as the computer's installed RAM is not sufficient for the solver.

The remainder of this paper is organized as follows: in Section 2 we discuss related work, and in Section 3 we provide a number of different models for the maximum distance-2 problem. In Section 4 we present the design of our algorithm and the prototype system implementation details, and in Section 5 we present the computational results from our approach and make comparisons with existing algorithms. Finally, in Section 6 we discuss the conclusions of our findings.

2. Related work

The maximum distance-k matching and largest k-packing problems described in the introduction have been investigated in the literature under a number of different names and for various special cases. The maximum distance-1 matching problem can be solved in polynomial time in generic graphs [11]. For $k \geq 2$ the complexity of the maximum distance-k matching problem was first investigated in [11] (where it is called the *maximum k-separated matching problem*). The problem was shown to be NP-hard in [11]. The D2EMIS problem ($k=2$) is known to be solvable in polynomial time for a number of special classes of graphs, including trees, chordal graphs, circular arc graphs, interval graphs, k-interval-dimension graphs, trapezoid graphs, and co-comparability graphs [4]. The problem was shown to be APX-complete in [12] for regular graphs. It has also been studied for random graphs (see [13,14], and references therein). Random graphs are formally defined as sub-graphs of a complete graph over n vertices that either have exactly M edges and uniform probability measure [14] or have an edge connecting each pair of vertices with probability p [13]. The focus of these papers is on proving analytical bounds and sharp concentration results for the cardinality of maximum distance-2 matchings (which in this case is an integer random variable).

The term D2EMIS was introduced in [1] where a special case of graphs (geometric intersection graphs) relevant to capacity in radio networks is considered. These graphs are derived by placing nodes in a 2-D plane and connecting nodes with edges based on their geometric properties. A number of special cases are described in [1], the simplest one being *unit disk graphs* in which two nodes are connected by an edge if and only if their Euclidean distance is less than or equal to a given range (which can be normalized to 1). *Disk graphs*, *directed disk graphs* and (r,s)-*civilized graphs* are other special categories of geometric intersection graphs. Although the D2EMIS problem on generic geometric disk graphs is NP-hard, their special geometric characteristics have been exploited in [1,15] to develop polynomial time approximation algorithms (PTAS) for solving this problem. Although the PTAS family of algorithms guarantees that for any fixed ε there is an algorithm that approximates the optimal solution within $(1-\varepsilon)$, the price to pay is that the polynomial exponent gets too large.

Suboptimal distributed algorithms for the maximum 2-packing and k-packing problems have been proposed in [2,5,16,17], respectively. They differ in the number of required steps and amount of locally stored information. A suboptimal distributed algorithm that takes advantage of the broadcast nature of the wireless channel is proposed in [1] for the D2EMIS problem in unit disk graphs. All the above algorithms generate maximal packings (matchings) in the sense that no more nodes (edges) can be added to the solution $S_V(S_E)$ without violating the distance constraints, but the cardinality of the solution is not necessarily optimal.

In [7], we have proposed an efficient algorithm for finding the optimal solution of a generic D2EMIS problem which when run on the state-of-the-art Constraint Integer Programming solver SCIP [10] allowed us to obtain exact solutions for fairly large problem sizes. In this paper, we propose efficient near-optimal algorithms for the generic D2EMIS problem and compare them with existing approximate and exact algorithms. The proposed algorithms fall in the category of the Greedy Randomized Adaptive Search Procedures (the GRASP framework) for solving combinatorial optimization problems which have been successfully applied in a great number of application domains (see [18] for a long list of such application domains and the references therein). The first paper on GRASP is [3], while the interested reader can refer to [19] for a detailed and thorough discussion.

3. Different models of the D2EMIS problem

Given an (undirected) graph $G(V,E)$, we define for each node $u \in V$ the set $N(u) \triangleq \{v \in V \mid d_v(v,u)=1\}$. Further, for each node $n \in V$, $i \in \mathbb{N}$, we define the set $N_i(n) \triangleq \{v \in V-\{n\} \mid d_v(n,v) \leq i\}$. We extend the above definition of node-neighborhood to include set-neighborhoods $N_i(S), S \subseteq V$ in the following manner: $N_i(S) \triangleq \{v \in V-S \mid \exists u \in S : d_v(u,v) \leq i\}$. The expression "distance-i nodes from a node$x$" then refers to all nodes $y \neq x$ that satisfy $d_v(x,y) \leq i$.