Matching in Geometric Uniform Hypergraphs

Matthew J. Katz
The Stein Faculty of Computer and Information Science, Ben-Gurion University of the Negev, Israel

Yuval Nidam
The Stein Faculty of Computer and Information Science, Ben-Gurion University of the Negev, Israel

Rachel Saban
The Stein Faculty of Computer and Information Science, Ben-Gurion University of the Negev, Israel

Micha Sharir
School of Computer Science, Tel Aviv University, Israel

Abstract

Let P be a set of n points in $\mathbb{R}^d$, $d \geq 2$, and let $t \geq 2$ be an integer. Let $H_t(P)$ denote the t-uniform hypergraph on P, whose hyperedges consist of all t-tuples $T \subset P$ for which $\|p - q\| \leq 1$, for any two points $p, q \in T$. A matching in $H_t(P)$ is a collection of vertex-disjoint hyperedges. We present a PTAS for finding a maximum matching in $H_t(P)$. In particular, we present the first PTAS for the well-studied problem known as maximum (vertex-disjoint) triangle packing in unit disk graphs. Our approach consists of a sparsification stage, which replaces P by a subset Q with favorable properties, followed by an implementation of a PTAS for a maximum matching in $H_t(Q)$. The two stages follow the high-level machinery in [3] and [1], respectively, but are considerably more involved.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases Geometric hypergraphs, maximum matching, PTAS, sparsification, local search

Digital Object Identifier 10.4230/LIPIcs.ESA.2026.48

Funding *Matthew J. Katz*: Partially supported by Israel Science Foundation Grant 495/23.
Micha Sharir: Partially supported by Israel Science Foundation Grant 495/23.

1 Introduction

Let P be a set of n points in $\mathbb{R}^d$, $d \geq 2$, and let $t \geq 2$ be an integer. We define a t-uniform hypergraph $H_t(P)$, whose vertex set is P, and a subset $T \subset P$, $|T| = t$, determines a hyperedge of $H_t(P)$ if and only if $\|p - q\| \leq 1$, for any pair of points $p, q \in T$.

Consider the unit ball graph $\mathrm{UBG}(P)$ obtained by drawing a ball of radius $1/2$ around each point of P. That is, the vertex set of $\mathrm{UBG}(P)$ is P, and there is an edge between $p, q \in P$ if and only if $B(p) \cap B(q) \neq \emptyset$, where $B(p), B(q)$ are the balls centered at p and q. Then, $T \subset P$, $|T| = t$, determines a hyperedge of $H_t(P)$ if and only if T is a clique of size t in $\mathrm{UBG}(P)$.

A subset M of hyperedges of $H_t(P)$ is a *matching* in $H_t(P)$ if the hyperedges of M are pairwise disjoint. M is *maximum* if for any other matching M' in $H_t(P)$, $|M| \geq |M'|$. Notice that if $t = 2$, then $H_t(P)$ is "just" a simple graph and the notion of matching in a hypergraph coincides with the usual notion of matching in a graph. Another term for (maximum) matching in a (uniform) hypergraph is (maximum) *packing*.

We study the problem of computing a maximum matching in $H_t(P)$. Specifically, since the exact problem is NP-hard (see [6]), we present a polynomial-time approximation scheme (PTAS) for the problem. That is, given $0 < \varepsilon < 1$, we find in polynomial time (in n) a matching M in $H_t(P)$, such that $|M| \geq (1 - \varepsilon)|M^*|$, where M^* is a maximum matching in $H_t(P)$. In particular, we obtain the first PTAS for maximum (vertex-disjoint) triangle packing in unit disk graphs, see below.

© Matthew J. Katz, Yuval Nidam, Rachel Saban, and Micha Sharir;
licensed under Creative Commons License CC-BY 4.0

A high-level overview. Our PTAS consists of two main stages. In the first stage, referred to as *sparsification*, we reduce the problem to the problem of computing a maximum matching M^* in $H_t(Q)$, for a suitable subset $Q \subseteq P$ for which the *depth* of UBG(Q) is bounded by some constant $c = c(t, d)$. (Recall that the depth of UBG(Q) is the maximum, over all points $z \in \mathbb{R}^d$, of the number of balls, corresponding to points of Q, that contain z.) To obtain a maximum matching in $H_t(P)$, we add to M^* all hyperedges of a collection of matchings M_i, so that the following two properties are satisfied.

(i) Each M_i is a *perfect* matching in $H_t(P_i)$, for suitable pairwise-disjoint subsets P_i of P, such that each $|P_i|$ is divisible by t and any t vertices of P_i determine a hyperedge of $H_t(P_i)$. Thus, $H_t(P_i)$ is *complete* and it admits a perfect matching (any partition of P_i into subsets of size t is such a matching).

(ii) The sets P_i form a partition of $P \setminus Q$. (So in particular $|P \setminus Q|$ is divisible by t.)

In the second stage, we run a *local-search* algorithm to compute a matching M in $H_t(Q)$, such that $|M| \geq (1 - \varepsilon)|M^*|$, for some prespecified $\varepsilon > 0$. We then return the matching $M \cup (\cup_i M_i)$, which is clearly a $(1 - \varepsilon)$-approximation of a maximum matching in $H_t(P)$. The proof of the bound on $|M|$ relies crucially on the bounded-depth property of the set of balls defining UBG(Q). Specifically, this property guarantees that there exists a small *separator* in UBG(Q), namely, a subset $S \subset Q$ of size strictly sublinear in $|Q|$, such that $Q \setminus S$ can be partitioned into subsets A and B, each of size at most some constant fraction of $|Q|$, with no edge between A and B. The existence of such a separator enables us to adapt the local-search analysis of Aschner et al. [1] (which builds on that of Mustafa and Ray [15]) to our setting.

The first stage of our PTAS is influenced by the work of Bonnet et al. [3], who presented (among other related results) a randomized $O(n^{\omega/2})$-time algorithm for computing a (regular) matching, which is maximum with high probability, in a unit disk graph, where ω is the matrix multiplication exponent, that is, the infimum of all real numbers τ such that two $n \times n$ matrices can be multiplied in $O(n^\tau)$ time.

As just mentioned, the second stage of our PTAS is based on a generalization of a result of Aschner et al. [1], who present a PTAS, based on local search, for computing a maximum matching in $H_3(P)$, under the assumption that P is a set of points in the plane such that UBG(P) is of constant depth.

1.1 Related work

Computing a maximum matching in a hypergraph is a classical problem in graph theory. A matching in a 3-partite hypergraph is referred to as a *3-dimensional matching* (3DM). Karp [10] showed that computing a maximum 3DM is NP-hard. Since 3DM is a special case of maximum matching in 3-uniform hypergraphs, it follows that computing a maximum matching in a 3-uniform hypergraph is also NP-hard. More generally, and not surprisingly, the problem is NP-hard for all $t \geq 3$.

From the approximation perspective, Hurkens and Schrijver [8] showed that for every $\varepsilon > 0$, one can obtain a $(2/t - \varepsilon)$-approximation for maximum matching in general t-uniform hypergraphs. For the special case of maximum 3DM, Cygan [7] gave an improved algorithm achieving approximation ratio 3/4. For simple graphs of degree at most 4, Manić and Wakabayashi [13] obtained a $(1/(3 - \sqrt{13}/2) - \varepsilon) \approx (0.835 - \varepsilon)$-approximation algorithm for maximum (vertex-disjoint) triangle packing, while Caprara and Rizzi [4] proved that this case is already APX-hard. (Maximum triangle packing in a simple graph $G = (V, E)$

is equivalent to maximum matching in the 3-uniform hypergraph on V, in which (u, v, w) determine a hyperedge if and only if $(u, v), (v, w), (w, u) \in E$.) Finally, for planar graphs, Baker [2] used the *shifting strategy* to obtain PTASs for maximum triangle matching and, more generally, for maximum H-matching, where H is any fixed connected graph structure of constant size, such as a 4-clique or a 4-cycle.

In our setting, when $t = 2$ the problem becomes maximum matching in unit disk graphs. Bonnet et al. [3] designed an efficient algorithm for this problem using a two-phase framework that first sparsifies the vertex set of the input graph, so that the corresponding collection of disks centered at the vertices of the sparsified subset has bounded depth, and then computes a maximum matching in the resulting sparsified graph. Moreover, every vertex removed during the sparsification phase is, by construction, guaranteed to be matched, with some other removed vertex, allowing the matching in the sparsified graph to be extended to a maximum matching in the original graph. See [11] for an application of the algorithm of Bonnet et al. to bottleneck matching in the plane.

When $t = 3$, Hunt et al. [9] extended Baker's results to $H_3(P)$, under the assumption that the distance between any two points in P is at least some fixed $\lambda > 0$, and Aschner et al. [1] presented a PTAS, based on local search, for maximum matching in $H_3(P)$, under the assumption that the unit disk graph on P has bounded depth.

The separator theorems of Miller et al. [14] and Smith and Wormald [16] show that for any k-ply system $\mathcal{B}$ of n balls in $\mathbb{R}^d$ (i.e., a family $\mathcal{B}$ of balls with the property that any point in $\mathbb{R}^d$ is contained in at most k balls of $\mathcal{B}$), there exists a sphere σ that intersects only $O(k^{1/d}n^{1-1/d})$ balls of $\mathcal{B}$ and partitions the remaining balls of $\mathcal{B}$ into two roughly balanced parts (those that are contained in the interior of σ and those that are contained in the exterior of σ). Miller et al. also present a randomized algorithm to find such a sphere in linear time.

Building on such results, *local search* has become a central tool for approximation problems of this kind. In such a search we fix a parameter k, and, as long as possible, replace $k - 1$ hyperedges of the structure under construction by k other hyperedges, maintaining the pairwise disjointness of the hyperedges (or any other desired property). See below for more details. Mustafa and Ray [15] and Chan and Har-Peled [5] developed frameworks showing that, when combined with separator theorems, notably the planar separator theorem of Lipton and Tarjan [12], local search yields PTASs for a wide range of geometric optimization problems.

1.2 Our results

We extend, in a nontrivial manner, the sparsification technique of Bonnet et al. [3], and generalize the local-search framework of Aschner et al. [1], to obtain a PTAS for maximum matching in $H_t(P)$, for any $d \geq 2$ and $t \geq 3$ (Theorem 12). In particular, we present the first PTAS for maximum (vertex-disjoint) triangle packing in unit ball graphs, an important and widely-studied graph family. Previously, such results were known only for planar graphs and for a few restricted graph families defined by somewhat artificial constraints (see above). The extended sparsification technique is described in Section 2 (Theorem 11). Although its high-level ideas are inspired by those of the original technique, their implementation for larger values of t and d is highly nontrivial and requires numerous careful adjustments and structural enhancements. We anticipate that the extended technique will find further applications beyond the scope of this work. The generalized local-search approach is presented in Section 3, where the extension in this case is more direct yet still requires careful adaptation to our setting.

Finally, we show that our two-phase approach applies not only to the problem of finding a maximum number of disjoint t-cliques in a unit ball graph, but also to several related problems. For instance, given P, d, and t, one may wish to find as many vertex-disjoint t-cycles as possible in the corresponding unit ball graph, or as many disjoint t-subsets as possible such that each subset is contained in some unit ball. In Section 2.3, we show that our two-phase approach yields a PTAS for these problems, as well as for various similar ones.

2 Sparsification

In this section we describe the first stage of our PTAS, namely, the *sparsification* stage. Recall that the objective of this stage is to replace the input set P by a subset Q, such that $\mathrm{UBG}(Q)$ has bounded depth, and such that a (near) maximum matching in $H_t(Q)$ can be extended to a (near) maximum matching in $H_t(P)$. We first consider the case $d = 2$ and $t = 3$, so P is a set of points in the plane, and our edges are triangles. We then extend the result to $d = 2$ and any $t > 3$, and finally to the general case $d > 2$ and any t.

Terminology and notation. Since we are dealing only with hypergraphs, we often refer to hyperedges simply as edges. Also, we use the notation $e = (p_1, p_2, p_3)$ to denote an edge of $H_3(P)$, where p_1, p_2, p_3 are three distinct points of P, but it should be clear that there is no significance to the order in which the points are listed, e.g., (p_3, p_1, p_2) refers to the same edge e. Similar conventions apply to the general case $t > 3$ (and/or $d > 2$). Finally, as long as we are in the plane, we will use the notation UDG (for unit-disk graphs) to replace UBG.

2.1 $d = 2$ and $t = 3$

We begin by laying a regular grid $\mathcal{G}$ with side length $1/\sqrt{2}$. We may assume that every point of P lies in the *interior* of some grid cell. (Otherwise, we can impose this by a suitable translation of the grid.) For a cell v of $\mathcal{G}$, we denote the subset of points of P that lie in v by P_v, that is, $P_v = P \cap v$. Clearly, for each grid cell $v \in \mathcal{G}$, the subgraph of $\mathrm{UDG}(P)$ induced by P_v is a clique, and thus any three points in P_v form a hyperedge of $H_3(P)$.

For any three cells u, v, w such that not all three are the same (but two may be the same), we denote the corresponding set of hyperedges of $H_3(P)$ by $E(u, v, w)$, that is,

$$E(u, v, w) = \{(p, q, r) \in E(H_3(P)) \mid p \in P_u, q \in P_v, r \in P_w\}\,.$$

We now define the *pattern hypergraph* $\mathrm{PH} = \mathrm{PH}(\mathcal{G})$, whose vertex set is the set of all grid cells (that contain points of P). The hypergraph PH is 3-uniform and its hyperedges are all triples of cells, not all the same, such that each pair of them is within Euclidean distance 1 from each other. That is,

$$E(\mathrm{PH}) = \{(u, v, w) \mid u, v, w \text{ are grid cells, not all the same, and } \|u - v\|, \|v - w\|, \|u - w\| \le 1\}\,,$$

where $\|u - v\| = \min_{p \in u,\, q \in v} \|p - q\|$, and similarly for $\|v - w\|$, $\|u - w\|$. (Here too there is no significance to the order in which the cells are listed in a hyperedge of $E(\mathrm{PH})$, but unlike $H_3(P)$ a vertex (i.e., cell) may appear twice in a hyperedge of $E(\mathrm{PH})$.) The usefulness of the pattern hypergraph is demonstrated by the following obvious property. For any triple of grid cells u, v, w, not all the same, $E(u, v, w) \neq \emptyset \implies (u, v, w) \in E(\mathrm{PH})$.

We can now express the edge set of $H_3(P)$ as the union of pairwise disjoint sets (ignoring empty sets):

$$E(H_3(P)) = \bigcup_{(u,v,w) \in E(\mathrm{PH})} E(u, v, w) \cup \bigcup_{v \in \mathcal{G}} E(H_3(P_v))\,.$$

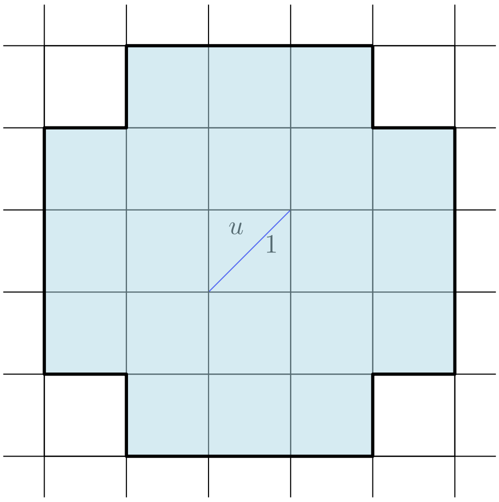

Figure 1 The grid cells within distance 1 from the cell u.

Let λ be the number of grid cells that are within Euclidean distance 1 from a given cell u, excluding u. It is easy to verify that $\lambda = 20$. See Figure 1. Let Λ denote the degree of a cell in the pattern hypergraph PH, where the degree of a cell is the number of hyperedges to which it belongs. Then $\Lambda \le \binom{\lambda}{2} + 2\lambda = O(\lambda^2)$.

▶ **Lemma 1.** *There exists a maximum matching M in $H_3(P)$ such that, for any hyperedge $(u, v, w) \in E(\mathrm{PH})$, M has at most two hyperedges from $E(u, v, w)$.*

Proof. Let M be a maximum matching in $H_3(P)$ that minimizes the number of hyperedges that span more than one cell, that is, a maximum matching for which the sum $\sum_{(u,v,w)\in E(\mathrm{PH})} |M \cap E(u, v, w)|$ is minimum.

Assume now that there exists a hyperedge $(u, v, w) \in E(\mathrm{PH})$, such that $|M \cap E(u, v, w)| \ge 3$. Then M has three disjoint hyperedges (p_1, q_1, r_1), (p_2, q_2, r_2), and (p_3, q_3, r_3), such that $p_i \in P_u$, $q_i \in P_v$, and $r_i \in P_w$, for $i = 1, 2, 3$. We can then define a new matching M' by replacing these hyperedges by the hyperedges (p_1, p_2, p_3), (q_1, q_2, q_3), and (r_1, r_2, r_3). That is, we define

$$M' = (M \setminus \{(p_1, q_1, r_1), (p_2, q_2, r_2), (p_3, q_3, r_3)\}) \cup \{(p_1, p_2, p_3), (q_1, q_2, q_3), (r_1, r_2, r_3)\}\,.$$

The matching M' is still a maximum matching, but the number of its edges that span more than one cell is smaller (by 3) than that of M, contradicting the choice of M. Therefore, no such hyperedge of PH exists. ◀

For each hyperedge $(u, v, w) \in E(\mathrm{PH})$, we construct a subset of hyperedges $\mathcal{W}(u, v, w) \subseteq E(u, v, w)$ by applying Algorithm 1, given below. The parameter Ψ in the algorithm will be determined later.

Here is a verbal sketch of the algorithm. It first constructs a matching $M = M(u, v, w) \subseteq E(u, v, w)$, by adding "legal" hyperedges one at a time, until either M consists of $3\Psi + 1$ hyperedges, or M consists of fewer hyperedges but is maximal, that is, all remaining hyperedges are illegal (each intersects a hyperedge already in M).

Now, if $|M| = 3\Psi + 1$, then set $\mathcal{W}(u, v, w) = M$; otherwise, construct $\mathcal{W}(u, v, w)$ incrementally as follows. For each point $p \in P_u \cup P_v \cup P_w$ that appears in an edge of M do the following. Associate a counter with each of the points in $(P_u \cup P_v \cup P_w) \setminus \{p\}$ and set it to 0. Additionally, set an *edge counter* c to 0. Now, for each edge $e \in E(u, v, w)$ in which p

appears, add it to $\mathcal{W}(u,v,w)$ if (i) so far $e \notin \mathcal{W}(u,v,w)$, (ii) $c < 2\Psi(\Psi+1)+1$, and (iii) the counters associated with the other two points in e are both less than $\Psi+1$. (If e is added, increase c and the other two relevant vertex counters by 1.)

Algorithm 1 Constructing the set $\mathcal{W}(u,v,w)$.

Require: Three cells u, v, w with $(u,v,w) \in E(\mathrm{PH})$; parameter Ψ
Ensure: A set $\mathcal{W}(u,v,w) \subseteq E(u,v,w)$
Construct a matching $M \subseteq E(u,v,w)$ greedily. Stop when the matching reaches size $3\Psi+1$ or earlier, if M becomes maximal.
if $|M| = 3\Psi+1$ **then**
    $\mathcal{W}(u,v,w) \leftarrow M$
else ▷ $|M| \le 3\Psi$
    $\mathcal{W}(u,v,w) \leftarrow \emptyset$
    for all vertices $p \in P_u \cup P_v \cup P_w$ that appear in an edge of M **do**
        Initialize counters $C[p'] \leftarrow 0$ for all $p' \in (P_u \cup P_v \cup P_w) \setminus \{p\}$
        Initialize edge counter $c \leftarrow 0$
        for all edges $(p,p_1,p_2) \in E(u,v,w)$ containing p **do**
            if $C[p_1] < \Psi+1$ and $C[p_2] < \Psi+1$ and $c < 2\Psi\cdot(\Psi+1)+1$ and $(p,p_1,p_2) \notin \mathcal{W}(u,v,w)$ **then**
                Add (p,p_1,p_2) to $\mathcal{W}(u,v,w)$
                Increment $C[p_1]$, $C[p_2]$, and c
            end if
        end for
    end for
end if
return $\mathcal{W}(u,v,w)$

We notice that there are at most $3\Psi+1$ edges in M, and for each vertex in these edges, we add at most $O(\Psi^2)$ edges, thus the size of $\mathcal{W}(u,v,w)$ is $O(\Psi^3)$. Note also that, unless $|M| = 3\Psi+1$, $\mathcal{W}(u,v,w)$ is in general not a matching in $E(u,v,w)$.

We now define the graph $G_{(u_1,u_2,u_3)}$, for each $(u_1,u_2,u_3) \in E(\mathrm{PH})$, as the graph whose vertex set is $P_{u_1} \cup P_{u_2} \cup P_{u_3}$ and edge set is $\mathcal{W}(u_1,u_2,u_3)$.

▶ **Lemma 2.** *A maximum matching in*

$$\tilde{G} = \left(\bigcup_{(u_1,u_2,u_3)\in E(\mathrm{PH})} G_{(u_1,u_2,u_3)} \right) \cup \left(\bigcup_{u\in\mathcal{G}} H_3(P_u) \right)$$

is a maximum matching in $H_3(P)$.

Proof. We first claim that the size of a maximum matching in $\tilde{G}$ is equal to the size of a maximum matching in $H_3(P)$. Therefore, any maximum matching in $\tilde{G}$ is also a maximum matching in $H_3(P)$, since $\tilde{G}$ is a subgraph of $H_3(P)$.

By Lemma 1, there exists a maximum matching in $H_3(P)$, such that for any hyperedge $(u_1,u_2,u_3) \in E(\mathrm{PH})$, the matching has at most two hyperedges from $E(u_1,u_2,u_3)$. Among all such maximum matchings, let M be one that minimizes the number of hyperedges that span more than one cell and do not belong to their corresponding subset $\mathcal{W}(u_1,u_2,u_3)$. That is, M is a matching that minimizes k, where

$$k = \sum_{(u_1,u_2,u_3)\in E(\mathrm{PH})} |M \cap (E(u_1,u_2,u_3) \setminus \mathcal{W}(u_1,u_2,u_3))|\,.$$

We show below that $k = 0$, and therefore M is also a matching in $\tilde{G}$. But $\tilde{G}$ is a subgraph of $H_3(P)$, so M, which is a maximum matching in $H_3(P)$, is clearly also a maximum matching in $\tilde{G}$. This completes the proof of the claim above, once we show that k is indeed zero.

Assume to the contrary that $k > 0$. Let M' be the set of hyperedges of M that span more than one cell, i.e.,

$$M' = M \cap \left(\bigcup_{(u_1,u_2,u_3)\in E(\mathrm{PH})} E(u_1, u_2, u_3) \right).$$

Since M has at most two hyperedges from $E(u, v, w)$, for any triple $(u, v, w) \in E(\mathrm{PH})$, and since each cell appears in at most Λ such triples, it follows that each cell x is "involved" in at most 2Λ hyperedges of M'. However, in some of these hyperedges (at most 2λ), x may be involved twice. We conclude therefore that at most $2\Lambda + 2\lambda$ vertices of P_x are incident to hyperedges of M'. We put $\Psi = 2\Lambda + 2\lambda$.

By assumption, there exists a hyperedge $(p, q, r) \in M'$ which is not in its corresponding subset $\mathcal{W}(u, v, w)$. We aim to replace (p, q, r) by a hyperedge from $\mathcal{W}(u, v, w)$ that avoids all other vertices that appear in M'. We refer to these vertices as *forbidden* vertices. By the preceding arguments, there are at most 2Λ hyperedges in M' that use vertices from, say, u (including (p, q, r)). The other vertices in these hyperedges are not necessarily from v or w, but at any rate the number of vertices from each of u, v, w is at most Ψ. Lemma 3, given below, guarantees that for any choice of at most Ψ forbidden vertices in each of the cells u, v, w, there exists a hyperedge $(p', q', r') \in \mathcal{W}(u, v, w)$, such that none of p', q', r' are among the forbidden vertices, that is, none of them is incident to a hyperedge of M'.

We now modify the matching M by removing the edge (p, q, r) and inserting (p', q', r') instead. In addition, if M contains an edge e that spans only u (i.e., $e \in E(H_3(P_u))$) such that p' is one of its vertices, then we replace p' by p in e. We do this also for v and q' and for w and r', if needed.

Finally, we observe that the resulting modified M is still a maximum matching, and moreover, it still has at most two edges from $E(u_1, u_2, u_3)$, for any $(u_1, u_2, u_3) \in E(\mathrm{PH})$. However, we managed to decrease (by 1) the number of edges in M that span more than one cell and do not belong to their corresponding $\mathcal{W}(u_1, u_2, u_3)$ subset. This latter fact contradicts our initial assumption about M. ◀

To complete the proof, we have:

▶ **Lemma 3.** *For any triple $(u, v, w) \in E(\mathrm{PH})$, the subset $\mathcal{W}(u, v, w)$ satisfies the following property. For any choice of Ψ forbidden vertices in each of the cells involved (so only 2Ψ vertices in total if only two cells are involved), if there exists an edge $(p, q, r) \in E(u, v, w) \setminus \mathcal{W}(u, v, w)$ that is not incident to a forbidden vertex, then there also exists such an edge $(p', q', r') \in \mathcal{W}(u, v, w)$.*

Proof. We distinguish between two cases, where M here is the matching constructed in the first step of Algorithm 1. Below, we say that an edge in $E(u, v, w)$ is *forbidden* if it is incident to a forbidden vertex.

Case 1: $|M| = 3\Psi + 1$. Recall that in this case $\mathcal{W}(u, v, w) = M$, so we need to show that there exists an edge of M that is not forbidden. Indeed, since there are at most 3Ψ forbidden vertices in $P_u \cup P_v \cup P_w$, there can be at most 3Ψ forbidden edges in M, and therefore at least one of the edges of M is not forbidden.

Case 2: $|M| < 3\Psi + 1$. In this case, M is maximal, that is, for any edge of $E(u, v, w) \setminus M$, and in particular for (p, q, r), there exists an edge of M, such that the two edges have a vertex in common. Assume without loss of generality that p is such a vertex. We examine the reasons why (p, q, r) was not added to $\mathcal{W}(u, v, w)$ in the edge-adding step of Algorithm 1.

Edge count exceeded. When processing the vertex p, $2\Psi \cdot (\Psi + 1) + 1$ other edges of the form (p, p_1, p_2) have already been added to $\mathcal{W}(u, v, w)$. However, at most $2\Psi \cdot (\Psi + 1)$ of these edges are forbidden, since there are 2Ψ forbidden vertices in the cells v and w (only Ψ if $v = w$) and each of them can appear in at most $\Psi + 1$ of the edges. We conclude that at least one of the edges that have already been placed in $\mathcal{W}(u, v, w)$ is not forbidden.

Vertex count exceeded. When processing the vertex p, the edge (p, q, r) was rejected since one of the two counters, say $C[q]$, has already reached the value $\Psi + 1$. In other words, $\Psi + 1$ edges of the form $(p, q, ?)$ have already been added to $\mathcal{W}(u, v, w)$. But at most Ψ of them are incident to a forbidden vertex in w. So once again we conclude that at least one of the edges that have already been placed in $\mathcal{W}(u, v, w)$ is not forbidden.

In both cases, there exists an edge (p', q', r') in $\mathcal{W}(u, v, w)$ that is not forbidden. ◀

With this lemma at hand, the proof of Lemma 2 is completed.

Let Q_0 be the set of all points that are incident to an edge of one of the subsets $\mathcal{W}(u_1, u_2, u_3)$, $(u_1, u_2, u_3) \in E(\mathrm{PH})$, that is,

$$Q_0 = \bigcup_{\substack{(u_1,u_2,u_3)\in E(\mathrm{PH}) \\ (p_1,p_2,p_3)\in \mathcal{W}(u_1,u_2,u_3)}} \{p_1, p_2, p_3\},$$

and consider the hypergraph $H_3(Q_0)$. Since the edge set of $H_3(Q_0)$ contains all the edges of the graphs $G_{(u_1,u_2,u_3)}$, for $(u_1, u_2, u_3) \in E(\mathrm{PH})$, we may conclude (by Lemma 2) that a maximum matching in

$$H_3(Q_0) \cup (\bigcup_{u\in\mathcal{G}} H_3(P_u))$$

is a maximum matching in $H_3(P)$.

Next, recall that $H_3(Q_0)$ is obtained from $\mathrm{UDG}(Q_0)$, by forming a hyperedge for each 3-clique in $\mathrm{UDG}(Q_0)$. We prove below that the *depth* of $\mathrm{UDG}(Q_0)$ is bounded by some constant, where the depth of $\mathrm{UDG}(Q_0)$ is the maximum, over all points $z \in \mathbb{R}^2$, of the number of disks of $\mathrm{UDG}(Q_0)$ that contain z. This property is crucial for the next stage.

▶ **Lemma 4.** *For any grid cell u, $|Q_0 \cap P_u| = O(\Lambda^4)$.*

Proof. Recall that Λ denotes the degree of a cell u in the pattern hypergraph PH and that $\Lambda = O(\lambda^2)$. Moreover, it follows from the construction in Algorithm 1 that $|\mathcal{W}(u_1, u_2, u_3)| = O(\Psi^3)$, for any triple $(u_1, u_2, u_3) \in E(\mathrm{PH})$, where $\Psi = 2\Lambda + 2\lambda$. So $|\mathcal{W}(u_1, u_2, u_3)| = O(\Lambda^3)$, and the total number of hyperedges in the subsets $\mathcal{W}(u_1, u_2, u_3)$ involving u (i.e., for which $u \in \{u_1, u_2, u_3\}$) is bounded by $O(\Lambda^4)$ (because the degree of u in $E(\mathrm{PH})$ is at most Λ). Since each of these hyperedges is incident to at most two points of P_u, we conclude that $|Q_0 \cap P_u| = O(\Lambda^4)$. ◀

▶ **Lemma 5.** *The depth of* $\mathrm{UDG}(Q_0)$ *is* $O(\lambda^9)$.

Proof. Fix a point $z \in \mathbb{R}^2$, and let $u \in \mathcal{G}$ be the cell containing z. Any point $p \in Q_0$ for which $\|p - z\| \le 1$ must lie either in one the λ cells surrounding u (see above) or in u itself. Since, by Lemma 4, the number of points of Q_0 in a grid cell is $O(\Lambda^4) = O(\lambda^8)$, we get that the number of points of Q_0 whose disk might contain z is $(1+\lambda) \cdot O(\lambda^8) = O(\lambda^9)$. ◀

We now set $Q = Q_0$ and proceed as follows. For each cell $u \in \mathcal{G}$, consider the set $P'_u := P_u \setminus Q_0$ and let $i = |P'_u| \bmod 3$. If $i > 0$, we remove i arbitrary points from P'_u and add them to Q. At this point, the size of each of the sets P'_u is divisible by 3, and we have added at most two points from u, which were previously not in Q_0, to Q.

We observe that the depth of $\mathrm{UDG}(Q)$ is still $O(\lambda^9)$, since the depth of a point $z \in \mathbb{R}^2$ can increase by at most $2(\lambda + 1)$ when moving from $\mathrm{UDG}(Q_0)$ to $\mathrm{UDG}(Q)$, see Lemma 5.

We can now establish the main result of this subsection.

▶ **Theorem 6.** *Let P be a set of n points in the plane. The problem of finding a maximum matching in $H_3(P)$ can be reduced to the problem of finding a maximum matching in $H_3(Q)$, for a subset Q of P for which the depth of $\mathrm{UDG}(Q)$ is $O(\lambda^9)$.*

Proof. We compute a matching M in $H_3(P)$ as follows. Recall that for each cell u, the graph $H_3(P_u)$ is complete, that is, any three points from P_u define a hyperedge of $H_3(P_u)$. We can therefore compute, for each non-empty cell u, a perfect matching in $H_3(P'_u)$, where P'_u is as defined above, and add its edges to M. Next we compute a maximum matching in $H_3(Q)$ and add its edges to M. (Notice that we are not claiming that we actually know how to compute such a matching. In Section 3, we present a PTAS for computing such a matching.) Since $Q \cap P'_u = \emptyset$, for each cell u, M is a maximum matching in $H_3(Q) \cup (\bigcup_{u \in \mathcal{G}} H_3(P'_u))$.

We claim that M is a maximum matching in $H_3(P)$. Indeed, let N be a maximum matching in $H_3(Q_0) \cup (\bigcup_{u \in \mathcal{G}} H_3(P_u))$ that minimizes the sum, over all cells $u \in \mathcal{G}$, of the number of points from $P_u \setminus Q_0$ that appear in *mixed* edges, where an edge is *mixed* if it is incident to both a point from $P_u \setminus Q_0$, for some u, and a point from Q_0. Notice that, by the definition of N, a mixed edge cannot span more than one cell, i.e., it is incident to three points from the same subset P_u, at least one from $P_u \setminus Q_0$ and at least one from $P_u \cap Q_0$.

Recall that (by Lemma 2) N is a maximum matching in $H_3(P)$. To prove that M is a maximum matching in $H_3(P)$, we show that N can be transformed, by replacing some of its edges, into a matching in $H_3(Q) \cup (\bigcup_{u \in \mathcal{G}} H_3(P'_u))$. But this implies that $|M| \ge |N|$ and therefore M is a maximum matching in $H_3(P)$ as claimed.

Let u be any non-empty cell. Notice that the number of points from $P_u \setminus Q_0$ that appear in mixed edges of N is at most two, since otherwise we can decrease the number of such points. (Indeed, if there are three or more such points, then we can form a non-mixed edge from any three of them, and then form new (mixed or non-mixed) edges from the three or six points remaining in the edges from which the three points were removed.) Moreover, if only one point $p \in P_u \setminus Q_0$ appears in a mixed edge of N, then the number of points of $P_u \setminus Q_0$ that do not appear in any edge of N is at most one, since if there were two such points, we could form a non-mixed edge from them and the point p, thus replacing the mixed edge of N by a non-mixed one. Similarly, if two points from $P_u \setminus Q_0$ appear in mixed edges of N, then all the points of $P_u \setminus Q_0$ appear in some edge of N.

We conclude that the number of points from $P_u \setminus Q_0$ that either appear in mixed edges or do not appear in any edge of N is at most two, and that this number is $|P_u \setminus Q_0| \bmod 3$. Thus, if $|P_u \setminus Q_0| \bmod 3 \neq 0$, we replace these one or two points with the one or two points, respectively, that we have removed from $P_u \setminus Q_0$ to obtain the set P'_u. After performing this replacement for each of the non-empty cells, N contains a perfect matching in P'_u, for each non-empty cell u, and its remaining edges consist of a matching in $H_3(Q)$. ◀

2.2 $d = 2$ and $t > 3$

We now extend our results for the 3-uniform case to the general t-uniform case, still in the plane, where $t > 3$ is a constant. We show how to find a subset $\mathcal{W}(u_1, \dots, u_t)$, of a controlled number of *cross-cell* hyperedges, for each t-tuple of cells (that are not all the same), with some favorable properties. This will allow us to bound the depth of the resulting sparse graph $\mathrm{UDG}(Q_0)$ (or $\mathrm{UDG}(Q)$) in terms of t and λ.

The various definitions introduced for the case $t = 3$ extend naturally, with obvious modifications, to the case $t > 3$.

We begin by generalizing Lemma 1.

▶ **Lemma 7.** *There exists a maximum matching in $H_t(P)$, such that, for any t-tuple $(u_1, \dots, u_t) \in E(\mathrm{PH})$, the matching has at most $t - 1$ hyperedges from $E(u_1, \dots, u_t)$.*

Proof. As in the proof of Lemma 1, let M be a maximum matching for which the sum $\sum_{(u_1,\dots,u_t)\in E(\mathrm{PH})} |M \cap E(u_1, \dots, u_t)|$ is minimum.

Suppose there exists a t-tuple $(u_1, \dots, u_t)$ such that $|M \cap E(u_1, \dots, u_t)| \geq t$. Then, for each cell u_i, $1 \leq i \leq t$, we can pick the vertex of P_{u_i} from each of t of these hyperedges, thereby obtaining t vertices of P_{u_i} or, in other words, a hyperedge of $H_t(P)$ consisting of vertices solely from u_i. Replacing the t cross-cell hyperedges with these t intra-cell hyperedges contradicts the minimality assumption. We conclude that each such t-tuple of cells corresponds to at most $t - 1$ hyperedges of M. ◀

For each $(u_1, \dots, u_t) \in E(\mathrm{PH})$, we construct the subset $\mathcal{W}(u_1, \dots, u_t) \subseteq E(u_1, \dots, u_t)$ using Algorithm 1 with obvious modifications. We stop the greedy construction once the matching M reaches size $t\Psi + 1$, where the parameter Ψ will be determined later. If this size has not been reached, we use the thresholds $(t-1)\Psi(\Psi^{t-2} + 1) + 1$ for the edge counter, and $\Psi^{t-2} + 1$ for the vertex counters.

If $|M| = t\Psi + 1$, then the number of edges in $\mathcal{W}(u_1, \dots, u_t)$ is clearly $O(t\Psi)$, since in this case $\mathcal{W}(u_1, \dots, u_t) = M$. If $|M| < t\Psi + 1$, then the number of vertices that appear in some edge of M is $O(t^2\Psi)$ and for each of them we add $O(t\Psi^{t-1})$ edges to $\mathcal{W}(u_1, \dots, u_t)$, so the number of edges in $\mathcal{W}(u_1, \dots, u_t)$ is $O(t^3\Psi^t)$. We conclude that for any $(u_1, \dots, u_t) \in E(\mathrm{PH})$ the number of edges in $\mathcal{W}(u_1, \dots, u_t)$ is $O(t^3\Psi^t)$.

We now state and prove the lemma analogous to Lemma 3.

▶ **Lemma 8.** *For any t-tuple $(u_1, \dots, u_t) \in E(\mathrm{PH})$, the subset $\mathcal{W}(u_1, \dots, u_t)$ satisfies the following property. For any choice of Ψ forbidden vertices in each of the cells involved (i.e., at least two cells and at most t), if there exists an edge $(p_1, \dots, p_t) \in E(u_1, \dots, u_t) \setminus \mathcal{W}(u_1, \dots, u_t)$ that is not incident to a forbidden vertex, then there also exists such an edge $(p'_1, \dots, p'_t) \in \mathcal{W}(u_1, \dots, u_t)$.*

Proof. We distinguish between two cases.

Case 1: $|M| = t\Psi + 1$. Recall that in this case, $\mathcal{W}(u_1, \dots, u_t) = M$, so we need to show that there exists an edge in M that is not forbidden, i.e., it is not incident to a forbidden vertex. Indeed, since there are at most $t\Psi$ forbidden vertices in $P_{u_1} \cup \dots \cup P_{u_t}$, there can be at most $t\Psi$ forbidden edges in M, and therefore at least one of the edges in M is not forbidden.

Case 2: $|M| < t\Psi + 1$. In this case, M is maximal, that is, for any edge of $E(u_1, \dots, u_t) \setminus M$, and in particular $(p_1, \dots, p_t)$, there exists an edge of M, such that the two edges have a vertex in common. Assume without loss of generality that p_1 is such a vertex for the edge $(p_1, \dots, p_t)$. We examine the reasons why $(p_1, \dots, p_t)$ was not added to $\mathcal{W}(u_1, \dots, u_t)$ in the edge-adding step of Algorithm 1.

Edge count exceeded. When processing the vertex p_1, there are already $(t-1)\Psi(\Psi^{t-2}+1)+1$ edges containing p_1. But, at most $(t-1)\Psi(\Psi^{t-2}+1)$ of them are forbidden, since there are $(t-1)\Psi$ forbidden vertices in the cells $u_2, \ldots, u_t$ and each of them can appear in at most $(\Psi^{t-2}+1)$ of the edges. We conclude that at least one of these added edges is not forbidden.

Vertex count exceeded. When processing the vertex p_1, the edge $(p_1, \ldots, p_t)$ was rejected since one of the $t-1$ counters, say $C[p_2]$, has already reached the value $\Psi^{t-2}+1$. In other words, $(\Psi^{t-2}+1)$ edges of the form $(p_1, p_2, ?, \ldots, ?)$ have already been added to $\mathcal{W}(u_1, \ldots, u_t)$. But, at most Ψ^{t-2} of them are incident to a forbidden vertex in $u_3, \ldots, u_t$. (Since the number of $(t-2)$-tuples that one can form by choosing a forbidden vertex from each of the cells $u_3, \ldots, u_t$ is at most Ψ^{t-2}.) So once again we conclude that at least one of the added edges is not forbidden. ◀

We now state the general version of Lemma 2. Recall that $G_{(u_1,\ldots,u_t)}$ is the graph with vertex set $P_{u_1} \cup \cdots \cup P_{u_t}$ and edge set $\mathcal{W}(u_1, \ldots, u_t)$.

▶ **Lemma 9.** *A maximum matching in*

$$\tilde{G} = \left(\bigcup_{(u_1,\ldots,u_t)\in E(\mathrm{PH})} G_{(u_1,\ldots,u_t)} \right) \cup \left(\bigcup_{u\in\mathcal{G}} H_t(P_u) \right)$$

is a maximum matching in $H_t(P)$.

We omit most of the proof, which is essentially the same as the proof of Lemma 2. We only present the part in which we fix the value of Ψ. Recall that M is a maximum matching in the hypergraph $\tilde{G}$ with the property (among others) that $|M \cap E(u_1 \ldots, u_t)| \le t-1$, for any t-tuple $(u_1 \ldots, u_t) \in E(\mathrm{PH})$. Moreover, as defined in the proof of Lemma 2, $M' \subseteq M$ is the set of cross-cell hyperedges. We aim to bound the number of vertices that are incident to hyperedges of M' (which are the forbidden vertices in Lemma 8).

As above Λ denotes the degree of a cell u in the pattern hypergraph PH, which is now $O(\lambda^{t-1})$. Since M has at most $t-1$ hyperedges from $E(u_1 \ldots, u_t)$, for any t-tuple $(u_1 \ldots, u_t) \in E(\mathrm{PH})$, and since each cell appears in at most Λ such t-tuples, it follows that each cell x is "involved" in at most $(t-1)\Lambda$ hyperedges of M'. However, in these hyperedges, x may be involved multiple times, up to $t-1$ times. We conclude that at most $(t-1)^2\Lambda$ vertices from the cell x can appear in hyperedges of M'. We thus set $\Psi = (t-1)^2\Lambda$.

As before, define

$$Q_0 = \bigcup_{\substack{(u_1,\ldots,u_t)\in E(\mathrm{PH}) \\ (p_1,\ldots,p_t)\in\mathcal{W}(u_1,\ldots,u_t)}} \{p_1, \ldots, p_t\},$$

and observe that a maximum matching in $H_t(Q_0) \cup (\bigcup_{u\in\mathcal{G}} H_t(P_u))$ is a maximum matching in $H_t(P)$.

▶ **Lemma 10.** *The depth of* UDG(Q_0) *is* $O(t^{2t+4} \cdot \lambda^{t^2})$.

Proof. The number of edges in each subset $\mathcal{W}(u_1, \ldots, u_t)$ is $O(t^3\Psi^t)$, and since $\Psi = O(t^2\Lambda)$, we get $|\mathcal{W}(u_1, \ldots, u_t)| = O(t^{2t+3}\Lambda^t)$. So the total number of hyperedges in the subsets involving u is bounded by $O(t^{2t+3}\Lambda^{t+1})$. Since each of these hyperedges is incident to at most $t-1$ points of P_u, we conclude that $|Q_0 \cap P_u| = O(t^{2t+4}\Lambda^{t+1})$.

Now, fix a point $z \in \mathbb{R}^2$, and let $u \in \mathcal{G}$ be the cell containing z. Any point $p \in Q_0$ for which $\|p - z\| \le 1$ must lie either in one of the λ cells surrounding u or in u itself. As mentioned above, the number of points of Q_0 in a grid cell is $O(t^{2t+4}\Lambda^{t+1})$, we get that the number of points of Q_0 that might cover z is $(\lambda+1) \cdot O(t^{2t+4}\Lambda^{t+1}) = O(t^{2t+4}\Lambda^{t+1}\lambda) = O(t^{2t+4}\lambda^{t^2})$. ◀

We now augment the set Q_0 as before: for each cell $u \in \mathcal{G}$, define the set $P'_u \coloneqq P_u \setminus Q_0$ and move $|P'_u| \bmod t$ arbitrary points from P'_u to Q_0. At this point, the size of each of the sets P'_u is divisible by t, and we have added at most $t-1$ points from u, which were previously not in Q_0, to Q_0. We denote the augmented set by Q. We observe that the above asymptotic bound on the depth of $\mathrm{UBG}(Q)$ does not change, as the number of points added is negligible.

▶ **Theorem 11.** *Let P be a set of n points in the plane. The problem of finding a maximum matching in $H_t(P)$ can be reduced to the problem of finding a maximum matching in $H_t(Q)$, where $Q \subseteq P$ and the depth of* $\mathrm{UDG}(Q)$ *is $O(t^{2t+4}\lambda^{t^2})$.*

Proof. We compute a matching M in $H_t(P)$ as follows. Recall that for each cell u, the graph $H_t(P_u)$ is a clique, that is, any t points from P_u define a hyperedge of $H_t(P_u)$. We first compute, for each non-empty cell u, a perfect matching in $H_t(P'_u)$ and add its edges to M. Next we compute a maximum matching in $H_t(Q)$ and add its edges to M. The same argument as above shows that M is indeed a maximum matching in $H_t(P)$. ◀

2.3 Extensions

The following two properties capture the essential requirements of the sparsification stage.

1. Any set of t points lying in the same grid cell u forms a t-clique in $\mathrm{UBG}(P)$.
2. The points forming a t-clique in $\mathrm{UBG}(P)$ belong to a connected portion of the grid of some maximum constant size.

Consequently, the sparsification technique described above applies to a variety of t-uniform hypergraphs beyond $H_t(P)$.

For example, assume we are interested in t-cycles in $\mathrm{UBG}(P)$. That is, we wish to find as many t-cycles as possible, where each point can participate in at most one t-cycle. Then, the two properties continue to hold when t-clique is replaced by t-cycle, and we construct the corresponding t-uniform hypergraph H, where a subset T of t points determines a hyperedge if and only if there exists a permutation $(q_1, \ldots, q_t)$ of T such that $\|q_i - q_{i+1}\| \le 1$, for $i = 1, \ldots, t-1$, and $\|q_t - q_1\| \le 1$. Once we have H, we can apply the sparsification technique above.

Furthermore, the sparsification technique extends naturally to higher-dimensional spaces. We only need to set the side length of a grid cell to $\frac{1}{\sqrt{d}}$ in order to preserve the first property.

More broadly, the framework accommodates a wide range of geometric properties beyond those found in unit disk graphs. For instance, it applies to the t-uniform hypergraph whose hyperedges are all subsets of t points that can be covered by a disk of radius $1/2$, since the two key properties are trivially satisfied in this setting.

The sparsification technique also extends to intersection graphs defined by balls of varying radii, assuming all radii lie within a fixed range $[1, \Xi]$, since we may scale otherwise. In this setting, we define $H_t(P)$ as in the unit ball case – consider the t-uniform hypergraph where a subset of t points forms a hyperedge if the corresponding t balls induce a clique in the intersection graph, that is, every pair of balls intersects. The two key properties required for sparsification continue to hold: (1) any t balls centered within the same grid cell necessarily form a clique, as we choose the grid side length proportionally to the minimum possible radius of a ball, ensuring that all centers are sufficiently close; and (2) any t-clique in the ball intersection graph must be localized within a constant-sized neighborhood of grid cells, where the constant depends only on Ξ. Therefore, the sparsification framework applies directly in this setting, enabling approximation schemes for problems involving cliques in intersection graphs of non-uniform balls.

The sparsification framework further extends to intersection graphs defined by more general geometric objects, such as α-fat shapes. A family of objects is called α-fat if, for each shape S in the family, there exists a ball of radius r contained in S and a ball of radius at most αr that contains S. Assuming all objects in the family have those r-values in a bounded range $[1, \Xi]$ and constant fatness parameter α, we define the t-uniform hypergraph $H_t(P)$ analogously: a subset of t points forms a hyperedge if the corresponding shapes (associated with these points) induce a clique in the intersection graph. To apply sparsification, we must assign each object to a grid cell based on a single reference point. We take this point to be the center of the contained ball. This choice ensures that the first sparsification property holds: by choosing the grid side length proportionally to the minimum diameter, any t shapes whose contained-ball centers lie in the same grid cell necessarily intersect pairwise, forming a clique. For the second property, note that any point inside each shape S is at most $2\Xi\alpha$ away from the chosen center, as this is the maximal diameter of a circle that contains it. Therefore, for any other shape S' that intersects S, the center point chosen are at most $4\Xi\alpha$ apart, so the localization argument still applies. Thus, the sparsification technique generalizes to a wide class of fat objects, including disks, ellipses with bounded aspect ratio, and convex polygons with bounded complexity.

3 PTAS via local search

Recall that our goal is to compute a maximum matching in $H_t(P)$. By Theorem 11, it is sufficient to compute a maximum matching in $H_t(Q)$, where $Q \subseteq P$ is the set of bounded depth obtained in the sparsification stage. (By "bounded depth" we mean that the depth of UBG(Q) is bounded by some constant.) In this section we describe the second stage, in which, given $\varepsilon > 0$, we compute a matching whose size is at least $(1-\varepsilon)$ times the size of a maximum matching in $H_t(Q)$.

We first consider the original setting where P is a set of points in the plane and the hyperedges of $H_t(P)$ correspond to the t-cliques in UDG(P), for some constant $t \geq 3$. In this setting, we can apply *local search* to obtain a PTAS, as observed by Aschner et al. [1]. That is, we set $k = O(\varepsilon^{-2})$ (where k also depends on the constant parameters t and λ) and let M be any matching in $H_t(Q)$. Now, as long as it is possible to increase the size of M by replacing some $j \leq k-1$ hyperedges of M by $j+1$ hyperedges (which are currently not in M), such that M remains a matching in $H_t(Q)$, do so. Finally, return the matching M.

For the purpose of the analysis, let M^* denote a maximum matching in $H_t(Q)$. Extending the beautiful analysis of Mustafa and Ray [15], Aschner et al. [1] showed the following. Suppose that one can define a geometric bipartite graph $G = (B, R, E)$, where the blue vertices (i.e., those in B) represent the edges of $M \setminus M^*$ and the red vertices (those in R) represent the edges of $M^* \setminus M$, such that (i) if the hyperedges corresponding to $b \in B$ and $r \in R$ intersect (i.e., have a point in common), then $(b, r) \in E$ (this condition is called the *locality property*), and (ii) G has a strictly sublinear separator. Then $|M| \geq (1-\varepsilon)|M^*|$.

In the spirit of Aschner et al. [1], we define G as follows. For each hyperedge $e = (p_1, \ldots, p_t) \in E(H_t(Q))$, we arbitrarily pick one of its points as its representative. Using these representative points, we obtain the sets B and R, where B (resp., R) consists of the representative points whose hyperedges belong to $M \setminus M^*$ (resp., to $M^* \setminus M$). The graph G is then $\mathrm{UDG}_1(B \cup R)$, that is, the unit-disk graph on $B \cup R$, whose disks are of radius 1. Clearly, G has the desired two properties. Indeed, if the hyperedges represented by b and r have a point in common, say p, then both b and r are within distance at most 1 from p, and therefore within distance at most 2 from each other; that is their disks intersect one another. The second property holds as well, since $G = \mathrm{UDG}_1(B \cup R)$ is a subgraph of $\mathrm{UDG}_1(Q)$. However, by Lemma 10 and the paragraph following it, the depth of $\mathrm{UDG}_{1/2}(Q)$ is bounded by some constant $c = c(t, \lambda)$, which is easily seen to imply that the depth of $\mathrm{UDG}_1(Q)$ and of any of its subgraphs, such as G, is bounded by a suitable multiple of c. Finally, the results of Miller et al. [14] and of Smith and Wormald [16] imply that G has a separator of size $O(\sqrt{cn})$.

All these considerations, put together, lead to the desired PTAS; see Theorem 12 below.

We note that instead of considering t-cliques, we may consider any connected structure of constant size, such as t-cycles or t-stars.

Higher dimensions. As mentioned in Section 2.3, the sparsification technique extends to higher-dimensional spaces (with obvious modifications), so we only need to verify that local search is still applicable. Indeed, the locality property holds in any dimension d, i.e., if the hyperedges represented by b and r have a point in common, then the distance between b and r is at most 2, so the edge (b, r) is in G. Moreover, by the result of Miller et al. [14], $G = \mathrm{UBG}_1(B \cup R)$, defined analogously to the preceding definition, has a separator of size $O((cn)^{1-1/d})$, where $c = c(t, \lambda, d)$ is a constant. It follows that local search computes a $(1 - \varepsilon)$-approximation of a maximum matching in $H_t(Q)$. In conclusion, we have:

▶ **Theorem 12.** *Let P be a set of n points in $\mathbb{R}^d$, for any $d \geq 2$, let t be an integer parameter, and let $\varepsilon > 0$. Then we can compute a matching M in the hypergraph $H_t(P)$, as defined above, that satisfies $|M| \geq (1 - \varepsilon)|M^*|$, where M^* is a maximum matching in $H_t(P)$. The running time of the algorithm is polynomial in n, where the exponent depends on $1/\varepsilon$.*

Other cases. Consider, for example, the case where the hyperedges of $H_t(P)$ (where now $P \subset \mathbb{R}^2$) correspond to all subsets of P of size t that can be enclosed by a unit disk (of radius 1/2). Again, we only need to verify that local search is applicable in this case, as we already observed in Section 2.3 that the sparsification technique extends to cases of this kind. We define G as above, i.e., we first obtain the sets B and R by picking an arbitrary point from each hyperedge and then set $G = \mathrm{UDG}_1(B \cup R)$. As above, it is easy to verify that both the locality property and the small separator property hold, and therefore local search outputs a $(1 - \varepsilon)$-approximation of a maximum matching in $H_t(P)$.

References

1. Rom Aschner, Matthew J. Katz, Gila Morgenstern, and Yelena Yuditsky. Approximation schemes for covering and packing. In *Proc. 7th International Workshop on Algorithms and Computation (WALCOM)*, volume 7748 of *Lecture Notes in Computer Science*, pages 89–100. Springer, 2013. doi:10.1007/978-3-642-36065-7_10.
2. Brenda S. Baker. Approximation algorithms for NP-complete problems on planar graphs. *J. ACM*, 41(1):153–180, 1994. doi:10.1145/174644.174650.
3. Édouard Bonnet, Sergio Cabello, and Wolfgang Mulzer. Maximum matchings in geometric intersection graphs. *Discret. Comput. Geom.*, 70(3):550–579, 2023. doi:10.1007/s00454-023-00564-3.
4. Alberto Caprara and Romeo Rizzi. Packing triangles in bounded degree graphs. *Inf. Process. Lett.*, 84(4):175–180, 2002. doi:10.1016/S0020-0190(02)00274-0.
5. Timothy M. Chan and Sariel Har-Peled. Approximation algorithms for maximum independent set of pseudo-disks. *Discret. Comput. Geom.*, 48(2):373–392, 2012. doi:10.1007/S00454-012-9417-5.
6. Brent N. Clark, Charles J. Colbourn, and David S. Johnson. Unit disk graphs. *Discret. Math.*, 86(1-3):165–177, 1990. doi:10.1016/0012-365X(90)90358-O.

7 Marek Cygan. Improved approximation for 3-dimensional matching via bounded pathwidth local search. In *Proc. 65th Symposium on Foundations of Computer Science (FOCS)*, pages 509–518. IEEE Computer Society, 2013. doi:10.1109/FOCS.2013.61.

8 Cor A. J. Hurkens and Alexander Schrijver. On the size of systems of sets every t of which have an SDR, with an application to the worst-case ratio of heuristics for packing problems. *SIAM J. Discret. Math.*, 2(1):68–72, 1989. doi:10.1137/0402008.

9 Harry B. Hunt III, Madhav V. Marathe, Venkatesh Radhakrishnan, S. S. Ravi, Daniel J. Rosenkrantz, and Richard Edwin Stearns. NC-approximation schemes for NP- and PSPACE-hard problems for geometric graphs. *Journal of Algorithms*, 26(2):238–274, 1998. doi:10.1006/JAGM.1997.0903.

10 Richard M. Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, The IBM Research Symposia Series, pages 85–103. Plenum Press, New York, 1972. doi:10.1007/978-1-4684-2001-2_9.

11 Matthew J. Katz and Micha Sharir. Bottleneck matching in the plane. *Comput. Geom.*, 112:101986, 2023. doi:10.1016/J.COMGEO.2023.101986.

12 Richard J. Lipton and Robert E. Tarjan. A Separator Theorem for Planar Graphs. *SIAM Journal on Applied Mathematics*, 36(2):177–189, 1979. doi:10.1137/0136016.

13 Gordana Manić and Yoshiko Wakabayashi. Packing triangles in low degree graphs and indifference graphs. *Discret. Math.*, 308(8):1455–1471, 2008. doi:10.1016/J.DISC.2007.07.100.

14 Gary L. Miller, Shang-Hua Teng, William P. Thurston, and Stephen A. Vavasis. Separators for sphere-packings and nearest neighbor graphs. *Journal of the ACM*, 44(1):1–29, 1997. doi:10.1145/256292.256294.

15 Nabil H. Mustafa and Saurabh Ray. PTAS for geometric hitting set problems via local search. In *Proc. 25th Symposium on Computational Geometry (SoCG)*, pages 17–22. Association for Computing Machinery, 2009. doi:10.1145/1542362.1542367.

16 Warren D. Smith and Nicholas C. Wormald. Geometric separator theorems & applications. In *Proc. 39th Symposium on Foundations of Computer Science (FOCS)*, pages 232–243. IEEE Computer Society, 1998. doi:10.1109/SFCS.1998.743449.