On Efficiently Solvable Cases of Quantum k -SAT

Marco Aldi¹, Niel de Beaudrap², Sevag Gharibian^{3,4} , Seyran Saeedi⁴

¹ Department of Mathematics and Applied Mathematics, Virginia Commonwealth University, Richmond, VA, USA. E-mail: maldi2@vcu.edu

² Department of Computer Science, University of Oxford, Oxford, UK. E-mail: niel.debeaudrap@cs.ox.ac.uk

³ Department of Computer Science, Paderborn University, Paderborn, North-Rhine-Westphalia, Germany.

E-mail: sevag.gharibian@upb.de

⁴ Department of Computer Science, Virginia Commonwealth University, Richmond, VA, USA.

E-mail: saeedis@mymail.vcu.edu

Received: 28 February 2019 / Accepted: 30 April 2020

© The Author(s) 2020

Abstract: Estimating ground state energies of local Hamiltonian models is a central problem in quantum physics. The question of whether a given local Hamiltonian is *frustration-free*, meaning the ground state is the simultaneous ground state of *all* local interaction terms, is known as the Quantum k -SAT (k -QSAT) problem. In analogy to its classical Boolean constraint satisfaction counterpart, the NP-complete problem k -SAT, Quantum k -SAT is QMA₁-complete (for $k \geq 3$, and where QMA₁ is a quantum generalization of NP with one-sided error), and thus likely intractable. But whereas k -SAT has been well-studied for special tractable cases, as well as from a “parameterized complexity” perspective, much less is known in similar settings for k -QSAT. Here, we study the open problem of computing satisfying assignments to k -QSAT instances which have a “dimer covering” or “matching”; such systems are known to be frustration-free, but it remains open whether one can *efficiently compute* a ground state. Our results fall into three directions, all of which relate to the “dimer covering” setting: (1) We give a polynomial-time classical algorithm for k -QSAT when all qubits occur in at most two interaction terms or clauses. (2) We give a “parameterized algorithm” for k -QSAT instances from a certain non-trivial class, which allows us to obtain exponential speedups over brute force methods in some cases. This is achieved by reducing the problem to solving for a single root of a single univariate polynomial. An explicit family of hypergraphs, denoted Crash, for which such a speedup is obtained is introduced. (3) We conduct a structural graph theoretic study of 3-QSAT interaction graphs which have a “dimer covering”. We remark that the results of (2), in particular, introduce a number of new tools to the study of Quantum SAT, including graph theoretic concepts such as transfer filtrations and blow-ups from algebraic geometry.

A short version of this article has appeared as a long abstract in Proceedings of the 43rd International Symposium on Mathematical Foundations of Computer Science (MFCS), volume 117 of Leibniz International Proceedings in Informatics (LIPIcs), pp. 38:1–38:16, 2018.

1. Introduction

Estimating ground state energies is a central problem in quantum physics. Specifically, given a k -local Hamiltonian $H = \sum_i H_i$ acting on n qubits, and where each local interaction term or *clause* H_i acts non-trivially on a subset of at most k qubits,¹ the aim is to estimate the smallest eigenvalue of H . In 1999, Kitaev showed that this problem, dubbed the k -local Hamiltonian problem (k -LH), is complete for a quantum generalization of NP known as Quantum Merlin Arthur (QMA) [1]. This not only demonstrated that k -LH is likely intractable in the worst case, but also spawned an entire area of research at the intersection of mathematics, condensed matter physics, and computational complexity theory known as *Quantum Hamiltonian Complexity* (see e.g. [2–4] for surveys).

Let us explore this connection to classical complexity theory further. In the canonical NP-complete Boolean constraint satisfaction problem MAX- k -SAT, one is given a set of k -local Boolean functions or *constraints/clauses* $\phi := \{\phi_i\}_{i=1}^m$, such that each clause acts on k out of all n bits in the system. (Moreover, for MAX- k -SAT, each ϕ_i is a logical OR of k literals, such as $\phi_i = x_2 \vee \bar{x}_3 \vee \dots \vee x_1$, for $\bar{x}$ the negation of x .) Given such a set of clauses ϕ , the aim is to compute the maximum number of clauses which can be simultaneously satisfied over all assignments to the bits $x_1, \dots, x_n$. Note the *maximization* flavor of this problem; the correct generalization of this to the quantum setting is k -LH, since in k -LH the ground space may also be *frustrated*, meaning it may be impossible to satisfy *all* constraints simultaneously. If we instead demand that the ground space be *frustration-free*, i.e. the ground state $|\psi\rangle \in (\mathbb{C}^2)^n$ lives in the simultaneous ground spaces of each constraint H_i , then we are more accurately considering a quantum analogue of the NP-complete problem k -SAT, in which the question is given ϕ , can we simultaneously satisfy *all* constraints? This quantum analogue, denoted Quantum k -SAT (k -QSAT), was introduced by Bravyi in 2006 [5], and is QMA₁-complete for $k \geq 3$ [5, 6]. (See “previous work” below for further known results.) (Here, QMA₁ is QMA with perfect completeness, meaning in a YES instance of a given problem, the quantum proof must be accepted with certainty by the quantum verifier.)

Given the importance of such constraint satisfaction problems (classical or quantum) and their intractability (up to standard complexity theoretic conjectures such as $P \neq NP$ and $BQP \neq QMA_1$), much effort has been devoted by the classical community to approaches for MAX- k -SAT and k -SAT, including approximation algorithms, heuristic algorithms, and exact algorithms.² A particular direction of research involving *exact* algorithms is what restrictions to the input are sufficient to allow polynomial-time solvability. In this paper, we focus on this theme, and ask:

Which special cases of k -QSAT can be solved efficiently on a classical computer?

Unfortunately (but also intriguingly), this problem appears to be markedly more difficult quantumly than classically. Consider the motivating example for this work: Classically, if each clause c of a k -SAT instance can be matched with a unique variable v_c (such that v_c appears in clause c), then clearly the k -SAT instance is satisfiable. For example, in the instance $\phi = \{\phi_1 = x_1 \vee \bar{x}_2, \phi_2 = \bar{x}_1 \vee x_2\}$, we can match x_1 with

¹ For clarity, in the general formulation of k -LH, there is no geometric restriction on where the interaction terms act (e.g. on a 1D chain or a 2D lattice).

² The distinction between these three approaches is as follows. Approximation algorithms run in polynomial time, and provably provide a solution which is within some fixed prescribed relative error of optimal. A heuristic typically is also polynomial time, but does not provide any guarantee as to the accuracy of the solution found. An exact algorithm produces the optimal solution, but either takes superpolynomial time (as expected) but better than brute force time, or exactly solves a restricted case of the problem in polynomial time.

clause ϕ_1 and x_2 with clause ϕ_2 . For any such “matched” k -SAT instance, finding a solution is trivial: Set variable v_c to satisfy clause c . In our example, this means setting $x_1 = 1$ and $x_2 = 1$. (Note that the matching can be found efficiently via a reduction to network flow, which can then be solved in polynomial time via, e.g., the Ford–Fulkerson algorithm [7].) Let us compare and contrast with the quantum setting. Quantumly, it has been known [8] since 2010 that k -QSAT instances with “matchings” (also called a “dimer covering” in [8]) are also satisfiable/frustration-free. Moreover, a ground state exists which is of a particularly nice form—it can be assumed to be a tensor product state. Despite this, *finding* said ground state efficiently has proven elusive (indeed, the proof of [8] is *non-constructive*). In complexity theoretic terms, we have a trivial NP *decision* problem (since the answer is always YES, the system is frustration-free, and the proof is a tensor product state which can be described via polynomially many bits) whose analogous *search* version is not known to be efficiently solvable. It should be pointed out here that classically, the question of whether decision versus search complexity coincide for NP problems is a long-standing open problem (see, e.g., [9]), and so it is perhaps surprising that this phenomenon seems to potentially also play a role in “matched” QSAT. This is the starting point of the present work.

Results and techniques. Our results fall under three directions, all of which are related to k -QSAT with matchings, and which can very broadly be summarized as follows: (1) We show how to efficiently solve k -QSAT with sufficiently bounded occurrence of variables, (2) we give a general framework for parameterized algorithms for Quantum k -SAT, which in principle applies to *any* Quantum k -SAT instance, (3) we attempt to characterize the set of all 3-uniform hypergraphs (which encode interaction graphs for 3-QSAT, formal definitions given shortly) with a “dimer covering”, or formally, a “system of distinct representatives”. Results (1) and (2) partially resolve the open problem of [8] by giving efficient classical algorithms for special cases of “matched QSAT”, and result (2) in particular gives (as far as we are aware) the first known parameterized algorithm for QSAT, obtaining exponential speedups over brute force in some cases (Sect. 4.4). However, we are not yet able to fully resolve whether the decision and search complexity of “matched QSAT” coincide; indeed, as the structural graph theoretic study of (3) suggests, the geometry of all 3-QSAT instances with matchings/dimer coverings appears quite complex.

Definitions and precise statements of results. To more precisely state our results, we now require formal definitions. For this, we first define Quantum k -SAT (k -QSAT) [5] and the notion of a system of distinct representatives (SDR). For k -QSAT, the input is a two-tuple $\Pi = (\{\Pi_i = |\psi_i\rangle\langle\psi_i|\}_i, \alpha)$ of rank 1 projectors or *clauses*³ $\Pi_i \in \mathcal{L}(\mathbb{C}^2)^{\otimes k}$, each acting non-trivially on a set of k (out of n) qubits, and non-negative real number $\alpha > 1/p(n)$ for some fixed polynomial p . The output is to decide whether there exists a satisfying assignment on n qubits $|\psi\rangle \in (\mathbb{C}^2)^{\otimes n}$, i.e. to distinguish between the cases $\Pi_i|\psi\rangle = 0$ for all i (YES case), or whether $\langle\psi|\sum_i \Pi_i|\psi\rangle \geq \alpha$ (NO case). Note that k -QSAT generalizes k -SAT. As for a *system of distinct representatives (SDR)* (which is formal terminology for the notion of “matching” for k -QSAT instances, as adopted from the combinatorics literature [10]), given a set system such as a hypergraph $G = (V, E)$, an SDR is a set of vertices $V' \subseteq V$ such that each edge in $e \in E$ is paired with a distinct

³ The original definition [5] of k -QSAT did not require each Π_i to be rank 1. As in [8], we require the rank-1 condition to make our definitions and results well-defined and valid, respectively. Nevertheless, we do allow one to “stack” multiple clauses on a fixed set of k qubits to simulate higher rank clauses.

vertex $v_e \in V'$ such that $v_e \in e$. In previous work on QSAT, an SDR has been referred to as a “dimer covering” [8].

1. Quantum k -SAT with bounded occurrence of variables. Our first result concerns the natural restriction of limiting the number of times a variable can appear in a clause. For example, 3-SAT with at most 3 occurrences per variable is known to be NP-hard. We now show this threshold is tight,⁴ in the following sense.

Theorem 1. *There exists a polynomial time classical algorithm which, given an instance Π of k -QSAT in which each variable occurs in at most 2 clauses, outputs a satisfying product state if Π is satisfiable, and otherwise rejects. Moreover, the algorithm works for clauses ranging from 1-local to k -local in size.*

To show this, our idea is to “partially reduce” the k -QSAT instance to a 2-QSAT instance. We then use the transfer matrix techniques of [5, 11, 12] (particularly the notion of *chain reactions* from [12]), along with a new notion of “fusing” chain reactions, to deal with the remaining clauses of locality at least 3 in the instance.

Whereas *a priori* this setting may seem unrelated to the open question of computing solutions to k -QSAT instances with SDRs, we connect the two settings as follows. Denote the *interaction hypergraph* $G = (V, E)$ of a k -QSAT instance as a k -uniform hypergraph (i.e. all edges have size precisely k), in which the vertices correspond to qubits, and each clause c acting on a set of k qubits S_c , is represented by a hyperedge of size k containing the vertices corresponding to S_c .

Theorem 2. *Let $G = (V, E)$ be a hypergraph with all hyperedges of size at least 2, and such that each vertex has degree at most 2. Then, G has an SDR.*

Thus, Theorem 1 resolves the open question of [8] for k -QSAT instances with SDRs in which additionally (1) each variable occurs in at most two clauses and (2) there are no 1-local clauses. (The latter of these restrictions is necessary, as allowing edges of size 1 easily makes Theorem 2 false in general; see Sect. 3.1.)

2. Parameterized algorithms for Quantum k -SAT. Our next result, and the main contribution of this paper, gives a parameterized algorithm⁵ for explicitly computing (product state) solutions for a non-trivial class of k -QSAT instances. As discussed in Sect. 4.4, this algorithm in some cases provides an *exponential* speedup over brute force diagonalization.

Transfer filtrations. To sketch the idea of the algorithm at a high level, we first introduce a new graph theoretic notion of a *transfer filtration* for a k -uniform hypergraph $G = (V, E)$ (see Definition 9 for a formal definition). Intuitively, a “transfer filtration of type $b > 0$ ” should be thought of as a set of b qubits (out of all n qubits the k -QSAT instance acts on) which form the “hard core” or *foundation* of the instance. Roughly, this is in the sense that *if* one could determine what states to assign to the qubits in the foundation, then

⁴ In addition to being NP-hard, whether k -QSAT with at most 3 occurrences of variables is also QMA₁-hard is not clear. For example, the QMA₁-completeness constructions of [5, 6] utilize more than 3 occurrences per variable.

⁵ Roughly, parameterized complexity characterizes the complexity of computational problems with respect to specific *parameters* of interest other than just the input size. For example, for an NP-complete graph theoretic problem Π on n vertices, a parameterized algorithm might run in time polynomial in n , but potentially exponentially with respect to some other graph parameter, such as the treewidth of the input graph. In this example, Π would be tractable on input graphs of logarithmic treewidth. Parameterized complexity is the focus of an entire research community classically; see “previous work” for further details.

Fig. 1. (Top left) A chain. (Top right) A cycle. (Bottom middle) A semicycle. Note that, in contrast to the cycle, the semicycle is missing the blue edge $\{v_1, v_2, v_5\}$

resolving the remainder of the instance would be “easy” (this is not entirely accurate, hence the quotes on “easy”).

Algorithm sketch. With the notion of transfer filtration in hand, our framework for attacking k -QSAT can be sketched at a high level as follows.

1. (Decoupling step) First, given a k -QSAT instance Π on G with transfer filtration of type b , we “blow-up” Π to a larger, *decoupled* instance Π^+ (Decoupling Lemma, Lemma 18).
2. (Solve decoupled instance) The decoupled nature of Π^+ makes it “easier” to solve (Transfer Lemma, Lemma 28), in that any assignment to the b “foundation” qubits can be *extended* to a solution to all of Π^+ . This raises the question—how does one map the solution of Π^+ back to a solution of Π ?
3. (Map solutions back to original instance) We next give a set of “qualifier” constraints $\{h_s\}$ (Qualifier Lemma, Lemma 30) acting on *only* the b foundation qubits, with the following strong property: If a (product state) assignment $\mathbf{v}$ to the b foundation qubits satisfies the constraints $\{h_s\}$, then not only can we extend $\mathbf{v}$ via the Transfer Lemma to a full solution for Π^+ as in Step 1 above, but we can also map this extended solution back to one for the *original* k -QSAT instance Π . (For clarity, there is an additional required technical condition on the *transfer functions* g_i in the Transfer Lemma, which we circumvent in our main Theorem 38.)

Once the framework above is developed, we show that it applies to the non-trivial family of k -QSAT instances whose k -uniform hypergraph $G = (V, E)$ has a transfer filtration of type $b = |V| - |E| + 1$. This family includes, e.g., the semi-cycle of Fig. 1, (a slight modification of) the tiling of the torus (Fig. 2), and “fir tree” (Fig. 3). Our main result (Theorem 38) says the following: For any k -QSAT instance Π on such a G and whose constraints are generic (see Sect. 4), computing a (product state) solution to Π reduces to solving for a root of a *single univariate* (see Remark 41) polynomial P —*any* root thereof (which always exists if the field $\mathbb{K}$ is algebraically closed) can then be extended back to a full solution for Π .

Advantages of framework. The key advantage of this approach, and what makes it a parameterized algorithm, is the following—the *degree* of polynomial P , and hence the runtime of the algorithm, scale exponentially only in the foundation size b and a “radius” parameter r of the transfer filtration (assuming $k \in O(1)$, see Eq. (10) for an explicit bound on runtime). Thus, given a transfer filtration where b and r are at most logarithmic,

Fig. 2. A 3×3 tiling of the torus. Note the closed boundary conditions, i.e. we identify two vertices with the same label as being the same vertex

Fig. 3. (Fir tree) A 3-uniform hypergraph G of transfer type b to which Theorem 38 applies. G has $m = 31$ edges, $n = 42$ vertices, and foundation size $b = 12$ (the vertices marked with a star are foundation vertices). The hypergraph's name was chosen due to its (vague) resemblance to a fir tree

finding a (product state) solution to k -QSAT reduces to solving for a single root over $\mathbb{C}$ for a single univariate polynomial P of polynomial degree, which can be done in polynomial time [13, 14]. Indeed, in Sect. 4 we give a non-trivial family of k -uniform hypergraphs, denoted Crash (Fig. 4), for which our algorithm runs in polynomial time, whereas brute force diagonalization would require exponential time.

If, on the other hand, the foundation size b and radius r of the input k -QSAT instance are *superlogarithmic*, then our algorithm requires superpolynomial time. Even here, however, the framework yields a new result: it gives a *constructive* proof that all k -QSAT instances satisfying the preconditions of Theorem 38 have a (product state) solution. In particular, in Corollary 43 and Theorem 50, we observe that such hypergraphs must have

Fig. 4. Depiction of 3-uniform crash hypergraph $C_{3,3}$. Generally, $C_{r,k}$ has an exponential separation between the filtration radius and foundation size versus number of vertices and edges

SDRs, and so we constructively reproduce the result of [8] that any 3-QSAT instance with an SDR is satisfiable by a product state (again, assuming the additional conditions of Theorem 38 are met).

Finally, although this result stems primarily from tools of projective algebraic geometry (AG), the presentation herein avoids any explicit mention of AG terminology (with the exception of defining the term “generic” in Sect. 4.3) to be accessible to readers without an AG background. For completeness, a brief overview of the ideas in AG terms is given at the end of Sect. 4.

3. A study of 3-uniform hypergraphs with SDRs. Given that computing ground states for k -QSAT instances with SDRs appears challenging, in the hope of helping guide future studies on the topic, our final contribution aims to better understand the set of all k -QSAT instances with SDRs, particularly in the boundary case when $|E| = |V|$. Unfortunately, this in itself appears to be a difficult task (if not potentially impossible, see comments about a “finite characterization” below). Our results here are as follows. We first give various characterizations involving intersecting families (i.e. when each pair of edges has non-empty intersection). We then study the setting of *linear* hypergraphs (i.e. each pair of edges intersects in at most one vertex), which are generally more complex. (For example, the set of edge-intersection graphs of 3-uniform linear hypergraphs is known not to have a “finite” characterization in terms of a finite list of forbidden induced subgraphs [15].) We study “extreme cases” of linear hypergraphs with SDRs, such as the Fano plane and “tiling of the torus”, and in contrast to these two examples, demonstrate a (somewhat involved) linear hypergraph we call the iCycle which also satisfies the *Helly property* (which generalizes the notion of “triangle-free”).

A main conclusion of this study is that even with multiple additional restrictions in place (e.g. linear, Helly), the set of 3-uniform hypergraphs with SDRs remains non-trivial. To complement these results, we show how to fairly systematically construct large linear hypergraphs with $|E| = |V|$ without SDRs. We hope this work highlights the potential complexity involved in dealing with even the “simple” case of 3-QSAT with SDRs.

Discussion. Regarding our parameterized algorithm, we stated earlier that our framework applies to the non-trivial family of k -QSAT instances whose k -uniform hypergraph $G = (V, E)$ has a transfer filtration of “type $b = |V| - |E| + 1$ ”. Let us elaborate on this further, and in particular stress which parts of our framework require this additional assumption on b , and which parts of the framework apply to arbitrary instances of k -QSAT.

First, our notions of transfer filtrations and blow-ups apply to *any* instance of k -QSAT (and thus also⁶ k -SAT), including QMA_1 -complete instances. (For example, every k -uniform hypergraph has a trivial foundation obtained by iteratively removing vertices until the resulting set contains no edges. A key question⁷ is how *small* the foundation size b and radius r of the filtration can be chosen for a given hypergraph, as our algorithm’s runtime scales exponentially in these parameters; see Eq. (10).) More precisely, our techniques in Sect. 4, up to and including the Qualifier Lemma, apply to arbitrary k -QSAT instances. If one also assumes that constraints are *generic* (but still on an arbitrary hypergraph), then the Surjectivity Lemma (Lemma 36) also holds.

The remaining question is Phase 3 of the framework (“Map solutions back to original instance” above)—when can *local* solutions to the qualifier constraints be extended to *global* solutions for the entire k -QSAT instance? Answering this requires solving for a common root of a set of high-degree multi-variate polynomials (hence the connection to algebraic geometry). This is where we consider the restriction to input k -uniform hypergraphs of transfer type $b = |V| - |E| + 1$, which allows us to reduce the entire framework to the solution of a *single univariate* (high-degree) polynomial (to which one can now apply the root-finding algorithm of Schönhage [14]; see Sect. 4.4). We stress that the family of k -QSAT instances with $b = |V| - |E| + 1$ is non-trivial, and includes the semi-cycle (Fig. 1), tiling of the torus (Fig. 2), fir tree (Fig 3), and crash (Fig. 4) families of hypergraphs. In particular, for the latter family we obtain exponential speedups over brute force diagonalization in Sect. 4.4. For clarity, all runtimes in this paper are rigorous worst-case runtimes.

Previous work. Quantum k -SAT was introduced by Bravyi [5], who gave an efficient (quartic time) algorithm for 2-QSAT, and showed that 4-QSAT is QMA_1 -complete. Subsequently, Gosset and Nagaj [6] showed that Quantum 3-SAT is also QMA_1 -complete, and independently and concurrently, Arad, Santha, Sundaram, Zhang [17] and de Beaudrap, Gharibian [12] gave linear time algorithms for 2-QSAT. The original inspiration for this paper was the work of Laumann, Läuchli, Moessner, Scardicchio and Sondhi [8], which showed *existence* of a product state solution for any k -QSAT instance with an SDR. Thus, the decision version of k -QSAT with SDRs is in NP and trivially efficiently solvable. However, whether the search version (i.e. compute an explicit satisfying

⁶ For the special case of k -SAT, note that it is not *a priori* clear that having a transfer filtration with a small foundation suffices to solve the system trivially. This is because the genericity assumption on constraints, which k -SAT constraints do not satisfy, is required to ensure that any assignment to the foundation propagates to all bits in the instance. Thus, the brute force approach of iterating through all 2^b assignments to the foundation does not obviously succeed.

⁷ Whereas our algorithm’s runtime benefits greatly from small foundation size, a natural question is whether a “small” versus “large” foundation generally should make k -QSAT easier or harder. Intuitively, indeed one expects a smaller foundation to yield “easier” instances of k -QSAT, as recall the foundation captures the hard “core” of the k -QSAT instance (in that, roughly speaking, assignments to the foundation qubits generically force assignments onto the remaining, non-foundation, qubits). In the “Open Questions” paragraph of Sect. 1, we conjecture that computing the minimum foundation size is NP-hard of a given input hypergraph is actually NP-hard (similar to how computing common parameterized algorithmic parameters such as treewidth is NP-hard [16]).

assignment) is also polynomial-time solvable remains open. The question of whether the decision and search complexities of NP problems are the same is a longstanding open problem in complexity theory; conditional results separating the two are known (see e.g. Bellare and Goldwasser [9]).

In terms of approaches to k -QSAT, a complementary and well-studied tool which should also be mentioned is the Quantum Lovász Local Lemma (QLLL), introduced by Ambainis, Kempe, and Sattath [18]. The QLLL non-constructively guarantees frustration-freeness of the input Hamiltonian under certain conditions. Subsequently, Schwarz, Cubitt and Verstraete [19] and Sattath and Arad [20] provided *constructive* (i.e. algorithmic) versions of QLLL under the assumption of commuting local terms [19, 20], and more recently Gilyén and Sattath [21] gave an algorithm QLLL for non-commuting terms by introducing a uniformly gapped assumption. The QLLL has since been extended to Shearer’s bound by Sattath, Morampudi, Laumann and Moessner [22], in that a quantum analogue of Shearer’s bound is a sufficient criterion for frustration-freeness. Most recently, He, Li, Sun, and Zhangit showed [23] that in fact Shearer’s bound is tight in the quantum setting. We remark that the approaches in these works appear distinct from those taken here; for example, the QLLL allows certain cases of k -QSAT to be resolved even when it is not clear the ground space contains a tensor product state, whereas our algorithms search for product state solutions (guaranteed to exist for any k -QSAT instance with an SDR, which recall was the motivation for this work). This also means algorithmic QLLL results such as [21] use *quantum* algorithms to prepare a ground state, whereas our framework here is classical. As a result, our work should be viewed as a complementary approach for tackling k -QSAT instances, which introduces genuinely distinct tools such as parameterized algorithms and blowups from algebraic geometry.

Finally, in terms of classical k -SAT, in stark contrast to k -QSAT, solutions to k -SAT instances with an SDR can be efficiently computed. As for parameterized complexity, classically it is a well-established field of study (see, e.g., [24] for an overview). The parameterized complexity of SAT and #SAT, in particular, has been studied by a number of works, such as [25–31], which consider a variety of parameterizations including based on tree-width, modular tree-width, branch-width, clique-width, rank-width, and incidence graphs which are interval bipartite graphs. Regarding parameterized complexity of *Quantum* SAT, as far as we are aware, our work is the first to initiate a “formal” study of the subject. However, we should be clear that existing works in Quantum Hamiltonian Complexity [3, 4] have long implicitly used “parameterized” ideas. For example, Markov and Shi [32] give a classical simulation algorithm for quantum circuits whose runtime scales polynomially in the size of the circuit, but exponentially in its treewidth. Their algorithm is based on *tensor networks* (e.g. [33]), whose bond dimension can be viewed as a parameter constraining the complexity of the contraction. In particular, 1d tensor networks whose bond dimension is at most polynomial in the length of the chain can be contracted efficiently (a fact used, for example, in the polynomial time algorithm of Landau, Vidick and Vazirani [34] for solving the 1D gapped Local Hamiltonian [1] problem). This is in stark contrast to 2D tensor networks (often denoted PEPS, for Projected Entangled Pair States), which are #P-hard to contract even for constant bond dimension [35].

Open questions. We close with a number of open questions. We showed that k -QSAT with at most 2 occurrences per variable is efficiently solvable (Theorem 1), but 3 occurrences per variable is well-known to be at least NP-hard (for 3-SAT, which is a special case of 3-QSAT). Is 3-QSAT with at most 3 occurrences per variable similarly QMA₁-complete? Can ideas from classical parameterized complexity be generalized to the

quantum setting? We have developed a number of tools herein for studying Quantum SAT—can these be applied in more general settings, for example beyond the families of k -QSAT instances considered in Theorem 38? The “parameters” in our results of Sect. 4 include the radius of a transfer filtration—whether a transfer filtration (of a fixed type b) of *minimum* radius can be computed efficiently, however, is left open for future work. Similarly, it is not clear that given $b \in \mathbb{N}$, the problem of deciding whether a given hypergraph G has a transfer filtration of type at most b is in P. We conjecture, in fact, that this latter problem is NP-complete. Finally, the question of whether solutions to arbitrary instances of k -QSAT with SDRs can be computed efficiently (recall they are guaranteed to exist [8]) remains open.

Organization. Section 2 gives basic notation and definitions. Section 3 gives an efficient algorithm for 3-QSAT with bounded occurrence of variables, which also introduces the notions of transfer matrices (which are generalized via transfer functions in Sect. 4). Our main result is given in Sect. 4, and concerns a new parameterized complexity approach for solving k -QSAT. The algorithm and its runtime, along with the study of asymptotic speedups, are given in Sect. 4.4. Finally, Sect. 5 conducts a structural graph theoretic study of hypergraphs with SDRs.

2. Preliminaries

Notation. For complex Euclidean space $\mathcal{X}$, $\mathcal{L}(\mathcal{X})$ denotes the set of linear operators mapping $\mathcal{X}$ to itself. For unit vector $|\psi\rangle \in \mathbb{C}^2$, the unique orthogonal unit vector (up to phase) is denoted $|\psi^\perp\rangle$, i.e. $\langle\psi|\psi^\perp\rangle = 0$.

We now give some definitions from graph theory, some of which are used primarily in Sect. 5.

Definition 3 (*Hypergraph*). A *hypergraph* is a pair $G = (V, E)$ of a set V (vertices), and a family E (edges) of subsets of V . If each vertex has degree d , we say G is d -regular. Alternatively, when convenient we use $V(G)$ and $E(G)$ to denote the vertex and edge sets of G , respectively. A *simple* hypergraph has no repeated edges, i.e. if $e_i \subseteq e_j$, then $i = j$. We say G is k -uniform if all edges have size k .

Definition 4 (*Chain* [36], Fig. 1). A k -uniform hypergraph $G = (V, E)$ is a *chain* if there exists a sequence $(v_1, v_2, \dots, v_l) \in V^l$ for $l \geq n$ such that (1) the sequence contains all elements of V at least once, (2) $v_1 \neq v_l$, and (3) $E = \bigcup_{1 \leq i \leq l-k+1} e_i$ for distinct edges $e_i = \{v_i, v_{i+1}, \dots, v_{i+k-1}\}$. The length of the chain G is $m = l - k + 1$.

Definition 5 (*Cycle* [36], Fig. 1). A k -uniform hypergraph $G = (V, E)$ is a *cycle* if there exists a sequence $(v_1, v_2, \dots, v_l) \in V^l$ for $l \geq n$ such that (1) the sequence contains all elements of V at least once, (2) for all $1 \leq i \leq l$, $e_i = \{v_i, v_{i+1}, \dots, v_{i+k-1}\}$ are distinct edges in E , where indices are understood modularly. The length of the cycle G is $m = l$.

Definition 6 (*Semicycle* [36], Fig. 1). A k -uniform hypergraph $G = (V, E)$ is a *semicycle* if there exists a sequence $(v_1, v_2, \dots, v_l) \in V^l$ for $l \geq n$ such that (1) the sequence contains all elements of V at least once, (2) $v_1 = v_l$, and (3) for all $1 \leq i \leq l - k + 1$, $e_i = \{v_i, v_{i+1}, \dots, v_{i+k-1}\}$ are distinct edges of G . The length of the semicycle is $m = l - k + 1$.

Fig. 5. A 3-uniform hypergraph which is a tight star of size 3

Definition 7 (*Tight star* [36], Fig. 5). A k -uniform hypergraph $G = (V, E)$ is a *tight star* if there exists a set of $k - 1$ vertices S , such that the intersection of any pair of edges is precisely S .

Definition 8 (*t -stacked set*). Let $G = (V, E)$ be a k -uniform hypergraph. A subset $S \subseteq E$ of t edges is called a *t -stacked set* if every edge in S is incident to the same k vertices in V , i.e. $|\bigcup_{e \in S} e| = k$.

3. Quantum SAT with Bounded Occurrence of Variables

In this section, we study k -QSAT when each qubit occurs in at most two constraints. For this, we first recall tools from the study of Quantum 2-SAT [5, 11, 12]. Recall that throughout this paper, each clause is assumed to be rank 1 (in general, however, one can “stack” multiple rank 1 clauses on a set of vertices to simulate a higher rank clause).

Transfer matrices, chain reactions, and cycle matrices. For any rank-1 constraint $\Pi_i = |\psi\rangle\langle\psi| \in \mathcal{L}((\mathbb{C}^2)^{\otimes k})$, consider Schmidt decomposition $|\psi\rangle = \alpha|a_0\rangle|b_0\rangle + \beta|a_1\rangle|b_1\rangle$, where $|a_i\rangle \in (\mathbb{C}^2)^{\otimes(k-1)}$ lives in the Hilbert space of the first $k - 1$ qubits and $|b_i\rangle \in \mathbb{C}^2$ the last qubit. Then, the *transfer matrix* $T_\psi : (\mathbb{C}^2)^{\otimes k-1} \mapsto \mathbb{C}^2$ is given by $T_\psi = \beta|b_0\rangle\langle a_1| - \alpha|b_1\rangle\langle a_0|$. In words, given any assignment $|\phi\rangle$ to the first $k - 1$ qubits, if $T_\psi|\phi\rangle \in \mathbb{C}^2$ is non-zero, then it is the *unique* assignment to qubit k (given $|\phi\rangle$ on qubits 1 to $k - 1$) which satisfies Π_i .

In the special case of $k = 2$, transfer matrices are particularly useful. Consider first a 2-QSAT interaction graph (which is a 2-uniform hypergraph, or just a graph) $G = (V, E)$ which is a path, i.e. a sequence of edges $e_1 = (v_1, v_2)$, $e_2 = (v_2, v_3)$, $\dots$, $e_m = (v_{m-1}, v_m)$ for distinct $v_i \in V$, and where edge e_i corresponds to constraint $|\psi_i\rangle$. Then, any assignment $|\phi\rangle \in \mathbb{C}^2$ to qubit 1 *induces a chain reaction (CR)* in G , meaning qubit 2 is assigned $T_{\psi_1}|\phi\rangle$, qubit 3 is assigned $T_{\psi_2}T_{\psi_1}|\phi\rangle$, and so forth. If this CR terminates before all qubits labelled by V receive an assignment, which occurs if $T_{\psi_i}|\phi'\rangle = 0$ for some i , this means that constraint i (acting on qubits i and $i + 1$) is satisfied by the assignment $|\phi'\rangle$ to qubit i alone, and no residual constraint is imposed on qubit $i + 1$. Thus, the graph G is reduced to a path $e_{i+1}, \dots, e_m$. In this case, we say the CR is *broken*. Note that if G is a path, then it is a satisfiable 2-QSAT instance with a product state solution.

Finally, consider a 2-QSAT instance whose interaction graph G is a cycle $C = (v_1, \dots, v_{m+1})$ with m . Then, a CR induced on vertex v_1 with any assignment $|\psi\rangle \in \mathbb{C}^2$ will in general propagate around the cycle and impose a consistency constraint on v_1 . Formally, denote the product $T_C = T_{\psi_m} \cdots T_{\psi_1} \in \mathcal{L}(\mathbb{C}^2)$ as the *cycle matrix* of C . Then, if the cycle matrix is not the zero matrix, it can be shown that the satisfying assignments for the cycle are precisely the eigenvectors of T_C . (If $T_C = 0$, any assignment on v_1

Fig. 6. (Top left) The initial 3-QSAT instance, denoted $H(G)$ for the purposes of this example, and where each hyperedge is an arbitrary (rank 1) constraint. (Top middle) The hypergraph obtained after applying Step 2 to vertices $v_1, v_2, v_3, v_4, v_5, v_6$. (Top right) The hypergraph obtained by applying Step 2 to vertex v_7 . (Bottom left) The hypergraph obtained by applying Step 3bii to edge (v_8, v_9) . Applying Step 3bii to edge (v_{10}, v_{11}) would thus reduce us to a 2-QSAT cycle, which is then solved via Step 3bi. Note the order of vertices/edges processed in this example is arbitrary, and was chosen here simply for symmetry. (Bottom right) A pseudo-line graph corresponding to the hypergraph in the top left image

will only propagate partially around the cycle, thus decoupling the cycle into two paths.) Thus, if G is a cycle, then it also has a product state solution.

In this section, when we refer to “solving the path or cycle”, we mean applying the transfer matrix techniques above to efficiently compute a product state solution to the path or cycle.

k-QSAT with bounded occurrence of variables. We now restate and prove Theorem 1. Subsequently, we demonstrate the algorithm on an example (Fig. 6) and discuss its applicability more generally. Section 3.1 shows the connection between Theorem 1 and k -QSAT with SDRs.

Theorem 1. *There exists a polynomial time classical algorithm which, given an instance Π of k -QSAT in which each variable occurs in at most two clauses, outputs a satisfying product state if Π is satisfiable, and otherwise rejects. Moreover, the algorithm works for clauses ranging from 1-local to k -local in size.*

Proof. We begin by setting terminology. Let Π be an instance of k -QSAT with k -uniform interaction graph $G = (V, E)$. For any clause c , let Q_c denote the set of qubits acted on c , i.e. Q_c is the edge in G representing c . We say c is *stacked* if Q_c is contained within another edge/clause $Q_{c'}$, i.e. if there exists $c' \neq c$ such that $Q_c \subseteq Q_{c'}$. For a qubit v , we use shorthand $|v\rangle$ to denote the current assignment from $\mathbb{C}^2$ to v . For a clause c , $|c\rangle$ denotes the bad subspace of c , i.e. clause c is given by rank-1 projector $I - |c\rangle\langle c|$. The set of clauses vertex v appears in is denoted C_v . For any assignment $|v\rangle$, we introduce shorthand $S_{|v\rangle} = \{|v\rangle\langle c| \mid c \in C_v\} \subseteq \bigcup_{i=0}^{k-1} \mathbb{C}^{2^i}$, where recall c can be a clause on $1, \dots, k$ qubits,

and we implicitly assume $\langle v|$ acts as the identity on the qubits of c which are not v . For example, if c acts on qubits v, w, x , then $\langle v|c\rangle = (\langle v| \otimes I_{w,x})c \in \mathbb{C}^4$ is the residual constraint on qubits w, x , given assignment $|v\rangle$ to v . Thus, $S_{|v\rangle}$ is the set of constraints we obtain by taking the clauses in C_v , and projecting down qubit v in each clause onto assignment $|v\rangle$. As a result, the clauses in $S_{|v\rangle}$ no longer act on v . The algorithm we design will satisfy that the only possible element of $\mathbb{C}$ in $S_{|v\rangle}$ is 0, which can only be obtained by projecting a constraint $|c\rangle \in \mathbb{C}^2$ onto its orthogonal complement to satisfy it; thus, we can assume without loss of generality that $S_{|v\rangle} \subseteq \{\langle v|c\rangle \mid c \in C_v\} \subseteq \bigcup_{i=1}^{k-1} \mathbb{C}^{2^i}$. Finally, we say two 1-local clauses $|c\rangle, |c'\rangle \in \mathbb{C}^2$ *conflict* if $|c\rangle$ and $|c'\rangle$ are linearly independent (i.e. $|c\rangle\langle c| + |c'\rangle\langle c'|$ has an empty null space, meaning there is no satisfying assignment).

The algorithm proceeds as follows. Let Π satisfy the conditions of our claim. For clarity, any time a CR on a path is broken by a transfer matrix T_ψ on edge (u, v) , i.e. $T_\psi|u\rangle = 0$, we shall implicitly assume we continue by choosing assignment $|0\rangle$ on v to induce a new CR and continue solving the path. (This is important for the correctness analysis later in which Step 3 must not create a 1-local constraint.)

Statement of algorithm. The intuitive idea behind the algorithm is to “partially reduce” Π to a 2-QSAT instance, and use the transfer matrix techniques outlined above to solve this segment. Combining this with a new notion of fusing CRs, the technique can be applied iteratively to reduce k -local constraints to 2-local ones until the entire instance is solved. More informally, the algorithm exploits the extra degrees of freedom which arise when each qubit appears in at most two clauses. In particular, by considering clauses in the “right” order below, we are able to (iteratively) greedily satisfy a given clause locally, and propagate the implications of such an assignment via transfer matrix techniques.

Algorithm A.

1. While there exists a 1-local constraint c acting on some qubit v :
 - (a) If c conflicts with another 1-local clause on v , reject. Else, set $|v\rangle = |c^\perp\rangle \in \mathbb{C}^2$. Set⁸ $C_v = S_{|v\rangle}$, and remove v from Π .
2. While there exists a qubit v appearing only in clauses of size at least $k' \geq 3$:
 - (a) Set $|v\rangle = |0\rangle$ and $C_v = S_{|v\rangle}$. Remove v from Π .
3. While there exists a 2-local clause:
 - (a) If there exists a stacked 2-local clause c , i.e. $c' \neq c$ such that $Q_c \subseteq Q_{c'}$:
 - i. If $Q_c = Q_{c'}$, remove the qubits c acts on, and set their values to satisfy c and c' .
 - ii. Else, $Q_c \subset Q_{c'}$. Thus, c' is k' -local for $3 \leq k' \leq k$. Set the values of the qubits in Q_c so as to satisfy c . This collapses c' to a $(k' - 2)$ -local constraint on qubits $Q_{c'} \setminus Q_c$.
 - A. If $k' - 2 = 1$, then c' has been collapsed to a 1-local constraint on some vertex $v \in Q_{c'} \setminus Q_c$, creating a path rooted at v . Set v so as to satisfy c' , and use a CR to solve the resulting path until either the path ends, or a k'' -local constraint is hit for $3 \leq k'' \leq k'$. In the latter case (Fig. 7, Left), the k'' -local constraint is reduced to a $(k'' - 1)$ -local constraint and we return to the beginning of Step 3.
 - (b) Else, pick an arbitrary 2-local clause c acting on variables v_1 and v_2 . Then, v_1 (v_2) is the start of a path h_1 (h_2) (e.g., Fig. 7, Middle).

⁸ Setting $C_v = S_{|v\rangle}$ means we update all clauses c acting on v by projecting qubit v of c onto $|v\rangle$. Note that each v does not keep a local copy of C_v , i.e. all clauses are referenced in a global fashion.

Fig. 7. (Left) Solving the path rooted at v_1 via CR allows us to satisfy clauses (v_1, v_2) and (v_2, v_3) , and since v_3 receives an assignment in this process, the clause (v_3, v_4, v_5) is projected onto a 2-local residual clause on (v_4, v_5) . The CR then stops. (Middle) Letting c denote the clause on (v_2, v_3) , v_2 is the start of a path $(v_2, v_1, \dots)$ to the left, and v_3 is the start of a path $(v_3, v_4, \dots)$ to the right. (Right) Inducing CRs on v_1 and v_7 , we eventually assign values to v_3 and v_5 . This collapses the 3-local clause on (v_3, v_4, v_5) into a 1-local clause on v_4 with a unique satisfying assignment, which in turn induces a new CR starting at v_4 . Thus, the two CR's are “fused” into one CR through the 3-local clause

- i. If the path forms a cycle from v_1 to v_2 , use the cycle matrix to solve the cycle. Remove the corresponding qubits and clauses from Π .
- ii. Else, set v_1 and v_2 so as to satisfy c . Solve the resulting paths h_1 (h_2) until a k' -local (k'' -local) constraint l_1 (l_2) is hit for $3 \leq k' \leq k$ ($3 \leq k'' \leq k$). If both l_1 and l_2 are found:
 - A. If $l_1 = l_2$ (i.e. $k' = k''$) and $k' - 2 = 1$, then fuse the paths h_1 and h_2 into a new path beginning at the qubit in l_1 which is not in h_1 or h_2 (e.g., Fig. 7, Right). Iteratively solve the resulting path until a k' -local constraint is hit for $3 \leq k' \leq k$.

4. If any qubits are unassigned, set their values to $|0\rangle$.

An illustration of algorithm A on an example input is given after this proof (see also Fig. 6). It is clear that the algorithm runs in polynomial time. We now prove correctness.

Correctness. In Step 1, any qubit v acted on by a 1-local constraint c has only one possible satisfying assignment $|c^\perp\rangle$ (up to phase). Thus, if there are conflicting 1-local clauses acting on v , we must reject; else, we must set v to $|c^\perp\rangle$, and subsequently simplify all remaining clauses acting on v (i.e. map C_v to $S_{|v\rangle}$). Note that since v is removed at this point, we can never have a conflict on it in a future iteration of the algorithm.

If we reach Step 2, we claim that Π is satisfiable. To show this, note first that each time Step 3 is run, if there exists a 2-local clause, then at least one new clause is satisfied and subsequently removed from Π . Thus, in order to prove that Π is satisfiable, it suffices to show that the following loop invariant holds.

Loop Invariant: Before each execution of Step 3, either Π contains some 2-local clause and no 1-local clauses, or Π contains no clauses.

Two notes are in order here. First, the invariant's constraint on the absence of 1-local clauses is necessary, as otherwise we are not guaranteed that paths can be iteratively solved as in Step 3(a)(ii)(A). Second, the invariant implies that once the algorithm reaches Step 4, all clauses must have been satisfied. Thus, all unused qubits at that point can be set arbitrarily.

We now show that the loop invariant holds throughout the execution of the algorithm. First, since after Step 1, Π contained no 1-local constraints, and since in Step 2 we only reduce k' -local clauses to k'' -local ones for $k'' \geq 2$, it follows that the invariant holds the first time Step 3 is run.

We next show that if the invariant holds just before Step 3 is run, then it also holds after Step 3 is run. We divide the analysis into cases depending on which line of Step 3 is run.

- (Step 3(a)(i)) Since each qubit appears in at most 2 clauses by assumption, it holds that c and c' must be disjoint from all other clauses in Π . Informally, the

joint satisfiability of c and c' is hence independent of the satisfiability of the rest of the instance. In the language of our loop invariant, the removal of c and c' does not create any 1-local constraints. Moreover, if subsequently Π is non-empty, then it is not the case that all remaining clauses are k' -local for $3 \leq k' \leq k$. This is because otherwise the set of clauses disjoint from c and c' must contain a variable involved only in clauses of size at least $k' \geq 3$, contradicting Step 2. Thus, the loop invariant holds. (Aside: If c and c' act on vertices u and v , a satisfying assignment to c and c' can be computed by thinking of (u, v) and (v, u) as a cycle, and subsequently solving the cycle using the cycle matrix technique.)

- (Step 3(a)(ii)) Once c and the variables it acts on are removed, we have a $(k' - 2)$ -local constraint on c' . Thus, if $k' - 2 = 1$, this induces a CR on the incident path of 2-local constraints on c' (this path may have length 0). Note that this path can never loop back to c' , nor can it contain a cycle, as otherwise a variable would appear in more than 2 clauses. Thus, either the path consists solely of 2-local constraints, which are all satisfied via the CR (recall we assume that any broken CR's are continued automatically by assignment $|0\rangle$ to the next vertex in the path to induce a new CR), or eventually we hit a k'' -local constraint for $k'' \geq 3$, which we now collapse to a $(k'' - 2)$ -local constraint; call the latter ϕ . Thus, since the CR removes any possible 1-local constraint on c' and creates no further 1-local constraints, no 1-local constraints exist after Step 3.

Also, on the existence of a 2-local clause after Step 3 (assuming clauses remain): If the path above consisted solely of 2-local constraints, then the CR removed a connected component of the interaction graph, in which case the remaining connected components must contain a 2-local clause (otherwise, we again contradict Step 2). If the path encountered a k'' -local clause for $k'' = 3$, on the other hand, the CR itself created a 2-local clause. Similarly, if $k'' \geq 4$, then ϕ is at least 3-local, and we claim that at least one of the vertices of ϕ must be incident on a 2-local edge. Indeed, otherwise we again contradict Step 2. Thus, the loop invariant holds.

- (Step 3(b)(i)) Since any variable occurs in at most 2 clauses, the cycle must be disjoint from all other clauses in Π . Thus, its removal does not create a 1-local constraint. That there must exist a 2-local constraint if Π is non-empty now follows by similar arguments as in Steps 3(a)(i) and 3(a)(ii).
- (Step 3(b)(ii)) There are 2 possible scenarios: Both h_1 and h_2 are disjoint paths, or they intersect on a k' -local clause. The first case is analogous to Step 3(a)(ii)'s analysis. In the second case, suppose h_1 and h_2 intersect on clause c . Then, since the paths do not form a cycle (otherwise, we would be in Step 3(b)(i)), c must be k' -local for $k' \geq 3$. If $k' - 2 = 1$, let v denote the qubit in c not acted on by h_1 or h_2 . Then, solving the paths h_1 and h_2 collapses c into a 1-local constraint on v . But this, in turn, creates a new path rooted at v , whose analysis follows from Step 3(a)(ii) above. If $k' - 2 > 1$, the CR stops once c is collapsed to a new clause of size at least 2. That a 2-local clause still exists at this point (assuming the remaining instance is non-empty) follows by the argument of Step 3(a)(ii).

This concludes the proof. □

An example and applicability. Having shown Theorem 1, we make two remarks.

1. (Demonstration) We illustrate in Fig. 6 how algorithm A repeatedly partially reduces parts of a 3-QSAT instance to 2-QSAT instances (this particular example does not require fusing of chain reactions).

2. (Applicability) Whereas the class of hypergraphs considered in Theorem 1 may seem *a priori* rather restricted, as a “proof of concept” we observe that the example of Fig. 6 can be generalized to an entire family of non-trivial examples to which Theorem 1 applies. (We focus on 3-uniform hypergraphs for simplicity.) Specifically, encode the hypergraph $H(G)$ in the top left of Fig. 6 by a graph $G = (V_d, V_c, E)$ reminiscent of a line graph, which we call a *pseudo-line graph*, as depicted in the bottom right of Fig. 6. G is defined as follows. Each disc vertex $s_i \in V_d$ corresponds to a red hyperedge (color online; e.g. hyperedge $\{v_7, v_8, v_9\}$ is red). Each cross vertex $t_i \in V_c$ corresponds to a pair of green hyperedges (e.g. the hyperedges containing v_6) which intersect on a set of size 2. An edge in G means the two hyperstructures corresponding to each end point of the edge intersect in precisely one element (since each vertex has degree 2, this is well-defined). That each disc and cross vertex have degree 3 and 2, respectively, ensures the corresponding hypergraph, $H(G)$, is 2-regular (we assume G has no self-loops or parallel edges). It is not difficult to see that any such pseudo-line graph G gives rise to a 2-regular 3-uniform hypergraph $H(G)$, thus yielding an entire family of 3-QSAT instances to which Theorem 1 applies. Finally, we connect this example with the notion of *transfer filtrations* in Sect. 4. Namely, the hypergraph $H(G)$ in Fig. 6 is of transfer type at most 12—for all pairs of distinct hyperedges e_i, e_j with $e_i \cap e_j = 2$, simply add $e_i \cap e_j$ to the foundation.

3.1. Connection to k -QSAT instances with a SDR. Recall that the initial motivation for this study was the open question [8] of constructing satisfying assignments to k -QSAT instances which have a system of distinct representatives (SDR). We now show that the class of k -QSAT instances solved by Theorem 1 is closely related to this motivation.

Specifically, note that the only line of algorithm A which can reject is 1a, and this is due to the presence of 1-local clauses. Since any k -QSAT instance with an SDR is satisfiable (by a product state) [8], it follows that hypergraphs with hyperedges of size 1 in general cannot have SDRs. Indeed, consider the hypergraph with edge set $E = \{\{1, 2, 3\}, \{1\}, \{2\}, \{3\}\}$ —clearly, this has no SDR, and if we set the clauses to (respectively) project onto $|000\rangle$, $|1\rangle$, $|1\rangle$, and $|1\rangle$, then the corresponding 3-QSAT instance is unsatisfiable.

However, if our k -QSAT instance has no 1-local clauses in Theorem 1, then algorithm A always produces a (product-state) solution. We now show that this is no coincidence — such hypergraphs must have an SDR. Thus, Theorem 1 answers the open question of [8] for all k -QSAT instances in which (1) there are no 1-local clauses and (2) each qubit occurs in at most 2 clauses.

Theorem 2. *Let $G = (V, E)$ be a hypergraph with all hyperedges of size at least 2, and such that each vertex has degree at most 2. Then, G has an SDR.*

Proof. The claim is vacuously true if G is empty; thus, assume G is non-empty. We claim the following (poly-time) algorithm constructs an SDR.

1. While there exists a pair of distinct edges $e_i, e_j \in E$ such that $|e_i \cap e_j| \geq 2$:
 - (a) Pick an arbitrary pair of distinct vertices $v_i, v_j \in e_i \cap e_j$. Match v_i to e_i and v_j to e_j , and remove v_i, v_j, e_i, e_j from G .
2. While there exist vertices v of degree 1 in some hyperedge c_v , match v to c_v and remove v and c_v .
3. Remove any vertices of degree 0.
4. Let G' denote the remaining hypergraph. (We assume G' is connected; if not, simply repeat the following steps on each connected component.) Construct its line graph L

(defined in proof of correctness below). Find a cycle C in L , and let v be an arbitrary vertex on C .

5. (Build a “reverse SDR” on L) Run a depth-first search (DFS) rooted at v on L , with the following two modifications: (1) Each time an edge $x = (u, w)$ is used to get to a previously unseen vertex w , add 2-tuple (w, x) to set $M \subseteq V(L) \times E(L)$. (2) When the DFS begins, do not mark the root v as having been “seen”. (In other words, v should only be marked as “seen” once some edge (w, v) is used in the DFS.)
6. (Convert the “reverse” SDR on L to an SDR on G') For each $(w, x) \in M$:
 - (a) Suppose $x = (u, w)$ for some $u \in V(L)$. Then, let e_u and e_w denote the hyperedges in G' corresponding to u and w , respectively. Match e_w to the unique element of $e_u \cap e_w$.

Correctness. We analyze each step in sequence.

1. (Step 1) Since by assumption v_i and v_j have degree at most 2, they only appear in edges e_i and e_j . Thus, we can safely match them to e_i and e_j , respectively, and remove v_i, v_j, e_i, e_j .
2. (Steps 2 and 3) Since v occurs only in c_v , we can safely match it to c_v and remove v and c_v .
3. (Step 4) After Steps 1, 2, and 3, G' has two properties: (a) each pair of edges intersects in at most one vertex, and (b) each vertex has degree precisely 2. By property (a), the “usual” (i.e. we do not need to distinguish between different intersection sizes between edges) definition of line graph can be applied to G' to construct L : Each hyperedge of G' is a vertex in L , and two distinct vertices u and v in L are neighbors if and only if their corresponding hyperedges in G' intersect.

Moreover, by properties (a) and (b) of G' , each vertex of L has degree at least 2. It follows that L contains a cycle C (e.g. since summing degrees on each vertex yields there are at least as many edges in L as vertices, and L is connected by the assumption that G' is connected).

4. (Step 5) Since a DFS sees each vertex (at least) once, and since the first time a vertex w is seen it must be that the edge $x = (u, w)$ just used to get to w must have been used for the first time, it follows that no other 2-tuple in M contains x or w . Moreover, as stated in Step 5, note that the root v of the DFS is not marked until some edge (w, v) is used; the latter is guaranteed to happen since (1) we chose v to lie on a cycle C , and (2) if the two edges on the cycle C which v is incident on are $\{w, v\}$ and $\{v, w'\}$, then by definition of a DFS at least one of them will be traversed in direction (w, v) or (w', v) (as opposed to both being traversed in directions (v, w) and (v, w') ; note this property holds since we use a DFS instead of a breadth-first search). It follows that each vertex w of L appears in precisely one 2-tuple of $(w, x) \in M$. We conclude that M is a “reverse SDR” for L , by which we mean each *vertex* of L is matched with a unique *edge* of L (whereas our desired SDR for hypergraph G is supposed to match each *hyperedge* to a unique *vertex*⁹).
5. (Step 6) This step converts each two-tuple in M to a matched hyperedge e_w and vertex in $v \in e_u \cap e_w$. Note that $|e_u \cap e_w| = 1$ by property (a) of G' . Since M is a “reverse” SDR, it follows that the matching produced by Step 6 completes the desired SDR. $\square$

⁹ For further clarity, the distinction is that in a reverse SDR, there may be more edges in the graph than vertices, whereas in an SDR, there may be more vertices than edges.

```

1: input: A  $k$ -uniform hypergraph  $G = (V, E)$  (along with a constraint  $H_e$  for each hyperedge  $e$ ),
   transfer filtration  $F$  of  $G$  of type  $b$  and radius  $r$ .
2: output: An  $n$ -qubit tensor product assignment satisfying all constraints, if one exists. Else,
   the algorithm rejects.
3:
4: procedure BLOWUP( $G, F$ )
5:   Apply Decoupling Lemma (Lemma 18) to  $G$  to compute  $\tilde{G}$  and  $p : V(\tilde{G}) \mapsto V(G)$ .
6:   Apply Transfer Lemma (Lemma 28) to  $\tilde{G}$  to compute  $g_1, \dots, g_n : W^b \mapsto W$ .
7:   Apply Qualifier Lemma (Lemma 30) to  $g_1, \dots, g_n$  to compute  $h_1, \dots, h_{m-n+b} : W^b \mapsto \mathbb{K}$ .
8:   Compute foundation assignment  $v_1, \dots, v_b \subseteq W$  such that for all  $s \in [m - n + b]$ ,
        $h_s(v_1, \dots, v_b) = 0$  and  $g_{i_s}(v_1, \dots, v_b) \neq 0$ . If no such assignment found, reject.
9:   For each qubit  $i$  of  $G$ , return assignment  $g_i(v_1, \dots, v_b)$ .
10: end procedure

```

Fig. 8. Parameterized algorithm for Quantum k -SAT

4. Quantum SAT and Parameterized Algorithms

We next develop a parameterized algorithm for computing an explicit (product state) solution to a non-trivial class of k -QSAT instances (Theorem 38). Although the inspiration stems from algebraic geometry (AG), we generally avoid AG terminology to increase accessibility. For those versed in the topic, however, we include a brief overview of the ideas of this section in AG terms at the end of Sect. 4.3. The final algorithm and its runtime, along with the study of asymptotic speedups, are given in Sect. 4.4. For readers who wish to see the outline of the algorithm before reading the details, see Fig. 8.

4.1. The transfer type of a hypergraph.

Transfer filtrations. The core graph theoretic idea on which our framework will be built is the new notion of *transfer type* of a hypergraph, which we now introduce.

Definition 9. A hypergraph $G = (V, E)$ is of *transfer type b* if there exists a chain of subhypergraphs (denoted a *transfer filtration of type b*) $G_0 \subseteq G_1 \subseteq \dots \subseteq G_m = G$ and an ordering of the edges $E(G) = \{E_1, \dots, E_m\}$ such that

1. $E(G_i) = \{E_1, \dots, E_i\}$ for each $i \in \{0, \dots, m\}$,
2. $|V(G_i)| \leq |V(G_{i-1})| + 1$ for each $i \in \{1, \dots, m\}$,
3. if $|V(G_i)| = |V(G_{i-1})| + 1$, then $V(G_i) \setminus V(G_{i-1}) \subseteq E_i$,
4. $|V(G_0)| = b$, where we call $V(G_0)$ the *foundation*,
5. and each edge of G has at least one vertex not in $V(G_0)$.

In other words, a transfer filtration of type b builds up G iteratively by choosing b vertices as a “foundation”, and in each iteration adding precisely one new edge E_i and *at most* one new vertex. If a new vertex is added in iteration i , condition (3) says it must be in edge E_i added in iteration i . Note that under certain conditions (i.e. when the foundation size is bounded by $b \leq |V(G)| - |E(G)| + k - 1$), we later show (Theorem 50) that this iterative construction can be exploited to prove the hypergraph G must have an SDR. Also, recall that as stated in Sect. 1, we conjecture that given G as input, finding a foundation of minimum size is NP-hard. (In the context of our parameterized algorithms, this is reminiscent of common parameters such as treewidth being NP-hard to compute [16].) Finally, condition 5 in Definition 9 is a “minimality” condition. For

Fig. 9. The running example from Example 10. In clockwise order, beginning with the top left, are the subgraphs G_i in the specified transfer filtration of Example 10: G_0 (foundation), G_1 , G_2 , G_3 , $G_4 = G$. Note that G_2 is a chain, G_3 a semicycle, and G_4 a cycle

example, the entire vertex set is also a valid foundation, but contains more vertices than required to satisfy conditions 1 through 4. To prevent such cases, for convenience in our discussions we have added Condition 5 to Definition 9, which holds without loss of generality. In particular, given a filtration satisfying 1-4 but not 5, it is easy to compute a modified filtration with a smaller foundation satisfying 1-5 (if edge e fails 5, simply remove a single arbitrary vertex $v \in e$ from G_0 , and add an additional intermediate graph G_i to the filtration to first “pick up” vertex v).

Example 10 (Running example (Fig. 9)). We now introduce a hypergraph G which will serve as a running example to concretely illustrate the ideas of this section. Let $V(G) = \{1, 2, 3, 4\}$ with edges $E_1 = \{1, 2, 3\}$, $E_2 = \{1, 2, 4\}$, $E_3 = \{1, 3, 4\}$ and $E_4 = \{2, 3, 4\}$. By Definition 5, G is a 3-uniform cycle. Consider hypergraphs G_0, G_1, G_2, G_3 such that $V(G_0) = \{1, 2\}$, $V(G_1) = \{1, 2, 3\}$, $V(G_2) = V(G_3) = V(G_4) = V(G)$, $E(G_0) = \emptyset$ and $E(G_j) = \{E_1, \dots, E_j\}$ for $j = 1, 2, 3$. Then $G_0 \subseteq G_1 \subseteq G_2 \subseteq G_3 \subseteq G_4 = G$ is a transfer filtration of type 2. In particular G_2 is a chain and G_3 is a semicycle in the sense of Definitions 4 and 6, respectively.

Remark 11. Let G be a k -uniform hypergraph of transfer type b . Since G_1 has exactly one edge, then $b \geq k - 1$.

Example 12 The *tight line graph* [36] of a k -uniform hypergraph G is the (undirected) graph $L(G)$ with vertices labelled by $E(G)$ and edges $\{(E_i, E_j) \mid E_i, E_j \in E(G) \text{ and } |E_i \cap E_j| = k - 1\}$. (This generalizes the notion of a line graph; a yet more general definition, called an *edge intersection graph*, is mentioned in Sect. 5, which adds an edge (E_i, E_j) so long as $E_i \cap E_j \neq \emptyset$.) A k -uniform hypergraph G is *line graph connected* if its tight line graph is connected. Arguing as in the proof of Lemma 1 of [36], it is easy to show that every k -uniform line graph connected hypergraph with no isolated vertices has transfer type $k - 1$. (An *isolated* vertex is not contained in any hyperedge.) To see this, pick an edge E_{i_1} , set G_1 to be the hypergraph induced by E_{i_1} , and choose G_0 to be any subset of G_1 of cardinality $k - 1$. Since the vertex of $L(G)$ corresponding to E_{i_1} is not isolated, there exists another edge E_{i_2} such that $|E_{i_1} \cap E_{i_2}| = k - 1$. Let G_2 be the hypergraph with vertices $E_{i_1} \cup E_{i_2}$ and edges E_{i_1}, E_{i_2} . Since the induced subgraph with vertices E_{i_1} and E_{i_2} is not a connected component of $L(G)$ this procedure can be iterated until the required transfer filtration is constructed.

Example 13 (Running example). Let G be the 3-uniform cycle of Example 10. Then $L(G)$ is a cycle of length 4, in the usual graph-theoretical sense. In particular G is line graph connected.

Example 14 Let G be a k -uniform hypergraph with no isolated vertex. Suppose that $L(G)$ has c connected components and apply the construction of Example 12 to each connected component. Taking the union of the corresponding G_0 sets (one for each connected component) and adding the edges one at the time, we obtain a transfer filtration of type at most $c(k-1)$ on G . More generally, isolated vertices can be taken into account by adding them to the 0-th term of the filtration. Hence every k -uniform hypergraph has transfer type at most $c(k-1) + i$ where c is the number of connected components of $L(G)$ and i is the number of isolated vertices of G .

Example 15 Let $a_1, \dots, a_{k-1} \geq 2$ be integers. Consider the hypergraph G with vertices $(\mathbb{Z}/a_1\mathbb{Z}) \times \dots \times (\mathbb{Z}/a_{k-1}\mathbb{Z})$ and edges of the form

$$\{(i_1, \dots, i_{k-1}), (i_1 + 1, \dots, i_{k-1}), (i_1, i_2 + 1, \dots, i_{k-1}), \dots, (i_1, \dots, i_{k-1} + 1)\}$$

for each $(i_1, \dots, i_{k-1}) \in V(G)$. Let G_0 be the subhypergraph of G with vertices $(\mathbb{Z}/a_1\mathbb{Z}) \times \dots \times (\mathbb{Z}/a_{k-2}\mathbb{Z}) \times \{0\}$ and no edges. Let G_1 be the subhypergraph with vertices $V(G_0) \cup \{(0, \dots, 0, 1)\}$ and a single edge $\{(0, \dots, 0), (1, 0, \dots, 0), \dots, (0, \dots, 0, 1)\}$. Let G_2 be the subhypergraph obtained from G_1 by adding the vertex $(0, \dots, 0, 1, 1)$ and the edge

$$\{(0, \dots, 0, 1, 0), (1, 0, \dots, 0, 1, 0), \dots, (0, \dots, 0, 2, 0), (0, \dots, 0, 1, 1)\}.$$

Proceeding in this way, say by reverse lexicographic order, one constructs all the vertices with last coordinate equal to 1 and all the edges containing $k-1$ vertices with last coordinate equal to 0 and exactly one vertex with last coordinate equal to 1. The corresponding chain of subhypergraphs can be labeled as $G_0 \subseteq G_1 \subseteq \dots \subseteq G_{a_1 \dots a_{k-2}}$. Iterating this construction one can extend this chain to $G_0 \subseteq G_1, \dots, \subseteq G_{2a_1 \dots a_{k-2}}$ in such a way that all the vertices whose last coordinate is 0, 1, or 2 are accounted. Further iterations lead to a transfer filtration of type $|G_0| = a_1 \dots a_{k-2}$. Clearly there is nothing special about the particular choice of G_0 for instance we could have picked $V(G_0) = \{0\} \times (\mathbb{Z}/a_2\mathbb{Z}) \times \dots \times (\mathbb{Z}/a_{k-1}\mathbb{Z})$ in which case an obvious adaptation of the construction above leads to a transfer filtration of type $a_2 \dots a_{k-2}$. In particular, this shows that in general a hypergraph G can be of transfer type b for different choices of b .

Remark 16. Let G be a hypergraph admitting a transfer filtration $G_0 \subseteq G_1 \subseteq \dots \subseteq G_m = G$ of type b . Assume the edges of G are ordered in such a way that $E(G_i) = \{E_1, \dots, E_i\}$ for each $i \in \{1, \dots, m\}$. Since by construction each edge contains at least one vertex not in $V(G_0)$, there exists a function $r : \{1, \dots, m\} \rightarrow \{0, \dots, m-1\}$ such that $r(i) < i$ and $|E_i \setminus V(G_{r(i)})| = 1$ for all $i \in \{1, \dots, m\}$.

Example 17 (Running example). Let G be the 3-uniform cycle of Example 10. Then $r : \{1, 2, 3, 4\} \rightarrow \{0, 1, 2, 3\}$ can be chosen to be such that $r(1) = r(2) = 0$, $r(3) = 1$ and $r(4) = 1$.

The Decoupling Lemma. As the first step in our construction, we show how to map any k -uniform hypergraph G of transfer type b to a new k -uniform hypergraph G' of transfer type b whose transfer filtration *must* add a vertex in each step (this follows directly from the relationship between $|V(G)|$ and $|E(G)|$ below). This has two effects worth noting: First, G' is guaranteed to have an SDR. Second, it *decouples* certain intersections in the hypergraph, as illustrated in Fig. 10. For clarity, in the lemma below, for a function p acting on vertices, we implicitly extend its action to edges in the natural way, i.e. if $e = (v_1, v_2, v_3)$ then $p(e) = (p(v_1), p(v_2), p(v_3))$.

Fig. 10. For the hypergraph on the left, consider the transfer filtration in which G_0 contains vertices v_1, v_2 , and we iteratively add edges $\{v_1, v_2, v_3\}$, $\{v_1, v_2, v_4\}$, and $\{v_1, v_3, v_4\}$ to the filtration. Then, the Decoupling Lemma (Lemma 18) maps the hypergraph on the left to the hypergraph on the right, in the process decoupling the intersection on vertex v_4 . The surjective function p “undoes” the decoupling by mapping v_1, v_2, v_3 to themselves, and v_4, v_5 to v_4

Lemma 18. (Decoupling lemma). *Given a k -uniform hypergraph G of transfer type b , there exists a k -uniform hypergraph $\tilde{G}$ of transfer type b with $|E(G)| + b$ vertices and a surjective function $p : V(\tilde{G}) \rightarrow V(G)$ such $p(\tilde{E}) \in E(G)$ for every $\tilde{E} \in E(\tilde{G})$.*

Proof. Let $G_0 \subseteq G_1 \subseteq \dots \subseteq G_m = G$ be a transfer filtration such that $|V(G_0)| = b$ and let $r : \{1, \dots, m\} \rightarrow \{0, \dots, m-1\}$ as in Remark 16. Suppose the vertices of G are labeled as $\{1, \dots, n\}$ in such a way that $V(G_0) = \{1, \dots, b\}$. Furthermore, we may assume that the edges $\{E_1, \dots, E_m\}$ of G are labeled in such a way that $E(G_i) = \{E_1, \dots, E_i\}$ for each $i \in \{1, \dots, m\}$. Let $V(\tilde{G}) = \{1, \dots, m+b\}$ and define $p : \{1, \dots, m+b\} \rightarrow \{1, \dots, n\}$ so that $p(i) = i$ for all $i \in \{1, \dots, b\}$ and $\{p(i)\} = E_{i-b} \setminus V(G_{r(i-b)})$ for all $i \in \{b+1, \dots, m+b\}$ (where recall r is from Remark 16). By Remark 16, p is surjective. For each $j \in \{1, \dots, m+b\}$, let $\underline{j} = \min(p^{-1}(p(j)))$. If we define $E(\tilde{G}) = \{\tilde{E}_1, \dots, \tilde{E}_m\}$ by setting $\tilde{E}_i = \{i+b\} \cup \{\underline{j} \mid j \in p^{-1}(E_i \setminus p(i+b))\}$ for each $i \in \{1, \dots, m\}$, then $\tilde{G}$ is clearly k -uniform. Let $\tilde{G}_0$ be the hypergraph with vertices $\{1, \dots, b\}$ and no edges. For each $i \in \{1, \dots, m\}$, let $\tilde{G}_i$ be the subgraph of $\tilde{G}$ with vertices $\{1, \dots, i+b\}$ and edges $\{\tilde{E}_1, \dots, \tilde{E}_i\}$. Then $\tilde{G}_0 \subseteq \tilde{G}_1 \subseteq \dots \subseteq \tilde{G}_m = \tilde{G}$ is a transfer filtration and $\tilde{G}$ is a hypergraph of transfer type b . $\square$

Example 19. (Running example (Fig. 11)). Let G be the 3-uniform cycle of Example 10. The proof of Lemma 18 (appendix) produces a 3-uniform hypergraph $\tilde{G}$ with vertices $\{1, 2, 3, 4, 5, 6\}$ and edges $\tilde{E}_1 = \{1, 2, 3\}$, $\tilde{E}_2 = \{1, 2, 4\}$, $\tilde{E}_3 = \{1, 3, 5\}$, $\tilde{E}_4 = \{2, 3, 6\}$, and surjective function $p : \{1, 2, 3, 4, 5, 6\} \rightarrow \{1, 2, 3, 4\}$ defined by $p(1) = 1$, $p(2) = 2$, $p(3) = 3$, $p(4) = p(5) = p(6) = 4$. This choice is not unique: setting $\tilde{E}_4 = \{2, 4, 6\}$ and $p(6) = 3$ also satisfies Lemma 18.

Fig. 11. The “decoupled” hypergraph from the running example of Example 19. Note the order of listing v_4 and v_5 has been swapped for visual simplicity

Remark 20. If $k = 2$, then G (as defined in Lemma 18) is an ordinary graph and the construction of p can be understood topologically in terms of the geometric realization $|G|$ of the standard simplicial set defined by G . Assume for simplicity that G is connected. Then the universal cover $|\widetilde{G}|$ is the geometric realization of a (possibly infinite) tree. Moreover, the geometric realization of p is the restriction of the canonical covering map $|\widetilde{G}| \rightarrow |G|$ to the closure of a fundamental domain for the canonical action of $\pi_1(|G|)$ on $|\widetilde{G}|$.

Radius of transfer filtration. One of the “parameters” in our parameterized approach will be the *radius* of a transfer filtration, defined next. The concept is reminiscent of radii of graphs, and roughly measures “how far” an edge is from the foundation of b vertices with respect to the filtration.

Definition 21 (*Radius of transfer filtration*). Let G be a hypergraph admitting a transfer filtration $G_0 \subseteq \dots \subseteq G_m = G$ of type b . Consider the function (whose existence is guaranteed by Remark 16) $r : \{0, \dots, m\} \rightarrow \{0, \dots, m-1\}$ such that $r(0) = 0$ and $r(i)$ is the smallest integer such that $|E_i \setminus V(G_{r(i)})| = 1$ for all $i \in \{1, \dots, m\}$. The *radius of the transfer filtration* $G_0 \subseteq \dots \subseteq G_m = G$ of type b is the smallest integer β such that $r^\beta(i) = 0 \forall i \in \{1, \dots, m\}$ (r^β denotes the composition of r with itself β times). The *type b radius of G* is the minimum value $\rho(G, b)$ of β over the set of all possible transfer filtrations of type b on G .

Example 22 (Running example). For G the 4-cycle from Example 10, since function r described in Example 17 is non-constant and $r(r(i)) = 0$ for all $i \in \{1, 2, 3, 4\}$, then the transfer filtration of Example 10 has radius $\beta = 2$.

4.2. The main construction. To describe the main construction, we shall use rather general terminology to describe the most general setting in which the framework applies. For those versed in quantum information, we will attempt to note what special cases for the general definitions we give are most relevant to said community; these notes will be marked with terminology (QI) (color online).

Definitions. To begin, let W be a two dimensional vector space over a field $\mathbb{K}$. (QI) We may set $\mathbb{K} = \mathbb{C}$ and identify W with $\mathbb{C}^2$ if desired. We now define k -local functions; (QI) these capture the notion of a k -local Hamiltonian constraint evaluated against a tensor product of k input qubit states $|v_{i_1}\rangle \otimes \dots \otimes |v_{i_k}\rangle$.

Definition 23 (*k -local, interaction graph, product satisfiability set*). A function $H_i : W^n \rightarrow \mathbb{K}$ is k -local if there exists a subset $E_i = \{i_1, \dots, i_k\} \subseteq \{1, \dots, n\}$ and a non-zero functional $H_i^* : W^{\otimes k} \rightarrow \mathbb{K}$ such that

$$H_i(v_1, \dots, v_n) = H_i^*(v_{i_1} \otimes \dots \otimes v_{i_k})$$

for all $v_1, \dots, v_n \in W$, i.e. H_i acts non-trivially only on a subset of k indices. A collection $H = (H_1, \dots, H_m)$ of k -local functions $H_1, \dots, H_m : W^n \rightarrow \mathbb{K}$ is k -local. The corresponding subsets $\{E_1, \dots, E_m\}$ (i.e. on which H_1 through H_m act non-trivially, respectively) are the edges of a hypergraph G_H with vertices $\{1, \dots, n\}$ known as the *interaction graph of H* . The *product satisfiability set* of the k -local collection H is the set $\mathcal{S}_H$ of all $(v_1, \dots, v_n) \in W^n$ such that $v_i \neq 0$ for all $i \in \{1, \dots, n\}$ and $H_j(v_1, \dots, v_n) = 0$ for all $j \in \{1, \dots, m\}$.

Remark 24 ($\sharp$ isomorphism). Consider an isomorphism $\sharp$ between W and its dual $W^\vee$ that to each $v \in W$ assigns a functional $v^\sharp \in W^\vee$ such that $v^\sharp(v) = 0$. For instance, if a basis $\{w_1, w_2\}$ for W is chosen then we may define $\sharp$ by setting $((a_1 w_1 + a_2 w_2)^\sharp)(b_1 w_1 + b_2 w_2) = a_1 b_2 - a_2 b_1$ for all $a_1, a_2, b_1, b_2 \in \mathbb{K}$. Moreover, given any $v_1, v_2 \in W$, then $v_1^\sharp(v_2) = 0$ if and only if there exists $\lambda \in \mathbb{K}$ such that $\lambda v_2 = v_1$. (QI) The $\sharp$ isomorphism essentially maps a vector $|\psi\rangle \in \mathbb{C}^2$ to its unique (up to scaling) orthogonal state $|\psi^\perp\rangle \in \mathbb{C}^2$.

Definition 25 (Fibonacci numbers of order N). Let N be a non-negative integer. The Fibonacci numbers of order N are the entries of the sequence $(F_r^{(N)})$ such that $F_r^{(N)} = F_{r-1}^{(N)} + \dots + F_{r-N}^{(N)}$ for all $r \geq N$, $F_{N-1}^{(N)} = 1$ and $F_r^{(N)} = 0$ for all $r \leq N - 2$.

Remark 26. It is shown in [37] that there exists a monotonically increasing sequence (ψ_N) with values in the real interval $[1, 2)$ such that, for each $N \geq 1$, $F_r^{(N)} \sim \psi_N^r$ as $r \rightarrow +\infty$.

Definition 27 (Degree). A function f on W^l with values in a $\mathbb{K}$ -vector space has *degree* $(d_1, \dots, d_l)$ if $f(\lambda_1 v_1, \dots, \lambda_l v_l) = \lambda_1^{d_1} \dots \lambda_l^{d_l} f(v_1, \dots, v_l)$ for every $\lambda_1, \dots, \lambda_l \in \mathbb{K}$ and every $v_1, \dots, v_l \in W$.

The Transfer Lemma. Applying the Decoupling Lemma to an input k -uniform hypergraph G with transfer type b , we obtain a k -uniform hypergraph $\tilde{G}$ of type b with $m = n - b$, for m and n the number of edges and vertices, respectively. The next lemma shows that $\tilde{G}$ is “nice”, in that any global (product) solution to the k -QSAT system can be derived from a set of assignments to the b foundation vertices, and conversely, any (product) assignment to the latter can be extended to a global (product) solution. For the reader familiar with transfer matrices and CRs (see Sect. 3), with a little thought it can be seen that the latter of these claims is similar to picking an arbitrary assignment to the b foundation vertices, and then iteratively applying transfer matrices to satisfy all clauses (i.e. in step i of the filtration in which edge E_i and vertex v_i are added, the 4×2 transfer matrix of edge E_i is applied to the two pre-existing vertices of E_i in the filtration, yielding an assignment to v_i which satisfies clause E_i).

Lemma 28 (Transfer Lemma). Let $H = (H_1, \dots, H_{n-b})$ be a k -local collection of functions $H_i : W^n \rightarrow \mathbb{K}$ whose interaction graph is a k -uniform hypergraph of transfer type b . There exist non-zero (non-constant) functions, which we call “transfer functions”, $g_1, \dots, g_n : W^b \rightarrow W$ with the following properties.

1. (Global to local assignments) If $(v_1, \dots, v_n) \in \mathcal{S}_H$ (recall the v_i are non-zero by definition of $\mathcal{S}_H$) there exist non-zero $\lambda_1, \dots, \lambda_n \in \mathbb{K}$ such that, for every $i \in \{1, \dots, n\}$,

$$\lambda_i v_i = g_i(v_1, \dots, v_b). \quad (1)$$

2. (Local to global assignments) For any non-zero $v_1, \dots, v_b \in W$ there exist $v_{b+1}, \dots, v_n \in W$ such that $(v_1, \dots, v_n) \in \mathcal{S}_H$ and $v_i = g_i(v_1, \dots, v_b)$ for every i such that $g_i(v_1, \dots, v_b) \neq 0$.
3. (Degree bounds) g_i has degree $(d_{i1}, \dots, d_{ib})$ such that $d_{ij} \leq F_i^{(b)}$ for all $j \in \{1, \dots, b\}$.

(QI) Intuitively, a transfer function g_i takes in a tensor product assignment to the foundation qubits, and gives a closed formula for an assignment to qubit i .

Proof. For $i = 1, \dots, b$, define $g_i(v_1, \dots, v_b) = v_i$. Up to relabeling the n vertices, since the number of edges is $n - b$, we may assume that G_H admits a filtration $G_0 \subseteq G_1 \subseteq \dots \subseteq G_{n-b} = G_H$ such that each G_i is of transfer of type b and $V(G_i) \setminus V(G_{i-1}) = \{i + b\}$ for all $i \in \{1, \dots, n - b\}$. Therefore, we may work by induction on n (for each fixed b) and assume that $g_{b+1}, \dots, g_{n-1}$ have been constructed. Suppose that $E_{n-b} = \{n, i_1, \dots, i_{k-1}\}$ (where we shall assume vertex n is added in step $n - b$ of the filtration) and consider the function $g_n^\sharp : W^b \rightarrow W^\vee$ such that

$$(g_n^\sharp(v_1, \dots, v_b))(v) = H_{n-b}^*(g_{i_1}(v_1, \dots, v_b) \otimes \dots \otimes g_{i_{k-1}}(v_1, \dots, v_b) \otimes v) \quad (2)$$

for all $v_1, \dots, v_b, v \in W$. (QI) Intuitively, $g_n^\sharp$ evaluates the local Hamiltonian term (i.e. estimates its energy penalty) corresponding to H_{n-b}^* on the tensor product assignment prescribed by transfer functions g_i on the first $k - 1$ indices and argument v on the k -th index. For (1) of the claim, by induction, we have that $(v_1, \dots, v_n) \in \mathcal{S}_H$ implies that $g_1, \dots, g_{n-1}$ as (1) exist and $g_n^\sharp(v_1, \dots, v_b)(v_n) = 0$. Given an isomorphism $\sharp$ between W and $W^\vee$ as in Remark 24, we define a function $g_n : W^b \rightarrow W$ by setting $(g_n(v_1, \dots, v_b))^\sharp = g_n^\sharp(v_1, \dots, v_b)$ for all $v_1, \dots, v_b \in W$. Then $(v_1, \dots, v_n)$ is in $\mathcal{S}_H$ implies (1) holds for every $i \in \{1, \dots, n\}$. This proves (1).

To prove (2), suppose $v_{b+1}, \dots, v_{n-1}$ have been constructed. Let v be such that $H_{n-b}^*(v_{i_1} \otimes \dots \otimes v_{i_{k-1}} \otimes v) = 0$ (such a v exists, for example, since recall each QSAT constraint is rank 1 in our definition). Then $(v_1, \dots, v_{n-1}, \lambda v) \in \mathcal{S}_H$ for any $\lambda \in \mathbb{K}$. Moreover, if $g_n(v_1, \dots, v_b) \neq 0$ then by Remark 24 it has to equal μv for some non-zero $\mu \in \mathbb{K}$. Setting $v_n = \mu v$, proves (2).

To prove (3), we observe that the degree of g_n equals the degree of $g_n^\sharp$. Using induction and (2) it is easy to see that the latter satisfies the claimed bounds. $\square$

Example 29. (Running example). Let $H = (H_1, H_2, H_3, H_4)$ be a 3-local collection of functions $H_i : W^6 \rightarrow \mathbb{K}$ whose interaction graph is the 3-uniform chain $\tilde{G}$ described in Example 19 (obtained by plugging the 4-cycle $\tilde{G}$ of Example 10 into the Decoupling Lemma). For clarity, H_i is defined on hyperedge $\tilde{E}_i$, where the order of vertices in each edge is fixed by the transfer filtration chosen; in particular, we use the natural ordering $\tilde{E}_1 = (1, 2, 3)$, $\tilde{E}_2 = (1, 2, 4)$, $\tilde{E}_3 = (1, 3, 5)$, $\tilde{E}_4 = (2, 4, 6)$, so that the foundation is $\{1, 2\}$. Then the proof of Lemma 28 constructs transfer functions $g_1, \dots, g_6 : W^2 \rightarrow W$ which give assignments to qubits 1 through 6, respectively, as follows. Fixing a basis $\{w_1, w_2\}$ of W and unraveling (2) we obtain

$$\begin{aligned} g_1(v_1, v_2) &= v_1; \\ g_2(v_1, v_2) &= v_2; \\ g_3(v_1, v_2) &= H_1^*(v_1 \otimes v_2 \otimes w_2)w_1 - H_1^*(v_1 \otimes v_2 \otimes w_1)w_2; \\ g_4(v_1, v_2) &= H_2^*(v_1 \otimes v_2 \otimes w_2)w_1 - H_2^*(v_1 \otimes v_2 \otimes w_1)w_2; \\ g_5(v_1, v_2) &= H_3^*(v_1 \otimes g_3(v_1, v_2) \otimes w_2)w_1 - H_3^*(v_1 \otimes g_3(v_1, v_2) \otimes w_1)w_2; \\ g_6(v_1, v_2) &= H_4^*(v_2 \otimes g_4(v_1, v_2) \otimes w_2)w_1 - H_4^*(v_2 \otimes g_4(v_1, v_2) \otimes w_1)w_2. \end{aligned}$$

In particular the matrix of degrees d_{ij} is

$$\begin{pmatrix} 1 & 0 & 1 & 1 & 2 & 1 \\ 0 & 1 & 1 & 1 & 1 & 2 \end{pmatrix}^T$$

so that, in accordance to the bound given in Lemma 28, the i -th entry of each column is less or equal than the i -th (ordinary) Fibonacci number $F_i^{(2)}$.

The Qualifier Lemma. Thus far, we have seen how combining the Decoupling and Transfer Lemmas “blows up” an input k -QSAT system Π to a larger “decoupled” system Π^+ which is easier to solve due to its decoupled property. Now we wish to relate the solutions of Π^+ back to Π . This is accomplished by the next lemma, which introduces a set of “qualifier” constraints $\{h_s\}$ with the key property (follows from Lemma 30 and Remark 31): Any solution to $\{h_s\}$ can be extended to one for Π^+ , and then *mapped back* to a solution for Π . Importantly, the qualifier constraints *act only* on the b foundation vertices, as opposed to all n vertices!

Lemma 30 (Qualifier Lemma). *Let $H = (H_1, \dots, H_m)$ be a k -local collection of functions $H_i : W^n \rightarrow \mathbb{K}$ whose interaction graph is a k -uniform hypergraph of transfer type b such that $m > n - b$. Then there exist non-zero (non-constant) functions, called qualifiers, $h_1, \dots, h_{m-n+b} : W^b \rightarrow \mathbb{K}$ and $\pi : W^n \rightarrow W^b$ such that*

1. $h_s(\pi(\mathcal{S}_H)) = 0$ for all $s \in \{1, \dots, m - n + b\}$;
2. h_s has degree $(d_{s1}, \dots, d_{sb})$ with $d_{sr} \leq 2F_{\rho(G,b)+b+1}^{(b)}$ for all $s \in \{1, \dots, m + b\}$ and all $r \in \{1, \dots, b\}$.

(QI) Condition (1) essentially says that any satisfying product state to the k -QSAT instance H will, after application of an appropriate map π from n qubits down to the b foundation qubits, yield a common root to all qualifier functions h_s .

Proof. We begin by applying the Decoupling Lemma (Lemma 18) to G_H . This yields the larger hypergraph $\widetilde{G}_H$ and surjection $p : V(\widetilde{G}_H) \rightarrow V(G_H)$, constructed via a transfer filtration $G_0 \subseteq \dots \subseteq G_m = G$ of type b . Note that $\widetilde{G}_H$ is the interaction graph of a k -local collection $\widetilde{H} = (\widetilde{H}_1, \dots, \widetilde{H}_m)$ of functions $\widetilde{H}_i : W^{m+b} \rightarrow \mathbb{K}$ such that $\widetilde{H}_i^* = H_i^*$ for each $i \in \{1, \dots, m\}$. Let $\Delta : W^n \rightarrow W^{m+b}$ be such that $\Delta(v_1, \dots, v_n) = (\widetilde{v}_1, \dots, \widetilde{v}_{m+b})$, where $\widetilde{v}_i = v_{p(i)}$ for all $i \in \{1, \dots, m+b\}$. In other words, Δ “blows up” any assignment on the original n vertices to one on $m+b$ vertices (where by assumption $m+b > n$), such that any vertex j of $\widetilde{G}_H$ obtained by “decoupling” a vertex i of G_H (formally, $p(j) = i$) is assigned the same vector as i . It follows that $(v_1, \dots, v_n) \in \mathcal{S}_H$ if and only if $\Delta(v_1, \dots, v_n) \in \mathcal{S}_{\widetilde{H}}$.

Having applied the Decoupling Lemma and obtained a hypergraph $\widetilde{G}_H$ satisfying $m = n - b$, we now apply the Transfer Lemma (Lemma 28) to $\widetilde{G}_H$. This yields that there exist $g_1, \dots, g_{m+b} : W^b \rightarrow W$ with the property that $\Delta(v_1, \dots, v_n) \in \mathcal{S}_{\widetilde{H}}$ implies that $g_i(v_1, \dots, v_b)$ is a multiple of $v_{p(i)}$ for all $i \in \{1, \dots, m+b\}$. Borrowing notation from the proof of Lemma 18, let $\{i_1, \dots, i_{m-n+b}\}$ be the subset of all $i \in \{1, \dots, m+b\}$ such that $i < i$ (intuitively, these indices correspond to vertices of $\widetilde{G}_H$ obtained by decoupling a vertex of G_H). For each $s \in \{1, \dots, m - n + b\}$, define qualifier $h_s : W^b \rightarrow \mathbb{K}$ be such that

$$h_s(v_1, \dots, v_b) = (g_{i_s}^\#(v_1, \dots, v_b))(g_{i_s}(v_1, \dots, v_b)) \quad (3)$$

for all $v_1, \dots, v_b \in W$. If $(v_1, \dots, v_n) \in \mathcal{S}_H$, then for every $s \in \{1, \dots, m - n + b\}$ there exists $\lambda_{i_s}, \lambda_{i_s} \in \mathbb{K}$ such that $\lambda_{i_s} v_{p(i_s)} = g_{i_s}(v_1, \dots, v_b)$ and $\lambda_{i_s} v_{p(i_s)} = g_{i_s}(v_1, \dots, v_b)$. Therefore,

$$h_s(v_1, \dots, v_b) = \lambda_{i_s} \lambda_{i_s} v_{p(i_s)}^\#(v_{p(i_s)}) = 0$$

for every $s \in \{1, \dots, m - n + b\}$. Therefore, (1) holds if we define π to be the composition of Δ with the projection onto the first b entries. It follows from Lemma 28 and (3) that h_s has degree $(d_{s1}, \dots, d_{sb})$ with $d_{sr} \leq 2F_{i_s}^{(b)}$ for every $s \in \{1, \dots, m - n + b\}$ and $r \in \{1, \dots, b\}$. Finally (2) follows from Lemma 28, by choosing the transfer filtration $G_0 \subseteq \dots \subseteq G_m = G$ of type b to have radius $\rho(G, b)$. $\square$

Remark 31. (How to use Qualifier Lemma to extend local solutions to global solutions). In the notation of the proof of Lemma 30, suppose $v_1, \dots, v_b$ is an assignment to the foundation which satisfies the qualifiers, i.e. $h_s(v_1, \dots, v_b) = 0$ for all $s \in \{1, \dots, m - n + b\}$, and that $g_{i_s}(v_1, \dots, v_b) \neq 0$ for all $s \in \{1, \dots, m - n + b\}$. Then by Remark 24 we have that for each $s \in \{1, \dots, m - n + b\}$ there exist $\mu_s \in \mathbb{K}$ such that

$$\mu_s g_{i_s}(v_1, \dots, v_b) = g_{i_s}(v_1, \dots, v_b).$$

In other words, in a solution to the decoupled instance on $\widetilde{G}_H$, it must be the case that all decoupled copies of a vertex v receive the *same* assignment (up to scalars). Indeed, $g_{i_s}(v_1, \dots, v_b) \neq 0$, and hence by Lemma 28 there exist $v_{b+1}, \dots, v_n$ such that $(v_1, \dots, v_n) \in \mathcal{S}_H$, i.e. a satisfying assignment for H must exist. Thus, to solve the original k -QSAT instance Π , it suffices to: (1) apply the Decoupling Lemma to blow up the instance to a decoupled instance Π^+ , (2) apply the transfer functions from the Transfer Lemma to $v_1, \dots, v_b$ to obtain a solution on all $m + b$ vertices for Π^+ , and (3) (easily) map this solution back to one on n vertices for Π .

Example 32 (Running example). Let $H = (H_1, H_2, H_3, H_4)$ be a 3-local collection of functions $H_i : W^4 \rightarrow \mathbb{K}$ whose interaction graph is the 3-uniform cycle of transfer type 2 introduced in Example 10. If p is chosen as in Example 19, then the two qualifier functions are

$$h_1(v_1, v_2) = (g_5^\sharp(v_1, v_2))(g_4(v_1, v_2))$$

of degree (3, 2) and

$$h_2(v_1, v_2) = (g_6^\sharp(v_1, v_2))(g_3(v_1, v_2))$$

of degree (2, 3), where g_3, g_4, g_5, g_6 so that $d_{sr} \leq 3 \leq 10 = 2F_5^{(2)}$ for each $s, r \in \{1, 2\}$, in accordance to Lemma 30. Choosing a basis $\{w_1, w_2\}$ of W as in Example 32, we obtain via Condition 1 of Lemma 30 that $(v_1, v_2, v_3, v_4) \in \mathcal{S}_H$ implies

$$\begin{aligned} 0 = & H_1^*(v_1 \otimes v_2 \otimes w_1)H_2^*(v_1 \otimes v_2 \otimes w_1)H_3^*(v_1 \otimes w_2 \otimes w_2) \\ & + H_1^*(v_1 \otimes v_2 \otimes w_2)H_2^*(v_1 \otimes v_2 \otimes w_2)H_3^*(v_1 \otimes w_1 \otimes w_1) \\ & - H_1^*(v_1 \otimes v_2 \otimes w_1)H_2^*(v_1 \otimes v_2 \otimes w_2)H_3^*(v_1 \otimes w_2 \otimes w_1) \\ & - H_1^*(v_1 \otimes v_2 \otimes w_2)H_2^*(v_1 \otimes v_2 \otimes w_1)H_3^*(v_1 \otimes w_1 \otimes w_2) \end{aligned}$$

and

$$\begin{aligned} 0 = & H_1^*(v_1 \otimes v_2 \otimes w_1)H_2^*(v_1 \otimes v_2 \otimes w_1)H_4^*(v_2 \otimes w_2 \otimes w_2) \\ & + H_1^*(v_1 \otimes v_2 \otimes w_2)H_2^*(v_1 \otimes v_2 \otimes w_2)H_4^*(v_2 \otimes w_1 \otimes w_1) \\ & - H_1^*(v_1 \otimes v_2 \otimes w_1)H_2^*(v_1 \otimes v_2 \otimes w_2)H_4^*(v_2 \otimes w_2 \otimes w_1) \\ & - H_1^*(v_1 \otimes v_2 \otimes w_2)H_2^*(v_1 \otimes v_2 \otimes w_1)H_4^*(v_2 \otimes w_1 \otimes w_2). \end{aligned}$$

4.3. Generic constraints. Remark 31 outlined the high-level strategy for computing a (product-state) solution to an input k -QSAT system Π . For this strategy to work, however, we require an assignment to the foundation of the transfer filtration which (1) satisfies the qualifier functions from the Qualifier Lemma, and (2) causes the transfer functions g_i from the Transfer Lemma to output non-zero vectors. When are (1) and (2) possible? We now answer this question affirmatively for a non-trivial class of k -QSAT instances, assuming constraints are chosen generically. Aside: In this section alone, the terminology of projective algebraic geometry (AG) is used briefly to define generic constraints; otherwise, the content is intended to be accessible without a background in AG.

Remark 33 (Generic constraints). Let G be a k -uniform hypergraph. The set of k -local constraints H with interaction graph G is canonically identified with the projective variety $\mathcal{X}_G(\mathbb{K}) = (\mathbb{P}^{2^k-1}(\mathbb{K}))^m$. (We remark the same variety was used in [8].) We say that a property holds for the *generic constraint with interaction graph G* if it holds for every k -local constraint on a Zariski open set of $\mathcal{X}_G(\mathbb{K})$. In the important case $\mathbb{K} = \mathbb{C}$, this implies in particular that such a property holds for almost all choices of constraints (with respect to the natural measure on $\mathcal{X}_G(\mathbb{C})$ induced by the Fubini-Study metric). Note that we allow different (generic) constraints on each edge (i.e. the Hamiltonian is not necessarily translation-invariant). (QI) To give some intuition, a generic 2-qubit constraint $|\psi\rangle \in \mathbb{C}^2 \otimes \mathbb{C}^2$ would be any entangled $|\psi\rangle$ (since the set of two-qubit states with a Schmidt coefficient of zero forms a zero measure set). In the case of 2-QSAT, the upshot of this is that if $|\psi\rangle$ is a constraint on qubits i and j , then given any assignment $|\phi_1\rangle$ to qubit i , the transfer matrix T_{ij} corresponding to $|\psi\rangle$ is full-rank, meaning $T_{ij}|\phi_1\rangle \neq 0$ (and additionally, by definition of a transfer matrix, $\langle\psi|(I \otimes T_{ij})|\phi\rangle_i \otimes |\phi\rangle_j = 0$). In other words, an assignment on qubit i is guaranteed to propagate to a valid assignment on qubit j . Note that classical k -SAT constraints are not generic, as they represent a finite set in the variety of all constraints (and in particular, correspond to rank-1 projections onto product states). Nevertheless, as we prove in Theorem 50, the hypergraphs to which our main theorem applies (Theorem 38) have an SDR; and a satisfying assignment to any classical k -SAT instance with an SDR is easy to compute (as mentioned in Sect. 1).

Definition 34. Let $G_0 \subseteq G_1 \subseteq \dots \subseteq G$ be a transfer filtration on a k -uniform hypergraph G whose vertices $\{1, \dots, n\}$ are ordered in such a way that if $i \in V(G_j) \setminus V(G_{j-1})$, then $i-1 \in V(G_{j-1})$. Given $i, i' \in \{1, \dots, n\}$ we say that i' is a *successor* of i (and i is a *predecessor* of i') if there exists sequences $i_1, \dots, i_l \in V(G)$ and $E_1, \dots, E_{l+1} \in E(G)$ such that $i \leq i_1 \leq \dots \leq i_l \leq i'$, $i \in E_1, i' \in E_{l+1}$ and $i_j \in E_j \cap E_{j+1}$ for all $j \in \{1, \dots, l\}$.

Example 35. Let G be as in Example 35. Then 3 is a predecessor of 4 and a successor of both 1 and 2.

Before showing the main theorem of this section (Theorem 38), we first require the following lemma, which shows that under certain conditions, the transfer functions g_i of the Transfer Lemma satisfy a surjectivity criterion. (QI) This is roughly a generalization of the idea from 2-QSAT that when all constraints are entangled, then all transfer matrices are full rank (and hence are surjective maps).

Lemma 36 (Surjectivity Lemma). *Let G be a k -uniform hypergraph of transfer type b with n vertices and $n - b$ edges. Let H be a generic k -local constraint in $\mathcal{X}_G(\mathbb{K})$, let $g_1, \dots, g_n : W^b \rightarrow W$ be transfer functions as in the Transfer Lemma (Lemma 28) and*

choose non-zero $v_1, \dots, v_{b-1} \in W$. For each i , define $\gamma_i : W \rightarrow W$ such that

$$\gamma_i(w) := g_i(v_1, \dots, v_{b-1}, w) \quad (4)$$

for all $w \in W$. If i is a successor of b , then γ_i is surjective.

Proof. Up to a permutation of $\{1, \dots, b\}$, we may assume that n is a successor of b . By induction we may assume the claim has been proved for hypergraphs with $n - 1$ vertices (the base case $i = b$ follows since we can take g_b to be the identity, see the first line of the proof of Lemma 28). Therefore, a generic choice of $H_1, \dots, H_{n-1-b}$ leads to surjective maps γ_i for every successor of b . If the edge that contains vertex n is $\{i_1, \dots, i_{k-1}, n\}$, then at least one other vertex, say i_1 , is also a successor of b . Since γ_{i_1} is surjective by the induction hypothesis, we can choose $w', w'' \in W$ such that $\gamma_{i_1}(w')$ and $\gamma_{i_1}(w'')$ are linearly independent. Suppose that γ_n is not surjective. Since it has definite (i.e. well-defined) degree, this implies $\gamma_n(w')$ and $\gamma_n(w'')$ are not linearly independent. By (1) this means that H_{n-b}^* vanishes at two prescribed, linearly independent vectors. Since this condition is not generic, we have a contradiction. Thus, generically γ_n is surjective. $\square$

Remark 37. For each i, j , consider the function $\Gamma_{ij} : W \rightarrow \mathbb{K}$ defined by setting

$$\Gamma_{ij}(w) := (\gamma_i(w))^\sharp(\gamma_j(w))$$

for all $x \in \mathbb{K}$. It is easy to see by induction that if $w', w'' \in W$ are linearly independent, then $P_{ij}(x) = \Gamma_{ij}(w' + xw'')$ is a univariate polynomial in x with coefficients in $\mathbb{K}$. Furthermore, if the constraint H is chosen generically and either i or j is a successor of b then P_{ij} is not a non-zero constant.

The main theorem. We now show the main theorem of this section, which applies to k -uniform hypergraphs of transfer type $b = n - m + 1$. The latter includes non-trivial instances which we discuss in Sect. 4.4 (along with an example of exponential speedup of our parameterized algorithm over brute diagonalization). In words, the theorem says that for any k -uniform hypergraph of transfer type $b = n - m + 1$ (i.e. there is precisely one qualifier function h_1), if the constraints are chosen generically, then any zero of h_1 is the image under the map π (defined in Qualifier Lemma) of a satisfying assignment to the corresponding k -QSAT instance. The key advantage to this approach is simple: To solve the k -QSAT instance, instead of solving a system of equations, we are reduced to solving for the roots of just one polynomial— h_1 . Moreover, if both the foundation size b and the radius of the transfer filtration of G scale as $O(\log m + \log n)$, then h_1 has polynomial degree in m and n .

Theorem 38. *Let $\mathbb{K}$ be algebraically closed, and let $\mathcal{F}$ denote the set of k -uniform hypergraphs with n vertices, m edges, and transfer type $b = n - m + 1$. If H is a generic k -local constraint with interaction graph $G \in \mathcal{F}$ and h_1 and π are as in the Qualifier Lemma (Lemma 30), then $(h_1 \circ \pi)^{-1}(0) \cap \mathcal{S}_H$ is nonempty.*

Proof. Let $\tilde{G}$ and $g_1, \dots, g_{n+1}$ as in the proof of the Qualifier Lemma (Lemma 30). Since $\mathbb{K}$ is algebraically closed, by the discussion in Remark 41, there exist $\hat{v}_1, \dots, \hat{v}_b \in W$ such that $h_1(\hat{v}_1, \dots, \hat{v}_b) = 0$. As shown in the previous section, if additionally $g_{i_1}(\hat{v}_1, \dots, \hat{v}_b) \neq 0$, then $\pi^{-1}(\hat{v}_1, \dots, \hat{v}_b) \cap \mathcal{S}_H \neq \emptyset$, as desired. Thus, we are left to show that the same conclusion holds even if $g_{i_1}(\hat{v}_1, \dots, \hat{v}_b) = 0$ (which immediately implies $h_1(\hat{v}_1, \dots, \hat{v}_b) = 0$).

Let $j \in \{b+1, \dots, n+1\}$ be minimal among those predecessors of i_1 such that $g_j(\hat{v}_1, \dots, \hat{v}_b) = 0$. For a sufficiently generic choice of constraint we may assume that $g_l(\hat{v}_1, \dots, \hat{v}_b) \neq 0$ for any l that is not a successor of j . Our strategy is now to *remove* the edge in $\tilde{G}$ which was added in the same step of the filtration in which vertex j was added, namely edge $\widetilde{E_{j-b}}$. We then instead add vertex j directly to the foundation. Formally, consider hypergraph F obtained from $\tilde{G}$ by removing edge $\widetilde{E_{j-b}}$, and consider transfer filtration $F_0 \subseteq F_1 \subseteq \dots \subseteq F_{m-1} = F$ where:

- (Add j to the foundation) $V(F_i) = V(\tilde{G}_i) \cup \{j\}$ for all $i \in \{0, \dots, j-b-1\}$ and $V(F_i) = V(\tilde{G}_{i+1})$ for all $i \in \{j-b, \dots, m-1\}$,
- (Remove edge $\widetilde{E_{j-b}}$) $E(F_i) = E(\tilde{G}_i)$ for all $i \in \{0, \dots, j-b-1\}$ and $E(F_i) = E(\tilde{G}_{i+1}) \setminus \{\widetilde{E_{j-b}}\}$ for all $i \in \{j-b, \dots, m-1\}$.

The hypergraph F thus has $n_F := n$ vertices, $m_F := m-1 = n-b$ edges, and foundation of size $b_F := b+1 = n-m+2$. Thus, $m_F > n_F - b_F$, implying we can apply the Qualifier Lemma (Lemma 30) and the Surjectivity Lemma (Lemma 36) to obtain transfer functions $f_1, \dots, f_{n_F} : W^{b_F} \mapsto W$ and $\gamma_{1,F}, \dots, \gamma_{n_F,F} : W \mapsto W$ (based on f_i instead of g_i). Two comments are in order here: First, each $\gamma_{i,F}$ depends on $\hat{v}_1, \dots, \hat{v}_b$ (which are fixed) and parameter w , which corresponds to the vertex j added to the foundation in the construction of F . Second, for $i \in \{0, \dots, j-b-1\}$, we have by construction that for all $w \in W$, $f_i(\hat{v}_1, \dots, \hat{v}_b, w) = g_i(\hat{v}_1, \dots, \hat{v}_b)$ —this is because each corresponding $\tilde{G}_i$ (i.e. the first $j-b$ “steps” of the original transfer filtration) does not depend on vertex j .

Now recall from Remark 37 the polynomial

$$P_{i_1, \underline{i}_1}(x) := \Gamma_{i_1, \underline{i}_1}(w' + xw'') = (\gamma_{i_1, F}(w' + xw''))^\sharp (\gamma_{\underline{i}_1, F}(w' + xw'')) \in \mathbb{K}[x].$$

Since $\mathbb{K}$ is algebraically closed, $P_{i_1, \underline{i}_1}(x)$ has a root $\hat{x}$ in $\mathbb{K}$. Thus, by Remark 24, there exists $\lambda \in \mathbb{K}$ such that $\gamma_{i_1, F}(w' + \hat{x}w'') = \lambda \gamma_{\underline{i}_1, F}(w' + \hat{x}w'')$ (i.e. both decoupled vertices i_1 and $\underline{i}_1$ can be given the same assignment). Therefore, if we set $v_i = \gamma_{i, F}(w' + \hat{x}w'')$ for all $i \in \{1, \dots, n\}$ (note we may assume by Remark 39 that each $v_i \neq 0$), then (1) vertices i_1 and $\underline{i}_1$ receive consistent assignments, (2) edge $\widetilde{E_{j-b}}$ is satisfied by the second remark above that $f_i(\hat{v}_1, \dots, \hat{v}_b, w) = g_i(\hat{v}_1, \dots, \hat{v}_b)$ for $i \in \{0, \dots, j-b-1\}$ (here we also use the fact that by Eq. (2), edge $\widetilde{E_{j-b}}$ was satisfied by foundation assignment $(\hat{v}_1, \dots, \hat{v}_b)$ regardless of the assignment to vertex i_1), and (3) $(v_1, \dots, v_n)$ satisfies all constraints in hypergraph F . It follows that $(v_1, \dots, v_n) \in \pi^{-1}(\hat{v}_1, \dots, \hat{v}_b) \cap \mathcal{S}_H$, as required. $\square$

Remark 39 (The maps $\gamma_i(w)$ have no zeroes). The maps $\gamma_i(w)$ from the Surjectivity Lemma (Lemma 36) can without loss of generality be taken to have no zeroes, i.e. $\nexists w \in W$ such that $\gamma_i(w) = 0$. To see this, consider $\gamma_i(w' + xw'') : W \mapsto W$ for $x \in \mathbb{K}$ as in Remark 37. Since w' and w'' are linearly independent, we can write $\gamma_i(w' + xw'') = p(x)w' + q(x)w''$ for polynomials $p, q \in \mathbb{K}[x]$. In order for the right hand side to vanish on x_0 , since $\mathbb{K}$ is algebraically closed, p and q must share a common factor $(x - x_0)$. Factoring out $(x - x_0)$, we obtain new polynomials p' and q' , and redefine $\gamma'_i(w' + xw'') = p'(x)w' + q'(x)w''$. (Note we are implicitly working in projective space, since solutions to the Quantum SAT instance remain solutions under (non-zero) scaling.)

Example 40 (Running example). We illustrate the proof of Theorem 38 by specializing the construction to the 3-uniform semicycle G_3 from Example 10. Then $\tilde{G}_3$ is

the hypergraph with vertices $\{1, 2, 3, 4, 5\}$ and edges $\tilde{E}_1 = \{1, 2, 3\}$, $\tilde{E}_2 = \{1, 2, 4\}$, $\tilde{E}_3 = \{1, 3, 5\}$. Moreover, the transfer functions $g_1, \dots, g_5 : W^2 \rightarrow W$ can be chosen as in Example 29. Let h_1 be as in Example 32 and suppose $v_1, v_2 \in W$ are such that $h_1(v_1, v_2) = 0$. If none of the $g_i(v_1, v_2)$ are zero, then a solution of the form (v_1, v_2, v_3, v_4) can be found by Remark 31. Else, suppose (say) $g_3(v_1, v_2) = 0$ (generically, only one $g_i(v_1, v_2)$ will be zero in this case). Then the hypergraph F has vertices $\{1, 2, 3, 4, 5\}$ and edges $\tilde{E}_2, \tilde{E}_3$. Furthermore, F has a transfer filtration $F_0 \subseteq F_1 \subseteq F_2 = F$, where F_0 has vertices $\{1, 2, 3\}$ and no edges, while F_1 has vertices $\{1, 2, 3, 4\}$ and a single edge $\tilde{E}_2$. Unraveling (2) we obtain

$$\begin{aligned}\gamma_{1,F}(w) &= v_1 ; \\ \gamma_{2,F}(w) &= v_2 ; \\ \gamma_{3,F}(w) &= w ; \\ \gamma_{4,F}(w) &= H_2^*(v_1 \otimes v_2 \otimes w'')w' - H_2^*(v_1 \otimes v_2 \otimes w')w'' ; \\ \gamma_{5,F}(w) &= H_3^*(v_1 \otimes w \otimes w'')w' - H_3^*(v_1 \otimes w \otimes w')w''\end{aligned}$$

and thus

$$\Gamma_{54}(w) = H_3^*(v_1 \otimes w \otimes w')H_2^*(v_1 \otimes v_2 \otimes w'') - H_3^*(v_1 \otimes w \otimes w'')H_2^*(v_1 \otimes v_2 \otimes w')$$

for all $w \in W$. Hence $P_{54}(x) = P_{54}^0 + P_{54}^1 x$ where

$$P_{54}^0 = H_3^*(v_1 \otimes w' \otimes w')H_2^*(v_1 \otimes v_2 \otimes w'') - H_3^*(v_1 \otimes w' \otimes w'')H_2^*(v_1 \otimes v_2 \otimes w')$$

and

$$P_{54}^1 = H_3^*(v_1 \otimes w'' \otimes w')H_2^*(v_1 \otimes v_2 \otimes w'') - H_3^*(v_1 \otimes w'' \otimes w'')H_2^*(v_1 \otimes v_2 \otimes w').$$

In this case it is immediately clear that generically $P_{54}(x)$ is non-constant and, being linear, has a root over any (non-necessarily algebraically closed) field. However, in general, $P_{ii}(x)$ may have arbitrarily high degree and assuming that $\mathbb{K}$ is algebraically closed becomes necessary.

Remark 41 (Reduction to univariate polynomials). Theorem 38 reduces us to solving a single polynomial equation,

$$h_1(v_1, \dots, v_b) = 0,$$

which is multi-variate. In this case, we reduce it further to a univariate polynomial by fixing arbitrary vectors $w_1, \dots, w_b \in W$ and $w'_b \in W$ linearly independent from w_b . Then

$$P(x) = h_1(w_1, \dots, w_b + xw'_b)$$

is a *univariate* polynomial in $\mathbb{K}[x]$, which has a root $x \in \mathbb{K}$ since $\mathbb{K}$ is algebraically closed. Thus, we can reduce the problem of computing a root of the multivariate polynomial h_1 to computing a root of the univariate polynomial P .

Before explicitly stating our algorithm and formally analyzing runtimes in Sect. 4.4, we finally tie Theorem 38 back to SDRs, and give an AG overview of our approach.

Remark 42 (Connection to SDRs). Theorem 38 gives a *constructive* proof that all hypergraphs in $\mathcal{F}$ have a satisfying (product state) assignment if the constraints are chosen generically. This implies each hypergraph in $\mathcal{F}$ has an SDR, as we now observe.

Corollary 43. *If G is a k -uniform hypergraph of transfer type b and such that $|E(G)| = |V(G)| - b + 1$, then G has an SDR.*

Proof. By Theorem 38, $\mathcal{S}_H$ is nonempty for generic $H \in \mathcal{X}_G(\mathbb{C})$. Theorem 2 of [8] then implies that G has an SDR. $\square$

Thus, Theorem 38 constructively recovers the result of [8] (that any k -QSAT instance with an SDR has a (product-state) solution) in the case when the additional conditions of Theorem 38 are met (recall [8] works on all graphs with an SDR, but is not constructive). Finally, in Theorem 50, we show a generalization of Corollary 43 above with a self-contained proof based on Hall’s marriage theorem.

Remark 44 (AG view). As stated at the outset of Sect. 4, we now give a brief sketch of the ideas discussed here in AG terms. It is natural to represent product states using the Segre map $(v_1, \dots, v_n) \rightarrow v_1 \otimes \dots \otimes v_n$ which describes an embedding of the n -fold product of projective lines into a projective space of dimension $2^n - 1$. Using this representation, the product satisfiability $\mathcal{S}_H$ can be naturally thought of as an algebraic set in $(\mathbb{P}^1(\mathbb{K}))^n$. In the language of algebraic geometry, the Transfer Lemma (Lemma 28) says that if the interaction graph is k -uniform of transfer type $b = |V(G)| - |E(G)|$, then $\mathcal{S}_H$ can be realized as an iterated blow-up of $(\mathbb{P}^1(\mathbb{K}))^b$, so in particular it is a rational variety. Moreover, the blow-ups are taken along the loci where the transfer functions g_i vanish or equivalently where their projectivizations (which are by construction rational functions in suitable coordinates) are indeterminate. On the other hand, the Qualifier Lemma (Lemma 30) says that the projection of $\mathcal{S}_H$ along the coordinates corresponding to G_0 is contained in the algebraic set X_H cut out by the regular functions h_s . In general, we cannot guarantee that the projection of $\mathcal{S}_H$ coincides with X_H . However, Theorem 38 says that in the non-trivial case when X_H is a hypersurface, then the projection of $\mathcal{S}_H$ does indeed coincide with X_H .

4.4. Explicit algorithm statement, runtime, and speedups. Figure 8 states our algorithm explicitly. We now give a runtime analysis, as well as families of inputs for which a speedup (polynomial to exponential) is obtained over brute force methods.

Runtime analysis. We bound the complexity of each step of Fig. 8. Note these bounds are rather naive; improved polynomial factors can likely be attained by an appropriate use of data structures and a more careful analysis. Below, we assume $\mathbb{K} = \mathbb{C}$, use of the Random Access Machine model (e.g. looking up entries in arrays costs constant time), and we count the number of field operations over $\mathbb{C}$ (as opposed to bit operations).

- (Decoupling Lemma, Line 5) Let $(E_1, \dots, E_m)$ denote the order in which edges of G are added in the filtration F . In order to compute $\tilde{G}$ and p , the first step is to compute $r : \{1, \dots, m\} \mapsto \{0, \dots, m - 1\}$ from Remark 16. Recall this was defined such that $r(i)$ is the smallest integer satisfying $|E_i \setminus V(G_{r(i)})| = 1$, i.e. the first point in the filtration at which E_i could be added next via the introduction of precisely one vertex. Since $k \in O(1)$, one can compute all vertex sets $V(G_i)$ in $O(m)$ time. With all $V(G_i)$ precomputed, computing $r(i)$ now takes $O(i)$ time, and so precomputing all values of r takes $O(m + m^2) \in O(m^2)$ total time.

Moving on to p , since $p(i) = i$ when $i \in [b]$ and $p(i) = E_{i-b} \setminus V(G_r(i-b))$ when $i \in \{b+1, \dots, b+m\}$, each entry $p(i)$ can now be computed in $O(1)$ time, implying all values of p can be computed in time $O(m+b)$.

Since the preimage $p^{-1}(j)$ can be of size up to m , computing it takes $O(m)$ time, implying computing j and $\tilde{E}_i$ take $O(m)$ time. Therefore, computing all edges $\tilde{E}_i$ takes $O(m(m+b))$ time.

The total runtime for this step is hence $O(m^2) + O(m+b) + O(m(m+b)) \in O(m(m+b))$.

- (Transfer Lemma, Line 6) To analyze the cost of computing a transfer function (where by “computing”, we mean explicitly writing out all coefficients of the corresponding multi-variate polynomial), $g : W^b \mapsto W$, it is helpful to use the construction of Example 29 as a guide. Namely, since W is two-dimensional, it suffices to analyze the maximum cost of outputting a component of g ’s two-dimensional output vector. This requires bounding the cost of expanding an expression of the form

$$H^*(g_{i_1}(v_1, \dots, v_b) \otimes \dots \otimes g_{i_{k-1}}(v_1, \dots, v_b) \otimes w_l) =: H^*(\mathbf{u})$$

(for w_l some fixed basis vector for W) as a multivariate polynomial in the $2b$ components of $v_1, \dots, v_b$. For convenience, let $\mathbf{u} = g_{i_1}(v_1, \dots, v_b) \otimes \dots \otimes g_{i_{k-1}}(v_1, \dots, v_b) \otimes w_l$ above, and set $S = \{i_1, \dots, i_{k-1}\}$. Observe now that if g_{i_t} has degree $(d_{i_t,1}, \dots, d_{i_t,b})$, then since g_{i_t} is a multivariate polynomial, it has at most $O(d_{i_t,1}d_{i_t,2} \dots d_{i_t,b})$ distinct monomials. It follows that computing an arbitrary entry of $\mathbf{u}$ costs $O(\prod_{i \in S} \prod_{j=1}^b d_{i,j})$ time (for this simple analysis, we naively multiply out all monomials and add like terms). Note next that if H denotes the k -local Hamiltonian term corresponding to H^* , then $H^*(\mathbf{u}) = \mathbf{u}^\dagger H \mathbf{u}$. Thus, assuming $k \in O(1)$, computing $H^*(\mathbf{u})$ costs

$$O\left(\prod_{i \in S} \prod_{j=1}^b d_{i,j}^2\right) \quad (5)$$

time in the worst case. By the degree bounds of Lemma 28 and Remark 26, we have $d_{i,j} \leq F_i^{(b)} \in o(2^i)$. Equation (5) can hence be bounded by

$$O\left(\prod_{i \in S} d_{i,j}^{2b}\right) \in O\left(2^{2b \sum_{i \in S} i}\right) \in O\left(2^{2bk \max_{i \in S} i}\right), \quad (6)$$

where the last statement follows since $|S| \leq k$. A naive analysis would now suggest $\max_{i \in S} i \leq n$, as there are n transfer functions g . However, this analysis is too loose—it assumes the worst case that each tensor factor $g_{i_t}(v_1, \dots, v_b)$ depends recursively on all previous transfer functions $g_1, \dots, g_{i_t-1}$ (hence the degree of $g_{i_t-1} \leq F_{i_t-1}^{(b)}$ for $i \in [n]$). By taking the *radius* of the filtration into account, $\rho(G, b)$, we can obtain a much smaller bound when the radius is small. Specifically, similar to the analysis of Lemma 30, one can instead show $d_{i,j} \leq F_{\rho(G,b)+b+1}^{(b)} \in O(2^{\rho(G,b)+b})$. Thus, we can improve Eq. (6) to

$$O\left(2^{2kb(\rho(G,b)+b)}\right).$$

Since there are n transfer functions in total, the cost of Line 6 is hence $O(n2^{2kb(\rho(G,b)+b)})$, where recall we assume $k \in O(1)$.

- (Qualifier Lemma, Line 7). We consider the cost of computing an arbitrary qualifier h_s . Since in Line 5 we already computed $\underline{i}$ for all $i \in \{1, \dots, m+b\}$, we may compile set $\{i_1, \dots, i_{m-n+b}\}$ in time $O(m-n+b)$. Since W is a two-dimensional space, computing $g_{i_s}^\#$ takes time linear in the number of monomials in each amplitude/component of g_{i_s} , i.e. $O(d_{i_s,1} \cdots d_{i_s,b}) \in O(2^{b(\rho(G,b)+b)})$ time. It follows that computing h_s requires $O(2^{2b(\rho(G,b)+b)})$ time. Since there are $m-n+b$ qualifier functions, the total cost of this step is

$$O\left((m-n+b)2^{2b(\rho(G,b)+b)}\right).$$

We may bound the cost of Lines 5–7 (which apply to any instance of k -QSAT) by

$$O\left((m-n+b)2^{2kb(\rho(G,b)+b)}\right). \quad (7)$$

For arbitrary k -QSAT instances, the runtime of Lines 8–9 is less clear (e.g. how to find a common root to all h_s while also ensuring g_{i_s} to be non-zero?); however, these steps can indeed be solved for the family of instances with generic constraints and transfer type $b = n - m + 1$, as demonstrated in Sect. 4.3. As noted in Remark 41, in this case we may reduce Lines 8–9 to solving for the roots of a single univariate polynomial (moreover, the Surjectivity Lemma 36 ensures g_{i_s} outputs a non-zero value). We now analyze this reduction’s runtime.

- (Lines 8–9, generic constraints, $b = n - m + 1$) When $b = n - m + 1$, there is precisely one qualifier function $h_s(v_1, \dots, v_b)$, which is a multivariate polynomial in $2b$ complex variables, and has degree $(d_{s1}, \dots, d_{sb})$ for $d_{sr} \leq 2^{\rho(G,b)+b+2}$ (by Lemma 30) for all $r \in [b]$. Following Remark 41, fix $w_1 = \dots = w_b = |0\rangle \in \mathbb{C}^2$ and $w'_b = |1\rangle \in \mathbb{C}^2$ and consider variable $x \in \mathbb{C}$. Then, $q(x) := h_s(w_1, \dots, w_{b-1}, w_b + xw'_b)$ is a univariate non-constant¹⁰ polynomial over x with degree at most

$$d \leq 2^{\rho(G,b)+b+2}. \quad (8)$$

Observe that h_s has at most $M = \prod_s d_{sr} \leq 2^{b(\rho(G,b)+b+2)}$ monomials. Thus, computing $q(x)$ takes $O(M)$ time since each of the $2b$ complex variables plugged into h_s is given a value from set $\{0, 1, x\}$ (i.e. simplifying monomials with large exponents is trivial). To find a root of q , which exists since $\mathbb{C}$ is algebraically closed, we can now apply an algorithm of Schönhage [14] as follows.

Schönhage’s univariate polynomial factorization algorithm. Let us write $q(x) = \sum_{i=1}^d c_i x^i \in \mathbb{C}[x]$, and define one-norm $\|q\|_1 := \sum_{i=1}^d |c_i|$. Then, Eq. (3.2) of [14] says “approximate factorization” of q is possible within error¹¹ $\epsilon = 2^{-N} \|q\|_1$ in

¹⁰ Note that, in principle, $q(x)$ could be constant. However, since h_s is non-constant by the Qualifier Lemma, there must be at least one coordinate $j \in \{1, \dots, b\}$ such that redefining $q(x)$ to have its j th argument to equal $w_j + xw'_j$, and all other arguments $j' \neq j$ equal to $|\psi\rangle$ for some appropriate constant vector $|\psi\rangle$, yields a non-constant polynomial over x . (For example, set the entries of $|\psi\rangle$ to be sufficiently large in magnitude to ensure the monomials of different degrees in the expansion of h_s cannot cancel out. Recall here that since we are essentially working in projective space, $|\psi\rangle$ need not be normalized.)

¹¹ For clarity, Eq. (3.2) of [14] assumes $\|q\|_1 \leq 1$, and hence the error is stated therein as $\epsilon = 2^{-N}$. In our setting, however, we first need to rescale our q to ensure $\|q\|_1 \leq 1$. This does not affect the number of field operations required for Schönhage’s algorithm, but results in the multiplicative $\|q\|_1$ term (which is absent in [14]) in the resulting bound on accuracy, Eq. (9). We hence define the error as $\epsilon = 2^{-N} \|q\|_1$ here for convenience.

$O(d^3 \log d + d^2 N)$ field operations. Here, “approximate factorization” within error ϵ means computing linear factors $L_j(x) = u_j x + v_j$ for $1 \leq j \leq d$ such that

$$\|q(x) - L_1(x) \cdots L_d(x)\|_1 < 2^{-N} \|q\|_1. \quad (9)$$

Let $w_j = -v_j/u_j \leq 1$ denote the root of linear factor $L_j(x)$. Given Eq. (9), the next question is: How closely do the w_j approximate the roots z_j of $q(x)$? As done in Sect. 3.4 of [14], one can apply the perturbation bound of Theorem 2.7 of [13] to obtain the following. If $\|q(x) - L_1(x) \cdots L_d(x)\|_1 \leq \eta \|q(x)\|_1$ for $\eta \leq 2^{-7d}$, then there exists an ordering of the w_j such that for all j ,

$$|w_j - z_j| < 9\eta^{1/d}.$$

It follows that in order to compute the roots z_j within additive error $2^{-p(n)}$ for some polynomial p , it suffices to approximately factorize q within error $\epsilon = \|q\|_1 / (9^d 2^{p(n)d})$, which can be accomplished in time

$$O\left(d^3 \log d + d^2 \log\left(9^d 2^{p(n)d}\right)\right).$$

- (Line 9) We now have a root $(v_1 := |0\rangle, \dots, v_{b-1} := |0\rangle, v_b := |0\rangle + x|1\rangle)$ of h_s . To compute $g_i(v_1, \dots, v_b)$ for any $i \in [n]$ involves evaluating a multivariate polynomial of degree $(d_{s1}, \dots, d_{sb})$ for $d_{sr} \leq 2^{\rho(G,b)+b+2}$ for all $r \in [b]$. Naively substituting in our root and using square-and-multiply to compute monomials with large powers, evaluating each g_i hence requires at most time

$$O((\rho(G, b) + b)2^{b(\rho(G,b)+b+2)}).$$

Combining the cost of all steps, we find k -QSAT instances with generic constraints and $b = n - m + 1$ require total time at most (for $d \leq 2^{\rho(G,b)+b+2}$)

$$O(mn) + O\left(2^{2kb(\rho(G,b)+b)}\right) + O\left(d^3 \log d + d^2 \log\left(9^d 2^{p(n)d}\right)\right) + O((\rho(G, b) + b)2^{b(\rho(G,b)+b+2)}). \quad (10)$$

Thus, the algorithm is polynomial in m and n , and exponential in k (the locality of the constraints), $\rho(G, b)$ (the radius), and b (foundation size).

Aside: Error propagation. Note that in Line 8, we compute a root z of q within additive error $2^{-p(n)}$. In Line 9, we then substitute this root into our transfer functions g_i to obtain our final assignment on qubit i . This substitution can propagate the additive error from Line 8. However, since the degrees of the transfer functions are bounded by $\exp(n)$ (specifically, $d_{sr} \leq 2^{\rho(G,b)+b+2}$), it follows that by choosing p to be a sufficiently large fixed polynomial, one can make the additive error for the output of g_i exponentially small in n .

On exponential speedups via Theorem. 38 Recall Theorem 38 applies to k -uniform hypergraphs of transfer type $b = n - m + 1$, such as the semicycle. From a parameterized complexity perspective, however, most interesting are hypergraphs for which the foundation size b and filtration radius r satisfy $b, r \in o(n + m)$, for which we might obtain an asymptotic speedup over brute force diagonalization of the Quantum SAT system. We now discuss various hypergraph families and analyze their parameters m and n versus b and r . In particular, we obtain quadratic (tiling of torus, fir tree) to *exponential* (crash) separations between these parameters. Note that for the runtime of Eq. (10), a quadratic separation is unfortunately not enough for an asymptotic speedup. However, an *exponential* separation in parameters, as for the crash family of hypergraphs, implies our parameterized algorithm runs in polynomial time, whereas brute force diagonalization would require time exponential in m and n .

Example 45 (Semicycles (Fig. 1)). We begin with the study of semicycles, for which there is no separation in parameters. Namely, let $S_{t,k}$ denote the k -uniform hypergraph with vertices $V(S_{t,k}) = \mathbb{Z}/t\mathbb{Z}$ and edges $E(S_{t,k}) = \{E_0, \dots, E_{t-k+1}\}$ such that $E_i = \{i, i+1, \dots, i+k-1\}$ for every $i \in \{0, \dots, t-k+1\}$. By construction, $|V(S_{t,k})| - |E(S_{t,k})| + 1 = k - 1$ for every t . On the other hand, $S_{t,k}$ is line graph connected for every t and thus has transfer type $k - 1$ by Example 12. Depending on how the transfer filtration is chosen, the transfer radius is at least $|E(S_{t,k})|/2$ and at most $|E(S_{t,k})|$. Thus, the radius r satisfies $r \in \Theta(|E(S_{t,k})|)$, i.e. no separation in parameters exists.

Example 46 (Modified tiling of the torus (Fig. 2)). We next discuss a slight modification of the tiling of torus which yields a quadratic separation in parameters. Let $k \geq 3$ and let t be a positive integer. Let $V_{t,k}$ be the set of pairs (p, q) of integers such that $0 \leq (k-2)q \leq p \leq (k-2)t$. Let $G_{t,k}$ be the k -uniform hypergraph with vertices $V(G_{t,k}) = V_{t,k}$ and edges $E_{r,s} = \{(r, s), (r+1, s), \dots, (r+k-2, s), (r+k-2, s+1)\}$ for each pair $(r, s) \in V_{t-1,k}$. Then $G_{t,k}$ has $\frac{(t+1)(2+t(k-2))}{2}$ vertices and $\frac{t(2+(t-1)(k-2))}{2}$ edges. All vertices except for $(0, 0)$, $(t(k-2), 0)$ and $(t(k-2), t)$ have degree at least 2. There exists a transfer filtration with foundation $\{(0, 0), \dots, (t(k-2), 0)\}$ obtained by adding the edges in the following order

$$\begin{aligned} &E_{0,0}, E_{1,0}, \dots, E_{(t-1)(k-2),0}, E_{k-2,1}, E_{k-1,1}, \\ &\dots, E_{(t-1)(k-2),1}, E_{2(k-2),2}, \dots, E_{(t-1)(k-2),t-1}. \end{aligned}$$

In particular, this filtration has radius t .

Now, consider the k -uniform hypergraph $T_{t,k}$ obtained by identifying the degree 1 vertices of $G_{t,k}$. (This can be visualized similar to the tiling of the torus as in Fig. 2, except one keeps only the “lower triangular” portion of Fig. 2, i.e. vertices v_2, v_3, v_4, v_5, v_7 from Fig. 2 are discarded. Unlike the tiling of the torus, however, the example here has a transfer filtration type which satisfies the preconditions of Theorem 38, as stated below.) Then $T_{t,k}$ has a transfer filtration of type

$$\begin{aligned} t(k-2) &= \frac{(t+1)(2+t(k-2))}{2} - 2 + \frac{t(2+(t-1)(k-2))}{2} + 1 \\ &= |V(T_{t,k})| - |E(T_{t,k})| + 1, \end{aligned}$$

satisfying the preconditions for Theorem 38. Since $T_{t,k}$ has radius t , it follows that the radius and foundation size are $\Theta(t)$ (for $k \in O(1)$, which is the typical assumption), whereas $|V(T_{t,k})|, |E(T_{t,k})| \in \Theta(t^2)$, i.e. there is a quadratic gap between these parameters.

Example 47 (Fir Tree (Fig. 3)). We next give another example of a hypergraph family (Fig. 3) with a quadratic separation in parameters. For each $k \geq 3$, $t \geq 1$ and $i \in \mathbb{Z}/k\mathbb{Z}$ let $G_{t,k,i}$ be a copy of the k -uniform hypergraph $G_{t,k}$ introduced in Example 46. Let $H'_{t,k}$ be obtained by identifying the vertex $(t(k-2), 0)$ of $G_{t,k,i}$ with the vertex $(0, 0)$ of $G_{t,k,i+1}$ for each $i \in \mathbb{Z}/k\mathbb{Z}$ and let $H_{t,k}$ be the graph obtained from $H'_{t,k}$ by adding the edge whose elements are the vertices $(t(k-2), t)$ for each of the $G_{t,k,i}$. For example, $H_{4,3}$ is illustrated in Figure 3. Then $H_{t,k}$ has a filtration of type

$$\begin{aligned} tk(k-2) &= k \left(\frac{(t+1)(2+t(k-2))}{2} - 1 \right) - k \frac{t(2+(t-1)(k-2))}{2} - 1 + 1 \\ &= |V(H_{t,k})| - |E(H_{t,k})| + 1 \end{aligned}$$

and radius $t+1$. Thus, we have another family of hypergraphs which satisfy the preconditions of Theorem 38, and have a quadratic gap between the radius and foundation size (both $\Theta(t)$) versus the number of vertices and edges (both $\Theta(t^2)$).

Example 48 (Crash (Fig. 4)). Finally, we give a k -uniform hypergraph family with an exponential separation in parameters (Fig. 4). For $k \geq 2$ an integer, let $\Sigma = \{1, 2, \dots, k-1\}$ be an alphabet of size $k-1$. For $t \geq 1$ an integer, consider the hypergraph $C_{t,k}$ with vertices $V(C_{t,k}) = \bigcup_{j=0}^t V_j$ where $V_j = \Sigma^{t-j+1}$ for all $1 \leq j \leq t$, and $V_0 = \{(0, x) \mid x \in \Sigma\}$. The edge set of $C_{t,k}$ is the union of all edges of the following three forms:

1. For every $x \in V_1$, $E_x = \{x\} \cup V_0$;
2. for every $2 \leq j \leq t$ and every $x \in V_j$, $E_x = \{x\} \cup \{xa \mid a \in \Sigma\}$;
3. $E_0 = \{(0, 1)\} \cup V_t$.

Then $C_{t,k}$ has a transfer filtration with foundation V_0 obtained by first adding all the edges E_x with $x \in V_1$, then adding all the edges E_x with $x \in V_2$ etc. with E_0 added last. Clearly this transfer filtration has radius t and type

$$\begin{aligned} k-1 &= (k-1) \left(1 + \frac{(k-1)^t - 1}{k-2} \right) - \frac{(k-1)((k-1)^t - 1)}{k-2} - 1 + 1 \\ &= |V(C_{t,k})| - |E(C_{t,k})| + 1. \end{aligned}$$

In particular, this family satisfies the preconditions for Theorem 38 and yields an exponential separation between radius and foundation size ($\Theta(t)$ and $O(k)$, respectively) versus the number of vertices and edges ($\Theta[(k-1)^t]$ for constant $k \geq 3$). Thus, via the runtime of Eq. (10), our algorithm runs in polynomial time on crash (i.e. time $\text{poly}(|V(C_{t,k})|, |E(C_{t,k})|)$), whereas brute force diagonalization would require exponential time in $|V(C_{t,k})|$. Intuitively, one may think of crash hypergraphs as a hypergraph-version of trees; a crash hypergraph “starts small” (has a small foundation, corresponding to a tree root), and then rapidly expands “horizontally”, resulting a hypergraph whose “depth” is short relative to the number of vertices (just as the depth of a complete binary tree is logarithmic in the size of the tree). Whether such hypergraphs are physically relevant is not clear; here, this family is presented primarily as a proof of concept demonstrating that exponential separations (versus brute force) are indeed possible for our algorithm.

5. On the Structure of 3-Uniform Hypergraphs

In this section, we take a structural graph theory perspective, and make steps towards understanding the set of 3-uniform hypergraphs $G = (V, E)$ in which each edge is matched to a unique vertex. In the combinatorics literature, such a matching for a set system is called a *System of Distinct Representatives (SDR)* [10]. Formally, an SDR for a sequence of sets $S_1, \dots, S_m$ is a sequence of distinct elements $x_1, \dots, x_m$ such that $x_i \in S_i$ for all $i \in [m]$. Hall's well-known Marriage Theorem says that a set system $\{S_i\}_{i=1}^m$ has an SDR if and only if for any $I \subseteq [m]$, $|\bigcup_{i \in I} S_i| \geq |I|$.

To begin, let us state a known corollary of the Marriage Theorem, which will help foreshadow the complexity involved in attempting to characterize the set of 3-uniform hypergraphs with SDRs. Let $m = |E|$ and $n = |V|$ for brevity.

Corollary 49 (See, e.g., Corollary 5.2 of [10]). *Let $G = (V, E)$ be a d -regular r -hypergraph. If $m \leq n$, then the hypergraph has an SDR.*

Thus, the class of hypergraphs with SDRs is, in the sense above, not necessarily small. In this section, we focus our attention on the edge case $m = n$. In this case, any 2-uniform hypergraph (i.e. a *graph*) must be a cycle (up to attaching disjoint paths to each vertex of the cycle). In contrast, we will see that the set of 3-uniform hypergraphs with $m = n$ and with SDRs is more complex.

Aside. In Corollary 49, note via doublecounting that we must have $nd = mr$ (i.e. add up all degrees on the left, and take the union of all edges on the right); thus, if $r = 3$ and $m = n$, we must have regularity $d = 3$. This yields that any 3-regular 3-uniform hypergraph with $m = n$ has an SDR.

Finally, using the Marriage Theorem we obtain an elementary proof of a generalization of Corollary 43.

Theorem 50. *Let G be a k -uniform hypergraph of transfer type $b \leq |V(G)| - |E(G)| + k - 1$. Then G has an SDR. Moreover, this bound is tight.*

Proof. Consider a transfer filtration $G_0 \subseteq G_1 \subseteq \dots \subseteq G_m = G$ so that G has m edges and transfer type $b = |V(G_0)|$. Suppose the edges $E(G) = \{E_1, \dots, E_m\}$ are ordered in such a way that $E(G_i) \setminus E(G_{i-1}) = \{E_i\}$ for all $i \in \{1, \dots, m\}$. By the Marriage Theorem, it suffices to prove that $|E_{i_1} \cup \dots \cup E_{i_l}| \geq l$ for any increasing sequence $1 \leq i_1 < i_2 < \dots < i_l \leq m$. For each $i \in \{1, \dots, m\}$, let $\mu_i = |V(G_{i-1})| - |V(G_i)| + 1$. By definition of transfer filtration, $\mu_i \in \{0, 1\}$ for each $i \in \{1, \dots, m\}$. More precisely, $\mu_i = 0$ if and only if E_i contains a vertex that does not belong to any E_j with $j < i$. Hence, given any increasing sequence $1 \leq i_1 < i_2 < \dots < i_l \leq m$, we have $|E_{i_1}| = k$, $|E_{i_1} \cup E_{i_2}| \geq k + 1 - \mu_{i_2}$ and, iteratively adding one edge at the time,

$$|E_{i_1} \cup E_{i_2} \cup \dots \cup E_{i_l}| \geq k + l - 1 - (\mu_{i_2} + \dots + \mu_{i_l}). \quad (11)$$

Since by assumption we have

$$\mu_{i_2} + \dots + \mu_{i_l} \leq \mu_1 + \mu_2 + \dots + \mu_m = b - |V(G)| + |E(G)| \leq k - 1. \quad (12)$$

Substituting (12) into (11) we obtain

$$|E_{i_1} \cup E_{i_2} \cup \dots \cup E_{i_l}| \geq k + l - 1 - (k - 1) \geq l$$

which, by the Marriage Theorem, proves the bound. to prove that the bound is tight, consider a k -uniform hypergraph G with $|V(G)| = k$ and $|E(G)| = k + 1$ so that the same edge is repeated $k + 1$ times. Clearly there exist a transfer filtration of type $k - 1 > k - 2 = |V(G)| - |E(G)| + k - 1$ and (since $|V(G)| < |E(G)|$), G has no SDR. Hence the bound is tight. $\square$

5.1. Intersecting families. We begin with the study of 3-uniform hypergraphs which are *intersecting families*. Here, a set system S is an *intersecting family* if any pair of sets has non-empty intersection. We say the system is *k-intersecting* if any distinct pair of sets $A, B \in S$ satisfies $|A \cap B| = k$.

5.1.1. 2-Intersecting families For 2-intersecting families, we actually characterize a larger class: The set of 3-uniform hypergraphs with the weaker condition that if a pair of distinct edges e_i and e_j intersect, then either $|e_i \cap e_j| = 2$ (Lemma 51) or $|e_i \cap e_j| \geq 2$ (Lemma 53). The characterization for 2-intersecting families follows as an immediate corollary, and is given in Corollary 52.

Lemma 51. *Let $G = (V, E)$ be a 3-uniform hypergraph such that:*

1. $m = n$ and G has an SDR, and
2. *If a pair of distinct edges e_i and e_j intersects, then $|e_i \cap e_j| = 2$.*

Then G is a cycle of length 4.

Proof. We proceed by case analysis. We begin with a pair of intersecting edges (WLOG we may assume one such pair exists) as in Fig. 12.

Fig. 12. A pair of edges with intersection size 2

We now analyze how one can “grow” the graph by adding hyperedges. Consider first the formation of a tight star as in Fig. 5. We argue that this structure is impossible, given conditions 1 and 2. To see why, denote vertices $\{v_1, v_4, v_5\}$ as “outer” vertices and $\{v_2, v_3\}$ as “inner” vertices. By symmetry, there are seven options for adding a new edge to the tight star:

- (i) Add an edge containing a new vertex v_6 , an outer vertex, and an inner vertex, e.g. $\{v_1, v_2, v_6\}$. This violates property 2 since $|\{v_1, v_2, v_6\} \cap \{v_2, v_3, v_4\}| = 1$.

- (ii) Add an edge containing one new vertex v_6 , and two outer vertices, e.g. $\{v_1, v_5, v_6\}$. This violates property 2, since $|\{v_1, v_5, v_6\} \cap \{v_2, v_3, v_5\}| = 1$.

- (iii) Add an edge containing one new vertex v_6 , two inner vertices, e.g. $\{v_2, v_3, v_6\}$, which is the only case not violating property 2.

- (iv) Add an edge containing two outer vertices, and one inner vertex, e.g. $\{v_1, v_2, v_5\}$, which violates property 2, since $|\{v_1, v_2, v_5\} \cap \{v_2, v_3, v_4\}| = 1$.

- (v) Add an edge containing two inner vertices and one outer vertex, e.g. $\{v_2, v_3, v_4\}$, which clearly violates property 2.

- (vi) Add an edge containing three outer vertices, e.g. $\{v_1, v_5, v_4\}$, which violates property 2 since $|\{v_1, v_5, v_4\} \cap \{v_2, v_3, v_4\}| = 1$.

- (vii) Add two new vertices, and one inner or outer vertex. This violates property 2.

We conclude that the only possibility is case (iii), which grows the tight star. In other words, once we have a tight star, our only option is to continue to build a larger tight star (with distinct edges). But a tight star cannot satisfy property 1, since it has $n = m + 2$. Thus, G cannot contain a tight star (consisting of more than 2 edges).

Let us continue our case analysis of how Fig. 12 can be extended:

- (a) Add a new vertex and two outer vertices, which violates property 2.

- (b) Add a new vertex, one inner vertex, and one outer vertex, which violates property 2, e.g. in:

- (c) The only remaining option is to add an edge such as the blue one below. This temporarily violates property 1, since $m \neq n$.

An analogous case analysis shows that the only way to proceed is to add the yellow edge below, which yields a cycle of length 4, and which satisfies properties 1 and 2. No further edges can be added, completing the proof.

□

Lemma 51 yields the following immediate corollary, obtained by observing that the cycle of length 4 is in fact a 2-intersecting family.

Corollary 52. *The unique 3-uniform hypergraph which (1) is a 2-intersecting family, (2) has $m = n$, and (3) has an SDR is the cycle of length 4.*

More generally, we can relax the statement of Lemma 51 to enforce only $|e_i \cap e_j| \geq 2$ (as opposed to $|e_i \cap e_j| = 2$) and still obtain a characterization. Note that although the statement of Lemma 53 below does not require G to *a priori* be an intersecting family, the graphs obtained in the characterization are indeed intersecting families (but not necessarily k -intersecting for some fixed k).

Lemma 53. *Let $G = (V, E)$ be a 3-uniform hypergraph with property 1 in Lemma 51, and assume each pair of edges e_i and e_j which intersects satisfies $|e_i \cap e_j| \geq 2$. Then G is one of the following:*

- (i) A 3-stacked set, or
- (ii) a tight star with two 2-stacked sets, or
- (iii) a tight star with one 3-stacked set, or
- (iv) a cycle of length 4.

Proof. Case (i) trivially holds, and case (iv) is given by Lemma 53. For cases (ii) and (iii), we return to the case analysis in the proof of Lemma 51 involving Fig. 5, i.e. the tight star. Following the proof of Lemma 51, once we have a tight star, we only have two choices each time we add an edge: (1) Grow the tight star (case (iii) in Lemma 51), or (2) create a parallel edge (case (v) in Lemma 51); note this latter case is now possible since we can have $|e_i \cap e_j| = 3$. The former choice adds a new vertex and new edge, preserving the invariant $n - m$. The latter choice keeps n fixed but increments m by 1, implying $n - m$ decreases by 1. It follows that one must add precisely two parallel edges, yielding cases (ii) and (iii). Finally, analyzing cases (a), (b), and (c) in Lemma 51 yields that no other choices of G are possible. □

Fig. 13. The Fano hypergraph

5.1.2. 1-Intersecting families We next consider the set of 3-uniform hypergraphs in which each pair of edges intersects in precisely one vertex. In this case, it turns out there is a unique 3-uniform hypergraph with $m = n$ and an SDR—the well-known *Fano plane*.

Definition 54 (*Fano plane* (see, e.g., [10])). The Fano plane is the unique projective plane of order 2, and can be represented by the *Fano hypergraph* $G = (V, E)$ (see Fig. 13) with $V = \{v_1, v_2, v_3, v_4, v_5, v_6, v_7\}$ and

$$E = \{\{v_1, v_2, v_3\}, \{v_1, v_4, v_5\}, \{v_1, v_6, v_7\}, \{v_2, v_4, v_6\}, \\ \times \{v_2, v_5, v_7\}, \{v_3, v_4, v_7\}, \{v_3, v_5, v_6\}\}.$$

Lemma 55. *The unique 3-uniform hypergraph which is (1) a 1-intersecting family, (2) has $m = n$, and (3) has an SDR is the Fano hypergraph.*

Proof. Let $G = (V, E)$ satisfy the properties of the claim. WLOG, we begin with an arbitrary pair of intersecting edges, $e_1 = \{v_1, v_2, v_3\}$ and $e_2 = \{v_1, v_4, v_5\}$, pictured below:

We proceed by adding edges one at a time while maintaining property (1). To add the third edge e_3 , since the graph is 1-intersecting, there are two cases:

- (a) $e_3 = \{v_1, v_6, v_7\}$ intersects both e_1 and e_2 at vertex v_1 .

- (b) $e_3 = \{v_3, v_5, v_6\}$ does not intersect e_1 and e_2 at v_1 . By symmetry, we may represent this as in Fig. 14.

Fig. 14. A (linear) 3-cycle

Let us now add the fourth edge, which can be done in one of two ways: Either we add a new vertex, or we do not add a new vertex.

- If we add a new vertex: For case (a), we obtain a graph isomorphic to

This does not have $m = n$ and cannot be extended into a 1-intersecting family since any further edge (of size 3) would have to intersect 4 existing edges. For (b), a case analysis reveals that the only way to extend by adding a new vertex while maintaining the 1-intersection property is via a graph isomorphic to:

- If we do not add a new vertex:
 1. Case (a) can only be extended while maintaining the 1-intersecting property via a graph isomorphic to that in Fig. 15.
 2. For case (b), either we obtain a graph isomorphic to Fig. 15, or we add $e_4 = \{v_2, v_4, v_6\}$ to obtain Fig. 16.

From the fifth edge onward, an analogous case analysis yields at each step that only two types of edges can be added at this point which preserve the 1-intersecting property: A “crossing edge” (as in Fig. 15) $\{v_1, v_6, v_7\}$, $\{v_3, v_4, v_7\}$, or $\{v_2, v_5, v_7\}$, or a “circular edge” (as in Fig. 16) $\{v_2, v_4, v_6\}$. When all such edges are added, we arrive at 7 edges total and obtain the Fano hypergraph. No fewer edges satisfies the $m = n$ property, and clearly

Fig. 15. How to add a fourth edge to extend case (b) via a new vertex

Fig. 16. One of two ways to add a 4th edge to extend case (b) without adding a new vertex

one cannot add another edge at this point while satisfying the 1-intersecting property. Finally, that G has an SDR can either be verified directly or by applying Corollary 49. $\square$

5.2. Linear hypergraphs. In Sect. 5.1.1, we studied the set of 3-uniform hypergraphs in which if a pair of distinct edges intersect, then the intersection size is at least 2. We now take the complementary approach by asking that any pair of distinct edges intersect on *at most one* vertex. Such hypergraphs are called *linear*; note that linear hypergraphs are not necessarily 1-intersecting.

To set the context, linear hypergraphs in general (i.e. not necessarily with an SDR) are a complicated, but well-studied, class. Let L_3^l denote the set of edge intersection graphs of linear 3-uniform hypergraphs. (An *edge intersection graph* (EIG) is the generalization of a line graph to the setting of hypergraphs G ; namely, the vertices of the EIG are the hyperedges of G , and two vertices of the EIG are neighbors if and only if their hyperedges in G intersect. Note that an EIG is a graph, i.e. a 2-uniform hypergraph.) It is known that L_3^l has no “finite characterization” in terms of a finite list of forbidden induced subgraphs [15]. However, the same paper does give the following characterization: A graph G is in L_k^l if and only if G has so-called *Krausz dimension* k . Unfortunately, it was later shown [38] that deciding if Krausz dimension is at most 3 is NP-complete (even on planar graphs of degree at most 5); thus, determining if a graph is an EIG of a linear 3-uniform hypergraph is NP-complete, suggesting the class of linear 3-uniform hypergraphs is quite complex.

In our setting, we study linear 3-uniform hypergraphs G with the additional guarantees that $m = n$ and that G has an SDR. Nevertheless, we are not able to complete a structural characterization. Instead, we find various interesting “canonical-looking” examples. For brevity, we henceforth denote a linear 3-uniform hypergraph with $m = n$ and an SDR as an LGraph.

Minimum size and the Fano plane. What is the minimum size an LGraph can have?

Theorem 56. *The minimum size of an LGraph is $m = n = 7$, and the Fano hypergraph (Fig. 13) is the unique LGraph of this size.*

Proof. The lower bound on minimum size follows from an old lemma of Corrádi.

Lemma 57 (Corrádi 1969 [10]). *Let $A_1, \dots, A_N$ be r -element sets and X be their union. If $|A_i \cap A_j| \leq k$ for all $i \neq j$, then $|X| \geq (r^2 N)/(r + (N - 1)k)$.*

Plugging in our LGraph parameters $k = 1$, $r = 3$, $E = \{A_1, \dots, A_m\}$, $V = X = \{1, \dots, n\}$, yields that $|E| = |V| \geq 7$. Note that the Fano hypergraph is an LGraph with

$m = n = 7$, which also happens to be 1-intersecting. To show uniqueness, we hence show that any LGraph with $m = n = 7$ must be 1-intersecting; the claim then follows from Lemma 55.

To show that any LGraph with $m = n = 7$ must be 1-intersecting, suppose, for sake of contradiction, that $G = (V, E)$ is an LGraph with $m = n = 7$ and edges e_1, e_2 with $e_1 \cap e_2 = \emptyset$. Then, the subgraph of G induced by $\{e_1, e_2\}$ is isomorphic to

Since G is linear, any edge we add must contain v_4 , and in turn, must also contain a *unique* vertex from each of $\{v_1, v_2, v_3\}$ and $\{v_5, v_6, v_7\}$. Thus, the maximum number of edges we may add is 3, giving a total of 5 edges, which is a contradiction. $\square$

The tiling of the torus. There is a sense in which the Fano hypergraph is the “most compact” LGraph. On the opposite extreme, an example of what seems the “least compact” LGraph is the “tiling of the torus”, given in Fig. 2. More precisely, this is the specialization of Example 15 to the case $k = 3$ and $a_1 = a_2 = 3$.

The Helly property. Another well-studied hypergraph property is the *Helly* property, which requires that any intersecting family F of hyperedges of the hypergraph G have non-empty intersection. An equivalent, perhaps more geometrically intuitive, characterization is the following:

Corollary 58 (Page 23 of [39], stated as Corollary 5.1 of [40]). *A hypergraph H has the Helly property if and only if for any three vertices a_1, a_2, a_3 , the family of hyperedges which contains at least two of these vertices has non-empty intersection.*

In the special case of linear hypergraphs, the Helly property forbids a “linear 3-cycle”, as illustrated in Fig. 14. In this sense, linear Helly hypergraphs roughly generalize triangle-free graphs.

Observe now that the two “canonical-looking” LGraphs we have discussed thus far, the Fano plane and the tiling of the torus, do *not* satisfy the Helly property. This raises the question: Do all LGraphs violate the Helly property? In Lemma 59, we answer this question negatively by giving an LGraph with the Helly property. Thus, even with numerous restrictions (i.e. $m = n$, linear, Helly), the set of 3-uniform hypergraphs with SDRs seems non-trivial.

Lemma 59. *There exists a linear 3-uniform hypergraph $G = (V, E)$ with $|E| = |V| = 28$ such that G satisfies the Helly property, and G has an SDR.*

e	v	e	v	e	v	e	v	e	v
1	(1,2, <u>28</u>)	7	(12, <u>13</u> ,14)	13	(24,25, <u>26</u>)	19	(8,14, <u>19</u>)	25	(3,15, <u>27</u>)
2	(2,3,4)	8	(14,15,16)	14	(26, <u>27</u> ,28)	20	(6,12, <u>21</u>)	26	(1, <u>6</u> ,24)
3	(<u>4</u> ,5,6)	9	(16,17,18)	15	(<u>5</u> ,22,28)	21	(4,10, <u>23</u>)	27	(1,10,20)
4	(6,7, <u>8</u>)	10	(18,19,20)	16	(7,20,26)	22	(2,8,25)	28	(1,13, <u>17</u>)
5	(8,9, <u>10</u>)	11	(20,21,22)	17	(<u>9</u> ,18,24)	23	(10,15, <u>20</u>)		
6	(10,11, <u>12</u>)	12	(22,23,24)	18	(<u>11</u> ,16,22)	24	(6, <u>15</u> ,24)		

Fig. 17. A formal specification of the iCycle. The columns labelled e and v index the edges and 3-tuples corresponding to those edges, respectively. The underlined element of each 3-tuple denotes the qubit matched to that edge

Fig. 18. The iCycle

Proof. Since the Helly property forbids a linear 3-cycle, the idea of the construction is to build a longer linear cycle, with added “links” between vertices to ensure $m = n$. We call this the “interlinked cycle” (iCycle) for short; a formal specification is given in Fig. 17, and a graphical depiction in Fig. 18. The claimed properties are tedious, but straightforward, to verify. $\square$

Do all linear 3-uniform hypergraphs with $m = n$ have an SDR? The examples we have studied in this section are all LGraphs, even with constraints such as the Helly property. This raises the question: Could *all* linear 3-uniform hypergraphs with $m = n$ have SDR’s? Intuitively, this is plausible, as the linear property forces a 3-uniform hypergraph to be “spread out” in a rough sense. We now answer this question in the negative, and the construction of the counterexample will teach us something about the structure of LGraphs.

For this, we first require a definition of a *block* and useful lemma.

Definition 60 (*Block*). For hypergraph $G = (V, E)$, let $G_{E'}$ denote the edge-induced subgraph of G by $E' \subseteq E$. Then, we call $G_{E'}$ a *block* if it has an equal number of vertices and edges.

Lemma 61. *Let $G = (V, E)$ be a hypergraph with an SDR. Let A denote a block in G . Then, the SDR matches the vertices in A to the edges in A . Moreover, for any other block B in G , A and B are vertex-disjoint.*

Fig. 19. A linear hypergraph with $m = n$ and no SDR. The three large triangles denote any block with an SDR, such as the Fano hypergraph

Proof. Let $A = (V_A, E_A)$ be a block in G . Since A is edge-induced, the only vertices which can match to the edges in E_A are those in V_A . Since $|V_A| = |E_A|$ by definition of a block, our first claim follows. The second claim now follows as a corollary, since if A and B intersect on vertex $v \in V$, then *both* A and B require v as part of their respective SDR's by our first claim, which is impossible since each vertex can only be matched to a single edge. $\square$

Theorem 62. *There exists a linear hypergraph with $m = n$ which does not have an SDR.*

Proof. The intuitive idea is to “stitch together” three copies of (say) the Fano hypergraph with an extra edge, so that there is no way to match the new edge to a vertex. This loses the $m = n$ property; we then add further edges and vertices to recover the property. The resulting hypergraph is depicted in Fig. 19. That the graph lacks an SDR follows from Lemma 61; specifically, let e denote the unique edge incident on all three Fano hypergraphs. Since the latter are blocks, none of their vertices can be matched to e , implying e cannot be matched. $\square$

As an aside, note that the construction of Theorem 62 can be modified to also satisfy the Helly property—namely, replace each copy of the Fano hypergraph with a copy of the iCycle, discard the inner two circular edges $\{v_2, v_8, v_5\}$ and $\{v_3, v_4, v_7\}$ and the two “straight-line” edges $\{v_9, v_8, v_7\}$ and $\{v_4, v_5, v_6\}$, and discard vertices v_4, v_5, v_7, v_8 . The resulting hypergraph is linear, Helly, and satisfies $m = n$, but does not have a matching. Thus, the construction of Theorem 62 appears to yield a fairly systematic approach for constructing hypergraphs without SDR's, but still satisfying other desirable properties such as being linear or Helly.

Acknowledgements. The first result of this project was partially completed while NdB was affiliated with the Centrum Wiskunde & Informatica; SG thanks Ronald de Wolf and Centrum Wiskunde & Informatica for their hospitality. SG thanks Howard Barnum and David Reeb regarding discussions on algebraic geometry, and David Gosset for discussions on Quantum SAT. NdB acknowledges support from the EPSRC National Quantum Technology Hub in Networked Quantum Information Processing. SG acknowledges support from NSF Grants CCF-1526189 and CCF-1617710.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

References

1. Kitaev, A., Shen, A., Vyalii, M.: Classical and Quantum Computation. American Mathematical Society, Providence, RI (2002)
2. Aharonov, D., Naveh, T.: Quantum NP—a survey (2002). [arXiv:quant-ph/0210077v1](https://arxiv.org/abs/quant-ph/0210077v1)
3. Osborne, T.J.: Hamiltonian complexity. Rep. Progr. Phys. **75**(2), 022001 (2012)
4. Gharibian, S., Huang, Y., Landau, Z., Shin, S.W.: Quantum Hamiltonian complexity. Found. Trends® Theor. Comput. Sci. **10**(3), 159–282 (2014)
5. Bravyi, S.: Efficient algorithm for a quantum analogue of 2-SAT (2006). [arXiv:quant-ph/0602108v1](https://arxiv.org/abs/quant-ph/0602108v1)
6. Gosset, D., Nagaj, D.: Quantum 3-SAT is QMA1-complete. In: 54th IEEE Symposium on Foundations of Computer Science (FOCS 2013), pp. 756–765 (2013)
7. Ford Jr., L.R., Fulkerson, D.R.: Maximal flow through a network. Can. J. Math. **8**, 399–404 (1956)
8. Laumann, C.R., Läuchli, A.M., Moessner, R., Scardicchio, A., Sondhi, S.L.: On product, generic and random generic quantum satisfiability. Phys. Rev. Lett. **81**, 062345 (2010)
9. Bellare, M., Goldwasser, S.: The complexity of decision versus search. SIAM J. Comput. **23**(1), 97–119 (1994)
10. Jukna, S.: Extremal Combinatorics with Applications in Computer Science, second edn. Springer, Berlin (2011)
11. Laumann, C.R., Moessner, R., Scardicchio, A., Sondhi, S.L.: Phase transitions and random quantum satisfiability. Quantum Inf. Comput. **10**, 1–15 (2010)
12. de Beaudrap, N., Gharibian, S.: A linear time algorithm for quantum 2-SAT. In: 31st Conference on Computational Complexity (CCC 2016), vol. 50, pp. 27:1–27:21 (2016)
13. Schönhage, A.: Quasi-GCD computations. J. Complex. **1**(1), 118–137 (1985)
14. Schönhage, A.: Equation solving in terms of computational complexity. In: International Congress of Mathematicians, pp. 131–153 (1986)
15. Naik, R.N., Rao, S.B., Shrikhande, S.S., Singhi, N.M.: Intersection graphs of k -uniform linear hypergraphs. Eur. J. Comb. **3**(2), 159–172 (1982)
16. Arnborg, S., Corneil, D.G., Proskurowski, A.: Complexity of finding embeddings in a k -tree. SIAM J. Algebr. Discrete Methods **8**(2), 277–284 (1987)
17. Arad, I., Santha, M., Sundaram, A., Zhang, S.: Linear time algorithm for quantum 2SAT. In: 43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016), Volume 55 of Leibniz International Proceedings in Informatics (LIPIcs), pp. 15:1–15:14 (2016)
18. Ambainis, A., Kempe, J., Sattath, O.: A quantum Lovász local lemma. J. ACM **59**(5), 1–24 (2012)
19. Schwarz, M., Cubitt, T.S., Verstraete, F.: An information-theoretic proof of the constructive commutative Quantum Lovász Local Lemma (2013). [arXiv:1311.6474v1](https://arxiv.org/abs/1311.6474v1) [quant-ph]
20. Sattath, O., Arad, I.: A constructive quantum Lovász Local Lemma for commuting projectors. Quantum Inf. Comput. **15**(11&12), 0987–0996 (2015)
21. Gilyén, A.P., Sattath, O.: On preparing ground states of gapped Hamiltonians: an efficient Quantum Lovász Local Lemma. In: 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS), pp. 439–450 (2017)
22. Sattath, O., Morampudi, S.C., Laumann, C.R., Moessner, R.: When a local Hamiltonian must be frustration-free. Natl. Acad. Sci. **113**(23), 6433–6437 (2016)
23. Sun, X., Zhang, J., He, K., Li, Q.: Quantum Lovász Local Lemma: Shearer's bound is tight (2018). [arXiv:1804.07055v4](https://arxiv.org/abs/1804.07055v4) [cs.CC]
24. Downey, R.G., Fellows, M.R.: Parameterized Complexity. Springer, Berlin (2012)
25. Szeider, S.: On Fixed-Parameter Tractable Parameterizations of SAT, pp. 188–202. Springer, Berlin (2004)
26. Fischer, E., Makowsky, J.A., Ravve, E.V.: Counting truth assignments of formulas of bounded tree-width or clique-width. Discrete Appl. Math. **156**(4), 511–529 (2008)

27. Samer, M., Szeider, S.: Algorithms for propositional model counting. *J. Discrete Algorithms* **8**(1), 50–64 (2010)
28. Ganian, R., Hliněný, P., Obdržálek, J.: Better algorithms for satisfiability problems for formulas of bounded rank-width. *Fundam. Inf.* **123**(1), 59–76 (2013)
29. Sæther, S.H., Telle, J.A., Vatschelle, M.: Solving #sat and maxsat by dynamic programming. *J. Artif. Intell. Res.* **54**(1), 59–82 (2015)
30. Paulusma, D., Slivovsky, F., Szeider, S.: Model counting for CNF formulas of bounded modular treewidth. *Algorithmica* **76**(1), 168–194 (2016)
31. Gaspers, S., Papadimitriou, C.H., Sæther, S.H., Telle, J.A.: On satisfiability problems with a linear structure. In: *IPEC* (2016)
32. Markov, I., Shi, Y.: Simulating quantum computation by contracting tensor networks. *SIAM J. Comput.* **38**(3), 963–981 (2008)
33. Peng, C., Chen, X., Tagliacozzo, L., Su, G., Lewenstein, M., Ran, S., Tirrito, E.: *Tensor Network Contractions: Methods and Applications to Quantum Many-Body Systems*. Springer, Berlin (2020)
34. Landau, Z., Vazirani, U., Vidick, T.: A polynomial-time algorithm for the ground state of 1D gapped local Hamiltonians. *Nat. Phys.* **11**, 566–569 (2015)
35. Schuch, N., Wolf, M., Verstraete, F., Cirac, J.I.: The computational complexity of PEPS. *Phys. Rev. Lett.* **98**, 140506 (2007)
36. Katona, G.Y., Szabó, P.G.N.: Bounds on the number of edges in hypertrees. *Discrete Math.* **339**(7), 1884–1891 (2016). (**7th Cracow Conference on Graph Theory, Rytko 2014**)
37. Wolfram, D.A.: Solving generalized Fibonacci recurrences. *Fibonacci Q.* **36**(2), 129–145 (1996)
38. Hliněný, P., Kratochvíl, J.: Computational complexity of the Krausz dimension of graphs. In: *WG* (1997)
39. Berge, C.: *Graphs*. North-Holland Publishing, Amsterdam (1985)
40. Bretto, A., Cherifi, H., Ubéda, S.: An efficient algorithm for Helly property recognition in a linear hypergraph. *Electron. Not. Theor. Comput. Sci.* **46**, 177–187 (2001). (**IWCIA 2001, 8th International Workshop on Combinatorial Image Analysis**)

Communicated by M. M. Wolf