Available online at www.sciencedirect.com
ScienceDirect
Discrete Mathematics 306 (2006) 2481 – 2491
DISCRETE MATHEMATICS
www.elsevier.com/locate/disc

Odd cycles and Θ-cycles in hypergraphs

András Gyárfás[a], Michael S. Jacobson[b], André E. Kézdy[c], Jenő Lehel[d,1]

[a]Automation and Computing Research Institute of the Hungarian Academy of Sciences, 1111 Budapest Kende u. 13–17, Hungary
[b]Department of Mathematics, University of Colorado at Denver, Denver, CO 80217, USA
[c]Department of Mathematics, University of Louisville, Louisville, KY 40292, USA
[d]Department of Mathematical Sciences, The University of Memphis, Memphis, TN 38152, USA

Received 17 February 2003; received in revised form 19 September 2003; accepted 12 December 2005
Available online 8 August 2006

Dedicated to Professor Berge

Abstract

A Θ-cycle of a hypergraph is a cycle including an edge that contains at least three base points of the cycle. We show that if a hypergraph $H=(V,E)$ has no Θ-cycle, and $|e|\geqslant 3$, for every edge $e\in E$, then $\sum_{e\in E}(|e|-1)\leqslant 2|V|-2$ with equality if and only if H is obtained from a hypertree by doubling its edges.

This result reminiscent of Berge's and Lovász's similar inequalities implies that 3-uniform hypergraphs with n vertices and n edges have Θ-cycles, and 3-uniform simple hypergraphs with n vertices and $n-1$ edges have Θ-cycles. Both results are sharp. Since the presence of a Θ-cycle implies the presence of an odd cycle, both results are sharp for odd cycles as well. However, for linear 3-uniform hypergraphs the thresholds are different for Θ-cycles and for odd cycles. Linear 3-uniform hypergraphs with n vertices and with minimum degree two have Θ-cycles when $|E|\geqslant 5n/6-c_1\sqrt{n}$ and have odd cycles when $|E|\geqslant 7n/9-c_2\sqrt{n}$ and these are sharp results apart from the values of the constants.

Most of our proofs use the concept of edge-critical (minimally 2-connected) graphs introduced by Dirac and by Plummer. In fact, the hypergraph results—in disguise—are extremal results for bipartite graphs that have no cycles with chords.

Keywords: Hypergraph; Odd cycles; Extremal problem; Bipartite graph; Minimal blocks

1. Introduction

In this paper we shall prove inequalities similar to those of Berge, Lovász, and others pertaining to substantial generalizations of bipartite graphs. In particular, we address hypergraphs without odd cycles, more generally hypergraphs without Θ-cycles. The proofs are based on the natural bipartite graph representation of hypergraphs. In fact, we prove extremal results for bipartite graphs that have no cycles with chords. Yannakakis [10] also used these graphs to provide a decomposition algorithm for incidence matrices of hypergraphs without odd cycles. The basic properties of arbitrary graphs that do not contain cycles with chords were developed independently in Dirac [5] and in Plummer [8]. The first extremal result was obtained by Pósa [9] several years earlier than the work of Plummer and Dirac: graphs with n

[1] On leave from the Automation and Computing Research Institute of the Hungarian Academy of Sciences.
E-mail address: jlehel@memphis.edu (J. Lehel).

0012-365X/$ - see front matter
doi:10.1016/j.disc.2005.12.037

vertices that contain no cycles with chords have at most $2n-4$ edges, furthermore, for $n \geqslant 4$, the complete bipartite graph $K_{2,n-2}$ is the unique extremal graph.

For $k \geqslant 2$, a k-cycle of a hypergraph is an alternating sequence, $C=(x_1, e_1, x_2, e_2, \dots, x_k, e_k)$, of distinct vertices $x_1, x_2, \dots, x_k$ and distinct edges $e_1, e_2, \dots, e_k$ such that $x_k, x_1 \in e_k$, and $x_i, x_{i+1} \in e_i$, for $i=1,\dots,k-1$. The vertices $x_1, \dots, x_k$ are the *base points* of C. A cycle C is called an odd (even) cycle if k is odd (even).

A connected hypergraph with no cycles at all is called a *hypertree*. The following characterization of hypertrees is from Berge [1]: a connected hypergraph $H=(V,E)$ is a hypertree if and only if $\sum_{e\in E}(|e|-1)=|V|-1$. Hypergraphs without cycles of length at least three are treated in Lovász' inequality (see [7, Exercise 13.2(b)], a related inequality is in [11]): if $H=(V,E)$ is a hypergraph containing no k-cycle with $k \geqslant 3$, and $|e \cap f| \leqslant p$, for any distinct $e, f \in E$, then $\sum_{e\in E}(|e|-p) \leqslant |V|-p$. Our aim is to prove similar inequalities for hypergraphs without odd cycles or Θ-cycles. To define the latter, an edge of a cycle is called a *diagonal edge* if it contains at least three base points of the cycle. A cycle is a *Θ-cycle* if it includes a diagonal edge.

The main interest in introducing Θ-cycles here lies in their role in conjunction with odd cycles. A diagonal edge of a Θ-cycle together with an appropriate piece of the cycle form an odd cycle. Thus the existence of a Θ-cycle implies the existence of an odd cycle (of course, the converse is not true). Note that the counterpart of a Θ-cycle, the notion of a cycle with no diagonal edge, is usually referred to as a special cycle in the hypergraph literature (see [2]). Our first result is the following inequality.

Theorem 1.1. *If $H=(V,E)$ is a hypergraph with no Θ-cycle, and $|e| \geqslant 3$, for each $e \in E$, then*

$$\sum_{e\in E}(|e|-1) \leqslant 2|V|-2.$$

Furthermore, equality holds if and only if H is obtained by doubling each edge of a hypertree.

We have a few remarks concerning the theorem. Since the extremal hypergraphs characterized in Theorem 1.1 have no odd cycles (have only 2-cycles), the inequality remains tight for hypergraphs without odd cycles. In other words, both odd cycles and Θ-cycles emerge in hypergraphs violating the inequality of Theorem 1.1. The next remark concerns the condition $|e| \geqslant 3$. As usual, the *rank* of an edge of H is the number of vertices of the edge and the *lower rank* of H is the minimum rank of its edges. A hypergraph is *r-uniform* if every edge has rank r. Theorem 1.1 is stated for hypergraphs of lower rank at least 3, otherwise it is not true. Indeed, complete bipartite graphs (considered as 2-uniform hypergraphs) show this: they have no odd cycles and clearly violate the inequality whenever both sides have at least four vertices. For 3-uniform hypergraphs Theorem 1.1 has two immediate corollaries.

Corollary 1.2. *A 3-uniform hypergraph with n vertices and with n edges has a Θ-cycle (consequently an odd cycle). If n is even, the same conclusion holds for $n-1$ edges.*

A hypergraph is *simple* if it has no repeated edges. Since the extremal hypergraphs in Theorem 1.1 are not simple we obtain

Corollary 1.3. *A simple 3-uniform hypergraph with n vertices and with $n-1$ edges has a Θ-cycle (consequently an odd cycle).*

Both corollaries are obviously sharp. Corollary 1.3 generalizes Exercise 13.4 in Lovász [7] stating that the hypothesis implies the existence of a cycle of length at least three.

The linear bound concerning odd cycle-free hypergraphs in Theorem 1.1 cannot be expected to carry over to natural extensions. The most widely investigated extension of odd cycle-free hypergraphs is the family of unimodular hypergraphs (see [2]). A simple example of a 3-uniform unimodular hypergraph with n^2 edges and $2n+1$ vertices is obtained from the complete bipartite graph $K_{n,n}$ by extending all its edges with a (common) new vertex x. Forbidding odd cycles of fixed length or small length does not give linear upper bound either for the number of edges, because there are 3-uniform hypergraphs with n vertices and $n^{1+\varepsilon(t)}$ edges that contain no cycles of length less than $t \geqslant 2$. (This follows by using standard probabilistic arguments, see e.g. in [4].) The only sharp result known to us along this line is proved by Győri [6]: the maximum number of edges in a triangle-free 3-uniform hypergraph of order n is at most $n^2/8$ plus a small constant.

Our next aim is to sharpen Corollary 1.2 for *linear* hypergraphs, i.e. for hypergraphs without 2-cycles. The following example shows that odd cycle-free 3-uniform linear hypergraphs with n vertices can have almost n edges. Extend each edge of a balanced complete bipartite graph with a (separate) new vertex. However, in this hypergraph almost all vertices are of degree one. Imposing the condition of minimum degree two, Corollary 1.2 can be improved significantly. Moreover, we get different results for odd cycle-free and Θ-cycle-free hypergraphs.

Theorem 1.4. *Let H be a 3-uniform linear hypergraph of order n with minimum degree at least two. If H has no Θ-cycle, then it has at most $5n/6 - c_1\sqrt{n}$ edges, where c_1 is a positive constant.*

Theorem 1.5. *Let H be a 3-uniform linear hypergraph of order n with minimum degree at least two . If H has no odd cycle, then it has at most $7n/9 - c_2\sqrt{n}$ edges, where c_2 is a positive constant.*

Examples will show that these theorems are asymptotically sharp. The hypergraph results in Theorems 1.1, 1.4, and 1.5 will be formulated and proved in subsequent sections through the natural bigraph representation of hypergraphs.

The *bigraph representation* $G[A, B]$ of a hypergraph $H = (V, E)$ is the bipartite graph (*bigraph*) with partite sets $A = E$, $B = V$, and edges in $G[A, B]$ indicating the edge/vertex incidences in H. We refer to $G[A, B]$ as the bigraph of H. A k-cycle of H corresponds to a $2k$-cycle in its bigraph. Thus, H has no odd cycles if and only if its bigraph has no cycles of length congruent to 2 mod 4 (we shall call those bigraphs 0-bigraphs). A Θ-cycle of H corresponds to a cycle with a chord in its bigraph. Therefore, H contains no Θ-cycles if and only if its bigraph has no cycles with a chord. A hypergraph is linear if and only if its bigraph is C_4-free (has no cycles of length four).

To illustrate the use of bigraphs note that a hypergraph $H = (V, E)$ is a hypertree if and only if its bigraph is a tree. This observation results in the characterization of hypertrees in terms of the identity $\sum_{e\in E}(|e| - 1) = |V| - 1$ mentioned above. For another example, note that Pósa's extremal result implies immediately the following inequality similar to that of Lovász: if $H = (V, E)$ has no Θ-cycles, then $\sum_{e\in E}(|e| - 1) \leqslant 2|V| + |E| - 4$. Actually, our main result is a stronger version of this inequality together with a few variants restricted to various hypergraphs of rank 3.

The bigraph representation of Θ-free hypergraphs requires the study of *bipartite* edge-critical graphs. Their fundamental properties will be summarized in Section 2. In Section 3 bigraph analogs of Theorems 1.1, 1.4, and 1.5 will be proved, first for bipartite blocks. Then these auxiliary results will be used in Section 4 for deriving the full bigraph versions of these theorems (formulated as Theorems 4.1 and 4.2).

The auxiliary results, Theorems 3.2 and 3.3, on bipartite edge-critical blocks have some independent interest. To translate them to hypergraphs it is convenient to introduce a nonstandard terminology: a hypergraph will be called *2-connected* if its bigraph is 2-connected (i.e. a block). Notice that 2-connectivity of H means two conditions: the removal of any of its edges does not disconnect H, and that the removal of any vertex v together with replacing each edge e with $e\setminus\{v\}$ also does not disconnect H. Using this terminology, we state the 2-connected versions of Theorems 1.4 and 1.5 from Section 3 in abridged form.

Theorem 1.6. *If H is a 2-connected 3-uniform linear hypergraph of order n with no Θ-cycles, then H has at most $\frac{4}{5}(n - 1)$ edges.*

Theorem 1.7. *If H is a 2-connected 3-uniform linear hypergraph of order n with no odd cycles, then H has at most $\frac{3}{4}(n - 1)$ edges.*

It is worth noting that both results have unique extremal hypergraphs: the duals of the subdivisions of K_4 and $K_{3,3}$ (defined in the next section).

Our inductional proofs are based on conceptually rather simple reduction processes on bigraphs. However, the inductions are loaded with technical generalizations, and they contain involved details, especially in Section 3. Therefore, in the concluding Section 5 we restate Corollary 1.2 as Theorem 5.1, and we give a quick proof using a quite different approach.

2. Rings, blocks, bipolar graphs

We consider graphs without loops and multiple edges. A block is a graph with no cut vertex. A single edge is called a trivial block. A nontrivial block G is called *edge-critical* if the removal of any edge from the edge set results in a graph with a cut vertex. A graph is edge-critical, provided that each of its blocks is edge-critical. For convenience, a single edge is considered edge-critical, too. In this section we introduce our terminology and summarize the basic properties of edge-critical graphs and bigraphs for further reference. Properties (1)–(3) were found and used by Dirac [5] and by Plummer [8] (they are also in [3, Chapter 1.3]).

(1) A graph is edge-critical if and only if it has no cycle with a chord.

Definition of a chain of blocks: A connected graph with at least one cut vertex is called a chain if each of its cut vertices belongs to exactly two blocks, and each of its blocks contains at most two cut vertices.

(2) If G is a nontrivial edge-critical bloc, then $G - e$ is a chain of edge-critical blocks, for every edge e.

Let G be a nontrivial edge-critical block, and let $e = xy$ be an arbitrary edge of G. Vertices x, y, and the cut vertices of the chain $G - e$ are called *entry points* for the blocks. By (2), every block of the chain contains exactly two entry points.

Definition of a ring: For a nontrivial edge-critical block G, and for an edge e of G, the ring $G(e)$ is the chain $G - e$ completed with the trivial block e. Notice that $G(e) = G$ but the notation $G(e)$ emphasizes the structure of G that depends on e.

(3) The entry points of a nontrivial block of $G - e$ are not adjacent. Moreover, no path between any two entry points of $G - e$ has a chord.

Property (3) motivates the introduction of the notion of bipolar graphs. From now on we are concerned with bigraphs. A *bigraph* $G[A, B]$ is a simple bipartite graph with partite sets A and B. Let $G = G[A, B]$ be a bigraph such that no cycle has a chord. Let u and v be nonadjacent vertices of G such that no u, v-path has a chord. We call this pair of vertices *admissible entries* of G, and the triplet (u, G, v) is called a *bipolar graph*. A bipolar graph $(u, G[A, B], v)$ has type *AA* if $u, v \in A$, it has type *BB* if $u, v \in B$, and it is of type *AB* otherwise.

Examples. The only pair of admissible entries of $G[A, B] \cong K_{2,r}$ $(r \geqslant 3)$ is the smaller partite set A, thus $K_{2,r}$ is a bipolar block of type *AA*. By (3), if G has no cut vertex, and $e = xy$ is any edge of G, then $(x, G - e, y)$ is a bipolar graph of type *AB*.

For easier reference, the term edge-critical block will be used for bipartite edge-critical graphs without cut vertices. Edge-critical blocks of type *AA*, type *AB*, type *BB* refer to edge-critical bipolar blocks of the corresponding type.

Definition of 0*-graphs*: A bipartite graph such that the length of each cycle is divisible by four is clearly edge-critical, it will be called an edge-critical 0-graph. Similarly, an edge-critical 0-block is a bipartite block which is an edge-critical 0-graph. A bipartite ring which is an edge-critical 0-block is also called a 0-ring. We need two important properties of edge-critical 0-blocks.

(4) If (u, H, v) is a subchain in a 0-ring, then all uv-paths of H have the same length(mod 4).

The vertices u, v can be connected by a path P on the ring so that internal vertices of P are not in H. If H has two uv-paths of different length (mod 4) then the union of one of them with P results in a cycle of length 2(mod 4). Hence (4) follows.

(5) If (u, H, v) is a nontrivial edge-critical block of a 0-ring, then it is not of type AB.

Assume that (u, H, v) is an edge-critical block of type *AB*. The entry points can be connected with three internally vertex disjoint paths, two within H and a third along the ring. Two of the three paths have the same (even) number of internal vertices (mod 4), so their union is a cycle of length 2(mod 4).

Definition of subdivision and reverse subdivision graphs: An edge-critical bigraph, $\mathrm{SD}(H)$ can be obtained from any graph H (that can have multiple edges) as follows. Each edge of H is subdivided by a single new vertex. The resulting bipartite graph is $\mathrm{SD}(H)$, the new vertices give partite set B and the vertices of H form partite set A. The graph $\mathrm{SD}(H)$ is called the *subdivision graph* of H. Exchanging the role of A and B, we get the *reverse subdivision graph*, $\mathrm{RSD}(H)$, in which every vertex in A has degree two. Note that if H is 2-connected, then so are $\mathrm{SD}(H)$ and $\mathrm{RSD}(H)$. These constructions play an important role in our subsequent extremal results.

3. Edge-critical blocks

Let d_G be the usual degree function in graph G. For a bigraph $G = G[A, B]$, define $A(3) = \{a \in A : d_G(a) \geqslant 3\}$.

Theorem 3.1. *If $G = G[A, B]$ is a nontrivial edge-critical block, then*

$$\sum_{a\in A(3)} (d_G(a) - 1) \leqslant 2(|B| - 1),$$

with equality if and only if $G \cong K_{2,r}$, for some $r \geqslant 3$ (i.e. G is the subdivision graph of an edge of multiplicity at least three).

Theorem 3.2. *If $G = G[A, B]$ is a C_4-free edge-critical block and $d_G(a) \leqslant 3$, for each $a \in A$, then G is either the subdivision graph of a 3-regular graph with 4 or 6 vertices or*

$$|A(3)| \leqslant \tfrac{4}{5}(|B| - 2).$$

Theorem 3.3. *If $G = G[A, B]$ is a C_4-free edge-critical 0-block and $d_G(a) \leqslant 3$, for each $a \in A$, then G is either the subdivision graph of a 3-regular bipartite graph with 6, 8, or 10 vertices or*

$$|A(3)| \leqslant \tfrac{3}{4}(|B| - 2).$$

The three extremal results pertaining to edge-critical bipartite blocks will be proved simultaneously.

3.1. Reverse subdivision blocks

The theorems are trivial for reverse subdivision blocks. We will use the following properties of reverse subdivision blocks later. In general, $|A|$ cannot be bounded in terms of $|B|$. This is shown by the subdivision of an edge of multiplicity m (where $|B| = 2$ and $|A| = m$ is arbitrarily large). Clearly, C_4-free reverse subdivision blocks are obtained by subdividing simple graphs. Thus, we get the trivial bound $|A| \leqslant \binom{|B|}{2}$ (because for fixed $|B|$ the complete simple graph has the maximum number of edges). Finally, for C_4-free edge-critical 0-blocks, the subdivided graphs must be simple and bipartite, hence we get easily the bound $|A| \leqslant |B|^2/4$.

3.2. Subdivision blocks

If $G = G[A, B]$ is a subdivision block with t vertices of degree two in A, then we have

$$\sum_{a\in A(3)} (d_G(a) - 1) = |E(G)| - |A(3)| - 2t$$

$$= 2|B| - |A(3)| - 2t = 2|B| - |A| - t \leqslant 2|B| - 2$$

with equality if and only if $t = 0$ and $|A| = 2$. This proves Theorem 3.1 for subdivision graphs.

We need the following property discussed earlier as an example.

Proposition 3.4. *The subdivision graph $K_{2,r}$ $(r \geqslant 3)$ is not of type BB.*

In Theorems 3.2 and 3.3 2-connectivity and the condition $d_G(a) \leqslant 3$, for $a \in A$, imply

$$3|A(3)| + 2t = 2|B|.$$

Thus the inequalities reduce to $0 \leqslant |B| + 5t - 12$ and $0 \leqslant |B| + 8t - 18$, respectively.

These inequalities are obviously valid except for $t = 0$, 1, or 2. Note that G has no C_4, thus G is the subdivision of a simple graph. Moreover, in case of Theorem 3.3, G is the subdivision of a bipartite graph. Using these observations we can easily screen out those six exceptional edge-critical blocks and edge-critical 0-blocks that satisfy inequalities slightly weaker than the ones in Theorems 3.2 and 3.3. We summarize the results omitting the straightforward case analysis.

Exceptional blocks: Either SD(K_4), that satisfies $|A(3)| = \frac{4}{5}(|B| - 1)$, or the subdivision graph of a cubic graph with six vertices (there are two: $K_{3,3}$ and the prism graph $K_2 \times K_3$), for which we have $|A(3)| = \frac{4}{5}(|B| - \frac{3}{2})$.

Exceptional 0-*blocks*: The subdivision graph of a bipartite cubic graph with 10, 8, or 6 vertices. For these graphs we obtain $|A(3)|=\frac{3}{4}(|B|-\frac{5}{3})$, $|A(3)|=\frac{3}{4}(|B|-\frac{4}{3})$, or $|A(3)|=\frac{3}{4}(|B|-1)$, respectively. Note that the unique exceptional edge-critical 0-block with six vertices is $\mathrm{SD}(K_{3,3})$.

The basic properties of exceptional blocks we need later are formulated in the next two propositions.

Proposition 3.5. *If $G=G[A,B]$ is an exceptional edge-critical block different from $\mathrm{SD}(K_4)$ or an exceptional edge-critical* 0-*block different from $\mathrm{SD}(K_{3,3})$, then*

$$|A(3)|+1\leqslant \lambda|B|,$$

where $\lambda=\frac{4}{5}$ or $\frac{3}{4}$.

Proposition 3.6. *Exceptional edge-critical blocks or* 0-*blocks are not of type BB or AB.*

Proof. Let H be one of the six cubic graphs whose subdivision graph is an exceptional edge-critical block or 0-block. It is easy to check that H has the following property: for any pair of distinct edges e and f, e is a chord in a cycle containing f. Therefore, if b_e and b_f are the vertices subdividing e and f, respectively, then $\mathrm{SD}(H)$ has a path with a chord from b_e to b_f. Thus H is not of type *BB*. A similar argument shows that H cannot be of type *AB*. □

3.3. Removal of a suspended BB-path

A path in which every internal vertex has degree two is called a *suspended path*. A suspended path of a bigraph with both endpoints in B is called a suspended *BB-path*. Let G be an edge-critical block containing a suspended *BB*-path, and assume that the removal of its internal vertices and incident edges from G leaves a graph that is still an edge-critical block. The end vertices of the removed *BB*-path form an admissible pair, i.e. the block we obtain can have type *BB*.

The removal of suspended *BB*-paths will be used in the inductional proof of Theorems 3.1–3.3. Note that the operation does not change $A(3)$, thus induction might work even if the number of B-vertices does not reduce. For Theorem 3.1, the resulting edge-critical *BB*-block is not extremal, by Proposition 3.4, thus induction yields strict inequality. In case of Theorems 3.2 and 3.3 the resulting edge-critical *BB*-block is not exceptional, by Proposition 3.6, so the induction hypothesis is applicable. Therefore, in further reduction steps we shall assume that G has no removable suspended *BB*-paths. One useful consequence of this assumption is formulated in the next proposition.

Proposition 3.7. *If (b,H,v) is an edge-critical block without removable suspended BB-paths, and $b\in B$, then b is not on any* 4-*cycle of H.*

Proof. Assume that $C=(b,a_1,b_1,a_2)$ is a 4-cycle in H. If $b_1=v$, define $H^*=C$. Otherwise, consider two vertex disjoint paths from v to C and define H^* as the union of C and these paths. The only case that does not violate property (3) is when the paths run to b and to b_1. Because there is no removable suspended *BB*-path in H, $d_G(a_1)\geqslant 3$ follows. Let b_1' be a third neighbor of a_1 (different from b and b_1). Since a_1 is not a cut vertex in H, there is a shortest path in $H-a_1$ from b_1' to H^*-a_1. This contradicts property (3) of H. □

3.4. Ring reductions

Assume that G is an edge-critical block without removable suspended *BB*-paths and G is not a subdivision block. Then we can select an edge $e=xb$ of G with $b\in B$ and with $d_G(b)\geqslant 3$. Let (b,H_1,v) be the nontrivial block of the ring $G(e)$ containing b, and let $C=(v,H_2,x)$ be the subchain complementing H_1 in the ring.

3.4.1. Reduction of (b,H_1,v) (the conclusion of Theorem 3.1)

Let $a_1,\dots,a_k\in V(H_1)$, $k\geqslant 2$, be the neighbors of b in H_1. Identify these k vertices with a single new vertex a, and replace the resulting multiple edge between b and a with a single edge. Proposition 3.7 ensures that this reduction gives a simple graph G'.

Clearly, $d_{G'}(a)\geqslant 3$, and $d_{G'}(a)-1=\sum_{i=1}^{k}(d_G(a_i)-1)$. Because G' is an edge-critical block smaller than G, and it is not extremal, Theorem 3.1 follows easily, by induction.

Note that the operation used here may create a C_4, therefore it applies only for Theorem 3.1. We will need further reductions to conclude the proof of Theorems 3.2 and 3.3.

3.4.2. Reduction of (b, H_1, v), for $v \in B$

Observe that the subchain $C = (v, H_2, x)$ must contain a nontrivial block, otherwise C extended with b would be a removable suspended BB-path of G. When C is the union of an edge va and a block (a, H, x) of type AA, then C is called a *special chain* and G is called a *special ring* during the reduction. Thus, a special ring is the disjoint union of a block (b, H_1, v) of type BB and a block (a, H, x) of type AA with two additional edges, va and xb, between them.

Set $G_1 = G_1[A_1, B_1] = (b, H_1, v)$ and let $G_2 = G_2[A_2, B_2]$ be defined as the chain $C = (v, H_2, x)$ extended to a ring by adding a v, x-path P of length one or three. The choice of the length of P is specified as follows.

In Theorem 3.2, P is defined to have length three. Thus, G_2 is a C_4-free edge-critical block, just like G.

In case of Theorem 3.3, G is a C_4-free edge-critical 0-block, and so is C. Thus, by (4), all v, x-paths of C have equal (odd) length (mod 4). Hence, one can choose the length of P appropriately (one or three) to make G_2 a 2-connected 0-graph. Note that C is different from a trivial block. Moreover, because C is of type AB in the 0-ring $G(e)$, it cannot be a single nontrivial block, by (5). Therefore, G_2 is a simple edge-critical 0-bigraph, i.e. an edge-critical 0-block. Furthermore, it is also C_4-free, except when $P = (xv)$, G is a special ring, and the distance between the entries a and x in the block (a, H, x) is equal to two. The reduction of this case shall be done separately in 3.4.3, here we assume that G_2 is C_4-free.

Note that the order of G_1 and G_2 is smaller than that of G. This is obvious for G_1, and easily follows for G_2. Indeed, G_2 has at most $|V(G)| - |V(H_1)| + 3$ vertices, and $|V(H_1)| > 3$, because in any edge-critical block of type BB there are at least four vertices. Note that $|A(3)| = |A_1(3)| + |A_2(3)|$ and $|B_1| + |B_2| \leqslant |B| + 2$. Thus, we may apply induction, provided G_1 and G_2 are not exceptional, to get

$$|A(3)| = |A_1(3)| + |A_2(3)| \leqslant \lambda(|B_1| - 2) + \lambda(|B_2| - 2) \leqslant \lambda(|B| - 2),$$

where $\lambda = \frac{4}{5}$ or $\frac{3}{4}$, as required in Theorems 3.2 and 3.3, respectively.

Because G_1 is an edge-critical block of type BB, it is not exceptional, by Proposition 3.6. In case of Theorem 3.3, G_2 cannot be an exceptional C_4-free edge-critical block, because P has length three containing a vertex of A with degree two. We conclude that G_2 can be an exceptional C_4-free edge-critical 0-block only in the proof of Theorem 3.3.

By the definition of exceptional C_4-free edge-critical 0-blocks, we have $P = (xv)$, G must be a special ring, so that G_2 is the union of an edge-critical 0-block (x, H, a) and the path (x, v, a). In this case $|B_1| + |B_2| = |B| + 1$ and $|A_2(3)| \leqslant \frac{3}{4}(|B_2| - 1)$. Thus we obtain, by induction,

$$\begin{aligned} |A(3)| &= |A_1(3)| + |A_2(3)| \leqslant \tfrac{3}{4}(|B_1| - 2) + \tfrac{3}{4}(|B_2| - 1) \\ &= \tfrac{3}{4}(|B_1| + |B_2| - 3) = \tfrac{3}{4}(|B| - 2). \end{aligned}$$

3.4.3. Reduction of the special ring (the conclusion of Theorem 3.3)

Recall that in Theorem 3.3 $G(e)$ is a C_4-free 0-ring. We use the notation of step 3.4.2 above, thus (a, H, x) is an edge-critical block of type AA with a path (a, b^*, x) between its entry points. This implies that $d_G(b^*) = 2$. Let b_1 denote the neighbor of a different from b^* and let A_1 denote the set of neighbors of b_1 different from a. Similarly, let b_2 denote the neighbor of x different from b^* and let A_2 denote the set of neighbors of b_2 different from x. Since H is an edge-critical 0-block, $A_1 \cap A_2$ is empty. Define $G^* = [A^*, B^*]$ as the bigraph obtained from G by removing the vertices a, x, b^*, b_1, b_2 and by adding all edges from b to A_2 and all edges from v to A_1.

It is easy to see that G^* is a C_4-free edge-critical 0-block, and it is not exceptional. Using that $|B| = |B^*| + 3$ and $|A(3)| = |A^*(3)| + 2$, induction gives

$$|A(3)| = |A^*(3)| + 2 \leqslant \tfrac{3}{4}(|B^*| - 2) + 2 = \tfrac{3}{4}(|B| - 5) + 2 = \tfrac{3}{4}(|B| - 2) - \tfrac{1}{4}.$$

This concludes the proof of Theorem 3.3. If $G(e)$ is not a 0-ring, then the block (b, H_1, v) is not always of type BB. In the last step we discuss this case.

3.4.4. Reduction of (b, H_1, v)*, for* $v \in A$ *(the conclusion of Theorem 3.2)*

Let b_1, b_2 denote the neighbors of v in H_1. We try to apply the following reduction: remove v from G, and identify its three neighbors. Because the number of vertices in A and in B decreases by 1 and 2, respectively, we are done by induction, provided the reduction does not create a C_4. Otherwise, we modify the reduction as follows.

Case 3.4.4.1: There is a cycle $C_6 = (v, b_1, a_1, b^*, a_2, b_2, v)$ in H_1. If $b = b^*$, then either $H_1 = C_6$ or one of b_1 or b_2 has degree at least three. In the second case assume that $d_G(b_1) \geqslant 3$ and let $e_1 = vb_1$. It is easy to see that in the ring $G(e_1)$ the block containing b_1 is of type BB, because it has entries b_1 and b. Then the reduction 3.4.2 or 3.4.3 works. Thus we may assume that $H_1 = C_6$. In this case either $C = (v, H_2, x)$ is just a path of length two when the theorem is trivial, or we can replace H_1 by a single edge bv and this reduction is good for the induction.

If $b \neq b^*$, there exist two paths from b to C_6 sharing only one endpoint b. Since $d_{H_1}(v) = 2$ these paths must end in a_1 and a_2, respectively. Furthermore, both paths are of length at least three, and $d_G(b^*) = 2$. Now remove v and b^* from G and identify b_1, b_2 and the third neighbor of v. Thus, we obtain a C_4-free edge-critical block that is not exceptional. Because the number of vertices in A and B reduces by 1 and 3, respectively, the inequality of Theorem 3.2 follows by induction.

Case 3.4.4.2: There is no 6-cycle in H_1 containing v, but the ring $G(e)$ has a 6-cycle through v. In this case $C = (v, H_2, x)$ is a path of length two (C is a subchain of two trivial blocks), and H_1 contains a path vb_1a_1b. Since both b_1 and a_1 are of degree two, we can remove the vertices v, b_1, a_1 and identify b_2 with the third neighbor of v. We obtain a C_4-free edge-critical block that is not exceptional. Because the number of vertices in A and in B reduces each by 2, we are done by induction. This concludes the proof of Theorem 3.2.

4. Edge-critical bigraphs

In this section we deal with arbitrary edge-critical bigraphs, and we extend the results on edge-critical blocks obtained in Section 3. In particular, Theorem 4.1 below generalizes Theorem 3.1, and it also implies the hypergraph version stated as Theorem 1.1. Theorem 4.2 below shows that the bounds in Theorems 3.2 and 3.3 change significantly when extended from edge-critical blocks to arbitrary edge-critical bigraphs. The proof of Theorem 4.2 also yields the verification of the hypergraph versions stated as Theorems 1.4 and 1.5.

Theorem 4.1. *If* $G = G[A, B]$ *is an edge-critical bigraph of order at least three, then*

$$\sum_{a \in A(3)} (d_G(a) - 1) \leqslant 2(|B| - 1),$$

with equality if and only if every block of G is the subdivision graph of an edge of multiplicity at least three and each cut vertex of G is in B.

Proof. If G is a (nontrivial) block, then we are done by Theorem 3.1. Assume that G is the union of two bigraphs $G_1 = [A_1, B_1]$ and $G_2 = [A_2, B_2]$ with one common vertex v. For $v \in B_1 \cap B_2$, we obtain, by induction,

$$\sum_{a \in A(3)} (d_G(a) - 1) = \sum_{a \in A_1(3)} (d_{G_1}(a) - 1) + \sum_{a \in A_2(3)} (d_{G_2}(a) - 1)$$
$$\leqslant 2(|B_1| - 1) + 2(|B_2| - 1) = 2(|B| - 1).$$

Furthermore, there is equality if and only if the pieces G_1 and G_2 satisfy equality.

Assume now that $v \in A_1 \cap A_2$. If v has degree three or more in both graphs G_1 and G_2, then we obtain, by induction,

$$\sum_{a \in A(3)} (d_G(a) - 1) = 1 + \sum_{a \in A_1(3)} (d_{G_1}(a) - 1) + \sum_{a \in A_2(3)} (d_{G_2}(a) - 1)$$
$$\leqslant 1 + 2(|B_1| - 1) + 2(|B_2| - 1) < 2(|B| - 1).$$

In case $v \in A(3)$, but v has degree less than three in one or in both of G_1 and G_2, then one or both of the inequalities in the induction hypothesis is strict. Thus strict inequality follows for G as above. □

Theorem 4.2. *If $G=G[A,B]$ is a connected C_4-free edge-critical bigraph satisfying $2\leqslant d_G(a)\leqslant 3$ and $2\leqslant d_G(b)$, for each $a\in A$ and $b\in B$, then*

$$|A(3)|\leqslant \tfrac{5}{6}|B|-c_1\sqrt{|B|}. \qquad (*)$$

If, in addition, G is also a 0-bigraph, then

$$|A(3)|\leqslant \tfrac{7}{9}|B|-c_2\sqrt{|B|}. \qquad (**)$$

(*In the inequalities c_1 and c_2 are absolute constants.*)

Proof. We prove the two parts of the theorem simultaneously using a common procedure. If G is a block, then (*) or (**) is true, by Theorem 3.2 or 3.3. Assuming that $G_0=G$ is not a block, let E_1 be an endblock of G_0, and let v_1 be the cut vertex of G_0 contained by E_1. We set $w_1=v_1$ if $v_1\in B$ is incident with at least two edges not in E_1. Otherwise, let $P_1=(v_1,\ldots,w_1)$ be the longest suspended path leaving E_1 (i.e. P_1 is the unique subchain of trivial blocks between cut vertices v_1 and w_1). Let H_1 be the union of E_1 and P_1.

Note that no vertex of $(P_1\backslash\{v_1,w_1\})\cap A$ contributes to the left-hand side of the inequality (*) or (**). If v_1 or w_1 is in A, then it is also in $A(3)$, thus contributing to the left-hand side of the inequalities. Remove from G_0 all edges and all vertices of H_1 different from w_1. The graph G_1 we obtain in this way satisfies the conditions of Theorem 4.2. If G_1 is not a block, then we shall continue the procedure by repeating the previous step on G_1.

In the ith step, E_i is an endblock of G_{i-1} containing the cut vertex v_i. The graph H_i is defined as the union of E_i and the longest suspended path $P_i=(v_i,\ldots,w_i)$ leaving E_i (with $w_i=v_i$ if and only if $v_i\in B$ is incident with at least two edges not in E_i). Furthermore, G_i is the graph that remains after removing H_i (but not vertex w_i) from G_{i-1}. The procedure ends when G_i is a block. If G_k is the last block in the procedure, then we set $E_{k+1}=G_k$, and we define $v_{k+1}=w_{k+1}$ to be any vertex of $E_{k+1}\cap A$, so that $H_{k+1}=E_{k+1}$.

For $i=1,\ldots,k+1$, set $n_i=|(H_i\backslash\{w_i\})\cap B|$, and define

$$A_i(3)=(\{w_i\}\cup\{x\in H_i\mid d_{H_i}(x)=3\})\cap A.$$

Note that $n=\sum_{i=1}^{k+1}n_i=|B|$, and $A(3)=\bigcup_{i=1}^{k+1}A_i(3)$. By the degree condition, each E_i is a nontrivial block. Thus Proposition 3.5 and Theorem 3.2 or 3.3 can be applied to E_i. Case analysis shows that the contribution of H_i to $|A(3)|$ satisfies $|A_i(3)|\leqslant\frac{4}{5}n_i$ or $|A_i(3)|\leqslant\frac{3}{4}n_i$, unless $P_i=v_iw_i$ with $v_i\in B$, $w_i\in A$, and E_i is isomorphic to $\mathrm{SD}(K_4)$ or $\mathrm{SD}(K_{3,3})$. Call a step exceptional if this situation happens.

In case of an exceptional step we obtain that $|A_i(3)|=\frac{4}{5}n_i+\frac{1}{5}$ or $|A_i(3)|=\frac{3}{4}n_i+\frac{1}{4}$. Therefore, $|A(3)|\leqslant(4n/5)+(t/5)$ or $|A(3)|\leqslant(3n/4)+(t/4)$, where t is the number of exceptional steps made in the reduction procedure. We shall prove that $t\leqslant(n/6)-c_1'\sqrt{n}$ or $t\leqslant(n/9)-c_2'\sqrt{n}$, where c_1' and c_2' are absolute constants. From these inequalities Theorem 4.2 follows.

Consider the partition of the edge set of G into the subgraphs H_i defined by the reduction procedure. Color red the edges of all exceptional H_i (i.e. belonging to an exceptional step), and color blue the remaining edges of G. Let G_j, $j=1,\ldots,p$, be the connected components of the subgraph of all blue edges. It is easy to verify that each G_j has at most one vertex in B incident with some red H_i. Call this vertex exceptional (if it exists at all), and let b_j be the number of vertices of $G_j\cap B$ different from the exceptional vertex. Note that $b_j>1$, for each $1\leqslant j\leqslant p$, and $n=6t+\sum_{j=1}^{p}b_j$ or $n=9t+\sum_{j=1}^{p}b_j$.

Because w_i belongs to some G_j, for every red H_i, and these vertices are t distinct vertices in $A(3)$, we obtain easily that

$$t\leqslant\sum_{j=1}^{p}\binom{b_j+1}{2}\leqslant\binom{\sum_{j=1}^{p}b_j}{2}=\binom{n-6t}{2}$$

or, by using Turán's theorem for each G_j, $1\leqslant j\leqslant p$,

$$t\leqslant\sum_{j=1}^{p}\frac{(b_j+1)^2}{4}\leqslant\frac{\left(\sum_{j=1}^{p}b_j\right)^2}{4}=\frac{(n-9t)^2}{4}.$$

From these inequalities it follows that $t \leqslant (n/6) - c_1'\sqrt{n}$ or $t \leqslant (n/9) - c_2'\sqrt{n}$, with absolute constant c_1' or c_2'. This concludes the proof of Theorem 4.2. □

The following constructions show that Theorem 4.2 is tight. Let $t \geqslant 3$ be an arbitrary integer. Let $G[A, B]$ be the disjoint union of the reverse subdivision graph $\mathrm{RSD}(K_t)$ and $\binom{t}{2}$ copies of the exceptional edge-critical block $\mathrm{SD}(K_4)$. One B-vertex of each copy of $\mathrm{SD}(K_4)$ is joined to an A-vertex of $\mathrm{RSD}(K_t)$ so that distinct B-vertices are joined to distinct A-vertices. Then $|B| = t + 6\binom{t}{2}$, and G clearly satisfies $|A(3)| = |A| = 5\binom{t}{2} \geqslant \frac{5}{6}|B| - c_1'\sqrt{|B|}$, where c_1' is a constant independent of t.

Let $G^0[A, B]$ be the union of the reverse subdivision graph $\mathrm{RSD}(K_{t,t})$ and t^2 disjoint copies of the exceptional edge-critical block $\mathrm{SD}(K_{3,3})$. One B-vertex of each copy of $\mathrm{SD}(K_{3,3})$ is joined to an A-vertex of $\mathrm{RSD}(K_{t,t})$ so that distinct B-vertices are joined to distinct A-vertices. Then $|B| = 2t + 9t^2$, and G^0 satisfies $|A(3)| = |A| = 7t^2 \geqslant \frac{7}{9}|B| - c_2'\sqrt{|B|}$, where c_2' is a constant independent of t.

5. A quick proof

The inequality in Theorem 1.1 has the immediate corollary that a 3-uniform hypergraph with n vertices and with n edges has Θ-cycles (consequently odd cycles). Because this particular case might have some independent interest, and the proof of Theorem 1.1 via Theorems 3.1 and 4.1 is quite technical, we include here a quick proof of this result based on Hall's matching theorem.

Theorem 5.1. *If H is a 3-uniform hypergraph with at least as many edges as vertices, then H has a Θ-cycle.*

Proof. Assume that H has m edges and n vertices, $m \geqslant n \geqslant 3$. Let $G = G[A, B]$ be the corresponding bigraph, $|A| = m$ and $|B| = n$. We show that G has a cycle containing a chord. The proof is by induction on n. The claim is true for $n = 3$, because $G \cong K_{m,3}$. W.l.o.g. we may assume that G is *balanced*, that is, $m = n$, and G is an $n \times n$ bigraph with $2n$ vertices and $3n$ edges. Let $n \geqslant 4$, and assume that the claim is true for any $m' \times n'$ bigraph with $m' \geqslant n'$ and $n' < n$.

Step 1: Suppose that there is a set $A' \subset A$ such that if B' is the set of all neighbors of the vertices in A', then $|A'| > |B'|$. Because $G'[A', B']$ satisfies the properties of the theorem and $|A'| < n$, the claim follows by induction.

Step 2: We may assume that $|A'| \leqslant |B'|$, for every $A' \subseteq A$ (where B' is the set of all neighbors of the vertices in A'). By Hall's theorem, G has a perfect matching $M = \{g_1, g_2, \ldots, g_n\}$. Because $G - M$ has $2n$ vertices and $3n - n = 2n$ edges, $G - M$ has a cycle $C = (y_1, x_1, \ldots, y_k, x_k)$, where $y_i \in B$, $x_i \in A$, for $i = 1, \ldots, k$.

Step 3: Let $Y = \{y_1, \ldots, y_k\}$, and consider a maximal alternating forest F such that all component trees are rooted at Y and each path starting from a root is alternating with respect to M (that is, any such path starts with an edge in M and it is alternately taking on edges not in M and in M). Note that, by definition, each connected component of $F = F[A', B']$ is a balanced tree, thus $|A'| = |B'|$. Moreover, all neighbors of each vertex of $F \cap A$ belong to F. Thus, if F does not span G, then $|B'| < n$, and the claim follows by induction.

Step 4: We may assume that F spans G. In particular, every vertex x_i $(1 \leqslant i \leqslant k)$ is joined to some root $y_j \in Y$ by an alternating path of F we call F-path. Choose a pair x_i, y_j joined by an F-path P such that the subpath of C between them, (x_i, C, y_j), is shortest possible. If $x_i y_j$ is an edge of C, then $P \cup C$ contains a cycle with a chord.

Step 5: Since $x_i y_j$ is not an edge of C, $x_{j-1} \neq x_i$. Let Q be an F-path from x_{j-1} to some root y_ℓ. By the choice of the pair x_i, y_j, we have $y_\ell \neq y_j$. Moreover, the (cyclic) order of these vertices along C is $x_i, x_{j-1}, y_j, y_\ell$. Thus, $P \cup Q \cup C$ contains a cycle in which $x_{j-1}y_j$ is a chord. This proves the claim and concludes the proof of the theorem. □

References

[1] C. Berge, Graphs and Hypergraphs, North-Holland, London, 1973.

[2] C. Berge, Hypergraphs, Combinatorics of Finite Sets, North-Holland, Amsterdam, 1989.

[3] B. Bollobás, Extremal Graph Theory, Academic Press, New York, 1978.

[4] W.G. Brown, P. Erdős, V.T. Sós, Some extremal problems on r-graphs. New Directions in the Theory of Graphs, in: Proceedings of the Third Ann Arbor Conference on Graph Theory, University of Michigan, Academic Press, New York, 1973, pp. 53–63.

[5] G.A. Dirac, Minimally 2-connected graphs, J. Reine Angew. Math. 228 (1967) 204–216.

[6] E. Győri, Personal communication.
[7] L. Lovász, Combinatorial Problems and Exercises, North-Holland, Amsterdam, 1979.
[8] M.D. Plummer, On minimal blocks, Trans. Amer. Math. Soc. 134 (1968) 85–94.
[9] L. Pósa, Problem no. 127, Mat. Lapok 12 (1961) 254 (in Hungarian).
[10] M. Yannakakis, On a class of totally unimodular matrices, Math. Oper. Res. 10 (1985) 280–304.
[11] C. Zhang, W. Li, A class of hypergraphs satisfying an inequality of Lovász, J. Combin. Theory Ser. B 34 (1983) 224–228.