Decomposition and l_1–embedding of weakly median graphs

Hans–Jürgen Bandelt [1] and Victor Chepoi [2]

[1] Mathematisches Seminar, Universität Hamburg, Bundesstr. 55, D-20146 Hamburg, Germany
[2] SFB 343 Diskrete Strukturen in der Mathematik, Universität Bielefeld, D-33615 Bielefeld, Germany

Abstract. Weakly median graphs, being defined by interval conditions and forbidden induced subgraphs, generalize quasi–median graphs as well as pseudo–median graphs. It is shown that finite weakly median graphs can be decomposed with respect to gated amalgamation and Cartesian multiplication into 5–wheels, induced subgraphs of hyperoctahedra (alias cocktail party graphs), and 2–connected bridged graphs not containing K_4 or $K_{1,1,3}$ as an induced subgraph. As a consequence one obtains that every finite weakly median graph is l_1–embeddable, that is, it embeds as a metric subspace into some $\mathbb{R}^n$ equipped with the 1–norm.

In this paper we continue to elaborate on a structure theory of graphs based on two fundamental operations, viz., Cartesian multiplication and gated amalgamation. While Cartesian multiplication is a standard operation, gated amalgamation seems to appear only in the context of median graphs and their generalizations; cf. [4, 6, 8, 23, 27]. An induced subgraph H of a graph G is called *gated* if for every vertex x outside H there exists a vertex x' (the *gate* of x) in H such that each vertex y of H is connected with x by a shortest path passing through the gate x'; cf. [18]. Then G is a *gated amalgam* of two graphs G_1 and G_2 if G_1 and G_2 are (isomorphic to) two intersecting gated subgraphs of G whose union is all of G. A graph with at least two vertices is said to be *prime* if it is neither a proper Cartesian product nor a gated amalgam of smaller graphs. For instance, the only prime median graph is the two–vertex complete graph K_2; see Isbell [21] and van de Vel [26]. More generally, the prime quasi–median graphs are exactly the complete graphs; quasi–median graphs were introduced by Mulder [23] and further studied in [14, 28, 8]. The pseudo–median graphs form yet another class of graphs for which the prime members are known; see Bandelt and Mulder [6]. Unfortunately, the latter class is not closed under Cartesian multiplication and does not include all quasi–median graphs. In order to overcome these deficiences we consider here the somewhat larger class of weakly median graphs, previously studied by Chepoi [10, 11] under the name "locally median" graphs. First, we say that a graph G is *weakly modular* if its shortest–path metric $d = d_G$ satisfies the following two conditions:

– for any three vertices u, v, w with $1 = d(v,w) < d(u,v) = d(u,w)$ there exists a common neighbour x of v and w such that $d(u,x) = d(u,v) - 1$;

– for any four vertices u, v, w, z with $d(v,z) = d(w,z) = 1$ and

$$2 = d(v,w) \leq d(u,v) = d(u,w) = d(u,z) - 1,$$

there exists a common neighbour x of v and w such that $d(u,x) = d(u,v) - 1$.

For an illustration of these conditions see [9, Fig.2]. A *weakly median* graph is a weakly modular graph that does not contain any two vertices with an unconnected triple of common neighbours; see Figure 1.

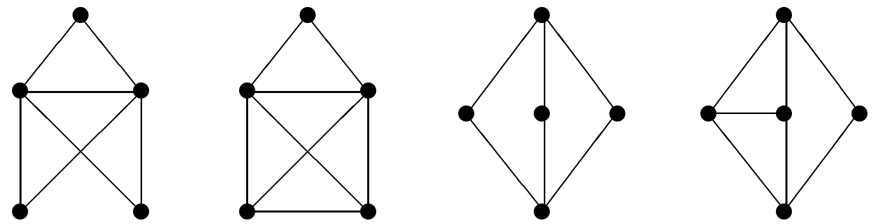

$K_{1,1,3}$ $K_{2,3}$

FIGURE 1. Forbidden induced subgraphs.

Note that all bridged graphs are weakly modular; cf. [10, 7]. A graph is called *bridged* if G does not contain any isometric cycle of length greater than 3, that is, each cycle of length greater than 3 has a shortcut in G; see Soltan and Chepoi [25] and Farber and Jamison [20]. Bridged graphs can easily be constructed since, according to [1, Corollaries 2.4 and 2.6], they admit certain vertex elimination schemes (relaxing simplicial elimination for chordal graphs). Now, a weakly median bridged graph (i.e., a bridged graph in which the first two graphs of Fig. 1 are forbidden induced subgraphs) is prime exactly when it has at least two vertices and does not have any cut vertex, that is, it is either K_2 or two–connected. This is an immediate consequence of the following two elementary facts: first, a gated subgraph cannot intersect any triangle just in a single edge, and second, in a two–connected bridged graph any two non–adjacent vertices having a common neighbour w are connected by another path within the neighbourhood of w; in consequence, two–connected bridged graphs do not have any proper gated subgraphs other than singletons. Furthermore, all wheels (whether bridged or not) are prime weakly median graphs; an n*–wheel* ($n \geq 4$) consists of a cycle of length n and a "central" vertex adjacent to all vertices of the cycle. Finally, all multipartite graphs of the form $K_{i_1,i_2,i_3,\ldots}$ (with $1 \leq i_j \leq 2$) different from $K_1, K_{1,2}$, and $K_{2,2}$ are prime weakly median graphs. A particular instance is the α*–octahedron* $K_{2,2,2,\ldots}$ (or *hyperoctahedron,* for short), which is the complement of the disjoint union of $\alpha \geq 3$ copies of K_2. For convenience, we refer to induced subgraphs of hyperoctahedra as to *subhyperoctahedra* when they contain either K_4 or an induced 4–wheel $K_{1,2,2}$ (that is, whenever they constitute 1–skeletons of at least 3–dimensional polyhedra).

THEOREM 1. *Every finite weakly median graph G (with more than one vertex) is obtained by successive applications of gated amalgamations from Cartesian products of two–vertex complete graphs, 5–wheels, subhyperoctahedra, and two–connected, K_4– and $K_{1,1,3}$–free bridged graphs. The latter building stones are exactly the graphs which can be realized as plane graphs such that all inner faces are triangles and all inner vertices have degrees larger than 5.*

Particular instances of this result are the decomposition theorems for quasi–median graphs (with prime graphs being complete) [23, 8] and pseudo–median graphs [6].

An important feature of weakly median graphs is that they embed in rectilinear space. A finite graph G with shortest–path metric d_G is said to be *l_1–embeddable* if there exists a distance–preserving embedding φ into some $\mathbb{R}^n$ equipped with the 1–norm $\|\cdot\|_1$, that is,

$$d_G(x,y) = \|\varphi(x) - \varphi(y)\|_1$$

for all vertices x, y of G. Assouad and Deza [3] have shown that a graph G is l_1–embeddable if and only if for some integer $\eta \geq 1$ it admits a *scale η embedding ψ* in some hypercube Q (being a Cartesian power of K_2), that is,

$$\eta \cdot d_G(x,y) = d_Q(\psi(x), \psi(y)) \quad \text{for all} \quad x, y;$$

see Figure 2 for an instance of a scale 2 embedding. The graphs with scale 1 embeddings in hypercubes are thus the isometric subgraphs of hypercubes (characterized in [17]). Recently, Shpectorov [24] has proved that a finite graph is l_1–embeddable if and only if it is an isometric subgraph of the Cartesian product of hyperoctahedra and "half–cubes" (which are obtained from one parity half of a hypercube, with two vertices being adjacent exactly when their distance in the hypercube equals 2). To decide whether a given graph G is l_1–embeddable it suffices to check its *prime components,* i.e., those prime (gated) subgraphs from which G can be built up by successive Cartesian multiplications and gated amalgamations:

PROPOSITION. *A finite graph G is l_1–embeddable if and only if every prime component of G is such. G has a scale η embedding in a hypercube if and only if every prime component does. G is l_1–rigid if and only if every prime component is such.*

The particular instance of amalgamations along single vertices has already been dealt with in [16, Proposition 7.6.1].

An l_1–embeddable graph G is called *l_1–rigid* if it has an essentially unique l_1–representation, that is, its metric d can uniquely be expressed in the form $d = \sum_{i=1}^{m} \lambda_i \cdot \delta_i$ $(\lambda_i > 0)$ as a positive linear combination of "split" (alias "cut") metrics δ_i that are associated with the *splits* $\{A_i, B_i\}$ of G, i.e., partitions of the vertex–set into two parts, according to

$$\delta_i(x,y) = \begin{cases} 0 & \text{if either} \quad x, y \in A_i \quad \text{or} \quad x, y \in B_i, \\ 1 & \text{otherwise,} \end{cases}$$

for $i = 1, \ldots, m$ (cf. [5, 16]). Necessarily, for each i the sets A_i and B_i constitute complementary *halfspaces* of G, that is, either set includes all shortest paths of G between any two of its vertices. Indeed, suppose the contrary: let x and y be vertices at minimal distance such that x, y lie in A_i, say, but some shortest x, y–path $x = x_0, x_1, x_2, \ldots, x_k = y$ does not belong to

A_i. Then $x_1, \ldots, x_{k-1}$ are in B_i by minimality of $d(x,y)$. Moreover, every split $\{A_j, B_j\}$ that separates x and y either separates x_1 and x_{k-1} or one of the edges xx_1, $x_{k-1}y$. Since $\{A_i, B_i\}$ separates both edges xx_1, $x_{k-1}y$ but not x from y, we obtain the inequality

$$d(x,y) \leq d(x,x_1) + d(x_1,x_{k-1}) + d(x_{k-1},y) - 2\lambda_i,$$

which however contradicts that x_1 and x_{k-1} lie on a shortest x,y–path. Therefore a necessary condition for l_1–embeddability is that there is an ample supply of halfspaces (separating the vertices). Namely, a graph G has a scale η embedding in a hypercube if and only if there is a collection $\mathcal{Z}$ of complementary halfspaces A_i, B_i such that any two adjacent vertices of G are separated by exactly η splits of $\mathcal{Z}$. Necessity follows from the preceding observation. As to the converse, we will show that the metric d of G is the sum of split metrics, associated to the splits from $\mathcal{Z}$. Pick an arbitrary shortest path $x = x_0, x_1, x_2, \ldots, x_k = y$ between two vertices x and y. Every split A_i, B_i, separating the vertices x and y, separates exactly one pair of adjacent vertices on this path. Conversely, if $x_j \in A_i$ and $x_{j+1} \in B_i$, then necessarily $x_0, \ldots, x_j \in A_i$ and $x_{j+1}, \ldots, x_k \in B_i$. Hence the split A_i, B_i does not separate any other pair of adjacent vertices from the selected shortest xy–path. Therefore x and y are separated by exactly $\eta d(x,y)$ splits from $\mathcal{Z}$. In order to construct all pairs of complementary halfspaces separating two adjacent vertices u and v, note that any such split necessarily separates the sets

$$W(u,v) = \{x \in V : d(u,x) < d(v,x)\}, \quad W(v,u) = \{x \in V : d(v,x) < d(u,x)\}.$$

So, it remains to partition the complement of $W(u,v) \cup W(v,u)$. For bipartite graphs this is trivial, because $W(u,v)$ and $W(v,u)$ cover the whole vertex–set. Therefore, for $\eta = 1$ we obtain the Djoković characterization of isometric subgraphs of hypercubes [17]: G is isometrically embeddable in a hypercube if and only if G is bipartite and $W(u,v), W(v,u)$ are halfspaces for any adjacent vertices u and v. The specific information on halfspaces of weakly median graphs obtained in [11, 12] together with Theorem 1 and the Proposition then enables us to prove the concluding result:

THEOREM 2. *Every finite weakly median graph G is l_1–embeddable. G has a scale 2 embedding in a hypercube if and only if it does not contain $K_{1,1,1,1,2}$ (that is, K_6 minus an edge) as an induced subgraph. In particular, G is l_1–rigid if and only if it is K_4–free.*

A finite weakly median graph G which contains some induced $K_{1,1,1,1,2}$ has scale η embeddings only for $\eta \geq 4$. The minimum scale η then depends solely on the maximal induced subhyperoctahedra. To determine this number η for a particular subhyperoctahedron is not a trivial task; see Deza and Laurent [16, chapter 7.4].

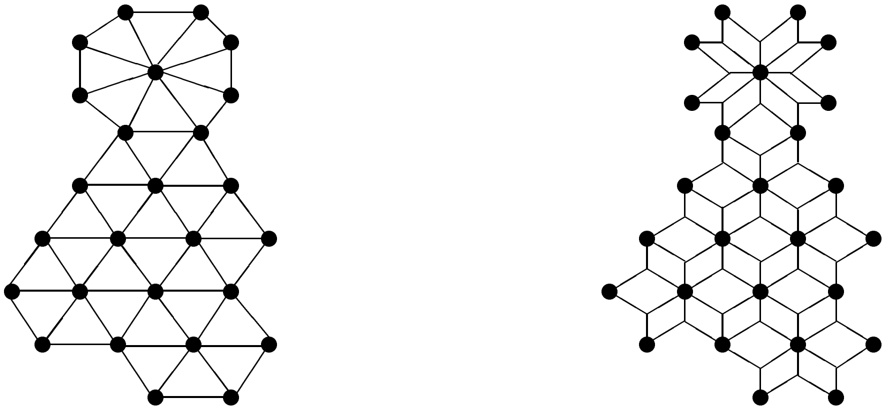

FIGURE 2. A weakly median bridged graph H and a fragment of a hypercube indicating the scale 2 embedding of H.

Proof of Theorem 1

We commence by establishing a number of auxiliary results. Unless stated otherwise, G is always a weakly median graph (not necessarily finite). Recall that the *interval* $I(u,v)$ between two vertices u and v in G consists of all vertices x on shortest paths from u to v, that is, $d(u,x)+d(x,v)=d(u,v)$. An induced subgraph (or a subset of vertices) H is called *convex* if H includes every interval $I(u,v)$ of G between two vertices u,v from H. Halfspaces are thus the nonempty convex sets with nonempty convex complement. The smallest convex (or gated, respectively) subgraph containing a given subgraph S is the *convex hull* (or *gated hull,* respectively) of S. A subgraph H is said to be Δ*–closed* if, for every triangle having two vertices in H, the third vertex belongs to H as well; then the smallest Δ–closed subgraph containing S is the Δ*–closure* of S. In order to check whether a given subgraph of G is convex or gated the following lemma is useful. This essentially coincides with Theorem 7 of Chepoi [10] and can be proved quite easily by induction.

LEMMA 1. *A connected subgraph H of G is convex if and only if for every pair of nonadjacent vertices in H all common neighbours belong to H whenever at least one common neighbour does. Moreover, a convex subgraph is gated if and only if it is Δ–convex.*

The next four lemmas provide the necessary information on gated hulls and isometric cycles in G. A *prism* is the Cartesian product $K_2 \square K_3$ of K_2 and K_3, and a *house* is obtained from a prism by deleting one vertex.

LEMMA 2. *The convex hull of an induced (i) house, (ii) 5–cycle, (iii) 4–wheel, respectively, in G is a (i) prism, (ii) 5–wheel, (iii) subhyperoctahedron, respectively.*

PROOF. Assertion (i) is obtained from [6, Lemma 4] and its proof.

Let C be an induced 5–cycle. Since G is weakly median there exists a common neighbour z of three vertices of C, two of which are adjacent with the third one being opposite to them. Then C and z constitute a 5–wheel, for otherwise, an induced house arises whose convex hull would not be a prism. If there is yet another vertex y in G adjacent to two non-adjacent vertices on C, then we would get either one of the forbidden induced subgraphs (see Figure 1) or an induced house whose convex hull is not a prism.

An induced 4–wheel can be extended to a maximal induced subhyperoctahedron H in G. Clearly H is included in the convex hull of this wheel. Suppose that H is not convex: then there exist two non–adjacent vertices u and v in H with a common neighbour y outside H. As H together with y does not induce a subhyperoctahedron, there exist two distinct vertices w and x in H which are not adjacent to y. Since both u and v are common neighbours of w and x, the five vertices induce the third or fourth graph of Figure 1, giving a contradiction. □

LEMMA 3. *Induced 5–wheels and convex subhyperoctahedra containing an induced 4–cycle are gated in G and do not contain any proper gated subgraph.*

PROOF. In view of the preceding lemmas it suffices to show that induced 5–wheels and subhyperoctahedra are Δ–closed. Clearly they do not have any proper gated subgraphs.

Suppose that W is an induced 5–wheel which is not Δ–closed. Then two adjacent vertices u and v of W have a common neighbour y outside W. If, say, u is the central vertex of W, then y must also be adjacent to the two common neighbours of u and v in W in order to avoid forbidden subgraphs. This, however, contradicts the fact that induced 5–wheels are convex. Therefore u and v are peripheral vertices of the wheel. Let z be the central vertex of W, and let t be the vertex opposite to the edge uv on the cycle. If $d(t, y) = 3$, then as G is weakly median the two vertices of W different from t, u, v, z would have a common neighbour with y, which is necessarily outside W, contrary to convexity. Hence $d(t, y) = 2$, and so (again as G is weakly median) there exists a common neighbour x of t, y, z, which is impossible by what has just been shown.

Let H be a convex subhyperoctahedron which is not Δ–closed. Then there exists a vertex y outside H such that the neighbours of y in H form a complete subgraph K of size at least 2. Consider any induced 4–cycle C in H. If two vertices of K belong to C, then we would obtain an induced house the convex hull of which includes induced 4–wheel, which is impossible. Therefore y and any vertex pair from K together with two (suitably chosen) non-adjacent vertices on C induce the first graph of Figure 1, a contradiction. □

LEMMA 4. *There are no isometric odd cycles in G of length greater than 5.*

PROOF. Suppose the contrary, and choose a cycle C having minimal length among all isometric odd cycles of length at least 7. Let C consist of the vertices $x_0, \ldots, x_{2n}$ and edges $x_i x_{i+1}$ ($i = 0, \ldots, 2n$, indices modulo $2n + 1$). Since x_0 and x_1 are at distance n from x_{n+1}, they have a common neighbour y_1 with $d(y_1, x_{n+1}) = n - 1$ (because G is weakly median). Further, x_2 and y_1 have a common neighbour y_2 with $d(y_2, x_{n+1}) = n - 2$. Continuing this way, we eventually obtain a shortest path $x_0, y_1, \ldots, y_{n-1}, x_{n+1}$ such that each y_i is adjacent to x_i ($i = 1, \ldots, n - 1$). Observe that each y_i is actually different from x_{i+1} because $d(x_0, x_{i+1}) = i + 1$ but $d(x_0, y_i) = i$. Now, a shortest path $y_{n-1}, y_n, \ldots, y_{2n-3}$ is constructed as follows: let y_i be a common neighbour of y_{i-1} and x_{i+2} with $d(x_0, y_i) = 2n - i - 2$ for

$i = n, \ldots, 2n-3$. Again, each y_i ($i \geq n-1$) must be different from x_{i+3}. We claim that the cycle induced by $x_0, x_1, \ldots, x_{n-1}, y_{n-1}, y_n, \ldots, y_{2n-3}$ is isometric. Indeed, suppose that for some $i = n, \ldots, 2n-3$ the distance from y_i to one of x_{i-n+1}, x_{i-n+2} is smaller than $n-1$. Then $d(x_{i+2}, x_{i-n+1}) < n$ or $d(x_{i+2}, x_{i-n+2}) < n$ would follow, conflicting with the isometry of C. This proves the claim. Since the new isometric cycle has length $2n-1$, we must have $n = 3$ by virtue of the initial minimality assumption. Thus, x_0, x_1, x_2, y_2, y_3 induce a 5–cycle with y_1 in its convex hull. Hence, by Lemma 2, x_2 and y_1 are adjacent. Now, by interchanging the roles of x_i and x_{8-i} for $i = 1, 2, 3$ we obtain yet another 5–wheel with central vertex y_1, so that x_6 and y_1 must be adjacent as well. This, however, implies $d(x_2, x_6) = 2$, a final contradiction. □

LEMMA 5. *If G does not contain any induced 4–wheel or 5–wheel, then the gated hull of any triangle is a two–connected bridged graph. No two–connected, weakly median bridged graph H has a proper gated subgraph.*

PROOF. First suppose that H contains a proper gated subgraph S. Pick a vertex x outside S having a neighbour w within S. Then there are shortest paths from x to the vertices of S passing through w. Since H is two–connected, each vertex of S different from w is connected to x by a path avoiding w. Take such a path of minimal length. Not all of its interior vertices can have w as their gate in S. Hence there exists an edge yz of this path such that w is the gate of y and some vertex $v \neq w$ is the gate of z in S. It is not difficult to see that v and w are adjacent with $d(w, y) = d(v, z)$. Choose a common neighbour u of v and x at distance $d(v, z) - 1$. Then u, v, w, x induce a 4–cycle, which however is impossible.

Now, there exists a maximal isometric two–connected bridged subgraph H of G that contains a given triangle. Indeed, directed unions preserve isometry, two–connectedness as well as the property of being bridged. To prove the lemma it thus suffices to show that H is gated (by what has just been proven).

First suppose that H is not convex. Then by Lemma 1, we can find non-adjacent vertices x, y in H having a common neighbour z in H and another one, v, outside H. We may assume that v, x, y, z are chosen so that the smallest cycle of H containing x, y, z (guaranteed by two–connectedness) has minimal length. Then choose a path P from x to y avoiding z such that the sum of all distances from z to the vertices of P is minimal. It is not difficult to see that, as H is bridged, the path P lies entirely in the neighbourhood of z. Since induced 4–cycles in G are necessarily convex (as 4–wheels are forbidden), P has length at least 3. By the initial minimality assumption x and y are the unique vertices on P adjacent to v. We claim that v and z must be adjacent. Suppose the contrary: then v, x, y, z together with the neighbour t of x on P induce a house. This house extends to an induced prism (according to Lemma 2); let w denote the common neighbour of t, v, y. Then w does not belong to H because H is bridged. Now, the cycle formed by w and the subpath of P from t to y is shorter than the one formed by v and P, thus contradicting the choice of v, x, y, z. This proves the claim.

Next we show that the larger subgraph H' induced by H together with v is also isometric in G. Suppose the contrary: let u be a vertex of H such that the distance of u and v in H' exceeds the distance $d(u, v) = k \geq 2$ in G. As no shortest path from u to v in G can pass through one of the three neighbours x, y, z of v (because H is isometric), the distances $d(u, x), d(u, y), d(u, z)$ are necessarily between k and $k+1$. So, we distinguish two cases.

Case 1: $d(u,z) = k+1$.

If $d(u,x) = d(u,y) = k$, then x and y have a common neighbour t in the isometric bridged subgraph H such that $d(u,t) = k-1$, thus yielding a forbidden 4–cycle (induced by t,x,z,y) in H. Therefore $d(u,y) = k+1$, say. Then H contains a common neighbour t of y and z at distance k to u. Since G is weakly modular, there exists a common neighbour w of t and v with $d(u,w) = k-1$. Now, t,v,w,y,z induce the fourth graph of Figure 1, which is impossible.

Case 2: $d(u,z) = k$.

Since now the vertices v and z are equidistant to u, they have a common neighbour w in G with $d(u,w) = k-1$. As the distance between u and v in H' is larger than k, the vertex w lies outside H. In order to avoid forbidden induced subgraphs, w must be adjacent to both x and y. Now, replacing v by w and k by $k-1$ we are back in Case 1, thus leading to a contradiction.

We conclude that the extended subgraph H' is indeed isometric. From what has been shown above we know that H' cannot include any induced 4–cycle. Therefore H' does not have any isometric even cycle at all: for otherwise, such a cycle C would contain v, and the two neighbours x', y' of the vertex z' opposite to v on C would admit a second common neighbour v' in G (because of weak modularity) which satisfies $d(v,v') = d(v,z') = 2$, so that a forbidden 4–cycle arises (whether v' belongs to H or not). By the hypothesis of the lemma (together with Lemma 2), H' is without induced 5–cycles, and hence by Lemma 4, it must be bridged. Clearly H' is two–connected, and therefore we arrive at a final contradiction to the maximality of H. This proves that H is in fact convex.

It remains to verify that H is Δ–closed as well. Suppose there exists a vertex v outside H having at least two neighbours x, y in H. Then, as H is convex, the neighbours of v in H form a complete subgraph. So, v is a simplicial vertex of the extended subgraph H' induced by H and v. Hence H' is a bridged graph, which is evidently two–connected. As to isometry, consider any vertex u of H. If $d(u,y) < d(u,v)$, then u and v are at distance $d(u,v)$ in H', too. We may therefore assume that $d(u,y) = d(u,v)$ because H is convex. Since G is weakly modular, we can find a common neighbour w of v and y at distance $d(u,v)-1$ to u. As H is convex, w belongs to H. This shows that there is a shortest path in H' between u and v of length $d(u,v)$. It follows that H' is also a two–connected, isometric bridged subgraph, thus conflicting with the choice of H. Therefore H is Δ–closed, concluding the proof. □

A two–connected, weakly median bridged graph is actually either a subhyperoctahedron (and thus a complete graph or a complete graph minus an edge) or K_4–free, as the following lemma confirms.

Lemma 6. *Every prime weakly median graph H containing some K_4 is a subhyperoctahedron.*

Proof. Suppose the contrary: then (by Lemmas 2,3, and 5) H is two–connected and bridged but is not a subhyperoctahedron although it contains some K_4. Extend this K_4 to a maximal induced subhyperoctahedron H', which is necessarily convex. By the hypothesis, we can find a vertex z outside H' which forms a triangle together with two vertices from H'. Now, however, we arrive at a contradiction in that either $H' \cup \{z\}$ would induce a subhyperoctahedron or a forbidden induced subgraph (from Fig. 1) would arise. □

To characterize the two–connected, K_4– and $K_{1,1,3}$–free bridged graphs via their planar embeddings, we will make use of the following counting argument.

LEMMA 7. *Let G be a finite two–connected plane graph in which all inner faces are triangles and all inner vertices (i.e., the vertices not incident with the outer face) have degrees larger than 5. Then the numbers n_2 and n_3 of vertices with degrees 2 and 3 satisfy the inequality $2n_2 + n_3 \geq 6$.*

PROOF. Let f denote the number of inner faces of G, m the number of edges, n the number of vertices, and b the number of vertices incident with the outer face. Then

$$f - m + n = 1 \qquad \text{and} \qquad 3f + b = 2m$$

hold according to Euler's theorem and the hypothesis that all inner faces are triangles. Eliminating f yields

$$3n - b - m = 3.$$

The information on the vertex degrees is turned into the inequality

$$\begin{aligned} 2m &\geq 6(n-b) + 4(b - n_2 - n_3) + 3n_3 + 2n_2 \\ &= 6n - 2b - 2n_2 - n_3, \end{aligned}$$

whence

$$2n_2 + n_3 \geq 6n - 2b - 2m = 6,$$

as required. □

We are now in position to identify the finite two–connected, K_4– and $K_{1,1,3}$–free bridged graphs with the plane graphs described in the preceding lemma, when chosing a planar embedding such that the outer face is bounded by the edges contained in exactly one triangle.

For the bridged graphs with the additional properties we construct a planar embedding recursively by employing the dismantling scheme of Anstee and Farber [1] (for a short proof, see [13]): there exists a vertex z dominated by some neighbour y in the sense that every vertex adjacent to z is also adjacent or identical to y. If the degree of z was larger than 3, then y and z together with three common neighbours would either induce $K_{1,1,3}$ or include some K_4, contrary to the hypothesis. Therefore z has degree 2 or 3. We assume that G has at least four vertices and that the desired planar embeddings can be realized for all proper induced subgraphs which are two–connected.

CASE 1: z has exactly two neighbours x and y (which are adjacent).

Then $G - z$ is a graph of the same kind, to which the induction hypothesis applies. Note that the edge xy belongs to exactly one triangle of $G - z$ because G is K_4– and $K_{1,1,3}$–free. Therefore we have chosen a planar embedding of $G - z$ with xy on the boundary of the outer face. Attaching the triangle x, y, z to $G - z$ so that z lies in the outer face of $G - z$, we obtain a plane graph with the required properties.

CASE 2: z has exactly three common neighbours w, x, and y (such that y is adjacent to w and x).

Then w and x are not adjacent. If $G - z$ is not two–connected, then y is the unique cut vertex. Moreover, $G - \{y, z\}$ comprises exactly two components, which together with y induce either K_2 or two–connected subgraphs of G. In any case we can transform and combine the planar embeddings of these subgraphs so that wy and xy lie on one line for which one of the associated closed halfplanes includes $G - z$. Placing z onto the complementary open halfplane and linking it with w, x, y produces the desired embedding. If $G - z$ is two–connected, then we could chose the planar embedding of $G - z$ right away, with wy and xy lying on the boundary of the outer face (since both edges belong to exactly one triangle of $G - z$). Locating z in this outer face we can extend the planar embedding to G, thereby creating two new triangles and turning y into an inner vertex. Consider a minimal path P in $G - z$ connecting w and x but not passing through y. Then P together with y and z induce a k–wheel with $k \geq 6$, whence y satisfies the degree constraint.

As to the converse, let G be a plane graph satisfying the hypothesis of Lemma 7. We may assume that G has at least four vertices. Consider any triangle of G : together with its interior in the plane it constitutes a plane graph H to which Lemma 7 equally applies. We infer that each vertex of the boundary triangle must have degree 2 in H, that is, H includes no inner vertex. Hence all triangles of G constitute inner faces (and vice versa). In particular, G does not include any K_4 or $K_{1,1,3}$ as an induced subgraph. To show that G is bridged, we proceed by induction.

Case 1: There exist two adjacent vertices u and v separating G.

Then necessarily u and v both lie on the boundary of the outer face. We can thus decompose G into two plane subgraphs G_1 and G_2 whose boundaries intersect in the edge uv and cover the boundary of G. Certainly, G_1 and G_2 fulfill the hypothesis of Lemma 7 and hence are bridged by the induction hypothesis. Then G is bridged as well.

Case 2: G does not have any separating edge.

Then, by Lemma 7, the boundary contains some vertex v of degree 3 in G. Let w, x, y be the neighbours of v. One of the edges vw, vx, vy does not lie on the boundary, say vx. Since vx is thus contained in two triangles, w and y must be adjacent to x. Further, x cannot be a boundary vertex, for otherwise, the edge vx would separate G. Therefore the plane subgraph $G - v$ is two–connected and inherits its inner faces and inner vertices from G. It follows from the induction hypothesis that $G - v$ is bridged. Suppose that G contains some isometric cycle C of length $2k$ or $2k + 1$ with $k \geq 2$. Then C includes v, w, y but not x. Substituting v by x creates a cycle in $G - v$ of the same length. This cycle must have a short cut, so that some vertex z on C is at distance k to v but $k - 1$ to x in G (as $G - v$ is clearly isometric in G). If w and y are at distance $k - 1$ to z, then there exist common neighbours w' of w and x and y' of x and y, both at distance $k - 2$ to z, such that either $w' = y'$ or w' and y' are adjacent because $G - v$ is bridged. Then v, w, x, y together with w', y' induce a 4– or 5–wheel, so that x would become an inner vertex of G with degree smaller than 6, a contradiction. Therefore C must be an odd cycle such that exactly one of w, y is at distance k to z, say w. Then the neighbour $u \neq v$ of w on C must be adjacent to x because G is bridged. Hence, as C is isometric, we infer that C is a 5–cycle comprising v, w, u, z, y, which together with x induces a 5–wheel, again a contradiction. We conclude that G is bridged.

This completes the proof of the second statement in Theorem 1, characterizing the bridged building stones.

The subsequent Lemmas 8 and 10 are needed to detect amalgams or products within G. Any gated subset S of G gives rise to a partition W_a ($a \in S$) of the vertex–set of G; viz., the *fibre* W_a of a relative to S consists of all vertices x (including a itself) having a as their gate in S. For adjacent vertices a, b of S, let U_{ab} be the set of vertices from W_a which are adjacent to vertices from W_b.

LEMMA 8. *Let S be a gated subgraph of G. Then each fibre W_a relative to S is gated. There exists an edge between two distinct fibres W_a and W_b if and only if a and b are adjacent. Moreover, for any two adjacent vertices a, b of S, the sets U_{ab} and U_{ba} constitute isomorphic gated subgraphs of G under the canonical isomorphism $f_{ab} : U_{ab} \to U_{ba}$ that maps each vertex in U_{ab} to its unique neighbour in U_{ba}.*

PROOF. In the case that S consists only of a single edge one could simply copy the proof of Theorem 12 from Bandelt and Mulder [6]. Still, we can use some of their arguments here. For instance, as in Step 3 (ibidem), one can easily prove that $d(a, x) = d(b, y)$ and $d(a, b) = d(x, y)$ for any two adjacent vertices $x \in W_a$ and $y \in W_b$. Trivially, each fibre W_a is connected because $I(a, x) \subseteq W_a$ for all $x \in W_a$. So, in order to show that W_a is gated, we may apply Lemma 1: assume that two vertices x and y of W_a have a common neighbour v in some fibre W_b. Then b must be adjacent to a, as noted above. Further, x and y belong to $I(a, v)$ and hence have a common neighbour z in $I(a, x) \cap I(a, y)$. Then $v, z \in I(b, x)$, whence there exists a common neighbour w of v and z in $I(b, v) \cap I(b, z)$. Clearly v is not adjacent to z and w is not adjacent to x or y. The five vertices v, w, x, y, z thus induce a forbidden subgraph. Therefore W_a is gated.

Now, for any two adjacent vertices a and b, the set U_{ab} includes all intervals $I(a, x)$ with $x \in U_{ab}$. This is readily proven by induction on $d(a, x)$, exactly as in Step 5 (ibidem). For $x, y \in U_{ab}$, let x' and y' denote the unique neighbour of x and that of y, respectively. Suppose x, y are adjacent but x' and y' are not. Since W_b is convex, x' and y' must have a common neighbour v in W_b. Then x, y, y', v, x' induce a 5–cycle because W_a and W_b are gated. By Lemma 2, these vertices have a common neighbour, which must belong to W_a as well as to W_b, being absurd. We conclude that the neighbours map f_{ab} is an isomorphism from U_{ab} onto U_{ba}. To prove that U_{ab} is gated, let z be a common neighbour of two vertices $x, y \in U_{ab}$. Necessarily, z belongs to W_a and is at distance 2 to the neighbours of x, y in U_{ba}. The gate of z in W_b is therefore adjacent to z, showing that $z \in U_{ab}$, as desired. □

The *cycle space* of a graph with edge–set E is the subspace of $(\mathbf{GF(2)})^E$ comprising all unions of closed walks. The isometric cycles clearly generate this space. In the presence of weak modularity the triangles and induced 4–cycles generate all isometric cycles, as is easily seen by induction. Recall from Duchet et al. [19] or Jamison [22] that a graph is *null–homotopic* if its cycle space admits a basis constituted solely of triangles. We then record the following elementary fact.

LEMMA 9. *A weakly modular graph is null-homotopic whenever every induced 4–cycle extends to a 4–wheel.*

In case that a proper gated subgraph S of G is two–connected and null–homotopic we can say more about the associated sets U_{ab} : the following lemma constitutes the tool for detecting proper decompositions of non–bipartite weakly median graphs.

LEMMA 10. *Let S be a gated two–connected and null–homotopic subgraph of G. Then the gated subgraphs U_{ab} (with a, b adjacent in S) are all isomorphic, and their union induces a gated subgraph H isomorphic to a Cartesian product $S \Box U$ (where U may be any U_{ab}). If W_a and U_{ab} ($b \in S$) do not coincide for some $a \in S$, then G is the gated amalgam of W_a and $G - (W_a - U_{ab})$.*

PROOF. First we show that $U_{ab} = U_{ac}$ whenever a, b, c form a triangle in S. Let xy be an edge of G with $x \in U_{ab}$ and $y \in U_{ba}$. According to Lemma 8 and its proof, c is equidistant to x and y, whence there is a common neighbour z of x and y on shortest paths to c. Then $d(c, z) = d(c, x) - 1 = d(a, x)$, and hence as a is the gate of x in S we infer that z belongs to W_c. Therefore $z \in U_{ca}$ as well as $x \in U_{ac}$. Interchanging the roles of a, b, c, this proves

$$U_{ab} = U_{ac}, \quad U_{ba} = U_{bc}, \quad U_{ca} = U_{cb}.$$

Now assume that q, r are any distinct neighbours of a in S. Then, as S is two–connected, there exists a path P from q to r not passing through a. By C denote the closed walk from q to r along P and then back to q via the vertex a. To prove that $U_{aq} = U_{ar}$ we proceed by induction on the minimal number k of triangles whose sum gives C (thereby using the null–homotopy of S). By what has just been shown we can assume that $k \geq 2$. Since P does not include a, there must be some common neighbour s of a and r such that the closed walk from q to s, via a, and then back to q is the (modulo 2) sum of $k - 1$ triangles. Then $U_{aq} = U_{as} = U_{ar}$ by virtue of the induction hypothesis. This justifies the shorthand U_a for the sets U_{ab} (with b being adjacent to a in S).

We can verify that all sets U_a ($a \in S$) are actually isomorphic using the same kind of argument: we claim that for every closed walk $a_0, a_1, \ldots, a_n, a_0$ the composition $f_{a_n a_0} \circ f_{a_{n-1} a_n} \circ f_{a_{n-2} a_{n-1}} \circ \ldots \circ f_{a_0 a_1}$ is the identity map. Indeed, this is evidently true for triangles ($n = 2$), by the first part of the proof. The general case is settled again by induction on the minimal number of triangles adding up to C. In particular, we get a unique isomorphism f_{rs} from U_r to U_s for any two (not necessarily adjacent) vertices $r, s \in S$, obtained by composing the isomorphisms f_{ab} along the edges ab of any path from r to s.

The product representation of H (the subgraph induced by the union of all sets U_a) is now immediate: pick any vertex a in S and consider the mapping

$$\begin{aligned} f : S \Box U_a &\rightarrow H \\ (s, x) &\mapsto f_{as}(x). \end{aligned}$$

This constitutes the desired isomorphism since (i) each mapping f_{as} is an isomorphism from U_a to U_s, (ii) the sets U_s ($s \in S$) partition H, (iii) there is an edge between U_r and U_s only if r and s are adjacent, and (iv) the isomorphism f_{rs} maps each vertex x onto a neighbour whenever r and s are adjacent.

Finally, assume $W_a \neq U_a$ for some $a \in S$. Clearly every common neighbour in G of any two vertices of H belongs to H, whence H is gated by Lemma 1. Then the subgraph $G - (W_a - U_a)$ obtained by deleting the vertices of W_a outside U_a is gated, and therefore G is the gated amalgam of this subgraph and the fibre W_a. □

Now, we have collected all the information that is necessary to conclude the proof of Theorem 1. Assume that G is a finite weakly median graph with at least three vertices which is neither a 5–wheel, nor a subhyperoctahedron, nor a two–connected bridged graph. We have to show that G decomposes as either a gated amalgam or a Cartesian product. If G includes an induced 4– or 5–wheel, then by Lemmas 2 and 3 it has a proper gated subgraph S, which is a subhyperoctahedron or a 5–wheel. Since these graphs are null–homotopic and two–connected, Lemma 10 provides us with the required decomposition. So, we can assume that G is without induced 4– or 5–wheels. If G still contains some triangle, then by Lemma 5 we obtain a proper gated subgraph S which is bridged and two–connected. Since bridged graphs are null–homotopic (cf. Lemma 9), Lemma 10 applies again, yielding a proper decomposition of G. It remains to consider the case where G is trianglefree. Then, by Lemmas 2 and 4, there are no odd cycles at all, and thus G is a median graph, for which the desired decomposition into copies of K_2 is well established; see Isbell [21] or Mulder [23].

In conclusion, note that Cartesian multiplication distributes over gated amalgamation, viz., the Cartesian product of a graph H with a gated amalgam of two graphs G_1 and G_2 equals the gated amalgam of $H \Box G_1$ and $H \Box G_2$. This completes the proof of the Theorem.

Proof of the Proposition

In the case that $G = G_1 \Box G_2$ is the Cartesian product of two nontrivial graphs G_1 and G_2 the assertions of the Proposition are evident; see [16, Proposition 7.5.2]. In fact, we may regard G_1 and G_2 as gated subgraphs of G intersecting in a single vertex. Then the halfspaces of G correspond to the pairs H_1, G_2 and G_1, H_2 where each H_i is a halfspace of G_i.

As to gated amalgamation, the following observation is instrumental.

LEMMA 11. *Let G be a graph having a scale η embedding φ in a hypercube Q. Then every scale η embedding ψ of a gated subgraph S of G in some hypercube R extends to a scale η embedding of G in some hypercube containing R.*

PROOF. Let T be the convex hull of the image $\varphi(S)$ in the hypercube Q. For each vertex x of G let x' be the gate of x in S. We claim that $\varphi(x')$ is the gate of $\varphi(x)$ in the subhypercube $\varphi(S)$. Suppose it is not: then some vertex z from $T - \varphi(S)$ is this gate. Choose any halfspace H of Q with $\varphi(x') \in H$ but $z \notin H$. Since $\varphi(x')$ is in the interval between $\varphi(x)$ and $\varphi(y)$ for each $y \in S$ (as φ is a scale embedding), it follows that H includes $\varphi(S)$ and hence T, yielding a contradiction. This proves the claim. In particular, the distance between the gates in T of two vertices $\varphi(w), \varphi(x)$ from the image of G equals $\eta d_G(w', x')$. Let U be a subhypercube of Q intersecting T in a single vertex such that the convex hull of T and U is all of Q. Letting φ_U denote the scale embedding φ of G in Q followed by the gate map onto U, we thus have

$$d_U(\varphi_U(w), \varphi_U(x)) = \eta(d_G(w, x) - d_G(w', x')).$$

Now, the required scale η extension of ψ is given by

$$x \;\mapsto\; (\psi(x'), \varphi_U(x)) \in R \,\Box\, U.$$

Indeed, for vertices w, x of G,

$$\begin{aligned} \eta d(w, x) &= \eta d_G(w', x') + d_U(\varphi_U(w), \varphi_U(x)) \\ &= d_R(\psi(w'), \psi(x')) + d_U(\varphi_U(w), \varphi_U(x)). \quad \Box \end{aligned}$$

Assume that G is the gated amalgam of two graphs G_1 and G_2, which admit scale η embeddings φ_1 and φ_2 in hypercubes Q_1 and Q_2, respectively. Let T be the convex hull of $\varphi_1(G_1 \cap G_2)$ in Q_1. By virtue of Lemma 11 we can extend the restriction $\varphi_1|_{G_1 \cap G_2}$ to a scale η embedding ψ of G_2 in a hypercube R such that R intersects Q_1 only in T. The median graph $Q_1 \cup R$ extends isometrically to a hypercube Q. We can then regard the union $\varphi_1 \cup \psi$ as a mapping from G to Q, yielding the required scale η embedding.

As to l_1–rigidity, observe that H is a halfspace of G exactly when either H or its complement is a halfspace of G_1 or G_2 not intersecting $G_1 \cap G_2$, or H is a gated amalgam of halfspaces H_i of G_i $(i = 1, 2)$. This obviously implies that G is l_1–rigid whenever G_1 and G_2 are such. Conversely assume that the l_1–embeddable graph G contains a gated subgraph S which is not l_1–rigid. Then G has some scale η embedding φ in a hypercube, while S admits yet another scale ξ embedding ψ in a hypercube such that $\varphi|_S$ and ψ induce different weighted systems of pairs of complementary halfspaces on S. Without loss of generality assume that the scales ξ and η are the same (since scales can be enlarged to arbitrary multiples). Then the extension of ψ to G guaranteed by Lemma 11 is essentially different from φ, showing that G is not l_1–rigid.

Proof of Theorem 2

In view of the Proposition it suffices to verify the assertions of the theorem only for the prime components of the given graph G. We may therefore assume that G is prime. First note that the 1–cube K_2 is trivially l_1–rigid.

LEMMA 12. *In a prime K_4–free weakly median graph G other than K_2, any two adjacent vertices u and v are separated by exactly two distinct pairs of complementary halfspaces.*

PROOF. Since G is two–connected, null–homotopic, and K_4–free, every edge belongs to exactly one or two triangles. Recall from [12, Lemmas 9 and 10] that for every triangle u, v, w the sets $W(u, v)$ and $W(v, u) \cup W(w, u)$ are convex. Let V denote the vertex–set of G. Consider any two adjacent vertices u and v.

CASE 1: u and v have a unique common neighbour w.

Then every vertex x equidistant to u and v is closer to w than u and v (by weak modularity), whence the convex sets $W(u, v), W(v, u)$, and $W(w, u) \cap W(w, v)$ partition V. It follows that $W(u, v)$ and $W(v, u)$ are non–complementary halfspaces. If H is any halfspace with $u \in H$ and $v, w \in V - H$, then necessarily $W(u, v) \subseteq H$ and $W(v, u) \cup W(w, u) \subseteq V - H$, thus yielding $W(u, v) = H$.

CASE 2: u and v have exactly two common neighbours w and w'.

Then every vertex x equidistant to u and v belongs to $W(w, u) \cap W(w, v)$ or $W(w', u) \cap W(w', v)$. Therefore $W(u, v) \cup W(w, v), W(v, u) \cup W(w', u)$ and $W(u, v) \cup W(w', v), W(v, u) \cup W(w, u)$ constitute two distinct pairs of complementary halfspaces. Let H be any halfspace with $u \in H$ and $v \in V - H$. Since $u, v \in I(w, w')$, the vertices w and w' are separated by $H, V - H$, say $w \in H$ and $w' \in V - H$. Then necessarily $W(w, v) \subseteq H$ and $W(w', u) \subseteq V - H$, whence $H = W(u, v) \cup W(w, v)$. □

From Lemma 12 and the observations preceding Theorem 2 we immediately infer that the graphs G of Lemma 12 have scale 2 embeddings in hypercubes. The proof further shows that the three splits of any triangle u, v, w uniquely extend to pairs of complementary halfspaces of G. This implies that the associated split metrics are linearly independent, thus establishing l_1–rigidity.

To conclude the proof of Theorem 2, we can assume (by Theorem 1) that G is a subhyperoctahedron containing K_4. The l_1–embeddability of hyperoctahedra has been established by Assouad [2]. It is easy to see that the 4–octahedron $K_{2,2,2,2}$ has a scale 2 embedding in a 4–cube, but the minimum scale for $K_{1,1,1,1,2}$ equals 4; see [16, Lemma 7.4.6]. The scale 2 embeddable subhyperoctahedra containing K_4 are thus the subhyperoctahedra K_{i_1,i_2,i_3,i_4} with $1 \leq i_j \leq 2$ $(j = 1, 2, 3, 4)$, all of which fail to be l_1–rigid; cf. [16, Proposition 7.4.3]. This completes the proof of Theorem 2.

Acknowledgement. A part of this research was done when the second author was supported by the Alexander von Humboldt Stiftung.

References

[1] R.P. Anstee and M. Farber, On bridged graphs and cop–win graphs, *J. Combin. Theory Ser. B* **44** (1988), 22–28.

[2] P. Assouad, Embeddability of regular polytopes and honeycombs in hypercubes, in: The geometric vein: The Coxeter Festschrift, 1981, pp.141–147.

[3] P. Assouad and M. Deza, Espaces métriques plongeables dans un hypercube: aspects combinatoires, *Annals Discrete Math.* **8** (1980), 197–210.

[4] H.-J. Bandelt and V. Chepoi, Cellular bipartite graphs, *Europ. J. Combin.* **17** (1996), 121–134.

[5] H.-J. Bandelt and A.W.M. Dress, A canonical decomposition theory for metrics on a finite set, *Advances Math.* **92** (1992), 47–105.

[6] H.-J. Bandelt and H.M. Mulder, Pseudo–median graphs: decomposition via amalgamation and Cartesian multiplication, *Discr. Math.* **94** (1991), 161–180.

[7] H.-J. Bandelt and H.M. Mulder, Cartesian factorization of interval–regular graphs having no long isometric odd cycles, in: Y. Alavi, G. Chartrand, O.R. Oellermann, A.J. Schwenk, eds., Graph Theory, Combinatorics, and Applications, Vol. 1 (Wiley, 1991) 55–75.

[8] H.-J. Bandelt, H.M. Mulder, E. Wilkeit, Quasi–median graphs and algebras, *J. Graph Theory* **18** (1994), 681–703.

[9] H.-J. Bandelt and E. Pesch, Dismantling absolute retracts of reflexive graphs, *Europ. J. Combin.* **10** (1989), 211–220.

[10] V. Chepoi, Classifying graphs by means of metric triangles (in Russian), Metody Diskret. Analiza **49** (1989), 75–93.

[11] V. Chepoi, d–Convexity and local conditions on graphs (in Russian), *Issledovania po Prikladnoj Matematike i Informatike*, Chişinău, Stiinţa, 1990, pp.184–191.

[12] V. Chepoi, Separation of two convex sets in convexity structures, *J. Geometry* **50** (1994), 30–51.

[13] V. Chepoi, Bridged graphs are cop–win graphs: an algorithmic proof, *J. Combin. Theory Ser. B* **69** (1997), 97–100.

[14] F.R.K. Chung, R.L. Graham, and M.E. Saks, A dynamic location problem for graphs, *Combinatorica* **9** (1989), 111–131.

[15] M. Deza and M. Laurent, l_1–Rigid graphs, *J. Alg. Combin.* **3** (1994), 153–175.

[16] M. Deza and M. Laurent, *Geometry of Cuts and Metrics.* Springer–Verlag, Berlin, 1997.

[17] D. Ž. Djoković, Distance–preserving subgraphs of hypercubes, *J. Combin. Theory Ser. B* **14** (1973), 263–267.

[18] A.W.M. Dress and R. Scharlau, Gated sets in metric spaces, *Aequationes Math.* **34** (1987), 112–120.

[19] P. Duchet, M. Las Vergnas, and H. Meyniel, Connected cutsets of a graph and triangle bases of the cycle space, *Discr. Math.* **62** (1986), 145–154.

[20] M. Farber and R.E. Jamison, On local convexity in graphs, *Discr. Math.* **66** (1987), 231–247.

[21] J.R. Isbell, Median algebra, *Trans. Amer. Math. Soc.* **260** (1980), 319–362.

[22] R.E. Jamison, On the null–homotopy of bridged graphs, *Europ. J. Combin.* **8** (1987), 421–428.

[23] H.M. Mulder, The interval function of a graph, *Math. Centre Tracts* **132** (Amsterdam), 1980.

[24] S. V. Shpectorov, On scale embeddings of graphs into hypercubes, *Europ. J. Combin.* **14** (1993), 117–130.

[25] V.P. Soltan and V.D. Chepoi, Conditions for invariance of set diameters under d–convexification in a graph, *Cybernetics* **19** (1983), 750–756.

[26] M. van de Vel, Matching binary convexities, *Topology Appl.* **16** (1983), 207–235.

[27] M.L.J. van de Vel, *Theory of Convex Structures.* North–Holland, Amsterdam, 1993.

[28] E. Wilkeit, The retracts of Hamming graphs, *Discr. Math.* **102** (1992) 197–218.