Europ. J. Combinatorics (1996) **17,** 121–134

Cellular Bipartite Graphs

HANS-JÜRGEN BANDELT AND VICTOR CHEPOI

In this paper we investigate the graphs that are obtained from single edges and even cycles by successive gated amalgamations. These 'cellular' graphs are characterized among bipartite graphs by having a totally decomposable shortest-path metric, and can be recognized by a quadratic time algorithm.

1. AVANT-PROPOS

Graphs with their shortest-path metrics are particular instances of finite metric spaces, and may thus be investigated from the metric point of view. Although the theory of finite metric spaces is not yet fully developed, some facets are already well studied; cf. [4, 10]. The l_1-embeddability question for metric spaces prompted the investigation of l_1-graphs [9, 13]. A particular class of finite l_1-spaces, possessing a rich theory, are formed by the totally decomposable spaces, in which the summands in a (canonical) l_1-decomposition would obey a certain compatibility rule [4]. It is the purpose of this note to demonstrate that the bipartite graphs with totally decomposable metric have a convenient decomposition scheme, the ingredients of which are gated amalgamation as a fundamental operation and even cycles and single edges as building stones.

2. TOTAL DECOMPOSABILITY

The simplest (pseudo-)metrics on a finite set X are the 'split', alias 'cut', metrics δ_S associated with the *splits* $S = \{A, B\}$ of X, i.e. partitions of X into two non-empty subsets A and B:

$$\delta_S(x, y) = \begin{cases} 0 & \text{if } x, y \in A \quad \text{or} \quad x, y \in B, \\ 1 & \text{otherwise.} \end{cases}$$

By definition, an l_1-metric d on X is any positive linear combination of split metrics:

$$(*) \qquad d = \sum_{S \in \mathscr{X}} \lambda_S \cdot \delta_S$$

with $\lambda_S > 0$ for all S from a collection $\mathscr{X}$ of splits. A *totally decomposable* metric d on X is a particular l_1-metric, which admits a 'feasible' representation $(*)$ with a collection $\mathscr{X}$ consisting of triplewise 'weakly compatible' splits. Three splits $\{A_i, B_i\}$ $(i = 1, 2, 3)$ are said to be *weakly compatible* if $A_1 \cap A_2 \cap A_3 \neq \varnothing$ implies $B_1 \cap B_2 \cap B_3 = B_i \cap B_j$ for some i, j; or, equivalently, there is no 4-subset Y of X such that the three splits $\{A_i, B_i\}$ would induce the three distinct splits on Y separating two from two points. The weight

0195-6698/96/020121 + 14 $18.00/0

λ_S of a split $S = \{A, B\}$ participating in a feasible representation (*) is then determined as the *isolation index*

$$\alpha_{A,B} = \tfrac{1}{2} \cdot \min_{\substack{a,a' \in A \\ b,b' \in B}} (\max\{d(a, b) + d(a', b'), d(a, b') + d(a', b), d(a, a') + d(b, b')\} - d(a, a') - d(b, b'));$$

see [4]. The metrics obtained from trees or cycles are simple instances of totally decomposable metrics.

Two facts concerning totally decomposable metrics are needed here. First, recall that a subset A of a metric space (X, d) is *convex* if the (metric) *interval*

$$I(u, v) = \{x \in X: d(u, x) + d(x, v) = d(u, v)\}$$

between any two points u and v of A lies entirely in A. The *convex hull* $\operatorname{conv}(Y)$ of a subset Y is the smallest convex set containing Y.

FACT 1 ([4, Proposition 3 and its proof]). *If d is a totally decomposable metric on X, then*

$$\operatorname{conv}(Y) = \bigcup_{u,v \in Y} I(u, v)$$

for every set $Y \subseteq X$.

A condition on subspaces stronger than 'convex' is the following. A subset Y of X is *gated* if for every point $x \in X$ there exists a (unique) point $x' \in Y$ (the *gate* for x in Y) such that $x' \in I(x, y)$ for all $y \in Y$ (cf. [12]).

FACT 2 ([4, Proposition 2]). *If X is covered by two intersecting proper gated subsets Y and Z such that the restrictions of the metric d to Y and Z are totally decomposable, then d is totally decomposable.*

In the preceding situation we say that X is a *gated amalgam* of Y and Z (along $Y \cap Z$). Gated amalgamations play an important role in structure theories of classes of graphs generalizing median graphs (see [6, 7, 14]).

3. MAIN RESULTS

We can now state the characterization of bipartite graphs with totally decomposable metrics. Since they are built up from cycles (their 'cells'), we dub them *cellular* graphs. All graphs considered here are assumed to be finite.

THEOREM 1. *For a bipartite graph $G = (V, E)$ with at least two vertices, the following conditions are equivalent*:
(1) *the metric d of G is totally decomposable*;
(2) $\operatorname{conv}(u, v, w) = I(u, v) \cup I(v, w) \cup I(w, u)$ *for all $u, v, w \in V$*;
(2′) $\operatorname{conv}(X) = \bigcup_{x,y \in X} I(x, y)$ *for all $X \subseteq V$*;
(3) *every isometric cycle of G is gated, and G does not contain any three isometric*

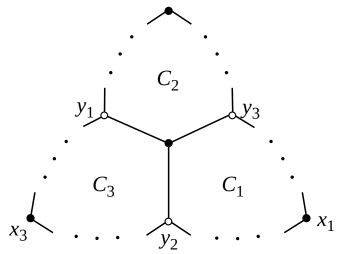

FIGURE 1. An obstruction to total decomposability.

cycles C_1, C_2, C_3 and three distinct edges e_1, e_2, e_3 sharing a common vertex such that e_i belongs to C_j exactly when $i \neq j$ (see Figure 1);
(4) *G can be obtained from a collection of single edges and even cycles by successive gated amalgamations;*
(5) *the splits $S(u, v) = \{W(u, v), W(v, u)\}$ for $uv \in E$ are triplewise weakly compatible, where $W(u, v) = \{x \in V: d(u, x) < d(v, x)\} = \{x \in V: u \in I(v, x)\}$ and $W(v, u) = V - W(u, v)$.*

Condition (2′) entails the *Peano property* ('join-hull commutativity', cf. [14]), stating that the convex hull of the union of a convex set Y with a vertex x equals the union of all intervals $I(x, y)$ with y from Y. According to [8], a bipartite graph G having the Peano property also enjoys the *Pasch property* [14], which then guarantees that any two disjoint convex sets A and B are separated by some 'convex' split $S(u, v)$; that is, there exists an edge uv in G such that $W(u, v)$ and $W(v, u)$ are convex and include A and B, respectively. Summarizing, a cellular bipartite graph G is a Pasch–Peano graph *sensu* [14] (alias 'join space'). As a simple consequence of this in conjunction with condition (5) note that for any four vertices x_1, x_2, x_3, x_4 at least one intersection $I(x_1, x_i) \cap I(x_j, x_k)$ with $\{i, j, k\} = \{2, 3, 4\}$ is non-empty. In particular, the Radon number (as defined in [14]) is at most 3.

The graphical representation of a totally decomposable metric by a network as described in [4, 5] would recover a cellular bipartite graph from its metric. In general, however, not every network corresponding to a totally decomposable metric would yield a cellular bipartite graph (after disregarding edge lengths): the obstruction shown in Figure 1 would often arise (cf. [5, Figures 0, 1 and 3]).

We say (by slight abuse of language) that a gated cycle of length k is *pendant* in G if it includes a path of length $\frac{1}{2}k - 2$, all vertices of which have degree 2 in G. Removing this path from G then results in an isometric subgraph and hence again a cellular bipartite graph. The structural information on cellular bipartite graphs that is gathered in the proof of Theorem 1 allows to immediately derive the following fact, which establishes an elimination scheme for cellular bipartite graphs.

THEOREM 2. *Every cellular bipartite graph with at least two vertices has either two pendant vertices or a pendant gated cycle.*

In each step the number of edges is decreased by at most twice the number of deleted vertices until one reaches a path with three vertices. Therefore the number m of edges of a cellular bipartite graph with $n \geq 3$ vertices is bounded above by $m \leq 2n - 4$. We conjecture that actually $\lfloor 2(n - \sqrt{n}) \rfloor$ is an upper bound for m, which is sharp for every n (being attained in a subclass of cube-free median graphs). Cube-free median graphs [3] are built up by gated amalgamations from single edges and 4-cycles

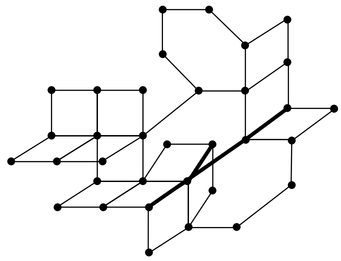

FIGURE 2. A cellular bipartite graph and (in bold) a gated cutset.

exclusively; that is, they are the cellular bipartite graphs in which all gated cycles have length 4. In particular, Theorem 2 generalizes the vertex elimination scheme for cube-free median graphs, established in [1, Corollary 2].

We can actually further specify the gated amalgamations needed to build up all cellular bipartite graphs. A *cutset* R of a connected graph G is any subset (or subgraph) for which $G - R$ is disconnected. It is evident that every gated cutset R induces a representation of G as a gated amalgam of two gated subgraphs G_1 and G_2 along R (and vice versa): the union of R and any component of $G - R$ may serve as G_1, while R together with the other component(s) of $G - R$ then gives G_2.

THEOREM 3. *Every cellular bipartite graph either is indecomposable* (*i.e., comprises a single vertex, or a single edge, or an even cycle*) *or possesses a gated cutset that is a tree.*

The proof of Theorem 3 (being rather constructive) entails a polynomial time algorithm for finding a gated tree cutset. A gated amalgam of two cellular bipartite graphs along a tree is illustrated in Figure 2.

The Cartesian product of a 3-star with itself (being a cube-free median graph) is a cellular graph which is not planar (since it contains a homeomorph of $K_{3,3}$). Nevertheless, we have a kind of Euler formula when one considers gated cycles instead of 'faces', as follows.

COROLLARY 1. *Let G be a cellular bipartite graph with n vertices, m edges, and g gated cycles. Then*

$$n - m + g = 1.$$

PROOF. We proceed by induction. If G is indecomposable, then the equality trivially holds. Therefore, let G be a gated amalgam of two gated subgraphs, G_1 and G_2, along a tree G_0 ($=G_1 \cap G_2$). Let n_i, m_i, g_i be the numbers of vertices, edges and gated cycles, respectively, of G_i ($i = 0, 1, 2$). Then, by the induction hypothesis, we obtain

$$n - m + g = n_1 + n_2 - n_0 - m_1 - m_2 + n_0 - 1 + g_1 + g_2 = 1,$$

using the fact that $g_0 = 0$ and $m_0 = n_0 - 1$. □

The *cycle space* of any graph G is the linear space over **GF(2)** having all Eulerian subgraphs (in which the vertices have even degrees) as its elements, with symmetric difference as (Boolean) addition. A simple fact is that every cycle is a Boolean sum of isometric cycles (proof by induction on the length). Therefore one would always find a basis of the cycle space comprising only isometric cycles. It is well known that $m - n + 1$ is the dimension of this space. For a cellular bipartite graph G, this number equals g, whence we obtain the following fact.

Corollary 2. *The gated cycles of a cellular bipartite graph G constitute a basis of the cycle space of G.*

4. Proof of Theorems 1 and 2

The implication $(2') \Rightarrow (2)$ is trivial, while $(1) \Rightarrow (2')$ and $(4) \Rightarrow (1)$ are covered by Facts 1 and 2.

$(2) \Rightarrow (3)$: We claim that every isometric cycle C is gated. Suppose the contrary: then there is a vertex x outside C having no gate in C; that is, a vertex w of C at minimum distance to x cannot lie in the interval between x and the vertex w' opposite to w on C. Let v and y be the two neighbours of w on C. Since w belongs to $I(v, x)$, we have

$$y \in C \subseteq \operatorname{conv}(v, w', x) = I(v, w') \cup I(w', x) \cup I(x, v)$$

by the hypothesis (2). Since $w \in I(v, y) \cap I(x, y) - I(w', x)$, we cannot allocate y to any of the three intervals in question. This conflict settles the claim.

Next, suppose that we could find three gated cycles C_1, C_2, C_3 sharing a common vertex x_0 such that each pair C_i, C_j intersects in an edge $x_0 y_k$ for $\{i, j, k\} = \{1, 2, 3\}$ with y_1, y_2, y_3 being distinct (see Figure 1). Let x_i be the vertex opposite to x_0 on C_i $(i = 1, 2, 3)$. Then the union $I(x_1, x_2) \cup I(x_2, x_3) \cup I(x_3, x_1)$ contains y_1, y_2, y_3 but not x_0. However, this violates (2) as $x_0 \in I(y_j, y_k)$ for $j \neq k$, thus finally establishing (3).

To prove that (3) implies (4), we need an auxiliary result concerning the Djoković relation Θ of a bipartite graph $G = (V, E)$. Define, for any edges uv and xy of G,

$$uv\,\Theta\,xy \Leftrightarrow \text{either } x \in W(u, v) \text{ and } y \in W(v, u), \quad \text{or} \quad y \in W(u, v) \text{ and } x \in W(v, u),$$

which is, of course, equivalent to either $u \in W(x, y)$ and $v \in W(y, x)$, or $v \in W(x, y)$ and $u \in W(y, x)$. The relation Θ is transitive (and hence an equivalence relation on E) iff G can be represented as an isometric subgraph of some hypercube [11]. We may compare Θ to the following relation Ψ^*. First say that two edges uv and xy are in relation Ψ if they either are equal or constitute opposite edges on some gated cycle C of G. Then let Ψ^* be the transitive closure of Ψ on the edge set E.

Lemma 1. *Let G be a bipartite graph in which every isometric cycle is gated. Then the relations Θ and Ψ^* on the edge set of G coincide. In particular, G is isometrically embeddable into a hypercube.*

Proof. We first claim that $\Theta \circ \Psi \subseteq \Theta$. Let xy and $y'x'$ be opposite edges of some gated cycle C (with x and x' being opposite on C), and assume $uv\,\Theta\,xy$ for some edge uv of G, say $x \in I(u, y)$ and $y \in I(v, x)$. Then the gate of u in C belongs to the path $I(x, y')$, whence $y' \in W(u, v)$. Similarly, we infer $x' \in W(v, u)$, and therefore $uv\,\Theta\,y'x'$, which settles the claim. A trivial induction yields

$$\Psi^k = \Psi^{k-1} \circ \Psi \subseteq \Theta \circ \Psi \subseteq \Theta,$$

and hence

$$\Psi^* = \bigcup_k \Psi^k \subseteq \Theta.$$

In order to show that Ψ^* includes Θ, proceed by induction on the distance k between two distinct edges uv and xy in the relation Θ:

$$k = \min(d(u, x), d(u, y), d(v, x), d(v, y)) \geqslant 1;$$

say, $d(u, x) = d(v, y) = k$. Select any shortest paths P and Q joining the pairs u, x and

v, y, respectively. If the cycle C formed by P, Q together with the edges uv, xy is isometric (and hence gated), then, trivially, $uv\ \Psi\ xy$. Otherwise, some vertex w of P is connected by a path R to some vertex z of Q such that R is shorter than both paths connecting w and z along C. Necessarily, R includes some edge rs with $r \in W(u, v)$ and $s \in W(v, u)$, so that $rs\ \Theta\ uv$. By the choice of R we obtain $d(u, r) < d(u, x)$, and therefore $uv\ \Psi^*\ rs$ by the induction hypothesis. In view of symmetry, we have $xy\ \Psi^*\ rs$ as well. Then, as Ψ^* is an equivalence relation, we conclude that $uv\ \Psi^*\ xy$. This yields $\Theta \subseteq \Psi^*$, and in conjunction with $\Psi^* \subseteq \Theta$ establishes the desired equality. According to the theorem of Djoković [11], transitivity of Θ ensures isometric embeddability into a hypercube. This completes the proof of the lemma. □

LEMMA 2. *Assuming condition* (3) *of Theorem* 1, *let D be a non-trivial block of the equivalence relation $\Theta = \Psi^*$, and let F be the union of all gated cycles containing edges from D. Then for any edge uv of D, both*

$$F_1 = F \cap W(u, v) \qquad \textit{and} \qquad F_2 = F - F_1 = F \cap W(v, u)$$

are convex trees in G. Moreover, F_1 is gated in the convex subgraph $W(u, v)$ and F_2 is gated in $W(v, u)$.

PROOF. Each gated cycle containing an edge of D shares exactly two opposite edges with D, and thus attributes a path to F_1 and F_2 each. Therefore F_1 and F_2 constitute connected subgraphs of G. Suppose that F_1 is not a convex tree in G: choose a path P in F_1 of smallest length k which is not convex in G. Then, necessarily, there is a shortest path R in G of length k intersecting P only in the end vertices y and z. We claim that the cycle C_1 formed by P and R is isometric and hence gated. Indeed, otherwise we could find interior vertices p of P and r of R that are connected by a shortest path having no interior vertex in common with P and R. Let q be the neighbour of p on this path. Since G is an isometric subgraph of a hypercube, all intervals are convex. Consequently, $q \in I(p, r) \subseteq I(y, z)$, and thus q belongs to either $I(p, y)$ or $I(p, z)$. As these two intervals constitute (convex) subpaths of P, we arrive at a contradiction. This proves the claim that C_1 is gated. Then C_1 does not contain any edge from D because P is included in F_1. On the other hand, P is composed of subpaths of gated cycles that share edges with D. Therefore there must exist an interior vertex x of P such that the two edges xx_2 and xx_3 from P incident with x belong to gated cycles C_3 and C_2, respectively, that also have edges with D in common. If C_2 and C_3 share yet another vertex x_1, then the three cycles C_1, C_2, C_3 and edges xx_1, xx_2, xx_3 would contradict the hypothesis (3). Hence x (being the unique common vertex of C_2 and C_3) is the gate in C_2 of each vertex of C_3, and vice versa. This, however, conflicts with the fact that some edge of C_2 is in relation Θ to another edge of C_3. This contradiction finally proves that F_1 and (analogously) F_2 are convex trees in G.

Suppose that F_1 is not gated in $W(u, v)$. Choose a vertex $z \in W(u, v)$ at minimum distance to F_1 having no gate in F_1. Let r be a vertex of F_1 closest to z, and let t be a vertex of F_1 such that the interval $I(z, t)$ does not contain r, where $d(r, t)$ is as small as possible. Then the neighbour s of t on the path $I(r, t)$ satisfies $r \in I(z, s)$. By minimality of $d(r, z)$, any neighbour y of z in $I(r, z)$ is closer to s than to t. Therefore the edges st and yz are in relation Θ. Applying the first part of the proof to the block of Θ containing these two edges, we infer that the intervals $I(s, y)$ and $I(t, z)$ are (disjoint) convex paths. Then the cycle C_1 formed by these two paths together with yz and st must be induced for, otherwise, either $I(r, t)$ could not be a convex path, or the intersection of $I(z, r)$ and $I(z, t)$ would contain a vertex different from z, thus violating minimality of $d(r, z)$. If C_1 is not isometric, then there exists a shortest path of length at least 2 between a vertex $p \neq y$ of $I(s, y)$ and a vertex w of $I(t, z)$ which intersects C_1 only in p and w. Since intervals in G are convex and G is bipartite, the neighbour q of p on this path would necessarily belong to $I(s, z)$. Then, as $p \neq y$ and $I(s, y)$ is a convex path, it follows that $q \in I(p, z) - I(p, y)$. Therefore $pq\,\Theta\,yz$. Since then $pq\,\Theta\,st$ and $I(r, t)$ is a convex path, we must have $p \in I(r, z)$, whence $q \in I(r, z) \cap I(t, z)$, contradicting the minimality choice of z. So, C_1 is indeed a gated cycle (included in the convex subgraph $W(u, v)$ by construction). Since $I(r, t)$ is a path of length at least 2 shared by F_1 and C_1, there must be a subpath x_2, x, x_3 of $I(r, t)$ such that the edges x_2x and xx_3 belong to some gated cycles C_3 and C_2, respectively, which each have two edges in common with D. Then, as in the first part of the proof, we eventually obtain a forbidden configuration (Figure 1). We finally conclude that F_1 and F_2 are gated subgraphs of $W(u, v)$ and $W(v, u)$, respectively. □

Now, we have the essential prerequisites at hand to accomplish the proof of $(3) \Rightarrow (4)$. Assume that G is neither a single edge nor an even cycle. We may further assume that G is 2-connected (i.e. without cut vertex) for, otherwise, G could be decomposed via gated amalgamation along a single vertex. In particular, every edge of G lies on some cycle—actually on some gated cycle because $\Theta = \Psi^*$. Hence there are at least two distinct gated cycles in G.

CASE 1. There exist two disjoint gated cycles C_1 and C_2. Take an edge uv of a path joining a vertex from C_1 with one from C_2 which has smallest length. Then $W(u, v)$ includes C_1, say, while $W(v, u)$ includes C_2. Let F_1 and F_2 be defined for the Θ-block containing uv, as in the previous lemma. Since C_i is not included in the tree F_i $(i = 1, 2)$, F_i $(i = 1, 2)$ is a cutset. Then, by virtue of Lemma 2, both $W(u, v) \cup F_2$ and $W(v, u) \cup F_1$ are proper gated subgraphs of G amalgamated along $F = F_1 \cup F_2$.

We may therefore assume that the gated cycles of G pairwise intersect. Then necessarily, they all have some vertex v in common (cf. [2, Proposition 2.4]). Note that G is covered by its gated cycles since every edge lies on a gated cycle.

CASE 2. All gated cycles intersect in a single edge. Then G is the gated amalgam of any gated cycle and the union of the other gated cycles (along that common edge).

CASE 3. The gated cycles intersect in the vertex v but do not share a common edge. Since G is 2-connected, two distinct gated cycles C_1 and C_2 of G intersect in an edge which is necessarily incident with v, say, uv. Let $C_1, \ldots, C_k$ $(k \geqslant 2)$ be the gated cycles containing uv, and let w_i be the neighbour of v on C_i different from u $(i = 1, \ldots, k)$. Consider the Θ-block D containing uv, and define F, F_1, F_2 as in Lemma 2. Then the union $C_1 \cup \cdots \cup C_k$ coincides with the gated subgraph F. The union F' of all other gated cycles of G intersects F in v and a subset of $\{w_1, \ldots, w_k\}$. Evidently, $F \cap F'$ is a cutset of G, and induces a star. Since F_2 is gated in $W(v, u)$, it easily follows that $F \cap F'$ is a gated cutset of G, whence G is the gated amalgam of F and F' in this case. Therefore the implication $(3) \Rightarrow (4)$ is settled.

It remains to prove that (5) is equivalent to (1), say. This can actually be derived from [9, Proposition 3.1 and its proof]. Alternatively, one could verify the implications $(1) \Rightarrow (5) \Rightarrow (3)$ directly as follows. Assuming (1), for every edge uv of G, any d-split $\{A, B\}$ with $u \in A$ and $v \in B$ must coincide with $S(u, v)$ because its parts A and B are necessarily convex. This establishes (5), since d-splits are weakly compatible. Finally, assume that (5) holds. We will show that G satisfies condition (3). Suppose that there exists an isometric cycle C which is not gated. Then we select vertices x, w, w', v, y as in the proof of the implication $(2) \Rightarrow (3)$. Choose a neighbour x' of x in $I(w, x)$. Then the three splits $S(v, w)$, $S(w, y)$, $S(x, x')$ and the four vertices v, w', x, y constitute an obstruction to weak compatibility. We conclude that all isometric cycles in G are gated. Next suppose that three gated cycles C_1, C_2, C_3 pairwise intersect in edges x_0y_1, x_0y_2, x_0y_3, as indicated in Figure 1. Let x_i be the vertex opposite to x_0 on C_i ($i = 1, 2, 3$). Then the splits $S(x_0, y_1)$, $S(x_0, y_2)$, $S(x_0, y_3)$ and the vertices x_0, x_1, x_2, x_3 yield an obstruction. The proof of Theorem 1 is now complete.

It remains to verify Theorem 2. Take any edge yz of G and select edges uv and $u'v'$ with $W(u, v) \subseteq W(y, z)$ and $W(v', u') \subseteq W(z, y)$ such that $W(u, v)$ and $W(v', u')$ have as fewest vertices as possible. Let F comprise the edge uv and (if there are any) the gated cycles containing edges Θ-equivalent to uv (cf. Lemma 2). In a similar way, the edge $u'v'$ gives rise to a gated subgraph F'. In case $F_1 = F \cap W(u, v)$ would not equal $W(u, v)$, we could find a neighbour s of some vertex $t \in F_1$ outside F. Then $W(s, t) \subseteq W(u, v) - \{u\}$, contrary to the choice of $W(u, v)$. Therefore $G = F \cup W(v, u)$ and, analogously, $G = W(u', v') \cup F'$. If G has at most one pendant vertex, at least one of the disjoint sets F_1 and $F' \cap W(v', u')$ does not include any pendant vertex of G; say, the former. Then F_1 is a tree with at least two vertices. Any pendant vertex of this tree lies on a unique gated cycle of G, which is necessarily pendant in G. This concludes the proof.

5. Proof of Theorem 3

Two auxiliary results are established first.

Lemma 3. *Let u, v and w be vertices of a cellular bipartite graph G such that the intervals $I(u, v)$, $I(v, w)$, $I(w, u)$ intersect each other only in the common end vertices. Then the union*

$$C = I(u, v) \cup I(v, w) \cup I(w, u)$$

constitutes a gated cycle of G.

Proof. We already know from Theorem 1 that C is a convex subset. Suppose that x and y are adjacent vertices of C which do not together lie in any one of the three constituent intervals, say, $x \in I(u, v)$ and $y \in I(u, w)$, with $x, y \neq u$. Then, as G is bipartite, either x or y would belong to $I(u, v) \cap I(u, w)$, contrary to the hypothesis of the lemma. It follows that any shortest path between two vertices x and y from different constituent intervals $I(u, v)$, $I(u, w)$, $I(v, w)$ of C passes through either the common end vertex of the two intervals containing x and y or the other two end vertices. Therefore, any cycle formed by three shortest paths from u to v, from v to w, and w to u, respectively, is necessarily isometric and thus gated, whence it must coincide with C. □

The following result, being of independent interest, expresses that the gated cycles determine the system of all gated sets in a cellular bipartite graph (just as the 4-cycles do in the case of median graphs).

Proposition 1. *A connected subgraph A of a cellular bipartite graph G is gated iff every gated cycle of G intersecting A either is included in A or intersects A in a single vertex or edge.*

PROOF. Necessity is evident: if a gated cycle C intersects a gated subgraph A in at least two non-adjacent vertices, then C must entirely be included in A.

As to the converse, we will first show that A is convex. Suppose the contrary: then there are vertices v and x of A such that A does not include $I(v, x)$, where v and x are chosen so that their distance in the connected graph A is as small as possible. Then there exists a shortest path Q between v and x which is not included in A. Let P be any path of minimal length joining v and x within A, and let u be the neighbour of v on P. By the choice of v, x the paths P and Q only intersect in the end vertices. If P were longer than Q then, by virtue of the minimality assumption, both the subpath of P from u to x and the path composed of uv and Q would be shortest paths lying entirely in A. Therefore P must have the same length as Q. By the choice of v and x, the neighbour w of v on Q does not belong to A, and the intervals $I(u, x)$ and $I(w, x)$ have no other vertex in common than x. If x is adjacent to u and w, then $P \cup Q$ (comprising u, v, w, x) would be a gated 4-cycle. Otherwise, $I(u, w) \cap I(u, x) = \{u\}$ and $I(u, w) \cap I(w, x) = \{w\}$. Hence, by Lemma 3, $P \cup Q$ is a gated cycle (of length at least 6). Thus in either case, we obtain a gated cycle which properly intersects A in more than two vertices. Therefore A is indeed convex.

Finally, suppose that A is not gated. Choose a vertex z at minimum distance to A having no gate in A. Let x be a vertex of A closest to z, and let y be a vertex of A such that the interval $I(y, z)$ does not contain x, where $d(x, y)$ is as small as possible. Then the intervals $I(x, y)$, $I(y, z)$ and $I(z, x)$ intersect each other only in the common end vertices. By Lemma 3, the union of these intervals is a gated cycle. The intersection of this cycle with A is the path connecting x and y in A. Since x and y are not adjacent, we obtain a contradiction to the initial hypothesis. □

We next record an immediate consequence of Proposition 1 to be used below. We say that a graph G is a *convex amalgam* of two convex subgraphs G_1 and G_2 if $G = G_1 \cup G_2$ and $G_1 \cap G_2 \neq \emptyset$. Note that the convex amalgam of cellular bipartite graphs need not be cellular, since the obstruction constituted by three cycles, as displayed in Figure 1, can be generated from two cellular graphs by a single convex amalgamation (as long as one of the three cycles has length larger than 4).

COROLLARY 3. *Let G_1 and G_2 be convex subgraphs of a cellular bipartite graph G such that G is the convex amalgam of G_1 and G_2. If T_1 and T_2 are gated cutsets of G_1 and G_2, respectively, such that $G_1 \cap T_2 = G_2 \cap T_1 \neq \emptyset$, then $T = T_1 \cup T_2$ is a gated cutset of G.*

PROOF. T is evidently a connected cutset of G. Since every gated cycle in G is entirely included in either G_1 or G_2, we infer from Proposition 1 that T is gated. □

In order to accomplish the proof of Theorem 3 we will actually verify a stronger assertion. Note that G contains at least one edge belonging to two distinct gated cycles whenever G is 2-connected.

CLAIM. Every edge xy that belongs to at least two gated cycles extends to a gated tree cutset of G.

We proceed by induction on the number of vertices. If G has a cut vertex, then the claim is settled by virtue of the induction hypothesis. Therefore we can assume that G is 2-connected. If all gated cycles intersect in a common vertex, then the proof of Theorem 1 (Cases 2 and 3) shows that the given edge xy extends to a gated star cutset.

Hence (by Case 1 of that proof) we can find an edge uv (not necessarily distinct from xy) such that G is a gated amalgam of the subgraphs $G_1 = W(u, v) \cup F$ and $G_2 = W(v, u) \cup F$, where F consists of uv and the union of all gated cycles containing edges Θ-equivalent to uv (cf. Lemma 2). We may assume, without loss of generality that the edge xy is included in G_2 but not in $F_1 = W(u, v) \cap F$. Then, the gated cycles containing xy necessarily lie entirely in G_2, and thus, by the induction hypothesis, we can extend xy to a gated tree cutset T_2 of G_2. If T_2 is disjoint from F_1, then T_2 is also a gated cutset of G, as required. Therefore, assume that T_2 shares some vertex s with F_1. Then T_2 (being connected) must contain some vertex t of the cutset $F_2 = W(v, u) \cap F$ as well. Since T_2 is a gated tree in G_2, it follows that st is an edge constituting the intersection of T_2 with F. If s is a pendant vertex of F_1, then s has at most one neighbour in $G_2 - T_2$, whence $T_2 - \{s\}$ would be a gated tree cutset of G_2 avoiding F_1, and thus constituting a cutset of G. Otherwise, st lies on at least two gated cycles (from G_1). By the induction hypothesis, st is included in some gated tree cutset T_1 of G_1. Clearly, $T_1 \cap F = T_2 \cap F$ holds, and therefore $T = T_1 \cup T_2$ is the desired tree cutset of G (according to Corollary 3).

6. Quadratic Time Recognition

Polynomial time algorithms for testing cellularity could employ conditions (1), (2) and (5) of Theorem 1. A more efficient approach is based on the elimination scheme set up by Theorem 2. Quadratic time complexity can be guaranteed, since the total length of all gated cycles is bounded linearly in the number of vertices.

Lemma 4. *Let G be a cellular bipartite graph with $n \geqslant 3$ vertices, m edges, and g gated cycles. Then $m \leqslant 2n - 4$ and $g \leqslant n - 3$. The sum of lengths of all gated cycles does not exceed $4n - 12$.*

Proof. We already know that the first two inequalities must hold. To verify the bound on the total length of the gated cycles, proceed by induction. For $n = 3$ the assertion holds. So let $n \geqslant 4$. If G has a pendant vertex, the conclusion is trivial. Otherwise, select a pendant cycle C of length k and remove a maximal subpath consisting of $p \geqslant k/2 - 1$ vertices of degree 2 in G. We have then erased exactly one gated cycle having at most $2p + 2 \leqslant 4p$ edges. This concludes the induction. □

Now, the algorithm proceeds as follows. Given a graph G with $n \geqslant 3$ vertices, we first check whether the number of edges is at most $2n - 4$. If so, we continue to compute the distance matrix d of G in quadratic time, and further test whether G is bipartite. With the distance matrix in hand, we can compute for each pair u, v of vertices the number of neighbours of v in $I(u, v)$; this can be achieved in quadratic time. Then we start a recursion with $G_0 = G$ in order to dismantle G, thereby determining all gated cycles. Assume that the (isometric) subgraph G_i is under processing. Check whether G_i has a pendant vertex x. If so, put $G_{i+1} = G_i - \{x\}$ and continue. Otherwise, we search for maximal paths all vertices of which have degree 2 in G. This search can be organized so that no vertex is visited more than once. We accept such a path P with p vertices if, for the neighbours u and v in $G_i - P$ of the end vertices x and y of P, either $d(u, v) \leqslant p - 1$ (so that P is not a shortest path), or v has a second neighbour in $I(u, v)$ (that is, P is a shortest path but not the unique one). Select any shortest path Q joining u and v in G_i disjoint from P. Then P, Q together with the edges ux and vy constitute a cycle C. For each vertex v of G, determine a vertex v^* of C at minimum distance to v, and check whether v^* serves as the gate of v in C. If C is found and proven to be gated (otherwise, G would not be cellular), C enters the list of gated cycles, and G_{i+1} is set equal to $G_i - P$. This recursive step has a complexity of order n times the length of C. The recursion eventually stops when the number of vertices of G_{i+1} has reached 3. The total number of steps is of the order n times the total length of the gated cycles, and hence of quadratic order in view of Lemma 4.

Finally, for each vertex x_0 of a gated cycle C_1, check whether x_0 has a neighbour y_0 such that $x_0 y_0$ belongs to two distinct gated cycles C_2 and C_3, which share with C_1 the two neighbours y_3 and y_2, respectively, of x. Again this search requires a number of steps bounded by a quadratic polynomial in n (by Lemma 4).

Thus, with the help of the elimination scheme of Theorem 2, we can test condition (3) of Theorem 1 in quadratic time.

7. Median Vertices and Cycles

The cellular structure of a cellular bipartite graph G is also reflected by a median property for triplets of vertices. Recall that any three vertices u, v, w in a median graph admit a unique *median vertex* x (hence the name 'median graph'); that is,

$$I(u, v) \cap I(v, w) \cap I(w, u) = \{x\}.$$

For cellular bipartite graphs which have constituent gated cycles other than 4-cycles, gated cycles serve as substitutes for median vertices: we say that a gated cycle C of length k is a *median cycle* for a triplet u, v, w of vertices if the gates x, y, z of u, v, w, respectively, satisfy

$$\max\{d(x, y), d(y, z), d(z, x)\} < k/2,$$

$$x, y \in I(u, v), \qquad y, z \in I(v, w), \qquad z, x \in I(w, u)$$

(see Figure 3). The cycle C is the union of the three intervals $I(x, y)$, $I(y, z)$ and $I(z, x)$. It is determined by the pairwise intersections of the intervals $I(u, v)$, $I(v, w)$, $I(w, u)$, giving rise to the gates x, y, z first, according to the next proposition.

Proposition 2. *For any three vertices u, v, w of a cellular bipartite graph G, there exists a (necessarily unique) vertex x such that*

$$I(u, v) \cap I(u, w) = I(u, x).$$

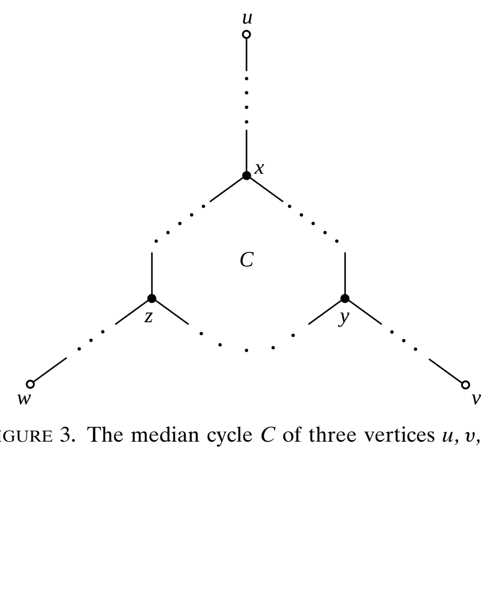

Figure 3. The median cycle C of three vertices u, v, w.

In other words, the vertex set of G is a semilattice with respect to the base-point order $\leqslant_u$ defined by $y \leqslant_u z$ iff $y \in I(u, z)$.

Proof. Suppose the contrary: then select a triplet u, v, w violating the assertion of the proposition such that $d(u, v) + d(u, w)$ is as small as possible. Then we can find two distinct vertices x and y such that both $I(u, x)$ and $I(u, y)$ are properly contained in $I(u, v) \cap I(u, w)$, while $I(x, v) \cap I(x, w) = \{x\}$ and $I(y, v) \cap I(y, w) = \{y\}$. We claim that

$$I(x, v) \cap I(x, y) = \{x\} \qquad \text{and} \qquad I(y, v) \cap I(y, x) = \{y\}.$$

Indeed, if the first equality were not true, then we could find a vertex z in $I(x, v)$ different from x, which would also belong to $I(x, y) \subseteq I(u, w)$ (as intervals are convex in G) and hence to $I(x, w)$, thus contradicting the choice of x. Similarly, we obtain the second equality and the analogous pair of equalities

$$I(x, w) \cap I(x, y) = \{x\} \qquad \text{and} \qquad I(y, w) \cap I(y, x) = \{y\}.$$

The minimality of $d(u, v) + d(u, w)$ guarantees that

$$I(v, x) \cap I(v, y) = \{v\} \qquad \text{and} \qquad I(w, x) \cap I(w, y) = \{w\}.$$

In view of the above six equalities, we infer from Lemma 3 that the unions $I(v, x) \cup I(x, y) \cup I(y, v)$ and $I(w, x) \cup I(x, y) \cup I(y, w)$ constitute two distinct gated cycles, which intersect in two non-adjacent vertices, viz. x and y, thus yielding the desired contradiction. □

Proposition 3. *Every triplet u, v, w of vertices of a cellular bipartite graph G admits either a unique median vertex of a unique median cycle.*

Proof. Let x, y, z be the vertices successively determined by

$$I(u, v) \cap I(u, w) = I(u, x), \quad I(v, x) \cap I(v, w) = I(v, y), \quad I(w, x) \cap I(w, y) = I(w, z),$$

according to Proposition 2. If $x = y = z$, then this vertex is the unique median vertex. Otherwise, the three vertices are different, and the intervals $I(x, y)$, $I(y, z)$ and $I(x, z)$ pairwise intersect only in the common end vertices. Thus, by Lemma 3, $C = I(x, y) \cup I(y, z) \cup I(z, x)$ is a gated cycle with $x, y \in I(u, v)$, $y, z \in I(v, w)$, and $z, x \in I(w, u)$, as required.

In order to prove uniqueness of the median cycle, we first show that

$$I(u, v) \cap I(v, w) = I(v, y) \qquad \text{and} \qquad I(u, w) \cap I(v, w) = I(w, z).$$

As to the former, suppose that there exists a neighbour t of y in $I(u, y) \cap I(w, y)$. Then u and w belong to the convex subgraph $W(t, y)$. Since $x, z \in I(u, w) \subseteq W(t, y)$, it follows that $t \in I(x, y) \cap I(y, z) = \{y\}$, yielding a contradiction. This settles the first equality, and the second one is proved analogously. Finally, assume that C' is any median cycle for u, v, w. Let x', y', z' be the gates of u, v, w, respectively, in C'. Then, necessarily,

$$C' = I(x', y') \cup I(y', z') \cup I(z', x')$$

contains x, y, z and thus includes C; that is, $C' = C$. □

It would be interesting to characterize the bipartite graphs in which every vertex triplet admits a unique median vertex or cycle.

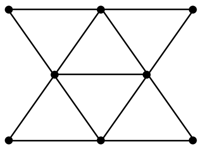

FIGURE 4. A non-bipartite graph with totally decomposable metric.

8. CONCLUSIONS

We have established a number of characterizations and features of cellular bipartite graphs, the most important of which are the decomposition along gated trees and the elimination scheme of pendant edges and pendant gated cycles. In view of this, one may expect that the class of cellular bipartite graphs has many nice algorithmic properties.

For a cellular bipartite graph G with metric d, the isolation indices of the splits $S(u, v)$ separating edges uv equal 1, so that the canonical decomposition of d into split metrics immediately yields an isometric embedding of G into a hypercube. If the requirement that G be bipartite is dropped, then total decomposability of d ensures that all positive isolation indices equal $\frac{1}{2}$ or 1, so that $2d$ would be a restriction of a hypercube metric, i.e. G would be 'scale 2 embeddable' [13] into a hypercube.

Each condition of Theorem 1 lends itself to a generalization of cellularity to the non-bipartite case. The feasible amalgamations would then be those performed along special convex sets—but not necessarily gated sets, if one wishes to capture graphs such as the one displayed in Figure 4.

Another possible generalization of cellular bipartite graphs would incorporate the median graphs as well by departing from cycles and hypercubes and using gated amalgamations.

ACKNOWLEDGEMENTS

The research of the second author, who is on leave from the Universitatea de stat din Moldova, Chişinău, was supported by the Alexander von Humboldt Stiftung.

REFERENCES

1. H.-J. Bandelt, Hereditary modular graphs, *Combinatorica,* **8** (1988), 149–157.
2. H.-J. Bandelt, Graphs with intrinsic S_3 convexities, *J. Graph Theory,* **13** (1989), 215–218.
3. H.-J. Bandelt and J. P. Barthélemy, Medians in median graphs, *Discr. Appl. Math.,* **8** (1984), 131–142.
4. H.-J. Bandelt and A. W. M. Dress, A canonical decomposition theory for metrics on a finite set, *Adv. Math.,* **92** (1992), 47–105.
5. H.-J. Bandelt and A. W. M. Dress, Split decomposition: a new and useful approach to phylogenetic analysis of distance data, *Mol. Phylogenet. Evol.,* **1** (1992), 242–252.
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, H. M. Mulder and E. Wilkeit, Quasi-median graphs and algebras, *J. Graph Theory,* **18** (1994), 681–703.
8. V. Chepoi, Geometric properties of d-convexity in bipartite graphs (in Russian), *Modelirovanie Informacionnych Sistem,* Chişinău, Stiinţa (1986), 88–100.
9. M. Deza and M. Laurent, l_1-rigid graphs, *J. Alg. Combin.,* **3** (1994), 153–175.
10. M. Deza and M. Laurent, Applications of cut polyhedra, *J. Comp. Appl. Math.* (to appear).
11. D. Ž. Djoković, Distance-preserving subgraphs of hypercubes, *J. Combin. Theory, Ser. B,* **14** (1973), 263–267.
12. A. W. M. Dress and R. Scharlau, Gated sets in metric spaces, *Aeq. Math.,* **34** (1987), 112–120.

13. S. V. Shpectorov, On scale embeddings of graphs into hypercubes, *Europ. J. Combin.*, **14** (1993), 117–130.
14. M. L. J. van de Vel, *Theory of Convex Structures.* North-Holland, Amsterdam, 1993.

Received 20 March 1995 and accepted 20 August 1995

HANS-JÜRGEN BANDELT AND VICTOR CHEPOI
Mathematisches Seminar,
Universität Hamburg,
Bundesstr. 55,
D-20146 Hamburg, Germany