

Navigating in a Graph by Aid of Its Spanning Tree^{*}

Feodor F. Dragan¹ and Martin Matamala²

¹ Algorithmic Research Laboratory, Department of Computer Science
Kent State University, Kent, OH 44242, USA^{**}

`dragan@cs.kent.edu`

² Departamento de Ingeniería Matemática, Universidad de Chile
Centro de Modelamiento Matemático UMR 2071CNRS, Santiago, Chile

`mmatamal@dim.uchile.cl`

Abstract. Let $G = (V, E)$ be a graph and T be a spanning tree of G . We consider the following strategy in advancing in G from a vertex x towards a target vertex y : from a current vertex z (initially, $z = x$), unless $z = y$, go to a neighbor of z in G that is closest to y in T (breaking ties arbitrarily). In this strategy, each vertex has full knowledge of its neighborhood in G and can use the distances in T to navigate in G . Thus, additionally to standard local information (the neighborhood $N_G(v)$), the only global information that is available to each vertex v is the topology of the spanning tree T (in fact, v can know only a very small piece of information about T and still be able to infer from it the necessary tree-distances). For each source vertex x and target vertex y , this way, a path, called a greedy routing path, is produced. Denote by $g_{G,T}(x, y)$ the length of a longest greedy routing path that can be produced for x and y using this strategy and T . We say that a spanning tree T of a graph G is an *additive r -carcass* for G if $g_{G,T}(x, y) \leq d_G(x, y) + r$ for each ordered pair $x, y \in V$. In this paper, we investigate the problem, given a graph family \mathcal{F} , whether a small integer r exists such that any graph $G \in \mathcal{F}$ admits an additive r -carcass. We show that rectilinear $p \times q$ grids, hypercubes, distance-hereditary graphs, dually chordal graphs (and, therefore, strongly chordal graphs and interval graphs), all admit additive 0-carcasses. Furthermore, every chordal graph G admits an additive $(\omega(G) + 1)$ -carcass (where $\omega(G)$ is the size of a maximum clique of G), each 3-sun-free chordal graph admits an additive 2-carcass, each chordal bipartite graph admits an additive 4-carcass. In particular, any k -tree admits an additive $(k + 2)$ -carcass. All those carcasses are easy to construct.

1 Introduction

As part of the recent surge of interest in different kind of networks, there has been active research exploring strategies for navigating synthetic and real-world

^{*} This work was partially supported by CONICYT through grants Anillo en Redes ACT08 and Fondap.

^{**} These results were obtained while the first author was visiting the Universidad de Chile, Santiago.

networks (modeled usually as graphs). These strategies specify some rules to be used to advance in a graph (a network) from a given vertex towards a target vertex along a path that is close to shortest. Current strategies include (but not limited to): routing using full-tables, interval routing, routing labeling schemes, greedy routing, geographic routing, compass routing, etc. in wired or wireless communication networks and in transportation networks (see [19,20,27,33,25,40] and papers cited therein); routing through common membership in groups, popularity, and geographic proximity in social networks and e-mail networks (see [2,3,27,32] and literature cited therein).

In this paper we use terminology used mostly for communication networks. Thus, navigation is performed using a routing scheme, i.e., a mechanism that can deliver packets of information from any vertex of a network to any other vertex. In most strategies, each vertex v of a graph has full knowledge of its neighborhood and uses a piece of global information available to it about the graph topology – some "sense of direction" to each destination, stored locally at v . Based only on this information and the address of a destination vertex, vertex v needs to decide whether the packet has reached its destination, and if not, to which neighbor of v to forward the packet.

1.1 Some Known Strategies

In *routing using full-tables*, each vertex v of G knows for each destination u the first edge along some shortest path from v to u (so-called *complete routing table*). When v needs to send a message to u , it just sends the message along the edge stored for destination u . While this approach guarantees routing along a shortest path, it is too expensive for large systems since it requires to store locally $O(n \log \delta)$ bits of global information for an n -vertex graph with maximum degree δ .

Unfortunately, if one insists on a routing via shortest paths, $\Omega(n \log \delta)$ bits is the lower bound on the memory requirements per vertex [24] (this much each vertex needs to know at least). To obtain routing schemes for general graphs that use $o(n)$ of memory at each vertex, one has to abandon the requirement that packets are always delivered via shortest paths, and settle instead for the requirement that packets are routed on paths that are relatively close to shortest. The efficiency of a routing scheme is measured in terms of its additive stretch, called *deviation* (or multiplicative stretch, called *delay*), namely, the maximum surplus (or ratio) between the length of a route, produced by the scheme for a pair of vertices, and the shortest route. There is a tradeoff between the memory requirements of a routing scheme (how much of global information is available locally at a vertex) and the worst case stretch factor it guarantees. Any multiplicative t -stretched routing scheme must use $\Omega(n)$ bits for some vertices in some graphs for $t < 3$ [21] (see also [17]), and $\Omega(n \log n)$ bits for $t < 1.4$ [24]. These lower bounds show that it is not possible to lower memory requirements of a routing scheme for an arbitrary network if it is desirable to route messages along paths close to optimal. Therefore, it is interesting, both from a theoretical and a practical view point, to look for specific routing strategies on graph families with certain topological properties.

One specific way of routing, called *interval routing*, has been introduced in [36] and later generalized in [31]. In this method, the complete routing tables are compressed by grouping the destination addresses which correspond to the same output edge. Then each group is encoded as an interval, so that it is easy to check whether a destination address belongs to the group. This approach requires $O(\delta \log n)$ bits of memory per vertex, where δ is the maximum degree of a vertex of the graph. A graph must satisfy some topological properties in order to support interval routing, especially if one insists on paths close to optimal. Routing schemes for many graph classes were obtained by using interval routing techniques. The classical and most recent results in this field are presented in [19,20].

Recently, so-called *routing labeling schemes* [33] become very popular. A number of interesting results for general graphs and particular classes of graphs were obtained. These are schemes that label the vertices of a graph with short labels (describing some global topology information) in such a way that given the label of a source vertex and the label of a destination, it is possible to compute efficiently the edge from the source that heads in the direction of the destination. In [18,40], a shortest path routing scheme for trees with $O(\log^2 n / \log \log n)$ -bit labels is described. For general graphs, the most general result to date is a multiplicative $(4k - 5)$ -stretched routing labeling scheme that uses labels of size $\tilde{O}(kn^{1/k})$ bits¹ is obtained in [40] for every $k \geq 2$. For planar graphs, a shortest path routing labeling scheme which uses $8n + o(n)$ bits per vertex is developed in [22], and a multiplicative $(1 + \epsilon)$ -stretched routing labeling scheme for every $\epsilon > 0$ which uses $O(\epsilon^{-1} \log^3 n)$ bits per vertex is developed in [39]. Routing in graphs with doubling dimension α has been considered in [1,10,37,38]. It was shown that any graph with doubling dimension α admits a multiplicative $(1 + \epsilon)$ -stretched routing labeling scheme with labels of size $\epsilon^{-O(\alpha)} \log^2 n$ bits. Recently, the routing result for trees of [18,40] was used in designing additive $O(1)$ -stretched routing labeling schemes with $O(\log^{O(1)} n)$ bit labels for several families of graphs, including chordal graphs, chordal bipartite graphs, circular-arc graphs, AT-free graphs and their generalizations, the graphs with bounded longest induced cycle, the graphs of bounded tree-length, the bounded clique-width graphs, etc. (see [12,13,14,15] and papers cited therein).

In wireless networks, the most popular strategy is the *geographic routing* (sometimes called also the *greedy geographic routing*), where each vertex forwards the packet to the neighbor geographically closest to the destination (see survey [25]). Each vertex of the network knows its position (e.g., Euclidean coordinates) in the underlying physical space and forwards messages according to the coordinates of the destination and the coordinates of neighbors. Although this greedy method is effective in many cases, packets may get routed to where no neighbor is closer to the destination than the current vertex. Many recovery schemes have been proposed to route around such voids for guaranteed packet delivery as long as a path exists [4,26,30]. These techniques typically exploit planar subgraphs (e.g., Gabriel graph, Relative Neighborhood graph), and packets traverse faces on such graphs using the well-known right-hand rule.

¹ Here, $\tilde{O}(f)$ means $O(f \text{ polylog } n)$.

All earlier papers assumed that vertices are aware of their physical location, an assumption which is often violated in practice for various of reasons (see [16,28,35]). In addition, implementations of recovery schemes are either based on non-rigorous heuristics or on complicated planarization procedures. To overcome these shortcomings, recent papers [16,28,35] propose routing algorithms which assign virtual coordinates to vertices in a metric space X and forward messages using geographic routing in X . In [35], the metric space is the Euclidean plane, and virtual coordinates are assigned using a distributed version of Tutte's "rubber band" algorithm for finding convex embeddings of graphs. In [16], the graph is embedded in R^d for some value of d much smaller than the network size, by identifying d beacon vertices and representing each vertex by the vector of distances to those beacons. The distance function on R^d used in [16] is a modification of the ℓ_1 norm. Both [16] and [35] provide substantial experimental support for the efficacy of their proposed embedding techniques – both algorithms are successful in finding a route from the source to the destination more than 95% of the time – but neither of them has a provable guarantee. Unlike embeddings of [16] and [35], the embedding of [28] guarantees that the geographic routing will always be successful in finding a route to the destination, if such a route exists. Algorithm of [28] assigns to each vertex of the network a virtual coordinate in the hyperbolic plane, and performs greedy geographic routing with respect to these virtual coordinates. More precisely, [28] gets virtual coordinates for vertices of a graph G by embedding in the hyperbolic plane a spanning tree of G . The proof that this method guarantees delivery is relied only on the fact that the hyperbolic greedy route is no longer than the spanning tree route between two vertices; even more, it could be much shorter as greedy routes take enough short cuts (edges which are not in the spanning tree) to achieve significant saving in stretch. However, although the experimental results of [28] confirm that the greedy hyperbolic embedding yields routes with low stretch when applied to typical unit-disk graphs, the worst-case stretch is still linear in the network size.

1.2 Our Approach

Motivated by the work of Robert Kleinberg [28], in this paper, we initiate exploration of the following strategy in advancing in a graph from a source vertex towards a target vertex. Let $G = (V, E)$ be a graph and T be a spanning tree of G . To route/move in G from a vertex x towards a target vertex y , use the following rule:

*from a current vertex z (initially, $z = x$), unless $z = y$,
go to a neighbor of z in G that is closest to y in T
(break ties arbitrarily).*

In this strategy, each vertex has full knowledge of its neighborhood in G and can use the distances in T to navigate in G . Thus, additionally to standard local information (the neighborhood $N_G(v)$), the only global information that is available to each vertex v is the topology of the spanning tree T . In fact, v

can know only a very small piece of information about T and still be able to infer from it the necessary tree-distances. It is known [23,34] that the vertices of an n -vertex tree T can be labeled in $O(n \log n)$ total time with labels of up to $O(\log^2 n)$ bits such that given the labels of two vertices v, u of T , it is possible to compute in constant time the distance $d_T(v, u)$, by merely inspecting the labels of u and v . Hence, one may assume that each vertex v of G knows, additionally to its neighborhood in G , only its $O(\log^2 n)$ bit distance label. This distance label can be viewed as a virtual coordinate of v .

For each source vertex x and target vertex y , by this routing strategy, a path, called a *greedy routing path*, is produced (clearly, this routing strategy will always be successful in finding a route to the destination). Denote by $g_{G,T}(x, y)$ the length of a longest greedy routing path that can be produced for x and y using this strategy and T . We say that a spanning tree T of a graph G is an *additive r -carcass* for G if $g_{G,T}(x, y) \leq d_G(x, y) + r$ for each ordered pair $x, y \in V$ (in a similar way one can define also a *multiplicative t -carcass* of G).

In this paper, we start investigating the problem, given a graph family \mathcal{F} , whether a small integer r exists such that any graph $G \in \mathcal{F}$ admits an additive r -carcass, and give our preliminary results. We show that rectilinear $p \times q$ grids, hypercubes, distance-hereditary graphs, dually chordal graphs (and, therefore, strongly chordal graphs and interval graphs), all admit additive 0-carcasses. Furthermore, every chordal graph G admits an additive $(\omega(G)+1)$ -carcass (where $\omega(G)$ is the size of a maximum clique of G), each 3-sun-free chordal graph admits an additive 2-carcass, each chordal bipartite graph admits an additive 4-carcass. In particular, any k -tree admits an additive $(k + 2)$ -carcass. All those carcasses are easy to construct.

2 Preliminaries

All graphs occurring in this paper are connected, finite, undirected, unweighted, loopless and without multiple edges. In a graph $G = (V, E)$ ($n = |V|, m = |E|$) the *length* of a path from a vertex v to a vertex u is the number of edges in the path. The *distance* $d_G(u, v)$ between the vertices u and v is the length of a shortest path connecting u and v . The *neighborhood* of a vertex v of G is the set $N_G(v) = \{u \in V : uv \in E\}$ and the *closed neighborhood* of v is $N_G[v] = N_G(v) \cup \{v\}$. The *disk* of radius k centered at v is the set of all vertices at distance at most k to v , i.e., $D_k(v) = \{u \in V : d_G(u, v) \leq k\}$. A set $S \subseteq V$ is a *clique* (an *independent set*) of G if all vertices of S are pairwise adjacent (respectively, nonadjacent) in G . A clique of G is *maximal* if it is not contained in any other clique of G .

Next we recall the definitions of special graph classes mentioned in this paper (see survey [8]). A graph is *chordal* if it does not have any induced cycle of length greater than 3. A *p -sun* ($p \geq 3$) is a chordal graph on $2p$ vertices whose vertex set can be partitioned into two sets, $U = \{u_0, \dots, u_{p-1}\}$ and $W = \{w_0, \dots, w_{p-1}\}$, such that W is an independent set, U is a clique, and every w_i is adjacent only to u_i and $u_{i+1} \pmod p$. A chordal graph having no induced subgraphs isomorphic

to p -suns (for any $p \geq 3$) is called a *strongly chordal graph*. A chordal graph having no induced subgraphs isomorphic to 3-sun is called a *3-sun-free chordal graph*. A graph is *chordal bipartite* if it is bipartite and has no induced cycles of length greater than 4. A *dually chordal graph* is the intersection graph of the maximal cliques of a chordal graph (see [8,7] for many equivalent definitions of dually chordal graphs and strongly chordal graphs). A graph is *interval* if it is the intersection graph of intervals of a line. It is known that interval graphs are strongly chordal and strongly chordal graphs are dually chordal (see [8,7]). A graph G is *distance-hereditary* if every induced path of G is shortest (see [8] for many equivalent definitions of distance-hereditary graphs). The *k -trees* are defined recursively: a clique of size k (denoted by K_k) is a k -tree; if G is a k -tree, then a graph obtained from G by adding a new vertex v adjacent to all vertices of some clique K_k of G is a k -tree. It is known (see [8]) that all k -trees are chordal graphs and that maximal cliques of a k -tree have size at most $k + 1$.

Let now $G = (V, E)$ be a graph and T be a spanning tree of G . In what follows, we will use the following notations. For vertices v and u from V , denote by vTu the (unique) path of T connecting vertices v and u . For a source vertex x and a target vertex y in G , denote by $R_{G,T}(x, y)$ a greedy routing path obtained for x and y by using tree T and the strategy described in Subsection 1.2. Clearly, for the same pair of vertices x and y , breaking ties differently, different greedy routing paths $R_{G,T}(x, y)$ can be produced. Denote, as before, by $g_{G,T}(x, y)$, the length of a longest greedy routing path that can be produced for x and y . If no confusion can arise, we will omit indexes G and T , i.e., use $R(x, y)$ and $g(x, y)$ instead of $R_{G,T}(x, y)$ and $g_{G,T}(x, y)$.

For $r \geq 0$ and $t \geq 1$, a spanning tree T of a graph G is called an *additive r -carcass* (a *multiplicative t -carcass*) for G if $g_{G,T}(x, y) \leq d_G(x, y) + r$ (respectively, $g_{G,T}(x, y) \leq t d_G(x, y)$) for each ordered pair $x, y \in V$.

Let x^* be the neighbor of x in $R_{G,T}(x, y)$ and x' be the neighbor of x in xTy . Since both x^* and x' are in $N_G(x)$ and $d_T(x', y) = d_T(x, y) - 1$, according to our strategy $d_T(x^*, y) \leq d_T(x', y) = d_T(x, y) - 1$ must hold. Furthermore, any subpath of a greedy routing path $R_{G,T}(x, y)$ containing y is a greedy routing path to y as well. Hence, one can conclude, by induction, that the length of any greedy routing path $R_{G,T}(x, y)$ never exceeds $d_T(x, y)$. It is clear also, that a greedy routing path $R_{G,T}(x, y) := (x := x_0, x_1, x_2, \dots, y := x_\ell)$ cannot have a chord $x_i x_j \in E$ with $j > i + 1$ (since $d_T(x_{i+1}, y) > d_T(x_j, y)$), i.e., any greedy routing path is an induced path. Thus, we have the following.

Observation 1. *Let G be an arbitrary graph and T be its arbitrary spanning tree. Then, for any vertices x, y of G ,*

- (a) $g_{G,T}(x, y) \leq d_T(x, y)$,
- (b) any greedy routing path $R_{G,T}(x, y)$ is an induced path of G ,
- (c) a tale of any greedy routing path is a greedy routing path.

Since in distance-hereditary graphs each induced path is a shortest path, by Observation 1(b), we conclude.

Corollary 1. *Any spanning tree of a distance-hereditary graph G is a 0-carcaass of G .*

There are well-known notions of additive tree r -spanners and multiplicative tree t -spanners. For $r \geq 0$ and $t \geq 1$, a spanning tree T of a graph G is called an *additive tree r -spanner* (a *multiplicative tree t -spanner*) of G if $d_T(x, y) \leq d_G(x, y) + r$ (respectively, $d_T(x, y) \leq t d_G(x, y)$) for each pair $x, y \in V$ [9]. By Observation 1(a), we obtain.

Corollary 2. *Any additive tree r -spanner (multiplicative tree t -spanner) of a graph G is an additive r -carcaass (multiplicative t -carcaass) of G .*

Note that the converse of Corollary 2 is not generally true. As we will see in next sections, there are many families of graphs which do not admit any tree r -spanners (additive as well as multiplicative) for any constant r , yet they admit very good carcaasses. For example, there is no constant r such that any 2-tree or any chordal bipartite graph has a tree r -spanner (additive or multiplicative), but both these families of graphs admit additive 4-carcaasses (see Section 5 for details).

In what follows, in a rooted tree T , by $f(v)$ we will denote the father of a vertex v .

3 Rectilinear Grids and Hypercubes

In this section we show that the rectilinear grids and the hypercubes admit additive 0-carcaasses.

Consider a rectilinear $p \times q$ grid G and assume that it is naturally embedded into the plane such that all inner faces of G are squares (see Fig. 1). First we notice that G does not admit any good tree spanner. For this, consider an arbitrary spanning tree T of G , and assume that p and q are odd integers and $p \leq q$. Since T is a planar graph with only the outer face, we can connect by a Jordan curve \mathcal{C} a point of the plane inside the central square of G with a point




Fig. 1. Rectilinear grids do not admit any (additive or multiplicative) tree r -spanners with a constant r , but have additive 0-carcaasses

in the outer face of G without intersecting the tree T . Let R be the first square of G crossed by \mathcal{C} and x and y be two opposite vertices of R (see Fig. 1(a) for an illustration). Clearly, for x and y , $d_T(x, y) \geq p + 1$ holds, while $d_G(x, y) = 2$. Here, we considered nonadjacent vertices of G since adjacent vertices are of no interest in our greedy routing. Thus, there are no good tree spanners for rectilinear grids. On the other hand, G admits an additive 0-carcaass. Consider a Hamiltonian path of G depicted on Fig. 1(b), called *column-wise Hamiltonian path*. This path is an additive 0-carcaass of G . We leave verification of this fact to the reader.

Now we turn to the hypercubes. Let $H_q = (V, E)$ be the q -dimensional hypercube whose vertices are binary words of length q and two vertices are adjacent if they differ in exactly one letter. Let $a \in \{0, 1\}$ and $i \in \{1, \dots, q\}$. Let $H_q^{a,i}$ be a subgraph of H_q induced by vertices having letter a in the position i . Then, $H' := H_q^{a,i}$ is isomorphic to the $(q - 1)$ -dimensional hypercube and $d_{H_q}(x, y) = d_{H'}(x, y)$, whenever the letter in the position i of x and y is a .

Let T be the *Gray-Hamiltonian* path of H_q defined recursively as follows. If $x_1, \dots, x_{2^{q-1}}$ is the Gray-Hamiltonian path for H_{q-1} , then T is given by T_0T_1 , where $T_0 = x_10, \dots, x_{2^{q-1}}0$ and $T_1 = x_{2^{q-1}}1, \dots, x_11$. By applying two steps of previous recursion, it is clear that T can be decomposed into four consecutive subpaths $T = T_{00}T_{10}T_{11}T_{01}$, where the subpath T_w contains all the vertices of H_q ending with w . Notice that T_0 is the Gray-Hamiltonian path for $H_q^{0,q}$, T_1 is the reverse of the Gray-Hamiltonian path for $H_q^{1,q}$ and $T_{10}T_{11}$ is the Gray-Hamiltonian path for the hypercube $H_q^{1,q-1}$.

By using induction on q , we prove that $g(x, y) := g_{G,T}(x, y) = d_G(x, y)$, where $G = H_q$ and T is the Gray-Hamiltonian path of H_q . When x and y belong to T_0 (resp. T_1), conclusion is obtained by applying induction hypothesis to vertices x and y in the hypercube $H_q^{0,q}$ (resp. $H_q^{1,q}$) with the Gray-Hamiltonian path T_0 (resp. T_1). Similarly, when x belongs to T_{10} and y belongs to T_{11} , we can apply induction in the hypercube $H_q^{1,q-1}$ with the Gray-Hamiltonian path $T_{10}T_{11}$.

For the remaining cases, let x^* denote the vertex next to x in a greedy routing path $R_{G,T}(x, y)$. If x^* and y belong to T_1 , and x belongs to T_0 , then $d_G(x^*, y) = d_G(x, y) - 1$, since x^* and y agree in their last letters. By applying induction to x^* and y using T_1 , we get that $g(x^*, y) = d_{H_q^{1,q}}(x^*, y) = d_G(x^*, y)$. Hence, $g(x, y) = 1 + g(x^*, y) = 1 + d_G(x^*, y) = d_G(x, y)$. Let us now consider the case when x and x^* belong to T_0 , and y belongs to T_1 . As x has a neighbor in T_1 and we are assuming that x^* belongs to T_0 , vertex y must belong to T_{11} . We have already considered the case when x belongs to T_{10} . Hence, let us assume that x belongs to T_{00} . As x has a neighbor in T_{10} , vertex x^* must belong to T_{10} . Since in this case $d_G(x^*, y) = d_G(x, y) - 1$, we can conclude as before, by applying induction to x^* and y in $H_q^{1,q-1}$ with the Gray-Hamiltonian path $T_{10}T_{11}$.

So, we can state the following theorem.

Theorem 1. *Every rectilinear grid and every hypercube admits an additive 0-carcaass (which is a Hamiltonian path) constructible in linear time.*

4 Locally Connected Spanning Trees Are Additive 0-Carcasses: Dually Chordal Graphs

In this section, we show that every dually chordal graph $G = (V, E)$ admits an additive 0-carcass constructible in linear time. Recall that every dually chordal graph has an additive tree 3-spanner and there are dually chordal graphs without any additive tree 2-spanners (see [6]). Clearly, those additive tree 3-spanners are additive 3-carcasses, but it is not hard to see that they are not necessarily additive 0-carcasses (see, e.g., Fig. 2).




Fig. 2. A dually chordal graph with an additive tree 3-spanner (on the left) and an additive 0-carcass (on the right). This dually chordal graph does not have any additive tree 2-spanner. A greedy routing path from x to y with respect to the corresponding tree is shown on both pictures.

Let G be a graph. We say that a spanning tree T of G is *locally connected* if the closed neighborhood $N_G[v]$ of any vertex v of G induces a subtree in T (i.e., $T \cap N_G[v]$ is a connected subgraph of T). See the right picture on Fig. 2 for an example of a locally connected spanning tree.

Theorem 2. *If T is a locally connected spanning tree of a graph G , then T is an additive 0-carcass of G .*

Proof. Assume that we want to route from a source vertex x to a target vertex y in G . Let v be an arbitrary vertex of G and v^* be a vertex from $N_G[v]$ closest to y in T . Since $T \cap N_G[v]$ is a connected subgraph of T , for each vertex $v \in V$ such a neighbor v^* is unique (any subtree of a tree has only one vertex closest in T to a given vertex y). Moreover, $v^* \neq v$, unless $v = y$. In what follows, we will assume that the tree T is rooted at vertex y .

Claim 1. *For any vertex $v \in V$, the vertex v^* belongs to a shortest path of G connecting v and y .*

Proof. We prove by induction on $d_G(v, y)$. If $d_G(v, y) \leq 1$, then $v^* = y$ and therefore v^* belongs to any shortest path between v and y . So, assume that $d_G(v, y) \geq 2$. Consider a shortest path $P(v, y) := (v, a, b, \dots, y)$ in G connecting v and y , where a and b are the first two (after v) vertices of this path. They exist since $d_G(v, y) \geq 2$.

To obtain the conclusion, we prove that v^* and b are adjacent. For sake of contradiction, let us assume that they are not adjacent. Since $a, b^* \in N_G(b)$ and $T \cap N_G[b]$ is a connected subgraph of T , we get that v^* is not on path aTv^* . We also know that $a \in N_G[v] \cap N_G[b]$ and therefore both v^* and b^* are ancestors in T of a (recall that we have rooted T at y). Hence, b^* is on the path aTv^* . As $a, v^* \in N_G(v)$ and $T \cap N_G[v]$ is a connected subgraph of T , we get that v and b^* are adjacent. By induction, b^* belongs to a shortest path of G between b and y , which leads to the following contradiction: $d_G(v, y) \leq 1 + d_G(b^*, y) = d_G(b, y) < d_G(v, y)$. \square (of Claim)

Now we prove by induction on $d_T(x, y)$ that $g(x, y) = d_G(x, y)$. Indeed, $g(x, y) = 1 + g(x^*, y)$ and, by induction, $g(x^*, y) = d_G(x^*, y)$ as $d_T(x, y) > d_T(x^*, y)$. By Claim 1, we conclude $g(x, y) = 1 + d_G(x^*, y) = d_G(x, y)$. \square

It has been shown in [7] that the graphs admitting locally connected spanning trees are precisely the dually chordal graphs. Furthermore, [7] showed that the class of dually chordal graphs contains such known families of graphs as strongly chordal graphs, interval graphs and others. Thus, we have the following corollary.

Corollary 3. *Every dually chordal graph admits an additive 0-carcass constructible in linear time. In particular, any strongly chordal graph (any interval graph) admits an additive 0-carcass constructible in linear time.*

Note that, in [5,7], it was shown that dually chordal graphs can be recognized in linear time, and if a graph G is dually chordal, then a locally connected spanning tree of G can be efficiently constructed.

5 Additive Carcasses for Chordal Graphs and Chordal Bipartite Graphs

In this section, we just list our results for chordal graphs and chordal bipartite graph. The proofs can be found in the journal version of this paper.

Theorem 3. *Every chordal bipartite graph admits an additive 4-carcass constructible in linear time.*

Recall that chordal bipartite graphs do not have any tree r -spanners (additive or multiplicative) with a constant r (see, e.g., [11]).

Theorem 4. *Any shortest path tree of a chordal graph G is an additive $(\omega(G) + 1)$ -carcass of G . Here $\omega(G)$ is the size of a maximum clique of G .*

Since k -trees are chordal graphs with the size of a maximum clique at most $k + 1$, we conclude.

Corollary 4. *Every k -tree admits an additive $(k + 2)$ -carcass constructible in linear time. In particular, any 2-tree admits an additive 4-carcass constructible in linear time.*

Recall that 2-trees do not have any tree r -spanners (additive or multiplicative) with a constant r (see, e.g., [29]). We also have the following result.

Corollary 5. *Every 3-sun-free chordal graph admits an additive 2-carcass.*

References

1. Abraham, I., Gavoille, C., Goldberg, A.V., Malkhi, D.: Routing in Networks with Low Doubling Dimension. In: ICDCS 2006, p. 75 (2006)
2. Adamic, L.A., Lukose, R.M., Huberman, B.A.: Local Search in Unstructured Networks. Wiley, New-York (2002)
3. Adamic, L.A., Lucose, R.M., Puniyani, A.R., Huberman, B.A.: Search in power-law networks. *Physical Review E* 64, 046135, 1–8 (2001)
4. Bose, P., Morin, P., Stojmenovic, I., Urrutia, J.: Routing with guaranteed delivery in ad hoc wireless networks. In: Proceedings of the 3rd International Workshop on Discrete algorithms and Methods for Mobile Computing and Communications, pp. 48–55. ACM Press, New York (1999)
5. Brandstädt, A., Chepoi, V.D., Dragan, F.F.: The algorithmic use of hypertree structure and maximum neighbourhood orderings. *Discrete Appl. Math.* 82, 43–77 (1998)
6. Brandstädt, A., Chepoi, V.D., Dragan, F.F.: Distance approximating trees for chordal and dually chordal graphs. *Journal of Algorithms* 30, 166–184 (1999)
7. Brandstädt, A., Dragan, F.F., Chepoi, V.D., Voloshin, V.I.: Dually chordal graphs. *SIAM J. Discrete Math.* 11, 437–455 (1998)
8. Brandstädt, A., Le, V., Bang, S.J.P.: Graph Classes: A Survey. *SIAM Monographs on Discrete Mathematics and Applications*, Philadelphia (1999)
9. Cai, L., Corneil, D.G.: Tree spanners. *SIAM J. Disc. Math.* 8, 359–387 (1995)
10. Chan, H.T.-H., Gupta, A., Maggs, B.M., Zhou, S.: On hierarchical routing in doubling metrics. In: Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2005), pp. 762–771. SIAM, Philadelphia (2005)
11. Chepoi, V.D., Dragan, F.F., Yan, C.: Additive Sparse Spanners for Graphs with Bounded Length of Largest Induced Cycle. *Theoretical Computer Science* 347, 54–75 (2005)
12. Dourisboure, Y.: Compact Routing Schemes for Bounded Tree-Length Graphs and for k-Chordal Graphs. In: Guerraoui, R. (ed.) DISC 2004. LNCS, vol. 3274, pp. 365–378. Springer, Heidelberg (2004)
13. Dragan, F.F., Yan, C.: Collective Tree Spanners in Graphs with Bounded Genus, Chordality, Tree-width, or Clique-width. In: Deng, X., Du, D.-Z. (eds.) ISAAC 2005. LNCS, vol. 3827, pp. 583–592. Springer, Heidelberg (2005)
14. Dragan, F.F., Yan, C., Corneil, D.G.: Collective Tree Spanners and Routing in AT-free Related Graphs. *Journal of Graph Algorithms and Applications* 10, 97–122 (2006)
15. Dragan, F.F., Yan, C., Lomonosov, I.: Collective tree spanners of graphs. *SIAM J. Discrete Math.* 20, 241–260 (2006)
16. Fonseca, R., Ratnasamy, S., Zhao, J., Ee, C.T., Culler, D., Shenker, S., Stoica, I.: Beacon vector routing: Scalable point-to-point routing in wireless sensor networks. In: Proceedings of the Second USENIX/ACM Symposium on Networked Systems Design and Implementation (NSDI 2005) (2005)
17. Fraigniaud, P., Gavoille, C.: Memory requirements for univesal routing schemes. In: Proceedings of the 14th Annual ACM Symposium on Principles of Distributed Computing, Ontario, Canada, pp. 223–230 (1995)
18. Fraigniaud, P., Gavoille, C.: Routing in trees. In: Orejas, F., Spirakis, P.G., van Leeuwen, J. (eds.) ICALP 2001. LNCS, vol. 2076, pp. 757–772. Springer, Heidelberg (2001)

19. Gavoille, C.: A survey on interval routing schemes. *Theoretical Computer Science* 245, 217–253 (1999)
20. Gavoille, C.: Routing in distributed networks: Overview and open problems. *ACM SIGACT News - Distributed Computing Column* 32 (2001)
21. Gavoille, C., Gengler, M.: Space-efficiency of routing schemes of stretch factor three. *Journal of Parallel and Distributed Computing* 61, 679–687 (2001)
22. Gavoille, C., Hanusse, N.: Compact Routing Tables for Graphs of Bounded Genus. In: Wiedermann, J., Van Emde Boas, P., Nielsen, M. (eds.) *ICALP 1999*. LNCS, vol. 1644, pp. 351–360. Springer, Heidelberg (1999)
23. Gavoille, C., Peleg, D., Pérennès, S., Raz, R.: Distance labeling in graphs. *J. Algorithms* 53, 85–112 (2004)
24. Gavoille, C., Pérennès, S.: Memory requirements for routing in distributed networks. In: *Proceedings of the 15th Annual ACM Symposium on Principles of Distributed Computing*, Philadelphia, Pennsylvania, pp. 125–133 (1996)
25. Giordano, S., Stojmenovic, I.: Position based routing algorithms for ad hoc networks: A taxonomy. In: Cheng, X., Huang, X., Du, D. (eds.) *Ad Hoc Wireless Networking*, pp. 103–136. Kluwer, Dordrecht (2004)
26. Karp, B., Kung, H.T.: GPSR: greedy perimeter stateless routing for wireless networks. In: *Proceedings of the 6th ACM/IEEE MobiCom*, pp. 243–254. ACM, New York (2000)
27. Kleinberg, J.M.: The small-world phenomenon: an algorithm perspective. In: *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing (STOC 2000)*, Portland, OR, USA, pp. 163–170. ACM, New York (2000)
28. Kleinberg, R.: Geographic routing using hyperbolic space. In: *INFOCOM 2007*, pp. 1902–1909 (2007)
29. Kratsch, D., Le, H.-O., Müller, H., Prisner, E., Wagner, D.: Additive tree spanners. *SIAM J. Discrete Math.* 17, 332–340 (2003)
30. Kuhn, F., Wattenhofer, R., Zhang, Y., Zollinger, A.: Geometric ad-hoc routing: of theory and practice. In: *Proceedings of the 22nd Annual Symposium on Principles of Distributed Computing*, pp. 63–72. ACM Press, New York (2003)
31. van Leeuwen, J., Tan, R.B.: Interval routing. *The Computer Journal* 30, 298–307 (1987)
32. Liben-Nowell, D., Novak, J., Kumar, R., Raghavan, P., Tomkins, A.: Geographic routing in social networks. *PNAS* 102, 11623–11628 (2005)
33. Peleg, D.: Distributed Computing: A Locality-Sensitive Approach. In: *SIAM Monographs on Discrete Math. Appl.* SIAM, Philadelphia (2000)
34. Peleg, D.: Proximity-Preserving Labeling Schemes and Their Applications. *J. of Graph Theory* 33, 167–176 (2000)
35. Rao, A., Papadimitriou, C., Shenker, S., Stoica, I.: Geographical routing without location information. In: *Proceedings of MobiCom 2003*, pp. 96–108 (2003)
36. Santoro, N., Khatib, R.: Labeling and implicit routing in networks. *The Computer Journal* 28, 5–8 (1985)
37. Slivkins, A.: Distance estimation and object location via rings of neighbors. In: *PODC 2005*, pp. 41–50 (2005)
38. Talwar, K.: Bypassing the embedding: Algorithms for low dimensional metrics. In: *STOC 2004*, pp. 281–290 (2004)
39. Thorup, M.: Compact oracles for reachability and approximate distances in planar digraphs. *Journal of the ACM* 51, 993–1024 (2004)
40. Thorup, M., Zwick, U.: Compact routing schemes. In: *13th Ann. ACM Symp. on Par. Alg. and Arch.*, pp. 1–10 (2001)