LABORATOIRE

INFORMATIQUE, SIGNAUX ET SYSTÈMES
DE SOPHIA ANTIPOLIS
UMR 6070

BROADCASTING IN ONE ORTHANT

C. Garcia, P. Hell, C. Peyrat

Projet RECIF

Rapport de recherche
I3S/RR–2002-43–FR

Septembre 2002

LABORATOIRE I3S: Les Algorithmes / Euclide B – 2000 route des Lucioles – B.P. 121 –
06903 Sophia-Antipolis Cedex, France – Tél. (33) 492 942 701 – Télécopie : (33) 492 942 898
http://www.i3s.unice.fr/I3S/FR/

RÉSUMÉ :

MOTS CLÉS :

ABSTRACT:

Broadcasting is a method of communication in a digraph, in which one vertex broadcasts a message to all other vertices, by means of calls (along the edges of the digraph) that involve each vertex in at most one call at any one time. We consider broadcasting in the general grid digraph, with edges determined by an *n*-dimensional "template" of vectors. Some 20 years ago, Stout conjectered that the maximum number of vertices that can be informed in *t* time units is always a degree n polynomial in *t*. This has so far only been proved for *n*=1. We prove Stout's conjecture in the case when the template F contains a set F_E of linearly independent vectors such that each other vector of F is a convex combination of vectors in F_E . A typical exemple is the template containing the n binary unit vectors multiplied by some large constant k, together with a number of shorter vectors (of L_1-norm less than k) in the first orthant.

KEY WORDS :

Broadcasting in one orthant

C. Garcia
I3S-ESSI, BP 145
06903 Sophia Antipolis
France
e-mail : garcia@essi.fr

P. Hell
School of Comp. Science
S.F.U., Burnaby, B.C.
Canada V5A 1S6
e-mail : pavol@cs.sfu.ca

C. Peyrat
I3S-ESSI, BP 145
06903 Sophia Antipolis
France
e-mail : claudine@essi.fr

February 21, 2001

Abstract

Broadcasting is a method of communication in a digraph, in which one vertex broadcasts a message to all other vertices, by means of calls (along the edges of the digraph) that involve each vertex in at most one call at any one time. We consider broadcasting in the general grid digraph, with edges determined by an n-dimensional 'template' of vectors. Some 20 years ago, Stout conjectured that the maximum number of vertices that can be informed in t time units is always a degree n polynomial in t. This has so far only been proved for $n = 1$. We prove Stout's conjecture in the case when the template F contains a set F_E of linearly independent vectors such that each other vector of F is a convex combination of vectors in F_E. A typical example is the template containing the n binary unit vectors multiplied by some large constant k, together with a number of shorter vectors (of L_1-norm less than k) in the first orthant.

1 Introduction

We view a digraph G as a model of an interconnection network, with vertices corresponding to the members of the network, and edges corresponding to the communication links. *Broadcasting* in the digraph G is a communication scheme in which one vertex called *the originator* disseminates a message to all other vertices by means of *calls*. Each call takes a unit of time, and a vertex can only call its out-neighbours. *Shouting* is a method of broadcasting in which a vertex can call *all* its out-neighbours in one time interval. *Whispering* only allows one call for each vertex in each time interval. (These terms are due to Q. Stout [12].) Let $\omega_G(t)$, respectively $\sigma_G(t)$, denote the maximum number of vertices of G that can be informed in t time intervals by whispering, respectively shouting. (We shall omit the subscript when the digraph G is understood from context.) Clearly, $\sigma_G(t)$ is simply the number of vertices of G that can be reached from the originator by a directed path of length at most t.

For a general survey about broadcasting problems we refer the interested reader to [6], and for more details about the interconnection network model and the relevance of whispering and shouting to processor-bound and link-bound networks we refer to [3].

We denote by R the set of all real numbers, by Z the set of all integers, and by N the set of all non-negative integers.

Let F be a fixed finite set of vectors in Z^n, called the *template*. The *grid* with template F is the digraph $G(F)$ with the vertex set $\{\sum_{f \in F} c_f \cdot f : c_f \in N\}$ and edges uv such that $v = u + f$, for some $f \in F$. The edge uv with $v = u + f$ is said to be *of type* f. We shall write $G(f_0, f_1, \ldots)$ for $G(\{f_0, f_1, \ldots\})$. We shall also write ω_F and σ_F for the functions $\omega_{G(F)}$ and $\sigma_{G(F)}$ respectively. Note that $G(F)$ contains the origin, and a directed path from the origin to every other vertex. Suppose $F' \subseteq F$ and $v \in V(G(F))$; then $v + G(F')$ is the subgraph of $G(F)$ induced by all vertices of the form $v + w, w \in G(F')$. (We call $v + G(F')$ a *shift* of $G(F')$.)

Broadcasting in grids was first studied by Farley and Hedetniemi [2], in the special case when the template F consists of the $2n$ integer unit vectors $(1, 0, 0, \ldots, 0), (-1, 0, 0, \ldots, 0), \ldots, (0, 0, 0, \ldots, 1), (0, 0, 0, \ldots, -1)$. While shouting in such a simple grid is trivial, whispering is not, and [2] conjectured the exact value of $\omega_F(t)$. This conjecture was verified for $n = 2$ by C.S. Ko [9, 10] and G.W. Peck [11]; however for $n > 2$ it was shown false, by both these authors. Ko computed the exact value of $\omega_F(t)$ for $n = 3$, Peck computed the two highest order terms of $\omega_F(t)$ for general n. (The exact form of $\omega_F(t)$ for general n is still not known.)

For all the templates F for which $\omega_F(t)$ was computed exactly, it turned out to be a polynomial in t, and Q. Stout [12] conjectured that $\omega_F(t)$ is always (for any template F) a polynomial in t, for sufficiently large t. (The degree of the polynomial is the dimension of the space spanned by the template.) This has still not been proved, even for the above special template of unit vectors. However, it has been verified for general F when $n = 1$ [7].

We shall prove Stout's conjecture in the special case when the template F contains a subset F_E of linearly independent vectors such that each $f \in F_N = F - F_E$ is a *convex combination* of F_E, i.e., $f = \sum_{e \in F_E} r_e \cdot e$, where each r_e is a non-negative real number and $\sum_{e \in F_E} r_e < 1$. (In other words, each $f \in F_N$ is in the convex hull of $F_N \cup \{0^n\}$.) We shall call the vectors $e \in F_E$ *extreme*, and the vectors $f \in F_N$ *non-extreme*. A particularly simple situation when this assumption applies is when the template contains the n vectors $(k, 0, 0, \ldots, 0)$, $(0, k, 0, \ldots, 0), \ldots, (0, 0, 0, \ldots, k)$, for some positive integer k, together with any number of non-negative vectors with L_1-norm less than k.

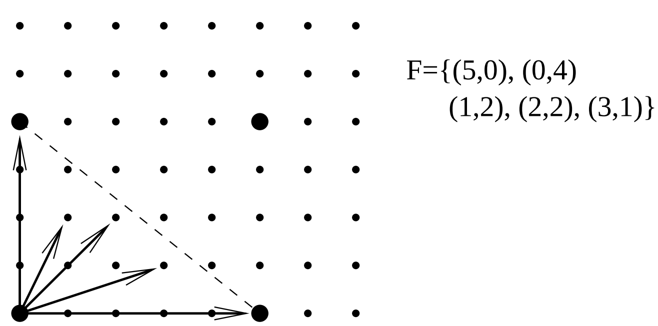

Figure 1: A typical template

Stout also conjectured that $\sigma_F(t)$ is always (for any template F) a polynomial in t, for sufficiently large t. This was proved by Klarner [8]. When $n = 2$, the exact form of this polynomial was determined in [4, 5]. For $n > 2$,

[5] determined the highest order term of the polynomial $\sigma_F(t)$. (Previous to [5], S. Djelloul [1] obtained these results for the special case where F consists of vectors with exactly one non-zero coordinate.) Stout has given an asymptotic estimate of the difference $\sigma_F(t) - \omega_F(t)$: it turns out to be $\sigma_F(t) - \omega_F(t) = O(t^{n-1})$. It follows that if F spans R^n, then the highest order term of $\sigma_F(t)$ is also the highest order term of $\omega_F(t)$. (Thus asymptotically grids can disseminate information just as quickly by whispering as by shouting.) Our proof of Stout's conjecture (in the special case detailed above) amounts to calculating the difference $\sigma_F(t) - \omega_F(t)$ more precisely. We prove the conjecture by verifying that this difference *is* a polynomial of degree at most $n-1$.

We call $\Lambda(F) = \{\sum_{f \in F} c_f \cdot f : c_f \in Z\}$ the *lattice* of F. We also write $\Lambda(f_0, f_1, \ldots)$ for $\Lambda(\{f_0, f_1, \ldots\})$. Two vectors of Z^n with difference in $\Lambda(F)$ are called *congruent modulo* $\Lambda(F)$. Note that $\Lambda(F)$ strictly contains all vertices of $G(F)$.

In a digraph G with originator o, we denote by $d_G(v)$ the length of a shortest directed path from o to v. If G is understood from the context we omit the subscript G. In $G(F)$, the grid with template F, the originator is the origin 0^n, and we write $d_F(v)$ for $d_{G(F)}(v)$. Note that $d_F(v)$ is the minimum value of $\sum_{f \in F} c_f$, with all $c_f \in N$, such that $v = \sum_{f \in F} c_f \cdot f$. We write $v \cong \sum_{f \in F} c_f \cdot f$ to signify that $v = \sum_{f \in F} c_f \cdot f$ and $d_F(v) = \sum_{f \in F} c_f$.

LEMMA 1 *For each non-extreme vector f there exists an integer k_f such that any shortest path starting from the origin 0^n has at most k_f edges of type f. In other words, if $v \cong \sum_{f \in F} c_f \cdot f$, then for each $f \in F_N$, we have $c_f \leq k_f$.*

Proof: Recall that each non-extreme vector can be written as $f = \sum_{e \in F_E} r_e \cdot e$, where each r_e is a non-negative real number and $\sum_{e \in F_E} r_e < 1$. Since the vectors e are in Z^n, the coefficients r_e may be chosen to be rational numbers, and thus there exists a positive integer k and positive integers c_e, $e \in F_E$, such that $k \cdot f = \sum_{e \in F_E} c_e \cdot e$ and $k < \sum_{e \in F_E} c_e$. We let k_f be the smallest integer k for which this is possible. It follows that a shortest path from 0^n can never use k_f (or more) edges of type f. □

It is clear from these definitions that any whispering scheme requires at least $d(v)$ units of time to inform v. A scheme that informs v after $d(v) + i$ units of time is said to have *delay i at v*; also the call that informed v at time $d(v) + i$ (and the edge on which it happened) is said to have *delay* i. The *delay of a whispering scheme* is a maximum delay of any vertex under the scheme, except possibly for a finite set of vertices v. Specifically, we say that a scheme has delay i if almost all vertices (all except for finitely many) have delay at most i.

Consider a whispering scheme in an arbitrary digraph G. The originator only made one call with delay 0, and the vertex so informed also only made one call with delay 0, and so on. Thus the edges with delay 0 form a path starting at the originator. We call this path the *delay*-0 *path*. A vertex lying on the delay-0 path possibly informed another one of its neighbours (in addition to the neighbour it informed with delay 0) by a call with delay 1. This neighbour may have, in turn, immediately made a call to one of its neighbours, who in turn immediately informed another vertex, and so on. These calls will all have delay 1, and form a *delay*-1 *path*. Any vertex on the delay-0 path can begin a delay-1 path. We can always organize the calls so that these paths (of delay 1) are disjoint from one another; thus we can select paths of delays 0 and 1 which form a tree. *Delay-i paths*, with $i > 1$, are defined analogously – any vertex on the tree formed by the delay-j paths with $j < i$ can begin a delay-i path.

We now derive an upper bound on $\omega_F(t)$, the number of vertices that can be informed in $G(F)$ in time t. Crucial for this estimation is the observation that at most $\binom{t}{j}$ vertices can be informed along paths of length $t - j, 0 \leq j \leq t$. Indeed, for each set S of j time periods from $1, 2, \ldots, t$, at most one vertex can be informed by a path of length $t - j$ on which a call is made for all time periods except those in S. Let now s be a fixed non-negative integer, and consider all $t \geq s$. There are $\sigma_F(t - s)$ vertices v with $d_F(v) \leq t - s$, and the vertices with $d_F(v) > t - s$ can only be informed along paths of length at least $t - s + 1$; thus there can be at most $\sum_{j \leq s-1} \binom{t}{j}$ such vertices. Thus we conclude:

LEMMA 2 *For any $s, 0 \leq s \leq t$,*

$$\omega_F(t) \leq \sigma_F(t - s) + \sum_{0 \leq j \leq s-1} \binom{t}{j}.$$

□

Moreover, we also see from the proof that any scheme that informs all vertices v with $d_F(v) \leq t - s$ and exactly $\binom{t}{j}$ vertices with $d_F(v) = t - j, 0 \leq j < s$ is optimal.

2 Basic Schemes

We assume from now on that $F \subseteq Z^n$ consists of vectors $f_0, f_1, \ldots, f_{p-1}$, where $F_E = \{f_0, f_1, \ldots, f_{m-1}\}$ are linearly independent, and $F_N = \{f_m, f_{m+1}, \ldots, f_{p-1}\}$ are convex combinations of vectors from F_E. (Note that $m \leq n$.)

We shall first explain how to broadcast in $G(F_E)$. This solves the problem (and proves Stout's conjecture) in the case when $F_N = \emptyset$, i.e., when all vectors of F are linearly independent.

Since the vectors are linearly independent, each vertex v of $G(F_E)$ has unique coordinates $(c_0, c_1, \ldots, c_{m-1})$ with $v = \sum_i c_i \cdot f_i$. The following will be our basic 'backbone' whispering scheme: The delay-0 path starts at the origin $o = 0^n$ (whose coordinates are $(0, 0, \ldots, 0)$), and makes only calls on edges of type f_0. For any $i, 0 < i < m$, each vertex with coordinates $(c_0, c_1, \ldots, c_{i-1}, 0, 0, \ldots, 0)$ begins a delay-i path which makes only calls on edges of type f_i. Let us now look at this whispering scheme from the perspective of an arbitrary vertex v; assume that v has coordinates $(c_0, c_1, \ldots, c_i, 0, 0, \ldots, 0)$ with $c_i > 0$. It receives its information, with delay i, along an edge of type f_i, from the vertex with coordinates $(c_0, c_1, \ldots, c_i - 1, 0, \ldots, 0)$. It begins whispering as soon as it is informed, and makes the calls to its neighbours on edges of the types $f_i, f_{i+1}, \ldots, f_{m-1}$, in that order. Clearly, in this scheme each vertex is informed with delay at most $m - 1$. Thus all vertices v with $d_{F_E}(v) \leq t - m + 1$ are informed within time t. Let us now consider vertices v with $d_{F_E}(v) = t - j$ where $0 \leq j \leq m - 2$. We claim that within time t the scheme informs precisely $\binom{t}{j}$ such vertices. Indeed, let S consist of the j time periods $s_1 < s_2 < \ldots < s_j$ from $1, 2, \ldots, t$. The path of length $t - j$ on which the information is passed during all times other than those in S informs the vertex with coordinates $(s_1 - 1, s_2 - s_1 - 1, s_3 - s_2 - 1, \ldots, s_j - s_{j-1} - 1, t - s_j, 0, 0, \ldots, 0)$. Thus different sets S result in informing different vertices v with $d_{F_E}(v) = t - j$.

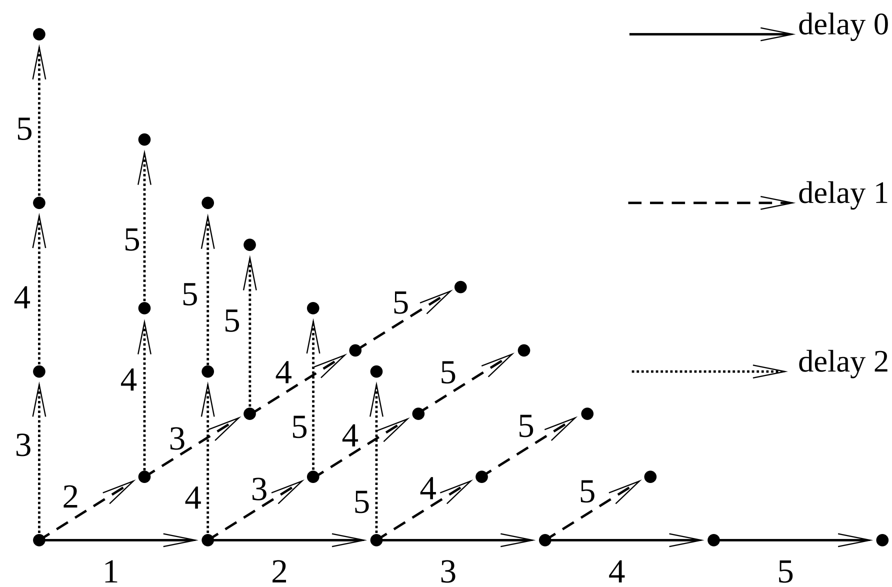

Figure 2: The backbone scheme of three independent vectors in Z^3

Thus, according to Lemma 2 (with $s = m - 1$) and the remark following it, we have an optimal scheme, i.e.,

$$\omega_{F_E}(t) = \sigma_{F_E}(t - m + 1) + \sum_{0 \leq j \leq m-2} \binom{t}{j}.$$

Since $\sigma_{F_E}(t)$ is a degree m polynomial (for large t) according to [8], the same is true of $\omega_{F_E}(t)$.

If F_N is not empty, we shall begin by whispering in $G(F_E)$ as explained above, and then proceed to inform the remaining vertices of $G(F)$ along a tree-like scheme using vectors of F_N. This is similar in principle to what is done in [7] for $n = 1$. In fact, our 'one-orthant' case can best be viewed as a generalization of the 'one-directional' case of [7] - this is the case of $n = 1$ with all $f \in F$ positive. (However, applying our techniques to $n = 1$ offers a simplification of the schemes from [7].)

Before proceeding with our proof, we give a few motivating examples.

First, we shall take a template F with exactly one extreme vector. We may assume that $n = 1$, i.e., that F consists of one-dimensional vectors. Thus let $F = \{5, 3, 2\}$, i.e., $f_0 = 5, f_1 = 3, f_2 = 2$.

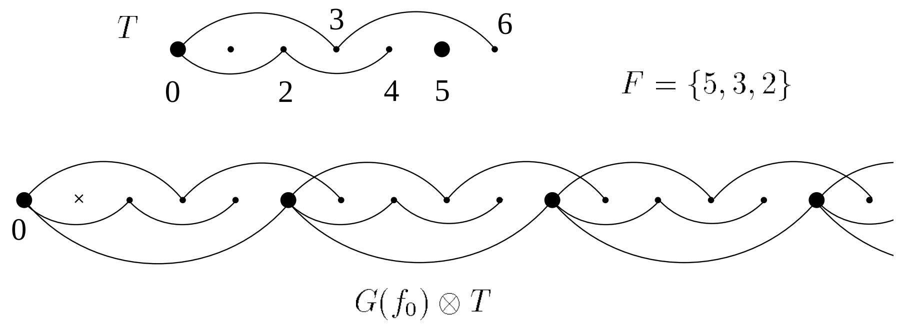

Figure 3: A one-dimensional example

Figure 3 shows the grid $G(f_0)$ in which we have attached at each vertex a copy of a particular tree T. The resulting graph, which we shall call $G(f_0) \otimes T$ (cf. below), is a subgraph of the grid $G(F)$, in which the distances from the origin (to all other vertices) are the same as in $G(F)$. Moreover, in this example the subgraph is spanning, i.e., contains all vertices of $G(F)$. (Note that 1 is not a vertex of $G(F)$.)

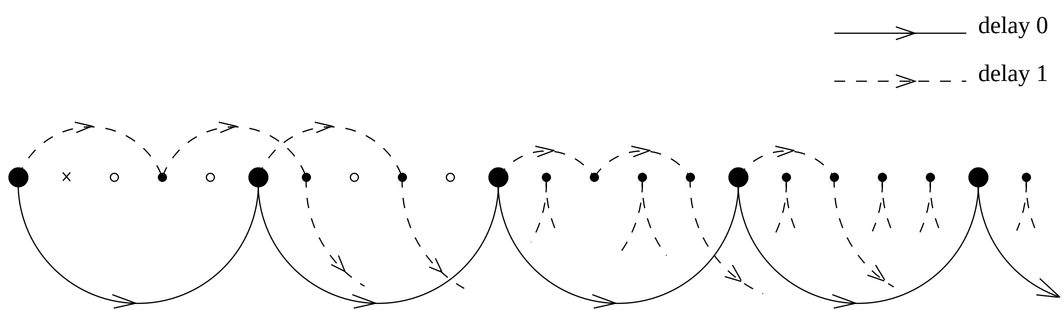

Figure 4: A whispering scheme

In Figure 4 we show a whispering scheme for $G(f_0, f_1, f_2)$ using the edges of $G(f_0) \otimes T$. The delay-0 path consists of all edges of $G(f_0)$ and is indicated by the solid lines. The delay-1 paths are depicted by interrupted lines. They consist of certain paths taken out from the tree T, followed by repeating edges of type f_0. (All the paths shown are continued by edges of type f_0.) The delay-1 paths are organized in such a way that all sufficiently large vertices i are covered by the tree of delay-0 and delay-1 paths. On the other hand, there is a finite number of vertices of $G(F)$ which are not informed by this scheme - namely $2, 4, 7$ and 9. Nevertheless, these vertices are easily informed by ad hoc paths when t is large enough; for instance, the origin 0 is only needed to make calls at times 1 and 2, and so is free to inform 2 at time 3, which can call 4 at time 4, and so on.

In what follows we shall explain how to construct a suitable tree T, and how to organize the broadcast in $G(F_E) \otimes T$. Unfortunately there are several complications not present in the case of just one extreme vector. We first give an example with two extreme vectors, which illustrates some (but not all) of the difficulties. We then give the general proofs, although we do emphasize the case of two extreme vectors, as the notation in the most general case tends to obstruct the ideas.

We first formally define the digraph $G(F_E) \otimes T$. Assume that T is a subtree of $G(F_N)$ containing the origin. If T contains at most one vertex from each congruence class modulo $\Lambda(F_E)$, then we say that T is *consistent* (with respect to F_E). A consistent tree is *complete* if it contains *exactly* one vertex from each congruence class. Given a consistent tree T we define $G(F_E) \otimes T$ to be the subgraph of $G(F)$ induced by vertices of the form $v + u$, where $v \in V(G(F_E))$ and $u \in V(T)$. (Whenever $G(F_E) \otimes T$ is mentioned, T is assumed to be consistent with respect to F_E.) Since T is consistent with respect to F_E, $v + u = v' + u'$, with $v, v' \in V(G(F_E))$ and $u, u' \in V(T)$, is only possible when $u = u'$ and $v = v'$. Thus each vertex of $G(F_E) \otimes T$ can be uniquely written as $u + \sum_{0 \leq i \leq m-1} c_i \cdot f_i$, i.e., has unique 'coordinates' $(u; c_0, c_1, \ldots, c_{m-1}), u \in V(T), c_i \in N$. Note that $G(F_E) \otimes T$ contains $G(F_E)$. In fact, $G(F_E) \otimes T$ can best be visualized as a cartesian product of $G(F_E)$ with T; in particular, it contains a copy of T 'planted' at each vertex v of $G(F_E)$ (this is the tree induced by the vertices with coordinates $(v; c_0, c_1, \ldots, c_{m-1})$ with fixed $c_0, c_1, \ldots, c_{m-1}$). We say that $G(F_E) \otimes T$ is *isometric* in $G(F)$ if the distances from the origin are the same in both these graphs (for all vertices of $G(F_E) \otimes T$). Sometimes we just say that T is isometric, meaning $G(F_E) \otimes T$ is isometric in $G(F)$.

If we are lucky enough to find an isometric tree T that is consistent with respect to F_E and covers all vertices of $G(F)$ (or at least leaves uncovered only a finite number of them), we can attempt to broadcast in $G(F_E) \otimes T$ as in the one-dimensional example (see the next Theorem). However, most of the time we shall have to be satisfied with a tree T which leaves uncovered infinitely many vertices of $G(F)$.

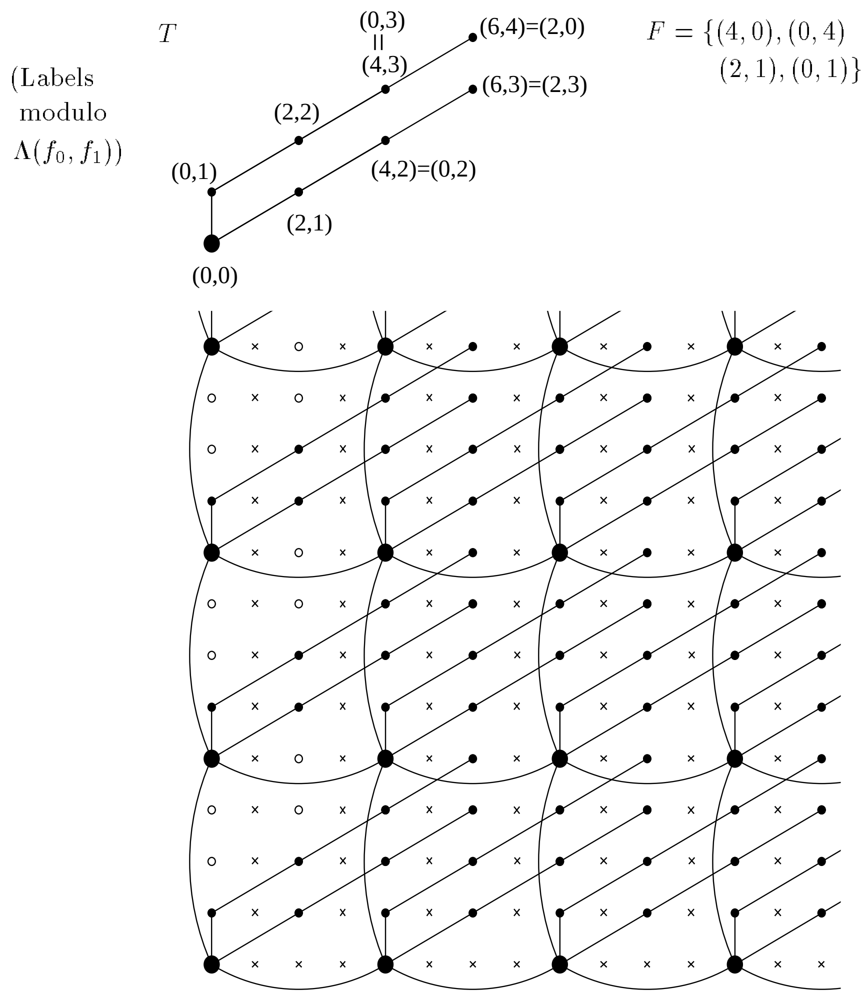

Figure 5: A tree that fails

In Figure 5, $f_0 = (4,0), f_1 = (0,4), f_2 = (2,1), f_3 = (0,1)$, where $m = 2$, i.e., f_0, f_1 are extreme, f_2, f_3 non-extreme. This digraph $G(F)$ has eight congruence classes modulo $\Lambda(f_0, f_1)$, and the tree T shown in the figure is isometric, consistent (and complete), with respect to $\{f_0, f_1\}$. Yet the figure shows that the graph $G(f_0, f_1) \otimes T$ does not cover any of the vertices $(0, 4\lambda+2), (0, 4\lambda+3), (2, 4\lambda+3), (2, 4\lambda+4), \lambda \geq 0$, of $G(F)$.

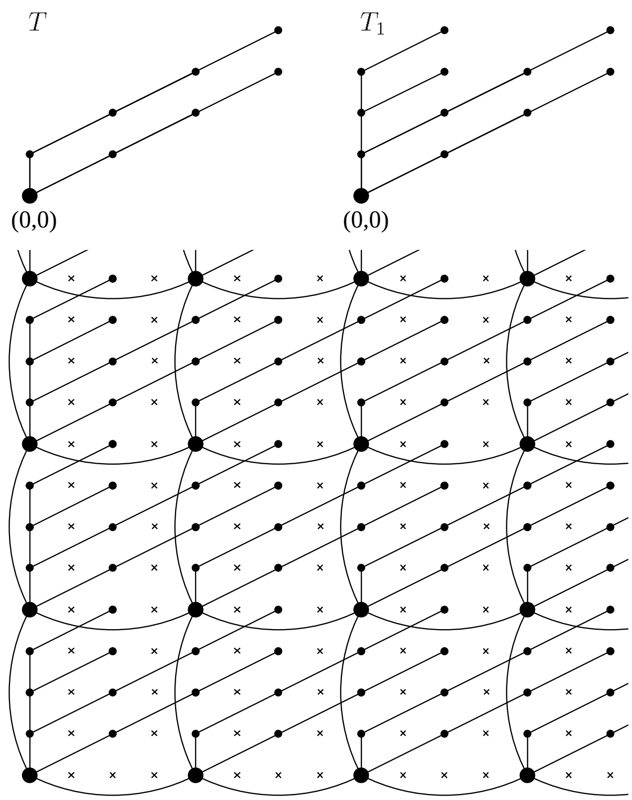

Figure 6: Tree T and T_1 planted in the appropriate places

In Figure 6 we suggest a remedy: A new, larger, isometric tree T_1 is given, this time consistent only with respect to $\{f_1\}$, and planted instead of T at all vertices of $G(f_1)$. The resulting digraph, shown in Figure 6, covers all vertices of $G(F)$.

In Theorem 5 we explain how to find the right trees and plant them at the right places. For the remainder of this section we concentrate on broadcasting in the resulting graph.

First we show an example of a whispering scheme for two extreme vectors.

To focus on the essence of the problem we shall take a simpler example than the one given above. In the example in Figure 7 we have the template with extreme vectors $f_0 = (2,0), f_1 = (0,2)$ and non-extreme vectors $f_2 = (1,0), f_3 = (0,1)$, and a consistent and isometric tree T. (In this example we do not have the complication discussed in Figures 5 and 6 - indeed, this time $G(f_0, f_1) \otimes T$ does cover all vertices.) We may try to broadcast in each shift $(\lambda, 0) + G(f_1), \lambda \geq 0$, using the one-dimensional strategy illustrated in Figure 4; this results in a broadcast shown in Figure 7:

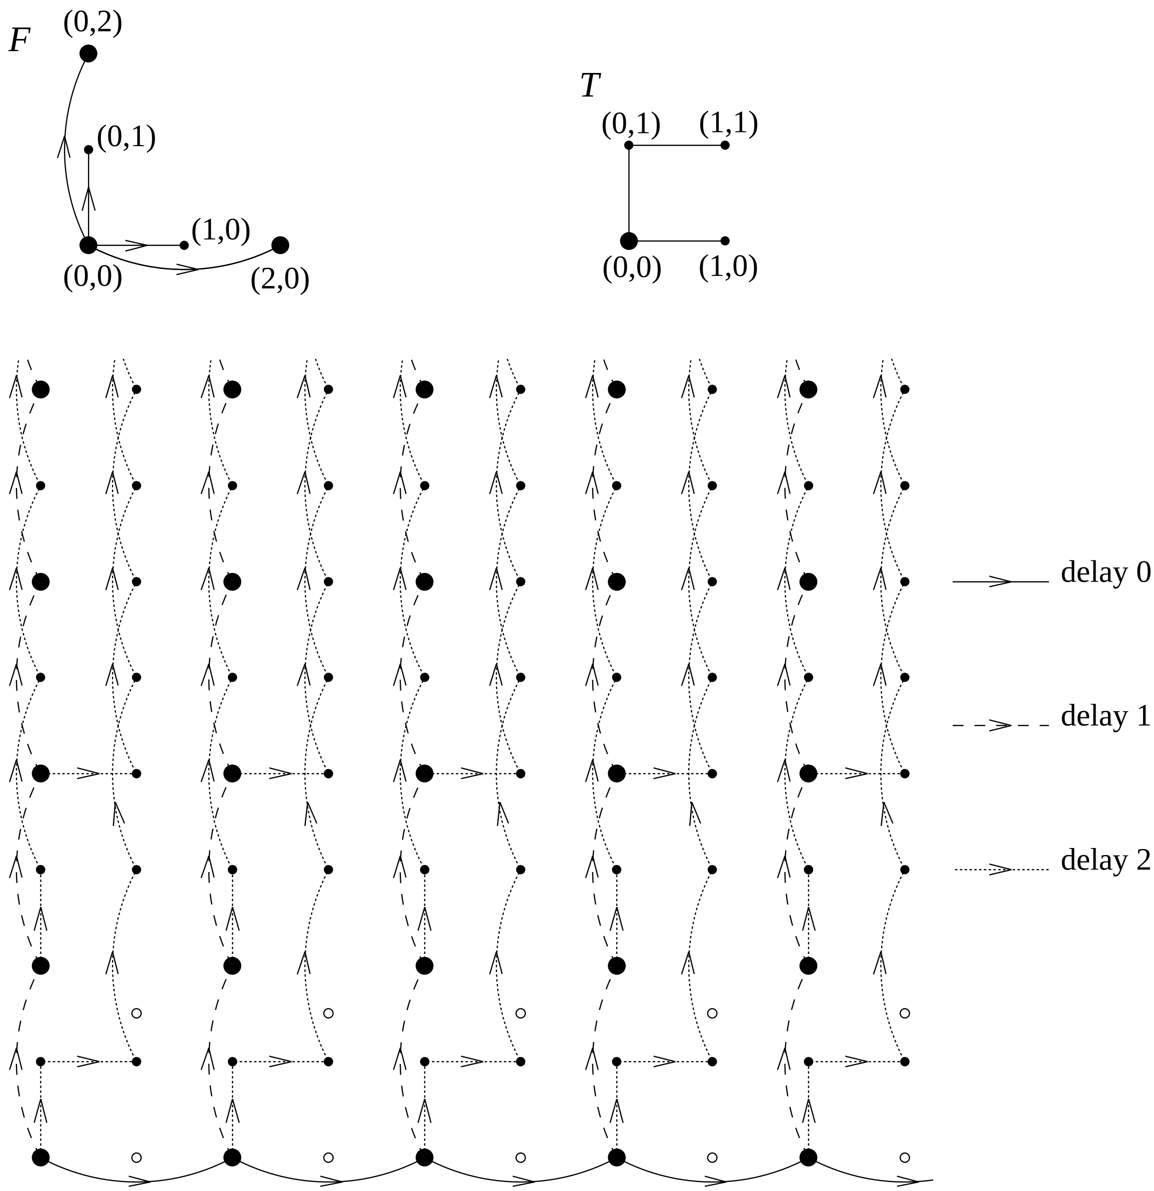

Figure 7: A broadcast that fails

This time, infinitely many vertices of $G(f_0, f_1) \otimes T$ were left uninformed. (They are marked by empty circles.) Figure 8 suggests a way to inform such vertices. By 'shifting' the scheme in $G(f_1)$ four steps upward (in other words, to $(0,4) + G(f_1)$), we were able to leave enough room to start horizontal broadcasts in $G(f_0)$ and $(0,2) + G(f_0)$, which inform all the missed vertices.

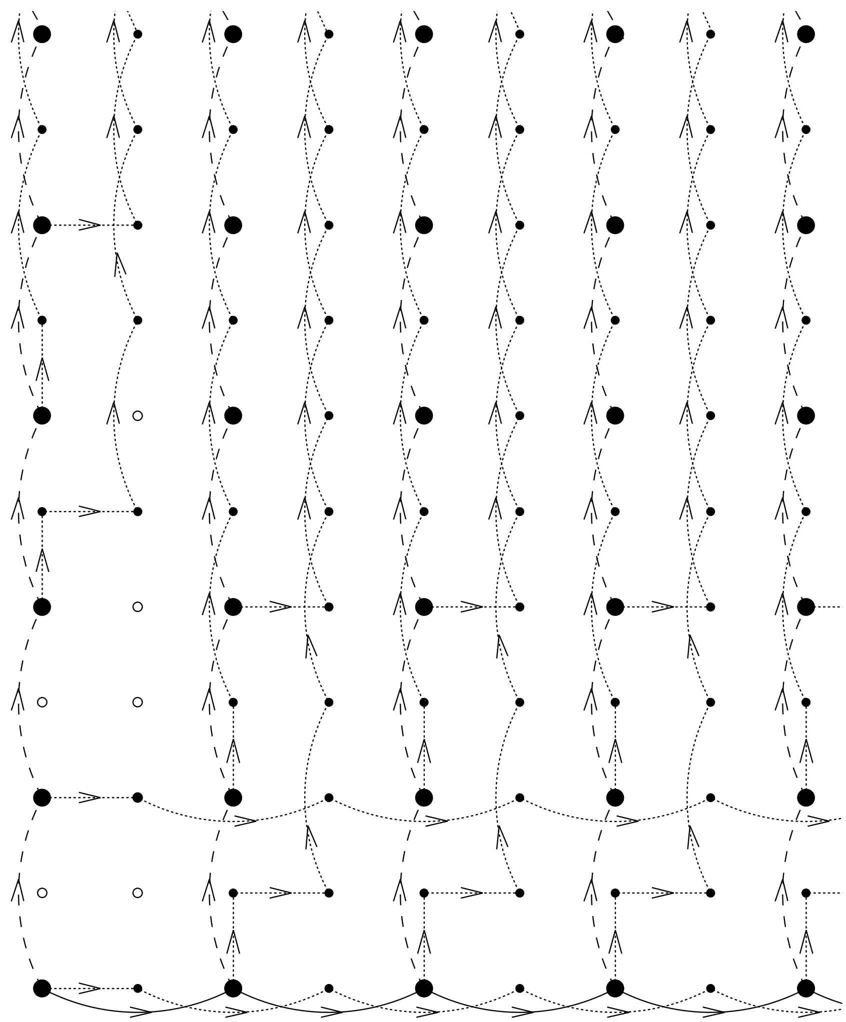

Figure 8: A scheme that informs almost all vertices

This kind of these whispering scheme can be constructed for any $G(F') \otimes T'$, where F' is a linearly independent set of vectors, and T' is consistent with respect to F'. (The example in Figure 8 takes once $F' = \{f_0, f_1\}$ with tree $T' = T$, and once $F' = \{f_1\}$ with tree $T' = T_1$. More general situations are discussed after the proof.)

THEOREM 3 *Let $F' \subseteq F_E$ and let T' be a subtree of $G(F)$ consistent with respect to F'. Let k be a positive integer such that each vertex of $G(F')$ has been informed with delay $k-1$, and is free to broadcast with delay k.*

Then almost all vertices of $G(F') \otimes T'$ can be informed with delay k.

Note that the standard backbone scheme in $G(F_E)$ does inform each vertex of $G(F_E)$ with delay at most $m-1$ and leaves it informed and free to broadcast with delay m. Thus we can apply the theorem with $F' = F_E$, $T' = T$, and $k = m$, assuring that in $G(F_E) \otimes T$ there is a whispering scheme of delay m. However, just as in Figure 8, we will need 'shifted' schemes which may have some other delays, and also we will have situations where some of the vertices of $G(F')$ have been informed in some other ways. The proof below does not depend on the particular way $G(F'')$ has been informed, as long as the hypotheses of the Theorem are fulfilled.

Proof. We number the vertices of T' in reverse breadth first order, i.e., $u_0, u_1, \ldots, u_r = 0^n$ so that $d_F(u_0) \geq d_F(u_1) \geq \ldots, d_F(u_r) = 0$. Let P_s be the path in T' from 0^n to u_s, for $s = 0, 1, \ldots, r$.

We will proceed by induction on $m = |F'|$. Assume first that $m = 1$ (cf. Figure 3): Thus $F' = \{f_0\}$, and the vertices of $G(F')$, assumed to have been informed with delay $k-1$, have the form $i \cdot f_0, i \in N$. The vertices of $G(F') \otimes T'$ have coordinates $(u_s; i)$, where $s = 0, 1, \ldots, r$, and $i \in N$. We claim that almost all of them can be informed with delay k. Each vertex $i \cdot f_0$ with $0 \leq i \leq r-1$ (this is the vertex with coordinates $(u_r; i)$) will begin a delay-k path which begins with edges of types found on the path P_i, and then repeatedly takes edges of type f_0 (cf. Figure 3). Because of the consistency of T' and the special numbering of the vertices of T' (and hence of the paths P_i), these delay-k paths are disjoint. It follows that all vertices with coordinates $(u_s; i)$, $s = 0, 1, \ldots, r$, $i \geq r$, are reached with delay k or less. This scheme misses only finitely many vertices, and we have a whispering scheme of delay k.

We note for future reference that by limiting the delay-k paths to those starting only at $i \cdot f_0, i \in I$ for some set $I \subseteq \{0, 1, 2, \ldots, r-1\}$, we will inform exactly the vertices with coordinates $(u_i; j), i \in I, j \in N$ (vertices corresponding to $u_i, i \in I$, in all of the trees planted in $G(f_0)$), except possibly for a finite set of vertices.

We also note that we can apply schemes to shifts of their original domain: A whispering scheme on $G(F')$ is applied to its shift $v + G(F')$ by changing each call from a to b into a call from $v + a$ to $v + b$. We shall say that we have *shifted the scheme on G' by the vector v*.

Before giving the general inductive step, we explain the idea for $m = 2$ (cf. Figure 8): thus let $F' = \{f_0, f_1\}$. The vertices of $G(F') \otimes T'$ consist of the union of the vertices of $[i \cdot f_0 + G(f_1)] \otimes T', i \in N$. They are assumed to have been informed with delay at most $k-1$. We proceed as follows: For each $i > r$ we shift the scheme from $G(f_1) \otimes T'$ by the vector $i \cdot f_0$. For each $0 \leq i \leq r$ we shift the same scheme by the vector $i \cdot f_0 + (r+1) \cdot f_1$. The shifted scheme in $[i \cdot f_0 + G(f_1)] \otimes T'$ informs all its vertices except possibly some of the vertices of the copies of T' planted at $i \cdot f_0 + j \cdot f_1$ with $j \leq r$. To take care of this infinite set of vertices, we shall shift, for each $0 \leq j \leq r$, the scheme from $G(f_0) \otimes T'$ by the vector $j \cdot f_1$, but limiting it to inform only those vertices $i \cdot f_0 + j \cdot f_1 + u, i \in N, u \in V(T')$, for which $(r+1) \cdot f_0 + j \cdot f_1 + u$ has not been informed by the scheme in $[(r+1) \cdot f_0 + G(f_1)] \otimes T'$. These calls are disjoint from other delay-k calls, and they assure that, except for pairs i, j with $i \leq r$ and $j \leq 2r+1$, the copy of T' planted at $i \cdot f_0 + j \cdot f_1$ has all its vertices reached with delay k or less.

We now prove by induction on m that there exists a constant r_m and a whispering scheme on $G(F') \otimes T'$ which reaches with delay at most k all vertices with coordinates $(u; i_0, i_1, \ldots, i_{m-1})$, where $u \in V(T')$ and $i_j > r_m$ for at least one j, and that the scheme can be restricted to inform only those vertices with coordinates $(u_i; i_0, i_1, \ldots, i_{m-1})$ with $i \in I$, for any subset I of $\{0, 1, \ldots, r-1\}$. We have just shown that this is so for $m = 1$, with $r_1 = r-1$ (and for $m = 2$, with $r_2 = 2r+1$).

Assuming the assertion for $m-1$, we define the whispering as follows: For each $i_0 > r_1$ we shift the scheme from $G(f_1, \ldots, f_{m-1}) \otimes T'$ by the vector $i_0 \cdot f_0$. For each $0 \leq i_0 \leq r_1$ we shift the same scheme by $i_0 \cdot f_0 + (r_{m-1} + 1) \cdot f_1$. For any $0 \leq i_0 \leq r_1, 0 \leq i_1 \leq r_{m-1}$, we shift the scheme from $G(f_2, \ldots, f_{m-1}) \otimes T'$ by the vector $i_0 \cdot f_0 + i_1 \cdot f_1 + (r_{m-1} + 1) \cdot f_2$. We continue this way, shifting for each $j \leq m-1$, and $0 \leq i_0 \leq r_1, 0 \leq i_1 \leq r_{m-1}, 0 \leq$

$i_2 \leq r_{m-1}, \ldots, 0 \leq i_{j-1} \leq r_{m-1}$, the scheme from $G(f_j, \ldots, f_{m-1}) \otimes T'$ by the vector $i_0 \cdot f_0 + \ldots + i_{j-1} \cdot f_{j-1} + (r_{m-1}+1) \cdot f_j$. At this point consider the copy of the tree T' planted at a vertex $\sum_{0 \leq j \leq m-1} i_j \cdot f_j$. If for some $j \geq 1$ we have $i_j \geq r_{m-1}$, then all vertices of the tree are reached with delay at most k. On the other hand, if all $i_j \leq r_{m-1}$, except for i_0, the tree may contain vertices not informed with delay k, even if i_0 is very large. Therefore we shall shift the scheme from $G(f_0) \otimes T'$ by each $\sum_{1 \leq j \leq m-1} i_j \cdot f_j$ with all $i_j \leq r_{m-1}$, but limiting it to inform only those vertices $i_0 \cdot f_0 + \sum_{1 \leq j \leq m-1} i_j \cdot f_j + u, i_0 \in N$, for which $(r_1+1) \cdot f_0 + \sum_{1 \leq j \leq m-1} i_j \cdot f_j + u$ has not been informed by the scheme in $[(r_1+1) \cdot f_0 + G(f_1, \ldots f_{m-1})] \otimes T'$. These calls assure that each vertex with coordinates $(u; i_0, i_1, \ldots, i_{m-1}), u \in V(T')$, is reached with delay k or less, unless $i_0 \leq r_1$ and all $i_j \leq 2r_{m-1}+1$.

Because of the way we have informed the vertices, by shifting lower dimensional schemes, it is still the case that, for any set $I \subseteq \{0, 1, \ldots, r-1\}$, we can restrict the scheme to inform only the vertices corresponding to $u_i, i \in I$, in all of the copies of T' planted in $G(F')$.

This completes the induction (with $r_m = 2r_{m-1}+1$), and the proof of Theorem 3. □

It is easy to see that this technique actually applies in the cartesian product of the graph $G(F')$ (with F' linearly independent) and any connected digraph H; all vertices can be informed with delay at most $m = |F'|$.

3 Whispering in $G(F)$

In the previous section we have shown how to broadcast in $G(F_E) \otimes T$, if T is consistent with respect to F_E. Suppose we could find a tree T consistent with respect to F_E, such that $G(F_E) \otimes T$ is isometric in $G(F)$, and contains almost all vertices of $G(F)$. Then Theorem 3 would produce a whispering scheme of delay $m = |F_E|$ in $G(F)$ (which would be optimal).

Unfortunately we cannot always find such a tree T; however, we shall still be able to use the whispering schemes from Theorem 3 on subgraphs of $G(F)$ so that almost all vertices are informed with delay at most n.

Our first observation is that there is a consistent tree T for which $G(F_E) \otimes T$ contains all those vertices of $G(F)$ which can be written as $\sum_{f \in F} x_f \cdot f$ so that all the coefficients x_e of the extreme vectors $e \in F_E$ are sufficiently large:

LEMMA 4 *There exists a tree T consistent with respect to F_E, such that $G(F_E)\otimes T$ is isometric in $G(F)$ and contains all vertices of $\sum_{e\in F_E} J\cdot e+G(F)$, for some constant J.*

Note that by the conclusion of the lemma we see that T is also complete.

Proof. By Lemma 2, there exist constants $k_f, f\in F_N$, such that each shortest path from the origin has at most k_f edges of type f. Let C be a congruence class modulo $\Lambda(F_E)$, and let v be an element of C in $G(F)$. Consider the infinite sequence $v+j\cdot\sum_{f\in F_E} f$, for $j\in N$, of elements of C in $G(F)$. For each $v+j\cdot\sum_{f\in F_E} f$, choose one shortest path (in $G(F)$) from 0^n, and consider the numbers x_f of edges of each type $f, f\in F_N$. Since each $x_f\leq k_f$, infinitely many $v+\sum_{f\in F_E} j\cdot f$ will have exactly the same numbers $x_f, f\in F_N$. We moreover assume that the choice of $v\in C$ and all the choices of x_f were made in such a way that $\sum_{f\in F_N} x_f$ is as small as possible.

Let P_C be a path starting at the origin which has x_f edges of type f, for each $f\in F_N$. There exists a subscript j_C such that all vertices $v+\sum_{f\in F_E} j_f\cdot f$ with all $j_f\geq j_C$ admit a shortest path from the origin which consists of P_C followed by a non-negative number of edges of types f with $f\in F_E$. Indeed let j_C be the first subscript such that $v+\sum_{f\in F_E} j_C\cdot f$ has a shortest path consisting of P_C and non-negative numbers j_f of edges of type $f, f\in F_E$. For any $v+\sum_{f\in F_E} j_f\cdot f$ with all $j_f\geq j_C$ there exists a $j'\geq j_f, f\in F_E$ for which $v+\sum_{f\in F_E} j'\cdot f$ also has a shortest path consisting of P_C and $j'_f\geq j_f$ edges of type $f, f\in F_E$. Since a restriction of a shortest path must also be a shortest path, it follows that $v+\sum_{f\in F_E} j_f\cdot f$ also has a shortest path consisting of P_C and non-negative numbers of edges of types $f, f\in F_E$. Suppose we found a path P_C and the corresponding value $j=j_C$ for each congruence class C modulo $\Lambda(F_E)$. Let J be the maximum of the values j_C, and let T be a tree obtained from the union of the P_C in which the distances from the origin are the same as on the paths P_C. (It is enough to recursively define T by starting with one P_C and then adding for each next P_C only that part which goes up to the first common vertex with the already constructed tree.) We now conclude that $G(F_E)\otimes T$ contains all vertices of $G(F)+\sum_{f\in F_E} J\cdot f$. □

Thus $G(F_E)\otimes T$ covers all vertices of $G(F)$ which can be written as $\sum_{f\in F} x_f\cdot f$, where $x_f\geq J$ for all f from F_E. Note that when $|F_E|=1$, this scheme leaves uninformed only finitely many vertices of $G(F)$, and thus represents a scheme of delay 1. In particular, when $n = 1$ we must have $|F_E| = 1$, and so we have a new, and simpler, proof of the one-directional result from [7]. For general F_E, we could have infinitely many uncovered vertices, cf. Figure 5. To deal with this situation, we shall partition the vertices of $G(F)$ into subgraphs which are shifts of other subgraphs $G(F')\otimes T'$ with suitably consistent trees T', which together cover almost all vertices of $G(F)$.

We shall first explain how this works in the case of two extreme vectors.

THEOREM 5 *There exists a constant K, such that for each $j = 0, 1, \ldots, K-1$, there exist trees T_j^0, T_j^1 (consistent with respect to $\{f_0\}, \{f_1\}$ respectively), and constants K_j^0, K_j^1, such that the digraphs*

- $G^{0,1} = [K \cdot f_0 + K \cdot f_1 + G(f_0, f_1)] \otimes T$,
- $G_j^0 = [K_j^0 \cdot f_0 + j \cdot f_1 + G(f_0)] \otimes T_j^0$, *for* $j = 0, 1, \cdots, K-1$, *and*
- $G_j^1 = [j \cdot f_0 + K_j^1 \cdot f_1 + G(f_1)] \otimes T_j^1$, *for* $j = 0, 1, \cdots, K-1$,

1. *are pairwise vertex-disjoint,*
2. *each is isometric in $G(F)$,*
3. *and their union U contains almost all the vertices of G.*

The vertices of $G^{0,1}$ are all those vertices of $G(F)$ which can be written as $x_0 \cdot f_0 + x_1 \cdot f_1 + v$ where v is in T and both $x_0 \geq K$ and $x_1 \geq K$. Similarly, the vertices of G_j^0 are all those vertices of $G(F)$ which can be written as $x_0 \cdot f_0 + j \cdot f_1 + v$, where v is in T_j^0 and $x_0 \geq K_j^0$, etc.

Proof. We denote by V the (finite) set of all vectors $\sum_{f \in F_N} x_f \cdot f$ with each $x_f \leq k_f$. Note that it follows from Lemma 2 that every vertex v of $G(F)$ satisfies $v \cong w + \sum_{f \in F_E} x_f \cdot f$ for some $w \in V$.

For any two vectors $v, v' \in V$ which are congruent modulo $\Lambda(f_0, f_1)$, we consider the unique integers l, m such that $v - v' = l \cdot f_0 + m \cdot f_1$. Let I be the maximum of all $|l|$ and $|m|$ over all such pairs of vectors v, v'. Since V is finite, I is well defined. This definition assures the following fact:

LEMMA 6 *If $v \in G(F)$ is not in $G(F_E) \otimes T$ then $v \cong w + \sum_{f \in F_E} x_f \cdot f$ for some $w \in V$, where at least one of the coefficients x_f is strictly smaller than $I + J$.*

Proof: Indeed, we know that $v \cong w + \sum_{f \in F_E} x_f \cdot f$ for some $w \in V$. There exists a vertex u in the tree T congruent to w modulo $\Lambda(F_E)$, i.e., $w = u + \sum_{f \in F_E} y_f \cdot f$ with all $|y_f| \leq I$. Thus

$$v = u + \sum_{f \in F_E} (x_f + y_f) \cdot f.$$

According to Lemma 4, one of the coefficients $x_f + y_f$ is smaller than J, and hence at least one of the coefficients x_f must be strictly smaller than $I + J$. $\square$

We shall now proceed to define the trees T^i_j and the constants K^i_j. The vertices of all the trees T^i_j will be from V, and their edges will be of type from F_N.

Initially, we take all trees T^i_j equal to T, and all constants K^i_j equal to $K = I + 2J$. Note that these choices satisfy the conditions *1,2* but not necessarily *3*. (However, all but finitely many vertices of $G(F_E) \otimes T$ are in U.) We shall show that as long as *3* is not satisfied, one of the trees T^i_j can be enlarged while maintaining the conditions *1,2*. (Since V is finite, this process will terminate, proving the Theorem.) Our enlargement technique consists of repeatedly applying the following claim:

CLAIM. If *3* is not satisfied, then there exists a tree T^i_j, a vertex u in T^i_j, and a vertex $w \in V$ adjacent to u by an edge of type $f \in F_N$ such that

- almost all vertices $v \cong w + j \cdot f_{1-i} + \lambda \cdot f_i$ are missing from U, and
- for each such λ, we have $v \cong j \cdot f_{1-i} + \lambda \cdot f_i + u + f$.

Since only finitely many vertices $w + j \cdot f_{1-i} + \lambda \cdot f_i$ are in U, all vertices $w + j \cdot f_{1-i} + \lambda \cdot f_i$ from a certain value of λ on, are not in U. We increase the current constant K^i_j (if necessary) to be greater than this λ, and enlarge the current T^i_j to include the vertex w (and the edge uw). It is now easy to see that *1* still holds. It is also clear that our second condition of the Claim implies that *2* still holds.

Proof of the Claim: Since there are infinitely many vertices v of $G(F)$ missing from U (and hence from $G(F_E) \otimes T$), we can use Lemma 6 to express each of them as $v \cong w + j \cdot f_{1-i} + \lambda \cdot f_i$, with $i = 0, 1$ and $j < I + J$. In fact, since V is finite and since $j \leq I + J$, we can assume that they all take the same w and the same i, j. Thus there are finitely many triples w, i, j such that all vertices $v \cong w + j \cdot f_{1-i} + \lambda \cdot f_i$, with λ sufficiently large, are missing from U. We now add an extremality condition to assure that w is properly adjacent to some u in T_j^i.

For each triple w, i, j we find a 'connector' vertex u as follows: if $w \cong \sum_{f \in F_N} w_f \cdot f$, we let $u \cong \sum_{f \in F_N} u_f \cdot f$ to be a vertex in the tree T_j^0 which maximizes $d_F(u)$ while each $u_f \leq w_f$. (Such a u must exist, as the origin 0^n is in T_j^0; ties can be broken arbitrarily.)

We now claim that for at least one triple w, i, j, the corresponding connector vertex u satisfies $w \cong u + f$, with $f \in F_N$ which, together with $v \cong w + j \cdot f_{1-i} + \lambda \cdot f_i$, proves the Claim. This is accomplished by choosing a triple w, i, j for which the difference $d_F(w) - d_F(u)$ is as small as possible. If we can show that this difference is equal to one, then, since each $u_f \leq w_f$, we shall have $w \cong u + f$.

If this difference were greater than one, then for some $e \in F_N$ we have $u_e < w_e$ and $u + e \in V, u + e \neq w$. Thus there is an $r \in V$ such that $w \cong u + e + r$. Consider now vertices $z = u + e + j \cdot f_{1-i} + \lambda \cdot f_i$. We have $z \cong u + e + j \cdot f_{1-i} + \lambda \cdot f_i$. By the choice of w, i, j, we know that almost all these vertices z are in U. Otherwise, the triple $(u + e), i, j$, with the same connector u, would be better than w, i, j. Since there are infinitely many z and only finitely many of the subgraphs forming U, we know that infinitely many z (and hence all z with large enough λ) belong to the same subgraph $(G^{0,1}, G_a^b)$. We begin by showing that this subgraph is $G_{j'}^i$, where i is defined above, and j' is some integer $j' \leq I + 2J$. Clearly, it cannot be any $G_{j'}^{1-i}$. Furthermore, if it were $G^{0,1}$, then some $z = t + z_0 \cdot f_0 + z_1 \cdot f_1$ with both z_0, z_1 greater than or equal to K would satisfy $z = u + e + j \cdot f_{1-i} + \lambda \cdot f_i$, i.e., $z_0 \leq j + I < J + 2I = K$, contradiction. Thus we may assume that there is a vertex t in the tree T, such that for all sufficiently large μ,

$$z \cong t + j' \cdot f_{1-i} + \mu \cdot f_i \cong u + e + j \cdot f_{1-i} + \lambda \cdot f_i.$$

¿From $v \cong w + j \cdot f_{1-i} + \lambda \cdot f_i$ and $w \cong u + e + r$ above, we conclude $v \cong u + e + r + j \cdot f_{1-i} + \lambda \cdot f_i$, which, together with the displayed fact, implies that $v \cong t + r + j' \cdot f_{1-i} + \mu \cdot f_i$. Therefore, the triple $(t + r), i, j'$ (with the corresponding connector vertex t) is a better choice than w, i, j, contrary to our choice. □

The above arguments can be extended to the general case. One proceeds by reverse induction, filling in first 'holes of higher dimensions'. What follows is an outline of the general proof.

THEOREM 7 *There exists a finite set A, and for each $a \in A$ a vector o_a, a subset F_a of F', and a finite tree T_a consistent with respect to F_a, such that the digraphs $[o_a + G(F_a)] \otimes T_a, a \in A$*

1. *are pairwise vertex-disjoint,*
2. *each is isometric in $G(F)$, and*
3. *their union U contains almost all the vertices of $G(F)$.*

Suppose $K_k, k \in N$ are constants. We denote by A^k the set of all m-tuples $a = (a_0, a_1, \ldots, a_{m-1})$ with each $a_i \leq K_k$. For each $a \in A^k$, we denote by F_a^k the set of all vectors $f_i \in F_E$ such that $a_i = K_k$.

LEMMA 8 *For any k with $0 \leq k \leq m$ there exists a constant K_k, a vertex o_a^k in $G(F_E)$, and a tree T_a^k, consistent with F_a^k, such that the digraphs $G_a^k = [o_a^k + G(F_a^k)] \otimes T_a^k$,*

1. *are pairwise vertex-disjoint,*
2. *each is isometric in $G(F)$, and*
3. *their union U contains almost all those vertices v of G_F, $v \cong w + \sum_{f \in F_E} x_f \cdot f$, $w \in V$, which have at least k of the coefficients x_f greater than K_k.*

Clearly chosing $k = 0$ in the Lemma proves the Theorem.

Proof of the Lemma. We will prove the Lemma by reverse induction on k. When $k = m$, we proceed as follows: Let T, I and J be in Lemma 4 and Lemma 6, and choose $K_m = I + J$, $o_a^m = \sum_{0 \leq i \leq m-1} a_i \cdot f_i$, and each $T_a^m = T$.

We now suppose the lemma is true for all $k, d+1 \leq k \leq m$, and prove it for $k = d$. We choose $K_d = K_{d+1} + I$. For each a in A^d we let a' be the element of A^{d+1} with $a'_i = Max(a_i, K_{d+1})$. We initially choose $T^d_a = T^{d+1}_{a'}$, and $o^d_a = \sum_{0 \leq i \leq m-1} a_i \cdot f_i$. These choices again satisfy *1,2*, but not necessarily *3*. To enforce *3*, we shall, as in the case $n = 2$, enlarge the trees T^d_a by adding elements of V. Instead of increasing the constants K^i_j, we will change the vertices o^d_a (to maintain the validity of the condition *1*); these changes will always consist of adding some vectors of F^d_a to o^d_a.

CLAIM. If *3* is not satisfied, then there exists some $a \in A^d$ for which F^d_a has cardinality d, and there exists a vertex u in T^d_a and a vertex w in V adjacent to u by an edge of type $f \in F_N$, and there exists a constant c, such that if all $f \in F^d_a$ have $\lambda_f \geq c$, then $v \cong w + \sum_{f \in F^d_a} \lambda_f \cdot f + o^d_a \notin U$.

Proof of the Claim: We shall show, as before, that if *3* fails, then there exists a fixed F', subset of F_E of cardinality d, a fixed $w \in V$, and $n - d$ fixed constants $x_f < K_d$, such that almost all vertices $v \cong w + \sum_{f \in F'} \kappa_f \cdot f + \sum_{f \in F_E - F'} x_f \cdot f$ are missing from U. It follows from the induction hypothesis that, in fact, all x_f are smaller than K_{d+1}.

Define $a = (a_1, a_2, \ldots, a_n)$ as follows: let $a_i = x_{f_i}$ if $x_{f_i} < K_d$, and let $a_i = K_d$ if $x_{f_i} \geq K_d$. Note that $F^d_a = F'$, and that, for all λ_f sufficiently large, we have $v \cong w + \sum_{f \in F^d_a} \lambda_f \cdot f + o_a$. (We maintain $o^d_a = \sum_{0 \leq i \leq m-1} a_i \cdot f_i + \sum_{f \in F^d_a} x_f \cdot f$.) Once again we find, for each pair w, a, a connector vertex u: if $w \cong \sum_{f \in F_N} w_f \cdot f$, we let $u \cong \sum_{f \in F_N} u_f \cdot f$ be a vertex in the tree T^d_a which maximizes $d_F(u)$, while each $u_f \leq w_f$. It is again the case that for at least one pair w, a, the connector vertex u satisfies $w \cong u + f$, with $f \in F_N$. This fact, together with $v \cong w + \sum_{f \in F^d_a} \lambda_f \cdot f + o^d_a$, proves the Claim. We choose a pair w, a for which $d_F(w) - d_F(u)$ is as small as possible. If this difference were greater than one, then for some $e \in F_N$, we would have $u_e < w_e$, and $u + e \in V$, $u + e \neq w$. Thus there is an $r \in V$ such that $w \cong u + e + r$.

Consider again all vertices $z \cong u + e + \sum_{f \in F^d_a} \lambda_f \cdot f + o^d_a$. Either there exists a constant c such that each z with all $\lambda_f \geq c$ belongs to U, or there exists a constant c' such that each z with all $\lambda_f \geq c'$ does not belong to U. (Indeed either the missing vertices, or the present vertices form a space of dimension d.)

In the first case, the pair $(u + e), a$, (with connector u) is better than w, a. In the second case, each z with all λ_f large enough, belongs to the same subgraph G_b^d. We show that this subgraph is some G_b^d, with $F_b^d = F_a^d$. We may assume that there is a vertex t in the tree T_b^d, such that for all sufficiently large μ_f, we have

$$z \cong t + \sum_{f \in F_b^d} \mu_f \cdot f + o_b^d \cong u + e + \sum_{f \in F_a^d} \lambda_f \cdot f + o_a^d.$$

Clearly, F_a^d is a subset of F_b^d. Moreover $F_a^d = F_b^d$, by the definition of I, as $K_{d+1} = K_d + I$.

¿From $v \cong w + \sum_{f \in F_a^d} \lambda_f \cdot f + o_a^d$ and $w \cong u + e + r$, we conclude that $v \cong u+e+r+\sum_{f \in F_a^d} \lambda_f \cdot f + o_a^d$ and $v \cong t+r+\sum_{f \in F_b^d} \mu_f \cdot f + o_b^d$. Therefore, the pair $(t+r), b$ (with connector t) is better than w, a, contrary to our choice. □

It remains to tie everything together. Recall that F_E consists of m linearly independent vectors and F_N of any number of vectors which are convex combinations of the vectors in F_N. Theorem 7 constructs a set of shifts of various subgraphs $G(F') \otimes T'$ with $F' \subseteq F_E$ and T' consistent with respect to T', which collectively cover almost all vertices of $G(F)$. Theorem 3 explains how we can whisper in each of them with delay at most m (using the backbone scheme in $G(F')$). Thus (cf. Figure 3) when t is large enough to allow for informing the finite number of uncovered vertices, we inform all vertices of $G(F)$ of distance at most $t - m$ from 0^n, and exactly $\binom{t}{j}$ vertices of distance $t - j$, with $0 \leq j \leq m - 1$. In view of our Lemma 2 we conclude:

THEOREM 9 *For all sufficietly large t,*

$$\omega_F(t) = \sigma_F(t-m) + \sum_{0 \leq j \leq m-1} \binom{t}{j}.$$

□

COROLLARY 10 *For all sufficietly large t, $\omega_F(t)$ is a polynomial of degree m in t.*

In closing, we remark that, at least in two dimensions, our techniques seem well suited to solve cases where the extreme vectors (those not equal to a convex combination of other vectors) are not linearly independent, as long as the vectors lie in the same halfspace (halfplane). On the other hand, it seems that new tools will be needed to solve situations where the extreme vectors do not lie in one halfspace. Even in the case of two-dimensional templates consisting of just three vectors, not in the same halfplane, we don't have a general technique to find $\omega_F(t)$, or prove that it is a polynomial in t.

References

[1] S. Djelloul, Broadcasting in m-dimensional grid graphs with a given neighborhood template, Discrete Appl. Math. 53 (1994) 25-36.

[2] A. Farley and S. T. Hedetniemi, Broadcasting in grid graphs, Congressus Numerantium 21 (1978) 275-288.

[3] P. Fraigniaud and E. Lazard, Methods and problems of communication in usual networks, Discrete Applied Math. 53 (1994) 79-134.

[4] C. Garcia, Réseaux arithmétiques pour réseaux d'interconnection, Thesis, University of Nice, 1993.

[5] C. Garcia and C. Peyrat, Broadcasting in an infinite n-grid, Proc. 28-th HICSS Hawaii, vol. 2, pp. 574-583.

[6] S. M. Hedetniemi, S. T. Hedetniemi and A. L. Liestman, A survey of gossiping and broadcasting in communication networks, Networks 18 (1988) 319-349.

[7] P. Hell and A. L. Liestman, Broadcasting in one dimension, Discrete Applied Math. 21 (1988) 101-111.

[8] D. A. Klarner, Mathematical crystal growth II, Discrete Applied Math. 3 (1981) 113-117.

[9] C. S. Ko, On a conjecture concerning broadcasting in grid graphs, AMS Notices (1979) A-196-A-197.

[10] C. S. Ko, Broadcasting, graph homomorphisms, and chord intersections, Ph.D. Thesis, Rutgers University, 1979.

[11] G. W. Peck, Optimal spreading in n-dimensional rectilinear grid, Stud. Appl. Math. 62 (1980) 69-74.

[12] Q. F. Stout, Information spread in mesh-connected computers, manuscript 1981.