Average Long-Lived Memoryless Consensus: The Three-Value Case ⋆

Ivan Rapaport[1] and Eric Rémila[2]

[1] DIM-CMM (UMI 2807 CNRS), Universidad de Chile,
2120 Blanco Encalada, Santiago - Chile,
`rapaport@dim.uchile.cl`

[2] Université de Lyon, LIP (UMR 5668 CNRS-ENS, Université Lyon 1),
46 allée d'Italie 69364 Lyon Cedex 7 - France,
`eric.remila@ens-lyon.fr`

Abstract. We study strategies that minimize the *instability* of a fault-tolerant consensus system. More precisely, we find the strategy than minimizes the number of output changes over a random walk sequence of input vectors (where each component of the vector corresponds to a particular sensor reading). We analyze the case where each sensor can read three possible inputs. The proof of this result appears to be much more complex than the proof of the binary case (previous work). In the binary case the problem can be reduced to a minimal cut in a graph. We succeed in three dimensions by using the fact that an auxiliary graph (projected graph) is planar. For four and higher dimensions this auxiliary graph is not planar anymore and the problem remains open.

1 Introduction

There are situations where, for fault-tolerant purposes, a number of sensors are placed in the same location. Ideally, in such cases, all sensor readings should be equal. But this is not always the case; discrepancies may arise due to differences in sensor readings or to malfunction of some sensors. Thus, the system must implement some form of fault-tolerant averaging *consensus function* ϕ that returns a representative *output value* of the sensor readings.

Let us consider n sensors which are sampled at synchronous rounds. In each round an *input vector* x of sensor readings is produced, where x_i is a value from some finite set S produced by the i-th sensor. Assuming that at least $t+1$ entries of the vector are correct, ϕ is required to return a value that appears in at least $t+1$ entries of x.

The sampling interval is assumed to be short enough in order to guarantee the sequence of input vectors to be *smooth*: exactly one entry of a vector changes from one round to the next one.

⋆ Partially supported by Programs Fondecyt 1090156, Basal-CMM, Instituto Milenio ICDB, Ecos C09E04 and IXXI (Complex System Institute, Lyon).

A natural function ϕ is the one that returns the most common value of vector x. However, the *instability* of such function is high. More precisely, the output value computed by such ϕ could change from one round to the next one unnecessarily often.

For tackling this stability issue a worst case complexity measure was introduced in two previous papers [6, 8]. The input sequence considered in those papers was assumed to be, in addition to smooth, *geodesic*: the i-th entry of the input vector was allowed to change *at most once* over the sequence. The instability of a consensus function was given by the largest number of output changes over any such sequence, called a *geodesic path*.

In [1] we introduced, as an alternative measure, a more natural (and subtle) notion called *average instability*. We removed the geodesic requirement and therefore the smooth sequences of input vectors we considered were random walks over the hypercube. If $P = x_0, x_1, \ldots$ is such a walk, then the average instability of a consensus function ϕ is given by the fraction of time ϕ changes its output over P.

We studied in [1] the case when the input is binary ($S = \{0, 1\}$). In particular, for the memoryless case, we proved that function ϕ_0, that outputs 1 unless it is forced by the fault-tolerance requirement to output 0 (on vectors with t or less 1's), is optimal.

In the present paper we analyze the 3-value case ($S = \{0, 1, 2\}$). We extend previous result proving that the natural extension of ϕ_0 is in fact optimal among all the anonymous strategies (a strategy is anonymous, or symmetric, when it only depends on the number of entries of each type, and not in the respective places of the entries).

More precisely, let $\#_b(x)$ be the number of entries of x that are equal to $b \in \{0, 1, 2\}$. Let ψ be the consensus function defined as follows:

- $\psi(x) = 2$ if $\#_2(x) \geq t + 1$,
- $\psi(x) = 1$ if $\#_2(x) \leq t$ and $\#_1(x) \geq t + 1$,
- $\psi(x) = 0$ if $\#_2(x) \leq t$ and $\#_1(x) \leq t$.

In this paper we prove that the consensus function ψ has optimal stability according to the average criterion (among all the consensus symmetric memoryless functions).

The proof of this result is much more complex than the proof of the binary case. In the binary case the problem can be reduced to a minimal cut in a graph. For higher dimensions this approach can not be used. In fact, the problem becomes a multiterminal cut problem, which is NP-hard [5].

We succeed in 3 dimensions by using the fact that an auxiliary graph (projected graph) is planar. Unfortunately, the auxiliary graph is not planar in 4 (and higher) dimensions. We are therefore facing a hard combinatorial problem. In fact, the multiterminal cut problem is polynomial for planar graphs and NP-complete in the general case.

We would like to point out that Davidovitch et al [6] faced the same change in the difficulty level when they extended their binary case result to a more general, multi-valued result. In fact, for that extension, they had to use topological techniques in high dimensional complexes.

As noted in [8], studying the instability of consensus functions may have applications in various areas of distributed computing, such as self-stabilization [7] (indeed, see [11]), Byzantine agreement [2], real-time systems [12], complexity theory [10] (boolean functions), and VLSI energy saving [3, 14, 17] (minimizing the number of transitions).

2 The model

Let n, t, k be non-negative integers such that $n \geq kt+1$. The set of *input vectors* is the set $V = \{0, 1, ..., k-1\}^n$. The *input space* is the graph (V, E) where the edges E are all the (unordered) pairs of vectors which differ in exactly one component. Notice that $|E| = n(k-1)k^n/2$. In this paper we focus in the case $k = 3$.

The *distance* $d(x_1, x_2)$ between two vectors x_1, x_2 is the distance in the input space, i.e., the number of entries in which x_1 and x_2 differ. We denote by $\#_b(x)$ the number of entries of x that are equal to $b \in \{0, 1, 2\}$. The *corners* of the input space are the vectors 0^n, 1^n and 2^n.

Let $x_0, x_1, x_2, \ldots$ be a *walk* in the input space, i.e., a sequence $(x_i)_{i \in \mathbb{N}}$ of vectors in V such that for each $i \in \mathbb{N}$, the pair $\{x_i, x_{i+1}\}$ is an element of E.

A consensus decision function ϕ assigns, to each x_i, an *output value* $d \in \{0, 1, 2\}$. In this paper we only consider such *memoryless* functions (for which the decision depends only on the current input vector x_i, and not on the past history).

Formally, $\phi : V \to \{0, 1, 2\}$ is the *consensus decision function*. The fault-tolerance requirement that ϕ must satisfy is the following: $\phi(x) = d \Rightarrow \#_d(x) \geq t+1$.

A consensus decision function ϕ is *symmetric* if $\phi(x)$ depends only on the three values $\#_0(x)$, $\#_1(x)$,$\#_2(x)$ (which correspond to the number of 0s, 1s and 2s of x).

An *execution* of the system is a sequence $(x_0, d_0) \to (x_1, d_1) \to \ldots$, where $d_i = \phi(x_i)$.

In the sequel we will often refer to the "consensus function" instead of the "consensus symmetric memoryless decision function".

2.1 Stability: random walk criterion

In [1] we proposed a new criterion that uses random walks for determining the instability of a consensus function. The idea is the following. The initial input vector x_0 is chosen according to some distribution λ. On the other hand, if x_i is the current input vector, then the next input vector x_{i+1} is chosen in a random uniform way among the vectors at distance one from x_i.

Formally, we have a Markov process (P, μ_0) whose set of states is V and there is a transition from x to x' if $\{x, x'\} \in E$. The probability of such transition is

$1/2n$ (this defines the transition matrix P). The initial distribution is $\mu_0 = \lambda$. Each state x_i has an associated output value $d_i = \phi(x_i)$. Therefore the random walk $x_0, x_1, x_2, \ldots$ defines an execution.

We use the classical Kronecker notation $\delta(i,j)$ where $\delta(i,j) = 1$ when $i = j$, and $\delta(i,j) = 0$ otherwise. Let $c_{\lambda,l}(\phi)$ be the random variable defined by: $c_{\lambda,l}(\phi) = \frac{1}{l}\sum_{k=0}^{l-1} \delta(d_k, d_{k+1})$. The *average instability* of a consensus function ϕ is defined by:

$$inst(\phi) = \mathbb{E}(\lim_{l\to\infty} c_{\lambda,l}(\phi)).$$

The average instability represents the frequency of decision changes along a random execution (and, as an easy consequence of the classical Ergodic Theorem [15], it is well defined).

Furthermore, the stationary distribution π of the random walk $x_0, x_1, x_2, \ldots$ is the uniform distribution, i.e., $\pi_x = 1/3^n$ for every $x \in V$ (for notation see [15])).

Since the instability of ϕ counts the number of times the function ϕ changes its decision along a random walk, by the Ergodic Theorem this value tends to the number of bicolored edges (where changes in the decision take place) divided by $|E| = n3^n$. In other words,

$$inst(\phi) = \frac{\sum_{\{x,y\}\in E} \delta_\phi(\{x,y\})}{n3^n} = \frac{|E_\phi|}{n3^n},$$

where $\delta_\phi(\{x,y\}) = 1$ when $\phi(x) \neq \phi(y)$ and $\delta_\phi(\{x,y\}) = 0$ otherwise, and E_ϕ denotes the set of edges e such that $\delta_\phi(e) = 1$ (i.e., E_ϕ is the set of bicolored edges). A detailed proof of the equality above was given for the binary case in [1]. The proof of the 3-value case is identical.

3 Basics for the symmetric case

We study here symmetric systems. For these systems it is convenient to use the *projected input space* (V', E', w), which is a weighted graph.

Each vertex v of the projected input space V' can be seen as a triple (i,j,k) with i,j,k nonnegative, and $i + j + k = n$. Each of these triplets represent the set of $\frac{n!}{i!j!k!}$ input vectors containing i 0-entries, j 1-entries and k 2-entries. We say that i, j and k are, respectively the *0-component*, the *1-component* and the *2-component* of v.

On the other hand, two distinct vertices (i,j,k) and (i',j',k') are linked by an edge of E' when $|i - i'| \leq 1$, $|j - j'| \leq 1$ and $|k - k'| \leq 1$ (remark that we necessarily have $i = i'$ or $j = j'$ or $k = k'$).

The weight $w(\{v, v'\})$ of the edge $\{v, v'\}$ is the number of edges $\{x, x'\}$ of E such that v is the projection of x and v' is the projection of x'. For $i < n$ and $j > 0$, the weight of the edge $\{(i,j,k),(i+1,j-1,k)\}$ is given by the equality:

$$w(\{(i,j,k),(i+1,j-1,k)\}) = \frac{n!}{i!(j-1)!k!} \quad (1)$$

For each symmetric function ϕ, we define $\phi(v) = \phi(x)$ where x is any input vector such that v is the projection of x. Therefore, in the symmetric case, the last equality of the previous section becomes:

$$inst(\phi) = \frac{\sum_{\{v,v'\}\in E'} w(\{v,v'\})\delta_\phi(\{v,v'\})}{n3^n} = \frac{\sum_{e'\in E'_\phi} w(e')}{n3^n},$$

where $\delta_\phi(\{v,v'\}) = 1$ when $\phi(v) \neq \phi(v')$, $\delta_\phi(\{v,v'\}) = 0$ otherwise, and E'_ϕ denotes the set of edges e' of E' such that $\delta_\phi(e') = 1$ (in other words, E'_ϕ is the set of bicolored edges in the projected input graph). It follows that, when we want to minimize the instability, we have to minimize the quantity $\sum_{e'\in E'_\phi} w(e')$.

The structure of the projected input space is easily understandable. It can be seen as a part of the triangular lattice delimited by an equilateral triangle of side length n, whose extreme points are the corners $(n, 0, 0)$, $(0, n, 0)$ and $(0, 0, n)$ In particular, this graph is planar. More precisely, we will use the following drawing: we take three points a, b, c, of the plane $\mathbb{R}^2$, usually $a = (0, 0), b = (n, 0)$ and $c = (\frac{n}{2}, \frac{n\sqrt{3}}{2})$. Each vertex (i, j, k) is identified with the point p_{ijk} which is the barycenter of (a, i), (b, j), and (c, k). Edges are classically represented by lines segments linking neighbor vertices. This representation is called the *canonical representation* of the projected input space.

We associate vectors to edges and faces: the vector corresponding to the edge $e = \{v, v'\}$ is the average between its endpoints vectors, i.e., the vector corresponding to the center of the line segment $[v, v']$. Hence, the vector of the edge $e = \{v, v'\}$ is formed by two semi-integer values and one integer value. The vector of a (finite) triangular face ϕ is the average between its vertices, i.e., the vector corresponding to the center of f. This vector is of the form (x, y, z), with $3x$, $3y$ and $3z$ being all integers (moreover, one can check that $3x$, $3y$ and $3z$ are equal modulo 3).

We recall that the dual graph of the projected input space (V', E') is the (multi)graph (F, E') such that F is the set of faces induced by (V', E'), and a pair $\{f, f'\}$ is an edge if there exists $\{v, v'\} \in E'$ such that the line segment $[v, v']$ is shared by both f and f'. As it is usually done, we refer indistinctly to the edge $\{f, f'\}$ of the dual graph and to the same edge $\{v, v'\}$ of the projected input space. Since the canonical representation of the projected input space is part of the triangular lattice of the plane, the canonical representation of the dual graph is part of the hexagonal lattice, with all pending edges linked to a particular vertex of the dual graph, which is the infinite face f_∞.

Let i be an integer. We define the edge set $E_{x=\frac{i}{2}}$ as the set of edges whose 0-component is $\frac{i}{2}$. In other words, $E_{x=\frac{i}{2}}$ is the set of edges intersecting the closed line segment $[(\frac{i}{2}, n - \frac{i}{2}, 0), (\frac{i}{2}, 0, n - \frac{i}{2})]$. Notice that an edge intersects this closed line segment if and only if the corresponding edge of the dual graph also intersects the same line segment (with the convention that the vertices of the dual graph are placed in the center of the triangular faces). One can define, in a similar way, for any integer j, $E_{y=\frac{j}{2}}$ and, for any integer k, $E_{z=\frac{k}{2}}$ (using respectively the 1-component and the 2-component).

Lemma 1. *Let e be an edge of $E_{x=\frac{i}{2}}$ with vector $(\frac{i}{2}, \frac{j}{2}, \frac{2n-i-j}{2})$.*

- *Let $sym(e)$ be the edge with vector $(\frac{i}{2}, \frac{2n-i-j}{2}, \frac{j}{2})$. We have $w(e) = w(sym(e))$.*
- *For i odd, and $j+1 \leq \frac{2n-i}{2}$, let e_+ denote the edge with vector $(\frac{i}{2}, \frac{j+1}{2}, \frac{2n-i-j-1}{2})$*
  - *If j is even, then $w(e) = w(e_+)$,*
  - *If j is odd, then $w(e) < w(e_+)$.*
- *For i even (which enforces j being odd) and $j + 2 \leq \frac{2n-i}{2}$, let e_+ denote the edge with vector $(\frac{i}{2}, \frac{j+2}{2}, \frac{2n-i-j-2}{2})$. We have $w(e) < w(e_+)$.*

Proof. This lemma is a direct consequence of equality 1.

The lemma above describes the evolution of weights of edges in $E_{x=\frac{i}{2}}$. First, it says that there is a weight symmetry with respect to the median axis, formed by points (x, y, z) such that $y = z$. Secondly, the lemma also says that the weight increases as you move towards the median axis.

Each consensus function ϕ divides the vertices of the projected input space (and therefore the plane since the graph is planar) into three *zones*, one for each output value. Each zone contains the corresponding corner. Notice that zones are not necessarily connected.

Consider the set E_ϕ of bicolored edges induced by ϕ or, more precisely, consider the corresponding edges in the dual graph. These edges form cycles, which surround connected components of the zones.

We call *network* a subset of edges. Given a network N, the weight $w^N_{x=\frac{i}{2}}$ is defined by $w^N_{x=\frac{i}{2}} = min\{w(e) \,|\, e \in E_{x=\frac{i}{2}} \cap N\}$ (we say that $w^N_{x=\frac{i}{2}} = \infty$ when $E_{x=\frac{i}{2}} \cap N$ is empty). This weight represents the minimal necessary weight for passing from one side of the segment $[(\frac{i}{2}, n - \frac{i}{2}, 0), (\frac{i}{2}, 0, n - \frac{i}{2})$ to the other part still remaining in the network N.

Lemma 2. *Consider a simple path p (of the dual graph) linking a finite face f_0 to the infinite face f_∞, remaining in a network N, whose last edge belongs to $E_{x=0}$. Let (i_0, j_0, k_0) be the vector corresponding to f_0. The sum $\sum_{e \in p} w(e)$ is denoted by $w(p)$. We have the inequality:*

$$w(p) \geq \sum_{0 \leq i < 2i_0} w^N_{x=\frac{i}{2}}$$

Proof. This is obvious since, for each integer i such that $0 \leq i < 2i_0$, the path must contain an edge of $E_{x=\frac{i}{2}}$ and the sets $E_{x=\frac{i}{2}}$ are pairwise disjoint.

4 Our result

The following theorem is the main result of the paper.

Theorem 1. *Consider the consensus function ψ defined by:*

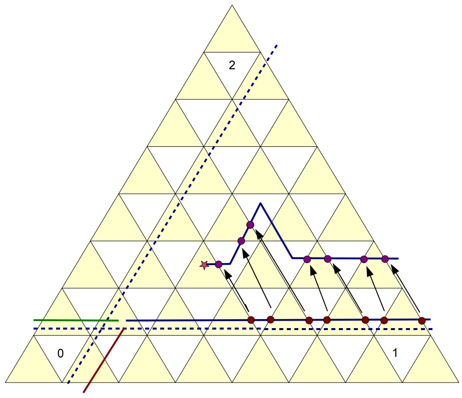

Fig. 1. An application of Lemma 2, with the network N_0 formed by edges whose 1-components and 2-components are both at least 3/2. The weight of the path issued from the star is larger than the weight of the marked edges in the horizontal path issued from the cell of vector $(n - 2t - \frac{2}{3}, t + \frac{1}{3}, t + \frac{1}{3})$ (notice that this path is a boundary path induced by the consensus function ψ).

– $\psi(x) = 2$ *if* $\#_2(x) \geq t + 1$,
– $\psi(x) = 1$ *if* $\#_2(x) \leq t$ *and* $\#_1(x) \geq t + 1$,
– $\psi(x) = 0$ *if* $\#_2(x) \leq t$ *and* $\#_1(x) \leq t$.

The consensus function ψ has optimal stability according to the average criterion, among all symmetric functions.

We decompose the proof into two lemmas from which the theorem is a direct consequence. We first limit ourselves to the case when each of the three zones formed from vertices with the same output value is connected in the merged input gragraph.

Lemma 3. *Let ϕ be a consensus function for which each of the three zones induced by ϕ is connected. We have $inst(\phi) \geq inst(\psi)$.*

Proof. In this case, E'_ϕ is just formed (in the dual graph) by three edge disjoint paths linking a face $f_0 = (i_0, j_0, k_0)$, with $i_0 > t, j_0 > t$, and $k_0 > t$ to the infinite face f_∞. We call p_0 the path linking f_0 to f_∞ and crossing the set $E_{x=0}$. This path is the boundary between the 1-zone and the 2-zone.

We denote by N_0 the network formed by edges whose vector (x, y, z) are such that $y > t$ and $z > t$. The path p_0 remains in N_0 thus, applying Lemma 1, we get:

$$w(p_0) \geq \sum_{0 \leq i < 2i_0} w^{N_0}_{x=\frac{i}{2}}.$$

In the same way we get $w(p_1) \geq \sum_{0 \leq j < 2j_0} w^{N_1}_{y=\frac{j}{2}}$ and $w(p_2) \geq \sum_{0 \leq k < 2k_0} w^{N_2}_{z=\frac{k}{2}}$.

Adding these inequalities, we get a lower bound for $w(E'_\phi) = w(p_0) + w(p_1) + w(p_2)$. But this bound is not sufficient to get our result. We need a refinement of Lemma 2 using other separating edge sets.

Up to symmetry, it can be assumed that $i_0 \geq j_0$. We define the integer j_1 as the lowest integer such that $j_1 \geq 2j_0$. For $2t < k < 2k_0$, we state $i_k = 2n - k - j_1$ in such a way that $L_k = (i_k/2, j_1/2, k/2)$ and $R_k = (j_1/2, i_k/2, k/2)$ have semi-integer or integer coordinates.

The set $E^{f_0}_{z=k/2}$ is formed by edges listed below (see Figure 2):

– the edges of $E_{z=k/2}$ whose whose 1-component is at least j_1 and whose 0-component is at least j_1 (i.e. the edges which meet the line segment $[L_k, R_k]$).
– the edges of $E_{x=i_k/2}$ whose 2-component is at most k (i.e., the edges which intersect the line segment $[L_k, (i_k/2, 0, n - i_k/2)]$).
– the edges of $E_{j=i_k/2}$ whose 2-component is at most k (i.e., the edges which intersect the line segment $[R_k, (0, i_k/2, n - i_k/2)]$).

Given a network N, the weight $u^{N\,f_0}_{z=k/2}$ is defined by: $u^{N\,f_0}_{z=k/2} = min\{w(e'), |e' \in E^{f_0}_{z=k/2} \cap N\}$. This weight represents the minimal necessary weight for passing from one side of the cut $E^{f_0}_{z=k/2}$ to the other part still remaining in the network N.

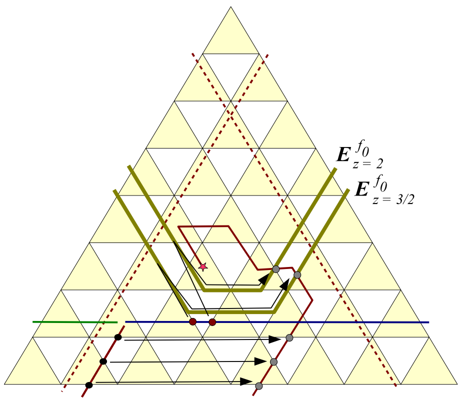

Fig. 2. The improvement of the argument of Lemma 2. We use some "broken lines". For each marked edge of E'_ψ, one can find an edge of larger weight on the path issued from the star.

Consider the path p_2 (of the dual of the projected input graph) linking the face f_0 to the infinite face f_∞ and remaining in the network N_2. We have:

$$w(p_2) \geq \sum_{0 \leq k \leq 2t} w^{N_2}_{z=\frac{k}{2}} + \sum_{2t < k \leq k_0} u^{N_2\,f_0}_{z=\frac{k}{2}},$$

because the considered cut sets of edges are pairwise disjoints and p_2 contains at least an edge of each of these sets.

From Lemma 1, we have: $u^{N_2\, f_0}_{z=\frac{k}{2}} = w^{N_0}_{x=\frac{i_k}{2}}$, since an edge of minimal weight for the two corresponding edge sets ($E^{f_0}_{z=\frac{k}{2}} \cap N_2$ and $E_{x=\frac{i_k}{2}} \cap N_0$) is the edge of vector $(\frac{i_k}{2}, \frac{2t+1}{2}, \frac{2n-i_k-2t-1}{2})$. Moreover, $2t < k < 2k_0$ if and only if $2n - 2k_0 - j_1 \leq i_k \leq 2n - 2t - j_1$. From the definition of j_1 this exactly means $2n - 2k_0 - 2j_0 \leq i_k \leq 2n - 2t - 2j_0$, i.e. $2i_0 \leq i_k \leq 2n - 2t - 2j_0$. Therefore, $\sum_{2t<k\leq k_0} u^{N_2\, f_0}_{z=\frac{k}{2}} = \sum_{2i_0 \leq i \leq 2n-2t-2j_0} w^{N_0}_{x=\frac{i}{2}}$, which gives:

$$w(p_2) \geq \sum_{0\leq k\leq 2t} w^{N_2}_{z=\frac{k}{2}} + \sum_{2i_0\leq i\leq 2n-2t-2j_0} w^{N_0}_{x=\frac{i}{2}}$$

On the other hand, for $2t < j < 2j_0$, we have, from Lemma 1, $w^{N_1}_{y=\frac{j}{2}} = w^{N_0}_{x=\frac{2n-2t-j}{2}}$: an edge of minimal weight in the edge set $E_{y=\frac{j}{2}} \cap N_1$ is the edge e of vector $(\frac{2n-j-2t-1}{2}, \frac{j}{2}, \frac{2t+1}{2})$, an edge of minimal weight in the edge set $E_{x=\frac{2n-2t-j}{2}} \cap N_0$ is the edge e' of vector $(\frac{2n-j-2t}{2}, \frac{j+1}{2}, \frac{2t+1}{2})$, and $w(e) = w(e')$.

Thus, $\sum_{2t<j<2j_0} w^{N_1}_{y=\frac{j}{2}} = \sum_{2t<j<2j_0} w^{N_0}_{x=\frac{2n-2t-j}{2}} = \sum_{2n-2t-2j_0<i<2n-4t} w^{N_0}_{x=\frac{i}{2}}$, which gives:

$$w(p_1) \geq \sum_{0\leq j\leq 2t} w^{N_1}_{y=\frac{j}{2}} + \sum_{2n-2t-2j_0<i<2n-4t} w^{N_0}_{x=\frac{i}{2}}.$$

We recall that

$$w(p_0) \geq \sum_{0\leq i<2i_0} w^{N_0}_{x=\frac{i}{2}}.$$

We have $w(E'_\phi) = w(p_0) + w(p_1) + w(p_2)$, thus, adding the three previous main inequalities, we get:

$$w(E'_\phi) \geq \sum_{0\leq i<2n-4t} w^{N_0}_{x=\frac{i}{2}} + \sum_{0\leq j\leq 2t} w^{N_1}_{y=\frac{j}{2}} + \sum_{0\leq k\leq 2t} w^{N_2}_{z=\frac{k}{2}}.$$

But $\sum_{0\leq i<2n-4t} w^{N_0}_{x=\frac{i}{2}}$ is exactly the weight of the path separating the 2-zone and the 1-zone for the function ψ, $\sum_{0\leq j\leq 2t} w^{N_1}_{y=\frac{j}{2}}$ is exactly the weight of the path separating the 2-zone and the 0-zone for the function ψ, and $\sum_{0\leq k\leq 2t} w^{N_2}_{z=\frac{k}{2}}$ is exactly the weight of the path separating the 1-zone and the 0-zone for the function ψ. Thus the second member of the equality is exactly $w(E'_\psi)$. We have: $w(E'_\phi) \geq w(E'_\psi)$.

Lemma 4. *Let ϕ be a consensus function. There exists a consensus function ϕ' for which each the induced zones induced by ϕ' are connected, such that $inst(\phi') \leq inst(\phi)$.*

Proof. Actually, we prove a (little bit) stronger fact: if a zone induced by ϕ is not connected, then there exists a consensus function ϕ' such that $inst(\phi') < inst(\phi)$.

If we apply this fact a sufficient number of times, then we will end up with consensus function with connected zones (obtaining the lemma), since the set of consensus functions is finite.

We call, for short, *0-domain* any connected component of the 0-zone (we define in a same way 1-domains and 2-domains). Assume that the 0-zone of ϕ is not connected and let D be a 0-domain which does not contain the vertex $(n, 0, 0)$. Let (i_1, j_1, k_1) be a vertex of D with j_1 minimal and (i_2, j_2, k_2) be a vertex of D with k_2 minimal.

First assume that $j_1 \geq t+1$. If there exists an edge $e = \{v, v'\}$, bicolor for ϕ, with $v \in D$ and $\phi(v') = 2$, then the function ϕ_2, defined by $\phi_2(v) = 2$ for $v \in D$, and $\phi_2(v) = \phi(v)$ otherwise, is a consensus function such that $inst(\phi_2) < inst(\phi)$: we are done. Otherwise, each edge $e = \{v, v'\}$, bicolor for ϕ, with $v \in D$ is such that $\phi(v') = 1$. This enforces that, for any such edge, the 2-component of v is at least $t+1$. Thus the function ϕ_1, defined by $\phi_1(v) = 1$ for $v \in D$, and $\phi_2(v) = \phi(v)$ otherwise, is a consensus function, such that $inst(\phi_1) < inst(\phi)$: we are done.

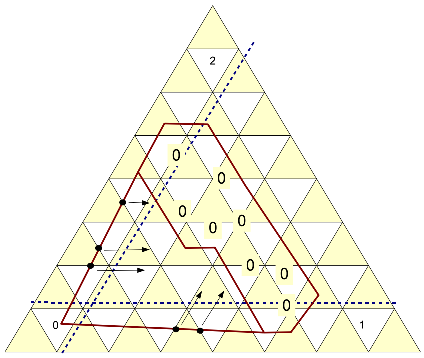

Fig. 3. the function ϕ' used in the tricky case of Lemma 4. The marked edges are some potential edges of $E_{\phi'} \setminus E_{\phi}$. Arrows indicate the way for fetching $i(e)$.

A similar argument can be used if it is assumed that $k_2 \geq t+1$. It remains to treat the tricky case when $j_1 \leq t$ and $k_2 \leq t$. We state $V_{j_1,k_2} = \{(i,j,k) \in \mathbb{N}^3 \mid i+j+k = n, j_1 \leq j, k_2 \leq k\}$. In this case, since D is connected, one can fix a simple path p in D from (i_1, j_1, k_1) to (i_2, j_2, k_2). Let $support(p)$ denote the set of vertices appearing in p. From the Jordan curve theorem [9, **?**], the path p divides $V_{j_1,k_2} \setminus support(p)$ into two parts, both connected in the projected input graph. Let D_p be the part which contains the vertex $(n-2t, t, t)$ (this vertex is not element of $support(p)$ since it is not element of D). Notice that the 0-component of any vertex of D_p is at least $t+1$. Thus, the function ϕ', defined by $\phi'(v) = 0$ for $v \in D_p$, and $\phi'(v) = \phi(v)$ otherwise, is a consensus function (see Figure 3).

The set $E'_{\phi'} \setminus E'_{\phi}$ can be partitioned into two sets E_j and E_k, where $e = \{v, v'\}$ is element of V_k if v is a vertex of D_p whose 2-component is k_2 and v' is a vertex whose 2-component is $k_2 - 1$ such that $\phi(v') = 1$, and, on the symmetric way, $e = \{v, v'\}$ is element of V_j if v is a vertex of D_p whose 1-component is j_1 and v' is a vertex whose 1-component is $j_1 - 1$ such that $\phi(v') = 2$.

Take an edge $e = \{v, v'\}$ of V_k, and let $(\frac{2n-j-2k_2+1}{2}, \frac{j}{2}, \frac{2k_2-1}{2})$ be its vector. Let k' be the lowest integer such that $k' \geq 2k_2 - 1$ and the edge of vector $(\frac{2n-j-k'}{2}, \frac{j}{2}, \frac{k'}{2})$ is bicolor for ϕ. This last edge is denoted by $i(e)$ (informally, $i(e)$ is the edge of $E_{y=\frac{j}{2}}$ which allows to go out of the connected component containing e for ϕ). We have $j \geq 2t + 1$, since $i(e)$ is bicolor, thus we have $k' \leq 2n - 2t - 1$. Thus $2k_2 - 1 < k' < 2n - 2t - 1$, which gives, from Lemma 1, $w(e) < w(i(e))$. Moreover, $i(e)$ is not element of $E_{\phi'}$ since, by definition, both vertices of $i(e)$ are elements of $D_p \cup support(p)$.

In a similar way, for each edge $e = \{v, v'\}$ of V_j, let $(\frac{2n-2j_1-k+1}{2}, \frac{2j_1-1}{2}, \frac{k}{2})$ be its vector. Let j' be the lowest integer such that $j' \geq 2j_1 - 1$ and the edge of vector $(\frac{2n-j'-k}{2}, \frac{j'}{2}, \frac{k}{2})$ is bicolor for ϕ. This last edge is denoted by $i(e)$ (informally, $i(e)$ is the edge of $E_{z=\frac{k}{2}}$ which allows to go out of the connected component containing e for ϕ). We have: $w(e) < w(i(e))$ and moreover, $i(e)$ is not element of $E_{\phi'}$.

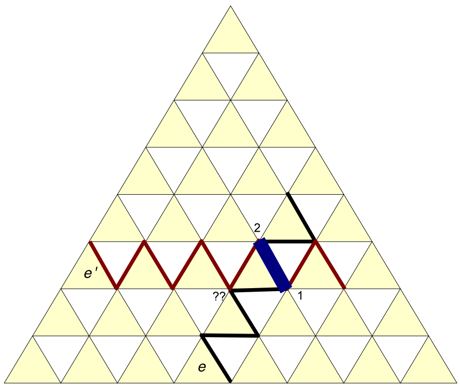

Fig. 4. Injectivity of the mapping i in Lemma 4, in the tricky case when i and j both are odd. If $i(e) = i(e')$, the $i(e)$ is the common edge. Each vertex v from e to the common edge is such that $\phi(v) = 1$ and, by symmetry, each vertex v' from e' to the common edge is such that $\phi(v') = 2$. Contradiction (at the vertex with question marks).

The other edges of $E_{\phi'}$ are contained in E_{ϕ}. For these edges, we state: $i(e) = e$. We claim that the mapping i is injective. To see it, we only have to check that if e is an edge of V_k and e' is an edge of V_j, then $i(e) \neq i(e')$. This can be easily done by a case by case analysis according to the parity of integers j and k such that $e' \in E_{y=\frac{j}{2}}$ and $e' \in E_{z=\frac{k}{2}}$ (see the argument in Figure 4).

In any case, we have: $w(e) \leq w(i(e))$. Moreover, we have: $V_k \neq \emptyset$ since, otherwise, we have $\phi(v) = 0$ for any v of the type (i, j, k_2), which contradicts the fact that the connected domain D does not contains $(n, 0, 0)$. Thus there is at least an edge e of $E'_{\phi'}$ for which $w(e) < w(i(e))$. Thus, by addition, we get $w(E'_{\phi'}) < w(E'_{\phi})$, which is the result.

References

1. F. Becker, S. Rajsbaum, I. Rapaport and E. Rémila, "Average Binary Long-Lived Consensus: Quantifying the Stabilizing Role Played by Memory," *Proc. SIROCCO*, Springer LNCS 5058, pp. 48-60, 2008.
2. P. Berman and J. Garay, "Cloture votes: $n/4$-resilient distributed consensus in $t+1$ rounds," *Math. Sys. Theory,* vol. 26(1): pp. 3–19, 1993.
3. A.P. Chandrakasan and R.W. Brodersen, *Low power digital CMOS design,* Kluwer Academic Publishers, 1995.
4. D. I. A. Cohen, "On the Sperner lemma," *J. Combin. Theory* vol. 2, pp. 585 – 587, 1967.
5. E. Dalhaus, D. S. Johnson, C. H. Papadimitriou, P. D.Seymour and P. Yannakakis, "The Complexity of Multiterminal Cuts," *SIAM J. on Computing,* vol. 23(4), pp. 864–894, 1994.
6. L. Davidovitch, S. Dolev and S. Rajsbaum, "Stability of Multi-Valued Continuous Consensus," *SIAM J. on Computing,* vol. 37(4), pp. 1057–1076, 2007.
7. S. Dolev, *Self-Stabilization*, the MIT Press, March 2000.
8. S. Dolev and S. Rajsbaum, "Stability of Long-lived Consensus," *J. of Computer and System Sciences*, Vol. 67, Issue 1, pp. 26-45, August 2003. Preliminary version in *Proc. of the 19th Annual ACM Symp. on Principles of Distributed Computing*, (PODC 2000), pp. 309–318, 2000.
9. C. Jordan, Cours d'analyse, Ecole Polytechnique, pp. 587-594, 1887.
10. J. Kahn, G. Kalai and N. Linial, "The Influence of Variables on Boolean Functions," *Proc. of the IEEE FOCS*, pp. 68–80, 1988.
11. S. Kutten and T. Masuzawa, "Output Stability Versus Time Till Output," *Proc. DISC*, Springer LNCS 4731, pp. 343–357, 2007.
12. H. Kopetz and P. Veríssimo, "Real Time and Dependability Concepts," chapter 16, pp.411–446 in Sape Mullender (ed.), *Distributed Systems,* ACM Press, 1993.
13. R. Maehara, "The Jordan Curve Theorem via the Brouwer Fixed Point Theorem", *American Mathematical Monthly*, Vol. 91, pp. 641–643, 1984.
14. E. Musoll, T. Lang and J. Cortadella, "Exploiting the locality of memory references to reduce the address bus energy," *Proc. of the Int. Symp. on Low Power Electronics and Design,* pp. 202–207, Aug. 1997.
15. J. R. Norris, *Markov Chains,* Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press, Oct. 1998.
16. E. Sperner, "Neuer Beweis f .ur die Invarianz der Dimensionszahl und des Gebietes, Abh." *Math. Sem. Univ. Hamburg*, Vol. 6, pp. 265 – 272, 1928
17. C-L. Su, C-Y. Tsui and A.M. Despain, "Saving power in the control path of embedded processors," *IEEE Design & Test of Comp.,* pp. 24–30, 1994.