

Fixation for Distributed Clustering Processes

M. R. Hilário*, O. Luidor†, C. M. Newman‡,
L. T. Rolla‡, S. Sheffield§, V. Sidoravicius¶*

January 13, 2010

Abstract

We study a discrete-time resource flow in \mathbb{Z}^d , where wealthier vertices attract the resources of their less rich neighbors. For any translation-invariant probability distribution of initial resource quantities, we prove that the flow at each vertex terminates after finitely many steps. This answers (a generalized version of) a question posed by van den Berg and Meester in 1991. The proof uses the mass-transport principle and extends to other graphs.

This preprint has the same numbering for sections, theorems, equations and figures as the published article “*Comm. Pure Appl. Math.* 63 (2010): 926–934.”

AMS Subject Classifications: 60K35, 90B15, 68M14.

Keywords: clustering, mass transport, cellular automaton.

1 The model and results

We consider the following model for distributed clustering. Initially each vertex $x \in \mathbb{Z}^d$ is assigned a random amount of resource, $0 \leq C_0(x) \leq \infty$, sampled according to a translation-invariant distribution. We denote $x \sim y$ if x, y are adjacent on \mathbb{Z}^d , write $x \simeq y$ if $x = y$ or $x \sim y$, and take $\mathcal{G}_x = \{y : y \simeq x\}$. At step $n = 0, 1, 2, \dots$, each vertex x holding some amount of resource will transfer all its resource to the vertex $a_n(x) \in \mathcal{G}_x$

*Instituto de Matemática Pura e Aplicada, Rio de Janeiro, Brazil

†Courant Institute of Mathematical Sciences, New York University, New York, NY USA

‡Département de Mathématiques et Applications, École Normale Supérieure, Paris, France

§Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA USA

¶Centrum voor Wiskunde en Informatica, Amsterdam, Netherlands

with the maximal amount of resource. All the vertices update simultaneously, leading to the state $(C_{n+1}(x), x \in \mathbb{Z}^d)$, where $C_{n+1}(x)$ is the sum of resources transferred to x at step n . More precisely, we take $a_n(x) = \operatorname{argmax}\{C_n(y), y \in \mathcal{G}_x\}$ if $C_n(x) > 0$, and otherwise $a_n(x) = x$. If there is more than one $y \in \mathcal{G}_x$ that attains the maximum, we say there is a *tie* at x and in this case $a_n(x)$ is chosen uniformly at random among the maximizing vertices. Finally we take $E_n(y) = \{x \in \mathcal{G}_y : a_n(x) = y\}$ and $C_{n+1}(y) = \sum_{x \in E_n(y)} C_n(x)$.

Note that, except for the possible tie breaking during the dynamics, all the randomness is contained in the initial data. Let \mathbb{P} and \mathbb{E} denote the underlying probability measure and its expectation, for both the initial resource quantities and possible tie breaks.

This model is a simple example of a self-organized structure that emerges from a disordered initial state. Instances of this type of phenomenon in several fields of science are mentioned in [4], where the model was introduced.

We are interested in the following phenomena, concerning the stability properties of this dynamics.

Question 1. Does each vertex transfer its resource eventually to the same fixed vertex?

Question 2. Does the the flow at each vertex terminate after finitely many steps?

Question 3. If the answer to the previous question is affirmative, is the expected value of the final resource equal to the expected value of the initial resource?

Van den Berg and Meester [3] considered this model on \mathbb{Z}^2 with continuously-distributed i.i.d. initial resource quantities and answered Question 1. Namely, they showed that $\mathbb{P}[a_n(x) \text{ is eventually constant}] = 1$. For the case of i.i.d. initial distributions supported on \mathbb{N} (see [3] for the precise hypotheses), they also answered Question 2 in any dimension, i.e., they proved that $\mathbb{P}[a_n(x) = x \text{ eventually}] = 1$.

Again for i.i.d. continuous initial distributions on the two-dimensional lattice, van den Berg and Ermakov [1] considered some percolative properties of the configuration after one step, and reduced Question 3 to a finite computation. As the calculation would be too big, even for the most powerful computers, they performed Monte Carlo simulations, obtaining overwhelming evidence that the answer to this question is positive.

In this paper we answer Question 2 in a rather general setting. We consider any dimension and allow any translation-invariant initial distribution. Question 3 remains open.

Theorem 1.1. *On \mathbb{Z}^d , for any translation-invariant distribution of the initial resources, the flow at each vertex almost surely stops after finitely many steps.*

The proof of local fixation essentially consists of ruling out, one by one, all the possibilities which could lead to a different scenario. For this goal we use the mass-transport principle in different ways.

In Section 2 we prove Theorem 1.1. In Section 3 we discuss generalizations of our results and conclude with a remark about Question 3.

2 Proof of Theorem 1.1

Theorem 1.1 will be proved using a combination of lemmas. We begin by introducing some extra notation.

Let $|\cdot|$ denote the cardinality of a set. For each $x \in \mathbb{Z}^d$ and $n = 0, 1, 2, \dots$, one and only one of the following events will happen (see Figure 2.1): $\mathcal{A}_n(x) = [C_n(x) = 0]$, $\mathcal{B}_n(x) = [C_n(x) > 0, E_n(x) = \{x\}]$, $\mathcal{C}_n(x) = [E_n(x) = \{z\} \text{ for some } z \neq x]$, $\mathcal{D}_n(x) = [|E_n(x)| > 1]$, $\mathcal{E}_n(x) = [E_n(x) = \emptyset]$. For $n = 0, 1, 2, \dots$ we also set

$$\begin{aligned} C'_n(x) &:= C_n(a_n(x)), \\ L_{n+1}(v) &:= a_n(L_n(v)), \text{ with } L_0(v) = v, \\ S_n(w) &:= \{v : L_n(v) = w\}. \end{aligned}$$

Observe that $C'_n(x) \geq C_n(x)$. $L_n(v)$ is the location at time n of the resource that initially started at vertex v , and $S_n(w)$ denotes the set of all vertices whose initial resource is located at vertex w at time n .

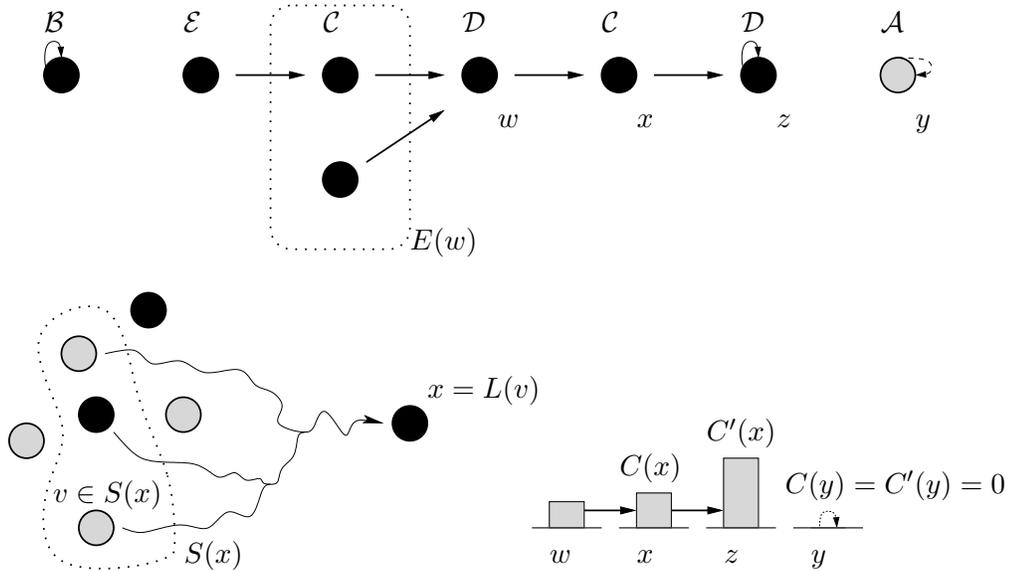


Figure 2.1: Schematic diagram for the notation used in this paper, omitting the index n .

Lemma 2.1 ([3], p.337). *For each x , $\mathbb{P}(\mathcal{D}_n(x) \text{ i.o.}) = \mathbb{P}(\mathcal{E}_n(x) \text{ i.o.}) = 0$.*

Proof. We repeat the proof of [3] for the convenience of the reader.

First, $\mathcal{E}_n(x)$ can happen for at most one value of n , after which one always has $\mathcal{A}_n(x)$. By translation-invariance one has $\mathbb{E}(|E_n(x)|) = \sum_w \mathbb{P}[a_n(x+w) = x] = \sum_w \mathbb{P}[a_n(x) = x-w] = 1$. This is the most common use of the mass transport principle. Defining $D_n(x) = |E_n(x)| - 1$ gives $\mathbb{E}[D_n(x)] = 0$. So $\mathbb{P}[D_n(x) > 0] \leq \mathbb{E}[D_n(x)^+] = \mathbb{E}[D_n(x)^-] = \mathbb{P}[D_n(x) = -1]$. But the last event corresponds to $\mathcal{E}_n(x)$, therefore $\sum_n \mathbb{P}[D_n(x) > 0] \leq 1$. Since $[D_n(x) > 0]$ corresponds to $\mathcal{D}_n(x)$, the result follows by the Borel-Cantelli Lemma. \square

Note that $C_{n+1}(x) > C_n(x)$ can only happen on the event $\mathcal{D}_n(x)$. So, almost surely, $(C_n(x))$ is a nonnegative, eventually non-increasing sequence, and therefore $\lim_{n \rightarrow \infty} C_n(x)$ exists.

The following lemma permits us to extend the results of [3] to a general translation-invariant probability distribution for the initial resource quantities.

Lemma 2.2. *\mathbb{P} -a.s., ties cannot happen infinitely often at a fixed vertex.*

Proof. For a pair of vertices $y \simeq x$, let

$$\mathcal{T}_n(x, y) = [0 < C_n(x) = C_n(z) = \max\{C_n(w), w \simeq y\} \text{ for some } z \simeq y, z \neq x]$$

denote the event that there is a tie at y at time n , and x is one of the candidates for its resource. Denote by $\mathcal{F}_{y,n}$ the σ -field generated by

$$(C_m(z), m \leq n \text{ and } z \in \mathbb{Z}^d) \text{ and } (a_m(z), m < n \text{ and } z \in \mathbb{Z}^d \text{ or } m = n \text{ and } z \neq y),$$

i.e., $\mathcal{F}_{y,n}$ contains all the information up to step n , except for the possible tie breaking at vertex y at the n -th step. Write $M_{x \setminus y}^n = |E_n(x) \setminus \{y\}|$. Note that $M_{x \setminus y}^n$ and $\mathcal{T}_n(x, y)$ are $\mathcal{F}_{y,n}$ -measurable. Conditioning on $\mathcal{F}_{y,n}$, if $M_{x \setminus y}^n = 0$ and $\mathcal{T}_n(x, y)$ occurs, then with probability at least $\frac{1}{2}$, $a_n(y) \neq x$, in which case $\mathcal{E}_n(x)$ will happen and thus $C_m(x) = 0$ for all $m > n$ (therefore $\mathcal{T}_n(x, y)$ can never happen again). If $M_{x \setminus y}^n > 0$ and $\mathcal{T}_n(x, y)$ occurs, then with probability at least $\frac{1}{2d+1}$, $a_n(y) = x$, in which case $\mathcal{D}_n(x)$ will happen. Therefore, the occurrence of $\mathcal{T}_n(x, y)$ for infinitely many n 's implies almost surely the occurrence of $\mathcal{D}_n(x)$ infinitely often. But the latter event has probability 0 by Lemma 2.1. \square

Lemma 2.3 ([3], p.338). *\mathbb{P} -a.s. for each vertex x , only one of the events $\limsup_n \mathcal{A}_n(x)$, $\limsup_n \mathcal{B}_n(x)$ or $\limsup_n \mathcal{C}_n(x)$ will happen.*

Proof. Once we know that there are finitely many ties at most, the argument of [3] can be applied. We present the proof for the convenience of the reader.

By Lemmas 2.1 and 2.2 we can take n_0 so that neither $\mathcal{D}_n(x)$ nor $\mathcal{E}_n(x)$ will happen, nor will there be a tie at any $y \in \mathcal{G}_x$, for any $n \geq n_0$. If $\mathcal{A}_n(x)$ happens for some n , it also happens for all $m > n$. It is thus enough to prove that, if $\mathcal{C}_n(x)$ happens for some $n \geq n_0$, then $\mathcal{C}_{n+1}(x)$ also happens. Now, if $\mathcal{C}_n(x)$ happens, we have $C_{n+1}(z) \geq C_n(x) > C_n(y) = C_{n+1}(x)$, where $E_n(x) = \{y\}$ and $z = a_n(x)$. The strict inequality holds because there cannot be a tie at y . Therefore $a_{n+1}(x) \neq x$ and, since \mathcal{D} and \mathcal{E} have been ruled out, we must have $\mathcal{C}_{n+1}(x)$ again. \square

By Lemma 2.3 we can say that each vertex is uniquely either an \mathcal{A} -vertex or a \mathcal{B} -vertex or a \mathcal{C} -vertex.

Corollary 2.4. \mathbb{P} -a.s., $C'_n(x) - C_n(x) \rightarrow 0$ for all $x \in \mathbb{Z}^d$.

Proof. If x is an \mathcal{A} -vertex or a \mathcal{B} -vertex then $C'_n(x) - C_n(x) = 0$ eventually. So suppose x is a \mathcal{C} -vertex. Take n_0 so that $\mathcal{C}_n(x)$ happens for all $n \geq n_0$. By Lemma 2.1 we can further assume that $\mathcal{D}_n(z)$ will not happen for any $z \in \mathcal{G}_x$ and $n \geq n_0$, in particular $C_{n+1}(z) \leq C_n(z)$. We claim that $C'_{n+2d}(x) \leq C_n(x)$ for all $n \geq n_0$. Since in addition $C_{n+2d}(x) \leq C'_{n+2d}(x)$ and $C_n(x)$ converges, this will finish the proof.

Let $n \geq n_0$ be fixed and for $m \geq n$ let $V_m = \{y \sim x : C_m(y) > C_n(x)\}$. Since resource quantities no longer increase, $V_{m+1} \subseteq V_m$. Moreover, the occurrence of $\mathcal{C}_n(x)$ implies that $a_n(x) \neq x$ and, since $\mathcal{D}_n(a_n(x))$ cannot occur, $a_n(x)$ is a \mathcal{C} -vertex and $E_n(a_n(x)) = \{x\}$. Thus, $C_{n+1}(a_n(x)) = C_n(x)$ and therefore $a_n(x) \notin V_{n+1}$. If $V_{n+1} = \emptyset$ we are done, and if $V_{n+1} \neq \emptyset$ we find again that $a_{n+1}(x) \in V_{n+1} \setminus V_{n+2}$. Proceeding this way, and since $|V_n| < 2d$, we must find $V_{n+j} = \emptyset$ in less than $2d$ steps. \square

Lemma 2.5. For each fixed $k \in \mathbb{N}$ and $x \in \mathbb{Z}^d$,

$$\mathbb{P}(0 < C'_n(x) - C_n(x) < \delta, |S_n(x)| \leq k) \rightarrow 0$$

as $\delta \rightarrow 0$, uniformly in n .

Proof. For $v \in \mathbb{Z}^d$ consider $S_n(L_n(v))$, the set of vertices whose initial resource were joined with that of v by time n (according to the dynamics rules, once two or more initially distinct sources join, they never separate), and define $N_n(v) := |S_n(L_n(v))|$. For each fixed v , $N_n(v)$ is non-decreasing in n .

Let $A_n(v) = C_n(L_n(v))$ be the total amount of resource at $L_n(v)$ at time n and let $A'_n(v) = C'_n(L_n(v)) = C_n(L_{n+1}(v))$. For each vertex v , the value $A'_n(v) - A_n(v)$ is nonnegative and it can only decrease when $N_n(v)$ increases.

It is therefore the case that $\inf [A'_n(v) - A_n(v)]$, the infimum taken over all n such that $N_n(v) \leq k$ and $A'_n(v) > A_n(v)$, is a strictly positive random variable. In particular, for fixed k ,

$$\mathbb{P}(0 < A'_n(v) - A_n(v) < \delta, N_n(v) \leq k) \rightarrow 0 \text{ as } \delta \rightarrow 0 \quad \text{uniformly in } n. \quad (2.1)$$

We shall relate the above limit to the desired result via a quantitative use of the mass-transport principle, what we call an unlikelihood transfer argument.

Let us omit δ, k from now on and denote the event in (2.1) by $\mathcal{U}_n(v)$. For $x \in \mathbb{Z}^d$, let $\mathcal{V}_n(x)$ denote the event that $\mathcal{U}_n(v)$ occurs for some $v \in S_n(x)$. If $\mathcal{U}_n(v)$ happens for any such v then it happens for all of them (because the values of A_n, A'_n and N_n are constant within $S_n(x)$ and equal $C_n(x), C'_n(x)$, and $|S_n(x)|$, respectively). Writing $m(v, x) = \mathbb{1}_{L_n(v)=x} \mathcal{U}_n(v)$ we have that $\mathbb{1}_{\mathcal{U}_n(v)} = \sum_x m(v, x)$. It follows from the translation invariance of the process that $\mathbb{E} \sum_y m(w, w+y) = \mathbb{E} \sum_y m(w-y, w)$, which means

$$\mathbb{E} \sum_x m(w, x) = \mathbb{E} \sum_v m(v, w) \quad \forall w, \quad (2.2)$$

giving

$$\mathbb{P}[\mathcal{U}_n(v)] = \mathbb{E}[|S_n(x)| \mathbb{1}_{\mathcal{V}_n(x)}] \geq \mathbb{P}(\mathcal{V}_n(x)) \quad \forall v, x. \quad (2.3)$$

Now $\mathcal{V}_n(x)$ is exactly the event considered in the statement of the lemma and since $\mathbb{P}[\mathcal{U}_n(v)] \rightarrow 0$ as $\delta \rightarrow 0$ uniformly in n the result follows. \square

Proof of Theorem 1.1: By Lemma 2.3 it is enough to show that $\mathbb{P}(a_n(x) \neq x) \rightarrow 0$ as $n \rightarrow \infty$. For any choice of $\delta > 0$, if $a_n(x) \neq x$ then either $C'_n(x) = C_n(x)$, which implies that there is a tie at x , or $C'_n(x) - C_n(x) \geq \delta$ or $0 < C'_n(x) - C_n(x) < \delta$. By Lemma 2.2 and Corollary 2.4 the probabilities of the first two events tend to zero as $n \rightarrow \infty$.

For any choice of $k \in \mathbb{N}$, the last event can be split into two cases, according to whether $|S_n(x)| \leq k$ or not. By Lemma 2.5, the probability that this event happens in the first case tends to zero as $\delta \rightarrow 0$, uniformly in n . Now, with $m(v, x) = \mathbb{1}_{L_n(v)=x}$, (2.2) gives $\mathbb{E}(|S_n(x)|) = 1$ so $\mathbb{P}(|S_n(x)| > k) < \frac{1}{k}$.

So, if we consider $\limsup_{n \rightarrow \infty} \mathbb{P}(a_n(x) \neq x)$, let $\delta \rightarrow 0$ and then let $k \rightarrow \infty$ we get the desired limit. \square

3 Concluding remarks

We conclude this paper by discussing how Theorem 1.1 extends to more general settings and why a positive answer to Question 3 does not follow from the previous arguments.

Generalizations Our proof applies in other settings with much generality, as long as the mass transport principle (2.2) is true. It thus covers cases when the distribution of initial resources is invariant with respect to a transitive unimodular group of automorphisms. Examples include Cayley graphs, regular trees, etc. The only change in the proof is to replace $2d$ by the graph degree.

It also covers graphs that are locally finite and can be periodically embedded in \mathbb{R}^d . Namely, one can consider graphs whose vertex set may be written as $[J] \times \mathbb{Z}^d$, where $[J] := \{1, \dots, J\}$, and whose edge set is invariant under the mappings $(j, x) \mapsto (j, x + y)$ for all $y \in \mathbb{Z}^d$. Here translation invariance is understood as the distribution of the initial resources being invariant under the above mappings.

The changes in the proof for the above case are the following. In Lemma 2.1, notice that $\mathcal{E}_n(j, x)$ will happen at most once for each $j \in [J]$ and fixed x ; so $D_n(x) := (\sum_1^J |E_n(j, x)|) - J$ satisfies $\mathbb{E}[\sum_n D_n(x)^+] = \mathbb{E}[\sum_n D_n(x)^-] \leq J$, thus $D_n(x) > 0$ can occur for at most finitely many n 's; finally $\mathcal{D}_n(j, x)$ corresponds to $|E_n(j, x)| - 1 > 0$, which in turn implies that $D_n(x) > 0$ or that $\mathcal{E}_n(j', x)$ occurs for some other $j' \in [J]$. In Lemma 2.2 and Corollary 2.4 replace $2d$ by $\max_j \deg(j, x)$. In Lemma 2.5 we take $m(v, x) = \sum_{j, j'} \mathbf{1}_{L_n(j, v) = (j', x), \mathcal{U}_n(j, v)}$ and the mass-transport principle (2.2) holds with the same proof. In the proof of Theorem 1.1 we can write $\mathbb{E}(\sum_j |S_n(j, x)|) = J$, giving $\mathbb{P}(|S_n(j, x)| > k) < \frac{J}{k}$ for any (j, x) .

Open question It would be natural to try to answer Question 3 using an argument similar to the proof of Theorem 1.1, in a way that uses the result to reinforce itself.

First notice that a yes to Question 3 is equivalent to showing that every amount of initial resource eventually stops moving. Indeed, if $C_\infty(x) := \lim_n C_n(x)$ and F_v denotes the event that $L_n(v)$ is eventually constant, then $\mathbb{E}[C_\infty(x)] = \mathbb{E}[C_0(v)\mathbf{1}_{F_v}]$, so a yes to the question is equivalent to F_v happening almost surely. This equality follows from Theorem 1.1 and (2.2) with

$$m(v, x) = \begin{cases} C_0(v), & \text{if } L_n(v) = x \text{ eventually,} \\ 0, & \text{otherwise.} \end{cases}$$

Let us denote by $\tilde{\mathcal{U}}_n(v)$ the event that, for some $m > n$, $L_{m+1}(v) \neq L_m(v)$. This event means that the resource that started at v is still going to move after time n . Therefore Question 3 boils down to whether we can show that $\mathbb{P}[\tilde{\mathcal{U}}_n(v)] \rightarrow 0$ as $n \rightarrow \infty$.

Let $\tilde{\mathcal{V}}_n(x)$ denote the event that the resource found at vertex x at time n will leave x after time n , that is, $a_m(x) \neq x$ for some $m > n$. Note that this is equivalent to the event that $\tilde{\mathcal{U}}_n(v)$ occurs for some $v \in S_n(x)$, and again if it happens for any such v then it happens for all of them. This implies that the analogue of (2.3) holds.

By Theorem 1.1, $\mathbb{P}[\tilde{\mathcal{V}}_n(x)] \rightarrow 0$. The inequality in the analogue of (2.3) does not help here, but we could make use of the equality if the distribution of $|S_n(x)|$ satisfied an appropriate moment (or uniform integrability) condition. This is another quantitative use of the mass-transport principle, that we call unlikelihood sharing principle.

For instance, suppose one can prove that $E(|S_n(x)|^\alpha)$ stays bounded as $n \rightarrow \infty$ for

some $\alpha > 1$. Then $\mathbb{P}[\tilde{\mathcal{V}}_n(x)] \rightarrow 0$ implies that $\mathbb{P}[\tilde{\mathcal{U}}_n(v)] \rightarrow 0$ since

$$\mathbb{E}(S\mathbf{1}_{\tilde{\mathcal{V}}}) = \mathbb{P}(\tilde{\mathcal{V}})\mathbb{E}(S|\tilde{\mathcal{V}}) \leq \mathbb{P}(\tilde{\mathcal{V}})[\mathbb{E}(S^\alpha|\tilde{\mathcal{V}})]^{1/\alpha} \leq [\mathbb{P}(\tilde{\mathcal{V}})]^{1-\frac{1}{\alpha}}\mathbb{E}[S^\alpha]^{1/\alpha},$$

where $S = |S_n(x)|$ and $\tilde{\mathcal{V}} = \tilde{\mathcal{V}}_n(x)$.

Unfortunately, we do not know how to control the tail of the distribution of $|S_n(x)|$, and Question 3 remains an open problem.

While the final version of this paper was in preparation, it was found by van den Berg, Hilário and Holroyd [2] that Question 3 as stated has a negative answer: there do exist nonnegative translation-invariant initial distributions for which resources do escape to infinity.

Acknowledgments This work had financial support from ARCUS Program, CNPq grants 140532/2007-2 and 141114/2004-5, FAPERJ, FAPESP grant 07/58470-1, FSM-Paris, German Israeli Foundation grant I-870-58.6/2005 (O.L.) US NSF under grants DMS-08-25081 (O.L.), DMS-06-06696 (C.M.N.), DMS-06-45585 (S.S.), and OISE-07-30136 (O.L., C.M.N., and S.S.).

Part of this collaboration took place during the Fall, 2008 semester on *Interacting Particle Systems* at IHP, Paris; we thank the organizers and staff for providing a nice working environment. M. R. Hilário and L. T. Rolla thank CWI-Amsterdam and C. M. Newman and L. T. Rolla thank CRM-Montréal for hospitality. V. Sidoravicius thanks Chris Hoffman for many fruitful discussions.

References

- [1] J. VAN DEN BERG AND A. ERMAKOV, *On a distributed clustering process of Coffman, Courtois, Gilbert and Piret*, J. Appl. Probab., 35 (1998), pp. 919–924.
- [2] J. VAN DEN BERG, M. R. HILÁRIO, AND A. HOLROYD, In preparation.
- [3] J. VAN DEN BERG AND R. W. J. MEESTER, *Stability properties of a flow process in graphs*, Random Structures Algorithms, 2 (1991), pp. 335–341.
- [4] E. COFFMAN, P.-J. COURTOIS, E. GILBERT, AND P. PIRET, *A distributed clustering process*, J. Appl. Probab., 28 (1991), pp. 737–750.