

Towards a Theory of “Local to Global” in Distributed Multi-Agent Systems (I)

Daniel Yamins
Harvard University
33 Oxford St.
Cambridge, MA 02138 USA
yamins@fas.harvard.edu

ABSTRACT

There is a growing need for a theory of “local to global” in distributed multi-agent systems, one which is able systematically to describe and analyze a variety of problems. This is the first in a series of two papers that begins to develop such a theory. Here, we analyze one particular multi-agent problem – the “equigrouping problem,” in which multiple identical agents organize themselves into groups of equal size. We develop a formal model for describing the system and an notion of equivalence characterizing multi-agent algorithms in terms of the group behaviors induced by the algorithm. Our main result is a characterization of the space of all solutions to the equigrouping problem with respect to this group behavior equivalence. The result allows us to obtain infinitely many substantially different solutions to the Equigrouping problem, and to understand these different solutions in a qualitatively satisfying manner. The second paper in this series indicates how to develop and generalize the modeling method obtained here to other problems.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent Systems

General Terms

Algorithms, Theory

Keywords

Distributed Algorithms, Theory of Algorithms, Local-to-Global, Emergent Order, “Local-to-Global”

Introduction

Distributed multi-agent systems are abundant in the biological world, exhibit rich and interesting behavior, and have been inspiring to researchers in many other areas. Examples of such systems – both natural and man-made – that often come to mind include: Flocks, herds, and schools ([20], [24],[8],[18]); Bacterial colonies and chemotaxis ([2], [22]); Embryological and morphogenetic systems ([28], [25], [16]); Ant colonies and bee/wasp swarms ([3], [7]); Ant- and swarm-inspired algorithms ([5], [12]);

Flight formation in UAVs ([31], [17]); Neurons and neural networks ([14], [19]); The immune system and immune-inspired algorithms; ([6], [10]); Robot soccer and other multi-agent team sports ([15], [23]); Amorphous Computing, iRobots, SmartDust, and other many-agent robotic systems ([1], [13], [9], [21]); and Cellular Automata ([4], [27], [26]).

The dominant mode of studying these systems has been through simulation ([20], [16], [3], [5], [1], [27]). Such simulations have uncovered important mechanisms underlying distributed systems, including positive feedback, stigmergy, signal gradients, positional information, probabilistic symmetry breaking, division of labor, stability etc. Mathematical models have also been created to quantitatively analyze candidate mechanisms in many cases ([8], [18], [2], [22], [25], [3], [14], [6], [10], [21]).

More abstract questions, regarding solution optimization, group-level programming, task and algorithmic complexity, and the limits of (de)centralization, have also been posed. In the current state of the field, there is a great diversity of (mostly informal) descriptive methods, mirroring perhaps the diversity of the systems they describe. However, there is a growing need to study these abstract problems in a systematic way, as well as to provide a qualitative mathematical framework in which to compare various possible underlying mechanisms. It would therefore be useful to have an coherent theory of “local to global” in distributed multi-agent systems, one which is able to describe and to analyze a variety of problems.

In this series of two papers, we attempt to provide the beginnings of a theory which addresses this question. In the first paper, we focus on a simple illustrative problem – that of “equigrouping” mobile agents into groups of equal size from arbitrary initial positions. We introduce a formal model of the system and carefully analyze the set of *all* one-dimensional solutions to equigrouping with respect to a *group behavior equivalence*. We thereby obtain a spectrum of infinitely many solutions, together with necessary and sufficient conditions characterizing the possible combinations of group behaviors of any solutions. In the second paper we indicate how to generalize the modeling approach used here to complex and realistic multi-agent systems. We believe that analyzing a given problem’s *solution space* – distinct from the typical approach of describing and verifying a single solution – is a novel theoretical approach to multi-agent systems.

1. THE EQUIGROUPING PROBLEM

In this paper, we describe and analyze the simple but instructive “equigrouping” problem introduced in [30]. Consider a one-dimensional lattice. Two point-agents placed on this lattice are said to be *in the same group* if all lattice points between the two agents’ positions are occupied by other agents. Conversely, two agents are *separated* if there is at least one unoccupied lattice point between them. For each positive integer p , the one-dimensional p -Equigrouping problem consists of finding local algorithms which take any arbitrary initial configuration

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS’05, July 25-29, 2005, Utrecht, Netherlands
Copyright 2005 ACM 1-59593-094-9/05/0007 ...\$5.00.




Figure 1: a) A non-equigrouned state in \mathcal{C} . b) A 3-equigrouned state, in $\mathcal{C}_3 \subset \mathcal{C}$.

of $m \times p$ agents (for m a positive integer) into a configuration of m separated groups containing p agents each – that is, a “ p -equigrouned” configuration with m separate groups.

The equigrouning pattern – that of p -sized groups laid out separated from each other – is one of the intuitively simplest non-local patterns available in one dimension. But despite of its simplicity, the problem was shown in [30] to cleanly illustrate several key difficulties in the local construction of global patterns.

More formally, let \mathcal{L} be the one-dimensional integral lattice. Denote by X an initial configuration of point agents on \mathcal{L} , so that $Ag(X) = \{a_1, \dots, a_n\}$ is a listing of the agent-positions along \mathcal{L} . Identifying \mathcal{L} with the integers \mathbb{Z} , we denote by $pos(a, X)$ the integral value of the lattice point at which agent a is located in configuration X , under this identification. For example, $pos(a, X) > pos(b, X)$ means that a is to the right of b in X . (We will drop the argument X from the notation when the context is clear.) Let the set of all configurations of finitely many agents on \mathcal{L} be denoted by \mathcal{C} . For $X \in \mathcal{C}$ denote by $le(X)$ and $re(X)$ the left-most and right-most agents in X , respectively. Given $a \in ag(X)$, let a^\pm denote the agents adjacent to the right and left, respectively. (Of course, a^- (resp. a^+) will not exist if $a = le(X)$ (resp. $a = re(X)$)). Let \mathcal{C}_p denote the set of all p -equigrouned configurations (Figure 1).

We will use the notations \oplus and \ominus to refer to addition and removal of agents, respectively. A decomposition of agents $X = X_1 \oplus X_2$ is *consecutive* if $pos(re(X_1)) < pos(le(X_2))$. If X and Y are two configurations of agents, then $X \boxplus Y$ represents the set of all Z for which $X \oplus Y$ is a consecutive decomposition.¹ Henceforth, we will use \circ to denote an empty lattice position and \bullet to denote an occupied position, so that in the (sub)configuration in figure 1a) is denoted $\circ\bullet\bullet\bullet\bullet\circ\bullet\bullet$. We will use exponential notation (for which $g = \circ\bullet^4 \circ^2\bullet^2$) to efficiently denote configuration segments.

Dynamics are generated from (mostly) identical local operators associated with each identical agent. To describe this mathematically, for a given agent $a \in X$, let $b_r(a, X) \subset X$ be the ball of radius r around a in X – meaning, the r lattice points to the left and r lattice points to the right of a , together with whatever agents are at those points. In all, $b_r(a, X)$ contains $2r + 1$ lattice points (including the point where a itself is) and at most $2r$ other agents. Let $f(a_i, X)$ be any operator given by

$$f : b_r(a_i, X) \rightarrow s,$$

in which s is a lattice segment identical to $b_r(a, X)$ except that agent a can have moved either to the left or right by one lattice unit, or have stayed in place. We do not allow two agents to occupy identical positions, so that, for example, if an agent is directly to the right of another agent, the first agent cannot move left. We require f to be identical for all agents a_i , except the right and left most agents $re(X)$ and $le(X)$, respectively. In fact, we allow $f(\{re(X), le(X)\}, X)$ to be different from $f(a, X)$ where a is not an end-agent, corresponding to the idea of giving agents line-of-sight information about whether or not they

¹Hence, as defined here, \boxplus is *not* a commutative operation. We will use the notation $g^{\boxplus k}$ to denote the k -times \boxplus of g with identical copies of g .

have neighbors to their left and right (at whatever distance).² Denote the (possibly different) left and right maps by f_l, f_r . In the case that $f(a, X)$ “moves” the agent a to the left, we write $[f(a, X)] = L$; and use analogous notation for R, S to denote right and stationary movement.

We require f to have a finite well-defined information radius – the size of the largest ball $b_r(a, X)$ from which f can draw information. This is denoted $r(f)$. We will use the notation \mathcal{A} to denote the set of all local algorithm specifications of this kind; hence, any local algorithm f is an element of \mathcal{A} . The specification thus far formally defines f on a local ball around a given agent; we can “globalize” this action to all of X in an obvious way by taking X to a configuration in which $b_r(a, X)$ has been replaced X with s ; that is

$$f(a, X) = (X \ominus b_r(a, X)) \oplus s.$$

For each $s = (a_1, \dots, a_n, \dots) \in \mathcal{SEM}$, the sequence of compositions

$$\circ_i f(a, \cdot) = f(a_n, (\dots (f(a_1, X)) \dots))$$

will be denoted by f_n^s , and applies to initial condition x_0 to generate trajectories $\{f_n^s(x_0)\}$. We can allow f to be probabilistically specified by attaching to each possible configuration of the agent’s $b_r(a)$ probabilities p_l, p_r of moving to the left and right, and probability $1 - p_l - p_r$ to remaining still.³

Let \mathcal{SEM} be the set of all infinite sequences of agent-labels such that each agent a_i appears infinitely many times. These *allowable semantic strings* correspond exactly to the UNITY semantics described in [11].⁴ Let \tilde{S} be the set of all eventually periodic sequences. We say that f is a solution to the p -equigrouning problem if for all such x_0 with $m \times p$ agents for any m and each $s \in \mathcal{SEM}$, the trajectory $f_n^s(x_0)$ converges to a fixed configuration in \mathcal{C}_p in finite time with probability 1. If we let P_n be the probability that $f_n^s(x_0) \in \mathcal{C}_p$ and $f_m^{s'}(f_n^s(x_0)) = f_n^s(x_0)$ for all m and $s' \in \mathcal{SEM}$, then $\lim_{n \rightarrow \infty} P_n = 1$. Let \mathcal{F}_p denote the set of deterministic solutions to p -equigrouning.

Our goal in the rest of this paper is to learn about the structure of \mathcal{F}_p . The first and most obvious question to ask is: is \mathcal{F}_p non-empty? Can any solutions be found? In [30] and [29] we identified two qualitatively different deterministic solutions:

Algorithm 1 For each p , define the algorithm $F_1(p)$ with information radius $r(F_1(p)) = p$ locally to any given agent a by the rules:

1. Suppose $a \neq le(X)$, but $a = le(g)$ for some consecutive group $g \subset X$. Then if $|g| \neq p$, move left, i.e. $[F_1(p)(a, X)] = L$.
2. Conversely, suppose $a = le(X)$. If a is in a group of size greater than p , the action is L .
3. In all other cases, stay, i.e. $[F_1(p)(a, X)] = S$.

Algorithm 2 Define a local rule $F_2(p)$ with information radius $r(F_2(p)) = p + 7$ which, on the state X , is given locally to any given agent a by:

1. Suppose $a = le(X), re(X)$. Then $F_2(p)(a, X) = F_1(p)(a, X)$.

²It turns out that it is impossible to solve the equigrouning problem without making this allowance; this is shown in [30].

³If a given motion is unavailable since the adjacent position in that direction is already occupied, the probability associated with that motion is automatically 0. Deterministic algorithms are simply the special case in which one of $p_l, p_r, 1 - p_r - p_l$ is 1 and the others zero.

⁴Following standard notation, we will use $s_2 \circ s_1$ to denote composition of semantics with s_1 first, followed by s_2 . The notation $s^{\circ k}$ indicates the k -times composition of s with itself.




Figure 2: Illustration of the operation of $F_1(3)$ in a representative case. Timesteps proceed from top row to bottom row.




Figure 3: Illustration of the operation of $F_2(3)$ in a representative case.

2. Suppose $a \neq le(X)$ and a is part of a contiguous group of size at least p . Then $F_2(p)(a, X) = F_1(p)(a, X)$.
3. Suppose $a, b \neq le(X), re(X)$ are consecutive agents, with $dist(a, b) \geq 5$. Then if possible, both a and b move left, $[F_2(p)(\{a, b, \}, X)] = L$.
4. Suppose $a, b \neq le(x), re(x)$ are consecutive agents with $dist(a, b) \leq 2$. Then $[F_2(p)\{a, b\}] = L$, while ensuring that, unless part of a p -group, a and b remain within distance 2 of each other. If one is part of a p group, then a and b are allowed to separate.
5. Suppose a, b are as above with $3 \leq dist(a, b) \leq 4$. Then unless to the left of a group of size p , $[F_2(p)(\{a, b, \}, X)] = R$.
6. Suppose a, b are as above with $3 \leq dist(a, b) \leq 4$, to locally to the left of a group of size p . Then $F_2(p)(\{a, b\}, X)$ compresses $\{a, b\}$ toward each other.
7. In all other cases, $[F_2(p)(a, X)] = L$

To visualize the operation of the first algorithm $F_1(p)$, let $p = 3, m = 2$ and consider the initial configuration $X = \bullet^4 \circ^2 \bullet^2$. Under the semantic generated by reading off agents repeatedly from left to right, $F_1(3)$ converges to a solution in 7 steps (figure 2). Via the obvious “mirror algorithm” construction it is easy to define a “rightward version” of which mirrors $F_1(p)$ is the mirror.

To visualize $F_2(p)$, let’s consider the case in which $p = 3, m = 3$ and the initial configuration X is $\bullet \circ \bullet^3 \circ^5 \bullet \circ^3 \bullet \circ^6 \bullet^3$. Under the same semantic as for algorithm $F_1(3)$, X evolves under $F_2(3)$ as illustrated in figure 3. In this particular situation, all the agents are initially stationary with the exception of the 5th and 6th agents. This group of two initially moves to the right as per to rule 5; eventually interacts with the right-end 3-group and compresses into a left-moving 2-group as per rule 6; moves left as per rule 3; and then interacts with the left-end agents, entering an equigrouped state as per rules 1 and 2. The information radius is $p+7$ exactly so that the left-most agent in a 2-group can evaluate rules 5 and 6.

F_1 and F_2 differ from each other in that while under F_1 agents are either stationary or left-moving, F_2 can sometimes have agents moving to the right. In solving the initial condition X , F_1 would much more efficiently move the 5th and 6th agents to the left, instead of first having them move to the right before interacting with the right-end agents as F_2 did. Intuitively, F_2 thereby

differs in an important way from F_1 . The exact extent of this difference will be explored in the next few sections. In fact, we will characterize this difference with respect to an equivalence relation on the set of solutions; F_1 and F_2 will turn out to be in different equivalence classes.

Notice that in both algorithms the information radius $r(F(p)) \geq p$. Also, notice that the behavior of (at least one of) the end-most agents is different from the rest – that is, the left-most agent must *know* that it is an end-agent. It turns out that these are not coincidences:

Proposition 1 *Let $f \in \mathcal{F}_p$. Then $r(f) \geq p$.*

Proposition 2 *There are no $f \in \mathcal{F}_p$ such that $f_l = f = f_r$.*

These establish that the right or left end agent (or both) must know that it is an end-agent. In other words, some kind of extra information – requiring the ability to communicate (with extremely minimal low information content) across the infinitely long line – is required. Propositions 1 and 2, originally proved in [30], are simple examples of more general statements that can be made about deterministic solutions to equigrouping.

2. GROUP BEHAVIOR AND EQUIVALENCE

Intuitively, just as individual agents can (and must) have one of three “behaviors” under any given algorithm f – namely L, R or S – we will show that groups of agents can sometimes be said to as well. In this section, we will define a reasonable concept of “group behavior”, and quote results that show this concept is well-behaved and applicable. We will then define an equivalence relation between two deterministic algorithms f and g , in which $f \sim g$ whenever f and g induce the same group behaviors. In the next section, we will characterize the quotient of \mathcal{F}_p up to this equivalence relation.

Definition 1 [Generalized Group Behaviors] Let g be an isolated set of agents and f a deterministic algorithm in \mathcal{A} . Then, for a given $s \in \mathcal{SEM}$,

1. Define (f, s) as *moving g to the left* if for any positive integer n , there exists m_n , another positive integer, such that the left-end agent after action of (f, s) for m_n steps (that is, $le(f_{m_n}^s(g))$) has been translated to the left of its original configuration (that is, $le(g)$) by more than n places, i.e.

$$pos(le(f_{m_n}^s(g))) - pos(le(g)) < -n;$$

and similarly for the right-end agents. If this situation obtains then we write $[g, f, s] = L$.

2. Define (f, s) as *moving g to the right* analogously, with

$$pos(le(f_{m_n}^s(g))) - pos(le(g)) > n;$$

and similarly for the right-end agents. If this situation obtains then we write $[g, f, s] = R$.

3. Define (f, s) as *staying g* if there exists n such that for all m there is an $m' > m$ for which the end-points of $f_s^{m'}$ have moved no further than distance n from initial points. If this is the case, we write $[g, f, s] = S$.

Analogous definitions can made for left- and right-end agent groups, using f_l or f_r in place of f ; notation for behaviors are $[g, f, s]_l$ and $[g, f, s]_r$ respectively. There is a natural ordering of behaviors with respect to the direction along the line $\mathcal{L}, L < S < R$.

Though L, R, S are mutually exclusive and exhaustive behaviors for a single isolated agent, it is not *a priori* clear whether this remains true for groups of more than one agent. It turns out that for the right kind of group, it is.

Definition 2 Let g be a set of agents on the line \mathcal{L} and $f \in \mathcal{A}$ be a deterministic algorithm. Then g is an (f, s) -prekernel for a given semantic string s if there is an integer $l \in \mathbb{N}$ such that when isolated, given any $m \in \mathbb{N}$, there is an integer $n \geq m$ such that for all a , $\text{dist}(a, a^+) < l$ in $f_n^s(g)$.

In [30], we showed:

Proposition 3 [Pre-kernels Carry Behavior] Suppose that g is an (f, s) -prekernel for $s \in \hat{S}$. Then the ω -limit set $\Omega = \lim_{n \rightarrow \infty} \{f_n^s(g)\}$ is a periodic cycle with well-defined displacement $m[g, f, s] \in \mathbb{Z}$ and speed $n[g, f, s] \in \mathbb{Q}$ defined by

$$m[g, f, s] = \text{pos}(a, t + \tau) - \text{pos}(a, t)$$

and

$$n[g, f, s] = \frac{\text{pos}(a, t + \tau) - \text{pos}(a, t)}{\tau/|g|}$$

in which $\tau = |\hat{s}|$ is the length of the minimal period of Ω and $a \in g$ is any agent.

A positive speed indicates motion to the right, a negative speed indicates motion to the left, and zero speed indications that the pre-kernel is stationary. Analogous results are true for $[g, f, s]_l$ and $[g, f, s]_r$.

A prekernel is a unit that “remains” together, and thereby acts as a carrier of group behavior. A natural question to ask to aid further analysis is what the *minimal* such units are. To this end we seek a formulation for decomposition of pre-kernels into minimal sub-pre-kernels such that the behavior of the whole is a “trivial” combination of the behavior of the subunits. A naive attempt at such a definition is: for a pre-kernel g , to say that it “decomposes” into $g = g_1 \oplus g_2$ if and only if $f_n^s(g) = f_n^{s|g_1}(g_1) \oplus f_n^{s|g_2}(g_2)$ for all n . However, for various reasons this definition fails to be consistent. Because of this, we formulate the following (unfortunately less concise) technical definition:

Let g be an (f, s) -prekernel and let $g = g_1 \oplus g_2$ be a consecutive decomposition of g .⁵ For each n , let $a_n = \text{re}(f_n^s(g)|_{g_1})$ and $b_n = \text{le}(f_n^s(g)|_{g_2})$. Let s' be the semantic created from s by removing calls to a_n whenever $\text{dist}(a_n, b_n) = 1$ and $[f(a_n, X \ominus g_2)] = R$ and to b_n whenever $[f(b_n, X \ominus g_1)] = L$. Define $g'_1 = f^{s'|g_1}(g_1)$, $g'_2 = f^{s'|g_2}(g_2)$. Obviously $\{f_n^{s'}(g)\} = \{f_n^s(g)\}$ as trajectories, the latter being a delayed version of the former.

Definition 3 [Minimal Decomposition] The consecutive decomposition $g = g_1 \oplus g_2$ of the (f, s) -prekernel g for $s \in \hat{S}$, is a *kernel decomposition* if

1. $s' \in \mathcal{SEM}$, and g_1, g_2 are pre-kernels for $(f, s'|_{g_1})$ and $(f, s'|_{g_2})$ respectively.
2. $\Omega(g, f, s') = \Omega(g_1, f, s'|_{g_1}) \oplus \Omega(g_2, f, s'|_{g_2})$.
3. At any timestep n , $f_n^{s'|g_1}(g_1)$ and $f_n^{s'|g_2}(g_2)$ are accessible to and from the trajectories $f^{s|g_1}(g_1)$ and $f^{s|g_2}(g_2)$ respectively.⁶

We say g is an (f, s) -kernel if no consecutive decomposition $g = g_1 \oplus g_2$ is a kernel decomposition.

In other words, an interaction kernel is a group which, when isolated, “stays together” and which cannot be written as the direct sum of two subgroups which act equivalently under the consecutive decomposition. Similar definitions are made for left-kernels (with agents at the left-end of the line, using f_l in place

⁵That is, $g \in g_1 \boxplus g_2$.

⁶Following standard terminology, we say that a configuration y is accessible from x if there is a finite string of calls s for which $y = f^s(x)$.

of f) and right-kernels, (with agents at the right-end of the line, using f_r in place of f). Let $\text{Ker}(f, s)$ denote the set of configurations which are (f, s) kernels; and $\text{Ker}^l(f, s)$ and $\text{Ker}^r(f, s)$ be analogous notations for left- and right-kernels.

We now need to ensure that kernels exist. To understand the existence result, we need to consider the *interaction* of pre-kernels. Suppose g and g' are two (f, s) -pre-kernels. We can form various consecutive sums in $g \boxplus g'$. If g is placed sufficiently far to the left of g' and

$$n[g, f, s] \leq n(g', f, s)$$

then g will never come into contact with g' .⁷ On the other hand, if

$$n[g, f, s] > n(g', f, s)$$

then the two groups will come into contact – that is, there will be a time such that $\text{le}(g')$ and $\text{re}(g)$ will be within distance $r(f)$ of each other, and after which

$$f_t^s(g \boxplus g') \not\subseteq f_t^s(g) \boxplus f_t^s(g').$$

The resulting conglomerate might become one larger pre-kernel, break down into several others, or perhaps somehow not generate any pre-kernels at all. Similar, but more complicated scenarios are imaginable with more than two starting pre-kernels. Hence, whenever pre-kernels exist, we can track the “interaction pattern” of pre-kernel formation and break-up. Formally, a consecutive sum $g_1 \oplus \dots \oplus g_n$ of pre-kernels (g_i, f, s) is said to be *good* if for all i , $n[g_i, f, s] \leq n[g_{i+1}, f, s]$; a good decomposition is one in which the ordering of the behaviors is consistent with the ordering of the lattice. A *temporary kernel decomposition* is one in which all g_i survive as independent kernels for at least $l = \max_i \{p(g_i, f, s)\}$ timesteps, where $p(g_i, f, s)$ is the prekernel-period of $\Omega(g_i)$.

We showed in [30] that an interaction pattern always exists:

Proposition 4 [Kernels Exist] For any configuration g , there is a time t and temporary kernel decomposition $g_t^1 \oplus \dots \oplus g_t^m$ for which

$$f_t^s(g) = g_t^1 \oplus \dots \oplus g_t^m$$

such that for $1 \leq i \leq j \leq m$, $g_t^i \oplus g_{t+1}^i \oplus \dots \oplus g_t^j$ is a (permanent) kernel decomposition whenever it is good. If $s \in \hat{S}$, the maximal (finest) such decompositions are periodic, and the maximal such decompositions which are good are (permanent) kernel decompositions.

The existence of one such t implies the existence of an infinite sequence t_i of future such times (periodically if $s \in \hat{S}$). Survival times $l_i = \max_j \{p(g_j^{t_i}, f, s)\}$ can be defined analogously. In the l_i timesteps after each t_i , interaction occurs only within the temporary kernels; between $t_i + l_i$ and t_{i+1} , interaction occurs *between* the temporary kernels, eventually setting up a new decomposition of temporary kernels. The entire evolution of states under f can be thought of a series of such kernel interaction patterns.

The existence and well-behaved properties of kernels require the assumption of eventually periodic semantics. This, however is a justified assumption because of the following lemma, proved in the appendix:

Lemma 1 If an algorithm f solves p -equivgrouping on all periodic UNITY semantics, then it is a solution for all UNITY semantics; that is, $f \in \mathcal{F}_p$

Hence, associated with any algorithm f , then is the record of its kernel sizes and their well-defined behaviors, information we will use to define the group-behavior equivalence relation on the space of local algorithms \mathcal{A} .

⁷Of course, if they are placed close to each other, they might interact.

Definition 4 [Group Behavior Equivalence] Let

$$K(f) = \{(|X|, [X, f, s]) \mid X \in \mathcal{C}, \exists s \in S \mid X \in \text{Ker}(f, s)\}.$$

If we define

$$\mathcal{W}_k = \{(j, B) \mid 1 \leq j \leq k, B \in \{L, S, R\}\},$$

then obviously $K(f) \subset \mathcal{W}_\infty$ for any f . $K(f)$ is the *kernel structure* of f . Define left- and right-kernel structures, denoted $K^l(f) = K(f_l)$ and $K^r(f) = K(f_r)$, analogously.

Two algorithms $f, g \in \mathcal{A}$ are kernel-equivalent when

$$(K^l(f), K(f), K^r(f)) = (K^l(g), K(g), K^r(g)).$$

We write $f \cong_K g$, and $[f]$ for the class of f in the quotient $\tilde{\mathcal{A}} = \mathcal{A} / \cong_K$.

Example 1 To fix ideas, let's look at the kernel-structure of algorithm 1. Let a represent an arbitrary agent, and let $g(k)$ be an isolated group of k consecutive agents. It is easily shown that:

$$K(F_1(p)) = \{(1, L), (p, S)\} = K^r(F_1(p)),$$

and

$$K^l(F_1(p)) = \{(k, S) \mid k \leq p\}.$$

This calculation intuitively corresponds to the fact that an individual agent travels to the left; 2-groups, \dots , up to $p-1$ -groups decompose into consecutive \oplus sums of independent 1-groups; p -groups are stationary; and groups larger than p “break off” to the left. The left-most agent behaves differently, as a kind of “stopper” which prevents agents from translating to the left indefinitely. Sub-portions of a group join up to make a whole p -group by collecting at the left end together.

In a sense, for each p , $[F_1(p)]$ is a “trivial element” in $\tilde{\mathcal{A}}$: all the kernels go in the same direction. However, the equivalence relation \cong_K is able to distinguish the $F_1(p)$ from $F_2(p)$ as elements of \mathcal{F}_p / \cong_K . This captures the inherent difference between the two algorithms noticed above. $K(F_1(p)) \neq K(F_2)$, because whereas $K(F_1(p))$ is as calculated above with no right-moving kernels, we obviously have $(2, R) \in K(F_2(p))$, the kernel structure carried by the configurations $\{(\bullet \circ \circ \circ \bullet), (\bullet \circ \circ \circ \bullet \circ \bullet)\}$.

Hence the quotient space $\tilde{\mathcal{F}}_p = \tilde{\mathcal{F}}_p / \cong_K$ contains at least two different elements of $\tilde{\mathcal{A}}$.

2.1 Classifying \mathcal{F}_p up to Behavioral Equivalence

The ideas of the previous section indicate several fundamental questions: How many inequivalent solutions are there to p -equigrouping? What sets arise as $K^l(f), K(f)$ and $K^r(f)$ for some solution $f \in \mathcal{F}_p$? Can tractable necessary and/or sufficient conditions on $K(f)$ or $K(f)$ be formulated for f to be in \mathcal{F}_p ? Knowing the answers to these questions would go a long way to a systematic understanding of the deterministic solutions to equigrouping. Indeed, it was the formulation of the first question that lead to the discovery of (the admittedly unintuitive) Algorithm 2.⁸

To proceed, we first answer a more basic question: What sets arise as $K(f)$ for a well-defined algorithm $f \in \mathcal{A}$? Any triple $(K^l, K, K^r) \in 2^{\mathcal{W}_\infty} \times 2^{\mathcal{W}_\infty} \times 2^{\mathcal{W}_\infty}$ is a potential kernel structure. But which of these actually occur?

Theorem 1 [Well-Defined Algorithms] K is the kernel structure of a well-defined algorithm only if

1. $\exists!(1, B) \in K$. Similarly for K^l and K^r .

⁸Our hope was that F_1 and its mirror would be, up to kernel equivalence, the unique solutions to equigrouping. Failure to find a proof of this fact lead to the formulation of F_2 as a counterexample.

2. If $(2, L)$ or $(2, R) \in K$ then $\exists(j, B) \in K$ for $j > 2$ or $B = S$.

Furthermore, there are finite conditions on the relationship between K^l, K^r and K such that all finite potential kernel structures satisfying these conditions arise.

The full statement and proof of Theorem 1 is given in the appendix; one relevant fact from the proof to proceed here is that the construction of the $f[K^l, K, K^r]$ can be made so that all kernels with size divisible by p are p -equigrouped states.

We now use theorem 1, which applies to all algorithms, to understand the solutions to the equigrouping problem in specific. To this end, consider the following conditions one could place on the kernel structure of f :

- A For all $(k_l, B_l) \in K^l$, $(k, B) \in K$, and $(k_r, B_r) \in K^r$, such that

$$k_l + k_r \equiv 0 \pmod{\text{gcf}(k, p)}$$

and one of k_l, k_r, k is $\notin p * \mathbb{N}$, we have either $B_l > B$ or $B > B_r$.

- B For any $(k_l, B_l) \in K^l$, $B_l > L$ and for any $(k_r, B_r) \in K^r$, $B_r < R$.

Call these the p -equigrouping kernel conditions; let \mathcal{T}_p denote the set of all algorithms whose kernel structures satisfy these conditions. Let \mathcal{T}'_p denote the set of algorithms in \mathcal{T}_p whose kernel structures are finite sets in \mathcal{W}_∞ .

The key result is the following, giving a characterization \mathcal{F}_p for finite kernel structures:

Theorem 2 [Equigrouping Classification] For all p , $\tilde{\mathcal{F}}_p \subset \tilde{\mathcal{T}}_p$ and $\tilde{\mathcal{T}}'_p \subset \tilde{\mathcal{F}}_p$.

PROOF. [Theorem 2] We will prove the “only if” direction by contradiction. Let $f \notin \mathcal{T}_p$; we will construct an initial condition X and a semantic s such that $f_n^s(X)$ does not converge to C_p .

So suppose first that f violates the first p -equigrouping kernel condition. Let $(k_l, B_l) \in K^l(f)$, $(k, B) \in K(f)$ and $(k_r, B_r) \in K^r(f)$ be such that $k_l + k_r \equiv 0 \pmod{\text{gcf}(k, p)}$ and $B_l \leq B$ and $B \leq B_r$. Let g_l, g, g_r be carriers of these behaviors, i.e. $[g_l, f, s_l] = B_l$, $[g, f, s] = B$, $[g_r, f, s_r] = B_r$ for some semantics s_l, s, s_r , and with $|g_l| = k_l$, $|g| = k$, and $|g_r| = k_r$. Now, there must be an m such that $k_l + k_r + mk = np$ for some n . Let $X_{M,N,O} = g_l \oplus g^{\oplus m} \oplus g_r$ be a consecutive sum such that

$$\text{dist}(\text{re}(g_l), \text{le}(g^{\oplus m})) > M, \text{dist}(\text{re}(g^{\oplus m}), \text{le}(g_r)) > N$$

and distances between copies of g in $g^{\oplus m}$ are greater than O . Let $S = s_l \circ s \circ s_r$. When M, N, O are sufficiently large, the kernels present in the initial state never interact and therefore remain separated with their original behavior, that is

$$f_n^S(X) = (g_l, B_l) \oplus (g, B)^{\oplus m} \oplus (g_r, B_r)$$

is a consecutive kernel decomposition for all n . Now, if one of k_l, k or k_r is not a multiple of p then $\lim_{n \rightarrow \infty} \{f_n^S(X)\}$ either does not exist or is $\notin C_p$. But $|X|$ is a multiple of p so, this contradicts the assumption that f is a solution to p -equigrouping.

Suppose now that f violates the second p -equigrouping kernel condition, and wlog that $(k_l, L) \in K^l(f)$. Let (k_l, K) be carried by g_l under semantic s_l . Then denote by y any configuration in \mathcal{C} such that $|y| \in p\mathbb{Z}$ and let $y' = g_l \boxplus y$ where $\text{dist}(\text{re}(g), \text{le}(y)) > r(f)$. Let b denote the minimum of speeds of kernel decompositions of y under f , for some arbitrary semantic s and let $k \in \mathbb{N}$ be such that $k|n(g_l, f_l, s_l)| > b$. Let $\tilde{s} = s_l^{\circ k} \circ s$. The point is that $\text{Lim}(f^{\tilde{s}}(y')) \notin C_p$, a contradiction to the supposition that f is a solution to equigrouping. The analogous argument works to show that $(k_r, R) \notin K^r(f)$.

For the “if” direction, let (K^l, K, K^r) be a triple of finite subsets of \mathcal{W}_∞ satisfying the p -equigrouping kernel conditions. We must show that if there is any algorithm f with this kernel structure, then we can produce a solution $f \in \mathcal{F}_p$ with this kernel structure. Hence by theorem 1, given a finite potential kernel structure (K^l, K, K^r) satisfying the two p -equigrouping kernel conditions and the five conditions from theorem 1, we have to construct an algorithm with the given kernel structure. The construction $f = f[K^l, K, K^r]$ from the proof of Theorem 1 does this for us. The question is whether f thereby constructed is actually a solution. But:

Lemma 2 *If $f \in \mathcal{T}_p$, then for all periodic semantics s ,*

$$\lim_{n \rightarrow \infty} \{f_n^s(X)\} = \bigoplus_{i=1}^k g_i$$

*is a consecutive kernel decomposition, where for each i g_i is an (f, s) -kernel with $|g_i| \in p * \mathbb{N}$ and $[g_i, f, s] = S$.*

Combining this (proof given in the appendix) with lemma 1, to complete the proof of theorem 2 we must only show that $f[K^l, K, K^r]$ equigroups its stationary kernels whose size are a multiple of p . But recall that we can make the construction in theorem 1 with this property: one can see from the construction (in the appendix) that the only fixed kernels with size $m \times p$ are S_{mp} – which are already in C_p . \square

Unfolding the statement of the theorem, what we have shown is that for all k , a triple

$$(K^l, K, K^r) \in 2^{\mathcal{W}_k} \times 2^{\mathcal{W}_k} \times 2^{\mathcal{W}_k}$$

can be written as $(K^l(f), K(f), K^r(f))$ for an $f \in \mathcal{F}_p$ if and only if the triple satisfies the p -equigrouping kernel conditions and is the kernel structure of a well-defined algorithm.⁹ This result largely answers the questions posed at the beginning of the section. Its basic import is that any combination of group behavior that are not immediately ruled out by the p -equigrouping kernel conditions is a valid combination for at least one actual solution. The theorem provides the existence of a spectrum of infinitely many solutions to equigrouping all of which are qualitatively different with respect to the group behavior they induce on groups of agents of various sizes. We therefore have the following:

Corollary 1 *For each $p \geq 2$, $|\mathcal{F}_p| \geq |\tilde{\mathcal{F}}_p| = \infty$.*

The theorem has several useful corollaries for analyzing candidate solutions to equigrouping. For example, Proposition 2 follows immediately from the theorem:

PROOF. [Proposition 2] Suppose $f_l = f = f_r$ for a deterministic solution. Then $(1, B) \in K^l(f) \cap K(f) \cap K^r(f)$ for some B . But obviously $1 + 1 \equiv 0 \pmod{p}$; which then triggers the first p -equigrouping condition, which contradicts the existence of the identical behavior for the (unique) size-1 kernel in K^l, K , and K^r . \square

The notion of group behavior equivalence and characterization of solutions may be useful in helping to understand a variety of basic properties of distributed algorithms. For example, one can define the *kernel complexity* of f ,

$$KC(f) = \sum_{g \in K(f)} |g|,$$

as a potential complexity measure. We conjecture that $F_1(p)$ and its mirror algorithm minimize KC among solutions of p -equigrouping. $F_1(p)$ also minimizes $r(f)$ among p -equigrouping

⁹Actually, we show more than this: the “only if” part applies to all kernel structures, finite or otherwise.




Figure 4: Average time-to-solution for $F_1(4)$ plotted vs. density of agents in the initial conditions.

solutions. For these reasons, we believe that the simplest solution is perhaps an optimal solution.

Another notion of optimality is that of time-efficiency, the time-to-solution averaged over semantics and initial conditions. There is evidence that the efficiency profile – that is, efficiency as a function of agent-density in initial conditions – of a given f is a function of the equivalence class $[f]$, and that efficiency degrades with $KC([f])$. The efficiency profile of $F_1(p)$ is shown in figure 4.

In future work, we intend to characterize in much greater depth these other relevant questions about \mathcal{F}_p , including its scaling with $r(f)$, the information radius, and the structure of its *non-deterministic* part. A key question, however, is whether the modeling approach that made the kind of analysis done in this paper possible can be used for systems other than equigrouping. The next paper in this series addresses that question.

3. REFERENCES

- [1] H. Abelson, D. Allen, D. Coore, C. Hanson., G. Homsy, T. F. Knight, R. Nagpal, E. Rauch, G. J. Sussman, and R. Weiss. Amorphous computing. *Comm. ACM*, 43(5), 2001.
- [2] M. Brenner, L. Levitov, and E. Budrene. *Biophys. J.*, 74, 1998.
- [3] S. Camazine, J. L. Deneubourg, N. R. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau. *Self-Organizing Biological Systems*. Princeton Univ. Press, 2001.
- [4] J. Conway. *Scientific American*, March 1970.
- [5] M. Dorigo. ACM Dig. Lib., 1999.
- [6] S. Forrest, A. S. Perelson, L. Allen, and R. Cherukuri. In *Proc. IEEE Symp. RSP*, 1994.
- [7] D. M. Gordon. *Ants at Work*. Free Press, 1999.
- [8] A. Jadbabaie, J. Lin, and A. Morse. In *Proc. CDC03*, 2003.
- [9] J. M. Kahn, R. H. Katz, and K. S. J. Pister. Next century challenges: Mobile networking. In *Proc. 5th ACM/IEEE CMCN*, 1999.
- [10] E. F. Keller and L. A. Segel. Traveling bands of chemotactic bacteria. *J. Theor. Bio.*, 30, 1971.
- [11] E. Klavins. In *Proc. CDC03*, 2003.
- [12] C. Langton and K. Shimohara. *Artificial Life*. ACM Dig. Lib., 1989.
- [13] J. McLurkin. The ants. Master’s thesis, MIT, 1996.
- [14] M. Minsky. *Perceptrons*. MIT Press, 1988.

- [15] R. M. Murray. *Euro. J. Control*, 2003.
- [16] R. Nagpal. *MIT CSAIL AI Memo*, 1999.
- [17] P. Ogren, E. Fiorelli, and N. E. Leonard. In *Proc. MTNS*, 2002.
- [18] R. Olfati-Saber and R. M. Murray. In *Proc. CDC03*, 2003.
- [19] J. Pearl. *Probabilistic Reasoning*. Morgan Kaufmann, 1988.
- [20] C. Reynolds. Flocks, herds, and schools: A distributed behavioral model. In *Proc. SIGGRAPH*, 1987.
- [21] S. Sastry and M. Bodson. *Adaptive Control: Stability, Converges and Robustness*. Prentice Hall, 1989.
- [22] A. Stevens. *SIAM J. Appl. Math.*, 61(1), 2000.
- [23] P. Stone and M. Veloso. *Autonomous Robots*. Kluweronline, 2000.
- [24] J. Toner and Y. Tu. Flocks, herds, and schools. *Phys. Rev. E*, 58(4), 1998.
- [25] A. M. Turing. *Phil. Trans. Roy. Soc. B*, 1952.
- [26] J. von Neumann. *UIUC Press*, 1966.
- [27] S. Wolfram. *Rev. Mod. Phys.*, 55, 1983.
- [28] L. Wolpert. *J. Theor. Bio.*, 25(1), 1969.
- [29] D. Yamins and N. Khaneja. In *Submitted to Proc. ACC05*, 2005.
- [30] D. Yamins, S. Waydo, and N. Khaneja. In *Proc. CDC04*, 2004.
- [31] F. Zhang and P. Krishnaprasad. In *Proc. CDC02*, 2002.

4. APPENDIX

Suppose $X, Y \subset \mathcal{W}_\infty$. Let $X \uplus Y$ denote the set $\{x+y \mid \exists (x, B_1) \in X, (y, B_2) \in Y, B_1 \leq B_2\}$. Let $\uplus^n K$ be the n -times \uplus of K with itself and $\uplus^\mathbb{N} K = \bigcup_{n \in \mathbb{N}} \uplus^n K$. The full statement of theorem 1 is:

Theorem 1 (K^l, K, K^r) is the kernel structure of a well-defined algorithm only if

- 1 $\exists!(1, B) \in K$. Similarly for K^l and K^r .
- 2 If $(2, L)$ or $(2, R) \in K$ then $\exists(j, B) \in K$ for $j > 2$ or $B = S$.
- 3 The following equation holds:

$$\begin{aligned} \mathbb{N} &= \uplus^\mathbb{N} K = K^l \uplus (\uplus^\mathbb{N} K) = (\uplus^\mathbb{N} K) \uplus K^r \\ &= [K^l \uplus (\uplus^\mathbb{N} K) \uplus K^r] \cup [K^l \cap K^r]. \end{aligned} \quad (1)$$

- 4 If $(j, B) \in K^l$ (resp K^r) and $\nexists(j', B') \in K$ with $1 < j' \leq j$, then either $(1, B) \in K$ or $(1, L) \in K$ (resp $(1, R) \in K$) and $B = S$.

- 5 Suppose $(1, S) \notin K$ and let

$$\begin{aligned} j_B &= \max\{l \mid (\nexists(l', B') \in K \mid 1 < l' < l, B' < B) \\ &\quad \wedge (\exists(j, B'') \in K \mid 1 < j \leq l)\}. \end{aligned} \quad (2)$$

Then $\exists(i, C) \in K^r$ with either $i = l$ or $C \geq B$. Analogous conditions apply for K^l .

Conversely, all finite potential kernel structures satisfying these conditions arise.

PROOF. [Theorem 1] (\Rightarrow) We will procede by contradiction. Consider the first assumption, and define for any r the configuration

$$V_r = \circ \circ \dots \circ \bullet \circ \dots \circ = \circ^r \bullet \circ^r$$

in which \circ is repeated r times on either side. Under any $f \in \mathcal{A}_r$, V_r must be a kernel, for note that $[f(f(\bullet, V_r))] = [f(\bullet, V_r)]$ and therefore

$$[\circ^m f(\bullet, V_r)] = [f(\bullet, V_r)]$$

for all m . Hence $(1, [f(\bullet, V_r)]) \in K(f)$. Further more, it is the unique element $(1, B) \in K(f)$, since any configuration of size one is simply a single copy of \bullet itself, and of course when "isolated"

a single \bullet is exactly V_r , and so therefore has the same behavior. This proves the first condition.

Now consider the second condition and suppose it were otherwise. Let y be any configuration with $|y| > 4r$. Now, for all periodic s , from proposition 4 and the assumption that S is not a behavior of f ,

$$\lim_{n \rightarrow \infty} \{f_s^n(y)\} = \bigoplus_{i=1}^{m_1} k_i \oplus \bigoplus_{m_1+1}^{m_2} k_i$$

is a permanent decomposition in which $|k_i| \leq 2$ for all i , in which

$$[k_i, f, s] = \begin{cases} L & \text{for } 1 \leq i \leq m_1 \\ R & \text{for } m_1 + 1 \leq i \leq m_2 \end{cases}.$$

Let $x_l = \bigoplus_{i=1}^{m_1} k_i$ and $x_r = \bigoplus_{m_1+1}^{m_2} k_i$. Since $|y| > 4r$, either $|x_l|$ or $|x_r|$ is greater than $2r$. Suppose wlog that $|x_l| > 2r$. The set $\lim_{n \rightarrow \infty} \{f_s^n(x_l)\}$ must admit a function $\rho: B_r(a, X) \rightarrow \{1, 2\} \times \{1, 2\}$, such that $\rho(a) = (m, |k_i|)$, where m is the number in the unique k_i containing a as read from the left. ρ must be invariant under changes in $\text{dist}(\text{le}(k_i), \text{re}(k_{i-1}))$ and $\text{dist}(\text{re}(k_i), \text{le}(k_{i+1}))$ for all i . But now consider two consecutive agents a, b in x_l with more than r agents on either side in x_l . It is always possible to find a periodic s' such that the resulting x_l has the two properties: 1) if $|k_i| = 1$ then $|k_{i-1}| = |k_{i+1}| = 2$ if they exist and 2) there are n, m such that $b_r(a, f_n^{s'}(x_l)) = b_r(b, f_m^{s'}(x_l))$. But this is impossible. The first condition implies that $\rho(a) \neq \rho(b)$ when no 3-kernels or larger exist; and the second requires $\rho(a) = \rho(b)$; a contradiction.

The third condition follows easily from proposition 4; any initial configuration under a periodic semantic s eventually has a consecutive decomposition $\bigoplus_i^k g_i$ where $B_i = [g_i, f, s] \leq [g_{i+1}, f, s]$. But initial configurations can be made of any size $n \in \mathbb{N}$. Hence any n must be able to be written as $\sum_i |g_i|$. By allowing the initial condition to contain $\text{re}(X)$, $\text{le}(X)$, or both, we get the three equalities; in the first case, g_k is in $K^r(f)$; in the second, $g_1 \in K^l(f)$; in the third case, both care - and they could be identical, which gives the $K^l \cap K^r$ term. With this same argument applied to initial conditions not containing end-agents, we also get $\mathbb{N} = \uplus^\mathbb{N} K$. But this is already obvious from the fact that $(1, B) \in K$ - any n can be written as $1 + \dots + 1$, repeated n times, where the 1s are carried by $(1, B)$.

Consider the fourth condition and suppose it were otherwise. Let $(j', B) \in K^l$ be carried by a kernel g under periodic semantic s with $j' < j$. Let s' be the semantic created from s by removing calls to $\text{le}(g)$ whenever $\text{dist}(\text{le}(g), g \ominus \text{le}(g)) = 1$ and $[f(\text{le}(g), g \ominus \text{le}(g))] = R$ and to $b = \text{le}(g \ominus \text{le}(g))$ whenever $[f(b, g \ominus \text{le}(g))] = L$. s' is eventually periodic. Now there are two cases. If $s' \in \mathcal{SEM}$, then the behavior of agents $g \setminus \text{le}(g)$ be that unique B' such that $(1, B') \in K$ (otherwise there would have to be a non-trivial kernel to which these agent belonged). So therefore the whole kernel has that behavior as well, i.e. $B = B'$ and $(1, B)$ is the unique 1-kernel in K . If on the other hand $s' \notin \mathcal{SEM}$, then it must be that $B = S$ and $B' = L$.

Finally, suppose the fifth condition did not hold, and wlog that it doesn't hold for K^r . Then there is a configuration z all of whose limiting decompositions in $\uplus^\mathbb{N} K \uplus K^r$ are of the form $x \oplus y = (1, B)^{\oplus m} \oplus (j, B')$ in which $B \leq B'$; and for which

$$\begin{aligned} m &\geq \min\{l \mid (\nexists(l', B'') \in K \mid 1 < l' \leq l, B'' \leq B) \\ &\quad \wedge (\exists(j, C) \in K \mid 1 < j \leq l)\}. \end{aligned} \quad (3)$$

But then there would be an s' under which $x = (1, B)^{\oplus m} \rightarrow x' \oplus (j, C)$ in which $C > B'$. But therefore any kernel decomposition $x \oplus y$ can be disturbed and is not a limiting decomposition; hence not all limiting decompositions are of the form supposed above, contradicting the assumption.

(\Leftarrow) - (Proof sketch). Now let $(K^l, K, K^r) \in 2^{\mathcal{W}^k} \times 2^{\mathcal{W}^k} \times 2^{\mathcal{W}^k}$ for a minimal k that satisfies conditions i) - v). We need

to construct an algorithm $f[K^l, K, K^r]$ with this as its kernel structure. Let $r(f) = 2(17k + 20)$.

For each $n \geq 2$, define configurations L_n, R_n , and S_n of n agents such that if $x, y \in \mathcal{X} = \bigcup_n \{L_n, S_n, R_n\}$, then $x \cap y \neq \emptyset$ IFF $x = y$. This can be done with $r(f) = 2(17k + 20)$ and with $S_{np} \subset C_p$ for a fixed p . Furthermore, define \tilde{f} so that it traces out trajectories $T_n(L)$ on L_n that cycle back to L_n in relative position, but translated to the left; $T_n(R)$ similarly for R_n , to the right; and $T_n(S)$ which leave S_n stationary (but fixed IFF n is a multiple of p). This can be achieved with the properties that 1) if $x, y \in \mathcal{T} = \bigcup_n T_n(L) \cup T_n(R) \cup T_n(S)$ then $x \cap y \neq \emptyset$ iff $x = y$ and 2) the various T_n are mutually disjoint, given that $r(f) = 2(17k + 20)$. Define $\tilde{f}_{L,R,S}^l$ trajectories for K^l and K^r similarly, but making allowances for the kernel structure imposed by conditions 4 and 5.

Suppose $a \in X$ in which $b_r(a, X)$ is isolated single R_n, S_n or L_n , and in which $a \neq le(b_r(a, X)), re(b_r(a, X))$, and let $j = \min\{l \mid \exists (l, B) \in K, \exists (l', B') \in K, 1 < l' \leq l\}$. Then define f by $([g], n[g]) \in K \Rightarrow f(a, X) = \tilde{f}(g)$; and similarly for $(g, n[g]) \in K^l \cup K^r$ when $|g| \geq j$ and $a \neq le(X), re(X)$.

Finally, define

$$(l, B) \in K^l \Rightarrow f_l(le(X), X) = \begin{cases} \tilde{f} & \text{for } l \geq j \\ \tilde{f}_L^l & \text{for } l < j, \quad B = L \\ \tilde{f}_R^l & \text{for } l < j, \quad B = R \\ \tilde{f}_S^l & \text{for } l < j, \quad B = S \end{cases}$$

and analogously for K^r . These definitions are possible given that $r(f) = 2(17k + 20)$, larger than the length of any possible kernel.

For the map f as defined, $K^l \subset K^l(f), K \subset K(f), K^r \subset K^r(f)$. We have ensured that the required potential kernel structure is a subset of the kernel structure of f . However, we to ensure that no *other* kernels exist for f that arise via unexpected interactions of the kernels defined thus far. To prevent unwanted kernels from arising, we define several kernel interaction and transition rules so that in the limit a good decomposition of elements in the specified K s always occurs. An interaction rule specifies what happens when two distinct kernels come into contact with each other; to this end we need to ensure that there is a unique partitioning of any given configuration into constituent L_n, S_n and R_n s. The following rule is used:

For any $a \in X$: let $k'(a, X)$ be the largest L_n, R_n or S_n that a is contained in; if $a = le(g_1) = re(g_2)$ for two equally-sized L_n, R_n, S_n , then let

$$k'(a, X) = \begin{cases} g_1 & \text{if } [g_1] \leq B_1 \\ g_2 & \text{if } [g_1] > B_1 \end{cases}.$$

Finally, let

$$k(a, X) = \begin{cases} k'(a, X), & k'(le(k'(a))) = k'(a) = k'(re(k'(a))) \\ (1, B_1) & \text{otherwise} \end{cases} \quad (4)$$

where we've dropped the X from the notation $k(a, X)$ for simplicity. Again, these definitions are possible given that $r(f) = 2(17k + 20)$ is in fact more than two times larger than the length of any possible kernel.

Now, consider the following transition rules, in which we've used the partition given by $k(a, X)$:

1) First, for all $x, y, n, m \geq 2$,

$$x \oplus R_n \oplus L_m \oplus y \xrightarrow{f} x \oplus L_m \oplus R_n \oplus y.$$

That is, the ‘‘conflict’’ of having a left-moving kernel to the right of a right-moving kernel is resolved by having the kernels ‘‘switch places’’ (though no agent ever physically goes past another). Similarly for the other ‘‘conflicting’’ kernel combinations.

2) Suppose $(2, R) \in K$, then let $m = \min\{l \mid (l, B) \in K, l > 2 \text{ or } B = S\}$ and let g be the carrier of such a kernel. (Guaranteed by condition 2) Then

$$(2, R) \oplus^{\lceil \frac{m}{2} \rceil} \rightarrow (1, B_1)^{\text{odd}(m)} \oplus g.$$

in which the notation $\text{odd}(m)$ denotes 1 if m is odd and 0 if it is even. Analogously for $(2, L)$.

3) Similarly, if $(1, S) \notin K$, left-isolated groups of 1-kernels of the form $X = (1, B)^{\oplus j}$ transition to the kernels guaranteed by condition 4; and analogously for right-isolated groups $(1, B)^{\oplus j}$. We say that such a group is left-isolated if

$$\text{dist}(le(X)^-, re(X)) > 2(17k + 40) - \alpha$$

or $le(X)^-$ is not in a 1-kernel; and α is just large enough to ensure $re(X)$ can determine whether $le(X)^-$ is in a 1-kernel. Right-isolated defined analogously.

4) Consider

$$X = R_n \oplus K_1 \oplus \bigoplus_{i=1}^j K_i \oplus s$$

where the R_n is a right or stay kernel containing the left-most agent $le(X)$ and K_j is the right-most kernel that $le(X)$ can ‘‘see’’ entirely. The transition rule is: when possible, $X \rightarrow \bigoplus_{i=1}^m g_i$, a good decomposition, where $[g_i] = S_i$ or L_i for some i ; when not possible, X transitions to a good decomposition, but with as few 1-kernels as possible and $le(X)$ translated to the right, if space between K_j and s allows. Analogous rules apply to right end configurations.

5) Finally, for all n when $(n, L) \in K^l$ but not in K , then $L_n \xrightarrow{f} M$ where M is an allowed good decomposition of n agents in $K^{\oplus \mathbb{N}}$ alone, with the fewest 1-kernels. Similarly, for all n when $(n, L) \in K$ but not in K^l , then $L_n \xrightarrow{f} M$ where M is an allowed good decomposition of n agents in $K^l \oplus K^{\oplus \mathbb{N}}$. The same holds for the right-end kernels and the other behaviors S and R .

That these transitions rules can be made consistent is shown with simply but somewhat tedious arguments not given here. \square

PROOF. [Lemma 1] Suppose s is an aperiodic semantic and x_0 an initial condition such that $\{f_n^s(x_0)\}$ does not converge to a fixed element in C_p . Suppose that the sequence $d_n = \text{dist}(le(f_n^s(x_0)), re(f_n^s(x_0)))$ is bounded. Then the relative positions $\{f_n^s(x_0)\}$ form a finite set; hence there is an $x_m = f_m^s(x_0)$ which is not in C_p , such that for some $m' > m$, the relative positions of x_m and $f_{m'}^s(x)$ are the same, and such that $\hat{s} = (s_m, \dots, s_{m'})$ mentions every agent at least once. But then under the periodic semantic $s^* = (\hat{s}, \hat{s}, \dots)$, f does not solve x , a contradiction. Now suppose on the other hand that d_n is not a bounded sequence. If $\lim(g, f, s)$ is a good decomposition $\bigoplus_{i=1}^m g_i$ then for d_n to be unbounded, then there is a finite subsequence $t \subset s$ for which either g_1 or g_m is a (f, t^{∞}) -kernel with $[g_1, f, t] = L$ or $[g_m, f, t] = R$. But this was shown in the arguments of the proof of the first part of Theorem 2 to be impossible for any $f \in \mathcal{F}_p$, a contradiction. Similar (though slightly more complicated) arguments apply to the case that $\lim(g, f, s)$ is not a good decomposition. \square

PROOF. [Lemma 2] Appealing to proposition 4 we see that $\lim_{n \rightarrow \infty} \{f_n^s(x)\}$ has a unique permanent consecutive kernel decomposition $\bigoplus_1^k (g_i, B_i)$ with $B_i \leq B_{i+1}$. But since $f \in \mathcal{T}_p$, $B_1 = B_k = S$. Hence $B_i = S$ for all i . But then suppose that $|g_i| \neq mp$ for any $m \in \mathbb{N}$. In this case, we violate the second of the p -equigrouting kernel conditions, so $f \notin \mathcal{T}_p$ as presumed, a contradiction. \square