

AntPaP: Patrolling and Fair Partitioning of Graphs by A(ge)nts Leaving Pheromone Traces

G. Elazar, A. Bruckstein

Gidon Elazar

Department of Computer Science
Israel Institute of Technology-Technion, Haifa, Israel

Alfred M. Bruckstein*

Department of Computer Science
Israel Institute of Technology-Technion, Haifa, Israel

*Corresponding author: freddy@cs.technion.ac.il

Abstract

A team of identical and oblivious ant-like agents – a(ge)nts – leaving pheromone traces, are programmed to jointly patrol an area modeled as a graph. They perform this task using simple local interactions, while also achieving the important byproduct of partitioning the graph into roughly equal-sized disjoint sub-graphs. Each a(ge)nt begins to operate at an arbitrary initial location, and throughout its work does not acquire any information on either the shape or size of the graph, or the number or whereabouts of other a(ge)nts. Graph partitioning occurs spontaneously, as each of the a(ge)nts patrols and expands its own pheromone-marked sub-graph, or region. This graph partitioning algorithm is inspired by molecules hitting the borders of air filled elastic balloons: an a(ge)nt that hits a border edge from the interior of its region more frequently than an external a(ge)nt hits the same edge from an adjacent vertex in the neighboring region, may conquer that adjacent vertex, expanding its region at the expense of the neighbor. Since the rule of patrolling a region ensures that each vertex is visited with a frequency inversely proportional to the size of the region, in terms of vertex count, a smaller region will effectively exert higher “pressure” at its borders, and conquer adjacent vertices from a larger region, thereby increasing the smaller region and shrinking the larger. The algorithm, therefore, tends to equalize the sizes of the regions patrolled, resembling a set of perfectly elastic physical balloons, confined to a closed volume and filled with an equal amount of air. The pheromone based local interactions of agents eventually cause the system to evolve into a partition that is close to balanced rather quickly, and if the graph and the number of a(ge)nts remain unchanged, it is guaranteed that the system settles into a stable and balanced partition.

Keywords: multi-agent, graph partition, simple interactions

1 Introduction

Patrolling is continuously traveling through an environment in order to supervise or guard it. Although mostly used to refer to humans guarding an area, the term patrolling is also used to describe

surveying through a digital, virtual environment. Consider, for example, the task of repeatedly reading web pages from the world-wide-web in order to keep an updated database representing the links between pages, possibly for the purpose of later retrieval of pages in an accurate and prompt manner. These problems exhibit similarities, in the sense that they can be represented as traveling through the vertices of a graph. But there are also differences: a physical area is usually fixed in size, whereas the virtual area is, in general, prone to constant change. The number of human guards is, generally, fixed for the particular area being patrolled, while the number of software agents or “bots” performing a large scale patrolling task may be subject to change as well.

Partitioning a graph into similar sized components is an important and difficult task in many areas of science and engineering. To name few examples, we can mention the partitioning of a netlist of an electronic VLSI design [5], the need for clustering in data mining [4], and the design of systems that balance the load on computer resources in a networked environment [6].

The general graph partition problem is loosely defined as dividing a graph into disjoint, connected components, such that the components are *similar* to each other in some sense. Practical considerations impose additional constraints. For example, an important problem, known as the *graph k-cut*, requires a partition where the sum of the *weights* of vertices belonging to each component is more or less equal, and additionally, the number and/or the sum of weights of edges that *connect disjoint components* is minimized [7]. The *k-cut* problem can model the distribution of tasks between computers on a network, while minimizing communication requirements between them.

In this work we define a patrolling strategy that fairly divides the work of patrolling the environment among several a(ge)nts by partitioning it into regions of more or less the same size. We have no constraints on edges connecting different components, but we impose strict restrictions on our patrolling agents in search for a heuristic multi-agent graph partitioning algorithm that may continuously run in the background of a host application. We are interested in programming the same behavior for each individual ant-like agent, which should be very simple in terms of resources, hardware or communications. Furthermore the agents have no id’s, hence are part of a team of units that are anonymous and indistinguishable to each other. Our a(ge)nts should have very little knowledge about the system or environment they operate in, have no awareness on the size or shape of the graph, no internal memory to accumulate information, nor a sense of the number and locations of other agents active in the system. These limitations mean that such a multi agent process has inherent scalability; the environment might be large, complex, and subject to changes, in terms of vertices, edges and even the number of agents, and our simple agents should still be able to patrol it, while also evolving towards, and ultimately finding balanced partitions, if such partitions exist. To simplify the discussion, we will think of a graph to be partitioned as a planar area, and the task at hand will be to partition the area into regions of more or less the same size. The area is modeled as a grid, where each vertex is a unit area, thus a balanced partition should have components of roughly the same number of vertices. In our scheme, agents are each given the task to *patrol* and define a region of their own, and have the ability to expand their region via conquests. Like ants, our agents leave pheromone marks on their paths. The marks decay with time and are subsequently used as cues by all the agents to make decisions about their patrolling route and about the possibility to expand their region. By assumption, each agent operates *locally*, thus it can sense levels of pheromones or leave pheromone marks on the vertex it is located, on its edges and on adjacent vertices.

While patrolling its region, an agent visits a vertex and reads the intensity of pheromone marks that remain from previous visits. It then uses the reading, and the known rate of pheromone decay, to calculate the vertex’s idle-time – the time that passed since the previous visit. Using the decaying pheromone mark we can choose a patrolling rule according to which the agents visit the vertices of their region in repetitive cycles, each vertex being marked with a pheromone once on each cycle. The patrolling process hence ensures that the idle-time measured by agents on visits to their region’s vertices is the same, effectively encoding their region’s *cover time* – the time that takes for an a(ge)nt to complete a full patrolling cycle – and therefore it can also be used to estimate the region’s size: the shorter the cover time, the smaller the region.

We assume that each agent detects pheromones without being able to distinguish between them, except for recognizing its own pheromone. When an a(ge)nt hits a border edge – an edge that connects

its region with one that is patrolled by another agent - it can use the neighbor's idle-time (encoded in its pheromone marks) to calculate the size of the neighboring region, and thereby decide whether to try to conquer the vertex "on the other side of the border". This causes an effect that mimics pressure equalization between gas-filled balloons: at two vertices on opposing ends of a border edge, the agent that hits the border more frequently is the agent with a shorter cover time (patrolling the smaller region) hence it may attempt a conquest.

We define that in a balanced partition, any pair of neighboring regions have a size difference of at most one vertex. This means that for a graph \mathcal{G} and k agents, our partitioning heuristics ensures a worst case difference of $k - 1$ vertices between the largest and smallest of the regions, once a balanced partition is reached. For example in a graph of 1 million vertices (e.g. database entries, each representing a web page) and 10 agents (network bots patrolling the pages), this difference is truly negligible. Additionally, the length of the patrolling path is predetermined, and is proportional to the size of the region being patrolled, therefore when a balanced partition is reached, the algorithm guarantees that the idleness of any of the vertices of the graph is bounded by a number of steps equal to $2 \left(\left\lfloor \frac{|\mathcal{G}|}{k} \right\rfloor + (k - 1) \right) - 1$, about twice the size of the largest possible region (Note that $|\mathcal{G}|$ denotes the number of vertices in the graph \mathcal{G}).

In Figure 1 one can see a series of snapshots depicting 8 patrolling agents working to partition a 50x50 grid. The first snapshot shows an early phase of the joint patrolling algorithm, where agents already captured some of the vertices around their initial random locations, in the second, the area is almost covered and most of the vertices of the graph are being patrolled, the third exhibits a phase when all the area is covered but the regions are not balanced, and finally, the last snapshot shows a balanced partition that the system evolved into.

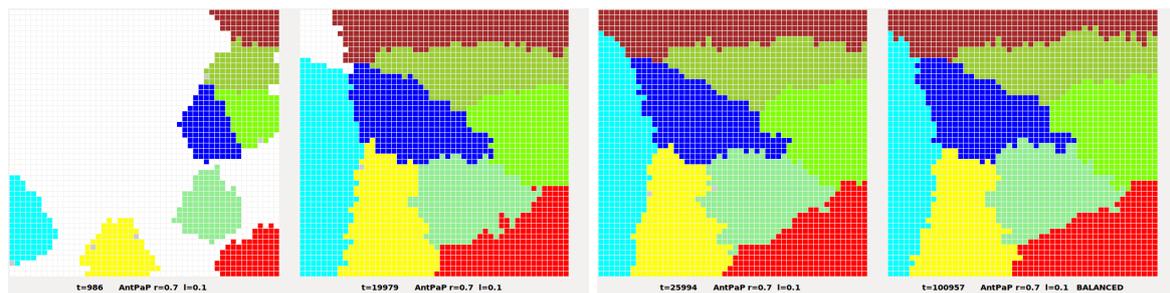


Figure 1: Evolution of a 50x50 Grid Graph Partitioning by 8 A(ge)nts

This Figure exhibits typical stages in the evolution of such a system, for which balanced partitions exist, and the environment graph remains stationary for a time long enough for agents to find one of them. Often, the agents will relatively quickly find a partition that covers the graph, and is *close* to being balanced. Then, on stationary graphs, they may spend a rather long time to reach a perfectly balanced partition. In a time varying environment the system will continuously adapt to the changing conditions.

2 Related Work

The concept of partitioning a graph with a(ge)nts patrolling a region and exerting pressure on neighboring regions was first presented by Elor and Bruckstein in [1]. They proposed a patrol algorithm named BDFS –Balloon DFS – and this work is a follow up research on this problem. According to BDFS, an agent patrolling a smaller region conquers vertices from a neighboring larger region. To achieve the goal of patrolling an area, BDFS uses a variation of Multi-Level Depth-First-Search (MLDFS), an algorithm presented by Wagner, Lindenbaum and Bruckstein in [2]. The task of the MLDFS too, was to distribute the work of covering "tiles on a floor" among several identical agents. The floor-plan mapping of the tiles is unknown and may even be changing, an allegory for moving furniture around while agents are busy cleaning the floor. MLDFS implements a generalization on DFS: agents leave decaying pheromone marks on their paths as they advance in the graph, and then

use them either to move to the vertex least recently visited or to backtrack. When none of the choices are possible, either when the graph covering ends, or following to changes in the graph or loss of tracks due to noise in the pheromone marks, agents *reset*, thus starting a new search. The time of reset, named “the time where new history begins”, is stored in the agents’ memory, as a **search-level** variable. After a *reset*, the cycle repeats, hence an agent traces pheromone marks left in an earlier cycle. The mere existence of a pheromone mark is, however, not sufficient for agents to choose a path not yet taken during the current search cycle. To select the next step, agents use the value stored in the **search-level** variable as a threshold: any pheromone that was marked on a vertex or edge prior to this time must have been the result of marking in an earlier cycle. In MLDFS, pheromones of all agents are the same, and agents are allowed to step on each other’s paths. For the task of partition a graph, in BDFS each agent has its own pheromone and it performs MLDFS cycles on its “own” region of the graph, leaving its particular pheromone marks. As long as the region is stationary, BDFS agents exactly repeat their previous route. If the region changes, either expands or shrinks, it will cause BDFS to look for a new and possibly substantially different route before settling into the next search cycle. This occurs due to a subtlety in the way that depth-first-search defines a spanning tree, a special type of tree called a *palm tree*, where each edge (v, w) not in the spanning tree connects a vertex v with one of its ancestors, see e.g. Tarjan [8]. The spanning tree defined during a BDFS search cycle does not consider all edges emanating from *all* of the region’s vertices, simply because some of the edges connect to vertices *on neighboring* regions. When BDFS conquers a vertex, it is possible that this vertex has more than one edge connecting to the region. All these edges will now be considered during the next search cycle, a process that may dictate a different palm tree. We call this event - *respanning*. In the algorithm we define here, named *Ant Patrolling and Partitioning, or AntPaP*, we use a different generalization of DFS that avoids *respanning*. Furthermore we reduce the requirements on the agent’s capabilities. For example, our agents have no memory, and also cannot control the levels of pheromones they leave, the pheromone level at the time of a marking is always the same. We further add the possibility for agents to *lose* a vertex if a conquest fails, and we provide a proof of convergence to balanced partitions, while experimentally observing much faster evolution towards such partitions.

The subject of multi-agent patrolling has been extensively studied. Lauri and Charpillet [12] also use an “ants paradigm”, where a method based on Ant Colony Optimization (ACO), introduced by Dorigo, Maniezzo, and Coloni in [13]. ACO provides multi-agent solutions for various problems, for example the Traveling Salesman Problem (TSP) in complete and weighted graphs by a so-called *ant-cycle* algorithm. Ants move to the next vertex according to a probability that is biased by two parameters: the closest neighbor vertex (corresponding to the lowest edge weight) and the level of pheromone on the edge. During their search, ants record their path to avoid visiting the same vertex twice. Since at each step all ants traverse one edge to a neighboring vertex, all ants complete their travel at the same time. Thereafter each ant leaves pheromone marks on the entire path it took. Due to the probability bias, shorter edges have a higher probability to be traversed, thus it is probable that multiple ants traversed them, hence they tend to accumulate stronger pheromone levels. The cycles repeat, and with each cycle the biasing gets stronger towards the shortest path. The process ends after a a-priori given number of cycles complete or when the ants all agree on the shortest path. For the patrolling problem, Lauri et. al. [12] used this method to find multiple paths, one for each agent, *before* the agents begin their joint work. Their algorithm employs multiple colonies of ants where each agent is assigned one ant on each colony. Ants in a colony cooperate (exchange information regarding their choices) to divide the exploration into disjoint paths, leading the agents to eventually cooperate in the patrolling task. Unlike for the TSP, the environment graph is not required to be complete, and ants are allowed to visit a vertex more than once when searching for a patrolling route.

Chevalyere, Yann, Sempe, and Ramalho [11] compared cyclic patrolling strategies, in which agents tend to follow each other, to partitioning strategies, in which agents patrol each its own region. By applying several algorithms on several graphs examples, they found that the choice of strategy should be based on the shape of the graph. The partitioning based strategy gets better results on “graphs having ‘long corridors’ connecting sub-graphs”, i.e. if there are high weight edges that are slow to traverse, it is better not traverse them at all by allocating them to connect disjoint partitions.

There is substantial research on heuristics for partitioning of a graph, and some of it even related to multi-agent scenarios. Inspired by ants, Comellas and Sapena [10] presented yet another *Ants* algorithm to find a k -cut solution to a graph. The system is initiated by randomly coloring all the graph vertices in a more or less even number of colors and positioning the agents randomly on the graph. Then a *local cost* value is calculated for each of the vertices, storing the percentage of neighbors that have the same color as its own. Agents will then iteratively move to a neighboring vertex v that has the lowest cost (i.e. with the most neighbors of a *different* color than its own), and then switch colors with a random vertex u on the graph, where the color of u is the one most suitable for v , i.e. similar to most of v 's neighbors. u is selected from a random list of vertices colored with the same color as u , by choosing from this list the one with lowest cost (most neighbors colored *differently* than u). Then the cost value is refreshed for both v and u . On each iteration the number of cuts, defined as the number of edges connecting vertices of different colors, is calculated over all the edges of the graph and the lowest value is stored. The choices of agent moves are stochastic, i.e. agents have a probability p to select the next vertex to move to by using the cost value, otherwise it selects another neighbor vertex at random. This allows the system to escape from local minima. Unlike our algorithm, agents of Comellas and Sapena's *Ants* aim to find a k -cut, and while doing so do not leave pheromones to be used as cues on vertices and edges they visit as our agents do. Also, their agents are assumed to have the ability to look at vertices that are anywhere in the graph and change their values, thus their sensing is not local as in our algorithm. In *Ants*, each iteration relies on a global calculation that involves access to values on *all* edges of the graph, in order to measure the quality of the partition so far determined, as well as storing the result.

Inspired by bee foraging, McCaffrey [9] simulates a bee colony in order to find a k -cut graph partition. Each of the agents, in this case called *bees*, is assumed to know in advance the size and shape of the graph, as well as the number of components desired. The agent must have an internal memory to store an array of vectors listing the vertices of all sub-graphs of a proposed solution, as well as the number of cuts this partition has, as a measure of its quality. In a hive, some 10% of the bees are considered *scouts*, all other agents being in one of two states, *active* or *inactive*. Emulated scouts select a random partition of the graph. If the selection is better than what the scout previously found, it stores it in its memory and communicates it to other bees in the hive that are in an *inactive* state. Some of those store the scout's solution in their own memory, change their state to *active* and begin to search for a better partition around this solution. If an active bee finds an improved solution it communicates it to the bees that are left in the hive. After looking at neighboring solutions for a long enough time, the *active* bee returns to the hive and becomes *inactive* again. The algorithm, therefore, is constantly searching for improvements in the quality of the partition that the bees collectively determine.

The partitioning and patrolling multi-agent algorithms that we have surveyed above, all assume that agents possess substantial internal memory. Some algorithms assume that the agents are able to sense and even change values of vertices and/or edges in graph locations that are distant from their position in the graph, and sometimes they can even sense and/or store a representation of the whole graph in their memory. Patrolling algorithms may be partition based, and then the task is divided into two stages. In the first stage the graph is partitioned into disjoint components, and at the second stage each of the agents patrols one of those components.

In our case, partitioning the graph, and thereby balancing the workload among our agents, is a requirement. Our algorithm does not have stages, the agents simply perform pheromone directed local steps thereby carrying out a *patrolling* algorithm, and while doing so also implicitly cooperate in partitioning the graph. Our agents have no internal memory at all. Their decisions are based on pheromone readings from vertices and edges alone, and they can only sense or leave pheromone marks around their graph location. One may view our solution for patrolling and partitioning the graph environment as using the graph as a shared distributed memory for our oblivious agents.

3 The AntPaP Algorithm and Empirical Results

The task analyzed here is the partition of an area or environment into regions of similar size by a set of agents with severe restrictions on their capabilities. The inspiration for the algorithm are gas filled balloons; consider a set of elastic balloons located inside a box, and being inflated at a constant and equal rate, until the balloons occupy the entire volume of the box. While inflating, it may be that one balloon disturbs the expansion of another balloon. This may cause a momentary difference of the pressure in the balloons, until the pressure difference is large enough to displace the disturbing balloon and provide space for the expansion of the other. Since the amount of gas is equal for all balloons, they will each occupy the same part of the volume, effectively partitioning it into equal parts. Our agents mimic this behavior by patrolling a region of the area “of their own”, while continuously aiming to expand it. The area is modeled by a graph and the region is a connected component of the graph. When expanding regions touch, the agent on the smaller region may conquer vertices of the larger region. We assume that initially a given number of agents are randomly placed in the environment, they start the process of expanding and this process goes on forever. Eventually the expansion is “contained” due to the interaction between the regions of the agent, hence the process will lead to an equalization of the sizes of the regions patrolled by the agents. In the discrete world of our agents a partition to regions of exactly the same size may not exist, therefore we define a balanced partition as such that any two neighboring regions may have a size difference of at most one vertex.

Agent Modeling and Implementation Details

For simplicity, a(ge)nts operate in time slots, in a strongly asynchronous mode, i.e. within a time interval every agent operates at some random time, so that they do not interfere with each other. During a given time slot, each agent may move over an edge to another vertex, and may leave pheromone marks on a vertex and/or edges. The marks, if made, are assumed to erase or coexist with the pheromone that remained there from the previous visit. Agents have no control over the amount of pheromone they leave, its initialization level being always the same. Thereafter, the pheromone level decays in time. Each agent has its own pheromone, thus pheromones are like colors identifying the disjoint components and hence the partitioning of the graph. The agents themselves can only tell if a pheromone is their own or not. Agents are oblivious, i.e. have no internal memory. On each time slot, an agent reads remaining pheromone levels previously marked on the vertex it is located and its surroundings, and bases its decisions solely upon these readings. The readings and decisions are transient, in the sense that they are forgotten when the time slot advances. Decaying pheromone marks on vertices and edges linger, serving both as distributed memory as well as means of communication. In our model, agents leave pheromones in two patterns: one pheromone pattern is marked when agents advance in their patrolling route, and the second pattern is used when agents decide to remain on the same vertex. Pheromones are decaying in time, thus once marked on a vertex or edge, their level on the vertex or edge decreases with each time step. A straightforward way for implementing such behavior in a computer program, is to use the equivalent “time markings”, i.e. stamping the *time* at which a pheromone is marked on the vertex or edge. We therefore denote by $\varphi_0(v)$ the time of pheromone marking on vertex v , hence $\varphi_0(v) = \tilde{t}$ means that an agent left a pheromone on vertex v at time \tilde{t} . As time advances, the “age” of the pheromone on vertex v , i.e. the time interval since it was marked, which can be calculated as $t - \varphi_0(v)$ where t is the current time, advances as well. This is equivalent to measuring the level of the temporally decaying pheromone on vertex v , and using its value along with the known rate of decay to calculate its “age”. Similarly, $\varphi(u, v) = t$ is a time marking, equivalent to the decaying pheromone level on the edge, where $\varphi(u, v)$ and $\varphi(v, u)$ are not necessarily the same. The use of time markings require the computer program implementation to know the current time t in order to be able to calculate the age of pheromones. However, the knowledge of current time is strictly limited to its use in the emulation of temporally decaying pheromones by equivalent time markings, thus it does not depart from our paradigm of obliviousness and local decisions based on decaying pheromone markings only. When an agent decides to leave a pheromone mark on a vertex, it may avoid erasing the pheromone that remains from the previous (most recent) visit. We denote the previous time marking as $\varphi_1(v)$, thus when an agent marks a pheromone on vertex v , the computer program implementation moves the value stored in $\varphi_0(v)$ to $\varphi_1(v)$ and afterwards sets the new time mark to $\varphi_0(v)$. Hence, the value, $\varphi_0(v) - \varphi_1(v)$ encodes the *idle time* of the vertex v .

The Patrol Algorithm

Agents patrol their region in a DFS-like route, in the sense that they advance into each vertex once and backtrack through the same edge once during a complete traversal of their region. When an agent completes traversing the region it resets (i.e. it stays at the same location for one time step and refreshes its pheromone mark), and subsequently starts the search again. The cycles repeat the same route as long as the region is unchanged. When the region does change – either expanding or shrinking – out agents persist on keeping advancing and backtracking into a vertex through the same edge that was used to conquer the vertex. This is implemented by marking “*pair trails*”, i.e. leaving pheromones over edges as well as vertices, when conquering and (subsequently) advancing into a vertex. A pair trail is a directed mark from a vertex u to an adjacent vertex v , of the form $\varphi_0(v) = \varphi(u, v) = \varphi(v, u)$, and is one of the two pheromone patterns that agents leave. This behavior results in a patrolling process that *follows the pair trails*, were agents advance through the earliest marked pair trail, refreshing the marks while doing so. When all pair trails to advance through are exhausted, it backtracks through the same pair trail it entered. An example of a route and the spanning tree it defines are depicted in Figure 2. The departure from the classic DFS is that edges that are not marked as pair trails are ignored. The pair trails mark a *spanning tree* (which is not necessarily a palm tree) of the region, where its root is the vertex where the search cycle begins, and each pair trail marks the path advancing upwards the tree. When an agent backtracks to the root, it has no untraveled pair trail to advance through, and it restarts the search cycle remaining one time slot in the root. It then uses the second marking pattern which is simply leaving a pheromone on the vertex, denoted as $\varphi_0(u) = t$, where u is the root.

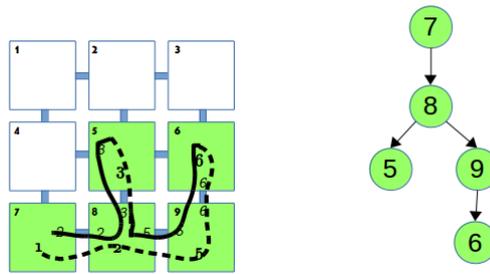


Figure 2: An example of a patrolling route in a region, and the spanning tree it defines. The arrows indicate *pair trail* directions: advance in the direction of the arrow, and (eventually) backtrack the other way

Since agents advance and backtrack once from each vertex in their region (except the root) and then restart a patrolling cycle in the root, the number of steps in one patrolling cycle, called the *cover time*, is $\Delta t_\alpha = 2|\mathcal{G}_\alpha| - 1$. \mathcal{G}_α denotes the region of agent α , the set of vertices that are part of α 's patrolling cycle. Patrolling cycles repeat the exact same route as long as the region remains unchanged, hence the *idle time* of any vertex v of the region is also the region's cover time $\varphi_0(v) - \varphi_1(v) = \Delta t_\alpha$. Thus the pheromone markings on the vertex can be used to calculate the size of the region $|\mathcal{G}_\alpha| = \frac{\Delta t_\alpha + 1}{2}$.

Conquest

To expand their region, agents may conquer vertices adjacent to (vertices of) their region. For agent α to attempt to launch a conquest from a vertex $u \in \mathcal{G}_\alpha$ to a target vertex v , the following conditions must apply:

1. v is not part of α 's region, $v \notin \mathcal{G}_\alpha$, let us then assume that $v \in \mathcal{G}_\beta$ of another agent β
2. u is subject to a *double visit* by α , i.e. α visits u leaving pheromone marks *twice*, while v was not visited even once by β during the same period of time. Since the time difference between the two visits by α is the *cover time* and the cover time is proportional to the size of the region, it means that α 's region is smaller than β 's, $|\mathcal{G}_\alpha| < |\mathcal{G}_\beta|$. An agent may check for this condition by evaluating if $\varphi_1(u) > \varphi_0(v)$.

3. If the double visit condition is met, thus β 's region is larger, allow a conquest attempt if it is not larger by *exactly one* vertex – since a difference of one vertex is considered balanced.
4. If the double visit condition is met and β 's region is larger by *exactly one* vertex, allow a conquest attempt if vertex v is stagnated – it's pheromones are older than their purported cover time. An agent checks this by comparing the idle time $t - \varphi_0(v)$ to the cover time $\varphi_0(v) - \varphi_1(v)$.

Depending on the above conditions, an agent may stochastically *attempt* the conquest of vertex v , with a predefined probability $0 < \rho_c < 1$. This mechanism works even if v is not part of any of the other agent's regions, $v \notin \mathcal{G}_i, \forall i$. In such case the pheromone marks on v will never be refreshed and the conquest conditions hold.

Temporary Inconsistency

When a region expands or shrinks as result of conquests, it becomes *inconsistent* in the sense that the size of the region changed, but at least some of the pheromone marks on its vertices encoding the *cover time* ($\varphi_0(v) - \varphi_1(v)$) do not reflect that immediately. To regain consistency on a vertex u , the pheromone marks on u must be *refreshed*, hence an agent must leave there a fresh pheromone, and that may occur only when the agent advances into vertex u through a pair trail. Therefore, there is a delay in the propagation of the change, thus there will be a temporary inconsistency between the actual size of the region and the cover time encoded on region's vertices. That inconsistency is certainly not desirable since it might result in a miscalculation of conquest conditions. Consider an agent α with a larger region than two of its neighbors β and γ . Both neighbors will be attempting to conquer vertices from α . Since all agents' awareness is local, β and γ have no means to know that α is shrinking due to the work of the other as well, and as β and γ repeatedly conquer vertices from α , the combined conquests may accumulate to “eat up” too much out of α 's region up to a point where the imbalance is reversed, and the areas of both β and γ are now larger than α 's. Nonetheless, the inconsistency is temporary. It is convenient to analyze this issue by considering the spanning tree of pair trails. When an agent conquers a vertex and expands its region, it results in adding a leaf to the spanned tree, and losing a vertex to another agent results in the *pruning* of the tree, the splitting of the tree into two or more branches, while the losing agent remains on one of them. In either case, the *follow the pair trails* strategy ensures that the new route remains well defined. It is therefore sufficient for an agent to patrol its region twice, to ensure that the region is consistent, as described in the following

Lemma 1. *A region is consistent if it has remained unchanged for a period of time which is twice its cover time.*

Proof. On the first cycle, the agent leaves a fresh pheromone, φ_0 , on each vertex, while the previous most recently visit, φ_1 , may reflect an inconsistent state. The second cycle repeats the exact same route as the first, since the region remains unchanged, and now both the most recent visit as well as the one preceding it, indicated by pheromone levels φ_0 and φ_1 , are updated, thus $\varphi_0 - \varphi_1$ reflects the cover time on all vertices of the route and the region becomes consistent. \square

Losing a Vertex

When balloons are inflated in a box, to the observer it looks as a smooth evolution where the balloons steadily grow and occupy more of the volume until the box is filled. But unlike gas inside a balloon that exerts pressure in all directions concurrently, our discrete agents work in steps, where at each step they attend one vertex of their region, while the other vertices may be subject to conquests by other agents. Since regions are defined by patrolling routes, an agent α , by conquest of a single vertex from β , may prune β 's region in a way that leaves β to patrol a much smaller region effectively rendering it smaller than α 's. Now the “balance tilts”, as the region that was larger prior to the conquest becomes the smaller. Pruning may cut a spanning tree into two or more sections, but in many cases the sections of the tree may still be *connected* by edges that are *not* marked by a pair trail. In such case, β has an opportunity to mark a pair trail over such an edge and regain access to a branch still marked with its own pheromones. Yet, sometimes the pruning divides the region into two unconnected components. We call these *balloon explosions*, and when these occur it is more difficult for the agent that lost part of its region to regain its loss. Therefore, when an agent launches

a conquest attempt it is not always clear if its success will advance or set-back the evolution towards convergence. It is then natural to add the following *vertex loss* rule: should an agent fail the conquest attempt, there is a predefined probability $0 < \rho_l < 1$ for *losing* the vertex from which the attempt was launched. *Losing* the vertex may indeed be a better evolution step than succeeding in that conquest, resembling actions of withdrawal from local minima used in *simulated annealing*. In fact, this property becomes instrumental in our convergence proof for the AntPaP algorithm. In order to prevent an agent from “cutting the branch it is sitting on”, we limit vertex loss events to steps of the patrolling process in which agents backtrack, and, symmetrically restrict conquests to steps in which the agents advance.

We next list the algorithm describing the work of each agent on the graph environment.

Rule AntPaP

Entry point of an agent α at time step t .

Upon entry, the agent is located on vertex u .

Pseudo Code

```

for each  $\exists v \in \mathcal{N}(u) \mid (\text{Self}(v) = \text{false}) \wedge (\varphi_1(u) > \varphi_0(v))$ 
  if  $(\varphi_0(u) - \varphi_1(u)) + 2 \neq (\varphi_0(v) - \varphi_1(v))$ 
    or  $(t - \varphi_1(v)) > (\varphi_0(v) - \varphi_1(v))$  then
      if  $\nexists v \in \mathcal{N}(u) \mid \varphi(u, v) = \varphi_0(u) + 1$  then
        if  $(\text{AgentPresent}(v) = \text{false}) \wedge (x < \rho_c)$  then
           $\varphi(u, v) = t; \varphi(v, u) = t; \varphi_0(v) = t$ 
           $\varphi_1(v) = 0$ 
          goto  $v$ ; return
           $\text{lose} = (y < \rho_l)$ 
        else  $\text{lose} = (y < \rho_l)$ 
      if  $\exists v = \underset{v \in \mathcal{N}(u)}{\text{argmin}} \{ \varphi_0(v), \text{s.t. } \varphi_1(u) > \varphi_0(v), \text{Self}(v) = \text{true} \}$ 
        then
           $\varphi(u, v) = t; \varphi(v, u) = t; \varphi_0(v) = t; \varphi_1(v) = 0$ 
          goto  $v$ ; return
      if  $\exists v = \underset{v \in \mathcal{N}(u)}{\text{argmin}} \{ \varphi(u, v), \text{s.t. } \varphi(u, v) \text{ in a pair trail} \}$  then
        if  $(\varphi(u, v) = \varphi(v, u) = \varphi_0(v)) \wedge (\varphi_0(u) > \varphi_0(v))$  then
           $\varphi_1(v) = \varphi_0(v)$ 
           $\varphi(u, v) = t; \varphi(v, u) = t; \varphi_0(v) = t$ 
          goto  $v$ ; return
        if  $(\varphi(u, v) = \varphi(v, u) = \varphi(u))$  then
          if  $\text{lose}$  then
             $\varphi_0(u), \varphi_1(u) = 0, \forall w \in \mathcal{N}(u), \varphi(u, w), \varphi(w, u) = 0$ 
            goto  $v$ ; return
   $\varphi_1(v) = \varphi_0(v); \varphi_0(u) = t$ 
return

```

Description

Explore border

for each neighbor v of u , not marked by α 's pheromone, and meeting the *double hit*.

if size difference $\neq 1$,

or v is stagnated (pruned or empty) then,

if not backtracked to u

if chance allows conquest (x random $\in [0, 1]$)

Conquer vertex v

mark pair-trail into v

move to v and exit this step

if no conquest, set *lose* flag (y random $\in [0, 1]$)

Rejoin Isolated

if v has α 's (self) mark but *double visit* is met, rejoin v . Move to v and exit this step

select v of the oldest pair-trail

Advance

if pair-trail points into v , and mark on v is older

keep previous time-mark

refresh the pair-trail pointing to v

move to v and exit this step

Backtrack

if pair-trail points into u

Lose

if *lose* flag is set

remove time marks of u and its pair-trails

move to v and exit this step

Reset

refresh vertex u

exit this step

4 Typical Evolution of Patrolling and Partitioning

During a patrolling cycle, an agent attempts conquests over all the border edges of all vertices of its region. Hence its region may expand with additional vertices bordering its route. This causes the spanning tree defined by its DFS-route to have an ever growing number of branches as the patrolling cycles continue, resulting in a tree shape resembling a “snow flake” .

Additionally, agents have a strong tendency to form “rounded” regions, if the environment and other agents’ regions allow it. This happens because vertices that are candidates for conquest and are adjacent to more than one of the region’s vertices, have a higher probability to be conquered and incorporated into the patrolled region.

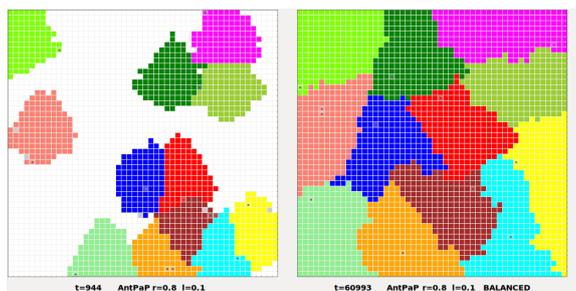


Figure 3: Rounded and Thick Regions build-up by 12 ants on a 50x50 grid

These two properties, the snow-flake like spanning tree and rounded build-ups, often assist in achieving a smooth evolution towards convergence. Thick regions make the possibility of major “balloon explosions” unlikely. The thinner branches at the ends of the snow-flake-like region, cause the pruning of the region by another agent to merely “shave” off small fragments from the region, and are less likely to cut out a large portion of its vertices. Moreover, if a greater portion was cut out by pruning, it is highly likely that the lost portion is connected by edges that were not pair trails, making it is easier for the losing agent to regain its vertices.

Figure 4 depicts a chart showing a typically observed evolution of 12 regions in a 50x50 square grid. The chart describes how the sizes of the regions change with time. The images of Figure 3 are two snapshots taken during the same evolution, the second snapshot being of the balanced partition that the agents reached. The chart shows that some regions grew faster than others, then at some point, the regions grew enough so that the graph was covered (or close to being covered due to continuous pruning), and the smaller regions began to grow at the expense of the larger regions until they all reach a very similar size, albeit not yet balanced. For practical purposes arriving to this state in the graph might suffice, especially if the graph is constantly changing and a partition that is balanced is not well defined. This process occurs quickly, then there is a much longer phase towards convergence. The “turbulence” seen at about $t = 30,000$ are due to some mild “balloon explosions” followed by recovery and convergence. The last *plateau* towards convergence is short in this particular example, but in some other simulated examples it appeared much longer. This is especially noticeable in cases which the evolution results in the system having two large regions that are adjacent and have very similar sizes but are not yet balanced, e.g. a size difference of 2 vertices. Such scenarios may increase the time required for a double visit, which is a condition for conquest.

Figure 5 is a chart depicting time to convergence for a system with 5 agents as a function of the graph size, overlaying results of multiple simulation runs. It shows that the spread in the time to convergence grows with size, but it also clear that the majority, depicted as dense occurrences, are not highly spread in value, indicative of runs of “typical evolution scenarios” on square grid, as was the environment tested in these simulations. The dotted line is an interpolation of average convergence time among the results achieved for each graph size.

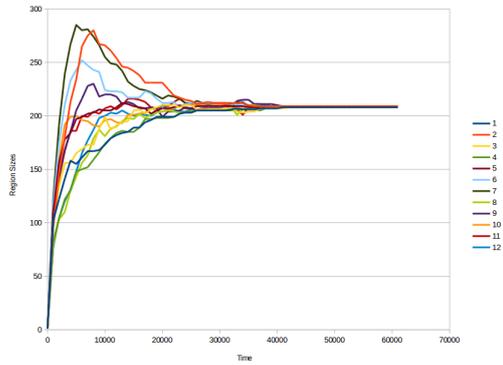


Figure 4: A Typically Observed Partition Evolution , 12 agents on a 50x50 grid

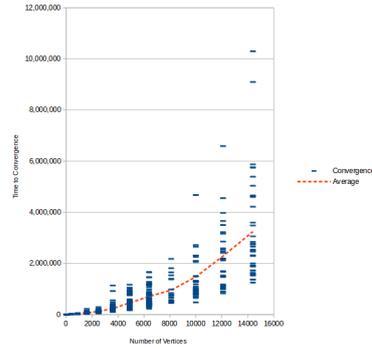


Figure 5: Time to Convergence as a Function of Graph Size

5 Proof of Convergence

The above discussed experimental evidence showcases an evolution of the system towards a balanced partition (when the topology of the environment/graph remains stationary), an evolution which is smooth without “dramatic” incidents, driven by the AntPaP algorithm mimicking pressure equalization. However, AntPaP is a heuristic process, and the experimentally observed smooth convergence is by no means guaranteed. In the evolution towards balanced partitions there are various events that may substantially alter the size difference between regions, and lead the system to longer and chaotic excursions. Chance dictates the way regions expand and, for example, a region may build up with less thickness in some areas allowing other agents to cut across it causing “major balloon explosions”. Furthermore, even a quite well-rounded region may be subject to an “unfair” probabilistic attack, driven to cut through its width and eventually succeeding to remove a large portion of its area. To make things even worse, the portion of the split region that ceased to be patrolled, becomes a readily available prey to neighboring agents. Therefore, although the system is relentlessly progressing towards a balanced partition due to the rules of “pressure equalization”, such “balloon explosions” are singular events that may significantly derail the smooth evolution towards convergence, slowing the process considerably. One may wonder if there are conditions where these events occur repeatedly, making the convergence into a balanced partition an elusive target, that may even never be reached.

Clearly, there are systems where a balanced partition cannot be reached, simply because one does not exist. An evident example is a graph having the shape of a star. Consider a graph of 7 vertices, one at the center and three branches of two vertices each. A system with 7 vertices and 2 agents should be partitioned into two connected components, one of 3 vertices and the other of 4. But such partition does not exist, thus repeated balloon explosions will forever occur. Interestingly, a balanced partition does exist for the same graph with 3 agents.

Therefore, our first step towards proof is to precisely define the systems of interest, which are based on environment-graphs for which a balanced partitions exist for *any* number of agents (up to a bound). The most general set of such graphs is an interesting question in itself. For our purposes, we shall limit our analysis to systems of the following type:

1. The environment is a graph \mathcal{G} that has a Hamiltonian path, a path that passes through every vertex in a graph exactly once, with any number of agents $n \leq |\mathcal{G}|$ patrolling it. Indeed, for any graph \mathcal{G} that has a Hamiltonian path, we can find multiple possibilities for a balanced partition for any number of agents $n \leq |\mathcal{G}|$. To name one, the partition where each region includes vertices that are all adjacent one to another along the Hamiltonian path, and the regions are chained one after another along the path. Some of the regions can be of size $\lfloor \frac{|\mathcal{G}|}{n} \rfloor$ and the others of size $\lfloor \frac{|\mathcal{G}|}{n} \rfloor + 1$. Note that for our purposes, the path does not need to be closed, so the existence of a *Hamiltonian cycle* is not required.
2. The environment is a *k-connected* graph, a graph that stays connected when any fewer than k of

its vertices are removed, with $n \leq k$ agents patrolling it. In [14], Gyori shows that a k -connected graph can always be partitioned into k components, including k different and arbitrarily selected vertices.

We shall analyze the evolution of the system as a stochastic process, and base our proof of convergence on the theory of Markov chains. The remainder of this section is organized as follows: we define a “system configuration” by considering a simple evolution example and show that there always exists a mapping from configurations to well defined “states”. We then look at more complex configurations and realize that although the set of configurations is not bounded, it can be divided into a finite set of equivalence classes, each class representing a state. Hence, we conclude that the number of states in the Markov chain is finite, and the evolution of the configurations maps into corresponding transitions between the states of the chain. Next we use the concept of consistency of a region, as presented in Lemma 1, to conclude that if a balanced partition is attained, it may persist indefinitely. This means that balanced partitions map to recurrent states in the Markov chain. We use this result to analyze the structure of the stochastic matrix that describes the chain. Then we turn to prove that it is only the balanced partitions that are mapped to recurrent states. We first abstract the complexity of the problem by classifying all possible graph partitions into mutually exclusive classes: uncovered, covered but unbalanced, balanced but unstable, balanced and stable. Then we proceed to analyze the changes that may cause the system to shift from a configuration in one class to a configuration in another. Finally we show that when the graph has a Hamiltonian path or is k -connected, despite the possibility that the system may repeatedly transition between these classes, it cannot do so indefinitely and will inevitably have to sink into a recurrent state that belongs to a set of states which are all assigned to the same balanced partition, forming a so-called “recurrent class”.

Definition 1. *The vertices and edges of the graph with their respective pheromone markings, and the agent locations, will be called a configuration of the system, and denoted by \mathcal{C} . Recall that, as discussed in section 3 (sub-section “Agents”), more straightforward time markings, (in which the current time, $\simeq_0(v) = t$, is marked), can be used to emulate their equivalent pheromone markings with temporal decay.*

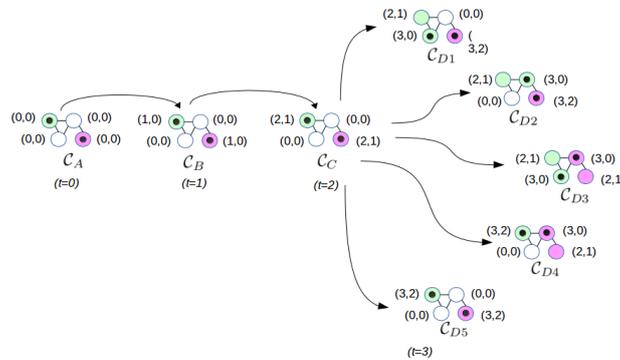


Figure 6: Starting at configuration \mathcal{C}_A , the system must move to \mathcal{C}_B and then to \mathcal{C}_C . Then it will stochastically transition to \mathcal{C}_{D1} , \mathcal{C}_{D2} , \mathcal{C}_{D3} , \mathcal{C}_{D4} or \mathcal{C}_{D5} . \mathcal{C}_C and \mathcal{C}_{D5} are equivalent since both have the same pheromone decay levels. The values shown are time markings (φ_0, φ_1) .

The diagram of Figure 6 depicts an example of transitions between configurations of a system of a 4-vertex graph on which two agents, the *green* and the *cyan*, are active. \mathcal{C}_A is the initial configuration at $t = 0$. Agents, shown as dots, are placed at some random initial vertices that are colored according to the agent patrolling them, *green* at the top-left vertex and the *cyan* at the bottom-right vertex. The pheromone markings on vertices are shown as ordered pairs of time markings (φ_0, φ_1) , where φ_0 is the most recent time that a vertex was marked with pheromone, and φ_1 is the previous time that the vertex was marked (so, generally, $\varphi_0 > \varphi_1$). When the two agents *wake up* at time slot $t = 1$, the readings of pheromone marks around are all zero, therefore the *double visit* condition $\varphi_1(u) > \varphi_0(v)$, is not met,

and conquests are prohibited. Hence the only possible action for the agents is *reset*, i.e. leaving a fresh pheromone mark $\varphi_0(u) = 1$ (the current time), and thus transitioning to the new configuration \mathcal{C}_B . At $t = 2$, the double visit condition is again not met, and the system transitions to \mathcal{C}_C (recall that according to the AntPaP algorithm, at the time that a pheromone is marked on a vertex, the previous mark is moved from φ_0 to φ_1). At $t = 3$, conquest conditions are met for both agents, and the system may now transition to any one of the configurations $\{\mathcal{C}_{D1}, \mathcal{C}_{D2}, \mathcal{C}_{D3}, \mathcal{C}_{D4}\}$, according to whether one or more conquests succeed, or to \mathcal{C}_{D5} , with a probability of $(1 - \rho_c)^3$, if all 3 conquest attempts fail. It is important to notice that \mathcal{C}_{D5} is equivalent to \mathcal{C}_C (and in fact the configurations are identical in terms of temporally decaying pheromone markings since $(3, 2)_{\text{at } t=3} \equiv (2, 1)_{\text{at } t=2} \equiv (t, t-1)_{\text{at } t}$). In both we have a φ_0 pheromone that has been freshly marked, so $(t - \varphi_0) = 0$, and a φ_1 which has been marked on the immediately preceding time slot, hence $(t - \varphi_1) = 1$. We can, therefore, map each configuration of Figure 6 to distinct “states”, states 1, 2, . . . , 7, as depicted in Figure 7, and group \mathcal{C}_{D5} and \mathcal{C}_C into an *equivalence class* of system configurations, were configurations in such equivalence class map to the same “state” (in this case the equivalence class that includes \mathcal{C}_{D5} and \mathcal{C}_C maps to state 3).

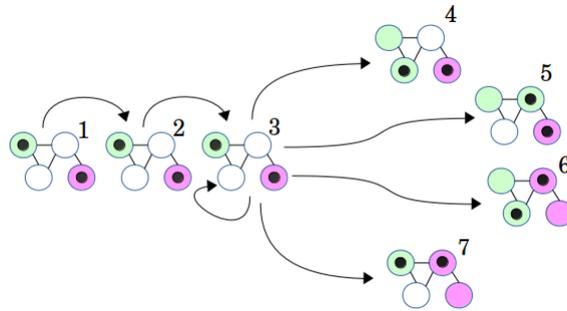


Figure 7: States Matching the Configurations of Figure 6

Calculating transition probabilities between “states” is straightforward, but sometimes subtleties arise, for example:

$$\mathcal{P}(\text{State 1} \rightarrow \text{State 2}) = \mathcal{P}(\text{State 2} \rightarrow \text{State 3}) = 1,$$

$$\mathcal{P}(\text{State 3} \rightarrow \text{State 3}) = (1 - \rho_c)^3 \text{ (both agents attempt conquests, but fail), and}$$

$\mathcal{P}(\text{State 3} \rightarrow \text{State 7}) = \frac{1}{2}(\rho_c)(1 - \rho_c) + \frac{1}{2}(1 - \rho_c)^2\rho_c$ (due to the strong asynchronous assumption, the *cyan* agent has a probability of $\frac{1}{2}$ to move first within the time slot, and if so, its attempted conquest has a probability ρ_c to succeed. The result is multiplied by a probability $(1 - \rho_c)$ that the *green* agent fails in its conquest. If the *green* moves first, there is a probability $(1 - \rho_c)^2$ that it would fail both attempts, and then a probability ρ_c that the *cyan* succeeds).

Definition 2. An edge (u, v) which is part of a pair-trail pattern marking, so that $\varphi(u, v) = \varphi(v, u) = \varphi(v)$, will be called a *pair-trail edge*.

Definition 3. For a time interval in which no conquests or losses of a vertex in a region occur, the region is considered “stable”.

Lemma 2. A system comprising a graph \mathcal{G} with time-invariant topology, and k agents, in a configuration \mathcal{C}_N , satisfying

1. The k regions marked by the agents are consistent
2. Pheromone marks exist only on vertices and pair-trail edges inside the k -regions, and no pheromone markings exist elsewhere on the graph,

will transition through a finite sequence of M of states, where M is the least common multiple of the cover times of the k -regions (i.e. $M = \text{lcm}(\Delta t_1, \Delta t_2, \dots, \Delta t_k)$), prior arriving to a configuration \mathcal{C}_M equivalent to \mathcal{C}_N (i.e. $\mathcal{C}_M \sim \mathcal{C}_N$), as long as the k regions are stable (no conquests or losses occur).

Proof. A consistent region \mathcal{G}_i is a region for which $\varphi_0(v) - \varphi_1(v) = \Delta t_i, \forall v \in \mathcal{G}_i$, where Δt_i is the cover time of region \mathcal{G}_i (see Lemma 1). Therefore, as long as the region is stable, all vertices and all pair-trail edges of \mathcal{G}_i , cyclically return to the exact same decaying pheromone levels, i.e. exactly the same temporal differences $t - \varphi_0(v)$ (where t is the current time) every Δt_i steps. We can thereby consider the k regions in the partition, each repeating its pheromone level markings independently, as k cyclic processes each with its own cycle time. Hence, all the processes complete an integer number of cycles every M steps, where M is the least common multiple of the cycle times, $M = lcm(\Delta t_1, \Delta t_2, \dots, \Delta t_k)$, which ensures that all vertices of all the k regions in the configuration \mathcal{C}_M that was reached have exactly the same temporal differences as in \mathcal{C}_N , and therefore $\mathcal{C}_M \sim \mathcal{C}_N$. \square

In the above Lemma, we required to have no pheromones at all on edges that are not pair-trail edges. But, if there were markings on such edges, the patrolling agents would simply ignore them, according to the AntPaP algorithm. Therefore such markings have no influence on the possible future evolutions of the system. We shall formally define states of the system by grouping together configurations that have “the same future evolutions”, i.e. same possible future configuration transition sequences with the same probabilities. For example, as seen above, configurations that differ only by levels of pheromones on non pair-trail edges form such equivalence classes, hence each class defines a distinct state. In systems theory, this is the classical *Nerode* equivalence way of defining states.

Accordingly, two configurations that do not have the exact same patrolling routes (either not having the same k regions, or the agents have developed different patrolling paths within the regions) cannot have the same future evolutions, since, even without any conquests or loses, the future sequences of configurations that the systems go through are different due to the different patrolling steps. Therefore these two configurations can not belong to the same equivalence class thus represent distinct states.

Next we turn to discuss pheromone markings that may exist on vertices and edges that are not part of any current patrolling route, hence outside of all the regions. Such scenarios may occur as result of a successful conquest by an agent that disconnects the region of a neighbor and hence prunes the spanning tree of that agent, splitting it into two or more disjoint branches. Clearly, the latter agent remains on one of these branches, while the others cease to be part of its patrolling route and remain “isolated”.

Definition 4. *A segment of a spanning tree (i.e. a set of vertices marked with pheromones and connected by pair trails) that is not part of a patrolling route, thus not included in any of the regions, forms what we shall call an isolated branch. For completeness, a single such vertex that is not connected by a pair-trail is also considered an isolated branch.*

In our pheromone marking model we have not limited the pheromone decay, thus, on an isolated branch, pheromones may decay indefinitely. This means that there is no bound to the set of configurations, and raises the question of whether there exists a bound to the set of equivalence classes to which they can belong, hence a bound on the number of states of the system. We shall, therefore, consider configurations that include isolated branches, and analyze the effect of pheromone decay in these branches on the evolution of the system, or more precisely, how such decay influences “future” system states. We have already seen that two configurations that do not have the exact same patrolling routes must represent different states, thus we shall verify that this distinction, by itself, does not produce an unbounded number of states. The number of permutations of possible k stable regions is finite (in a finite graph), and for each such permutation, the number of permutations of possible routes for the k agents must be finite too (since each of the k regions have a finite number of edges).

We are therefore left to show that starting at any arbitrary configuration with k agents patrolling k regions that also include isolated branches, all future evolutions in an arbitrarily large interval in which all k regions remain stable, can be grouped into a finite number of equivalence classes.

Let us consider a setup of k regions and an isolated branch, where the regions are stable in an arbitrarily large interval, and further assume that a vertex v of the isolated branch is adjacent to a vertex u in one of the regions (See figure 8). Since the regions are stable, every patrolling cycle the value $\varphi_0(u)$ is refreshed, thus its time-marking increases with each cycle. On the other hand, the time marking $\varphi_0(v)$ of the vertex in the branch remains unchanged. When the agent is on vertex u the following scenarios may arise:

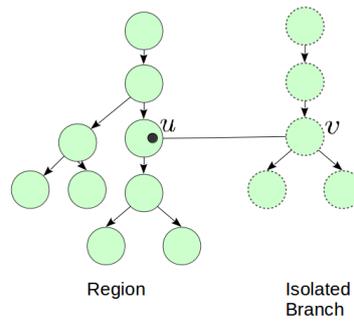


Figure 8: An agent (black dot) on a vertex u adjacent to a vertex v on an isolated branch that is marked with the agent’s pheromone. The solid line represents an edge, the arrows represent pair-trail edges. Rejoining the lost vertex v on the isolated branch results in pruning the isolated branch into disjoint parts. One with the upper two vertices, another with the lower three.

1. The vertex v on the branch might be marked with the same agent’s pheromone, and hence moving into the adjacent vertex consists of the action of *rejoining* a vertex previously lost. According to AntPaP, agents check for a double visit condition, i.e. $\varphi_1(u) > \varphi_0(v)$, prior to this action. When traversing onto the vertex, e.g. at time t , the agent marks there a fresh pheromone $\varphi_0(v) = t$. This may result in splitting the isolated branch into two or more disjoint branches. The agent will then follow the pair-trails emanating from vertex v at which it is presently located, oblivious to the fact that pheromone marks on pair-trails and vertices of the branch are old. Thereafter the agent traverses the section of the previously-isolated branch it is located on, thus refreshing its marks, until all the section is visited (e.g., in the example of Figure 8, the agent will visit all vertices on the lower section), then it returns to the vertex u from which the conquest was launched. The other disjoint branches (in the example of Figure 8, the upper section) remain isolated.
2. if the branch is marked with another agent’s pheromone, and conquest conditions are met (e.g. double visit ($\varphi_1(u) > \varphi_0(v)$) and the regions size difference is not exactly one vertex (or equivalently the difference in cover time is not 2 , i.e., $(\varphi_0(u) - \varphi_1(u)) + 2 \neq (\varphi_0(v) - \varphi_1(v))$), the agent may attempt a conquest on the vertex v and thereafter on all the vertices of the branch, one by one.
3. if the branch is marked with another agent’s pheromone but the double visit condition is not met (i.e. $\varphi_1(u) \not> \varphi_0(v)$), it may remain so only for at most two cycles of patrolling. Note that $\varphi_1(u)$ is growing with each agent’s visit, while $\varphi_0(v)$ remains unchanged and as a result a double visit condition will necessarily arise. Furthermore, meeting the double visit condition also ensures that all the additional conquest conditions are met at the same time, since either the cover time encoded in the vertex of the isolated branch indicates a region size conducive to conquests, or the agent recognizes that the neighboring region is stagnated (not being patrolled for too long, i.e., $(t - \varphi_1(v)) > (\varphi_0(v) - \varphi_1(v))$). Therefore any further decay of the pheromone mark on the isolated branch will not influence the future behavior of the system.

A double visit is, therefore, a sufficient condition for an agent to conquer or rejoin a vertex on an adjacent isolated branch, hence we conclude the following:

Lemma 3. *A system comprising a graph \mathcal{G} with time-invariant topology, and k agents, in a configuration \mathcal{C} that includes exactly one isolated branch will transition through a finite sequence of at most $N + M$ states, where $N = 2 * \max(\Delta t_1, \Delta t_2, \dots, \Delta t_k)$ and $M = \text{lcm}(\Delta t_1, \Delta t_2, \dots, \Delta t_k)$, where Δt_i is the cover time of region \mathcal{G}_i , as long as all the regions are stable (no conquests or losses occur).*

Proof. The completion of two patrolling cycles of a region by its patrolling agent ensures that the double visit condition is met at any vertex u of the region adjacent to a vertex v of the isolated branch (see discussion above). Therefore, after an interval of $N = 2 * \max(\Delta t_1, \Delta t_2, \dots, \Delta t_k)$ (i.e. when the

agent on the largest region completed two patrolling cycles) it is certain that the double visit condition is globally met (i.e. for any vertex u on any of the regions adjacent to any vertex v in the isolated branch). Moreover, it will be met on any time step that follows (as long as the regions are stable). Hence any two configurations on which the double visit condition is globally met, and have the same levels of pheromones on vertices and pair-trail edges that are in the k regions (but may differ in levels of pheromones on the isolated branch) are equivalent.

Since we also know, based on Lemma 2, that every $M = lcm(\Delta t_1, \Delta t_2, \dots, \Delta t_k)$ time steps, all pheromones in vertices and pair-trails included in the k regions cyclically return to the exact same decaying pheromone levels (i.e. exactly the same temporal differences), we conclude that a system in a configuration \mathcal{C} , will transition at most N distinct states to a configuration \mathcal{C}_N (on which the double visit condition is globally met) and then will cyclically transition through M states reaching, at each cycle, a configuration $\mathcal{C}_M \sim \mathcal{C}_N$. \square

Our next analysis is of the effect of multiple isolated branches on future evolutions of a system with stable regions. Consider two scenarios, both starting with the same configuration that has one isolated branch. An arbitrary time later, a conquest creates another isolated branch, the second branch being the same in both, only the *time* of its creation is different. Hence, there can be an arbitrarily large time difference between the creation of the second isolated branch in the two scenarios. Contemplating the case of an arbitrary number of isolated branches created each at an arbitrary time, the complexity of such presented scenarios may substantially grow. Nevertheless, in term of system states the above complexity does not matter. Once the decay of pheromones on an isolated branch is such that the double visit condition is globally met, the conquest or rejoin threshold is triggered, and afterwards no amount of further decay affects the future evolutions of the system. This insensitivity holds regardless of the presence of other isolated branches, simply because the double visit is a local condition, limited to the time difference encoded in pheromones on a vertex in a region and an adjacent vertex on the branch. Thus, any two configurations that differ only by level of pheromones on isolated branches for which the double visit condition is globally met, are equivalent. Particularly, there must exist a configuration such that the level of pheromones on isolated branches is at its “highest level”, i.e. the time marking on each vertex of each isolated branch is the highest that allows the double visit condition to be globally met. A branch with such “highest level” will have one vertex v , where $v = argmax(\varphi_0(w), w \in \text{isolated branch})$ with a time-mark value $\varphi_0(v) = t - (M + N)$ where t is the current time (i.e. a pheromone was left there $M + N$ steps before the current time), and all other vertices and edges with (lower) values that agree with the ordered directions of pair-trails.

Hence we conclude, again, that any configuration \mathcal{C} that includes multiple isolated branches will transition at most $M + N$ distinct states as long as all the regions remain stable.

Theorem 4. *For a system with a graph \mathcal{G} of stationary topology, and k agents, the set \mathcal{S} of states is finite.*

Proof. Based on the above analysis, we conclude:

1. any configuration \mathcal{C} is equivalent to a configuration \mathcal{C}_P identical to \mathcal{C} except for having no pheromone marking on edges that are not pair trails.
2. any configuration \mathcal{C}_P that includes isolated branches is equivalent to a configuration \mathcal{C}_A identical to \mathcal{C}_P except by the levels of pheromones on vertices and pair-trail edges in those isolated branches that globally meet the double visit condition on both configurations. Specifically, in \mathcal{C}_A , vertices and pair-trail edges on each such isolated branch, will be of a “highest level”, i.e. will have one vertex v , where $v = argmax(\varphi_0(w), w \in \text{isolated branch})$ with a time-mark value $\varphi_0(v) = t - (M + N)$ where t is the current time (i.e. a pheromone was left there $M + N$ steps before the current time), $N = 2 * max(\Delta t_1, \Delta t_2, \dots, \Delta t_k)$ and $M = lcm(\Delta t_1, \Delta t_2, \dots, \Delta t_k)$, and all other vertices and edges with values that agree to the directions of pair-trails.

Therefore any configuration \mathcal{C} is grouped in an equivalence class with a correspondingly “representative” configuration \mathcal{C}_A . To find out how many such classes exist, we observe that a \mathcal{C}_A includes the following elements: k regions with k patrolling routes, isolated branches that do not globally meet the

double visit condition and isolated branches of a highest level of time-markings that globally meet the double visit condition.

However, we have that:

- (a) The number of possible choices of k regions is finite (in a finite graph).
- (b) For any arbitrary set of k regions, the number of possible routes in the k regions is finite.
- (c) For any arbitrary set of k regions (with a particular selection of k routes) the number of vertices not included in these is finite, thus the number of possible isolated branches is finite (and their possible assignments to whether they meet the double visit condition or not is also finite).

Therefore the set of possible representative configurations \mathcal{C}_A is finite, each defines an equivalence class corresponding to a distinct state of the Markov chain, hence the set \mathcal{S} of system states is finite. \square

Concluding the above analysis we see that in spite of the infinite number of configurations possible for a system, the number of system states, though quite large, is finite. Let us denote the finite set of states of a system by S .

Definition 5. (Gallager [15]), A Markov chain is an integer-time process, $\{X_n, n > 0\}$ for which the sample values for each random variable $X_n, n > 1$ lie in a countable set S and depend on the past only through the most recent random variable X_{n-1} .

Clearly, any state of the system at time n , formally represented by $X_n \in S$, is dependent only on the previous state $X_{n-1} \in S$, since our agents have no memory, and their decisions are based solely on readings from vertices and edges of the configuration, which are completely described by X_{n-1} . We can, therefore, analyze the evolution of the system based on the theory of Markov chains. Our aim is to prove that the Markov chain is not irreducible (i.e. given enough time, the probability to reach some of its states tends to zero), and that all its recurrent states represent balanced partitions. To proceed with our analysis, we notice that the size of set S grows very fast with the size of the graph. Calculations show that even the simple example of Figure 7 develops to a surprisingly large chain. In order to be able to describe the evolution of the system in a simple manner, we also define a *partition* of the environment.

Definition 6. The coloring of each vertex of a configuration \mathcal{C} by its patrolling agent along with the set of unvisited vertices form a partition \mathcal{P} of the graph. Partitions are unconcerned about the levels of pheromones on the vertices and indifferent to agent locations, thus only exhibit the regions of \mathcal{C} .

Many different configurations (and hence states too) correspond to the same partition, therefore we can use the concept of a *partition* as an abstraction referring to all those configurations.



Figure 9: An example partition \mathcal{P} of the graph

Figure 9 is an example of a partition of the environment graph that the system we discussed above arrived to. From our previous discussion we know that it represents a set of states of the underlying Markov chain. One characteristic of that set of states is that it contains a cyclic path. This reflects the fact that agents may cyclically repeat their patrolling route for some period of time during which conquests or losses do not occur, and the partition remains stationary. In fact, having a cyclic path in the underlying Markov chain is characteristic of any reachable partition.

Eventually conquests or losses are stochastically enabled leading to a different partition, and, as a result, to a different set of underlying states. In Figure 10, the system may remain in partition A for a while, as the underlying chain cycles through the relevant states, but eventually it will probabilistically transition to one of the partitions B,C,D,E. Note that the transition from A to D means that both agents conquered one vertex each during the same time-slot.

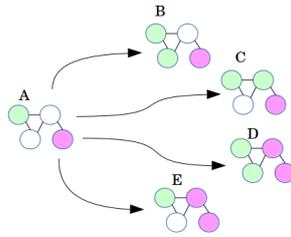


Figure 10: Transitions between partitions (note that these transitions may correspond to many different transitions between possible configurations or states, see figure 7)

A *recurrent class* in a Markov chain is a set of states which are all accessible from each other (possibly passing through other states), and no state outside the set is accessible (Gallager [15]). The following Lemma shows that the set of states corresponding to any balanced partition includes recurrent classes:

Lemma 5. *If a system remains in a balanced partition for a period of time equal to twice the cover time of its largest region, it will remain so indefinitely.*

Proof. We know from Lemma 1, that if a region remains unchanged for a period of time which is twice its cover time, it becomes consistent, so the pheromone levels in all of its vertices correctly indicate its cover time, $\varphi_0 - \varphi_1 = \Delta t$. Therefore if the system remains in a balanced partition for a period twice the largest cover time (the cover time of its largest region), it is guaranteed that all the regions are consistent. Hence, we conclude that no conquest attempts are subsequently possible, since the system is balanced and conquest conditions can not be satisfied across any border edge. \square

The conclusion of Lemma 5 is that a balanced partition with all its regions consistent, must correspond to a *recurrent* (and *periodic*) *class* in the Markov chain. The random process continuously repeats a series of states based on the individual agents’ patrolling cycles. Since agents may reach different patrolling routes for the same region, a balanced partition may correspond to multiple recurring classes. Additionally we can conclude the following:

Lemma 6. *A system may enter a state in which the partition is balanced, and then move into a state in which the partition is not balanced.*

Proof. Clearly, we see while the conditions for consistency are not satisfied for one or more regions in the partition, a conquest or loss may possibly happen, hence the partition may become unbalanced. \square

Since recurrent classes exist, the stochastic transition matrix of the Markov chain of the patrolling system must have the form:

$$\mathbf{M} = \begin{bmatrix} \mathbf{T} & \mathbf{TR} \\ \mathbf{0} & \mathbf{R} \end{bmatrix}$$

where \mathbf{T} is a matrix of transitions between transient states (for example those corresponding to non-balanced partitions), \mathbf{R} is a matrix of transitions between states in recurrent classes (for example those corresponding to balanced partitions with consistent regions), and \mathbf{TR} describes the transitions from transient to recurrent states (for example those corresponding to balanced partitions with inconsistent regions as described in corollary 6). Note that we assume an initial distribution given by a row vector $\bar{\pi}_0$ and hence the future distributions evolve according to $\bar{\pi}_{k+1} = \bar{\pi}_k \mathbf{M}$.

Each recurrent class representing one particular occurrence of k routes in the k regions contributes a section \mathbf{C}_i to \mathbf{R} , of the form of a shifted identity matrix,

$$\mathbf{C}_i = \begin{bmatrix} 0 & \mathbf{1} & 0 & \cdots & 0 \\ \vdots & & \mathbf{1} & 0 & \vdots \\ & \cdots & 0 & \mathbf{1} & 0 \\ 0 & & \cdots & 0 & \mathbf{1} \\ \mathbf{1} & 0 & & \cdots & 0 \end{bmatrix} \qquad \mathbf{P}_j = \begin{bmatrix} \mathbf{C}_1 & 0 & \cdots & 0 \\ 0 & \mathbf{C}_2 & & \\ \vdots & & \mathbf{C}_3 & \vdots \\ & & & \ddots & 0 \\ 0 & \cdots & 0 & & \mathbf{C}_r \end{bmatrix}$$

The rank of \mathbf{C}_i is $M = lcm(\Delta t_1, \Delta t_2, \dots, \Delta t_k)$ (i.e. the least common multiple of the cover times of the k regions), and is a function of the sizes of the regions in the partition. The contribution of a particular balanced partition would be a matrix \mathbf{P}_j that comprises of a set of \mathbf{C}_i matrices on its diagonal, where r here is the finite number of possible route combinations in the regions that form the partition. Therefore our goal is to show that the structure of \mathbf{R} is

$$\mathbf{R} = \begin{bmatrix} \mathbf{P}_1 & 0 & \cdots & 0 \\ 0 & \mathbf{P}_2 & & \\ \vdots & & \mathbf{P}_3 & \vdots \\ & & & \ddots & 0 \\ 0 & \cdots & 0 & & \mathbf{P}_n \end{bmatrix}$$

listing contributions from a finite number n of possible balanced partitions. To achieve this we must prove that recurring classes that are not representing balanced partitions do not exist. This is done next.

To visualize the problem, we will classify all possible partitions as shown in Table 1, recalling that a partition \mathcal{P} represents a set of states in the underlying Markov chain:

By construction of the classifications $\mathcal{CL}1, \mathcal{CL}2, \mathcal{CL}3, \mathcal{CL}4$ it is clear that these are mutually exclusive as well as complete, hence they divide the set of all possible partitions, so that any partition \mathcal{P} can belong to one and only one of the four defined classes. It also means that any state of the system $X \in S$ can be classified to one and only one of the above classes. $\mathcal{CL}4$ includes the set of partitions that comply with Lemma 5, and $\mathcal{CL}3$ includes the partitions of Lemma 6. Our goal is to show that from any initial state $X \in S$ the system will reach some state in $\mathcal{CL}4$.

Figure 11, which we will soon justify, depicts the possible *transitions* between the classes, i.e. from a state in one class, exists a transition sequence in the Markov chain to a state in another class as depicted in the diagram.

We first observe that clearly, any state in $\mathcal{CL}1$ or $\mathcal{CL}2$ must be *transient*, the system cannot remain in any one of them indefinitely. In $\mathcal{CL}1$ the partition includes free vertices (that are not part of any region). For an agent that persist in visiting a neighboring vertex, conquest conditions will be eventually met, and the agent will make repeated attempts to conquer it. Eventually all free vertices will be conquered and the graph becomes covered. Similarly, on $\mathcal{CL}2$, when an agent visits a vertex bordering a larger region there is a chance it will conquer the neighboring vertex. The conquest could cause the pruning of the spanning tree and even a balloon explosion of the larger region, and the system shifts to $\mathcal{CL}1$, or maybe, by chance, it will move to a balanced partition and end up in $\mathcal{CL}3$.

In $\mathcal{CL}3$ the system is balanced but inconsistent. We know from Lemma 5 that it can become consistent, thus shift to $\mathcal{CL}4$, since there is always a chance that no conquests or losses will occur in any finite period of time. If a conquest does occur, it may become unbalanced, and then we are back to $\mathcal{CL}1$ or $\mathcal{CL}2$. So any state in $\mathcal{CL}3$ is transient and may sink to a state in $\mathcal{CL}4$.

We are left to show that from any state in $\mathcal{CL}1$ and $\mathcal{CL}2$ there exists a path to $\mathcal{CL}3$. A sequence of states that switches the system back and forth between $\mathcal{CL}1$ and $\mathcal{CL}2$ is the scenario were the system repeatedly evolves to a covered graph only to retract by events such as a balloon explosion. We will show now that from any configuration there is a strictly positive probability to find a balanced partition, hence move to $\mathcal{CL}3$, and as result there is a strictly positive probability to get into $\mathcal{CL}4$. This means that $\mathcal{CL}4$, the set of balanced partitions, are the only partitions that map to recurrent classes. To do that we recall our restriction to Hamiltonian graphs (or k -connected graphs) and invoke the property of vertex loss.

Lemma 7. *Given a system with an environment graph \mathcal{G} having a Hamiltonian path (or a k -connected*

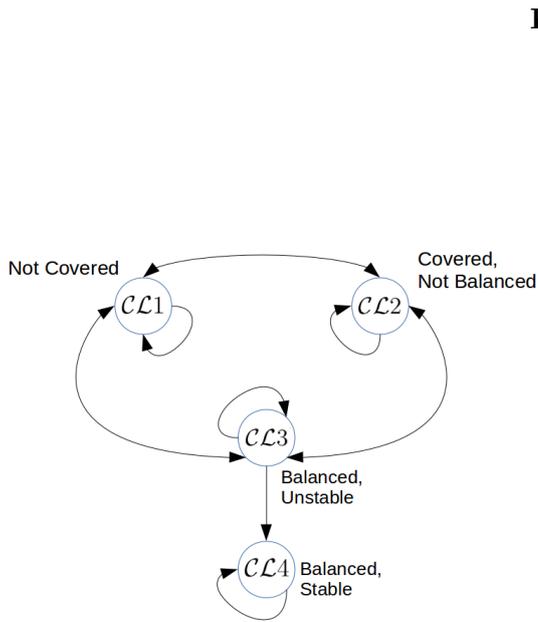


Figure 11:
Class Transitions Diagram

Partition Class	Name	Description
$\mathcal{CL1}$	Not Covered	Any partition \mathcal{P} that includes at least one vertex that is <u>not</u> part of a patrolling cycle.
$\mathcal{CL2}$	Covered, Not Balanced	Any partition $\mathcal{P} \notin \mathcal{CL1}$ that includes two adjacent regions with size difference greater than one.
$\mathcal{CL3}$	Balanced, Unstable	Any partition $\mathcal{P} \notin \{\mathcal{CL1}, \mathcal{CL2}\}$ where one or more inconsistent regions.
$\mathcal{CL4}$	Balanced, Stable (Convergence)	Any partition $\mathcal{P} \notin \{\mathcal{CL1}, \mathcal{CL2}\}$ where all regions are consistent.

Table 1: Classification of Partitions

graph), and n agents, in an arbitrary configuration \mathcal{C}_0 , there is a strictly positive probability for the system to evolve to a balanced partition.

Proof. Let us assume the contrary, that there exists a configuration \mathcal{C}_0 representing a state $s_0 \in \{\mathcal{CL1}, \mathcal{CL2}\}$ such that the probability of any arbitrarily long step sequence starting at \mathcal{C}_0 to arrive into a state in $\mathcal{CL3}$ is strictly zero. We first note that there is a strictly positive probability that the next change of a region is the loss of a vertex. Namely, assume that for a sequence of configurations $\{\mathcal{C}_0, \mathcal{C}_1, \dots, \mathcal{C}_N\}$, where N is finite, at configurations $\{\mathcal{C}_0, \mathcal{C}_1, \dots, \mathcal{C}_{N-1}\}$ there were no changes to any of the regions in the partition of the environment, and at \mathcal{C}_N one or more of the regions lost one vertex each, following failed conquest attempts. This could repeat and consequently, there is a strictly positive (though very small) probability for a sequence of configurations $\{\mathcal{C}_0, \mathcal{C}_1, \dots, \mathcal{C}_M\}$, where M is finite, to arrive to a configuration \mathcal{C}_M where each of the regions is of size 1. Now, we note that there is a strictly positive probability that the next change of a region is a conquest launched from one vertex to a neighbor along a given Hamiltonian path. This could repeat until the regions form one of the possible balanced partitions along the Hamiltonian path and the system is now in $\mathcal{CL3}$. This evolution contradicts our assumption regarding \mathcal{C}_0 , and we therefore conclude that $\{\mathcal{CL1}, \mathcal{CL2}\}$ is a transient set, thus from any arbitrary configuration there is a strictly positive probability to arrive to a balanced partition.

Note that a similar argument can be made for a k -connected environment graph. □

We therefore conclude:

Theorem 8. *A system with a Hamiltonian or k -connected graph \mathcal{G} with stationary topology, and n agents implementing AntPaP converges in finite expected time to a balanced and stable partition with probability 1.*

Proof. The conclusion from Lemma 7, is that any state $s \in \{\mathcal{CL1}, \mathcal{CL2}\}$ is transient. We also know that any state $s \in \mathcal{CL3}$ is transient by Lemma 5. This means that all recurrent classes included in the matrix \mathbf{R} of the stochastic matrix $\mathbf{M} = \begin{bmatrix} \mathbf{T} & \mathbf{TR} \\ \mathbf{0} & \mathbf{R} \end{bmatrix}$ represent balanced and stable partitions, classified as $\mathcal{CL4}$. A Markov chain described by a stochastic matrix of this form, will eventually enter a recurrent state, regardless of the initial state, and the probability that this takes more than t steps

approaches zero geometrically with t (see, for example, Gallager [15]). We conclude that a system with a graph \mathcal{G} (of stationary topology) and n agents implementing AntPaP converges with probability 1 and a finite expected time to a balanced and stable partition. \square

6 Experimental Results and Discussion

We presented and thoroughly analyzed the AntPaP algorithm for continuously patrolling a graph-environment with simple finite state automaton agents (or bots) using “pheromone traces”. The simulations presented so far were on an environment in the shape of a square. On such an environment, we know that many balanced partitions do indeed exist. Practical scenarios are seldom so simplistic. In many important cases, the environment graph is, in fact, uncharted and much more complex in its structure. Still, agents implementing AntPaP will find a balanced partition with probability one (almost surely), if such a partition exists, and will certainly divide their work fairly even when such partitions do not exist.

The shape of the environment considerably affects the time to convergence. The number of balanced partitions that the environment graph has is, naturally, one of the major factors. So is their diversity, i.e. how different the balanced partitions are from each other. If the balanced partitions are similar to one another, the dependency of the time to convergence on the initial locations of the agents tends to be higher than if the solutions are further apart.

Consider the system of Figure 12. The initial positions of 7 agents are shown in the first snapshot at $t = 1$. Next, at $t = 896$, the lower section becomes almost covered. At $t = 5995$, the upper section is almost covered, and the two agents there clearly have larger regions than the agents in the lower section. At $t = 12993$, the cyan agent is trapped in the upper section, and we witness a competition between the agents from the lower section to grow their regions into the prolonged section, that the cyan agent abandoned.

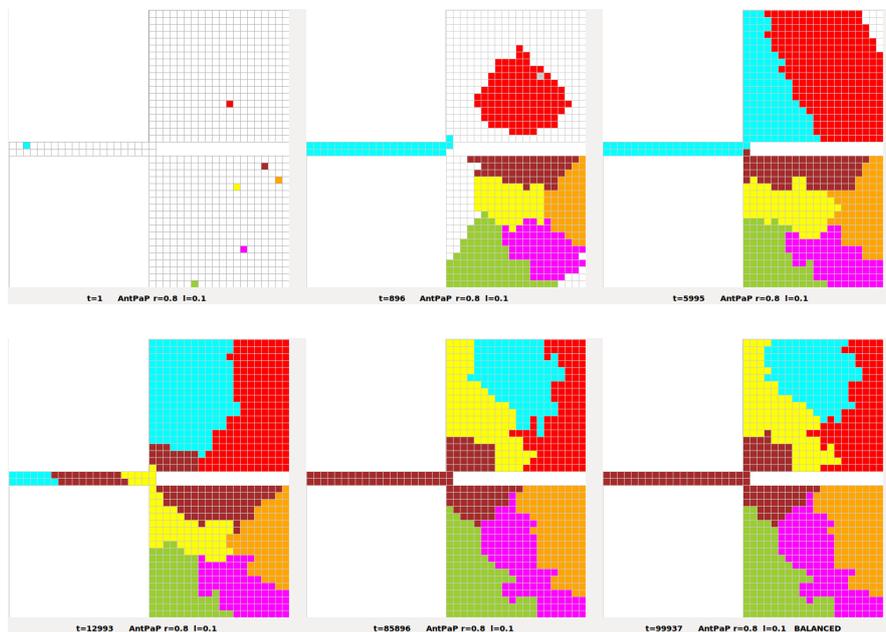


Figure 12: Evolution of a system with 7 a(ge)nts

At $t = 85896$, the competition ends after the yellow agent traversed into the upper section. Now we have 3 agents on each of the upper and lower sections, and one on the prolonged section. Soon after, the system reaches a balanced partition. Clearly, there are many balanced partitions for this system, but all of them have one agent on the prolonged section, and 3 agents on each of the upper and lower sections. Initial conditions with 3 agents on the upper and lower sections each will ensure faster convergence to a balanced partition. Following this experiment and discussion, it is interesting to consider a system with the same environment graph and an even number of agents.

We simulated a system with the same environment graph and 2 agents. This system has only one balanced partition, shown in Figure 13.

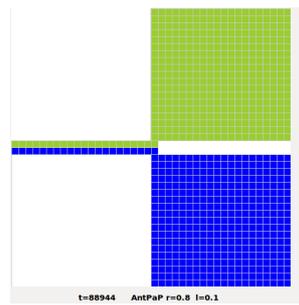


Figure 13: The only balanced partition of a system with 2 agents

Since only one balanced partition exists, it is reasonable to predict that the required time for convergence might be substantial. Figure 14 shows snapshots of an evolution of this system. Both the violet and yellow agents are initially located in the lower section. After a while, the violet agent expands its region so that part of it extends into the upper section. A while later, the violet region covers almost all of the upper section as well as the prolonged section. Then, the yellow agent begins to expand into the prolonged section, eventually causing a “balloon explosion” of violet’s region. Soon enough, the violet agent responds, and causes a balloon explosion of the yellow’s region.

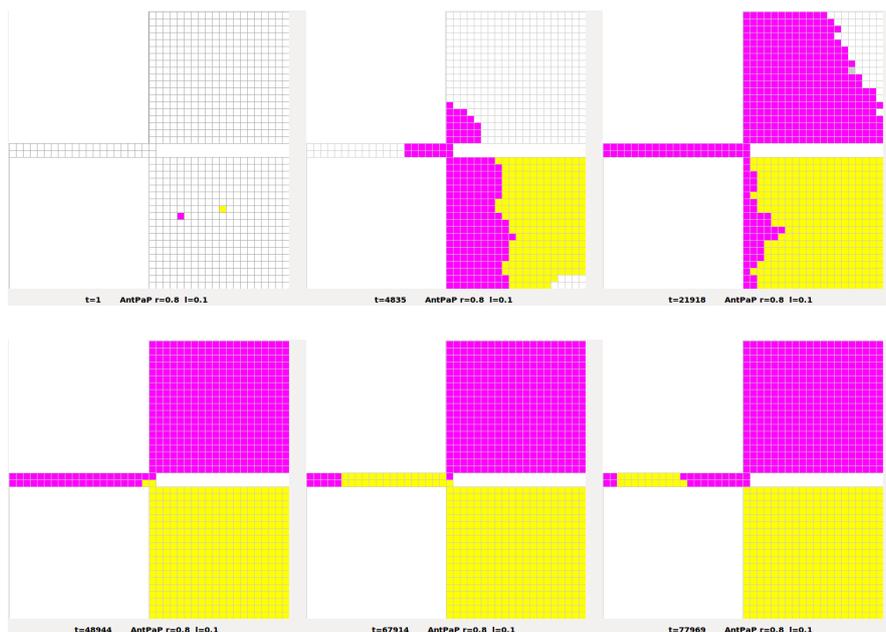


Figure 14: Evolution of a system with 2 a(ge)nts

Due to the shape of the graph, this cycle may repeat over and over again. It will stop only when the single possible balanced partition is reached, and subsequently the regions “lock-in”, and the system remains stable. For that to happen, an agent must conquer the appropriate half of the vertices of the prolonged section. We know that this will eventually happen with probability 1, however the time it will take can be very very long.

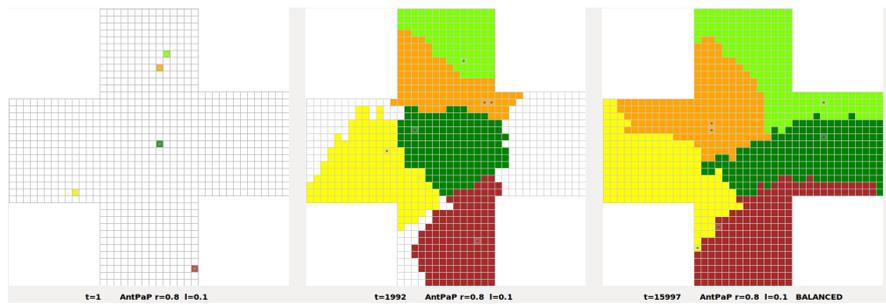


Figure 15: A system with a “Cross” graph and 5 agents

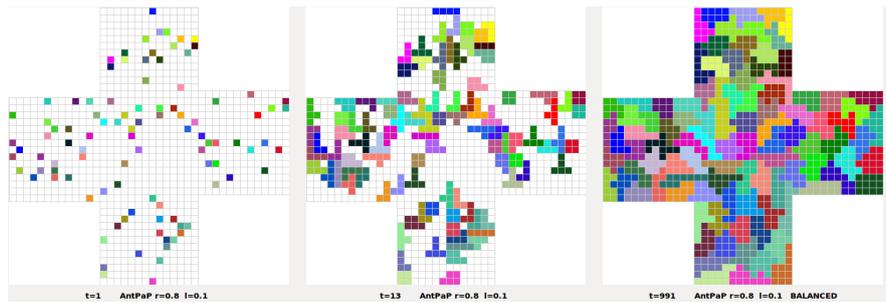


Figure 16: A system with the same “Cross” graph as in Figure 15, and 100 agents

The number of agents is also an important factor of convergence time. Generally, more agents hasten the convergence. Figure 15 shows an evolution of a system with 5 agents on a different environment. We shall call this environment graph the “Cross”. Figure 16 shows a system with the same “Cross” graph and 100 agents. Here, the “pressure” that an agent “feels” from other “balloons” quickly accumulates around its region, and the convergence is swift. Figure 17 depicts results of multiple simulation runs, of systems with the “Cross” graph of figure 15, exhibiting convergence time as a function of the number of agents.

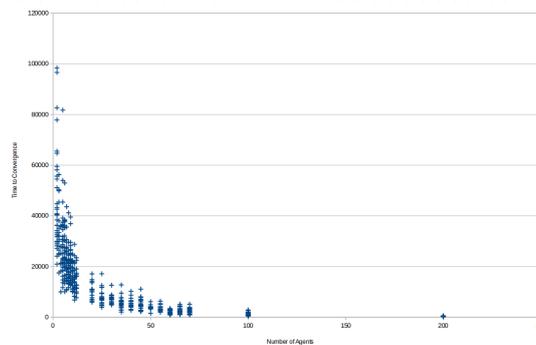


Figure 17: Convergence time of a system with a “Cross” graph as function of the number of agents

In some systems, particular numbers of agents may cause a substantially larger time to convergence. In Figure 18 we present a balanced partition in a graph environment that we call “6 Rooms”.

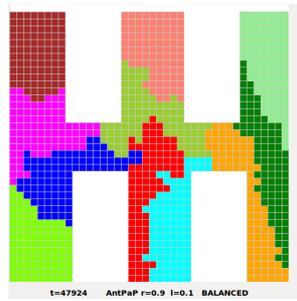


Figure 18: Balanced partition reached on a “6 Rooms” graph

Systems with a “6 Rooms” graph and 6 agents sometimes require a substantially longer convergence time, as shown in Figure 19. Ignoring the outliers at 6 agents, Figure 20 shows that the chart exhibiting convergence time as a function of the number of agents is similar in shape to the one we have seen for the “Cross” graph, in Figure 17.

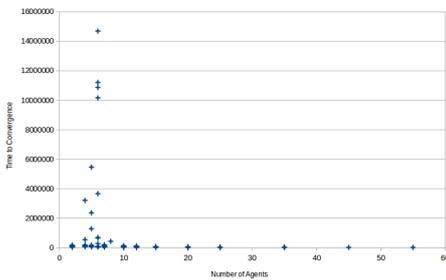


Figure 19: Convergence time of a system with a “6 Rooms” graph as function of the number of agents

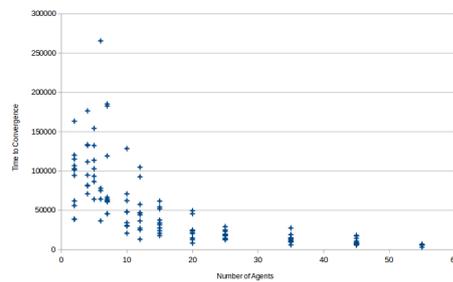


Figure 20: Zoom-In on the chart of Figure 19

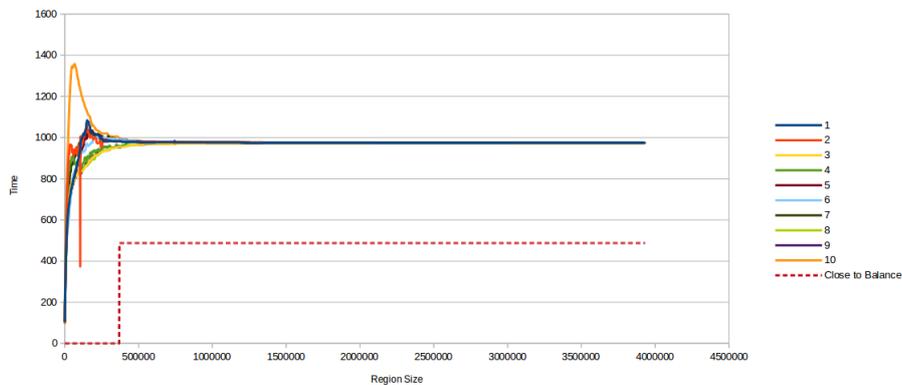


Figure 21: An evolution of a system of 10 agents and a “6 Rooms” graph of about 10,000 vertices. Convergence was achieved at approximately $t = 4,000,000$. But the system became “close to balanced” rather quickly.

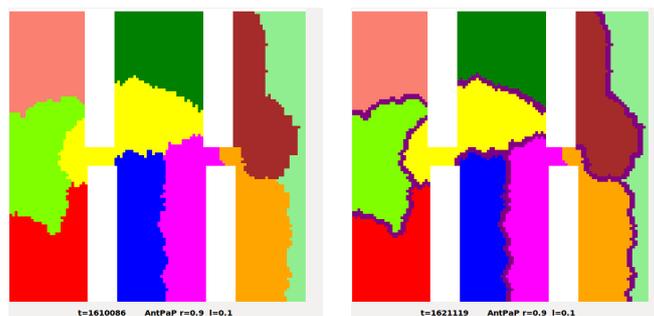


Figure 22: An evolution that reached “close to balanced” quickly, and remain so for long.

In the simulations described above, we tested the evolution of the multi-agent patrolling process until convergence to a stable and balanced partition. However remarkably, the system evolves rather quickly to close-to-balanced partitions due to the “balloon” forces implicitly driving the agents’ behavior. Therefore, for practical purposes we see that the AntPaP algorithm balances the work of the agents much earlier than its convergence time, and the partitioning becomes reasonably good rather quickly. This property is crucial in case of time varying topologies. Hence, AntPaP is a versatile and adaptive process. Considering again the “6 Rooms” example with 10 agents, we see in Figure 21 a temporal evolution of AntPaP until a stable and balanced partition is achieved. As is clear on the chart displaying the time evolution of the sizes of the 10 regions, the system reached convergence at approximately $t = 4,000,000$ steps. However it is also clear that after approximately 400,000 steps, the difference between the largest and smallest regions in the partition of the environment graph is already insignificant. In the chart, a system is defined as “close to balanced” when more than 99% of the graph is covered, and the difference between the largest and smallest regions is less than 5% the ideally balanced size (i.e. the graph size divided by the number of agents). Figure 22 exhibits a partition reached when the system was “close to balanced”. Both snapshots show the same partition (at two different times). The snapshot at the right also shows the borders between regions that “reached balance” (i.e. their size difference is at most 1 vertex) depicted in purple. There is only one border which is not yet balanced, between the magenta region and the dark yellow region, located in the right “corridor”. These regions are close in their sizes, and as a result the double visit condition does not occur very often. Despite the partition not being balanced yet, the division of work between agents is already fair, hence for practical purposes, a “close to balanced” condition is good enough.

We note in summary that AntPaP does not produce k -cut partitions [7], and generally assumes that there are no constraints on the grouping of vertices. Some important real-world problems impose such constraints, for example, the allocation of users in a social network to hosting servers, according to their interconnections. Other real-world problems, however, do not impose such constraints, for example, the division of work patrolling the world-wide-web for content analysis and classification. In view of the good properties discussed above, we envision that AntPaP could become a building block for distributed algorithms aiming to fairly divide between agents the labor of patrolling an environment, using very simple agents constrained to local interactions based on tiny “pheromone” marks left in the environment.

References

- [1] Y. Elor, A. M. Bruckstein, "Multi-a(ge)nt graph patrolling and partitioning". Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology-Volume 02. IEEE Computer Society, 2009.
- [2] I.A. Wagner, M. Lindenbaum, A. M. Bruckstein, “Efficiently searching a graph by a smell-oriented vertex process”. *Annals of Mathematics and Artificial Intelligence*, 24(1-4):211–223, 1998.
- [3] S. Even, “Graph Algorithms”. Rockville, MD: Comput. Sci. Press, 1979
- [4] P. Berkhin, "A survey of clustering data mining techniques." *Grouping multidimensional data*. Springer Berlin Heidelberg, pp. 25-71, 2006.

- [5] C. J. Alpert, A. B. Kahng, “Recent directions in netlist partitioning: a survey”, *Integration, the VLSI Journal* , 19(1–2):1-81,1995
- [6] K. Schloegel, G. Karypis, V. Kumar, “Parallel static and dynamic multi-constraint graph partitioning”, *Concurrency and Computation: Practice and Experience*, 14 (3):219–240,2002
- [7] R. G. Downey, V. Estivill-Castro, M. R. Fellows, E. Prieto, F. A. Rosamond, “Cutting Up Is Hard To Do: the Parameterized Complexity of k -Cut and Related Problems”, *Electronic Notes in Theoretical Computer Science*, 78():209–222,2003
- [8] R. Tarjan, “Depth-First Search and Linear Graph Algorithms”, *SIAM*, 1(2):146-160, 1972
- [9] J. D. McCaffrey. “Graph Partitioning using a Simulated Bee Colony Algorithm”, *IEEE Conference on Information Reuse and Integration (IRI) Las Vegas, USA*, pp. 400-405, 2011
- [10] F. Comellas, E. Sapena, “A multiagent algorithm for graph partitioning”. *EvoWorkshops, Lecture Notes in Computer Science 3907:279–285*, 2006
- [11] Y. Chevaleyre, F. Sempe, G. Ramalho, "A theoretical analysis of multi-agent patrolling strategies." *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 3. IEEE Computer Society*, pp. 1524-1525, 2004.
- [12] F. Lauri, F. Charpillet, "Ant colony optimization applied to the multi-agent patrolling problem." *IEEE Swarm Intelligence Symposium*, 2006.
- [13] M. Dorigo, V. M. Maniezzo, A. Colorni, "Ant system: optimization by a colony of cooperating agents." *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 26(1): 29-41, 1996.
- [14] E. Gyori, "On division of graphs to connected subgraphs." *Combinatorics (Proc. Fifth Hungarian Colloq., Keszthely)*. Vol. 1, 1976.
- [15] R. G. Gallager, *Stochastic processes: theory for applications*. Cambridge University Press, 2013.
- [16] A. Nerode, "Linear automaton transformations." *Proceedings of the American Mathematical Society* 9(4): 541-544, 1958



Copyright ©2025 by the authors. Licensee Agora University, Oradea, Romania.

This is an open access article distributed under the terms and conditions of the Creative Commons Attribution-NonCommercial 4.0 International License.

Journal's webpage: <http://univagora.ro/jour/index.php/ijccc/>



This journal is a member of, and subscribes to the principles of,
the Committee on Publication Ethics (COPE).

<https://publicationethics.org/members/international-journal-computers-communications-and-control>

Cite this paper as:

Elazar, G.; Bruckstein, A. (2025). AntPaP: Patrolling and Fair Partitioning of Graphs by A(ge)nts Leaving Pheromone Traces, *International Journal of Computers Communications & Control*, 20(1), 6907, 2025.

<https://doi.org/10.15837/ijccc.2025.1.6907>