Phermones versus Centralized control in an automated warehouse

a comparison of simulations

Marcel Flipse
1473379

Vrije Universiteit, Artificial Intelligence Department,
de Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands
marcel.flipse@gmail.com
Course: Master Orientation, supervisor: B.P.M. Weel

July 16, 2012

Abstract

In order to compare a system that uses pheromones to navigate the environment with a central controlled system, a total of 72 experiments were performed using 2 types of simulations. The results of these experiments show that the simulations that use pheromones perform better and show more interaction between agents than the central controlled simulations. This can be caused by the availability of different return routes and by the amount of resources spent on each simulation.

1 Introduction

Using robots for automated product pick up and placement in a warehouse environment is an interesting subject for research. Especially when systems like the Kiva system (http://www.kivasystems.com/), that uses a large number of relatively simple robots in combination with a centralized control system to place and retrieve products in a warehouse, are used more and more. It would be interesting to compare the efficiency of a central controlled and pheromone controlled system. In order to investigate this, two similar simulations were developed, that share a similar structure but have different methods for searching products, retrieving products, movement of robots and interaction between robots. This paper describes the simulations created and the results of experiments performed.

Some more about Kiva In the current Kiva system the environment is equipped with a barcode grid, which is used by the robots to know where they are and where they need to go. A central system tells the robot what to do on specific grid locations. Replacing the barcode grid with an Radio Frequency Identification (RFID)(Flipse[1]) tag grid will solve some maintenance issues of barcodes for the current Kiva system and will create new opportunities, like for instance enabling the robots to communicate with the grid and other robots at a greater distance. A pheromone system (Flipse[2]) can be used with the current Kiva environment, that still uses the current barcode grid. But when a barcode grid is used, each pheromone action requires communication between a robot and the central system. So now bandwidth can become an issue. When the pheromone system is introduced together with the RFID tag grid, this communication is mainly local, between a RFID tag and a robot, so bandwidth usage should not become problematic. Another feature of the robots in the Kiva system that is worth mentioning, is their ability to drive underneath product racks. This feature is made possible because the product racks have sufficient space underneath them so the robot is able to maneuver. This space also enables the robot to pick up a product rack with the help of a screw mechanism inside the robot. For the simulations developed, in order to keep the complexitiy of the simulations under control, the robot will only be allowed to move underneath a product when it wants to pick up a product.

The robots in the Kiva system are equipped with sensors. These sensors can detect when a collision has occurred and make sure the damage of the collision is limited. Robots might also be equipped with RFID tags that could be used in the environment. This will enable robots to identify other robots remotely and identify potential collisions or dangerous situations locally.

2 Research

The research focuses on comparing a central controlled system and a pheromone system for robots in a warehouse environment. In order to do this research two simulations were developed. This section continues with describing the requirements for the simulations and the actual implementations. The environment and agent structure will also be explained.

2.1 System description

The robots in the simulation, from now on called agents, will place and retrieve products according to an order list. How these products are retrieved or placed differs for each control method and order list combination.

Centralized control In this simulation a centralized system will execute all the complex actions, such as finding the location of the targeted product, while the robots only move to certain grid locations. The central system will search for the closest product that matches an item on the order list. The agent will then go to the grid location and performs a specific action before it returns using another path to deliver the product at the dropping location. This simulation will only contain basic actions. This means the agents will not avoid each other with more complex actions than stop and give priority to other agents.

Pheromone control In this simulation a central system only gives a combination of target product, that needs to be found, and an action, that needs to be performed such as retrieve or place a specific product. The agent uses the environment to find the target product and also updates the environment by communicating with the environment. For instance when an agent enters a product path, this will be communicated to the environment and this information is than stored in the environment. Also when an agent is retrieving or placing a product, this will also be communicated with the environment. One of the issues with this kind of control is that the basic implementation will require the agent to take the same path towards the product as towards the dropping locations. If the agent takes another route, the pheromones will not be updated with the correct information and the system stops working correctly. But when the agents take the same paths while searching or returning, they will encounter other agents often. In order to limit these encounters and associated actions, the number of agents for each path should be limited.

Environment The warehouse environment should support multiple agents and the ability to place products in the environment. Furthermore it should be possible to remotely scan the environment for closest products and the environment should be able to register certain actions, like moving a product through the environment. Also it should be possible to name parts of the environment and that this name can be read by agents. And last but not least, a central system should be able to communicate orders to the agents.

2.2 Implementation

Starlogo (http://education.mit.edu/starlogo/) is an existing simulation environment that can handle the control of multiple agents and provides a separated agent and patch environment that can be manipulated with syntax similar to the Java programming language. This similarity with Java, makes it easy to program, however it should be noted that Starlogo is far more limited than Java. Another reason for choosing Starlogo is that it is an existing simulation environment. This will probably save development time compared to developing an own simulation environment. Another reason is that there are clear rules about communication between agents and their environment. When creating a new simulation environment from scratch, such rules can too easily be adjusted. Although Starlogo has not been updated since 2006, it still provides a good enough simulation environment for this project.

2.3 Environment structure

The environment exists of patches and turtles. These turtles are the agents of this system and to avoid any confusion will be referred to as agents. Agents are able to read information from patches and write information to patches. They can also identify other agents, but for this simulation they cannot give other agents assignments. The environment is divided in certain zones. Agents start at the start location or wait in line at the bay zone. At the start location, the agent receives an order from an order list. After receiving an order, the agent takes the entry zone towards a path zone which has a connection to a product that the agent is searching. The path zone is connected to product zones, which contain either products or are empty. Next to the start and finish locations, the product zone is the only zone where products are allowed to be dropped off or picked up. When a product is found, the agent takes a path to a return road and delivers the product to the finish, where the product is delivered and the cycle begins again until the order lists are all empty.

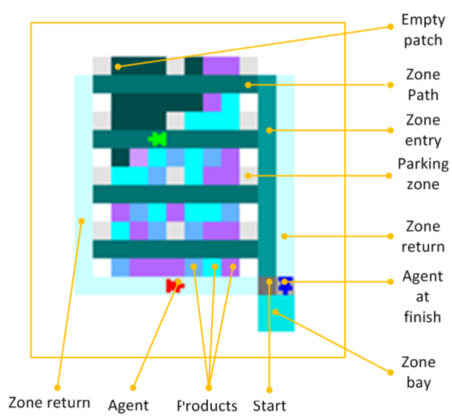

Figure 1: A screenshot of the used simulation together with the names of the most important visible parts of the system.

Centralized With the centralized approach, the agents are limited to ask patches which products they contain and which zone name the patches have. The only write actions the agents perform is when they drop a product at a location and when they are at the finish location when they are delivering a product.

Pheromone With the pheromone approach, the agents have more interaction with the environment. Patches contain all sorts of information such as how many of which product can be found. When agents move over a patch, the agent registers if it is retrieving or adding a product and updates the information of the patch accordingly. So if the pheromone first contained the information that the environment contained 12 of product "Alpha" and an agent with a retrieve "Alpha" order moved over the patch, the agent updates the patch that there will be 11 products "Alpha". For this type of control, 4 types of pheromones are used for products. "Alpha", "Beta", "Charlie" pheromones refer to the products and the "Zero" pheromone refers to an empty product zone. To somewhat prevent possible collisions in the pheromone approach, a traffic light system is introduced. When there are 2 agents in the same product path, the product path is closed until there are less then 2 agents in that product path. This is done by adding traffic information to the pheromone trail. When an agent sees a closed path it looks if there is another path available that helps the agent to fulfill its mission. One zone that is only used by the pheromone approach is the parking zone. This zone is used when one agent tries to avoid another agent, so the other agent can pass.

2.4 Agent structure

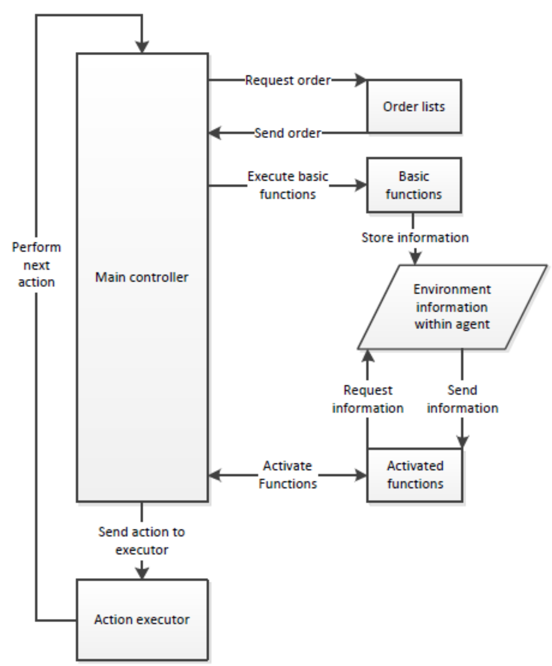

Figure 2: The model that is used by the agents to perform an action.

Agents reason with help of a certain structure to know what to do next. First the agent searches for and moves towards the start location. When the agent has reached the start location, it request an order from a central stored order list. When the agent receives an order, it activates a basic function that translates what the coordinates are for left, right, top and a few more coordinates that are needed for reasoning. Other basic functions are updating which zone the agent currently is present and collision detection. Collision detection detects if a collision might happen the next time the agent would move. All the information that is collected by the basic functions is stored within the agent. If a collision is detected, the agent will first avoid the collision before it does any other actions. Other functions that can be activated, that are non basic functions and use information that is stored within the agent, are the follow pheromone function, pick up product function, return function, deliver function and wait in line function.

Pheromone and Centralized The pheromone approach uses the same structure as the centralized approach. However a few non basic actions are different. For instance with the pheromone approach, the agent follows a pheromone track to find products, while in the centralized approach, the grid coordinates are given to the agent. With the pheromone approach the pheromone track is continuously updated, while the pheromone track does not even exist in the centralized approach. Furthermore with the pheromone approach, a more complex avoiding function is needed that not has been implemented in the centralized approach. A major difference between the simulations is that the centralized approach uses a different return route than was used to get to the targeted location, while the agents in the pheromone simulation use almost the same route to the target as from the target. This will result in a lower amount of total actions but a higher amount of avoid and stop actions for the pheromone approach compared with the centralized approach.

3 Experiment

3.1 Setup

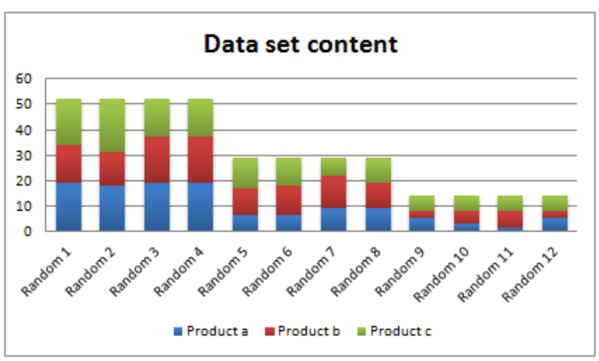

Figure 3: A visualization of the amount of each product sort in the randomly generated lists. This figure does not represent the order of products used.

Product lists In order to perform this research experiments needed to be created. Each experiment consisted out a product placement list and a retrieval list. Three different sizes were used for these lists. The full house experiments consisted out lists with 52 products, the half house experiments consisted out lists with 29 products and the quarter house experiments consisted out lists with 14 products. For each list size, 4 basic lists were randomly generated. These basic lists formed the first set of experiments were the product placement list was exactly the same as the product retrieval list. Figure 3 shows how many of each product is present in each list. In order to create more experiments, the basic lists were manipulated. The manipulation sorted the basic lists by product. These lists were then used as placement lists while the original lists were used as the retrieval lists. To create a third experiment set, this combination was reversed, so that the sorted lists were used as a product retrieval lists and the original lists were used as the product placement lists. A possible forth experiment set that used the sorted sets for both the product placement and the product retrieval was not used, because this would be similar to the first experiment set. To summarize the experiments, 2 methods were tested while using 3 combinations of product retrieval and placement lists, 3 different list sizes and 4 randomly generated sets. This resulted in a total of 72 experiments.

Starlogo All experiments used a Starlogo simulation speed of 98%. This was done because when running at full speed, the first agent will have done a retrieval or placement action before the other agents could perform one move. While running at 98% this problem did not exist.

Experiment All the experiments used a similar stop condition. When an agent finds an empty product retrieval list and the last agent in the environment has returned to the bay zone, the experiment would stop. During an experiment the number of all actions, wait actions and avoid actions are stored and these will be used to compare the centralized control method with the pheromone method.

3.2 Results

This section will describe the results of the 72 experiments performed.

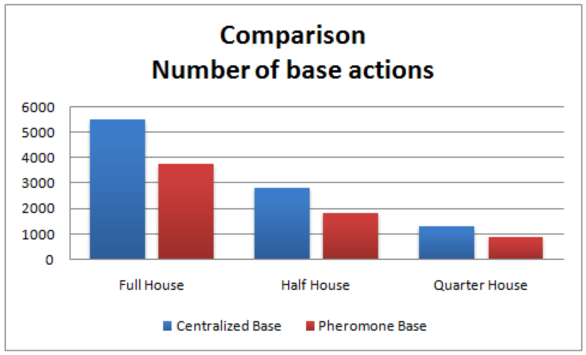

Figure 4: When avoid and stop actions are subtracted from the total actions, the number of "base" actions can be found. The x-axis shows the different size categories while the y-axis shows the number of base actions used in each size category.

Base Actions Figure 4 shows the base actions for both methods. The basic actions are calculated by subtracting the wait and return actions from the total number of actions. This resulted in the same number of base actions for experiments in a specific size category. For each experiment, the pheromone method needed a smaller amount of base actions than the central control method.

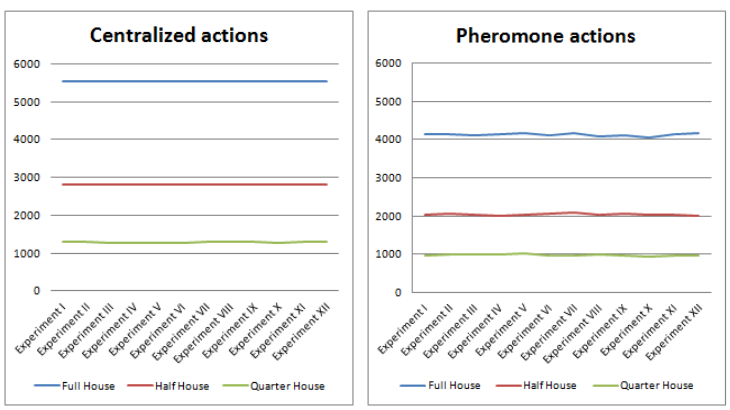

Figure 5: Left: A graph of the number of actions that were needed to complete each experiment while using centralized control. Right: A graph of the number of actions needed that were needed to complete each experiment while using pheromones.

All actions Figure 5 shows the total number of actions on the y-axis and each experiment performed on the x-axis. The experiments that used the pheromone control method needed less actions than the experiments that used the central control method. It must be noted, that at the current zoom level, the number of actions that were used for the centralized control method experiments, seem more stable than they actually are.

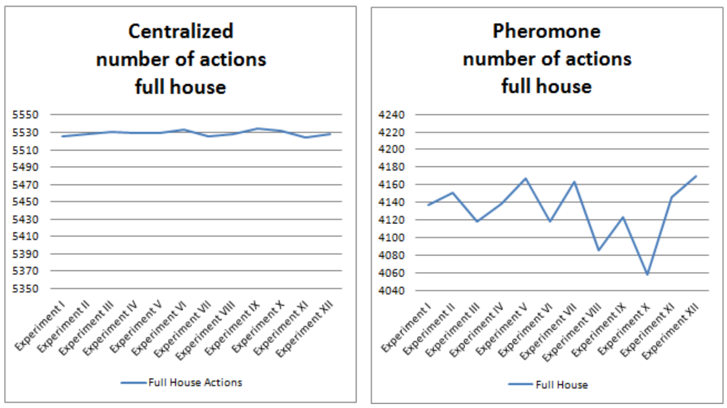

Figure 6: Left: A graph of the number of actions that were needed to complete each full house experiment while using centralized control. Right: A graph of the number of actions that were needed to complete each full house experiment while using pheromones.

Full house actions Figure 6 provides a closer look at the number of actions used for the central control method and the pheromone method while using full house experiments. The y-axis shows the number of actions used for each full house experiment on the x-axis. Note that the y-axis on both graphs spans 200 actions, but have different numbers for the range. This made comparing the two methods easier while preserving the characteristics of both systems. The centralized control method is not as stable as figure 5 shows, however the centralized control method is more stable than the pheromone method. The results of the other size categories were similar to the size category showed in figure 6.

Wait and avoid actions Figure 7 shows the number of wait and avoid actions used for each experiment and each experiment size category. The central control method does not use any avoid actions, so these are not included in the figure. The number of wait actions that are used for the central control method experiments is lower and more stable than the number of wait actions used for the pheromone method. Furthermore the number of avoid actions used for the pheromone method is lower than the number of wait actions used for the pheromone method.

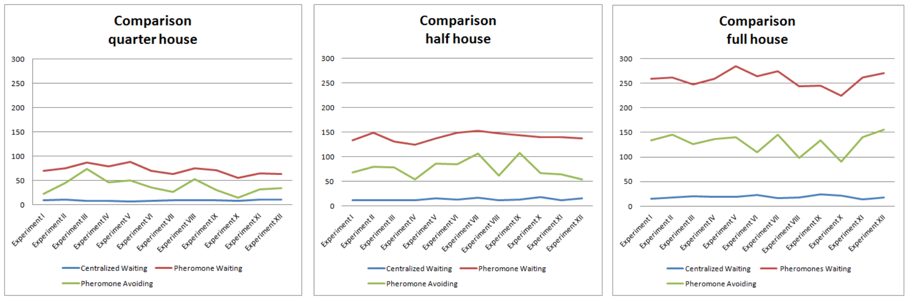

Figure 7: Number of wait and avoid actions for each system for the full house(right), half house(mid) and quarter house(left) experiments. For each graph in this figure, the y-axis shows the number of actions used for each experiment on the x-axis.

3.3 Additional experiment and results

An additional experiment was performed to research if the simulations would always give the same results by testing the same experiment several times. Figure 8 shows that this is not the case, however the variation is acceptable. The variation for the centralized simulation was smaller than the variation of the pheromone simulation. For the graphs in this figure the y-axis shows the number of actions needed for each instance of the same experiment on the x-axis. It should be noted that for this figure, the span of the number of actions for both graphs is 100, however both graphs use a different range.

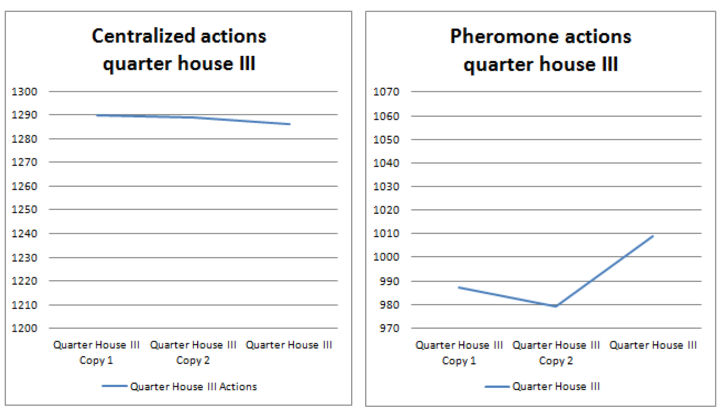

Figure 8: Left: number of actions used for 3 copies of the same experiment for the centralized approach. Right: number of actions used for 3 copies of the same experiment for the pheromone approach.

3.4 Discussion

The difference in number of actions used for the centralized and pheromone simulations can be caused by the difference in routes used. The agents in the centralized simulations use a different and longer return route than the agents in the pheromone simulations. The effect of the increased length for this different return route might have been underestimated. The route choice also explains the difference in avoid actions and wait actions used by the agents in the simulations and their stability. Because the shorter return route of the agents in the pheromone simulations use almost the same route to a location as from a location, this results in more potential collisions, together with more avoid and stop actions that are needed to prevent these potential collisions. The centralized simulations use a different path to a location as from a location, so there are less potential collisions. The most efficient avoid action in the centralized simulations would be just to wait until the collision is no longer possible, so separate avoid actions are not needed. Next to the availability of different routes, the more stable behavior of the centralized simulations in comparison with the pheromone simulations can also be explained by the usage of avoid actions. When agent "Delta" avoids another agent "Echo", agent "Delta" interrupts its planned route and removes itself from the path. Now it is possible that while undoing this avoid action another agent "Foxtrot" forces agent "Delta" to take another avoid action.

Expectations As expected, an increase in the number of products used in an experiment leads to an increased number of used actions for both the central control method as for the pheromone method. Furthermore the number of wait actions also increases when the number of products used in an experiment increases. That the number of wait actions are always higher than the number of avoid actions can be explained by the number of zones were these actions can be performed. The stop action is an action that can be used in almost all the zones while the avoid action is only available when the agent is in the path zone. It was not expected that when repeating the same experiment multiple times, different results would appear. It is suspected that Starlogo uses some sort of randomized order algorithm to determine which turtle has processing priority.

Resources The difference in performance between both systems can also be caused by the higher amount of resources that were spend on the pheromone simulation. The development of the pheromone simulation used so many resources that there were little resources left for the development of the centralized simulation. It should be possible to develop a better performing centralized simulation when more resources become available. However the same holds for the pheromone simulation.

3.5 Conclusion

The experiments performed, in order to compare a centralized control method with a pheromone method in a warehouse simulation while using agents to place and retrieve products, produced some expected and unexpected results. As expected, when an experiment contained more products, more actions were needed. For each experiment that used the same combination of placement and retrieval lists, the pheromone method needed a lower number of total actions than the centralized control simulations. The number of wait and avoid actions were higher and less stable for the pheromone method than for the centralized control method. It was unexpected that when the same experiment was repeated multiple times the results showed a small variation.

To summarize this research, based on the results of 72 experiments, the developed simulations using centralized control, show less interaction between agents and is found less efficient than the developed simulations using pheromone control. However it needs to be noted that most resources during this research were used for developing the pheromone simulations.

3.6 Future work

This research may be continued by implementing the same return route for both methods. The pheromone method can easily be adapted so that it uses the same return route as the centralized control method. The stoplight system used in the pheromone method can be adapted so that it can open a path remotely. However implementing this in a real world environment might become impossible due to technology limitations. Using the stoplight system for the centralized control method could also be an option, so the centralized control method could use the same return route as used in the pheromone method. However this probably would lead to challenges while retrieving products. Future work should also include a section about why there are different results while repeating the same experiment.

References

[1] Marcel Flipse, *RFID techniques for indoor warehouse location sensing.* "http://www.itwillwork.nl/rfid-paper.pdf", 2011.

[2] Marcel Flipse, *Altering and Improving Kiva.* "http://www.itwillwork.nl/kiva-paper.pdf", 2011.