Engineering Letters, 19:4, EL_19_4_02

A Closed-loop Bid Adjustment Approach to Dynamic Task Allocation of Robots

W. K. Zhu and S. H. Choi

***Abstract*—Dynamic task allocation is among the most difficult issues in multi-robot coordination, although it is imperative for a multitude of applications. Auction-based approaches are popular methods that aim to assemble robot team information at a single location to make practicable decisions on task allocation. However, a main deficiency of auction-based methods is that robots generally do not have sufficient information to estimate accurate and reliable bids to perform tasks, particularly in dynamic environments where there are operational uncertainties. While some techniques have been developed to improve bidding, they are mostly open-looped without feed-back adjustments to tune the bid prices for subsequent tasks of the same type. Robots' bids, if not assessed and adjusted accordingly, may not be trustworthy and would indeed impede team performance.**

To address this issue, we propose a closed-loop bid adjustment mechanism for auction-based multi-robot task allocation to evaluate and improve robots' bids, and hence enhance the overall team performance.

Each robot in a team maintains and uses its own track record as closed-loop feedback information to adjust and improve its bid prices. After a robot has completed a task, it assesses and records its performance to reflect the discrepancy between the submitted bid price and the corresponding actual cost of the task. A series of such performance records, with time-discounting factors, are taken into account to damp out fluctuations of bid adjustments. Adopting this adjustment mechanism, a task would be more likely allocated to a competent robot that submits a more accurate bid price, and hence improve the overall team performance. Simulation of task allocation of free-range automated guided vehicles serving at a container terminal is presented to demonstrate the effectiveness of the bid adjustment mechanism.

***Index Terms*—Multi-robot, task allocation, auction, bid adjustment, dynamic environments**

W. K. Zhu is with The Department of Industrial and Manufacturing Systems Engineering, The University of Hong Kong, Hong Kong, (e-mail: wenkaizhu@hku.hk).

S. H. Choi is with The Department of Industrial and Manufacturing Systems Engineering, The University of Hong Kong, Hong Kong, (e-mail: shchoi@hku.hk).

I. INTRODUCTION

Multi-robot task allocation addresses the problem of how to assign a set of tasks to the corresponding robots while achieving some specific objectives of team performance. Research in this field dates back to the late 1980s. Task allocation of multiple autonomous robots in dynamic environments is core to multi-robot control for a number of real world applications, such as military [1], transport services [2], search and rescue [3], etc., as shown in Fig. 1. Relevant research has flourished ever since. The European Union has sponsored several swarm robot projects. The I-SWARM project, for instance, aimed to develop a swarm of robots to perform cooperative tasks, such as foraging. An inter-university SWARMS initiative in the United States tried to develop a new system framework for controlling a swarm of mobile robots, synthesizing emergent behaviours for reactive response, and developing algorithms for decentralised transport task assignment [4].

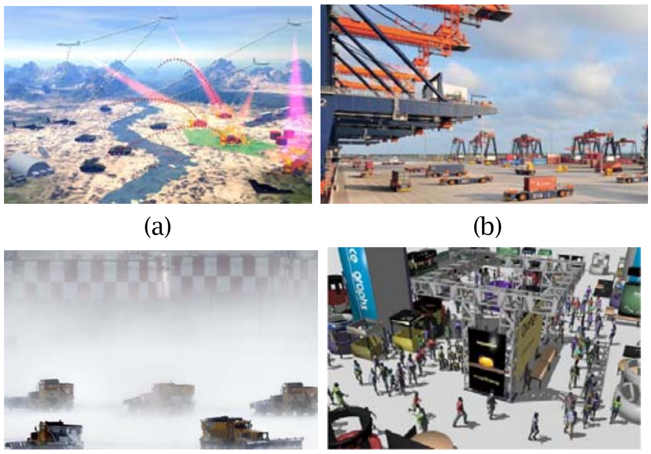

(a) (b) (c) (d)

Fig.1. (a) Distributed heterogeneous military systems [1]; (b) Automated guide vehicles at a container terminal [5]; (c) Snow removal at Philadelphia Airport [6]; (d) Simulation of multi-agents with respective task assignments in computer graphics [7].

The working environments of task allocation can be static or dynamic [8]. Static task allocation assumes completely known information about the environment, such as the number of tasks and robots, the arrival time of tasks, and the process of task execution. Traditionally, applications in multi-robot domains have largely remained in static

(Advance online publication: 12 November 2011)

scenarios, with an aim to minimise a cost function, such as total path length, or execution time of the team. Obviously, static approaches cannot adapt to changes in a dynamic and uncertain working environment. Dynamic task allocation, on the other hand, makes decisions based on real-time information and is therefore more adaptive to changes. This paper assumes a set of dynamically released tasks to be completed by a team of robots, and the conditions of the work process keep changing during task execution. This kind of dynamic working environment is ubiquitous in real-life applications, such as exploring and mapping by robots in an unknown environment, unexpected adversarial targets in a combat, stochastic pickup and delivery transport services, etc. [9].

According to the taxonomy of Gerkey and Mataric [10], multi-robot task allocation problems can be classified along three dimensions. In the dimension of robot, it can be a single-task robot or a multi-task one. A single-task robot is capable of executing exactly one task at a time, while a multi-task robot can handle more than one task simultaneously. In the dimension of task, it can be a single-robot task or a multi-robot task. A single-robot task requires only one robot to execute it, while a multi-robot task requires more than one robot to work on it at the same time. In terms of planning horizon, task allocation can be instantaneous assignment and time-extended assignment. Instantaneous assignment only considers the tasks currently available. Time-extended assignment elaborates the effect of current assignment on future assignment, involving task dependency and schedule. Instantaneous assignment is commonly used since it needs less computation on task sequencing algorithms, and is particularly practicable in dynamic situations where tasks are randomly released [11]. The task allocation problem in this paper is restricted as a single-task robot, single-robot task, and instantaneous assignment.

The problem of multi-robot task allocation is typically NP-hard. The challenges become even more complicated when considering operations in dynamic and uncertain environments, such as unexpected interference between robots, stochastic task requests, inconsistent information, and various component failures [8]. In these cases, it is not worth spending time and resources to secure an optimal solution, if that solution keeps changing as operations go on. Also, if there are time-window constraints, there may not be enough time to compute an exact and global solution. The basic objective of a multi-robot task allocation problem is to have tractable planning that produces efficient and practicable solutions. Auction-based, or market-based, approaches manage this by assembling team information at a single location to make decisions about assigning tasks over the team to produce practicable solutions quickly and concisely [12].

In an auction, a set of tasks are offered by an auctioneer in the announcement phase. Broadcast with the auction announcement, each of the robots estimates the cost of task separately and submits a bid to the auctioneer. Once all bids are received or a pre-specified deadline has passed, the auction is cleared. In the winner determination phase, the auctioneer decides with some selection criteria which robot wins which task [13]. This paper considers the case of cost minimisation in that an auctioned task is awarded to a robot offering the lowest bid price. A simple yet commonly used kind of auction is single-item auction in which only one task is offered at a time [14]. On the other hand, combinatorial auctions are more complex in that multiple tasks are offered and each participant can bid on any combination of these tasks. Since there are an exponential number of combinations to consider, auction administration such as bid valuation, communication, and auction clearing would soon become intractable [9]. Sequential single-item auction is a practicable approach when tasks are dynamically released, and is adopted in this paper.

The earliest example of auction-based multi-robot coordination appeared about thirty years ago, called contract net protocol [15]. Auction-based multi-robot coordination approaches have been growing in popularity in recent years. They have been successfully implemented in a variety of domains, such as robotic transport [16], [17], mapping and exploration [18], house cleaning [19], and reconnaissance [20]. Choi et al. [21] developed an auction-based decision strategy as the mechanism for decentralised task selection among a fleet of autonomous vehicles. They also used a consensus routine which was based on local communication as the conflict resolution mechanism to achieve agreement on the winning bid values. Numerical simulation showed that feasible solutions were robust against both inconsistency of information across the fleet and variations in the communication network topology. Jones et al. [22] designed an auction-based approach to multi-agent coordination for disaster response, with intra-path precedence constraints. A group of five fire-truck agents attempted to extinguish fires spreading throughout a simulated community. This auction-based approach was incorporated with two heuristic techniques, called clustering and opportunistic planning, to perform a bounded search of possible schedules and allocations. They tried to improve the approach so that it would handle uncertain task situations, where fires at new locations were constantly being discovered.

Auction-based approaches are preferable in on-line applications, in that they can quickly and concisely assemble team information at a single location to make decisions, significantly reducing the combinatorial nature of task assignment problems [23]. The solution quality, although not optimal, is guaranteed in most cases. Auction-based approaches are suitable for dynamic and uncertain applications since they can accommodate new information through frequent auctioning of tasks [9].

However, some issues of auction-based task allocation have yet to be further investigated [9]. Firstly, a clear conceptual understanding of auction-based coordination approaches is needed. Further works should be devoted to studying how components, such as performance assessment mechanism, bidding strategy, and auction clearing mechanism, can be implemented effectively in different multi-robot applications. Secondly, the fundamental premise of success in an auction relies on the ability of individual robots to make reasonable cost estimation and submit acceptably accurate bid prices. However, robots generally do not have sufficient information for reliable cost calculation, which requires an accurate model of the environment and computation-expensive operations. Thus, heuristics and approximation algorithms are commonly used, such as the first-come-first-served and the shortest-distance-first. Some progress has been made to improve the accuracy of bid price. In the work of [24], two physical robots executed distributed sensing tasks in a cell-based map. Path costs were estimated using the D* path planning algorithm with optimistic cost of unknown map-cells. It was demonstrated that auction-based approaches could improve team efficiency if cost estimation considered the environmental and mission characteristics. Duvallet and Stentz [25] applied an imitation learning technique to bias the bid prices in auctions to make better solutions. Two simulated scenarios were presented to demonstrate the applicability of this technique, including three fire-fighting agents putting out fires in buildings, and eight players in an adversarial game trying to score more points than their opponents. This approach needed a considerable amount of training samples and time to reach a reasonable solution, and the learning rate should be skilfully tuned. Khan et al. [26] developed a platform to simulate a team of five robots cooperating to move some objects to the specified goals, based on an auction-based task assignment method. The number of robots involved in an operation could be adaptively changed. They tried to improve the accuracy of cost estimation by an elaborated bidding function, which considered a set of environment conditions, such as the distance between a task and a robot, the velocity and orientation of a robot, obstacles in the path between a robot and a task, and the possible success rate of a task. While the simulation results were encouraging, it seemed that the selection criteria of environmental factors needed further justification.

Nevertheless, the above-mentioned approaches to improving cost estimation, like most of the current auction-based methods, are open-looped [27]. They cannot assess whether or not a bidder has kept its commitment to a task, because they do not have a mechanism to evaluate the bidder's performance after winning the task. Human bidders are self-interested in auctions, and would sometimes deliberately offer over-optimistic bid prices. Robots, on the other hand, are assumed to be honest in estimating the costs before offering the bid prices. However, there are often discrepancies between the bid prices and the actual costs in real-life applications, particularly in dynamic working environments. Discrepancies between the bid prices and the actual costs are usually caused by the uncertainties of a dynamic environment, such as unexpected task requests, changing traffic conditions, communication delay, inconsistent information, and stochastic component failures [9]. Unfortunately, these uncertainties are difficult to explicitly model in advance. By submitting either over-estimated or under-estimated bids, robots may not be able to deliver on their task promises. As a result, the overall team performance would be significantly hampered.

This paper therefore presents a closed-loop bid adjustment mechanism for auction-based multi-robot task allocation in light of operational uncertainties, with which a team of robots can evaluate and improve their bids respectively, and hence enhance the overall team performance. Each of the robots in a team maintains an array of track records for different corresponding types of tasks it has ever executed. After a robot has completed a specific type of task, it assesses and logs its performance to the corresponding track record, which reflects the discrepancy between the submitted bid price and the related actual cost of the task. These track records serve as closed-loop feedback information to adjust and improve the bid prices in future auctions. Moreover, when adjusting the bid price of a task, a series of performance records, with time-discounting factors, are taken into account to damp out fluctuations. As such, bid prices can be regulated and fine-tuned to alleviate some deviations of cost estimation due to operational uncertainties. Tasks are more likely allocated to competent robots that offer more accurate and reliable bids, resulting in significant improvement in the overall team performance.

Section 2 introduces the details of the proposed bid adjustment mechanism in auctions. Section 3 presents its implementation in a task allocation algorithm for simulation of free-range automated guided vehicles serving at a container terminal to demonstrate the effectiveness of the adjustment mechanism. Section 4 draws conclusion and discusses some future work.

II. The Closed-loop Bid Adjustment Mechanism

As concluded in Section 1, most of the traditional auction-based methods are open-looped. They cannot assess whether or not a bidder has kept its commitment to a task, and there is no mechanism to attain the feedback information to adjust and improve the bid prices for future auctions, as shown in Fig. 2.

Fig. 2. Traditional open-looped auction-based methods

A. *The Task Auction Architecture*

This paper instead proposes a closed-loop mechanism to adjust and improve the bid price, and hence to enhance the team performance. Fig. 3 shows the task auction architecture. Task allocation in this paper is restricted to a single-task robot, single-robot task, and instantaneous assignment problem. Hence, we adopt sequential single-item auction for situations where different types of tasks are stochastically released for auction during operation. A central processor is the auctioneer who auctions these tasks one by one. All the idle robots bid for a task being auctioned, and the one that submitted the lowest bid price wins the task.

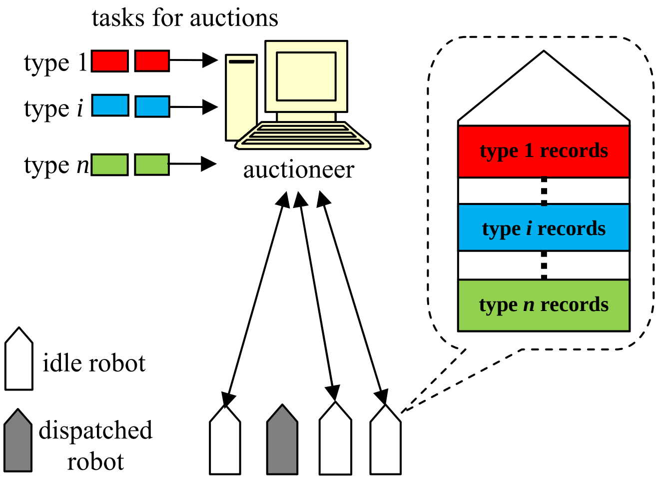

Fig. 3. Task auction architecture

B. The Closed-loop Bid Adjustment Mechanism

Each of the robots maintains an array of track records for the corresponding types of tasks that it has ever executed. Fig. 4 shows a block diagram of this bid adjustment mechanism. After a robot has completed a specific type of task, it evaluates its own performance and records a reward or a penalty accordingly. This track record facilitates adjustment of the bid price that the robot in question will subsequently submit for another task of the same type.

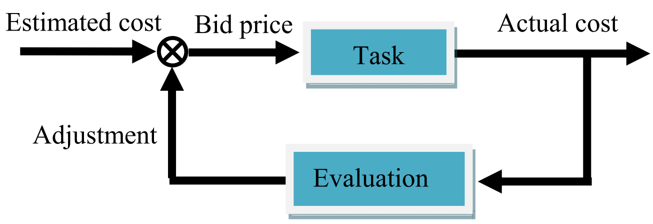

Fig. 4. The closed-loop bid adjustment mechanism

C. Task Auction Procedure

The procedure of a typical auction, with the proposed close-loop bid adjustment mechanism, is introduced as follows.

When a task is stochastically released and becomes ready for auction, the central auctioneer is requested to initiate an auction. The auctioneer broadcasts this specific type of task to the robot team, with detailed description of the task requirements. Some robots in the team may be occupied with their on-going tasks, and would simply ignore this auction announcement. Only the idle robots will take part in this auction. Each of the idle robots first estimates the cost of the task by some heuristics or approximation methods, such as assuming the route would be straight and there would be no obstacles in the way. Then, it adjusts the estimated cost, based on its track records for this type of tasks, to obtain a bid price. Subsequently, the resulting bid price is submitted to the auctioneer before the clearing time. After receiving all the bids from the robot team, the auctioneer assesses the bid prices and awards the task to the robot offering the lowest bid price. The robot awarded with the bid executes the task. After completion, the robot assesses its performance by comparing the actual cost with the bid price of this task to attain an adjustment in the form of either reward or penalty. This bid adjustment is then logged in the robot's track record for this type of tasks, which will serve as a closed-loop mechanism to adjust and improve the bid prices in future auctions.

D. Algorithm of the Bid Adjustment Mechanism

This section presents the detailed algorithm of the closed-loop bid adjustment mechanism. As shown in Fig. 5, each robot maintains an array of track records for each corresponding type of tasks it has ever executed.

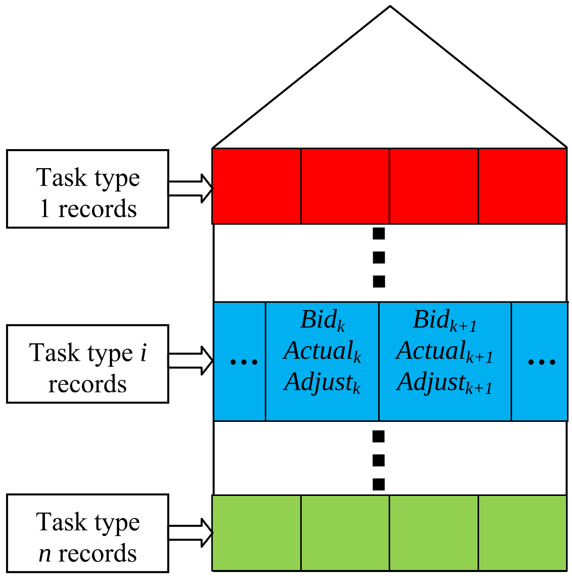

Fig. 5. Track records of a robot

For a specific type of tasks, such as type i in Fig. 5, we denote $Actual_k$ as the k^{th} record of actual cost, and Bid_k as the k^{th} record of bid price. Adjustments are in the form of either rewards or penalties:

$$Adjust_k = Actual_k - Bid_k \quad \text{..................(1)}$$

If the actual cost of the task is smaller than the proposed bid price, it means that the robot has kept its commitment to the task or even has performed better than its promise. The robot will get a reward with a negative adjustment value. Similarly, a penalty with a positive adjustment value will be imposed to this robot if the actual cost is larger than the proposed bid price, which means the robot has underestimated the cost of task and has not completed it as promised.

When a robot bids for a next task of the same type, it first estimates the cost, and then tunes the bid price based on the previous adjustment:

$$Bid_{k+1} = Cost_{k+1} + Adjust_k \quad \text{..................(2)}$$

where $Cost_{k+1}$ is the *(k+1)* th estimated cost, which can be acquired by other heuristics or approximation methods.

To damp out huge fluctuations and to reflect more reliable estimations, a series of previous adjustments should be taken into account. Moreover, since the working environment is dynamically changing, older track records are deemed relatively obsolete as time elapses. Hence, a time-discounting factor α, where $0 < \alpha < 1$, is introduced to weigh the track records. The averaged bid adjustment is:

$$\frac{\sum_{j=0}^{k-1} \alpha^j Adjust_{k-j}}{\sum_{j=0}^{k-1} \alpha^j}$$

In practice, the latest three terms are sufficient for adjustment of the bid price. The complete form of the proposed bid adjustment mechanism is given in equation (3).

$$Bid_{k+1} = Cost_{k+1} + \frac{\sum_{j=0}^{k-1} \alpha^j Adjust_{k-j}}{\sum_{j=0}^{k-1} \alpha^j} \quad \text{............(3)}$$

The task being auctioned is therefore assigned to the robot that submitted the lowest adjusted bid price, based on equation (3). As such, this closed-loop bid adjustment mechanism can improve bidding accuracy, considerably enhancing the overall team performance.

E. Robustness Analysis of the Algorithm

With the closed-loop feedback mechanism, the bid price submitted by a robot for a task being auctioned can be regulated and fine-tuned to mitigate deviations of cost estimation due to operational uncertainties. Moreover, a series of adjustment values are averaged with related time-discounting factors to damp out possible fluctuations of adjustments, further safeguarding the robustness of the closed-loop regulation mechanism. Therefore, the stability of the proposed approach can be effectively secured.

F. Workflow of Serial Auctions for all Tasks

The complete workflow of task auctions during operation is listed as follows:

Step 1: A task is released and a request for auction is sent to the auctioneer to announce;

Step 2: For each idle robot to participate in the auction:

- (2a): If this type of task has NOT been executed before, sets the bid adjustment to 0 and creates a track record for this type of tasks;

  else

  Reads the bid adjustment from the track record;
- (2b): Estimates the cost of the task;
- (2c): Adjusts the bid price by adding the bid adjustment to the estimated cost;
- (2d): Submits adjusted bid price before clearing time;

Step 3: The auctioneer assesses all the bid prices received and awards the task to the robot offering the lowest bid price;

Step 4: The winning robot executes its awarded task until it is finished;

Step 5: The robot compares the actual cost with the proposed bid price, and updates the bid adjustment;

Step 6: The robot logs the bid adjustment into its related track record and calculates the averaged adjustment value;

Step 7: Repeats from Step 1 until no more task is released.

III. Implementation and Case Studies

The closed-loop bid adjustment mechanism is incorporated with a multi-robot dynamic task allocation algorithm in a simulator, which also includes a module for motion planning of a fleet of robots.

A. Motion Planning for Robot Teams

In an operation involving a team of mobile robots, task allocation, which is the focus of this paper, is the first step to assign tasks among the robot team. After a robot is allotted a task, it plans the motion entailed to finish the task, such as picking up and delivering a cargo container. Multi-robot motion planning addresses the problem of how a team of autonomous mobile robots can share the same workspace, while avoiding interference with one another to achieve group motion objectives [4]. Motion planning for multiple autonomous robots to navigate safely and avoid moving obstacles in dynamic environments is still among the most difficult and important problems in multi-robot control [28].

The approach implemented for motion planning in this paper is a bio-inspired intelligent technique of motion planning for mobile objects in dynamic environments. The details of this motion planning approach can be found in [29]. In short, a motion planning cycle of each robot includes three stages, namely local sensing, detecting imminent neighbours, and real-time navigation. Fig. 6 shows the flow of such a cycle, which should be sufficiently short in order to take adaptive reactions to dynamic traffic conditions. The motion planning approach is inspired by the natural behaviours of creatures which tend to keep a safe distance between one another, and move towards their respective destinations. This approach does not require any central controller, and there is no communication between robots. It instead imitates the characteristics of creatures with local sensing for detection of imminent neighbours and navigation. Each robot is assumed to be driven by a virtual attractive force of its destination and repulsive forces of its imminent neighbours. In particular, it features a module to detect imminent neighbours, reducing computation overheads and eliminating redundant robot movements. Moreover, in comparison with other methods that are mostly based on a simple function for virtual force calculation, this approach adopts a more adaptive function to calculate repulsive forces to improve collision avoidance.

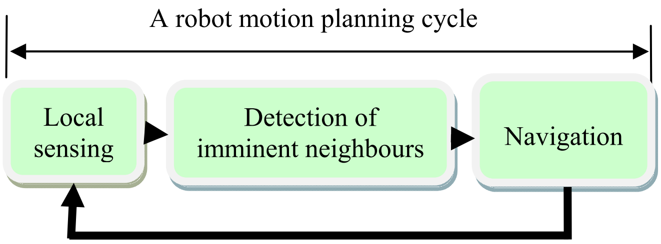

Fig. 6. Three stages in a motion planning cycle

B. Implementation of the Simulator

The simulator, which consists of task allocation and motion planning for a team of mobile robots, is developed in the Player/Stage [30] and C++ programming language. The Player/Stage is an open-source package widely used for multi-robot control and simulation. It runs in a Linux-based operating system called Fedora 13, and consists of two sub-packages, namely Player, and Stage. Player provides a network interface to a variety of physical robots and sensors. Player's client/server model allows robot control programs to be written in a number of programming languages and to run on any computer with a network connection to physical robots. Control program communicates with Player over TCP sockets, reads data from sensors, and writes commands to actuators. Stage is a Player plug-in simulation package which simulates a population of mobile robots moving and sensing in a 2D bit-mapped environment. Various sensor models are provided, including sonars, laser rangefinders, pan-tilt-zoom cameras and odometers. Virtual devices of Stage present a standard Player interface, and hence few or no changes are required to move between simulation and hardware. Controllers designed in Stage have been demonstrated to work on various physical robots.

C. Setup of the Scenario

An automated guided vehicle (AGV), with autonomous control and sensing devices, can be regarded as an autonomous robot. A team of AGVs at a container terminal transporting containers from the quay-side to the yard-side is used to verify the practicability of the proposed task allocation approach, as shown in Fig. 7. There are two vessels berthed at the quay-side. Each vessel is served by five quay cranes which unload the containers from the vessels. Small rectangles in black represent containers. Containers beside the vessels are ready to be picked up, while those being handled by the quay cranes are not shown in the figure. Racks at the quay-side are labelled as 1, 2, …, 10, while racks at the yard-side are labelled as A, B, …, J. The AGVs are transporting containers from the quay-side to the yard-side.

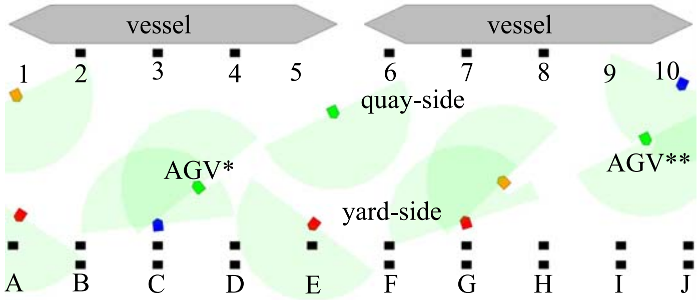

Fig. 7. Simulated working environment of free-range AGVs at a container terminal

There are two possibilities for the transfer of a container from a quay crane to an AGV. The first possibility is that the quay crane places a container directly onto the AGV. The second one, which we adopt in this paper, is that the quay crane places a container onto a buffer rack, from which an AGV will later picks up the container and transports it to the yard side [31].

Traditionally, most AGVs use fixed guide-paths, such as loops, and networks. The fixed routing approaches allow for reliable automation of vehicles. Such AGVs are however less manoeuvrable. Routes are unnecessarily long, incurring considerable transportation time and low system throughput. Route segments are shared for multiple vehicles, leading to potential congestion and deadlocks. With the advent of more powerful onboard processors and advanced sensors, it is now possible for AGVs to navigate without physical guide-paths. Some experimental systems have indeed been developed [32]. Preliminary simulation results showed that free-range routing was on average 19% shorter than traditional mesh-based routing, and 53% shorter than loop-based routing. Huge potentials are therefore seen for free-range routing to improve transport capacities of AGV systems at container terminals.

The AGVs work in an area of 600m × 150m. Each AGV measures 12m × 4.5m × 1.5m and weighs 25 tonnes. The maximum velocity and the maximum acceleration of an AGV are $V_{\max} = 7m\,s^{-1}$ and $a_{\max} = 1m\,s^{-2}$, respectively. Inertial measurement units (IMU) and sonars are used in this paper. An IMU is a device that utilises measurement systems such as gyroscopes and accelerometers to estimate the relative position, velocity, and acceleration of a vehicle in motion [33]. Sonars are common range sensors in mobile robotics. The general principle is that the system emits sound pulses and picks up the echoes bounced off from objects in range, if any. Knowing the transmission speed of sound in the medium and the time of flight, it is possible to compute the distance. This method is widely used due to the low cost of sensors with adequate performance [34]. The sensing field of view is 180°, and the range of sonar scan R is derived as follows. Consider an extreme case where two AGVs, heading directly towards each other without yaw steering, are braking from the maximum velocity $V_{\max}$ with the maximum deceleration $a_{\max}$. According to kinematic equations: $\frac{R}{2} = V_{\max} t - \frac{1}{2} a_{\max} t^2$, and $t = \frac{V_{\max}}{a_{\max}}$, R is derived as $R = \frac{V_{\max}^2}{a_{\max}}$, that is, R=49m, approximately four times the length of an AGV. With a proper yawing angle, this sensing range can sufficiently safeguard motion safety.

There are two major operational uncertainties for AGVs at container terminals, namely, dynamic task requirements, and uncertain traffic conditions. Dynamic task requirements are mainly due to the variation of vessel arrival time, the handling time of quay cranes, and the characteristic of containers to be transported. Uncertain traffic conditions are mainly due to stochastic interferences between AGVs [35]. It is assumed that each AGV can only carry one container at a time, and obviously a container should only be transported by one AGV. Whenever a container is put onto a rack from a quay crane, it is ready for auction. This is a single-task robot, single-robot task, and instantaneous assignment problem. Hence, this scenario of a team of decentralised free-range AGVs working at a dynamic container terminal is a good test-bed to validate the proposed bid adjustment mechanism for dynamic multi-robot task allocation.

A specific type of tasks is described by the pick-up location and the destination of delivery, as T(n, x), where n specifies the label of the pick-up location at quay-side (n=1, 2, …, 10), and x specifies the label of the destination at yard-side (x= A, B, …, J). For example in Fig. 7, T(3, F) is a type of tasks requiring to transport containers from rack 3 at quay-side to rack F at yard-side. The cost of a task in the simulation is the time consumed to transport and to handle a container, which is in the unit of minutes.

D. *Case Study One*

The first case study involves simulation of ten AGVs to transport 300 containers. Containers with different types of task requests are dynamically released from the quay cranes onto the racks. When a container is up for auction, the central auctioneer announces the specific type of task to the robot team, and all the idle AGVs participate in this auction. Each idle robot calculates and adjusts its bid price based on its track record for the same type of tasks, and submits the bid to the auctioneer before the clearing time. After receiving all the bids from the AGV team, the auctioneer assesses the bid prices and awards the task to the AGV that proposes the lowest bid price. The winning AGV proceeds to pick up the container, transports it to the specified location at the yard side, and subsequently drops it onto the rack. After completing this task, the AGV compares the actual cost with the corresponding bid price to attain an adjustment in term of a reward or a penalty. If the actual cost of the task is smaller than the proposed bid price, the AGV will get a reward with an adjustment of negative value. Similarly, a penalty with an adjustment of positive value will be issued to this AGV if the actual cost is larger than the proposed bid price. This bid adjustment is then logged in the track record for this type of tasks, which serves as a closed-loop mechanism to adjust the bid prices in future auctions. Each AGV maintains an array of track records for the corresponding type of tasks that it has ever completed.

Fig. 8 shows the track records of fifteen tasks of type T(3, F) ever performed by AGV*, which is the green one on the left side of Fig. 7. The time-discounting factor, α, was set to be 0.5. For the first time after AGV* had executed a T(3, F) type of task, the adjustment was a penalty of about two minutes. It means that AGV* under-estimated the cost of the task and submitted a bid price which turned out to be much lower than the actual cost incurred afterwards. In other words, AGV* did not keep its commitment to this task. Being imposed with this penalty, AGV* adjusted the bid price for task type T(3, F). It can be observed that the subsequent 2nd to 9th adjustment values of this task type were within the accuracy of ±1 minute band. Since these adjustment values indicate the discrepancies between the bid prices and the related actual costs, it verifies that, with the closed-loop bid adjustment mechanism in auctions, the discrepancies between the actual costs and the bid prices were effectively minimised. With the improved bid prices, tasks were assigned to the competent robots that proposed more reliable bid prices, accordingly enhancing the overall team performance.

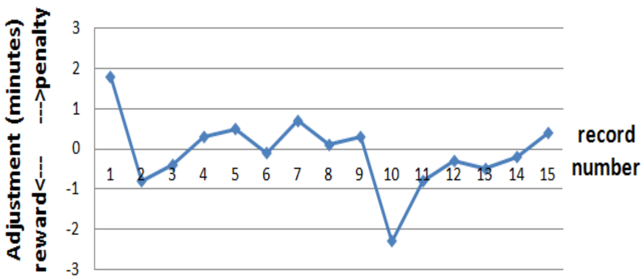

Fig. 8. Adjustment records of task type T(3, F) performed by AGV*

To verify the robustness of the proposed bid adjustment mechanism, the characteristic of task type T(3, F) was deliberately modified after some time in the operation, for example, to transport lighter containers. In this case, an AGV carrying a lighter container should move at a higher speed and be more flexible to evade the obstacles in the way. Hence, the actual cost of task fulfilment should be lower than that before the modification. Nevertheless, the bidding AGVs only received the task information with pick-up and drop-off locations, not knowing the relevant weigh of containers has been changed. The bidding AGVs still offered the previously adjusted bid prices. Therefore, a winning and dispatched AGV, like AGV*, was able to complete the task earlier than expected, and got a reward of about 2.3 minutes. With this reward, the AGV* adjusted the bidding price for task type T(3, F) in the future auctions. It can be noted in Fig. 8 that the subsequent 11th to 15th adjustments were within the accuracy of ±1 minute band again. It shows that, even in some dynamic situations during operation, the closed-loop bid adjustment mechanism can still stably minimise the discrepancies between the bidding prices and the actual costs.

With the closed-loop feedback mechanism, the bid price can be adjusted and fine-tuned to suppress some disturbances due to operational uncertainties. Moreover, a series of adjustment values are averaged with related time-discounting factors to damp out the fluctuations of adjustment values, further enhancing the robustness of the adjustment mechanism. Therefore, stability of the proposed approach can be effectively retained.

Fig. 9 shows a comparison of the overall team performances of ten AGVs to transport 300 containers, with and without the bid adjustment mechanism. It can be seen that the total operational time with bid adjustment is considerably shorter than without. With the proposed bid adjustment mechanism, the bid prices in auctions were adjusted and improved according to the operational conditions, and hence the discrepancies between the bid prices and actual costs could be significantly reduced. Containers were allocated to competent AGVs that submitted more reliable bid prices. As a result, a substantial improvement of 31% in overall team performance, in terms of operational time, was achieved.

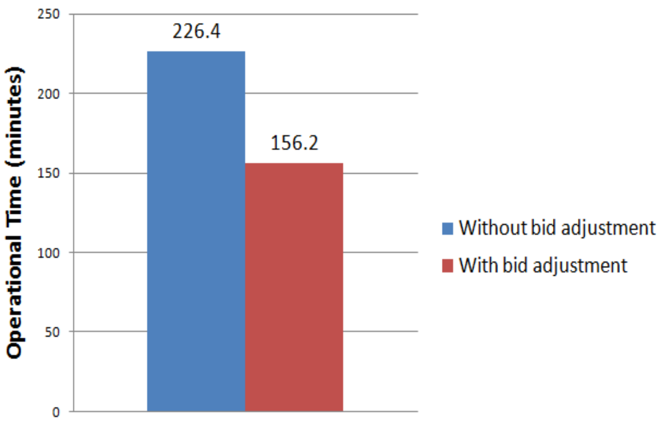

Fig. 9. Comparison of team performances, with and without bid adjustment, by ten AGVs transporting 300 containers

E. Case Study Two

The second case study is carried out to further test and verify the merit of the proposed closed-loop bid adjustment mechanism, which involves simulation of ten AGVs to transport 500 containers.

This time in Fig. 7, the green AGV on the right, which is labelled as AGV**, is taken into account. Fig. 10 shows the track records of twenty-four tasks of type T(7, J) ever performed by AGV**. The time-discounting factor, α, was set to be 0.5. After AGV** had finished a T(7, J) type of task for the first time, the adjustment record was a penalty of nearly three minutes. It indicates that AGV** under-estimated the cost of the task and offered a bid price which turned out to be about three minutes less than the actual cost incurred afterwards. In other words, AGV** could not keep its commitment to this task. Being imposed with this penalty, AGV** modified the bid price for the task type T(7, J) hereafter. We can observe that the subsequent 2nd to 14th adjustment values of this task type were within the accuracy of ±1.2 minute band. Since these adjustment values indicate the discrepancies between the bid prices and the related actual costs, it means that, with the closed-loop bid adjustment mechanism in auctions, the discrepancies between the actual costs and the bid prices were effectively reduced. With the improved bid prices, competent robots that proposed more accurate bid prices were dispatched to the awarded tasks, enhancing the overall team performance accordingly.

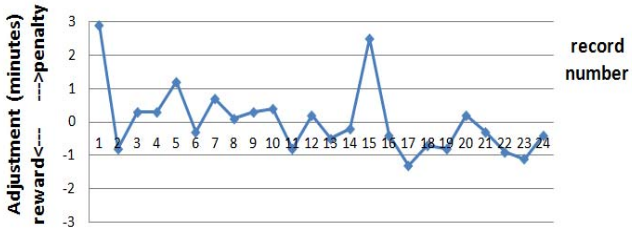

Fig. 10. Adjustment records of task type T(7, J) performed by AGV**

The robustness of the proposed bid adjustment mechanism was also tested by intentionally modifying the characteristic of task type T(7, J) after some time in the operation. This time, heavier containers were released to be transported. In this situation, an AGV carrying a heavier container should generally be more sluggish to evade other AGVs in the way and move at a lower speed. Hence, the actual cost of task completion should be higher than that before the modification. However, the bidding AGVs were only broadcast with the task information with pick-up and drop-off locations, having no idea that the containers are heavier. The idle AGVs involved in auctions still offered the previously adjusted bid prices. A winning and dispatched AGV, like AGV**, could not complete the task as promised, and got a penalty of about 2.5 minutes. Being imposed with this penalty, AGV** adjusted the bidding price for task type T(7, J) in future auctions. We can note in Fig. 10 that the subsequent 16th to 24th adjustment values were within the accuracy of ±1.3 minute band. This verifies that the closed-loop bid adjustment mechanism can stably reduce the discrepancies between the bidding prices and the actual costs, even with some uncertain disturbances during task executions.

Fig. 11 presents a comparison of the overall team performances, in terms of operational time. Simulations of ten AGVs to transport 500 containers, with and without the bid adjustment mechanism, were carried out and compared. It shows that the total operational time with bid adjustment is substantially shorter than without. With the proposed bid adjustment mechanism, the bid prices for different types of tasks in auctions were adjusted and improved according to the dynamic conditions during operation, therefore the discrepancies between the bid prices and the related actual costs could be effectively reduced. Competent AGVs that offered more reliable bid prices were awarded the auctioned containers to pick-up and deliver. A considerable improvement of 21% in overall team performance was achieved.

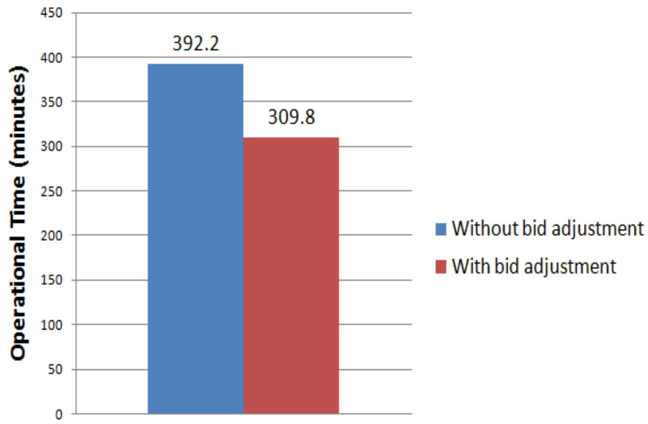

Fig. 11. Comparison of team performances, with and without bid adjustment, by ten AGVs transporting 500 containers

F. Implementation Issues

Two implementation issues are worthy of discussion. The first issue is the cycle processing time of the AGV team. The team of physical AGVs at a real-life container terminal are designed to be decentralised with independent and simultaneous motion planning. However, in computer simulations as demonstrated above, AGVs are inevitably processed sequentially by the computer program as shown in Fig. 12. Hence, the more AGVs involved in the team, the longer it takes to process and update all the AGVs. This would lead to a potentially dangerous situation where some AGVs that have not yet been updated with current information may collide with obstacles. Indeed, to simulate a large number of AGVs, the computer program design for the behaviours of each AGV and the program control logic should be optimised to render its execution as efficient as possible.

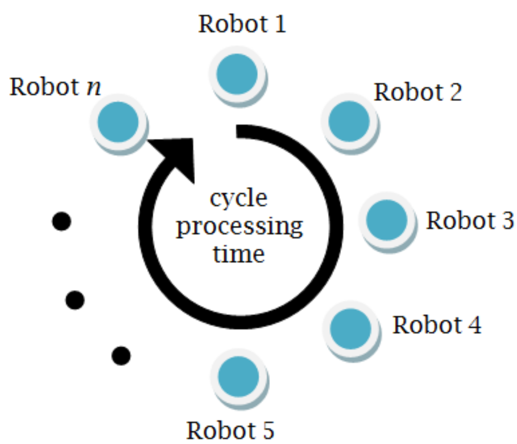

Fig. 12. Cycle processing time in simulation

Another issue is the sensing noises of a sensor. Both the inertial measurement units (IMU) and sonars used in this simulation have measurement noises. Indeed, they resemble the physical sensors in real-life applications. Hence, adequate tolerances should be incorporated accordingly during implementation to make the measurements of relevant sensors effective.

IV. Conclusion and Future Work

This paper has presented an auction-based approach with the proposed closed-loop bid adjustment mechanism to dynamic task allocation in robot teams. The bid adjustment mechanism evaluates and fine-tunes bid prices based on the performance track records of each robot in the team. A simulator has been developed, with case studies of AGVs transporting containers at a terminal, to verify and validate the proposed task allocation approach. Simulation results show that the bid adjustment mechanism can effectively minimise the discrepancies between the submitted bid prices and the corresponding actual costs of tasks. The stability of the approach has also been verified even in light of some operational uncertainties. With the proposed bid adjustment mechanism in task allocation, bid prices can be regulated and fine-tuned, reducing the discrepancies between the bid prices and the actual costs. This enhances the likelihood of allocating tasks to competent robots that are able to submit more accurate bids, and as a result, improves the overall team performance substantially.

Some issues of the auction-based task allocation approach are worthy of further study. Firstly, an allocated-but-not-yet-executed task cannot be re-auctioned even if the dispatched robot is locked in a heavy congestion or even fails. Future work will be devoted to incorporating some other market-based mechanisms, like task trading between robots. For example, it would be preferable if the locked robot can negotiate and trade its task to another robot which is more likely to fulfil the task, according to real-time working conditions. Nevertheless, adopting such a trade-based approach would cost more local communication overheads between robots. Moreover, the overall performance of the team would need further investigation, in comparison with the proposed auction-based approach.

Secondly, simulation results have shown the robustness and stability of the closed-loop bid adjustment mechanism, by introducing some stochastic conditions during operation. On the other hand, there still seems to be some room to improve the bid adjustment process, in terms of the overshot values in Fig. 8 and Fig. 10, as well as the transient fluctuations. These phenomena are inevitably caused by the stochastic operational conditions even though suppressed by our approach. It would be more fruitful if the overshot values and the transient fluctuations could be further alleviated so that the discrepancies between the bid prices and the actual costs could be accordingly reduced. Some techniques of learning theory and adaptive regulation are now under active consideration, although their responsiveness to dynamic situations is a concern.

REFERENCES

[1] A. Richards, J. Bellingham, M. Tillerson, and J. How, "Co-ordination and control of multiple UAVs," in: *Proceedings of AIAA Conference on Guidance, Navigation, and Control.* California, pp.251-256, 2002.

[2] R. M. Murray, "Recent research in cooperative control of multivehicle systems", *Journal of Dynamic Systems, Measurement and Control*, vol.129, no.5, pp. 571-583, 2007.

[3] L. E. Parker, "Path planning and motion coordination in multiple mobile robot teams". *Encyclopedia of Complexity and Systems Science 2009*, pp.5783-5800, 2009.

[4] Parker, L.E.: Path Planning and Motion Coordination in Multiple Mobile Robot Teams. In: Meyers, R., editor-in-chief: Encyclopedia of Complexity and System Science, Springer, The Netherlands, 2009.

[5] R. Stahlbock, and S. Vob, Vehicle Routing Problems and Container Terminal Operations: An Update of Research. In: B. Golden, et al.: *The Vehicle Routing Problem.* Springer, The Netherland, pp. 551-589, 2008.

[6] BBC News on snow removal by a team of trucks at Philadelphia Airport: http://www.bbc.co.uk/news/world-us-canada-12081267

[7] A. Sud, R. Gayle, E. Andersen, S. Guy, M. Lin, and D. Manocha, "Real-time Navigation of Independent Agents Using Adaptive Roadmaps", In: A. Majumder, L.F. Hodges, D. Cohen-Or, and S.N. Spencer, editors: *The ACM Virtual Reality Software and Technology 2007,* Newport Beach, California, USA, pp.99-106, 2007.

[8] L. E. Parker, Chapter 40: Multiple mobile robot systems. in: S. Bruno, and K. Oussama, editors. *Springer Handbook of Robotics. Springer*, pp. 921-941, 2008.

[9] M. B. Dias, R. Zlot, N. Kalra, and A. Stentz, "Market-based multi-robot coordination: a survey and analysis". *Proceedings of the IEEE,* vol.94, no.7, pp.1257-1270, 2006.

[10] B. P. Gerkey, and M. J. Mataric, "A formal analysis and taxonomy of task allocation in multi-robot systems". *The International Journal of Robotics Research*, vol.23, no.9, pp.939-954, 2004.

[11] K. Lerman, C. Jones, A. Galstyan, and M. J. Mataric, "Analysis of dynamic task allocation in multi-robot systems". *The International Journal of Robotics Research,* vol.25, no.3, pp.225-241, 2006.

[12] A.R. Mosteo, and L. Montano, "Comparative Experiments on Optimization Criteria and Algorithms for Auction based Multi-robot Task Allocation", In: *Proceedings of IEEE International Conference on Robotics and Automation 2007,* Roma, Italy, pp. 3345-3350, April 2007.

[13] E.G. Jones, M.B. Dias, and A. Stentz, "Learning-enhanced Market-based Task Allocation for Oversubscribed Domains", In: *Proceedings of the 2007 IEEE/RSJ International Conference on Intelligent Robots and Systems,* San Diego, CA, USA, pp.2308-2313, 2007.

[14] A. Schoenig, and M. Pagnucco, "Evaluating sequential single-item auctions for dynamic task allocation", in: *Lecture Notes in Computer Science*, Vol.6464, pp.506-515, 2010.

[15] R. G. Smith, "The contract net protocol: high-level communication and control in a distributed problem solver", *IEEE Transactions on Computing,* vol.29, no.12, pp.1104-1113, 1980.

[16] E. G. Jones, M. B. Dias, and A. Stentz, "Time-extended multi-robot coordination for domains with intra-path constraints", *Autonomous Robots*, vol.30, no.4, pp. 41–56, 2011.

[17] D. Herrero-P´erez, and H. Mat´ınez-Barber´a, "Decentralized Coordination of Autonomous AGVs in Flexible Manufacturing Systems", *IEEE/RSJ International Conference on Intelligent Robots and Systems 2008*, Nice, France, pp. 3674-3679, September 2008.

[18] J. Schneider, D. Apfelbaum, D. Bagnell, and R. Simmons, "Learning opportunity costs in multi-robot market based planners", in: *Proceedings of the 2005 IEEE International Conference on Robotics and Automation,* Spain, pp.1151-1156, 2005.

[19] N. Kalra, D. Ferguson, and A. Stentz, "Hoplites: a market-based framework for planned tight coordination in multirobot teams". in: *Proceedings of IEEE International Conference on Robotics and Automation,* Spain, pp.1170-1177, 2005.

[20] L. F. Bertuccelli, H.-L. Choi, P. Cho, and J. P. How, "Real-time multi-UAV task assignment in dynamic and uncertain environments". in: *Proceedings of The American Institute of Aeronautics and Astronautics Conference on Guidance, Navigation, and Control 2009*, Chicago, USA. pp. 389-395, 2009.

[21] H.-L. Choi, L. Brunet, and J. P. How, "Consensus-based decentralized auctions for robust task allocation", *IEEE Transactions on Robotics,* vol. 25, no.4, pp.912-926, 2009.

[22] E.G. Jones, M.B. Dias, and A. Stentz, "Time-extended multi-robot coordination for domains with intra-path constraints", *Autonomous Robots,* vol.30, no.1, pp.41-56, 2011.

[23] B.P. Gerkey, and M.J. Mataric, "Sold!: Auction Methods for Multirobot Coordination", *IEEE Transactions on Robotics and Automation,* vol. 18, no.5, pp.758-768, 2002.

[24] M. B. Dias, B. Ghanem, and A. Stentz, "Improving cost estimation in market-based coordination of a distributed sensing task", in: *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS),* Alberta, Canada, pp.3972–3977, 2005.

[25] F. Duvallet, and A. Stentz, "Imitation learning for task allocation", in: *Proceedings of The 2010 IEEE/RSJ International Conference on* Intelligent Robots and Systems, Taipei, Taiwan, pp.3568-3573, 2010.

[26] M.T. Khan, T. Imanuel, and C.W. de Silva, "Autonomous market-based multi-robot cooperation", In: *International Conference on Intelligent and Advanced Systems (ICIAS 2010),* Kuala Lumpur, Malaysia, pp.1-6. 2010.

[27] W.K. Zhu, and S.H. Choi, "An Auction-based Approach with Closed-loop Bid Adjustment to Dynamic Task Allocation in Robot Teams", Lecture Notes in Engineering and Computer Science: *Proceedings of The World Congress on Engineering 2011*, WCE 2011, 6-8 July, 2011, London, U.K., pp. 1016-1066.

[28] H. Choset, K.M. Lynch, S. Hutchinson, G. Kantor, W. Burgard, L.E. Kavraki, and S. Thrun, Principles of Robot Motion: Theory, Algorithms, and Implementations. MIT Press, Massachusetts, USA, 2005.

[29] S.H. Choi, and W.K. Zhu, "A bio-inspired intelligent approach to motion planning of mobile robots", *Computer-Aided Design and Applications*. vol.8, no.5, pp.773-783, 2011.

[30] Player/Stage Project, http://playerstage.sourceforge.net/

[31] Gottwald Port Technology. http://www.gottwald.com/gottwald/site/gottwald/en/products/agv.html

[32] M. B. Duinkerken, M. van der Zee, and G. Lodewijks, "Dynamic free range routing for automated guided vehicles", in: *Proceedings of IEEE International Conference on Networking, Sensing, and Control,* pp. 312-317, 2006.

[33] G. Dudek, M. Jenkin, Chapter 20: Inertial Sensors, GPS, and Odometry. In: S. Bruno, K. Oussama, editors: *Springer Handbook of Robotics,* Springer, The Netherlands, pp.477-490, 2008.

[34] L. Kleeman, and R. Kuc, Chapter 21 Sonar Sensing. In: S. Bruno, K. Oussama, editors: *Springer Handbook of Robotics*, Springer, The Netherlands, pp.491-519, 2008.

[35] I. F. A. Vis, "Survey of research in the design and control of automated guided vehicle systems". *European Journal of Operational Research,* vol.170, no.5, pp.677-709, 2006.