Evaluating the Impact of Perception and Decision Timing on Autonomous Robotic Exploration

Francesco Amigoni*, Alberto Quattrini Li*, Dirk Holz†
*Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Milano, Italy
†Autonomous Intelligent Systems Group, University of Bonn, Bonn, Germany

***Abstract*— Autonomous robotic exploration of initially unknown environments is at the basis of several applications, including map building and search and rescue. Despite the many recent works on robotic exploration, an issue that has not been adequately addressed in the literature so far is the evaluation of the impact of the perception (for map update) and decision (about where to go next) timing on the behavior of an exploring robotic system. In this paper, we contribute to fill this gap by providing a *quantitative* experimental analysis of how frequencies of perception and decision influence the performance of an exploring mobile robot. Results, obtained with an experimental simulation framework (implemented and made publicly available) based on ROS and Stage, confirm the intuitive idea that the best performance is obtained with fast-paced perceptions and decisions, but also suggest some trade-offs for the values of perception and decision frequencies in some settings.**

I. Introduction

Mobile robots that autonomously explore initially unknown environments are relevant for many applications, from map building [1] to search and rescue [2]. Arguably, two of the most important aspects of exploration are the integration of perceived data into the current map of the environment and the decision about where to move next in a partially known environment. These two activities are performed by robots either in an event-based fashion (e.g., when a destination location is reached) or in a frequency-based fashion (i.e., after a fixed amount of time). Usually, papers consider only a single combination of perception and decision modalities. This makes it difficult to assess the impact of the perception and decision timing on the performance of an exploring robotic system.

In this paper, we provide a *quantitative* analysis of the impact of timing of perception and decision on the performance of an exploring mobile robot. We consider these parameters as particularly important for autonomous exploration because they are mostly related to the robot control software, while other parameters, like robot speed, sensor speed, and sensor range are, although controllable to some degree, more related to the hardware equipment. The motivation of our work is to provide designers with insights on perception and decision frequencies to develop better exploring robotic systems.

We consider a single mobile robot equipped with a laser range scanner to discover the physical structure of an initially unknown environment. The basic exploration process goes as follows: (a) the robot perceives the surrounding environment, (b) it integrates the perceived data into a map representing the environment known so far, (c) it decides where to go next, and (d) it moves to the selected location and starts again from (a). We focus on the timing at which the activities (b) and (c) are performed. More precisely, we consider *perception* as the process of acquiring sensor data and integrating them into the current map and *decision* as the process of selecting the next destination location. Note that the timing of perception that we consider is exclusively related to map building. We do not consider the timing of acquisition of sensor data that are used to localize the robot, but the timing at which the map (the primary source of information for decision) is updated.

In this paper, whose nature is mainly experimental, we use standard platforms for the development of the robotic system (ROS [3]) and for its simulation in different environments (Stage [4]), according to the principles of emerging good experimental methodologies [5].

This paper is structured as follows: after a discussion of related work in Section II, we present our framework for evaluating the performance of exploration under different perception and decision timings in Section III. Our experimental methodology, the conducted experiments, and the obtained results are described in Section IV. Finally, Section V concludes the paper.

II. Related Work

Robotic exploration of initially unknown environments has been addressed in recent years by many works. The dominant approach is greedy [6], also called *Next Best View (NBV)*, and operates by cyclically evaluating some candidate destination locations and by assigning the best ones to robots. Candidate destination locations are usually selected on the *frontiers* between the known and the unknown portions of the environment [7]. They are valuated according to an *utility function* (exploration strategy) that considers some criteria like the distance of a location from the robot(s) and the expected amount of information that a robot can acquire from a location [8]. However, the impact of perception and decision timing on the exploration performance has not been considered and investigated so far. Explicit information about when perceptions and decisions are performed are rarely reported in papers. According to the available data, we can broadly classify the published works according to the following classes.

Event-based vs. frequency-based perception. In case of *event-based perceptions*, only the data acquired at the destination location is integrated into the current map of the environment. This means that a robot "blindly" navigates from its initial location to the destination location moving along a path inside the known free space, without updating the map. In *frequency-based perception*, data is continuously acquired and integrated into the map, at some fixed frequency, as the robot navigates toward the destination location. This frequency can depend either on the sensor acquisition frequency or, more often, on the update frequency of the map building method employed (e.g., the map is updated after a certain distance, after a certain rotation angle, or after a certain amount of time). Note that these timings are related only to map updates for exploration purposes. Indeed, depending on the methods used for map building and localization, data acquired during navigation can be used for localization and path following, but we do not consider the timing of these acquisitions if they do not enrich the map built by exploration.

Event-based vs. frequency-based decision. In the first case, a decision about the next destination location is made only when the previous destination location has been reached (or after that location has been declared unreachable, for example after a timeout expired). In the second case, decisions about the next destination location are continuously made, at some fixed frequency (for example, in [9] a new decision is basically made every 4 seconds), while the robot is moving toward the current destination location (implementing a sort of opportunistic behavior).

Table I classifies some of the most significant papers on robotic exploration according to the above dimensions. The aim of the table is not to exhaustively classify all works in the area, but only to provide an overview based on a significant sample of published papers. We mainly focus on works employing a single robot, but our considerations hold also for multi-robot exploration (e.g., [10]–[12]).

TABLE I: Perception and decision modalities employed by some papers

		perception	
		event-based	frequency-based
decision	event-based	[7], [11], [13]–[19]	[10], [12], [20]
	frequency-based	-	[9], [21], [22]

Most works consider event-based perception and decision. This is not surprising because of ease of implementation of such configuration. Although they use different exploration strategies and different map representations, these works basically consider reaching the current destination location as the event that triggers the acquisition of new data from sensors and their incremental integration into the current map as well as the decision making about the next destination location. Relatively less works use frequency-based perception and decision. Unsurprisingly, no work considers event-based perception and frequency-based decision. This is expected since it makes little sense to decide repeatedly about the next destination location on the basis of the same information about the environment.

Some recent works have partially addressed the study of the effects of perception and decision timing on exploration performance. For instance, [20] explicitly recognizes one of the main problems of event-based decision under frequency-based perception: since the map is continuously updated according to data coming from long range sensors, the robot might have fully explored a region (e.g., a dead-end corridor) before actually reaching the selected destination location. In [20], a heuristic is proposed to avoid to reach such destination locations if there is nothing the robot could discover there. This amounts to discard the old decision about the destination location. This approach is different from that evaluated in this paper, in which decisions are regularly revised at a given frequency, allowing the current destination location to be discarded and to opportunistically reach a more promising destination location.

A similar problem is recognized in [23] and some (fast) frontier detection algorithms are proposed as a support for making decisions about destination locations at a high frequency, up to $10\,\mathrm{Hz}$ (as it can be deduced from experiments).

Beyond these partial attempts, and to the best of our knowledge, a systematic analysis of the impact of perception and decision timing on the performance of exploration is still missing. In this paper, we aim at contributing to fill this gap.

III. Experimental Framework

We assume to have a mobile robot equipped with a $180°$ laser range scanner. Exploration is performed as a sequence of movements to destination locations, selected according to the exploration strategy. Decisions about destination locations can be made either in an event-based or in a frequency-based fashion, in this last case with a frequency f_d that can be set by the designer. During navigation, the robot perceives the surrounding environment either in an event-based or in a frequency-based fashion, also in this last case with a configurable frequency f_p.

Each perceived sensor reading is used to update a global map, that represents the environment. The global map is a finite two-dimensional grid, whose cells are identical squares, at a resolution of $0.2\,\mathrm{m}$. Each cell can be unknown, free, or occupied. Free and occupied cells are considered known.

Given a map represented as above, to calculate candidate destination locations, we consider reachable free cells that are on the boundary between known and unknown cells. Then, a set of 8-adjacent boundary cells are grouped in a cluster, called frontier. The centroid of each cluster is considered as a *candidate destination location* to reach. Each candidate destination location is represented as a cell position in the grid together with the orientation that the robot should take when the location is reached. The orientation is toward the unknown area along the perpendicular to the line tangent to the corresponding frontier and passing through the candidate destination location cell.

The utility $u(p, r)$ of a candidate destination location p for a robot r is evaluated according to two exploration strategies taken from the literature, which combine the following criteria in their utility functions:

- $A(p)$ is the expected information gain at p, computed according to the frontier size (i.e., the number of cells belonging to the cluster of p scaled with the cell resolution, to have the length of the frontier);

- $d(p, r)$ is the distance between p and current position of r; given p and r, this criterion is calculated using a path planner, based on Dijkstra algorithm, on the grid map that returns the length of the path;
- $o(p, r)$ is the cost related to the heading change that the robot should perform, computed according to the difference between the orientation required at p and the current orientation of r.

The first exploration strategy is based on a *weighted average* of the individual criteria (as for example in [10]):

$$u(p, r) = w_A A(p) - w_d d(p, r) - w_o o(p, r) \quad (1)$$

where w_j indicates the weight associated to criterion j.

The second exploration strategy is called *MCDM strategy* and combines the criteria of the set $N = \{A, d, o\}$ using the Multi-Criteria Decision Making (MCDM) approach. Please refer to [8] for a complete description. Here we just sketch the approach and its main parameters. We call $u_j(p, r)$, with $j \in N$, the utility value for candidate destination location p and robot r according to criterion j. Utilities are normalized to a common scale $I = [0, 1]$, using a linear relative normalization. Note that the larger $u_j(p, r)$, the better the location p for robot r.

Basically, the MCDM strategy replaces function (1) with:

$$u(p, r) = \sum_{j \in N} (u_{(j)}(p, r) - u_{(j-1)}(p, r))\mu(\mathscr{A}_{(j)}), \quad (2)$$

where $\mu : \mathcal{P}(N) \rightarrow [0, 1]$ ($\mathcal{P}(N)$ is the power set of set N) is a normalized *fuzzy measure* on the set of criteria N that associates a weight to each group of criteria. Slightly overloading the notation, $u_{(j)}$, with $(j) \in N$, indicates the j-th criterion according to an increasing ordering with respect to utilities, i.e., after the n criteria have been ordered to have, for candidate location p and robot r, $u_{(1)}(p, r) \leq \ldots \leq u_{(n)}(p, r) \leq 1$. It is assumed that $u_{(0)}(p, r) = 0$. Finally, the set $\mathscr{A}_{(j)}$ is defined as $\mathscr{A}_{(j)} = \{i \in N | u_{(j)}(p, r) \leq u_i(p, r) \leq u_{(n)}(p, r)\}$. Using (2) amounts to perform a sort of “distorted” weighted average of the utilities of criteria and is a more principled way than (1) to compute utilities, because it allows to consider importance of criteria and their mutual dependency relations. We consider these two exploration strategies because weighted average is less computationally expensive than MCDM. In this way, we can evaluate if the computational cost of making decisions has an impact on the exploration performance, given perception and decision timing.

In order to perform repeated tests under controlled conditions, we developed a ROS [3] package (publicly available at `http://sourceforge.net/projects/explorationeval`) for experimentally evaluating exploration strategies with different perception and decision timing using the Stage simulator [4]. The performed experiments considered simulated robots with realistically noisy odometry that affects both movement and sensing capabilities. Our package mainly depends on the following other ROS packages (some of which have been adapted or extended):

- *explore*. It implements frontier-based exploration (we added the MCDM exploration strategy to the default implemented strategy, which is weighted average). Next best frontier is selected at frequency f_d or when the current frontier is reached.
- *move_base*. It implements the action of movement to a destination location by following the trajectory returned by the planner.
- *gmapping*. It provides laser-based SLAM (Simultaneous Localization and Mapping) using a grid map. Map of the environment is updated when a perception is acquired (at frequency f_p or when the current frontier is reached).
- *costmap_2d*. It implements a two-dimensional costmap which takes in sensor data from the world, builds an occupancy grid from the data, and assigns costs to cells.
- *stageros*. It implements two-dimensional robot simulation using Stage.

IV. Experimental Results

For our experimental evaluation, we consider three indoor environments: *maze*, *fort*, and *open* (Fig. 1, the unit in the figure is 3 m), which are all publicly available as part of the ROS *bosch_common* package, of the *fort ap_hill_07b*, and of the *acapulco_convention_centre* data sets at Radish [24], respectively. The three environments present several challenges for exploration, including dead-ended corridors and intersections where the best decision about where to go next is not obvious if the environment is only partially known. Moreover, two environments are rather structured while the last one presents an empty large area.

The robotic platform used is a Segway-RMP robot equipped with a SICK LMS200 laser range scanner, with a maximum range of 8 m and angular resolution at 1°. We use the weights reported in the following tables for weighted average (left) and MCDM (right) exploration strategies, which have been set, after some preliminary experiments, in order to obtain good performance (we also experimentally verified that slightly different values provide similar performance).

criteria	$w_{()}$
A	1.0
d	0.005
o	0.0

criteria	$\mu()$	criteria	$\mu()$
A	0.5	A, d	0.9
d	0.3	A, o	0.7
o	0.1	d, o	0.4

We use six values for perception/decision frequencies f_p and f_d, namely 0.2, 0.4, 0.6, 1.0, 2.0, and 4.0 Hz. (smaller values make exploration too slow and larger values require too computational effort). We also use the combination of event-based perception and event-based decision. We refer to a combination of exploration strategy and perception/decision modalities as *setting*. For each environment, and for 10 randomly selected initial robot poses (shown by arrows in Fig. 1), we performed 5 runs per setting. For the runs correctly terminated, we measured the average (over initial locations and runs, namely over 50 values) travelled

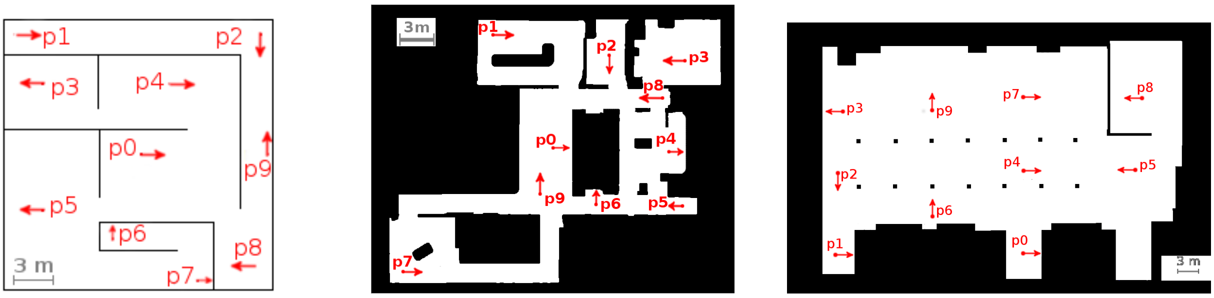

Fig. 1: Maze (left), fort (center), and open (right) environments

distance (in m) and the *time* required (in s) to map 90% of the free area of the environments. This termination criterion is of interest, for example, in rescue applications, for which knowing the general structure of an environment is more important than exploring completely the few last posts [22].

Figs. 2, 3, and 4 show results[1] for the maze, fort, and open environment, respectively. For each environment, the two leftmost graphs are relative to weighted average exploration strategy, while the two rightmost graphs are relative to MCDM. Moreover, for each environment and exploration strategy, the left-hand graph shows the average travelled distance with respect to the perception and decision frequencies, while the right-hand graph shows the average time to complete the exploration with respect to perception and decision frequencies. For ease of reading we discretize the values of travelled distance and exploration time in five bins, with the darker the better.

Some interesting trends emerge from the above results and are discussed below.

Distance vs. perception frequency. Given a decision frequency, there is an optimal interval of perception frequencies with respect to travelled distance. Increasing the number of perceptions in a given time interval reduces the travelled distance because the robot sees the environments at a higher pace. For example, in the maze environment, with MCDM and $f_d = 1.0$ Hz, average distance changes from 151.0 m to 113.8 m when f_p changes from 0.2 Hz to 1.0 Hz. This difference is statistically significant (p-value=$2.37 \cdot 10^{-11}$) according to an ANOVA analysis with a threshold for significance p-value < 0.05 [25]. In all the environments, the reduction of travelled distance tends to reach a plateau when the perception frequency grows, suggesting that there is some "optimal" travelled distance for an environment [26]. For example, in the open environment, with MCDM and $f_d = 1.0$ Hz, average distance changes from 299.8 m to 293.3 m when f_p changes from 2.0 Hz to 4.0 Hz, a difference that is not statistically significant (p-value=0.77).

Distance vs. decision frequency. Given a perception frequency, the travelled distance generally decreases when the decision frequency increases. For example, in the maze environment, with MCDM and $f_p = 1.0$ Hz, average distance changes from 123.3 m to 113.8 m when f_d changes from 0.2 Hz to 1.0 Hz (p-value=0.044). In a way, this is an expected behavior, because the more frequently decisions are revised, the better their outcome. Moreover, at a deeper level of analysis, this result also shows that the robot have not any "schizophrenic" behavior, namely it does not change destination location every time a new decision is taken, at least in the maze and fort environments. Indeed, if that was the case, the distance would have shown an increase with growing decision frequency. In other words, the two exploration strategies we tested have the nice property of estimating accurately the goodness of a destination location in the maze and fort environments and this estimate is usually not changed if more data about the environment is collected. In the open environment and considering MCDM (Fig. 4, third graph from left), the graph shows that there is a region in which increasing decision frequency can worsen the travelled distance, even if this is not statistically significant. For example, with $f_p = 2.0$ Hz, average distance changes from 263.9 m to 297.6 m when f_d changes from 2.0 Hz to 4.0 Hz (p-value=0.93).

Time vs. perception frequency. When the perception frequency increases, the robot collects data about the environment at a higher rate and the exploration time decreases. For example, in the maze environment, with weighted average and $f_d = 1.0$ Hz, average time changes from 306.8 s to 233.9 s when f_p changes from 0.2 Hz to 1.0 Hz (p-value=$1.34 \cdot 10^{-34}$). However, in some settings, when the perception frequency becomes too high, the robot spends a significant amount of time in updating the map and the exploration time slightly increases even if not significantly. For example, in the open environment, with weighted average and $f_d = 2.0$ Hz, average time changes from 804.2 s to 966.2 s when f_p changes from 1.0 Hz to 4.0 Hz (p-value=0.002).

Time vs. decision frequency. A similar behavior is encountered when looking at the impact of the decision frequency on the exploration time. With a very high decision frequency, the robot spends time in decision making (i.e., in evaluating the candidate destination locations) and the exploration time tends to increase with MCDM. This is more evident in the fort environment, which is more complicated than the maze environment, in which the exploration path is almost fixed. For example, in fort environment, with weighted average and $f_p = 2.0$ Hz, when f_d changes from 1.0 Hz to 4.0 Hz average time changes from 326.6 s

[1]Complete results are available at `http://sourceforge.net/projects/explorationeval`.

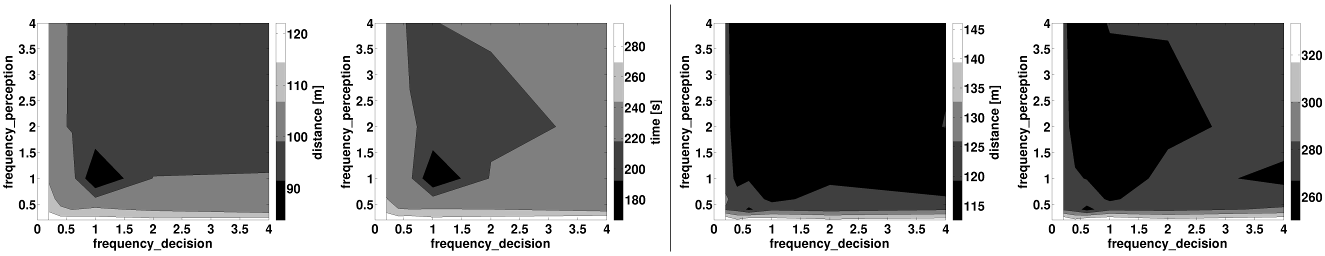

Fig. 2: Results for the maze environment, using weighted average (two leftmost graphs) and MCDM (two rightmost graphs)

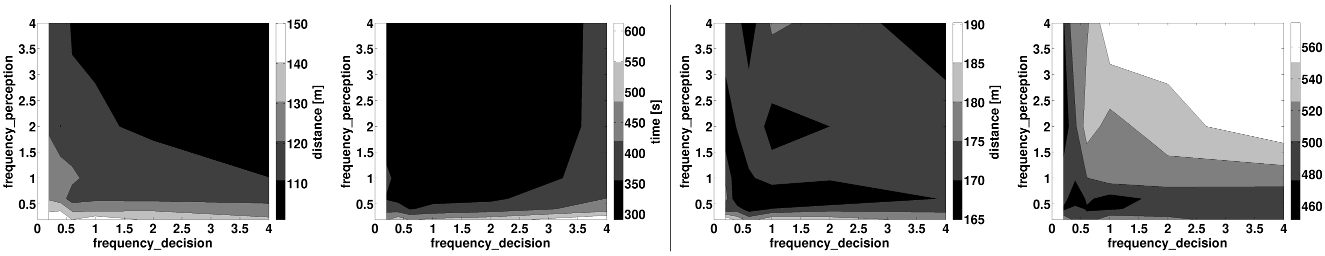

Fig. 3: Results for the fort environment, using weighted average (two leftmost graphs) and MCDM (two rightmost graphs)

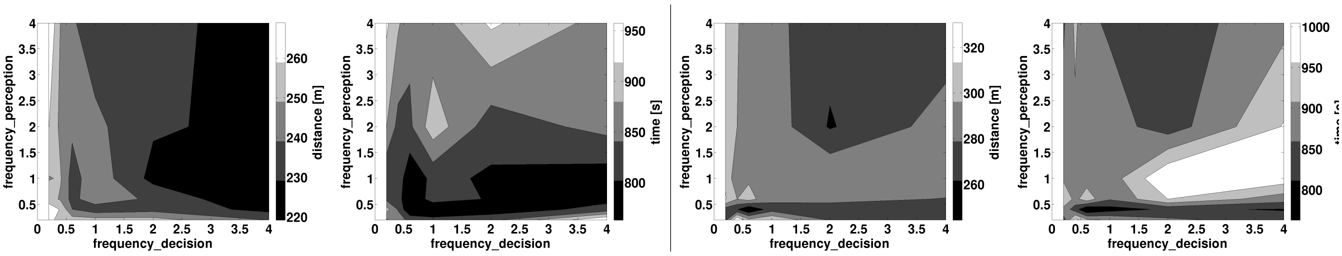

Fig. 4: Results for the open environment, using weighted average (two leftmost graphs) and MCDM (two rightmost graphs)

to $385.4\,\text{s}$ (p-value=$8.06 \cdot 10^{-4}$), whereas, in maze environment, it changes from $225.8\,\text{s}$ to $234.8\,\text{s}$ (p-value=0.034). With weighted average, instead, the increase of the decision frequency does not seem to affect too much the exploration time. This could be explained by the fact that MCDM exploration strategy requires more computational effort than weighted average exploration strategy. In the open environment, exploration time has different trends according to different perception frequencies for both exploration strategies. With weighted average, it seems that increasing the decision frequency does not affect too much the exploration time. Instead, with MCDM, for low perception frequencies, high decision frequency can worsen exploration time, while for high perception frequencies, settings with high decision frequencies obtains similar results to those with low decision frequencies. These results in the open environment suggest that decision making in unstructured environments needs an updated map to reliably choose a candidate location.

Distance and time vs. perception and decision frequencies. Looking at the impact of combined frequencies on travelled distance and time, it emerges that, in general, there is an interval of frequencies that guarantee the best performance. However, there are some differences related to exploration strategies and environments, as discussed below.

Weighted average and MCDM exploration strategies vs. environments. With weighted average, the travelled distance and the exploration time tend to reach a wide plateau, at almost the optimal "height", when the perception and decision frequencies increase. With MCDM, both the optimal travelled distance and the exploration time are obtained for a narrower interval of frequency values, in the fort and open environment. For example, in open environment, an optimal combination of decision and perception frequencies seems to be $f_p = 0.4$ Hz and $f_d = 0.6$ Hz, with average distance of $244.3\,\text{m}$ and average time of $763.0\,\text{s}$. For small changes in the values of the frequencies, results worsen, especially regarding exploration time. This can be explained by the higher computational effort required by MCDM, degrading the performance when decision frequency increases too much. In the open environment, the optimal travelled distance and the exploration time are in a narrower interval of frequency values compared to the case of maze and fort environments, even for the weighted average. This could be explained by considering that, in structured environments, the movements of the robot are "forced" by the presence of walls, lessening the impact of chosen frequencies values, while, in unstructured environments, the robot can possibly go in any direction and so carefully setting frequencies of perception and decision is more critical.

Finally, we present some results about event-based perception and decision. Fig. 5 shows results relative to the maze environment for event-based perception and decision, for frequency-based perception ($f_p = 5.0$ Hz) and event-based decision, and for frequency-based perception and decision ($f_p = 1.0$ Hz and $f_d = 1.0$ Hz). The left-hand graph shows the average travelled distance, while the right-hand graph shows the average exploration time for each of the above combinations.

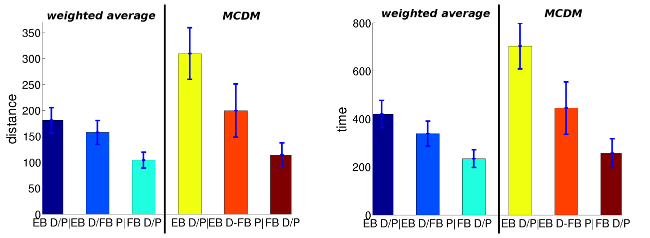

Fig. 5: Results for the maze environment (EB = Event-Based, FB = Frequency-Based, P = Perception, D = Decision)

Results confirm, as expected, that better performance is obtained with frequency-based perception and decision. For example, for MCDM, changing from event-based to frequency-based perception and decision, average distance changes from 309.4 m to 113.8 m (p-value=$4.65 \cdot 10^{-44}$) and average time changes from 703.0 s to 256.6 s (p-value=$5.51 \cdot 10^{-48}$). Note that MCDM performs worse than weighted average. This could be explained since the maze environment is rather simple and does not require any complex exploration strategy.

V. Conclusion

In this paper, we addressed the quantitative analysis of the impact of perception and decision timing on the performance of an exploring mobile robot. Our results aim at supporting the design of better robot systems for the exploration of unknown environments, by providing the designers of robot control systems some experimental evidence about trade-offs between perception and decision frequencies. Results obtained in our experimental settings show that, as expected, frequency-based perception and decision outperform event-based perception and decision. Moreover, although increasing perception and decision frequencies generally increases performance, when these frequencies become too high, performance starts to degrade due to increased computational effort. In unstructured environments, properly setting perception and decision frequencies seems to be more important than in structured environments, as decisions should be made with updated information and with right timing because the environment is not "driving" the robot.

The work presented in this paper enlightens some aspects of robotic exploration that have not received much attention so far; however, our results are not exhaustive nor definitive. For example, quality of the resulting map and the amount of area mapped over time could be included in the adopted metrics. Moreover, other environments, including outdoor areas, different map building methods (alternative to *gmapping*), and multiple robots should be considered in further experiments.

Acknowledgment

Authors gratefully acknowledge contributions of Daniele Tortorella and Simone Pignatelli to experimental activities.

References

[1] S. Thrun, "Robotic mapping: A survey," in *Exploring Artificial Intelligence in the New Millenium*, 2002, pp. 1–35.
[2] S. Tadokoro, *Rescue Robotics*. Springer-Verlag, 2010.
[3] M. Quigley, B. Gerkey, K. Conley, J. Faust, T. Foote, J. Leibs, E. Berger, R. Wheeler, and A. Ng, "ROS: an open-source robot operating system," in *ICRA Workshop on Open Source Software*, 2009.
[4] R. Vaughan, "Massively multiple robot simulations in stage," *Swarm Intelligence*, vol. 2, no. 2-4, pp. 189–208, 2008.
[5] F. Amigoni, M. Reggiani, and V. Schiaffonati, "An insightful comparison between experiments in mobile robotics and in science," *Auton Robot*, vol. 27, no. 4, pp. 313–325, 2009.
[6] C. Tovey and S. Koenig, "Improved analysis of greedy mapping," in *Proc. IROS*, 2003, pp. 3371–3378.
[7] B. Yamauchi, "A frontier-based approach for autonomous exploration," in *Proc. CIRA*, 1997, pp. 146–151.
[8] N. Basilico and F. Amigoni, "Exploration strategies based on multi-criteria decision making for searching environments in rescue operations," *Auton Robot*, vol. 31, no. 4, pp. 401–417, 2011.
[9] P. Newman, M. Bosse, and J. Leonard, "Autonomous feature-based exploration," in *Proc. ICRA*, 2003, pp. 1234–1240.
[10] W. Burgard, M. Moors, and F. Schneider, "Coordinated multi-robot exploration," *IEEE T Robot*, vol. 21, no. 3, pp. 376–378, 2005.
[11] B. Tovar, L. Munoz, R. Murrieta-Cid, M. Alencastre, R. Monroy, and S. Hutchinson, "Planning exploration strategies for simultaneous localization and mapping," *Robot Auton Syst*, vol. 54, no. 4, pp. 314–331, 2006.
[12] A. Visser and B. Slamet, "Including communication success in the estimation of information gain for multi-robot exploration," in *Proc. International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOPT)*, 2008, pp. 680–687.
[13] F. Amigoni and V. Caglioti, "An information-based exploration strategy for environment mapping with mobile robots." *Robot Auton Syst*, vol. 5, no. 58, pp. 684–699, 2010.
[14] D. Calisi, A. Farinelli, L. Iocchi, and D. Nardi, "Multi-objective exploration and search for autonomous rescue robots," *J Field Robot*, vol. 24, no. 8-9, pp. 763–777, 2007.
[15] L. Freda and G. Oriolo, "Frontier-based probablistic strategies for sensor-based exploration," in *Proc. ICRA*, 2005, pp. 3881–3887.
[16] H. Gonzáles-Baños and J.-C. Latombe, "Navigation strategies for exploring indoor environments," *Int J Robot Res*, vol. 21, no. 10-11, pp. 829–848, 2002.
[17] A. Makarenko, S. Williams, F. Bourgault, and H. Durrant-Whyte, "An experiment in integrated exploration," in *Proc. IROS*, 2002, pp. 534–539.
[18] R. Sim and G. Dudek, "Effective exploration strategies for the construction of visual maps," in *Proc. IROS*, 2003, pp. 3224–3231.
[19] N. Kwak, G.-W. Kim, S.-H. Ji, and B.-H. Lee, "A mobile robot exploration strategy with low cost sonar and tungsten-halogen structured light," *J Intell Robot Syst*, vol. 51, pp. 89–111, 2008.
[20] D. Holz, N. Basilico, F. Amigoni, and S. Behnke, "A comparative evaluation of exploration strategies and heuristics to improve them," in *Proc. ECMR*, 2011, pp. 25–30.
[21] T. Edlinger and E. von Puttkamer, "Exploration of an indoor-environment by an autonomous mobile robot," in *Proc. IROS*, 1994, pp. 1278 –1284.
[22] R. Simmons, D. Apfelbaum, W. Burgard, D. Fox, M. Moors, S. Thrun, and H. Younes, "Coordination for multi-robot exploration and mapping," in *Proc. AAAI*, 2000, pp. 852–858.
[23] M. Keidar and G. Kaminka, "Robot exploration with fast frontier detection: theory and experiments," in *Proc. AAMAS*, 2012, pp. 113–120.
[24] A. Howard and N. Roy, "The robotics data set repository (radish)," 2003. [Online]. Available: http://radish.sourceforge.net/
[25] W. Pestman, *Mathematical Statistics: an Introduction*. de Gruyter, 1998.
[26] A. Quattrini Li, F. Amigoni, and N. Basilico, "Searching for optimal off-line exploration paths in grid environments for a robot with limited visibility," in *Proc. AAAI*, 2012, pp. 2060–2066.