

Insights Toward Robot-Assisted Evacuation

Dylan A. Shell

Maja J Matarić

Interaction Laboratory

Center for Robotics and Embedded Systems

Department of Computer Science, University of Southern California

Los Angeles, CA USA 90089-0781

{shell|mataric}@usc.edu

Abstract

This paper discusses the application of robotic technologies to an evacuation assistance task. We describe how this kind of task differs from the more prevalent search and rescue efforts, and how this difference affects the approach taken in building systems focused on the evacuation phase. The paper describes the implementation and evaluation of a deployment algorithm intended to distribute audio navigational cues throughout an office building environment. A review of evacuation dynamics methods is presented and particular methods are used within the deployment algorithm in an on-line sense. We use a pedestrian simulation and a simple model of audio-evacuee interaction in order to show the effects of beacon deployment; this indicates that even a very few beacons can result in a large decrease in the mean and variance of egress distance (and time). We also advance possible uses for human motion based methods of environment evaluation more broadly for mobile robotics in general.

Keywords: Assistive Disaster Response, Evacuation and Pedestrian Dynamics, Crowd Models.

1 INTRODUCTION

It is likely that the roles robots play in future disaster recovery efforts will increase in both number and significance. Already a growing community of researchers is focusing on the important scientific and engineering challenges that must be overcome if the full potential of rescue robotics is to be realized. This paper describes our recent research centered on assistance of evacuees rather than the more traditional focus on victims. We describe important approaches used by evacuation dynamics researchers, and how this work can be used in a productive manner by robotists.

This paper addresses the task of autonomously deploying navigational aids in the minutes immediately following a disaster. Operating in the evacuation phase, requires that robots be deployed in a timely and non-interfering manner. In order to achieve this we have employed existing robotics approaches that are typically not used in rescue scenarios. Since evacuation assistance tasks have not been the subject of much attention from roboticists, we give a careful discussion of the research challenges, finding that they are distinct from and complementary to those addressed by much existing work.

In this paper, we first outline the existing rescue robotics work, the foundation enabling the identification of distinctive characteristics in evacuation-related tasks. Next we survey the key related concepts from pedestrian dynamics as they relate to the evacuation problem. We provide a broad description of available models and simulation techniques. As far as we are aware, this is the first robotics oriented paper that introducing this line of work, and we thus attempt to give a representative sample of existing

research. We then focus on our concrete evacuation approach: the use of directional sound beacons, and describe a method for automatic multi-robot-based beacon deployment strategy.

It has been shown [50, 51] that audio cues are an highly effective aid for evacuees, leading to the design of directional audio beacons specifically for that purpose [28]. Such beacons emit sound in a range of frequencies known to be easily identified and localized by the human auditory system. To be effective in guiding evacuation, beacons must be appropriately positioned in the environment. We thus focus our work on the problem of autonomously deploying a network of these audio beacons with a team of robots. This paper describes an effective distributed deployment algorithm that is robust with respect to both single robot failures and a class of communication failures within the robot team. Using a simple model of the effect of audio signals on navigational decisions, we demonstrate that the expected egress time is favorably affected by the deployed beacons, and further, that even a small number of beacons can be very effective in aiding evacuation. Evacuation dynamics insights are used in both design and evaluation parts of this work.

Although a simple model of evacuation is sufficient to provide new insight into the evacuation problem, there are a number of cases where more detailed models, which involve complex interaction dynamics such as panic, may be desirable. One important use of such models is the evaluation of environments and environmental constraints, useful in the process of building design. We describe how these methods of evaluation are useful in a robotics context, describing experiments we have conducted to compare different aspects of environment structure and human behavior. We then summarize and conclude the paper.

2 RELATED WORK IN ROBOTIC DISASTER RECOVERY

The idea of using robots to aid in human rescue is not new; see for example the early paper by Kobayashi and Nakamura [29]. Generally, the work focuses on urban search and rescue (USAR) tasks that require robots to overcome highly uncertain (i.e., potentially drastically altered) environments in order to locate, treat, and possibly transport incapacitated people [5]. These search and rescue types of tasks are a central focus of rescue robotics because the potential payoffs are tremendous, and the outstanding research challenges are daunting.

Robotic USAR could potentially save the lives of human emergency personnel by placing a degree of separation between them and the hazardous environment. This also results in an improvement on current performance, because robots can enable telepresence in structures previously considered unsafe and can reach otherwise inaccessible areas and difficult locations [35]. To reach this potential, significant advances are still needed both robot hardware and navigation, localization, and assistance software, thus opening up interesting avenues for research.

Another open research area is that of human-robot interaction (HRI) in the context of USAR. Murphy [33] describes the unique challenges involved, focusing on the robot-operator relationship. Systems in current use typically involve a team of people per robot; methods for reducing this ratio are needed.

In spite of the many outstanding challenges, initial progress has already been made in the USAR arena, by having search and rescue robots demonstrated in real disaster scenarios. Murphy [34] describes the development of a team of such robots, the initial observations of the Oklahoma City bombing site indicating their feasibility, and continued refinement of the systems. The same group deployed USAR robots at the site of the New York City World Trade Center disaster. The discovered locomotion challenges have already resulted in a number of improving innovations. Murphy [32] describes using both marsupial (small robot inside larger robot) and shape-shifting methodologies, while Wolf et al. [52] propose a hyper redundant ‘trunk’ mechanism to enable sensing in difficult-to-reach locations. Both shape-shifting and hyper redundancy properties are found in some self-reconfigurable robots, making

them also potentially useful USAR domains (see for example Støy et al. [45]).

USAR, however, is not the only interesting and rich area of rescue robotics. To visualize the larger picture, in Figure 1, we propose a space of rescue robotics tasks and problems. The three axes of the space are: *Urban Structural Disintegration*, *Diminishing Human Mobility*, and *Expected Response Time*. The two volumes shown each indicate a typical rescue task. The location and size of the volumes within this task space roughly indicate the expected task conditions. The larger of the volumes (shown on the upper right) represents the USAR task. There, the structural disintegration is typically very high, involving collapsed or partially collapsed buildings. The victims typically have very little mobility, as they are either trapped or disabled. Finally, the response time begins sometime after the disaster (e.g., rescue robots were deployed one day after the New York City disaster) and continues for an extended time period.

In contrast to USAR, we define the *evacuation assistance* task (represented by second, smaller volume in the diagram), whose conditions are significantly different. The environment is typically in much less affected state, evacuees are assumed to be self-mobile (but possibly impaired to an unknown degree and number), and the shorter, immediate evacuation phase occurs immediately, before the rescue interval. While the shown volumes occupied by each of each of these tasks clearly differ, the number of human lives saved in both scenarios are significant.

Given this newly defined rescue task, we now survey the relevant existing work in evacuation dynamics, which we can bring to bear toward automated, robotics solutions that aid in the evacuation assistance task.

3 RELATED WORK IN EVACUATION DYNAMICS

Any robot intended to assist during evacuation must be able to interact with different numbers of people under stressful conditions, from individuals to crowds. This section describes a number of relevant human crowd behavior models potentially useful in developing human-robot interaction modalities for this assistive evacuation domain.

There has been significant interest in ordinary pedestrian flow for quite some time [18], with research ranging from empirical studies to detailed modeling and simulation. Progress has been made with the discovery of an (inverse-proportionality) relationship between flow and pedestrian concentration [36]. Also, it has been amply demonstrated that simple systems can exhibit some of the characteristic properties of human traffic, such as self-organized lane forming [18, 19, 23, 44] and oscillatory behavior at bottlenecks [18, 19].

Generative models typically use empirical data to verify that expected behavior is in fact produced as output, and also as a basis for various input model parameters. Daamen et al. [8] describes the use such a pedestrian model in the process of architectural planning of railway transfer stations; it is postulated that the models will help to improve timetables and diminish waiting times.

Some researchers [17, 18, 39] have attempted to model panic and extreme circumstance in the behavioral dynamics of pedestrians. The work is intended to address questions regarding distressed crowds in evacuation scenarios, in the same way that pedestrian dynamics addresses more ordinary circumstances.

Proulx [37] and others have pointed out that people do not necessarily panic during evacuation scenarios, and often may remain rational, calm, and collected. Nevertheless, a number of recorded incidents of panic induced tragedies [19, 44, and references therein]. Identifying the key causal elements for collective panic is taken up by sociologists and social psychologists, in both experimental (e.g., Mintz [31]) and theoretical (e.g., Smelser [43]) work. We do not address these issues any further, but

now turn to the models and simulation mechanisms that have direct implications for the use of robots in evacuation scenarios.

3.1 Models and Simulations

Traditionally, models have focused more on analytical solutions to problems, and simulations on computational experiments. In the domain of pedestrian and evacuation dynamics, this distinction is less useful because there exist descriptions that seek to model course-grained macroscopic system properties requiring significant numerical optimization [40], as well simulation scenarios for which analytical relations have been found [27]. In our survey, we follow the same broad categorization as Hamacher and Tjandra [17], based on whether the focus is on individual behavior or on collective properties.

3.1.1 Macroscopic Approaches

Macroscopic approaches do not target the behavior of a single pedestrian but focus instead on aggregate properties; these models are commonly based on network flow concepts. We describe models used by Hamacher and Tjandra [17], which are largely representative of this approach. In those, a graph is constructed that represents the environmental constraints, and calculations are performed to evaluate aspects of the graph and hence the implications of the environmental constraints on evacuation time.

The graph as an abstraction of the environment is most often used to represent a single building. A typical construction would have a single graph vertex per room, with an edge placed between any two adjoining vertices; a single additional *goal* vertex then represents the area external to the building. The edges of the graph are labeled with various route details, such as *edge capacity* and *travel time*. Vertices are also labeled, such as with *initial contents* and *capacity*. Graphs that capture dynamics (e.g., spreading fire fronts) through duplicates for each time-step are called *dynamic networks*.

Flow takes place from vertex to vertex, but cannot exceed the capacities defined in either the edges or the destination. *Max flow* represents the maximum throughput in the represented environment, and its optimization is typically solved using specialized network flow algorithms. Max flow and other calculations, such as quickest flow, provide a characterization of the environment, identification of bottlenecks and potential trouble areas, and allow engineers to answer “what if” questions regarding building structure. Hamacher and Tjandra [17] claim that these macroscopic models are typically used to find good lower bound estimates for evacuation scenarios.

3.1.2 Microscopic Approaches

While the described macroscopic models are useful, they are unable to take into consideration complex effects related to interactions and interference among people. Microscopic approaches, in contrast, can capture those effects, and thus focus on the individuals and the rules that dictate their inter-personal behavior. Evacuees are simulated as simple role-governed entities in the environment, the collectively generate representative crowd behavior. From a scientific perspective, understanding the local rules that produce realistic collective dynamics is important, with implications toward proxemics [16], i.e., the personal use of space for highly dynamic and interactive activities.

One of the most widely recognized microscopic modeling approaches is the “social forces” model Helbing et al. [19]. In it, people are particles (sizes sampled from an appropriate distribution function) that move when acted upon by virtual forces; force terms act between people and also between people and the environment. Attractive forces bring groups of acquaintances together, while repulsive forces act to keep personal distances; the walls have a repulsive action, but a number of displays or other items of interest may have attractive forces. Additional frictional terms, inspired by granular flow, are used for high crowd densities.

The “social force” model has a single *agitation* parameter; by varying its value, a range of behavioral dynamics can be produced, from regular pedestrian flows to panicked behavior. The use of forces also allows the applied pressure to be calculated, and hence fatalities caused by crushing. In the most general case, force terms may depend on time, velocity, and may be anisotropic. This allows for fire-fronts and other dynamic elements to be modeled.

When discrete models of space are used in these models, Cellular Automata (CA) models and algorithms can be applied. These models are popular both for representing pedestrians [27] and evacuees [39, 40] due to their simplicity and simulation speed. Models with continuous space and time also exist, such as Hoogendoorn et al. [23].

3.2 Space Syntax Analysis

Models described above are targeted either at pedestrians or evacuation traffic. However, more general models of space may also be useful in the context of evacuation. Next we describe space syntax, a model developed by Hillier and Hanson [22], and aimed at architectural design.

Connectivity is a key notion in space syntax theory. Environments are described in terms of topological representations, and environmental complexity is measured through connections to neighbors. This makes the model scale-invariant, allowing the theory to be applied to structures as small as single building (which we will focus on) and as large as an entire city [20, 38].

First, the environment is skeletonized into a graph that captures the topological structure. This step is similar in many ways to the graph construction involved in network flow methods described above. In regular polygonal structures such as typical office buildings, individual rooms become vertices, and their adjacency relationships become edges. In spaces where the structure is less obvious, the free space is segmented into convex polygons, then skeletonized using a set of *axial maps*, the minimum number of maximal length lines that cross all adjacencies between convex polygons. Each line in the axial map is called an *axial line* and becomes a graph vertex; any intersecting lines share an edge in the graph. Figure 2 shows three environments with their axial maps; the underlying convex polygon segmentation is easily inferred. Other approaches to converting environmental structure into graph form have been proposed; the two described above are representative.

The mean depth of vertex i (calculated to radius r) is given as:

$$\mathbf{MD}_i^r = \frac{1}{\|\mathcal{N}_i\| - 1} \sum_{\substack{j \in \mathcal{N}_i \\ j \neq i}} d(i, j) \quad \text{where } \mathcal{N}_i = \{\forall k : d(i, k) \leq r\} \quad (1)$$

where $d(i, j)$ is the minimum distance from vertex i to j , measured in numbers of edges of the adjacency graph. The set \mathcal{N}_i represents all other vertices with a minimum distance no farther than radius r away. Frequently, r is taken to be as large as the entire graph, so that \mathcal{N}_i includes all vertices and \mathbf{MD}_i^∞ . Another common case is $r = 3$, easily represented by \mathbf{MD}_i^3 .

The mean depth of a vertex indicates how directly it is connected to the regions around it. The relative asymmetry of vertex i (calculated to radius r) is defined as:

$$\mathbf{RA}_i^r = \frac{2(\mathbf{MD}_i^r - 1)}{\|\mathcal{N}_i\| - 2} \quad \text{where } \mathcal{N}_i = \{\forall k : d(i, k) \leq r\} \quad (2)$$

Hillier and Hanson [22] define the relative asymmetry as a normalization of the mean depth, ensuring that the value falls between 0 and 1.

In the cases where the radius includes all elements in the graph (\mathbf{RA}_i^∞) this measure is termed the *global integration* of vertex i . A small constant (typically set to 3) is used to define the *local integration* (\mathbf{RA}_i^3) of vertex i .




Figure 1: A decomposition of the task space for disaster recovery robotics. The small and large blocks represent evacuation assistance and USAR tasks, respectively. The space spanned by the axes indicates conditions of the survivors; some regions are infeasible, and no survivors are assumed to exist outside the axes.




Figure 2: Three different environments with varying complexity, each shown with axial maps in dotted lines.

Hillier [20] described motorized traffic patterns in a portion of central London, showing that rush-hour and mid-day traffic flow along roads represented with axial lines significantly correlate to global and local integration values, respectively. This analysis has subsequently been performed in dozens of other cities (e.g., [38]), with similar results. The technique has also been successfully applied to pedestrian traffic at the city level, which is a much larger scale than typically studied.

Space syntax measures are useful at smaller scales as well. One study calculated local integration values for convex portions of a museum, and found a clear correlation with people’s observed routes in the first ten minutes spent in the museum [25].

Collectively, this empirical evidence gave rise to the central theories of space syntax [20]. Most relevant to our work is the idea that the structure of the free space is largely responsible for traffic flow.

As Hillier [21] states, the goal of space syntax research was never to accurately predict pedestrian or vehicular flows, but rather to understand the degree to which movement is influenced by the constraints that the spatial configuration imposes. Since people tend to minimize effort while navigating [44], it appears that the connectedness of urban structures inherently captures that cost minimization.

The fact that space syntax measures of built structures have predictive capabilities at multiple levels of detail and differing scales is very fascinating because it indicates the deep relationship between peoples’ behavioral patterns and the space they inhabit. Batty [2] has proposed that ideas from the generative models used by Helbing be combined with space syntax rules, enabling even wider range of applicability.

In Section 5, we describe a number of implications of space syntax for mobile robotics, a field that is still in need of quantitative measures of environmental factors.

4 AUDITORY BEACON STRATEGIES FOR EVACUATION

The importance of the evacuation problem continues to drive the development and elaboration of models and improvements to existing environment. In contrast, in this work we focus on the use of robots during evacuation in order to improve the egress time, namely in the context of the evacuation assistance task (as defined in Section 2).

Løvås [30] offers three reasons for why evacuation can take longer than it should: delayed initial response, non-optimal selection of escape routes, and congestion while traveling. In cases of fire, initial response is often delayed because of underestimation of the effects of smoke [37]. It has been widely noted that non-optimal selection of routes often occurs because people typically exit a building the same way they entered it. Thus, congestion occurs due to architectural shortcomings, identical choices of exit routes by people, and self-reinforcing herding effects.

Low visibility circumstances can be particularly dangerous because people tend to use visual cues as the primary means for spatial localization, and hence in escape route planning and navigation. The most tragic examples are of deaths that occur because victims are unable to locate suitable emergency exits in time. For example, at the Düsseldorf Airport Disaster [9], asphyxiated bodies were discovered only three meters away from an emergency exit. Visual signs remain the most frequently used means to mark fire escape routes, in spite of their known rapid deterioration of effectiveness with smoke accumulation.

4.1 Directional Audio

One solution to the above problems is the use of directional audio beacons [49, 50, 51]. These beacons (see Figure 3) emit sound in a range of frequencies known to be well-suited for localization by the human ear and composed of fused sounds that minimize ambiguous “blind spots.” The beacons are

called *directional* because their bearing can be determined with high certainty by a listener even when many beacons are active simultaneously. This technology is distinct from 3D-audio [3].

Multi-floored buildings allow beacons to play an enhanced role. In addition to the standard pulsing broadband sound, beacons located near stairwells can play a sweeping melodic tone, with ascending tonality indicating that subjects should travel upward, and descending tonality meaning the reverse. Withington [50] describes the use of such beacons to direct a group of people during a navigation task through one such complex environment. The subjects were not informed about the meaning of the sounds, but nevertheless not a single participant in any of the trials took a wrong turn or mistakenly ended up in the incorrect room.

Trials conducted with human participants (both with and without decreased visibility) established overwhelming effectiveness of the beacons in aiding evacuation. For instance, in one trial, people got lost in an environment with 100% visibility. Notably, they had successfully evacuated the same environment, while smoke-filled, only minutes before with the aid of audio beacons [50].

4.2 The Deployment Task

A significant body of work has focused on evaluating the effectiveness of the above audio beacons while assuming that they are appropriately located within the environment in advance. Withington’s intention, for example, is to have audio beacons fitted during the building construction period. However, this is not a reasonable assumption for all buildings, and fails to cover existing structures. Therefore, we [42] have studied the problem of dynamically deploying audio beacons when needed. Numerous arguments can be made for such dynamic deployment; the most pragmatic is that urban structures are not likely to be retrofitted with audio beacons in the immediate future. Other benefits of dynamic deployment include robustness; the failure of an important beacon would automatically result in the addition or redistribution of other beacons. Less critical uses of such robotic solutions include the added ability to have robots monitor traffic (and other variables) at particular locations, making important statistics available to emergency personnel in real-time.

Two different solutions to the dynamic beacon deployment problem present themselves, each inspired by existing work in sensor network deployment. The first involves some number of robots depositing multiple beacons as they move through the environment in a purposeful fashion (see for example the sensor network deployment strategy in Batalin and Sukhatme [1]). The second provides each robot with a single beacon to be activated when the robot arrives at an appropriate location. We focus here on the second method assuming sufficient robots are available and minimizing deployment time is a key goal. Therefore, since each robot is equipped with a beacon, heretofore the robot-beacon synergy will just be called a robot.

We describe a distributed algorithm that employs a number of properties of the broader class of evacuation assistance tasks. The algorithmic choice effects the set of associated assumptions and applicable technologies, discussed next.

4.3 Robotic Technologies

The Structural Disintegration axis, shown in Figure 1, indicates that our evacuation assistance task is intended largely for unharmed or nearly unharmed environments, which still may be difficult to navigate, due to smoke and/or dust. As Murphy [33] points out, these scenarios can be considered benign when compared the severe hazards in collapsed buildings typical of USAR. Hence, existing algorithms for autonomous indoor localization and mapping (e.g., Howard [24]) are feasible for addressing the beacon deployment problem and most likely for a large set of evacuation assistance tasks, while the same

solutions are inoperable in USAR. Techniques garnered from the ongoing work on outdoor mapping will expand the envelope of applicable tasks.

The Time axis (of Figure 1) indicates that evacuation tasks typically require a faster response than most of USAR. This has significant design implications: the approach we describe here focuses on parallelizing the deployment through the use of multiple robots in order to minimize deployment time. Adding more robots to a system does not necessarily improve performance, but, fortunately, this trade-off between parallelism and interference is well studied and a number of well understood methods exist in multi-robot coordination literature that address it (e.g., Goldberg and Matarić [13, 14, 15]).

Both the severity of victims’ injuries and the environmental constraints in USAR dictate that a significant portion of robot intervention involves human operators. In contrast, in the deployment task and other evacuation scenarios, a much higher level of autonomy is appropriate and desirable. The algorithm we describe requires very little interaction from operators, as described below.

4.4 The Deployment Algorithm

Consider the following motivating scenario: a fire has started in an office building, an alarm is triggered, firefighters are notified, and suitable vehicles deployed. In addition, a squad of robot operators arrive at the scene, with a team of robots (equipped with beacons). The squad deploys the robots into the building at the available entrances, including doors and windows. If a map of the building is known, it is provided to the robots ahead of time; if not, the robots perform an autonomous mapping phase.

Next, an operator uses a console to provide some useful information to the robot team. Specifically, she indicates which locations in the building constitute suitable exits, including existing doorways but also any emergency exits and newly-created exit points such as windows with ladders, etc. In addition, the operator also provides any necessary navigation information the robots may not have acquired, such as stairwell connections. The robots then compute suitable deployment positions and navigate towards them. Upon arrival, they activate the beacons they carry. The beacon sounds serve two purposes; they 1) provide information about escape routes, and 2) spur any occupants who have missed or ignored prior evacuation warnings.

The above scenario involves a set of fundamental capabilities of the robots: 1) localization, navigation, and mapping, 2) wireless communication, and 3) task allocation for team coordination. Fortunately, all of those capabilities are available in the current state-of-the-art, but naturally require adaptation to the particular domain of use.

Currently favored methods for indoor localization and mapping based on Bayesian formulations have been shown to be robust for coordinated multi-robot mapping. In our work, we use an implementation similar to Howard [24], which assumes that the robots have a map of the environment and are, initially at least, localized within that map. Since the building to be evacuated may have more than a single level, we employ a map for each floor. As specified above, a list of links between the floors (at stairwell locations) and needed supplementary information, such as the locations of fire exits, are also provided. What results is essentially a $2\frac{1}{2}$ dimensional representation for the multi-level building. We provided such floor maps and used an adaptive particle filter [10] approach to enable the robot to compute continuous pose estimates.

Figure 4 shows the map used for a test environment with two floors. The figure also shows the additional knowledge that would be imparted by the emergency operator. We assume this information is given as labels to a topological (graph) overlay. Thrun and Bücken [46] describe a method for automatic generation of this graph from metric maps that robot range sensors with scan matching generate.

For coordinated beacon deployment, we employ an approach from the multi-robot task allocation (MRTA) literature Gerkey and Matarić [12]. Typical deployment algorithms intended for sensor net-



Figure 3: Two directional beacons used to deliver audio signals so that the source's location is easily identifiable. Manufactured by Brigade PLC and Klaxon Signals PLC respectively, the shrieker on the right was explicitly designed for evacuation scenarios.




Figure 4: A representation of one of the multi-floored test environments used for beacon deployment. The dots mark pieces of information added by the operator, in this case emergency exits and stairwell links. The image on the left includes a topological overlay for the environment. On the right, the emergency exits are numbered 1-5 and the connection stairways 6 and 7 (the second floor is only partially shown).

works tend to focus on the notion of coverage, i.e., maximizing the total area collectively sensed. In the evacuation scenario, however, coverage is less important than providing useful information about the location of the exits. Therefore, we chose to use a MRTA approach as the primary coordination mechanism, previously used for multi-robot exploration [4]. Our approach is described next.

4.4.1 Task-Allocation for Coordination

Gerkey and Mataric [12] describe a wide variety of robot task allocation problems, their formal properties, and some existing algorithmic solutions. They classify the problem space along three dimensions: 1) single-robot tasks v. multi-robot tasks, 2) single-task robots v. multi-task robots, and 3) instantaneous v. time-extended task assignments [11]. Importantly, they identify that the problem of instantaneous assignment with single-task robots and single-robot tasks, the problem most frequently treated in the multi-robot coordination literature, can be treated as an instant of the well-known Optimal Assignment Problem (OAP), which has a known polynomial-time solution.

A direct way to cast the deployment problem into the MRTA framework as an OAP instance is to consider the potential deployment locations as navigational tasks for each of the robots. This requires the construction of a utility matrix, with each element being an estimate of the utility, or expected worth, of a particular assignment. The (i, j) th utility matrix entry is for robot i being deployed to location j , and can be considered the *reward* for having a beacon at location j , less the *cost* expected to be incurred by that robot during navigation that location. Once the full matrix has been constructed, robots can be assigned to tasks. To decentralize the process, we distribute the utility estimates and have the robots compute their own allocations locally.

Unfortunately, the OAP formalism requires fixed non-interrelated utilities, a requirement that is not satisfied by this example. The reward obtained by having a beacon at a particular location depends not only on the environmental constraints, but also on where the other beacons are positioned. Imagine that one exit is in an excellent location; assigning two robots close to that exit is inefficient, once a robot has been deployed to a location, the utility of having others nearby is greatly decreased.

To avoid this problem we permit a task (in the MRTA framework) to include a set of locations within some local neighborhood; a robot being assigned to that location is ensured that no other robots will be deployed within that neighborhood (since we are considering single-robot tasks). Next, we discuss the problem in three phases: 1) choosing a list locations suitable for deployment and clustering them into tasks, 2) assigning robots to those tasks, and 3) having robots navigate toward their appropriate locations. Each of these is discussed in turn.

4.4.2 Where to Deploy

Feasible locations are selected from the set of emergency exits and stairways which are strategic locations consistent with those used in validation experiments with the beacons¹. These locations are clustered based on distance. Assigning robots to entire clusters separates deployed beacons by a minimum threshold distance. The difficult problem involves finding the maximum such threshold.

Consider, for example, the two-floor environment shown in Figure 4. When deploying a single robot we consider the partition of feasible locations to be the set of all possible destinations, $S_1^1 = \{1, \dots, 7\}$. If another robot is added, then the environment is partitioned into two sets of locations, $S_0^2 = \{1, \dots, 4, 6\}$ and $S_1^2 = \{5, 7\}$. This is performed using an approach similar to Kruskal's Algorithm [7] for construction of a minimal spanning tree (MST). The last step of Kruskal's algorithm merges the two sets of vertices, S_0^2 and S_1^2 , to form S_1^1 . Rather than constructing the full MST, the last merger need not be performed, resulting in a forest of size two. More generally, the set merger step is applied only until the

¹Personal Correspondence, Oct 16, 2003.

correct number of tasks remains; the length of the last edge added to the MST is effectively the cluster threshold.

A complication arises due to the clustering of locations, since one requires not only that there be a sufficient number of tasks, but also that there exist a feasible assignment of tasks to robots. This is easily handled by calculating the S -sets as described above, thereafter relaxing the threshold until a feasible assignment is found for all robots; if none exists for a particular robot, then that robot is assigned no task, i.e., a *no-op* task.

The utility estimate for a cluster is calculated as the maximum utility over all the members of that cluster. A utility estimate for each cluster attempts to factor multiple determinants into a single scalar value. The complete value (i.e. utility for robot m being positioned at location n) is calculated as:

$$\mathcal{U}_{m,n} = \max(0, \alpha(1 - \mathbf{RA}_n^\infty) - \beta D_m^n) \quad (3)$$

where \mathbf{RA}_n^∞ is the global integration from space syntax analysis, in this case taken for the location. This gives us a value for how well connected the exit is, if the exit is well connected $\mathbf{RA}_n^\infty \approx 0$. We define D_m^n to be the distance from the current pose of robot m , to the final destination n (which implied $D_m^n \geq 0$). All distances used in the calculations are based on actual distances involved in navigating between locations; they incorporate the known obstacles. The values of α and β allow the importance of connectivity versus cost of travel to be traded-off with respect to one another; we found that $\alpha = 100$, and $\beta = 1$ (with D_m^n in meters) were appropriate.

4.4.3 Whom to Deploy

Each of the robots uses Kuhn’s Hungarian Method [12] to calculate the assignments; in our implementation the calculation of utility matrix entries was far more computationally intensive than actually performing the assignment.

Because robots are assumed to be constrained to a particular floor, the resulting utility matrix has a particular structure: two robots on adjacent floors have positive utility estimates for destinations on their own floor and 0 for the locations elsewhere, stairways between the floors are the only cases where both robots may have non-zero utilities. Stairwells are locations that exist on both floors – if a robot is tasked with one of these locations it moves to either the head or foot of the stairwell, as appropriate.

If we were to ignore the existence of stairs (and assume that sound does not permeate between floors) then the assignment problems could be solved independently on each floor. However, stairwells are extremely important during evacuation and need to be included in the assignment. Most exits occur on the ground floor, so in realistic environments robots deployed on higher floors tend to be assigned to stair duty, while those on the ground floor mark the emergency exits. This is a solution that our algorithm finds naturally.

4.4.4 Autonomous Navigation

Our navigation algorithm uses a simple graph-based planner, and an implementation of VFH+ [47] for local obstacle avoidance. The former generates way-points, the latter travels between them. The implementation is written in Python and C using the Player robot device server and the Stage simulator. Virtual device abstractions (described in Vaughan et al. [48]) permit effortless use of the VFH+ routines and adaptive particle-based localization [10].

As the robots move through the environment, they continually monitor the network for a signal indicating that a new robot is available to join the team. When this occurs, the current assignment is preempted and a new one (that includes the newcomer) is calculated. A timeout is used to ensure that all robots which were included in an assignment are still available – if one is no longer functional, the

assignment is recalculated. These two actions ensure that the positions occupied by beacons always (up to a time-out frequency) reflect the best possible configuration for the number and constraints of the available robots.

The nodes only communicate in order to share their global pose estimates at the beginning of the assignment step. Implicitly, the robots assume that all others are performing identical assignments; since the assignment algorithm is deterministic, all robots arrive at the same deployment agreement. This method of distributing the task allocation process is described in Gerkey and Matarić [12]. No single robot is responsible for assigning tasks to others, their simultaneous execution of the same algorithm results in a consistent allocation.

4.5 Evacuation Dynamics Models in Utility Calculation

The explicit function used to calculate the utility estimate for the robot location assignment problem is given in Equation 3 above. It consists of two terms, one for the estimated benefit that a beacon at a particular location would bestow, the other for the cost of achieving that assignment.

Within our formulation, the evacuation (or pedestrian) dynamic domain knowledge plays a role in only the first term. The space syntax calculation described above clearly has a number of useful properties, which include efficient real-time on-board calculation and the ability for almost all the essential computation to be performed on the topological overlay.

We, however, see the space syntax method as only a single representation among many; in fact, almost any of the methods described in Section 3 would be suitable given adequate information and sufficient computational resources. The resulting research questions are: 1) under what circumstances are pedestrian and evacuation models appropriate, and 2) are the differences in approach (between, say, microscopic and macroscopic models) indicative of assumptions regarding the type of flow? Section 5 returns to this topic.

The ability to include arbitrary factors into the calculation of abstract utilities is certainly one benefit of decision theoretic reasoning, and hence a strength of MRTA. It should be stressed, however, that the utility values are only estimates of the worth of a particular robot for a given task assignment. The world within which the robots operate is dynamic, and many unforeseen events may (and do) occur. In the current implementation, the utility calculation step in the deployment algorithm currently takes a significant portion of each robot’s time. Introducing additional computation into the utility estimation may improve the estimation at the prohibitive cost of real-time response.

4.6 Experimental Validation

We have evaluated our implementation along two dimensions: 1) robustness with respect to failure and 2) effectiveness in terms of impact on the evacuees. A desirable system must meet both of these criteria to a satisfactory degree; we attempt to demonstrate that this is the case for our algorithm.

4.6.1 Robustness

Inspired by traditional white-box testing methodologies, we would like to have all possible logic branches of the algorithm tested; this is obviously infeasible. We present a representative sample run where the system behaved reasonably. The run is one of ten instances in which beacons were deployed into the Figure 4 environment and a failure purposefully induced. Similar runs on three additional environments were also performed; in all cases the system demonstrated similar reliability – detection and reassignment to compensate for the failure. The forced failure was performed by shutting down a robot after it has been deployed. The exact time of this operation (and its detection) differed from run to run, which




Figure 5: The map used with physical robots. Exits marked 1, 2. The connection labeled 3 is a stairwell. The ActivMedia Pioneer DX2 is shown moving toward the top of the stairwell.




Figure 6: The distances of each robot from its assigned destination. The reassignment occurs at the 300 second mark, resulting in the discontinuity.

unfortunately makes it difficult to present a single visually meaningful summary of the data from all the runs.

Figure 6 shows the distance of each robot from its assigned goal destination for the sample run, over a range of times. (Robots 1 and 2 have line segments very close to one another, and are indistinguishable unless magnified.) By time $t = 200$, all of the beacons have been deployed and the system could remain in this state indefinitely, with activated sound beacons. Next, robot 5 was shut down. The system remained in the state, with only four active beacons, until a time-out elapsed resulting in a robot sending a request for another’s pose information. This caused a cascade of further requests. The remaining four discovered that robot 5 was no longer functioning, and that robot 4 would serve as a better beacon if it moved to a place previously assigned to the failed robot. This is observable as the spike in the graph around $t = 300$ for robot 4, because its destination changed.

The case of a network partition is dealt with similarly; those robots that can communicate do so, and coordinate their efforts accordingly. Two robots on different partitions cannot communicate, and thus cannot avoid being assigned destinations that are co-located. This, however, has the useful side-effect of bringing them together, and most likely within communication range. One can construct a scenario with communications links being added and removed in a particular order such that a continuous cycle of deployment activities occurs. Our empirical experience indicates that most examples involve some number of robots being stable and repeatedly assigned to the same location, and a few others switching between partitions; those robots that remain stable still act as effective beacons.

An additional ten runs were attempted on three physical robots on the first two floors of our laboratory building (see map in Figure 5). Our experimental platform consisted of three pioneer 2-DX robots equipped with laser range finders. The first two robots were deployed from the second floor; they moved toward the top of the stairs (bottom left corner, Figure 5) and the exit marked 2 (upper right corner, Figure 5), respectively. A third robot was deployed on the bottom floor, and moved toward the exit marked 1. A failure was induced in the robot located at exit 2; after being detected, the robot above the stairwell moved toward exit 2, resulting in a better spread. After six consecutive successful trials, the experiment was halted due to a hardware failure.



4.6.2 Effectiveness

There is no formal description for where the robots should best be deployed, or which beacons should be relocated when a beacon failure occurs. We have embraced the standard practice of using exits and stairways, but evaluation of our location selection mechanism is difficult.

We propose a metric that attempts to capture the efficacy of a given arrangement of beacons. The metric is based upon a simulation of expected evacuee flow with and without beacons. The authors are unaware of any other existing simulation of evacuees (or pedestrians) being aided by external signals intended to improve movement.



We define a *choice point* as a location where there is a fork in the topological structure of the environment, i.e., where the evacuee (or robot) must make a navigational choice. These are all the vertices with degree three or larger. We consider these choice points as the locations wherein the navigational effects of audio (or other signals) occurs. In order to simulate human behavior in an unknown environment, we assume that all unvisited options at any choice point is taken with equal probability. The audio signals *bias* the navigational decision taken at choice points.

A graph-based representation of the environment allows for the identification of choice points. In order to estimate macroscopic flow, we generated evacuees, placed them randomly on the graph, and



(a) Expected Distance traveled per evacuee, when the simulation uses uniform (over area) sampling for initial locations. Both variance and mean decrease significantly after the introduction of a single beacon. Beacon bias parameter = 0 with 1,000 simulated pedestrians for each of the 6 beacon values.

(b) Expected Distance traveled per evacuee, when the simulation uses \mathbf{RA}_i^∞ as a statistical weight for initial locations i . The trend is consistent with (a). Parameter values are identical with those in (a).



(c) The effect of the beacon bias parameter on the mean escape distance. Solid lines represent the extreme values of 0 and 1, as upper and lower bounds, respectively. The bias parameter captures the ability of an evacuee to pick out the loudest (nearest) beacon, as expected, when this is done flawlessly (value=1) egress distance and time is minimized.

(d) Same as (c), only difference being that the initial sampling was based on \mathbf{RA}_i^∞ instead of uniform. Even if exactly the nearest beacon can not be located, modeled as bias parameter value = 0, the mean egress time is not devastatingly effected.

Figure 7: All four of these plots give the mean distance an evacuee will travel (based on the simple flow model described in the text) as a function of the number of deployed beacons. This is a function of the *locations* of the beacons, too; this plot is for the five locations chosen by the described allocation algorithm. The two plots on the left have initial locations of simulated pedestrians assigned uniformly over the environmental area; the right plots use the value of the *global integration* as the statistical weight. The top two plots show generated data, mean and standard deviations (error bar width = one standard deviation) of simulated runs. The bottom two plots show the range of variation as the bias parameter is varied from 0 (solid line that is upper bound) to 1 (lower bounding solid line).

simulated their navigational decisions. The distances traveled by each simulated person were recorded, permitting analysis.

Figure 7 shows the relationship between expected evacuation distance (or time) and the number of robots deployed into a particular environment. The graph was plotted with 1000 navigational instances for each of 6 values of “number of beacons.” Error bars in the top two plots give a measure for the variance of the distances traveled.

The four plots in Figure 7 summarize two fundamental variations explored with the model. The first is stability with respect to the initial distribution of people. The plots on the left assume people are uniformly spread throughout the environment; this is quite far from true in everyday use of space. To address this fact, we used space syntax analysis and calculated the global integration throughout the environment. This was used as a probabilistic weight to generate locations (higher integrations resulting in more pedestrians); results are shown in the two plots on the right. As expected, this resulted in a slightly lower mean value for each of the evacuation times. However, the model has shown that this difference is insignificant when compared with the time variance.

The lower two plots demonstrate the effect of the beacon bias parameter on the mean distances. The bias parameter captures the ability of an individual to infer distances to the beacons, and choose the route toward the nearest one or, said another way, the ability of the nearest beacon to bias the evacuee’s choice of route. A value of zero represents the case when route choice is based entirely on distance from the appropriate choice point to the beacons. As the value is increased toward unity, this models the persons ability to pick the shortest route. The lower solid lines in the plots (c) and (d) are for the parameter set to 1, and represent a lower bound on mean distance.

In the plot the robots were deployed from the same initial locations as the sample run from the previous section. The environment shown in Figure 4 is used here because it is the most complex of the environments we considered, enabling more robots to be deployed and hence more resulting data. The trend is, however, consistent with the other simulated deployment runs mentioned in the previous section: a very rapid decrease in effectiveness per beacon. This indicates that even very few beacons are likely to be effective. Both the mean and variance in egress times decrease in the presence of audio beacons than without. This is significant, because a decrease in one without the other has limited value.

5 MEASURING ENVIRONMENTAL COMPLEXITY

In the deployment algorithm described above, space syntax measures were used to make informed calculations regarding local environmental structure and its effects on people’s route-planning behavior. We believe this is the first instance of a quantitative measure of environmental complexity used in an on-line fashion. The measure may seem closely tied to the application, or perhaps, more generally to the class of evacuation assistance tasks. Next we provide a condensed description of the ideas that appeared in Shell and Matarić [41], arguing that human motion provides the basis for objective measures of environmental intricacy, for a larger community of experimental robotists.

As is well known, the structure and complexity of the environment can have a major influence on robot task performance. Consider, for example, the task of target tracking. Jung and Sukhatme [26] note that “how obstructed [the environment] is, seems to be significant.” They then use a metric based on visibility between points in the environment to provide a measure of occlusion. Their measures (from the computer graphics community, originally proposed by [6]) are representative of a common focus on sensing as the baseline. The environment affects not only sensing, but also action, i.e., *motion*. Since motion is something that mobile robots must do, regardless of task, the effects of the environment that influence freedom of motion are particularly relevant.

Indoor environments are designed for human use, and so measures of human mobility are appropriate for assessing environmental structure. Of course, robots are not people, and so it is questionable how meaningful those measures are when applied to robot movement in the same environments. Nevertheless, focusing on the movements and capabilities people does provides a robot-neutral, and hence relatively objective measure.

Space syntax measures above were applied to compare different parts of a single environment. In those cases, the values of local (\mathbf{RA}_i^3) and global (\mathbf{RA}_i^∞) integration were directly suitable. When quantitatively comparing two entire environments with one another, the typical approach in architecture is to use a mean values of local and global integration taken over the entire environment. In Shell and Mataric [41], we attempted to experimentally compare space syntax measures and a representative microscopic fluid-like simulation, based on a subset of forces from Helbing et al. [19]. The microscopic simulation focused on ego-centric behavior in separation of pedestrians. We considered integration values (for parts of environments) and calculations involving the means. In all cases, little or no correlation could be found between the two.

Although any experimental approach would likely be inconclusive due to the large space of possible environments, we hypothesize that the observed disagreement is indicative of different types of collective behavior. Fluid simulation are more suitable for understanding congestion and interference effects, while the space syntax analysis is ideal for understanding expected navigational choices. Helbing [18] discusses the transition between order and disordered regimes in freeway traffic; perhaps the same holds true for pedestrians, presenting interesting research questions in evacuation dynamics.

6 SUMMARY AND CONCLUSION

This paper has defined a new rescue task, the evacuation assistance task, and contrasted it to USAR. We believe that there is significant potential for life-saving robotics applications beyond USAR. Careful analysis of task constraints (implicit or explicit) may result in useful exploitation of such robotics technologies.

The focus on evacuation rather than rescue changes the role of robots significantly. We have argued that audio beacon deployment is an assistive approach to the problem of evacuation, because the evacuees are assumed to able-bodied. This also opens up an avenue for interesting human-robot interaction in

the evacuation context.

We have presented an audio beacon-based evacuation algorithm implemented in the framework of multi-robot task allocation. We have used space syntax methods both in the design and evaluation of the algorithm, and a method similar to existing flow-based approaches for the assessment of the beacon deployment approach. Although our work has not used microscopic models directly, we have described how this could be incorporated in our algorithm in terms of utility calculation. We have also proposed that a number of human motion-based methods for evaluation of environmental characteristics could be used more generally within the mobile robotics community.

REFERENCES

- [1] M. Batalin and G. Sukhatme. Coverage, exploration and deployment by a mobile robot and communication network. In *Proceedings of the International Workshop on Information Processing in Sensor Networks*, Palo Alto Research Center (PARC), Palo Alto, CA, USA, April 2003.
- [2] M. Batty. Predicting where we walk. *Nature*, 388:19–20, July 1997.
- [3] D. Begault. *3D Sound for Virtual Reality and Multimedia*. Academic Press, San Diego, CA, USA, 1994.
- [4] M. Berhault, H. Huang, P. Keskinocak, S. Koenig, W. Elmaghraby, P. Griffin, and A. Kleywegt. Robot exploration with combinatorial auctions. In *Proceedings of the International Conference on Intelligent Robots and Systems (IROS'03)*, pages 1957 – 1962, Las Vegas, NV, USA, October 2003.
- [5] J. Casper, M. Micire, and R.R. Murphy. Issues in intelligent robots for search and rescue. *SPIE Ground Vehicle Technology II*, 4:41–46, April 2000.
- [6] F. Cazals and M. Sbert. Some integral geometry tools to estimate the complexity of 3D scenes. Technical Report RR-3204, Institut National de Recherche en Informatique et en Automatique (INRIA), July 1997.
- [7] T.H. Cormen, C.E. Leiserson, R.L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, Cambridge, Massachusetts, USA, 2nd edition, 2001.
- [8] W. Daamen, P.H.L. Bovy, and S.P. Hoogendoorn. Modelling pedestrians in transfer stations. In M. Schreckenberg and S.D. Sharma, editors, *Proceedings of the First International Conference on Pedestrian and Evacuation Dynamics*, pages 59–73, Duisberg, Germany, April 2001.
- [9] Düsseldorf Airport Disaster, 10 April 1997. Official Prosecutor Investigative Report, 1997.
- [10] D. Fox. KLD-sampling: Adaptive particle filters. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems 14*, pages 713–720, Cambridge, MA, USA, 2001. MIT Press.

- [11] B.P. Gerkey and M.J. Matarić. A framework for studying multi-robot task allocation. In *Multi-Robot Systems: From Swarms to Intelligent Automata*, volume 2, pages 15–26. Kluwer Academic Publishers, 2003. presented at 2003 International Workshop on Multi-Robot Systems.
- [12] B.P. Gerkey and M.J. Matarić. A formal framework for the study of task allocation in multi-robot systems, 2004. To appear in *International Journal of Robotics Research* 2004, also Technical Report CRES-03-013, Center for Robotics and Embedded Systems, University of Southern California.
- [13] D. Goldberg and M.J. Matarić. Interference as a tool for designing and evaluating multi-robot controllers. In *Proceedings of the AAAI Fourteenth National Conference on Artificial Intelligence*, pages 637–642, Providence, RI, USA, 1997.
- [14] D. Goldberg and M.J. Matarić. Coordinating mobile robot group behavior using a model of interaction dynamics. In Oren Etzioni, Jorg P. Muller, and Jeffrey M. Bradshaw, editors, *Proceedings, The Third International Conference on Autonomous Agents (Agents '99)*, pages 100–107, Seattle, WA, USA, May 1999. ACM Press.
- [15] D. Goldberg and M.J. Matarić. Design and evaluation of robust behavior-based controllers. In Tucker Balch and Lynne E. Parker, editors, *Robot Teams: From Diversity to Polymorphism*. A K Peters Ltd, April 2002.
- [16] E.T. Hall. *The Silent Language*. Doubleday, Garden City, NY, USA, 1959.
- [17] H.W. Hamacher and S.A. Tjandra. Mathematical modelling of evacuation problems: A state of the art. In M. Schreckenberg and S.D. Sharma, editors, *Proceedings of the First International Conference on Pedestrian and Evacuation Dynamics*, pages 227–266, Duisberg, Germany, April 2001.
- [18] D. Helbing. Traffic and related self-driven many particle systems. *Reviews of Modern Physics*, 73: 1067–1141, 2001.
- [19] D. Helbing, I. Farkas, and T. Vicsek. Simulating dynamical features of escape panic. *Nature*, 407: 487–490, September 2000.
- [20] B. Hillier. *Space is the machine: A Configurational Theory of Architecture*. Cambridge University Press, 1996.
- [21] B. Hillier. The Common Language of Space: a way of looking at the social, economic and environmental functioning of cities on a common basis., August 1998. <http://www.spacesyntax.org/publications/commonlang.html> last access May 26, 2004.
- [22] B. Hillier and J. Hanson. *The social logic of space*. Cambridge University Press, 1984.
- [23] S.P. Hoogendoorn, P.H.L Bovy, and W. Daamen. Microscopic pedestrian wayfinding ad dynamics modelling. In M. Schreckenberg and S.D. Sharma, editors, *Proceedings of the First International*

- Conference on Pedestrian and Evacuation Dynamics*, pages 123–154, Duisberg, Germany, April 2001.
- [24] A. Howard. Multi-robot mapping using manifold representations. In *Proceedings of the International Conference on Intelligent Robots and Systems (ICRA'04)*, pages 4198–4203, New Orleans, Louisiana, USA, April/May 2004.
- [25] B. Jiang and C. Claramunt. Extending Space Syntax towards an Alternative Model of Space within GIS. In *3rd European Agile Conference on Geographic Information Science*, May 2000.
- [26] B. Jung and G.S. Sukhatme. Tracking targets using multiple robots: The effect of environment occlusion. *Autonomous Robots*, 13(3):191–205, November 2002.
- [27] A. Keßel, H. Klüpfel, J. Wahle, and M. Schreckenberg. Microscopic simulation of pedestrian crowd motion. In M. Schreckenberg and S.D. Sharma, editors, *Proceedings of the First International Conference on Pedestrian and Evacuation Dynamics*, pages 193–200, Duisberg, Germany, April 2001.
- [28] Klaxton Signals, 2004. *Directional localizer product description*. Available on the World Wide Web at <http://www.klaxonsignals.com/directional-sounders.htm> last access May 22, 2004.
- [29] A. Kobayashi and K. Nakamura. Rescue robots for fire hazards. In *Proceedings of the first International Conference on Advanced Robotics (ICAR'83)*, pages 91–98, September 1983.
- [30] G.G. Løvås. Models of wayfinding in emergency evacuations. *European Journal of Operational Research*, 105:371–389, 1998.
- [31] A. Mintz. Non-adaptive group behavior. *The Journal of Abnormal and Normal Social Psychology*, 46:150–159, 1951.
- [32] R.R. Murphy. Marsupial and shape-shifting robots for urban search and rescue. *IEEE Intelligent Systems*, 15(2):14–19, March 2000.
- [33] R.R. Murphy. Human-Robot Interaction in Rescue Robotics. *IEEE Transactions on Systems, Man and Cybernetics, Part C*, 43(2):138–153, May 2004.
- [34] R.R. Murphy. Rescue robotics for homeland security. *Communications of the ACM, Special Issue on emerging technologies for homeland security*, 47(3):66–68, March 2004.
- [35] R.R. Murphy, J. Casper, and M. Micire. Potential tasks and research issues of mobile robots in robocup rescue. In P. Stone, T. Balch, and G. Kraetschmar, editors, *RoboCup-2000: Robot Soccer World Cup IV*, pages 339–344. Springer Verlag, Berlin, 2001.
- [36] V.M. Predtechenskii and A.. Milinski. *Planning for foot traffic flow in buildings*. Amerind Publishing Co., New Delhi, India, 1978. Translated from the original Russian (Stroiizdat Publishers, Moscow, 1969).

- [37] G. Proulx. Misconceptions about human behavior in fire emergencies. *Canadian Consulting Engineer*, pages 36–38, March 1997.
- [38] S. Read. The grain of space in time: the spatial/functional inheritance of amsterdam’s centre. *Urban Design International*, 5(3):209–220, December 2002.
- [39] A. Schadschneider. Cellular automaton approach to pedestrian dynamics – applications. In M. Schreckenberg and S.D. Sharma, editors, *Proceedings of the First International Conference on Pedestrian and Evacuation Dynamics*, pages 87–97, Duisberg, Germany, April 2001.
- [40] A. Schadschneider. Cellular automaton approach to pedestrian dynamics – theory. In M. Schreckenberg and S.D. Sharma, editors, *Proceedings of the First International Conference on Pedestrian and Evacuation Dynamics*, pages 75–85, Duisberg, Germany, April 2001.
- [41] D.A. Shell and M.J. Matarić. Human motion-based environment complexity measures for robotics. In *Proceedings of the International Conference on Intelligent Robots and Systems (IROS’03)*, pages 2559–2564, Las Vegas, NV, USA, October 2003.
- [42] D.A. Shell and M.J. Matarić. Directional audio beacon deployment: an assistive multi-robot application. In *Proceedings of the International Conference on Intelligent Robots and Systems (ICRA’04)*, pages 2588–2594, New Orleans, Louisiana, USA, April/May 2004.
- [43] N. J. Smelser. *Theory of Collective Behavior*. The Free Press, New York, 1963.
- [44] G.K. Still. *Crowd Dynamics*. PhD thesis, Warwick University, Mathematics department, August 2000.
- [45] K. Støy, W.-M. Shen, and P.M. Will. A simple approach to the control of locomotion in self-reconfigurable robots. *Robotics and Autonomous Systems*, 44(3–4):191–200, 2003.
- [46] S. Thrun and A. Bücken. Integrating grid-based and topological maps for mobile robot navigation. In *Proceedings of the AAAI Thirteenth National Conference on Artificial Intelligence*, pages 944 – 950, Portland, OR, USA, 1996.
- [47] I. Ulrich and J. Borenstein. VFH+: Reliable obstacle avoidance for fast mobile robots. In *Proceedings of the International Conference on Intelligent Robots and Systems (ICRA’98)*, pages 1572–1577, Leuven, Belgium, May 1998.
- [48] R.T. Vaughan, B.P. Gerkey, and A. Howard. On device abstractions for portable, reusable robot code. In *Proceedings of the International Conference on Intelligent Robots and Systems (IROS’03)*, pages 2421–2427, Las Vegas, NV, USA, October 2003.
- [49] D. J. Withington. Faster evacuation from ferries using sound beacons. *Fire: Journal of the fire protection profession*, 39, March 2000.

- [50] D. J. Withington. Life saving applications of directional sound. In M. Schreckenberg and S.D. Sharma, editors, *Proceedings of the First International Conference on Pedestrian and Evacuation Dynamics*, pages 277–298, Duisberg, Germany, April 2001.
- [51] D. J. Withington. Directional sound: an important new ship-evacuation aid. *The Naval Architect, Journal of the The Royal Institution of Naval Architects*, pages 6–7, July/August 2003.
- [52] A. Wolf, H. Brown, R. Casciola, A. Costa, M. Schwerin, E. Shamas, and H. Choset. A mobile hyper redundant mechanism for search and rescue tasks. In *Proceedings of the International Conference on Intelligent Robots and Systems (IROS'03)*, pages 2889–2895, Las Vegas, NV, USA, October 2003.