Journal of Network and Computer Applications 71 (2016) 118–129

Contents lists available at ScienceDirect

Journal of Network and Computer Applications

journal homepage: www.elsevier.com/locate/jnca

Adaptive event sensing in networks of autonomous mobile agents

Rodrigo R. Esch [a], Fábio Protti [b,*], Valmir C. Barbosa [a]

[a] Programa de Engenharia de Sistemas e Computação, COPPE Universidade Federal do Rio de Janeiro, Caixa Postal 68511, 21941-972 Rio de Janeiro, RJ, Brazil
[b] Instituto de Computação, Universidade Federal Fluminense, Rua Passo da Pátria, 156, 24210-240 Niterói, RJ, Brazil

ARTICLE INFO

Article history:
Received 24 July 2015
Received in revised form
10 November 2015
Accepted 18 April 2016
Available online 7 May 2016

Keywords:
Autonomous agents
Mobile agents
Sensor networks
Cooperative control
Distributed optimization

ABSTRACT

Given a connected region in two-dimensional space where events of a certain kind occur according to a certain time-varying density, we consider the problem of setting up a network of autonomous mobile agents to detect the occurrence of those events and possibly record them in as effective a manner as possible. We assume that agents can communicate with one another wirelessly within a fixed communication radius, and moreover that initially no agent has any information regarding the event density. We introduce a new distributed algorithm for agent control based on the notion of an execution mode, which essentially lets each agent roam the target region either at random or following its local view of a density-dependent gradient. Agents can switch back and forth between the two modes, and the precise manner of such changes depends on the setting of various parameters that can be adjusted as a function of the application at hand. We provide simulation results on some synthetic applications especially designed to highlight the algorithm's behavior relative to the possible execution modes.

1. Introduction

There exist several applications in which sensing a geographical region as accurately as possible can be crucial. Examples have traditionally included search and recovery operations, manipulation tasks in hazardous environments, surveillance, and environmental monitoring for pollution detection (Gusrialdi et al., 2011), as well as several mapping and monitoring tasks of more specific interest (Yamauchi et al., 1998; Eriksson et al., 2008). Often it is the case that monitoring the region in question with only one sensor to detect events all across it can be extremely expensive if not downright infeasible, as well as hardly fault-tolerant or scalable. Along the years, though, significant technological advances, especially those making communication, processing, and sensing more cost-effective, have enabled the use of autonomous mobile agents (Borenstein et al., 1997; Mondada et al., 2004; Martinez et al., 2007; Yu and George Lee, 2011; Hu et al., 2012; Huo et al., 2013).

If used in sufficient numbers given the size of the region to be monitored, these agents offer a viable alternative to overcome the limitations of the single-sensor scenario, provided only that they can perform local sensing and communicate with one another. Once put to work together on the same global sensing task, the agents can communicate to one another the sensing results obtained with their own relatively small sensing capabilities, thus generating a far greater sensing power for the detection of the majority of events occurring in the region. The key to realize such a tantalizing possibility lies, naturally, in how successfully the agents can cooperate toward a common goal.

One key additional ingredient is then agent mobility. Mobile agents can in principle behave in several different modes, such as moving to places where no events are currently occurring so that they can be detected when they do occur; moving to places where a great number of events is currently occurring and additional sensing power is needed; or simply moving in arbitrary directions aiming to discover unanticipated places where events of interest may come to occur or be already occurring. Given a specific sensing task, a successful ensemble of agents will be one whose agents behave in the appropriate manner often enough for the task to be accomplished accurately and with just enough agents.

Here we introduce a new distributed algorithm for the operation of such a group of autonomous mobile agents. Our algorithm is based on the notion of an execution mode to guide each agent toward participating in the overall sensing task as effectively as possible while switching back and forth from one mode to another as mandated by local circumstances and parameter values. As in so many cases in which agent adaptation has a key role to play, here too we aim to strike a fruitful balance between exploration and exploitation. We do so by using two execution modes exclusively. The first one, called the random mode, is designed so that the agent can contribute to the overall task at hand by exploring new territory that so far may have remained insufficiently monitored. The second mode, referred to as the gradient mode, lets the agent exploit the best sensing opportunities available to it by focusing on

* Corresponding author.
E-mail address: fabio@ic.uff.br (F. Protti).

http://dx.doi.org/10.1016/j.jnca.2016.04.022
1084-8045/

those places at which events are deemed more likely to occur insofar as the agent's local view of the global event density allows such a conclusion.

Our work is preceded by important related research providing both improvement opportunities and inspiration (Howard et al., 2002; Zou and Chakrabarty, 2003; Li and Cassandras, 2005; Madhani et al., 2005; Wang and Ramanathan, 2005; Nene et al., 2010; Gusrialdi et al., 2011, 2013; Sheng et al., 2012; Song et al., 2013). Some of this research has addressed the problem of maximizing coverage area when the agents are placed inside the target region either deterministically (Howard et al., 2002) or at random (Zou and Chakrabarty, 2003; Nene et al., 2010), and also the similar problem of adjusting the agents' initial locations and their speeds along a closed path so that all points inside such a perimeter are monitored as thoroughly as possible (Song et al., 2013). Another related problem addressed by such research has been the convergence of all sensed data onto a base station while being mindful of the energy spent on communication (Li and Cassandras, 2005).

Although all these problems are undoubtedly related to the one we tackle, the coverage-area maximization that lies at their core is only ancillary to us, in the sense that what we seek is full event coverage even if this means leaving some portions of the target region somewhat unmonitored. We share this goal with the works in Gusrialdi et al. (2011, 2013), but solve the more general problem in which event density is both time-varying and unknown to the agents initially. Moreover, the problem is to be solved subject to the further constraints that agents need not be within direct communication reach of one another, and that no centralized element (like the leader-following procedure from Ji and Egersted (2007) that the authors of Gusrialdi et al. (2011, 2013) adopt) is to be part of the solution. To the best of our knowledge, ours is the first approach that addresses event unpredictability while abiding by these further constraints.

Once a fully distributed solution with the desired level of event coverage is available and works well even in the face of initially unknown, possibly time-varying event-density conditions, the range of potential applications will in all likelihood get expanded in a way that is hard to overstate. For example, the use of mobile robots for planetary exploration that is already a reality[1] will have a chance to be stepped up from the current adoption of standalone robots (Chatila and Lacroix, 1995; Schenker et al., 2001) to that of the robot teams envisaged a few years ago (Kisdi and Tatnall, 2011) but still not implemented. In a related vein, it also seems reasonable to expect that other settings subject to similarly extreme environmental conditions will find themselves amenable to exploration by robot teams. Such settings include sites of nuclear waste (Nawaz et al., 2010), places occupied by hard-to-detect substances (Oyekan and Hu, 2013), and underwater locations (Bodi et al., 2015).

Other application areas in which such a solution can make a crucial difference are those of automated farming in general (Zecha et al., 2013; Anil et al., 2015), including weed management in cropping systems (Young et al., 2013) as well as the imitation of other forms of foraging behavior (Ferrante et al., 2015), and the deployment and operation of the so-called smart dust. The latter term is used to refer to typically very large ensembles of small-scale autonomous robots, whose realization has for nearly two decades been recognized as fraught with difficulties (Kahn et al., 1999) and only more recently is becoming feasible (Lücking et al., 2012; Tittl et al., 2013). In fact, the recent demonstration of self-assembly capabilities in swarms of relatively small robots (Rubenstein et al., 2014) is further evidence that the technology is at last becoming ripe.

We proceed in the following manner. First we lay down our assumptions regarding both the events to be monitored and the agents that will monitor them. We do this in Section 2. Then we move to a detailed description of our solution in Section 3. Simulation results on a few key scenarios are given in Section 4, which is then followed by conclusions in Section 5.

2. System model

The target region is modeled as a connected set $\Omega \in \mathbb{R}^2$. In this region, an event happens at the infinitesimal vicinity dq of point $q \in \Omega$ with probability proportional to $\phi(q, t)\mathrm{d}q$, where t is continuous time and $\phi(q, t)$ is a time-dependent event density function such that $\int_\Omega \phi(q, t)\mathrm{d}q < \infty$ for all $t \geq 0$. The function ϕ is unknown to all agents initially.

We assume that the occurrence of an event leaves a footprint that only disappears after a number of time units given by *VisTime*, a parameter, have elapsed. Upon coming across such a footprint, an agent is capable of estimating the corresponding event's time of occurrence. A useful example here is the indirect detection of a past fire at a certain point in Ω by means of the temperature at that point when it is reached by an agent. We refer to *VisTime* as an event's visibility time and assume it is the same for all events, given a specific domain of interest.

The probability that agent i, located at point $s_i \in \Omega$, detects an event (or its footprint) occurring at point $q \in \Omega$ is modeled as

$$p_i(q) = \begin{cases} \left(1 - \dfrac{d_i}{R_s}\right)^2 & \text{if } d_i \leq R_s \\ 0 & \text{otherwise,} \end{cases} \tag{1}$$

where d_i is the Euclidean distance between points s_i and q and R_s is the sensor's maximum sensing range. That is, we assume that the agent's sensing capability decays with the distance to the point of occurrence as a convex parabola, provided this point is located within the circle of radius R_s centered at the agent's location. No sensing is possible outside this circle. Whenever sensing is possible, we assume it to be instantaneous.

Agents can communicate with one another wirelessly. When agent i is at point s_i, any message it sends is assumed to be instantaneously received and successfully decoded by some other agent j, located at point s_j, with probability

$$r_i(s_j) = \begin{cases} 1 & \text{if } d_{ij} \leq R_c \\ 0 & \text{otherwise,} \end{cases} \tag{2}$$

where d_{ij} is the Euclidean distance between agents i and j and R_c is the maximum separation between two agents across which they can still communicate. We remark that, should an undirected graph be used to model the communication possibilities among all agents at a certain point in time, such a graph would follow what is known as the Boolean model. This model is widely used in theoretic studies (cf., e.g., Penrose, 2003; Peres et al., 2013) and mandates the existence of an edge between two geometrically positioned vertices if and only if, in our terms, they are no farther apart from each other than the distance R_c.

For simplicity's sake, in this work all agents are assumed to be identical, thence the values of R_s and R_c are the same for all agents. Additionally, we note that, as will become clear in the sequel, agents will tend to stand still at their current locations for far longer than the time they spend moving. For this reason, another simplifying assumption we make is that, whenever agents move, they do so instantaneously.

[1] http://solarsystem.nasa.gov/missions/target/mars.