Decentralized Control and Agent-Based Systems in the Framework of the IRVS

Valentino Crespi[1], George Cybenko[1]
Daniela Rus[2]

[1] Institute for Security Technology Studies
Thayer School of Engineering
Dartmouth College, NH 03755 USA
{valentino.crespi,gvc}@dartmouth.edu
[2] Computer Science Department
Dartmouth College, NH 03755 USA
rus@ruapeuh.dartmouth.edu

Abstract

We consider control problems that arise in the context of the IRVS. Specifically, we study the dynamics of systems of fully autonomous agents that are expected to optimize a global potential function of which they have only *partial* knowledge. Our ultimate goal is to characterize mathematically the global performance of agent-based systems in which the inter-agent cooperation is based on local exchange of information. First we introduce what we call the Opera Problem, a natural abstraction of a class of motion problems arising commonly in systems of autonomous vehicles. We then provide a *centralized* solution based on traditional control theory techniques. Finally, we present and analyze a decentralized solution.

1 Introduction

The development of decentralized control systems may be interpreted as an attempt to create systems of agents that somehow are able to perform reasonably well with respect to some global criteria even though they possess only a limited visibility over the entire system [3]. The suppression of the centralized "omniscient" authority, that typically characterizes classical control systems, forces the autonomous agents to decide on the basis of what they can locally perceive. We need to understand the interrelation between this global controller and what can be inferred after a sequence of local interactions between the agents in order to explain how and why the system will be globally well-behaved.

Similar questions about the general problem of understanding how locally exchanged information propagates through a system, have already been object of investigation in different areas for decades. Two problems discussed in this previous work are to be particularly relevant to our context: the theory of optimal convergence of load-balancing schemes for globally sparse interconnection networks (see for example [4]) and the theory of rapidly mixing Markov chains in connection with the Metropolis Algorithm [8].

The load balancing problem can be formulated as follows. Suppose we have a connected distributed network of processors which are assigned a work load. The goal is to reassign tasks so that each processor will perform an equitable share of the total work load. Moreover this redistribution of duties would be realized through a diffusion scheme, i.e., a discrete sequence of information exchanges between directly linked processors that, under proper conditions, converges to global uniformity.

The Metropolis algorithm [6] is a procedure for sampling from a typically large set of objects Ω according to a specified probability distribution π. The algorithm requires the user to define a connected, aperiodic Markov chain $K(x, y)$ on the set Ω, with the only constraint that $K(x, y) > 0$ implies $K(y, x) > 0$. The chain is then modified with auxiliary coin tossing to a new chain M with stationary distribution π.

The convergence of a load-balancing scheme is basically equivalent to the convergence of a Markov chain that derives from the multiprocessor network, *mutatis mutandis*. This means that the two traditions are complementary. Moreover, in both cases, it is vital, from the point of view of the applications, that the diffusion process under consideration, (i.e., the iterative exchange of work load between processors and the evolution of the Markov chain from a random initial state) converges as fast as possible to its limit.

In the framework of load-balancing schemes we would be talking about optimal diffusion [4] whereas in the framework of the Metropolis algorithms we would be talking about rapidly mixing Markov chains [7]. Those two problems can be reformalized in graph theoretic terms if we associate the nodes with the processors or the states of the Markov chain and the edges with the connection links in the network or with the transition arcs in the Markov chain. Several investigations, based on spectral analysis, reveal that the speed of convergence is heavily affected by the graph connectivity structure. In particular, it turns out that this velocity feature depends on the expansion properties of the graph.

We can view the convergence of such diffusion processes as the property that locally exchanged information will eventually spread out. As a consequence the system will somehow self-stabilize around a global optimum. The speed of convergence instead is more related to how fast this information spreads. And in the context of agent-based systems this not only implies efficiency issues but it might lead to unexpected misbehaviors due to temporal constraints. This corresponds to situations in which “good behavior” is displayed when a certain discrete process reaches its asymptotic stability. And this is required to happen in time before the occurrence of some events.

For example we would like a set of vehicles proceeding along a short link with a certain velocity to assume an ordered formation in a relatively short time. The study of formation control is quite common in the area of robotics even though at the moment there are mainly empirical results [1]. It would be interesting to provide mathematical grounds to several schemas that might have several applications in the IRVS [3, 9–11], considering that platooning is just an elementary and quite limited vehicle formation.

In this paper we discuss several flavors of centralized and decentralized control for moving agents such as the autonomous cars in the IRVS. In Section 2 we define the *Opera Problem* and a stochastic process. In the opera problem, a collection of agents move towards a single point, the exit from the opera house. The goal is to coordinate the exit of each agent through the door. This problem is a useful abstraction for a class of motion problems that arise in the IRVS. Several flavors of the general opera problem have direct impact on the IRVS. We observe that by changing the formulation in the opera problem so that the agents converge on a line direction rather than a point, we achieve the platooning behavior required by the IRVS. Additionally, by applying directly the Metropolis algorithm, we can encode the goal of establishing a desired interplatoon distance for vehicles already aligned and platooned. The idea is to set up a rapidly mixing Markov chain whose limit distribution represents the factions of the final interplatoon distances. In Section 2.2 we provide a classical centralized solution to the Opera problem, based on a global potential function to be optimized by the various agents. In Section 2.3 we develop a decentralized algorithm that is basically a variation of the classical solution and we prove its convergence. In Section 2.4 we show how the opera problem is related to the agent taxonomy that we defined in [3] and we discuss criteria in order to quantify the agents performance [5]. Finally, our discussion and future steps are presented in Section 3.

2 The Opera Problem

The Opera Problem can be defined as follows. Suppose n people wish to get out of a room through a small door which allows only few or even just one person to pass at each time. We would like to model the movements of the people in the process of approaching and crossing the gate.

We believe that this "Opera Problem" provides the mathematical foundation for a class of distributed algorithms that have an important role in the design of Planning and Information Agents for the IRVS. Planning Agents will synthesize an action that can be translated into a moving direction for the vehicle. We wish to control a set of vehicles towards platooning in a decentralized fashion. Specifically, the control will use only locally retrievable information gathered by Information Agents through a limited range interaction with neighboring vehicles or with the local link layer agents. Information Agents are required to be accurate enough in order for the planning algorithm to achieve a meaningful result.

In this section we begin with a classical formalization of the opera problem that leads to minimizing a global potential function with a gradient descent method. Then we derive a decentralized algorithm that allows each individual in the room to compute a descent direction for the potential using only local information. Finally we prove that this method converges.

2.1 Preliminaries

We want to formalize the Opera Problem as a nonlinear optimization problem of the type:

$$\min\{f(x) \mid x \in \mathcal{R}^n\}$$

where $\mathcal{R}^n$ denotes the n-dimensional Euclidean space and $f : \mathcal{R}^n \rightarrow \mathcal{R}$ is a continuously differentiable function verifying the Lipschitz condition, i.e.,

$$\exists L > 0 \; \forall x, y \in \mathcal{R}^n \; \|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\|.$$

One typical way to determine a solution is through an iterative procedure that starts from a initial point in the space and follows a descent direction until a local minimum. If we consider the sequence $\{x_t\}$ defined as:

$$\begin{cases} x_0 & = y_0 \in \mathcal{R}^n, \\ x_{t+1} & = x_t - \gamma_t \cdot \nabla f(x_t). \end{cases}$$

Then, under mild conditions on $\{\gamma_t\}$, it is possible to prove that either $f(x_t) \rightarrow \overline{y} \in \mathcal{R}$ and $\nabla f(x_t) \rightarrow 0$ or $f(x_t) \rightarrow -\infty$. We can assume only approximate knowledge of the gradient. In this case, the formulation is changed by replacing the term $\nabla f(x_t)$ in the equation with the term $s_t + w_t$ where s_t is an estimate of the gradient on x_t and w_t is a perturbation vector. This leads to the class of algorithms called *gradient methods with errors*, which have been studied by [2].

In [2], the following result is proven:

Theorem 1 (Bertekas and Tsitsiklis [2]). *Let x_t be the sequence generated by the method*

$$x_{t+1} = x_t + \gamma_t(s_t + w_t),$$

where

a) *the steps γ_t are long enough to guarantee a significant advance at each time even though as $t \to \infty$ they nullify:*

$$\sum_{t=0}^{\infty} \gamma_t = \infty, \quad \sum_{t=0}^{\infty} \gamma_t^2 < \infty;$$

b) *the descent direction s_t is in norm proportional to the gradient and it is not "too far" from it, i.e., there exist $c_1, c_2 > 0$ such that for all t:*

$$\|s_t\| \leq c_1(1 + \|\nabla f(x_t)\|);$$

$$-\nabla f(x_t)' \cdot s_t \geq c_2 \|\nabla f(x_t)\|^2;$$

c) *the perturbation w_t is proportional in norm to the gradient, i.e.:*

$$\exists p, q > 0 \quad \forall t \quad \|w_t\| \leq \gamma_t(p + q\|\nabla f(x_t)\|).$$

Then either $f(x_t) \to -\infty$ or there exists $\overline{y} \in \mathcal{R}$ such that $f(x_t) \to \overline{y}$ and $\nabla f(x_t) \to 0$.

2.2 A descent method for the Opera Problem

Consider an instantiation of the Opera Problem. Suppose n people $P_1 = (x_1, y_1), P_2 = (x_2, y_2), \ldots, P_n(x_n, y_n)$, scattered inside a bound environment, start moving altogether to approach the unique small exit door. We would like each person to walk toward the door and at the same time avoid impolite collisions with other people.

A suitable potential function $E : \mathcal{R}^{2n} \to \mathcal{R}$ might be the following (other artificial potential functions are being considered, but at this stage we can take just one demonstrative example):

$$E(x_1, x_2, \ldots, x_n, y_1, y_2, \ldots, y_n) = \sum_{i=1}^{n} \|P_i - O\| + 2\beta \cdot \sum_{i<j}^{n} \frac{1}{\|P_i - P_j\|^\alpha + c},$$

where $O = (0, 0)$ corresponds to the exit door and $c > 0$ causes the second term to remain bounded. The first summation clearly gets smaller as the points approach the origin whereas the second summation represents a *penalty* to ensure that people are not going to collide during the motion.

The classical gradient descent method for E would be given by

$$P(t+1) = P(t) - \gamma_t \nabla E(P(t))$$

where $P(t) = (x_1(t), x_2(t), \ldots, x_n(t), y_1(t), y_2(t), \ldots, y_n(t))$ is the state of the system at time t. We can immediately observe that, since

$$\frac{\partial E}{\partial x_i}(\underline{x}, \underline{y}) = \frac{x_i}{\sqrt{x_i^2 + y_i^2}} - \alpha\beta \sum_{j=1, j\neq i}^{n} (x_i - x_j) \cdot \frac{\|P_i - P_j\|^{\alpha-2}}{(\|P_i - P_j\|^\alpha + c)^2}$$

and

$$\frac{\partial E}{\partial y_i}(\underline{x}, \underline{y}) = \frac{y_i}{\sqrt{y_i^2 + y_i^2}} - \alpha\beta \sum_{j=1, j\neq i}^{n} (y_i - y_j) \cdot \frac{\|P_i - P_j\|^{\alpha-2}}{(\|P_i - P_j\|^\alpha + c)^2}$$

then the single agent A_i associated with person P_i, that at time t has to compute its new position $x_i(t+1), y_i(t+1)$, with

$$x_i(t+1) = x_i(t) - \gamma_t \cdot \frac{\partial E}{\partial x_i}(\underline{x}(t), \underline{y}(t)),$$
$$y_i(t+1) = y_i(t) - \gamma_t \cdot \frac{\partial E}{\partial y_i}(\underline{x}(t), \underline{y}(t)),$$

would need global knowledge about the positions of all the other people in order to evaluate its part of $\nabla E(\underline{x}, \underline{y})$.

2.3 A decentralized approach

The key point in decentralizing the gradient descent method described above is that the agent A_i associated with person P_i needs to estimate $(\frac{\partial E}{\partial x_i}(\underline{x}, \underline{y}), \frac{\partial E}{\partial x_i}(\underline{x}, \underline{y}))$ using only local information.

Suppose that each agent A_i can detect the position of all the people that are within a certain range, say ρ. This means that the computation of $(\frac{\partial E}{\partial x_i}(\underline{x}, \underline{y}), \frac{\partial E}{\partial x_i}(\underline{x}, \underline{y}))$ will be limited only to the points that are within distance ρ from P_i and this will be the core of our method.

In this section we discuss analytical techniques that might be applied to study the performance of our decentralized algorithm and we sketch an initial prove of convergence that applies for particular values of the parameters α and β.

First we may observe that our modified gradient descent method is indeed a gradient descent method with errors. So, results like theorem 1, mentioned in the preliminaries, might provide the kind of mathematical tools we need to achieve the strongest results, i.e., convergence for the largest possible range of the parameters.

Formally, each agent A_i, for $i = 1, 2, \ldots, n$, computes its new position as

$$x_i(t+1) = x_i(t) - \gamma_t \left(\frac{x_i(t)}{\|P_i(t)\|} - \alpha\beta \sum_{\|P_i(t)-P_j(t)\| \le \rho} (x_i(t) - x_j(t)) \frac{\|P_i(t) - P_j(t)\|^{\alpha-2}}{(\|P_i(t) - P_j(t)\|^{\alpha} + c)^2} \right)$$

$$y_i(t+1) = y_i(t) - \gamma_t \left(\frac{y_i(t)}{\|P_i(t)\|} - \alpha\beta \sum_{\|P_i(t)-P_j(t)\| \le \rho} (y_i(t) - y_j(t)) \frac{\|P_i(t) - P_j(t)\|^{\alpha-2}}{(\|P_i(t) - P_j(t)\|^{\alpha} + c)^2} \right)$$

This can be rewritten as

$$x_i(t+1) = x_i(t) - \gamma_t \left(\frac{\partial E}{\partial x_i}(\underline{x}(t), \underline{y}(t)) - \mathcal{E}_{x_i}(\underline{x}(t), \underline{y}(t)) \right),$$
$$y_i(t+1) = y_i(t) - \gamma_t \left(\frac{\partial E}{\partial y_i}(\underline{x}(t), \underline{y}(t)) - \mathcal{E}_{y_i}(\underline{x}(t), \underline{y}(t)) \right),$$

where

$$\mathcal{E}_{x_i}(\underline{x}(t), \underline{y}(t)) = \alpha\beta \sum_{\|P_i - P_j\| > \rho} (x_i(t) - x_j(t)) \frac{\|P_i(t) - P_j(t)\|^{\alpha-2}}{(\|P_i(t) - P_j(t)\|^{\alpha} + c)^2}$$

$$\mathcal{E}_{y_i}(\underline{x}(t), \underline{y}(t)) = \alpha\beta \sum_{\|P_i - P_j\| > \rho} (y_i(t) - y_j(t)) \frac{\|P_i(t) - P_j(t)\|^{\alpha-2}}{(\|P_i(t) - P_j(t)\|^{\alpha} + c)^2}$$

A second step would be to proceed and try to prove that the error that is introduced after the transformation that has turned the initial gradient descent algorithm into its decentralized version can still be dominated. The idea behind this is that, although it is true that each agent uses only local information to learn about the global state and so update its new destination, we may argue that as the system evolves such an error would ultimately vanish.

Let us now observe the following.

Lemma 1. *Let $\alpha > 1$ and $c > 0$. Then for all $\rho > 0$ and for all times t, the error vector is always bounded in norm.*

Proof. We are going to show that the perturbation vector is in norm bounded by a constant term independent on the norm of the gradient. Let $f(x) = \frac{\alpha x^{\alpha-1}}{(x^\alpha+c)^2}$ and

$$M(\alpha) = \max_{x>0}\{f(x)\} = \max_{x>0}\left\{\frac{\alpha x^{\alpha-1}}{(x^\alpha + c)^2}\right\}.$$

Notice that the value M is attained by $x_0 = (c(\alpha-1)/(\alpha+1))^{\frac{1}{\alpha}}$ and we assume $\rho > x_0$. By the properties of the moduli we can derive the following chain of inequalities:

$$\begin{aligned} \mid \mathcal{E}_{x_i}(\underline{x},\underline{y}) \mid &\leq \beta \sum_{\|P_i - P_j\| > \rho} f(\|P_i - P_j\|) \\ &\leq \beta \sum_{\|P_i - P_j\| > \rho} \min\{M(\alpha), f(\rho)\} \\ &\leq \beta(n-1) \cdot A(\alpha, \rho). \end{aligned}$$

It is immediate to see that also $\mid \mathcal{E}_{y_i}(\underline{x},\underline{y}) \mid \leq \beta(n-1)A(\alpha,\rho)$. Thus it follows that

$$\|\mathcal{E}(\underline{x}(t),\underline{y}(t))\| = \sqrt{\sum_{i=1}^{n} \mathcal{E}_{x_i}(\underline{x}(t),\underline{y}(t))^2 + \sum_{i=1}^{n} \mathcal{E}_{y_i}(\underline{x}(t),\underline{y}(t))^2} \leq \beta\sqrt{2n}(n-1)A(\alpha,\rho).$$

Now, we can introduce our result:

Theorem 2. *Let $\gamma_t = 1$ for all $t > 0$. Then there exists $\beta_0 = \beta_0(\alpha, n)$ such that all the agents reach the gate within a finite interval of time.*

Sketch The idea is the following. Since

$$\max_{0 \leq x \leq \rho}\{f(x)\} \leq M(\alpha),$$

we may easily establish a threshold β_0 such that for smaller values of β the vector determined by the main goal term: $P_i/\|P_i\|$, having unitary norm, will always prevail over the penalty term having norm upper bounded by $\beta(n-1)M(\alpha)$. This will cause each agent to reduce, at each step, its distance from the gate of a nonnegligible value.

The intuition behind the above result is that the global information each person needs pertains to only the penalty term in the potential function and this can always be dominated by the main goal term for particular values of β. Moreover, people far from each other do not contribute significantly to the respective penalty term in the gradient of the potential, because the chances that they will cause a collision in the short-term future are quite slim.

2.4 Planning and Information Agents

From the previous analysis we were able to conclude that our decentralized method converges for all possible positive values of the parameter ρ. We have not said anything about how the extent of the perturbation affects the motion of the people so far. It is clear that for very large values of ρ such a perturbation would be very small suggesting an efficient convergence process. On the other hand if we have too small values for ρ, the people will mass very quickly in front of the gate and eventually they would get stuck because of lack of maneuvering space.

Clearly, the perturbation $\mathcal{E}$ depends in general on both the range capacity ρ of the sensors and their precision, say ϵ. So measuring $\mathcal{E}$, through for example the parameters ρ and ϵ, will determine the quality of the Information Agents involved in the detection of the position of the person in question and its neighbors. This will, in turn, allow a quantification of the performance of the Planning Agent.

3 Future Work

Here is a list of steps that we believe are important for our deeper understanding of the opera problem and its connection to the IRVS.

- We are developing simulations of the decentralized algorithms in order to provide experimental evidence of their performance.
- We plan to study other interesting abstractions of motion problems. In particular we will investigate generalizations of the Opera Problem, in which vehicle alignment issues and non trivial formations are considered.
- We would like to investigate how to discretize the continuous Opera-like formulation of the motion problems for the vehicles.

References

1. Balch, T., Arkin, C.: Behavior-based Formation Control fot Multi-robot Teams. IEEE Transactions on Robotics and Automation, Dec. (1998)
2. Bertsekas, D.,P., and Tsitsiklis, J., N.: Gradient Convergence in Gradient Methods. SIAM J. on Optimization, Vol. 10, pp. 627-642 (2000)
3. Crespi, V., Cybenko, G.: Agent-Based Systems Engineering and Intelligent Vehicles and Road Systems. TASK Report (2001)
4. Cybenko G.: Dynamic Load Balancing for Distributed Memory Multiprocessors. Journal of Parallel and Distributed Computing, 7, 279–301 (1989)
5. Cybenko, G.: Agent-Based Systems Engineering. TASK Program Technical Proposal (2000)
6. Diaconis, P., Saloff-Coste, L.: What do we know about the Metropolis Algorithm? In Proceedings of STOC 95, Las Vegas, Nevada (1995)
7. Jerrum, M., Sinclair, A.: The Markov chain Monte Carlo Method: an Approach to Approximate Counting and Integration. In D.S. Hochbaum, ed., Approximation Algorithms for NP-hard problems, PWS Publishing, Boston (1996)
8. Metropolis, N., Rosenbluth, A., Rosenbluth, M., Teller, A., Teller, E.: Equations of state calculations by fast computing machines. J. Chem. Phys., 21, 1087–1092 (1953)
9. Frazier, T., Newman, A.: Intelligent Roadway Vehicle System. REF for the DARPA ISO TASK Program (2000)
10. Frazier, T., Newman, A., Roszman, L.: IVRS REF Scenarios. DARPA Progress Report, Dec (2000)
11. Varaiya, P. and Horowitz, R.: Control Design of an Automated Highway System. In Proceedings of the IEEE (2000)