Collective Locomotion

Pierre Arnaud
Laboratoire de Micro-Informatique, LAMI–DI–EPFL
Swiss Federal Institute of Technology, CH-1015 Lausanne
Pierre.Arnaud@di.epfl.ch — http://diwww.epfl.ch/lami/team/arnaud/

0. Abstract

In this talk, I will discuss how collective locomotion, inspired by biological models (in our case animals), can be applied to the navigation of so called agents (in our case autonomous mobile robots). I will compare the relative cost and performance of the behaviours observed in groups of animals, such as herds, flocks, schools, formations, and so on. Finally, I will conclude with an overview of a simulator currently under development at the LAMI, the Micro-computing lab at the Swiss Federal Institute of Technology in Lausanne.

1. Introduction

We are studying the evolution of collective behaviours involving numerous miniature autonomous mobile robots : our aim is to set up experiments with more than 10 Khepera robots.

[slide Khepera]

The Khepera was initially developed at the LAMI and is now being distributed by the K-team.

In our work, we study the collective behaviours displayed by animals which involve locomotion. We try to take advantage of these observations by applying them to mobile robots.

[slide R2D2]

We hope that the inspiration found in animal navigation strategies will help us reduce the complexity of tasks such as path planning in environments scattered with possibly moving obstacles.

2. Collective locomotion

Collective locomotion can be viewed as a subset of collective robotics.

[slide Group A-B]

It focuses on the specific tasks required to move a group of autonomous agents, either virtual or physical robots, from a point A to a point B.

We apply the concept of collective locomotion to robots, animals and human beings. It deals with the following topics :

[slide List of Topics]

- navigation
- distributed path planning
- inter-individual communication
- co-operation
- and auto-organisation

Collective locomotion can be instinctive or accidental :

- Individuals form groups on purpose or instinctively as the result of an evolutionary pressure : for example herds of grazing animals, schools of fish, flocks of birds, etc. get some benefit from gathering into groups with other members of the same species.

  [slide Sheep 1]

  These benefits are higher protection against predators,

  [slide Sheep 2]

  better food localisation

  [slide Sheep 3]

  or facilitated mating. Visual, tactile or auditory communication are often involved.

- Individuals may also display accidental — or unorganised — collective locomotion when crowded in a small area or when moving on a narrow path (for example ants following a pheromone trail or motorists going on holiday).

  [slide Boule & Bill]

  Each individual tries to take advantage of the flows naturally appearing in the group in order to get to its destination with the lowest cost, like people in crowds. This inevitably leads to involuntary inter-individual interactions and thus to some form of communication, either direct or through the modification of the environment.

Our work focuses on the study of the first category of collective locomotion, since we believe that it will provide interesting solutions to complex path planning problems involving numerous mobile robots.

Fleets of robots might be used in factories with unstructured environment or in open air. They will require specific solutions in order to avoid excessive computation time and to be able to face uncertainties found in dynamic environments.

[slide Télécommande]

We reject centralised control systems for collective locomotion, since these have many restrictions and drawbacks which are not acceptable for real-world robotics. In centralised control systems, a supervisor needs to communicate permanently with each individual.

With a growing number of robots, the limited bandwidth does no longer allow real-time reactions. Moreover, the system cannot be considered to be robust, since a failure of the supervisor usually produces a generalised breakdown of the system.

Our research is based on **d**istributed **a**utonomous **r**obotic **s**ystems : each robot interacts only with its close neighbours and operates on its own, without external supervision. We chose this system, since animals also behave in a distributed and autonomous way.

3. A taxonomy for collective locomotion

Animals exhibit many different collective locomotion patterns depending on the number of individuals in the group, on their size, speed, morphology or habits, on the environment and on the task to be accomplished.

We use the following taxonomy to name the different behaviours involving the locomotion of a group of individuals (let it be animals, humans or robots) :

- **Herds**.

  [slide Sheep]

  The individuals move approximately in the same direction, usually following some leader. This mainly applies to grazing or slowly moving herbivorous animals.

  Herds are quite robust configurations, since a superficial skin effect tends to keep the individuals together. They also allow, at rest, to keep watch in every direction.

  The reaction time for changing the group's direction is short. Herds cannot move very quickly in compact form, since collision avoidance would prove difficult without increasing the inter-individual distance. Flocks and schools offer a better solution to fast locomotion.

- **Flocks**.

  [slide Flock of birds]

  The individuals move in the same direction. A change of direction can be initiated by any individual. The reaction times are short and the co-ordination is very precise, as can be observed with birds — for example starlings (*Staren, étourneaux*).

  [slide Birds schema]

  Flocks are able to move very quickly in compact form thanks to shorter reaction times, gained through long distance sensing : an individual watching not only its direct neighbours, but also more distant ones, will be able to anticipate changes, thus reacting very efficiently to changes in velocity or direction. This provides an excellent group reaction time.

- **Schools**.

  [slide School]

  The individuals maintain the same orientation synchronously. This is called polarisation. They stay at a constant inter-individual distance. A change of orientation can be initiated by any individual. Geometric patterns may be observed with certain species of fish.

  Schools can be considered as a special kind of flocks in which the relative orientation of the individuals is primordial. Whereas individuals in flocks cannot change their direction on the spot, schools can orient themselves at rest. This is the reason for very fast and important changes in direction.

  [slide Flash explosion]

  Schools can easily change their direction, split, change their shape or explode in every direction.

- **Formations**.

  [slide Goose]

  The individuals maintain a constant geometric relationship while following a leader. Migrating birds often fly in formations to reduce their energy expenditures.

  In addition, formations might allow the rear individuals to keep watch of the preceding ones in hazardous environments. This property is used by the army.

  Formations are not well adapted to frequent changes of direction, because of the high angular inertia caused by the geometric arrangement. They are best suited for long distance linear travels (such as migrations) on open terrain (or in the air), where no large obstacles or predators have to be avoided.

- **Fronts**.

  [slide Caterpillar front]

  The individuals move side by side, maintaining the same direction. This behaviour is well adapted to grazing and harvesting. Groups moving as fronts cover a maximal area.

- **Processions**.

  [slide Processions]

  The individuals line up and follow each other. The first individual of the procession acts as the leader. This very simple behaviour is extensively used by some caterpillars and by some lobsters. Aquatic animals such as lobsters significantly reduce the underwater drag by lining up.

  Processions are also observed in groups of animals travelling on narrow paths.

  Only the leader needs to find an appropriate path, thus reducing the overall energy used for exploration and path planning. The leader may be a specialised individual (equipped with better sensors, more knowledge or higher autonomy).

A procession can easily be split, but the start of a new procession will take time : the first individual should not start moving too quickly until the whole queue is ready to move, or else individuals lagging behind would be lost in the operation.

The velocity changes among the group propagate like waves; a parallel can be drawn with the flows of automobiles driving on highways.

- **Swarms**.

  [slide Swarm]

  The individuals keep very close to each other in order to maximise inter-individual contact and minimise the volume occupied by the group. This forms balls of individuals. Honey bees move in swarms in order to protect their queen, located at the very centre of the swarm; fish shoals also form balls to protect themselves against predators.

  Swarms can only move slowly, because of the low inter-individual distance and the high collision probability.

  When a swarm of bees has to move, some scouts go out and explore the neighbouring environment in order to choose the best destination (or the best path). Each scout then tries to convince the swarm of its own choice being the best by broadcasting it. The most advertised choice wins. This complex behaviour finds the best possible path while spending only little energy in exploration. The scouts are the only individuals needing to have specialised sensors; the other individuals will just have to follow.

- **Unstructured groups**.

  [slide Ducklings]

  The individuals move around while keeping in touch with the other members of the group. This can be considered as a very loose herd and is observed mainly in small groups (for example offspring moving around a parent). Unstructured groups are not well adapted for the locomotion of large populations.

Our classification provides an idealised view of collective locomotion. Most natural behaviours do not fit exactly into a single category. The dynamic nature of these behaviours often lead to intermediate, possibly changing, solutions within the boundaries defined by these categories.

4. Collective locomotion and robotics

Research involved in collective locomotion applied to real world robots is still marginal. Most of the papers related to collective locomotion originate from scientists working on disciplines such as computer graphics (*for example Craig Reynolds*) and artificial life (*essentially works from Hodgins or Terzopoulos*), which are both involved with the simulation of artificial creatures moving in virtual environments.

[slide Artificial fishes]

In the past few years, distributed autonomous robotic systems have met an increasing interest in the scientific community. Problems such as distributed path planning, multiple robot co-ordination, co-operation and collaboration are being actively investigated, but most research is being conducted through simulation. Experiments involving groups of real mobile robots are still extremely rare.

The main contribution to collective locomotion involving real mobile robots, originates from Maja Mataric, who worked with groups of mobile robots exhibiting what she calls a *flocking behaviour*. We would rather name this *herding*, according to our definitions, since the robots only show a simple following behaviour.

Lynn Parker also studied group locomotion while developing a general behaviour-based architecture, called ALLIANCE, which results in a — quote — “fault tolerant, reliable, adaptive and coherent mechanism for co-operative robot control” — unquote.

5. Advantages and drawbacks of the different behaviours

I will now compare the behaviours mentioned in our taxonomy.

[slide Table 1]

These behaviours impose different requirements on the sensors, on the actuators and on the available computational power.

From an economical point of view, behaviours involving low sensor complexity, low actuator responsiveness and low computational power are the most interesting, since the corresponding robots are cheaper to manufacture.

Formations, fronts, processions and swarms match these criteria. It is interesting to note that these behaviours are mainly exhibited by insects, which are much more primitive than vertebrates, and have indeed less complex “sensors” and a lower “computational power”.

The behaviours I have just reviewed have advantages and drawbacks.

[Slide Table 2].

The **speed** of the group and its **angular responsiveness** are mainly determined by the reaction time and mobility of the individuals. Processions can be viewed to have a high angular responsiveness, since a change in direction is possible simply by changing the procession into a front.

The **ease of splitting** mainly depends on the shape and on the angular responsiveness of the group.

The **overall group responsiveness** gives some idea about its reaction time, which can be assimilated to the propagation speed of changes.

The **fitness to move in encumbered environments** is related to the group's ability to move around obstacles : it gives some indication about its fluidity.

A designer might want to choose the group's behaviour in advance and build robots with the appropriate sensors, actuators and computational power. Polyvalent robots will often prove to be a better choice than specialised ones — even if their construction is more expensive — because they will be able to adapt to their environment and choose the behaviour best suited to their situation.

6. Robustness

Applications involving numerous robots may require high speed, efficient obstacle avoidance, maximal flexibility, minimal energy consumption, low cost, etc. but they all expect a high robustness of the overall system.

The dysfunction of one or several individuals should not impair the group’s ability to carry out the assigned mission; furthermore, the overall performance should only degrade slightly when a partial system failure occurs.

Robustness can be achieved through distributed decision making — a partial failure will not immobilise the whole group — and by avoiding fixed hierarchical organisation; this means that no robot should have to rely on orders issued by a single individual, unless it can be dynamically replaced by any other individual.

Collective locomotion can be made more robust at the cost of some concessions, for example by diminishing the group's speed, increasing the complexity of the sensors or implementing smarter algorithms.

A system providing the desired robustness might either require inter-individual communication — in order to arbitrate local conflicts, detect the dysfunction of individuals, obtain information about the other individual's state, etc. — or implement some sense of “impatience”, as has been done by Lynn Parker.

It is interesting to note that the behaviours found in nature are not all intrinsically robust, as can be observed with processions of insect larvae (the caterpillar *thaumetopoea pinivora*). If the leading larva reaches the end of a procession, it simply follows it. This works well as long as the leader does not join its own procession's rear, which results in the formation of a loop. This has been reported with much detail by the entomologist Jean-Henri Fabre.

7. The simulator

We plan to set up collective locomotion experiments based on the Khepera. In order to test our algorithms and various scenarios, we need a computer simulation platform supporting multiple robots.

[slide Simulator]

The simulator has to be :

- generic
- highly flexible, because we want to easily add sensors, actuators and algorithms
- realistic, this means that the simulator uses physical laws, does not serialise simultaneous events, introduces randomness caused by noisy sensors and imperfect actuators, etc.
- portable.

Furthermore, it should not require too much computational power, in order to run both on workstations and on desktop computers. Simulators supporting multiple robots exist, but are either aimed at high-end workstations or are too simple to support our needs.

— the generic computer simulation platform for distributed robotic system development and experiments presented by Jing Wang is very similar to our simulator, but it runs on a network of SGI super computers —

We therefore chose to develop our own multiple robot computer simulation program.

We decided to take an object-oriented approach, since it allows the designer to build modular, reusable and highly flexible software components.

A first draft of the simulator was written in the increasingly popular Java language, thus allowing users of the Internet to run experiments directly on their web browser. A demo version can be found on my home page, at the Internet URL `http://lamicounter.epfl.ch/arnaud/`.

Because of some limitations related to the Java virtual machine used by the end of 1996 — mainly its low speed, poor browser uniformity and the overprotective sandbox model — we rewrote the simulator using C++.

This second version is currently running as a stand-alone Windows NT application.

High flexibility is obtained through modularity *(each sensor, actuator, algorithm, etc. is implemented as a separate class — a self-contained module in the object-oriented terminology — which can be derived, overridden and extended)* and special techniques such as dynamic object allocation and message passing.

[slide Simulator structure]

The simulator is based on several components :

- robots
- a position manager
- three databases which store :
    - measures from the sensors
    - robot, sensor and actuator parameters
    - and the positions and velocities of the objects moving around in the simulated arena
- further, there is a global system clock
- and a graphical front-end..

Each component will be able to run as a separate process. It is therefore possible to distribute the simulator among several computers in order to increase its performance.

The graphic user interface and the visualisation tools are not part of the simulator; they will be provided by external utilities, thus allowing an easy upgrade. 3D and virtual cameras feeding images to sensors of the simulator will be added afterwards.

The clock used by the simulator is incremented in discrete steps, but if real-time simulation is not required, the increments can be reduced as much as needed, thus achieving an almost continuous flow of time.

Real-time simulation is difficult to obtain with a large number $\underline{n}$ of robots, since the required computation time grows proportionally to $\underline{n}^2$.

The robots are modelled as physical objects governed by physical laws : they are not punctual, they move in a continuous space, they oppose inertia to any change of speed and direction, they support only limited acceleration and they bump into obstacles.

Noise is introduced artificially by adding a random signal to the output of the sensors or by modifying randomly their properties (for example their gain or sensitivity).

Imperfect actuators are simulated by the addition of a small drift to the robot’s direction and a slight variation of the speed.

The simulator’s clock is used as the time reference. The sensors and algorithms often rely on a slower clock, which can be provided on a per robot basis. With this implementation, simultaneous events may not be perceived simultaneously, as it is the case with real robots using discrete sampling rates.

The real position of the robots is maintained by the simulator, but it is not directly available to the robots : these have to rely on their limited odometry or on some external, possibly noisy, positioning system (for example a global positioning system).

Communication between the robots is made possible through the use of specialised sensors : the communicators. These are realistic simulations of local or global communication systems used on miniature mobile robots. The communicators have a limited range, bandwidth and reliability. Messages can be disturbed by interference or by obstacles.

Let me show the software model of a simulated robot.

[slide Robot]

Information originating from the virtual world is sampled by the sensors. They relay their — possibly filtered or pre-processed — measures to the actors.

The actors implement the algorithms needed to react to the robot's perceptions, in order to avoid obstacles or to follow other robots. The actors are responsible for the robot's low level instinctive intelligence. They output a general trend, characterised by a desired vector — usually the robot's velocity — and an associated priority; this can be viewed as a very low level behaviour.

The trends are then combined by a special module : the trend manager, which is responsible for weighting them and driving the actuators with reasonable values.

Internal state variables can be viewed as virtual actuator-sensor pairs. Their value is also determined by the trend manager.

Our suggested trend architecture is somewhat similar to the subsumption architecture developed by Rodney Brooks : a high priority actor inhibits the other ones. Whereas subsumption totally inhibits inappropriate behaviours, our trend architecture combines all the behaviours — even the less appropriate ones — and generates actions which probability is proportional to the trends' priorities. This can be viewed as some kind of *action fusion*.

We believe that the simulator we implemented is realistic enough to allow the development of experimental set-ups by simulation and to test them on real world robots without further modification.

8. Further research [slide]

We hope that our research will come up with solutions to the problems of auto-organisation and distributed decision making. We will try to answer questions such as :

- How do robots choose between different behaviours without any central supervision ?
- How can this choice be made with minimal inter-individual communication ?
- Is it possible to suppress completely inter-individual communication ?
- What would be the impact of this suppression on the system’s robustness and fault tolerance ?

9. Conclusion

Planning and supervising the navigation of fleets of robots will become increasingly important as robots are starting to work in teams. We believe that the behaviours exhibited by gregarious animals, obtained through a long evolutionary process, will be useful when applied to autonomous mobile robots. We will study these behaviours and try to develop methods which will let the robots choose the best ones based on local observations.

The simulator we presented in this paper can be used as a platform for experiments involving numerous autonomous mobile robots. Its high modularity make it easily extensible.

We will set up experiments involving numerous mobile robots based on the Khepera and hope that we will be able to move autonomous groups of 10–20 real robots, displaying one or more patterns of locomotion.