Self-organised Computational Structures for Real Time Analysis in Highly Distributed Environmental Monitoring

Rehan O'Grady[1], Anders Lyhne Christensen[2], Roderich Groß[3], and Marco Dorigo[1]

Abstract— **We propose the new concept of self-organising computational structures — simple distributed nodes that can pool their computational resources on the fly. We argue that distributed monitoring solutions with this ability will be able to achieve high density data sampling over a wide area whilst at the same time be capable of analysing the data in real-time. We identify associated research questions and problems, and explain how we plan to test our vision in the lab before deploying it in the field.**

I. Introduction

There are many benefits to using distributed systems in environmental monitoring scenarios. Compared to the traditional sensing station approach, sensor networks, for example, tend to offer much better spatial coverage, real-time data collection, rapid deployment, and redundancy [4], [5]. Nodes in so-called mobile sensor networks can actively move to areas of interest and form spatial arrangements with the express purpose of gathering relevant data points to produce more sophisticated analysis [1]. Distributed systems such as robot swarms have the potential go one step further and start actively manipulating the environment, for example helping to clean up pollution [2], [8].

There are two types of monitoring one can perform with distributed systems. One type of monitoring involves gathering as much data as possible — ideally with a high density distribution of nodes over a wide area. Existing systems that perform this type of monitoring tend to rely on decentralised networks comprised of very simple low cost units — thus allowing high numbers of nodes to be spread in the environment. Collected data is channelled through the network to collection points where the data can be subsequently analysed off-line [3]. The other type of monitoring involves real-time data analysis and decision making. Real-time analysis is particularly important in mobile sensor networks and robot swarms when decisions need to be made about node movement and/or how to manipulate the environment. Existing systems capable of real-time analysis and decision making tend to rely on complex nodes capable of rich modelling and analysis. In [6], for example, every node maintains a complete model of the environment.

To date, no existing distributed system is capable of performing both types of monitoring. Indeed, high density data gathering and real-time data analysis are usually treated as fundamentally incompatible requirements. The low cost nodes required to make sampling feasible at a high density over a wide area do not possess the computational resources required to analyse the data they are gathering. On the flip-side, the complex nodes required for real-time data analysis are too expensive to be distributed in significant numbers for dense sampling over a wide area.

In this paper, we propose a solution to resolve this perceived incompatibility. Our core idea is to have simple nodes pool their computational resources when necessary. Our solution is based on the novel concept of *self-organised computational structures*. We argue for the development of a new class of self-organisation based algorithms that will allow on-the-fly formation and dissolution of two types of computational structure — hierarchies and clusters. Computational hierarchies will allow data analysis calculations to be split up and performed in stages. Sensor data heavy computations can be performed close to the source of the data, reducing overall data volumes. Raw data can be discarded early, with only pre-analysed information passed further up the hierarchy for higher level analysis. Individual computations (at any level of the hierarchy) may still exceed the capacities of individual robots. Computational clusters will address this problem. Borrowing ideas from parallel computing, robots will use high

[1]IRIDIA-CoDE, Université Libre de Bruxelles, Avenue Franklin Roosevelt 50, 1050 Brussels, Belgium {rogrady,mdorigo}@ulb.ac.be
[2]Instituto de Telecomunicações & Instituto Universitrio de Lisboa (ISCTE-IUL), Lisbon, Portugal anders.christensen@iscte.pt
[3]Natural Robotics Lab, Department of Automatic Control and Systems Engineering, The University of Sheffield, UK r.gross@sheffield.ac.uk

bandwidth short range communication to form more powerful local computing clusters. We believe that our approach will enable new, massively scalable environmental monitoring systems that possess the benefits of both distributed self-organised systems with simple agents and of more centralised systems with complex agents.

II. Our Vision

We envision a computationally sophisticated robotic swarm that can not only monitor, but also analyse and react to dynamic environments. In dynamic environments, elements such as spills may appear or disappear, temperature changes may propagate through the environment depending on the wind direction, a spill may move or change shape, and so on. The swarm we envision will use self-organised computational structures to model the change of the environment over time, and thus predict the future state of the environment. This type of analysis can then feed back into the self-organised processes determining the spatial location of the robots and the logical organisation of the computational structures, ensuring that the distribution of robots remains well suited to the changing environment.

Imagine a distributed system of thousands or millions of robotic agents deployed to locate, analyse, and contain an environmental disaster such as an oil spill or a leakage of nuclear waste. Such a system of robots could first spot, for example, ten different locations where the spill is leaking. Local groups of robots could act in parallel to detect the rate at which spill is leaking in different locations. Each group could collectively observe how change in spill distribution is correlated with relevant environmental parameters (e.g., pressure, temperature, surface inclination, and so on). Together, the different robot groups could reason about the priorities of the different spill locations. The overall system could reason about how to get control over the spill at the high priority locations. For example, the system could compare the estimated impact of a range of action plans on the future distribution of the spill. The robots could manipulate the spill and/or environment according to action plans at several locations. Finally, the system could reason not only about the local impact of an action plan but also about interdependencies across different locations.

This long-term vision implies a system with an ability to perform computation and coordination at different degrees of centralisation when searching for spills, when analysing the spills and when collectively containing the spills, respectively. In the search phase, the system must engage in highly decentralised activity. The robots display independent and/or loosely coupled forms of computation — similar to how social insects like army ants search environments. To assess the extent of a spill and to reason about how its distribution may change over time, robots will have to pool their resources. In local computational structures, they can split up the problem and execute collective computation algorithms (similar to the operational principle of high performance computing clusters). Moreover, the robots must be able to form global computational structures to monitor and synchronise their activities, and to take prioritisation decisions (similar to information processing in bureaucracies). Finally, when attempting to contain a spill at a particular location, teams of robots must engage in highly coordinated activity to perform collective manipulation (similar to the function of the central nervous system in vertebrates).

III. Our Approach

Considering the required properties outlined in the previous section, it is evident that no system exists that fulfils all of the requirements. Current large-scale systems of physically embodied agents may leverage their spatial distribution through implicit reasoning, but they are unable to perform collective explicit computation. Current parallel computing algorithms are not designed to leverage spatial distributed agents.

Our approach relies on robots self-organising into spatial arrangements and logical relationships. These spatial arrangements and logical relationships are interdependent and mutually beneficial. The robots must distribute themselves in space so as to satisfy the needs both of data collection and of hierarchical computation. Figure 1 shows an example spill monitoring scenario, where the robots have organised themselves into a computational hierarchy. In the figure, the spatial organisation of robots allows them to monitor spill status and ocean currents, while also satisfying the spatial constraints imposed by the logical organisation of the computational hierarchy.

In general, our approach involves the formation of a computational hierarchy to process information from multiple locations in the environment. The information collected could be a count of the number of particular objects in the environment, the differences in temperature, or spatial locations defining the edge of a spill. Data points collected by robots in environmentally meaningful locations are propagated upwards in the

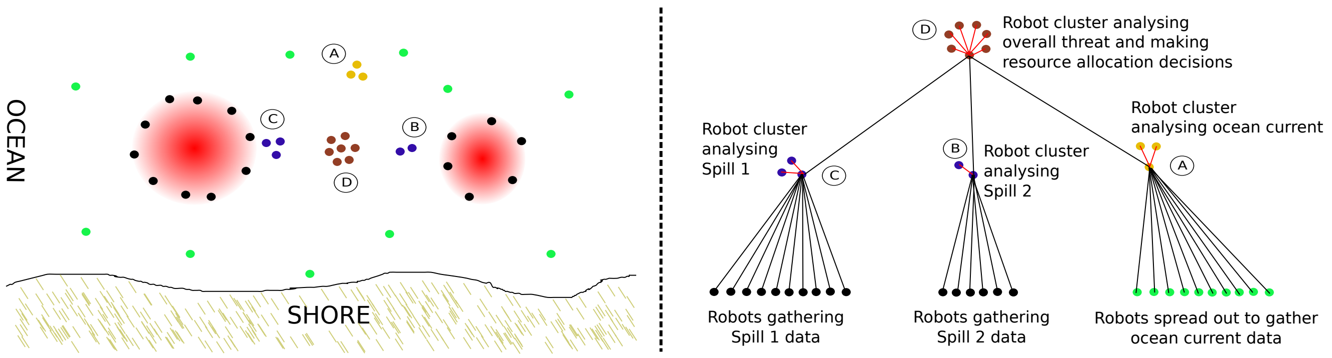

Fig. 1. Spill Monitoring Scenario. Red areas represent concentration of pollutant. Each small circle represents an individual robot. All robots are identical. Robot colours indicate different roles — roles are allocated by a self-organised process (no robots have a priori determined roles). Self-organised computational clusters are indicated by ringed letters A-D. Left: Spatial arrangement of the robots. Robots self-organise in space based both on environmental features (e.g., near a spill to monitor the spill's status) and on the needs of the self-organised computational structure (e.g., physically adjacent to form a cluster). Right: Logical arrangement of the same robots. The robots have formed a computational hierarchy (black lines). Robots responsible for analysis have recruited other robots (red lines) to form computation clusters — more complex analysis requires a larger cluster. **NB.:** This figure should be viewed in colour.

hierarchy via robots responsible for collating and processing the data. Parent robots will receive information from child robots in the hierarchy, collate and analyse that information and pass the processed results upwards in the hierarchy for further collation and analysis. The amount of data passed up the hierarchy thus does not accumulate but is kept relatively constant, which ensures the system does not suffer from data-overload however large it gets. Other approaches that require all data to be collected and analysed at a single point have tried to tackle scalability issues through more intelligent data routing [3]. But however well the data is routed, any approach based on centralised computation will inevitably hit scalability barriers once a certain number of robots are reached.

The analysis performed at any given point in the hierarchy may exceed the computational capacity of a single robot (e.g., combined camera feed analysis). In such cases, the responsible robot can recruit other robots to form an ad-hoc computational cluster. This type of clustering on demand ensures that individual robots in the system can remain simple and relatively inexpensive, thus removing cost based barriers to large scale deployment.

We hope that the combination of ad-hoc clustering, and hierarchical computation will enable environmental monitoring solutions based around robotic swarms that are orders of magnitude larger than anything feasible with today's wireless sensor network technology, and with active manipulation capabilities to boot.

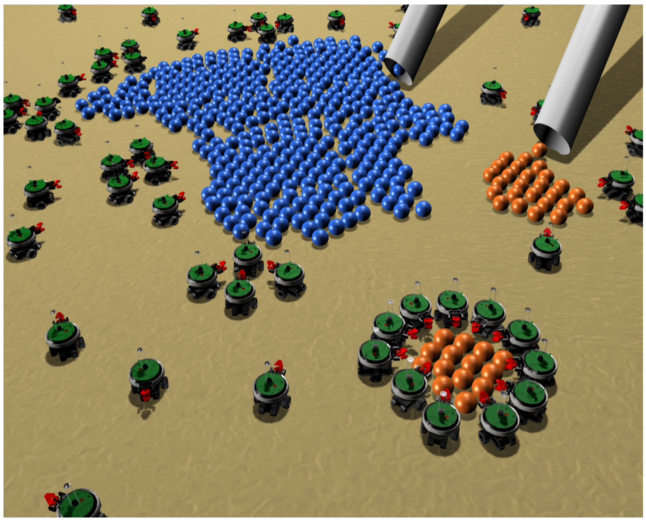

Fig. 2. Proof-of-concept environmental monitoring and containment scenario with two spill locations. In the short term, we aim to implement this scenario in the lab using robots that self-organise into ad-hoc computational structures.

IV. Validation in the Lab

A series of research questions still need to be answered before we can hope to implement our approach. Firstly, we need to develop self-organised phase transition mechanisms to enable the formation and reconfiguration of computational hierarchies. Secondly, we need to investigate how we can apply parallel computing results to the formation of ad-hoc robotic computing clusters. Finally, we need to study in application specific contexts how the relevant computations can be split up and made amenable to hierarchical analysis.

We plan to demonstrate collective computation in a swarm of robots through a proof-of-concept scenario illustrated in Fig. 2. The scenario mimics an environmental monitoring and containment scenario. Balls represent hazardous material and their colours represent the properties (toxicity) of the particular hazardous material. Additional material can be introduced at the spill locations during an experiment through pipes (see Fig. 2). To solve the task, the robot swarm must first engage in a phase of highly decentralised, implicit computation and coordination to locate spill locations. During this initial phase, the robots must maintain local communication connectivity through the system. The second phase of the task will require groups of robots near each spill to pool both their sensory resources and their computational resource to estimate the parameters of the spill: size, shape, rate of growth, environmental impact, and so on. The estimation will make use of the spatial location of the different agents in a local group and will take into account the properties of the objects (e.g., ball colour). In the third and final phase, the system will prioritise resources both globally between the different spills, and locally for each individual spill, to assign robots to locations and tasks to maximise monitoring and containment capabilities of the system. In our proof-of-concept scenario, robots will encapsulate and transport balls to a safe location to contain a spill, see Fig. 2.

The proof-of-concept scenario will essentially allow us to study the key aspects of self-organised computational structures: (i) during the spill localisation phase, the robots must display independent and/or loosely coupled forms of computation; (ii) during the estimation phase, the robots must form local hierarchies at each spill to estimate its parameters, using ad-hoc clusters where necessary to perform complex calculations; and (iii) during the containment phase, data from all local hierarchies must be used to prioritise and coordinate monitoring and clean-up efforts at the global level.

V. Conclusion and Future Work

In this paper, we proposed the new concept of self-organising computational structures in a robotic swarm to enable scalable, sophisticated environmental monitoring. A system based on this principle should combine the benefits of centralised and decentralised approaches to environmental monitoring. Our proposed concept has the desirable characteristic of being potentially revolutionary for future distributed monitoring systems, while at the same time being testable on existing robotic platforms. Our short term research will be based on a current robotic platform capable of high speed local communication between robots.

We are also keen to get self-organised computational structures into the field as soon as possible. Having checked our hypotheses in the lab, a first implementation might involve modifying existing sensor network platforms (e.g., [7]) to give them simple motion capabilities and computational sharing facilities.

In the longer term, simple robots comprising a swarm will need to coordinate their actions in order to actively manipulate the environment. In future work, we aim to use self-organised computational structures to control the physical coordination of tightly coupled robotic agents, perhaps through the formation of on-the-fly central nervous systems for composite agents.

Acknowledgments

This work was partially supported by the European Commission via ERC Advance Grant "E-SWARM: Engineering Swarm Intelligence Systems" (grant 246939), by FCT grants PTDC/EEACRO/104658/2008 and PEst-OE/EEI/LA0008/2011, and by the Marie Curie European Reintegration Grant within the 7th European Community Framework Programme (grant no. PERG07-GA-2010-267354). Marco Dorigo and Rehan O'Grady acknowledge support from the Belgian F.R.S.-FNRS, of which they are a research director and a postdoctoral researcher, respectively.

References

[1] I. Akyildiz and I. Kasimoglu. Wireless sensor and actor networks: research challenges. *Ad Hoc Networks*, 2(4):351–367, 2004.

[2] G. Beni. From swarm intelligence to swarm robotics. In E. Sahin and W. Spears, editors, *Swarm Robotics*, volume 3342 of *Lecture Notes in Computer Science*, pages 1–9. Springer, 2005.

[3] G. Di Caro, F. Ducatelle, and L. M. Gambardella. Anthocnet: an adaptive nature-inspired algorithm for routing in mobile ad hoc networks. *European Transactions on Telecommunications*, 16(5):443–455, 2005.

[4] J.K. Hart and K. Martinez. Environmental sensor networks: A revolution in the earth system science? *Earth-Science Reviews*, 78(34):177–191, 2006.

[5] F. Ingelrest, G. Barrenetxea, G. Schaefer, M. Vetterli, O. Couach, and M. Parlange. SensorScope: Application-Specific Sensor Network for Environmental Monitoring. *ACM Transactions on Sensor Networks*, 6(2):1–32, 2010.

[6] K.M. Lynch, I.B. Schwartz, P. Yang, and R.A. Freeman. Decentralized environmental modeling by mobile sensor networks. *IEEE Transactions on Robotics*, 24(3):710–724, 2008.

[7] P. Padhy, K. Martinez, A. Riddoch, R. H. L Ong, and J.K. Hart. Glacial environment monitoring using sensor networks. In *Real-World Wireless Sensor Networks*, 2005.

[8] M. Schwager, D. Rus, and J. Slotine. Decentralized, adaptive coverage control for networked robots. *International Journal of Robotics Research*, 28(3):357–375, 2009.