Decentralized approaches to SSA data gathering – implications for data policies

J. Dick, P. Herridge, Space Insight Ltd.
PO Box 2993, Eastbourne, BN21 9BD, United Kingdom
R. Tremayne-Smith, Out of Space
PO Box 2993, Eastbourne, BN21 9BD, United Kingdom
e-mail, all authors: info@spaceinsight.co.uk

Abstract

Any practical SSA system requires sensor assets, data centres, and data processing capabilities. For any future European SSA system funding and geographical constraints mean that many nations will be involved. So the organisation of the system elements to ensure operation in a coherent and efficient manner has to be balanced with the needs of a diverse user community: the presence of differing national military and civil interests within that community complicates issues by the introduction of additional layers of data security, facility access control, and cross-border data policies.

This paper discusses centralised and decentralised architecture options for an SSA system (with its network of sensors and other facilities). Decentralised multi-agent biological systems are introduced. The simple sets of rules by which multi-agent systems operate have direct counterparts in the efficient working of any distributed system. Using analogies with multi-agent system concepts, the paper explores the possibilities for simplifying sensor and data centre management within a multi-national, multi-interest SSA system.

In a decentralised architecture, self-scheduling sensors, multiple data centres, and data mining processes can be the foundation of knowledge production. The simplicity of such a system will keep low the cost of implementation and operation, while ensuring an openness to future reconfigurations and upgrades.

From a policy perspective, the decentralised architecture has important consequences because it is easier to disperse the SSA system's functional elements within a multi-nation user community. The dispersion of activity, devolution of responsibility, and the simplicity of the governing ruleset combine to provide a transparency which is important for building trust between the SSA partner nations and for ensuring an early success.

1. Introduction

The man-made space population (satellites, their associated rocket bodies, and fragments thereof) occupies a huge physical volume around the Earth. The current population is distributed throughout a volume of $\sim 3\times10^{14}$ km^3. Although some members of the population are concentrated into comparatively limited regions, such as the geostationary orbit (GEO), the diversity in the types of orbits imposes a wide spread in the range and relative motions of objects as viewed from the Earth.

Knowledge of the man-made space population is known as space situation awareness (SSA) and is based on the surveillance of space. For low Earth orbit (LEO) surveillance radar sensors are used; optical sensors are used for the higher Earth orbits (for which radars are less efficient). A space surveillance network (SSN) typically uses both types of sensor and while Earth-based sensors are common, space-based sensors are also used.

Although optical sensors are less costly to procure than their radar counterparts both are expensive to operate within a comprehensive SSN because of the number and geographical dispersion necessary. To maximise the benefit/cost ratio during operation, the data collection, assessment, validation and generation of information needs to be as autonomous and automatic (robotic) as possible. In the case of international collaborations the geographical dispersion would require different member states to host sensors, adding political and national security concerns to the set of issues which need to be addressed.

ESA has, for the last few years, been exploring the steps necessary for the development of a European SSA system, and an SSA programme was discussed at the 2008 ESA Ministerial meeting. The European SSA ambitions provide a good example of the political complexity of such a system: most ESA member states are physically small and so, for Europe, even a small and minimally-sufficient SSA network will need funding from, and co-operation between, many member nations.

The norm for many European projects is that multilateral exchange is common and to build trust between the participants open sharing of both cost and responsibility is important. Sharing imposes important content and process visibility requirements and these have to co-exist with, but have the capacity to compete with, the national security requirements of the participating states.

2. Architecture choices

In designing an SSA system, there is a choice of system architecture which ranges from the traditional highly centralised to the decentralised. A decentralised design seems to fit better a multi-partner system because its many components can be shared more easily between the partners. Alternatively, the centralised architecture appears to offer better control – a view often held by those who expect to host the control nodes of the system.

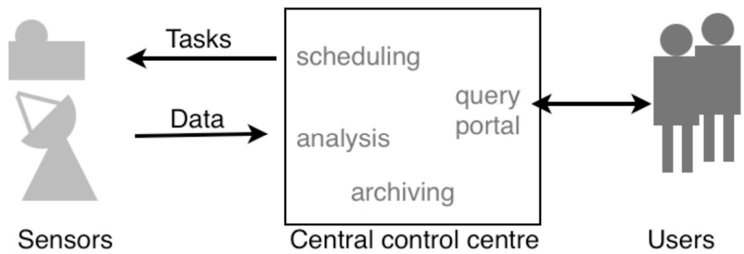

Figure 1. SSA system with centralised control

Whatever the political and financial influences are on the architecture, there are technical issues too. Bluntly, how do you make the data collection and analysis processes work? Are there technical methods and strategies to be found which are confluent with the political, national security, trust, and funding issues? And can consideration of these technical issues feed forward useful contributions which can help steer the higher levels of policy decisions?

3. Multi-agent architectures

Co-operating towards a common goal is an activity frequently found in many biological systems. Given that many human processes and artefacts have analogues in the behaviour of simpler species, are there lessons to be learned from multi-agent systems in biology?

Resnick [Ref. 1] gives some thought-provoking examples of decentralised co-operation in biological systems which are modelled in StarLOGO, a computer language for modelling massively parallel multi-agent communities. One example illustrates how termite avatars gather woodchips to build a termite house. In the computer simulation, the red avatars inhabit a world over which yellow woodchips have been scattered (Fig. 2a).

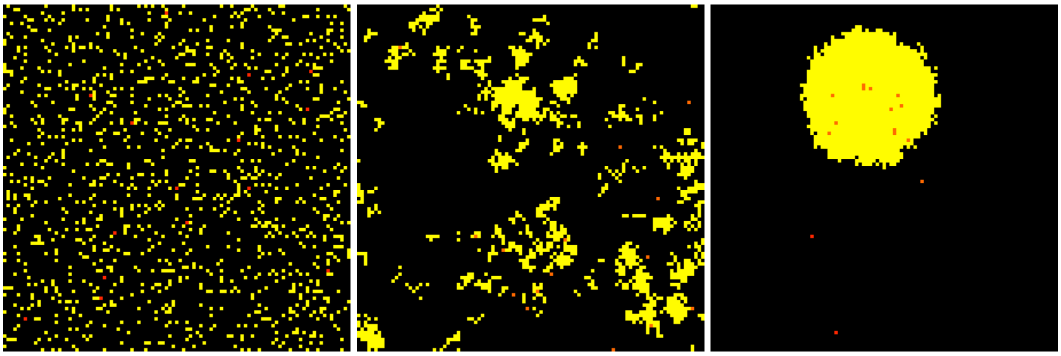

Figure 2. Resnick's termite simulation a), b), and c)

The termite avatars each obey a set of simple rules: *i*) wander around, *ii*) if not carrying a woodchip and you find one, pick it up, *iii*) if carrying a woodchip and you find another then put your chip down beside it and walk off. As time progresses, the avatars collate the woodchips into small mounds (Fig. 2b) and finally into one "house" mound (Fig. 2c).

The critical issues to note include: no avatar knows what the others are doing; there is no controller avatar; the avatars do not know they are building a house; there is no pre-defined woodchip distribution map (and the avatars are not told explicitly to search for woodchips); the avatars do not know *a priori* where the house will be situated. Additionally, there is no *a priori* information about which avatar will be handling which woodchip(s) nor need there be *post hoc* information about which avatar had handled particular woodchips. Although in Resnick's example all the termites come from one colony this unity is not a pre-requisite; the avatar termites could come from different colonies, or from different nations.

Resnick [*ibid*.] also gives an example of ant foraging in which ants hunt for food and stock their nest as a result of obeying simple behavioural rules, not as a result of any centralised control. The co-operative foraging is not organised but just happens as a by-product of obeying the rules. In this example, the ants' motion guidance rule is uncoupled from the food pick-up rule, no inter-ant (or ant-to-nest) communication is needed, and the only centralised aspect is that the ants need to know the nest location.

As with the termite example, different species of ants can be used to forage for different types of food (and they may have different or shared nests); information about food location is left by the avatars in the environment (that is, a stigmergic process is used [Ref. 2]) and can be selectively sensed by all or some of the participating ant species.

Although these simple biological communities are very different from the space world, the same principles of decentralised action and stigmergic information exchange are being applied in many diverse subject areas, including (for example) the problems of UAV control [Parunak, Ref. 3; Sauter, Ref. 4], and are applicable to the design of an SSA system.

4. Foraging in SSA

The foraging seen in biological communities has parallels with the surveillance activities of an SSN. For example, surveying sensors use a sequential area search strategy which is not target-oriented – they do not need an *a priori* picture of the target population or any knowledge of orbits, have no requirement for tasking from any source (centralised or decentralised), and so their data collection can be anonymous. Further, sensor-level anonymity can be ensured because there is no need for sensor-to-sensor communication and sensors need not know what their peers are observing.

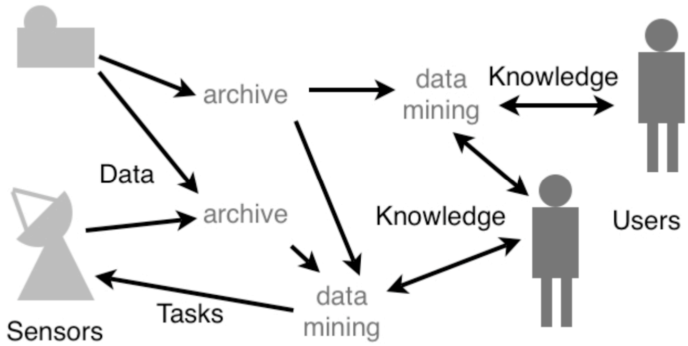

Figure 3. Decentralised SSA system

In a decentralised SSN, dispersed sensor-fed data mining processes can schedule tracking sensors without a centralised control system. Tracking sensors do not need to be linked with the surveying sensors because their tasking can come from a mining process carried out on the collectively-foraged datasets; the mining process can be specific (for high-interest known targets) or generalised (for unknown objects) and can be carried out on one or more data stores. The cueing process can use stigmergic principles (such as a demand-coded map of the sky) to ensure cue anonymity so a cued sensor need not know from whence a cue came. The transmission of tasks may be by broadcast to the whole SSN with devolution to each sensor of the target pick-up decision, or may be by sent only to certain sensors (for example, where national security concerns are present or there is a sensor capability issue).

Downstream from the sensors, data can be collated in different storage locations (where destination choice may be based on data type, or provenance) and mined to provider higher-order knowledge products. Importantly, the information extraction and knowledge creation work can happen in various locations, dispersed throughout the member states, and need not be co-located with the data stores.

5. Implementation and data policy issues in a decentralised SSN

The key in any multi-agent biological system simulation is to prescribe a small but sufficient ruleset in which the effects of obeying different rules interact minimally. A successful ruleset is likely to be one which, with reference to game theory, allows the resulting system to function either in a Nash equilibrium or a Pareto optimum. Testing a ruleset against these criteria can provide a prediction of success as judged qualitatively by the system partners and quantitatively by audits of the system's achievements.

For a multi-nation SSN, the ruleset needs to meet both the technical and policy requirements and would cover sensor usage, data distribution, and product generation. Simplicity in this ruleset is important because it minimises the implementation cost and time; simple rules ensure clarity of operation which helps to build trust between the SSN partners.

For example, sensors can self-task by mining the data centres to establish where there are observation needs; there is no requirement for a complex scheduling process, triggered by a sensor reporting its availability. Subsequently, an audit of availability can be carried out by the central co-ordinating entity to ensure that there is, overall, an adequate level of sensor coverage and availability.

A decentralised system architecture maps naturally onto a distributed processing architecture. The central co-ordinating (but not controlling) entity could host a master database but might provide only agreed application-level results (that is, not raw data) to a wide range of users. To meet specific user needs, various enterprise level solutions could be established on a common bearer system, with local storage of limited data as needed.

Any initial co-operative European SSA system would need to focus on the regular data to be shared and would need different levels of availability and archiving. Some data could be product specific; target-specific data and services would be handled separately from regular products and could be more experimental in any initial service offering.

The data policy related to any central system would also need to provide a research capability with timeliness and usage limits, and minimal data provision cost.

National processing centres, archiving, and data distribution would be quite separate from the user or service end of the SSA system; partner nations would be free to have their own processing and results delivery capability and may deal with security issues using their existing national laws.

6. Summary

Multi-agent biological systems provide examples of complex behaviours (such as collective foraging) which are based on simple rules and participant autonomy. The simplicity of such systems provides transparency of action and great freedom of choice. Many of the concepts contained in these systems can be applied to the operation of an SSN. The simplicity of organisation and transparency these concepts suggest are important for building trust and for reducing the administrative overhead common in more centralised architectures.

References

1. Resnick, M.: *Turtles, termites, and traffic jams*, ISBN 0-262-18162-2, MIT Press, USA, 1997
2. Grassé, P. P.: La Reconstruction du nid et les Coordinations Inter-Individuelles chez Bellicositermes Natalensis et Cubitermes sp. La théorie de la Stigmergie: Essai d'interprétation du Comportement des Termites Constructeurs, *Insectes Sociaux*, Vol. 6, pp. 41–84, 1959
3. Parunak, H. vD., Purcell, M., O'Connell, R.: Digital pheromones for autonomous co-ordination of swarming UAVs, *Proc. First AIAA Unmanned Aerospace Vehicles, Systems, Technologies, and Operations Conference*, 2002
4. Sauter, J. A., Matthews, R., Parunak, H. vD., Brueckner, S.A.: Performance of digital pheromones for swarming vehicle control, *Proc. AAMAS*, 2005