Autonomic resource management through self-organising agent communities

Mariusz Jacyno and Seth Bullock and Terry Payne and Nick Geard
School of Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK
{mj04r,sgb,trp,nlg }@ecs.soton.ac.uk

Michael Luck
Department of Computer Science, King's College London, Strand, London WC2R 2LS, UK
michael.luck@kcl.ac.uk

Abstract

In this paper, we analyse how autonomic resource management can be achieved within a system that lacks centralized information about current system demand and the state of system elements. Rather, regulation of service provision is achieved through local co-adaptation between two groups of system elements, one tasked to autonomously decide which services to offer and the other to consume them in a manner that minimises resource contention. We explore how varying the amount of information stored by agents influences system performance, and demonstrate that when the information capacity of individual agents is limited they self-organise into communities that facilitate the local exchange of relevant information. Such systems are stable enough to allocate resources efficiently and to minimise unnecessary reconfiguration, but also adaptive enough to reconfigure when resource demand changes.

1 Introduction

In this paper we describe a model of a decentralised system in which there may be no global information repository and the demand for particular resources changes over time. Agents must therefore organise the information that is available to them locally in order to adaptively respond to changes in demand in an efficient manner. Where previous studies have explored the role of system heterogeneity brought about by limiting knowledge [2] or using decision procedures that diversify agent behaviour [1], here we focus specifically on how varying the amount of information available to agents affects the flow of information. By conducting a thorough evaluation of our model, the results demonstrated that when agents possessed only a limited awareness of their peers within their neighborhood, they exhibited self-organising behaviour resulting in the emergence of community structures that support locally shared information. By limiting the quantity of information shared by constraining an agent's memory size, we found that a stable local community behavior emerged that was robust and efficient, and also adaptive when the whole agent population was exposed to global fluctuations in service demand.

2 Autonomic Computing management

Consider a population of servers that have limited capacity in offering requested resources (services) and thus can serve only a limited number of concurrent requests. Furthermore, each server can be configured at provisioning only one type of a service, but is capable of reconfiguring its provision at run-time. This involves a significant cost in the form of down time during which security compromising data is being removed and a new software module representing the new service type is being loaded. The management over the current server configuration is devolved to an autonomic software component that decides which service type should be provided and which service requests are served.

Service requests originate from a population of service consumers, which are a software manifestation of the infrastructure users and are responsible for acting on their behalf, such that the requested services are delivered reliably and 'on-demand'. The consumers may have different interests in service types and maintain only local information about the existing providers and their current state (whether they are busy and provide required service). Furthermore, within a system consisting of both types of autonomously acting elements, there exists no central queue maintaining existing demand or information repository about current state of providers. Given this, the resource management design goals are three fold:

1. which providers should be configured to offer what service types, in order to satisfy current demand;

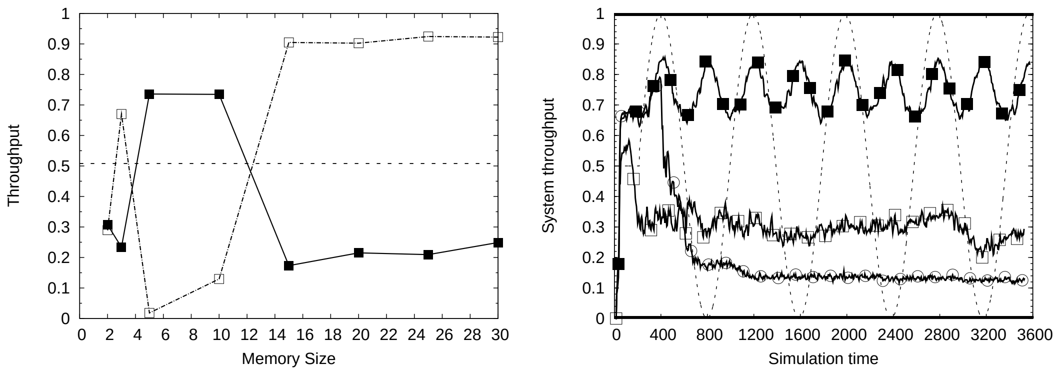

Figure 1. Left figure illustrates mean system throughput for centralised (dotted line) and decentralised (solid rectangles) system configurations. A proportion of consumer population updating their local knowledge for decentralised system is represented by the line with empty rectangles. Right figure depicts performance of three decentralised system configurations differing in memory size: $m_s = 2$ (empty squares), $m_s = 5$ (solid squares) and $m_s = 20$ (empty circles).

2. which of the providers known to a consumer should be utilised such that the system minimises resource competition; and
3. how should the system be organised such that it is robust to dynamism in changing load and demand for particular service types.

To approach this, we model resource management as a co-adaptive process between consumers and providers. We represent these elements in the form of agents comprising a decentralised multi-agent system. In the absence of any controllers or centralised information repositories, agents leverage their decision-making efficiency based on past transactions and observations of the local community. In this manner providers acquire information about the most demanded service type, and reconfigure accordingly, whereas consumers learn which providers are the most reliable and efficient at offering required service types. Such local observations maintained by consumer agents are then communicated to other consumers and constitute locally disseminated information about the current system state.

To achieve regulatory system response relying only on local information exchange between co-adapting elements, we focus our attention on the self-organising principles of natural systems and employ the idea of *gradients* these systems are capable to form and exploit for their organisation and regulatory function maintenance.

3 Results

Figure 1 (left) illustrates performance (in terms of successful service allocation, or *Throughput*) of both the decentralised (solid rectangles) and the centralised (dotted line) models measured as the mean throughput achieved by consumer agents during the simulation. The points are plotted against memory size (m_s) (for the decentralised model). The graphs are normalised with respect to the optimal system performance experienced by the system in equilibrium during steady state (i.e., when service-demand is satisfied by supply such than no reconfiguration of providers is necessary).

Analysis of the decentralised model's efficiency in successfully matching service requests with available service providers for different memory size configurations (Figure 1 (left)) reveals that the system efficiently manages resource allocation only for a certain subset of memory sizes (where $4 < m_s < 12$). In addition, the level of information flow between consumers is low under such conditions (as illustrated by the empty squares), suggesting that the allocation of knowledge held by consumers about the local providers is relatively stable. However, for memory size values outside this range, service throughput falls to a poor performance level, with a corresponding increase in the level of information flow between consumers.

4 Discussion

A critical factor in achieving the efficient performance demonstrated by some parameterisations of the decentralised model is the organisation of knowledge across the consumer and provider populations. A system where consumers successfully organise their local knowledge supports ongoing interaction between these consumers and service providers configured to offer the type of service that they require. Moreover, such an organisation also has the ability to smoothly reconfigure the provision of services in response to changing demand. By contrast, operating on the wrong information will result in degraded performance due to repeated rejected queries and time spent needlessly reconfiguring service provision.

References

[1] T. Hogg and B. A. Huberman. Controlling chaos in distributed systems. *IEEE Transactions on Systems, Man and Cybernetics*, 21:1325–1332, 1991.

[2] S. Sen, S. Roychowdhury, and N. Arora. Effects of local information on group behavior. In *Proceedings of the Second International Conference on Multi-Agent Systems*, pages 315–321. AAAI Press, Menlo Park, CA, 1996.