On Decentralizing Intelligence in Cyber-Physical Systems

Thesis submitted in partial fulfilment of the requirements for the award of the degree of

Doctor of Philosophy

in

Computer Science and Engineering

by

Tushar Semwal

Under the supervision of

Prof. Shivashankar B. Nair

Department of Computer Science and Engineering

Indian Institute of Technology Guwahati

Guwahati - 781039 Assam India

November, 2018

Abstract

Advancements in the low-cost computing and communication technologies have led to the mass proliferation of devices connected over a network. The networked devices have engendered a new era wherein they sense, compute and share information thereby forming loosely coupled Cyber-Physical Systems (CPS). Managing data and making intelligent decisions form the major domain of research in a CPS. Cloud-based centralized computation has always been the mainstream architecture due to its ease of implementation and enhanced control. However, data explosion, scalability and privacy issues, are certainly pointing toward the limits of such centralized systems. Decentralizing control and distributing the computing among the devices could be a better alternative for sharing intelligence. Investigating new decentralization mechanisms, thus, forms the major crux of this thesis. Realizing such *decentralized* Cyber-Physical Systems (*d*CPS) is fraught with challenges such as choosing the appropriate communication method, incorporating the right learning and knowledge sharing schemes, ensuring robustness and adaptivity, and the need for a proper middleware to cater to its functioning.

This thesis takes a bottom-up approach and presents its first contribution on extending the functionalities of the *Tartarus*, a multi-agent platform, in order to realize complex *d*CPSs. This section begins with the motivation behind *Tartarus* and discusses features which makes it a disparate environment for developing and deploying *d*CPS. A real-world CPS application comprising robots, a Raspberry Pi with a camera and a human administrator in-the-loop, described herein validates the feasibility of *Tartarus* for developing mechanisms to embed decentralized intelligence in CPSs.

The second contribution discusses an approach to transform a centralized system into a fairly distributed system with partial or no centralized control. The technique uses mobile agents to provide local sharing and in-network processing of data. A Location-Aware and Tracking Service (LATS) as a real-world *d*CPS

based application is used to portray the viability of the proposed approach. The technique allows users to embed queries, such as *where* and *when* about a specific person, into a mobile agent and get the results back to the user. A comparison of the proposed approach with its centralized version proved the former to be more efficient in terms of bandwidth utilization and energy efficiency thus, proving its applicability in practical scenarios.

Cooperative tasks involving autonomous entities such as robots are prone to the classical problem of Mutual Exclusion of shared Resources (MER). The next contribution presents a novel mechanism for ordering the task execution to mitigate the issue of MER in a *d*CPS of multiple robots. While the mobile agents achieve the computation, communication, and control, the physical execution of the tasks is performed by the robots in an asynchronous and pipelined manner without the use of a global clock. Unlike atomic commands which are bound to finish in a fixed period, the pipeline formed using robots as processing units is adaptive to the varying execution times experienced in real world robotic tasks. This makes the system flexible and versatile to the dynamics of the environment of the *d*CPS. The proposed mechanism provides unique features such as addition and deletion of both tasks and robots, on-the-fly. Mobile agents carrying the code or solution deliver the same to a robot to make it execute a task. Experiments conducted in an emulation environment validate the presented characteristics of the proposed mechanism. In addition, an application comprising ordering of jobs in a Warehouse Management System using real robots substantiates the feasibility of the approach in real-world decentralized and distributed systems.

The previous contribution assumed that the solutions carried by the mobile agents are optimal. However, in real systems, there could be several solutions to the same problem. The problem of selecting the best solutions, in a decentralized manner, for a given set of problems distributed across a network of nodes, motivates the next contribution. The proposed mechanism takes a non-conventional route and is inspired by the computational models derived from three immunological theories viz. clonal selection, danger theory, and the immune network theory. Whenever a problem occurs at a node, it releases danger signals to attract the concerned mobile agents carrying mappings for potential solutions. The agents that meet at the

distressed node stimulate and suppress each other thereby forming an immune network. The mobile agents carrying mappings of solutions with superior performance are rewarded and thus clone and grow in population while the remaining ones are penalized and thus diminish. The challenge is to find a generic mapping which can cater to the maximum number of similar problems. Extensive experiments conducted in the emulation environment created using the *Tartarus* platform validates the efficacy of the proposed approach. Besides, the results obtained by embodying the mechanisms in robots that discover the best path-following algorithm, substantiate their working in real-world scenarios.

With a selection mechanism in place, the next obvious step is to evolve the solutions to meet the demands of a varying stream of problems. The previous contribution assumed an infinite supply of solutions which is not entirely realistic when it comes to practical scenarios. In this contribution, the solutions, in the form of robot controllers, are thus evolved in an online and continuous fashion. Based on a single parameter, viz. the cross-reactivity threshold, the proposed approach divides the main task into subtasks and evolves separate sub-controllers for each of them. The better evolved solutions are rewarded and shared with other robots using mobile agents. A comparison of mobile agent-based communication with traditional broadcasting method portrays the former to be more energy efficient than the latter.

All the emulated experiments were conducted using the *Tartarus* platform. As an addendum, a brief review of the immunological terms and concepts employed in this thesis is also presented. Finally, the thesis concludes by summing up the contributions and providing avenues for future research and possible applications.