Passivity-Based Distributed Optimisation

Summary of the Master Thesis of Charles Muller

INTRODUCTION AND MOTIVATION

Many technical systems are comprised of numerous smaller subsystems, each pursuing different goals. Examples span across energy networks, cooperative multi-robot systems, communication systems, and smart manufacturing. These kinds of systems inherently require optimisation of system variables to achieve the best possible outcomes for the subsystems, or agents. This leads to distributed optimisation problems, such as distributed resource allocation in electrical distribution grids.

In distributed optimisation, the global objective function is typically broken down into a sum of local objective functions, each tied to a specific agent. These agents only have access to their local objectives and possibly some additional constraints. Through communication, they exchange information to collaboratively find the optimal network-wide solution.

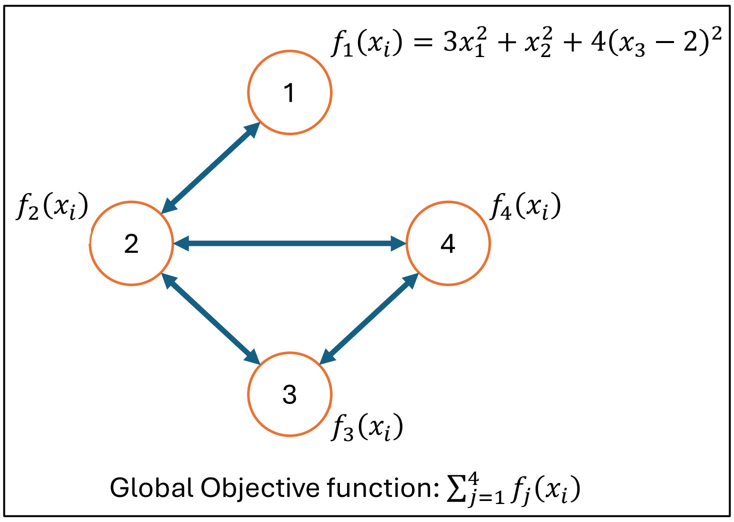

Communication graph with 4 agents, their objective functions f_i, and the global objective function

This method is especially useful for large-scale, complex problems that are too intensive for a single processor. Additionally, distributed optimisation helps maintain privacy, as local objectives don't need to be shared, unlike centralised optimisation algorithms that require the aggregation of all system information.

Existing distributed algorithms face challenges, including stability and adaptability when agents join or leave the network.

INNOVATIVE APPROACH AND KEY FINDINGS

This work leverages dissipativity theory, particularly passivity - concepts from control theory - to design a distributed optimisation algorithm that guarantees stability and convergence through locally verifiable conditions.

This algorithm tackles both unconstrained and constrained optimisation problems, ensuring that the system's consensus - where all agents agree on a solution - and optimality conditions - where the best possible solution is achieved - are met when the system reaches a steady state. This is guaranteed even in the case of newly connected or disconnected agents, and without the need for extensive system-wide information.

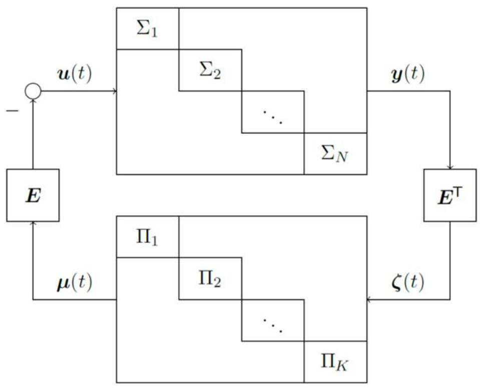

Structure of the proposed algorithm, with: the dynamics of the agents Σ_i, the dynamics between the agents Π_k, and the inncidence matrix of the communication graph E

The algorithm is verified through simulations across various scenarios, demonstrating its performance and applicability.

OVERVIEW OF THE WORK

The initial chapters introduce the fundamental concepts of dissipativity theory, convex optimisation, and graph theory. Following this, there's a review of relevant literature, specifically focusing on works in distributed optimisation and passivity-based control.

The heart of the thesis delves into the detailed design of the algorithm for distributed unconstrained optimisation, with an emphasis on ensuring stability through passivity properties. This design is then extended to handle constrained optimisation problems. The thesis further explores the generalisation of the algorithm for various network structures, assessing the implications for stability. Multiple scenarios are simulated to verify the algorithm, and its performance is compared with existing methods.

The final chapters summarise the strengths and weaknesses of the proposed algorithm and discuss potential future research directions.