Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 108C (2017) 1443–1452

Procedia Computer Science

International Conference on Computational Science, ICCS 2017, 12-14 June 2017, Zurich, Switzerland

Multi-objective optimisation in scientific workflow.

Hoang Anh Nguyen[1*], Zane van Iperen[1], Sreekanth Raghunath[2,4], David Abramson[1], Timoleon Kipouros[3], Sandeep Somasekharan[4]

[1]*Research Computing Centre, The University of Queensland, Australia.*

[2]*Centre for Hypersonics, The University of Queensland, Australia*

[3]*Propulsion Engineering Centre, Cranfield University, U.K.*

[4]*Zeus Numerix Pvt. Ltd. Pune, India.*

Abstract

Engineering design is typically a complex process that involves finding a set of designs satisfying various performance criteria. As a result, optimisation algorithms dealing with only single-objective are not sufficient to deal with many real-life problems. Meanwhile, scientific workflows have been shown to be an effective technology for automating and encapsulating scientific processes. While optimisation algorithms have been integrated into workflow tools, they are generally single-objective. This paper first presents our latest development to incorporate multi-objective optimisation algorithms into scientific workflows. We demonstrate the efficacy of these capabilities with the formulation of a three-objective aerodynamics optimisation problem. We target to improve the aerodynamic characteristics of a typical 2D airfoil profile considering also the laminar-turbulent transition location for more accurate estimation of the total drag. We deploy two different heuristic optimisation algorithms and compare the preliminary results.

Peer-review under responsibility of the scientific committee of the International Conference on Computational Science

Keywords: Multi-objective optimisation; Scientific Workflow; Engineering Design.

1 Introduction

Engineering design typically involves searching for good solutions that meet various performance criteria and constraints. These optimisation problems are complex because of several characteristics. First, they typically involve more than one, and often conflicting, objective functions. Although objective functions can be aggregated into a single one and thus simplifying the problem, multi-objective optimisation is generally considered to be more effective than single objective because it allows more flexible trading between objectives [1]. Second, the evaluations of objective functions might be computationally intensive and time consuming, which often requires access to supercomputers. Additionally, the optimisation domain can be large, making the whole optimisation

* Corresponding author: uqhngu36@uq.edu.au

1877-0509 © 2017 The Authors. Published by Elsevier B.V.
Peer-review under responsibility of the scientific committee of the International Conference on Computational Science
10.1016/j.procs.2017.05.213

process even more computationally intensive and time consuming. The third characteristic is related to the implementation details of the optimisation algorithms. As global optimisation is known to be NP-complete, heuristics are often required to find good solutions in a reasonable time [2]. The last several decades have seen much investment in meta-heuristics, especially nature-inspired ones such as simulated annealing and genetic algorithms [3]. This results in many different implementations of continuously appeared meta-heuristics, which creates another level of *complexity* for users when dealing with optimisation problems in terms of problem formulation, benchmarking, etc.

Scientific workflow technology has been shown to be effective for automating and encapsulating scientific processes [4]. Scientific workflow engines simplify the programming task by providing a high-level environment in which users connect a set of previously defined components, implementing a computational pipeline that meets their needs. These engines are typically integrated with distributed computing middleware, allowing workflows to distribute computation intensive jobs to high-performance computing platforms. There have been a large number of workflow engines produced in recent years, but all provide similar capabilities and functions. For a fairly recent view, we refer the reader to [3].

While optimisation codes are traditionally monolithic, solving optimisation problems with scientific workflows brings several benefits. First, the workflow can expose the components of the optimisation process to the user [5], making it is relatively straightforward to replace existing or add new components. Second, the optimisation can use distributed computing support embedded in scientific workflow engines, allowing computational intensive jobs to be off-loaded to high performance machines. With these benefits, scientific workflows have been used to formulate and solve optimisation problems [5]–[7]. However, only single objective optimisation is currently supported by these systems.

In this paper, we present our latest developments that enable multi-objective optimisation. Importantly, we aim to develop an *extensible* framework for future integration of different optimisation algorithms. We demonstrate the extensibility of our design by integrating two implementations of popular multi-objective optimisation algorithms. The paper then presents the solution of a state-of-the-art airfoil shape optimisation that employs a new methodology resulting in improved estimation of the total drag. Details of the methodology are presented later in Section 4.

2 Background

Optimisation is the process of searching a parameter space for the good solutions to a problem that is defined by one or multiple objective functions. In case of only one objective function, the solution is a point in the search space such that the objective function achieves its optimal (minimal or maximal) value [8]. When there is more than one objective function, it is unlikely to have a single point that is optimal for all the functions, especially if there are conflicting objectives. Instead, the solution for these problems is a set of Pareto optimal points [9]. Each Pareto optimal point is non-dominant to other point because the value of one object cannot be improved further without scarifying other objects.

As mentioned earlier, multiple objectives can be aggregated into single objective, often by a weighted sum of the objectives [8]. This allows single-objective optimisation techniques to be used to optimise the composite objective function. However, it requires the composite function to be pre-defined and the optimisation process generates a single vector as the optimal solution. This approach is not ideal in real-life use cases because the composite function reflects the designers' assumptions about the trade-off between objective functions, rather than the actual trade-offs [1]. As a result, multiple-objective optimisation is better to solve real-life optimisation problems.

One class of optimisation algorithm that attracts research interests is meta-heuristics. As opposed to problem-specific heuristics, meta-heuristics are algorithms designed to solve a range of problems