

Middleware for Data Intensive Analytics on HPC

Ioannis Paraskevakos
Rutgers University
Piscataway, New Jersey, USA

ACM Reference Format:

Ioannis Paraskevakos. 2018. Middleware for Data Intensive Analytics on HPC. In *Proceedings of 47th International Conference on Parallel Processing, August 13-16, 2018, University of Oregon, Eugene, Oregon, USA (ICPP)*. ACM, New York, NY, USA, 1 page. <https://doi.org/>

1 EXTENDED ABSTRACT

My research has focused on the design and development of Middleware for Data-intensive Analytics and Science (MIDAS). MIDAS provides the necessary resource management capabilities which enable the use of Big Data frameworks on HPC resources.

MIDAS has been used for two distinct application and analysis: First, it supports data-intensive analysis applications in conjunction with traditional HPC applications. It has enabled the analysis of biophysical simulations systems that were not feasible before. Second, MIDAS is being used to analyze high-resolution satellite imagery and derived products (such as digital elevation models). In addition to large volumes of data involved, the analysis of images is computationally intensive and requires high-performance and distributed resources.

My research focuses on providing a common workload management middleware abstraction, on distributed resources, that will capture the commonalities of the data analysis between different domain sciences, and a resource management abstraction that will allow Big Data style analytics on HPCs in conjunction with HPC style simulations. One important component of these abstractions is the Pilot abstraction. The Pilot abstraction operates as a resource placeholder which can be used to execute a set of heterogeneous tasks within the placeholder on the same resource. I am using the Pilot-Abstraction and its implementation RADICAL-Pilot to support different types of data analysis as well as use it as a building block for the middleware abstraction defined before.

Currently, I am a fourth year PhD candidate in the Electrical and Computer Engineering (ECE) department at Rutgers, the State University of New Jersey and a member of the RADICAL group, lead by Prof. Shantenu Jha. Prior to my PhD, I obtained my master degree in Integrated Hardware and Software Systems and a 5-year diploma (bachelor & masters of engineering) degree in Computer Engineering and Informatics from the University of Patras, Greece.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ICPP, 2018

© 2018 Copyright held by the owner/author(s).

ACM ISBN .

<https://doi.org/>