



**VINEYARD**

# Programming Language and Runtime System: Early Prototype

<b>DOCUMENT ID</b>	<b>D4.2</b>	<b>CONTRACT START DATE</b>	<b>1<sup>st</sup> FEBRUARY 2016</b>
<b>DUE DATE</b>	<b>31/03/2017 (extended)</b>	<b>CONTRACT DURATION</b>	<b>36 Months</b>
<b>DELIVERY DATE</b>	<b>31/03/2017</b>		
<b>CLASSIFICATION</b>	<b>Confidential</b>		
<b>AUTHOR/S</b>	<b>Hans Vandierendonck</b>		
<b>DOCUMENT VERSION</b>	<b>1.0</b>		

**PARTNERS**

Institute of Communications and Computer Systems, Maxeler Technologies, Bull Systems, Queen's University of Belfast, Foundation for Research and Technology, The Hartree Centre / Science and Technologies Facilities Council, Neurasmus BV, Neurocom Luxembourg, Hellenic Exchanges SA, Holding, Clearing, Settlement and Registry, LeanXcale, Loba



Co-funded by the Horizon 2020 Framework Programme of the European Union under Grant Agreement n° 687628

## 1 EXECUTIVE SUMMARY

This document describes the core functionality of the VINEYARD programming model and runtime system for accelerated data centres. We describe our approach to creating an abstract representation of accelerated kernels, such that application programmers can use these kernels without needing to worry about accelerator-specific calling conventions, or about the specific versions available in the VINEYARD accelerator library.

The second key contribution of this document is the description of our approach to virtualizing accelerators. We assume that accelerators are assigned to jobs only when they are really needed, and not at job allocation time. This raises issues that need to be addressed in the virtualization layer and also in the application's runtime. We describe these issues and our approach to solving.

Finally, we present an overview on the current state of the VINEYARD accelerator library.