

## Accepted Manuscript

Long-time simulations with complex code using multiple nodes of Intel Xeon Phi Knights Landing

Jonathan S. Graf, Matthias K. Gobbert, Samuel Khuvis

PII: S0377-0427(18)30018-9  
DOI: <https://doi.org/10.1016/j.cam.2017.12.050>  
Reference: CAM 11469

To appear in: *Journal of Computational and Applied Mathematics*

Received date: 24 May 2017  
Revised date: 22 December 2017

Please cite this article as: J.S. Graf, M.K. Gobbert, S. Khuvis, Long-time simulations with complex code using multiple nodes of Intel Xeon Phi Knights Landing, *Journal of Computational and Applied Mathematics* (2018), <https://doi.org/10.1016/j.cam.2017.12.050>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.



# Long-time Simulations with Complex Code Using Multiple Nodes of Intel Xeon Phi Knights Landing

Jonathan S. Graf<sup>a</sup>, Matthias K. Gobbert<sup>a,\*</sup>, Samuel Khuvis<sup>a</sup>

<sup>a</sup> Department of Mathematics and Statistics, University of Maryland, Baltimore County, 1000 Hilltop Circle, Baltimore, MD 21250, U.S.A.

## Abstract

Modern partial differential equation (PDE) models across scientific disciplines require sophisticated numerical methods resulting in complex codes as well as large numbers of simulations for analysis like parameter studies and uncertainty quantification. To evaluate the behavior of the model for sufficiently long times, for instance, to compare to laboratory time scales, often requires long-time simulations with small time steps and high mesh resolutions. This motivates the need for very efficient numerical methods and the use of parallel computing on the most recent modern architectures. We use complex code resulting from a PDE model of calcium dynamics in a heart cell to analyze the performance of the recently released Intel Xeon Phi Knights Landing (KNL). The KNL is a second-generation many-integrated-core (MIC) processor released in 2016 with a theoretical peak performance of over 3 TFLOP/s of double-precision floating-point operations for which complex codes can be easily ported because of the x86 compatibility of each KNL core. We demonstrate the benefit of hybrid MPI+OpenMP code when implemented effectively and run efficiently on the KNL including on multiple KNL nodes. For multi-KNL runs for our sample code, it is shown to be optimal to use all cores of each KNL, one MPI process on every other tile, and only two of the maximum of four threads per core.

**Keywords:** Intel Xeon Phi; Knights Landing; MPI; OpenMP; Parabolic partial differential equations; Calcium Induced Calcium Release.

**2000 MSC:** 35K61 65M08 65Y05 68U20 92C35

## 1. INTRODUCTION

The size and structure of modern processors has developed significantly in recent years. As the rapid processing speed increases of a single chip stalled in the presence of the physical issues of power consumption and heat generation, a shift to multi-core architectures occurred. Today, CPUs in consumer devices are dual- or quad-core. The iPhone 7 features a quad-core processor, as do most mainstream laptops. Typical state-of-the-art distributed-memory clusters contain two multi-core CPUs per node with, for instance, 8 to 16 cores. Recent developments in parallel computing architectures also include the use of graphics processing units (GPUs) as a massively parallel accelerator, with thousands of special purpose cores, in general purpose computing and many-integrated-core (MIC) architectures like the Intel Xeon Phi with more than 60 cores. Besides the larger number of computational cores in both GPU and Phi, the key difference to a CPU is each one's significant on-chip memory, on the order of several GB, which contributes significantly to their performance gain over CPUs. A difference between GPU and Phi is the x86 compatibility of each Xeon Phi core that makes porting of code from Intel CPUs to this architecture much more readily possible, typically by recompiling with the suggested addition of a compiler flag.

The recent emergence of the second-generation Intel Xeon Phi in 2016, codenamed Knights Landing (KNL), represents a significant improvement over the first-generation in 2012, codenamed Knights Corner (KNC). The KNL was announced in June 2014 [11] and began shipping in July 2016. The KNL itself is like a 'massively parallel' supercomputer from the early 2000s with dozens of nodes connected by a Cartesian network, all in a single chip now with a theoretical peak performance of over 3 TFLOP/s of double-precision floating-

\*Corresponding author Tel. +1 410 455 2404; fax +1 410 455 1066.

Email addresses: [jongraf1@umbc.edu](mailto:jongraf1@umbc.edu) (Jonathan S. Graf), [gobbert@umbc.edu](mailto:gobbert@umbc.edu) (Matthias K. Gobbert), [khsa1@umbc.edu](mailto:khsa1@umbc.edu) (Samuel Khuvis)

Download English Version:

<https://daneshyari.com/en/article/8902048>

Download Persian Version:

<https://daneshyari.com/article/8902048>

[Daneshyari.com](https://daneshyari.com)