



# DEVELOPMENT OF INTEL XEON PHI ACCELERATED ALGORITHMS AND APPLICATIONS AT IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTRE

Vít Vondrák, Lubomír Říha, Michal Merta, Jan Zapletal, Milan Jaroš, ...

# Introduction

## IT4Innovations National Supercomputing Centre – [www.it4i.cz](http://www.it4i.cz)

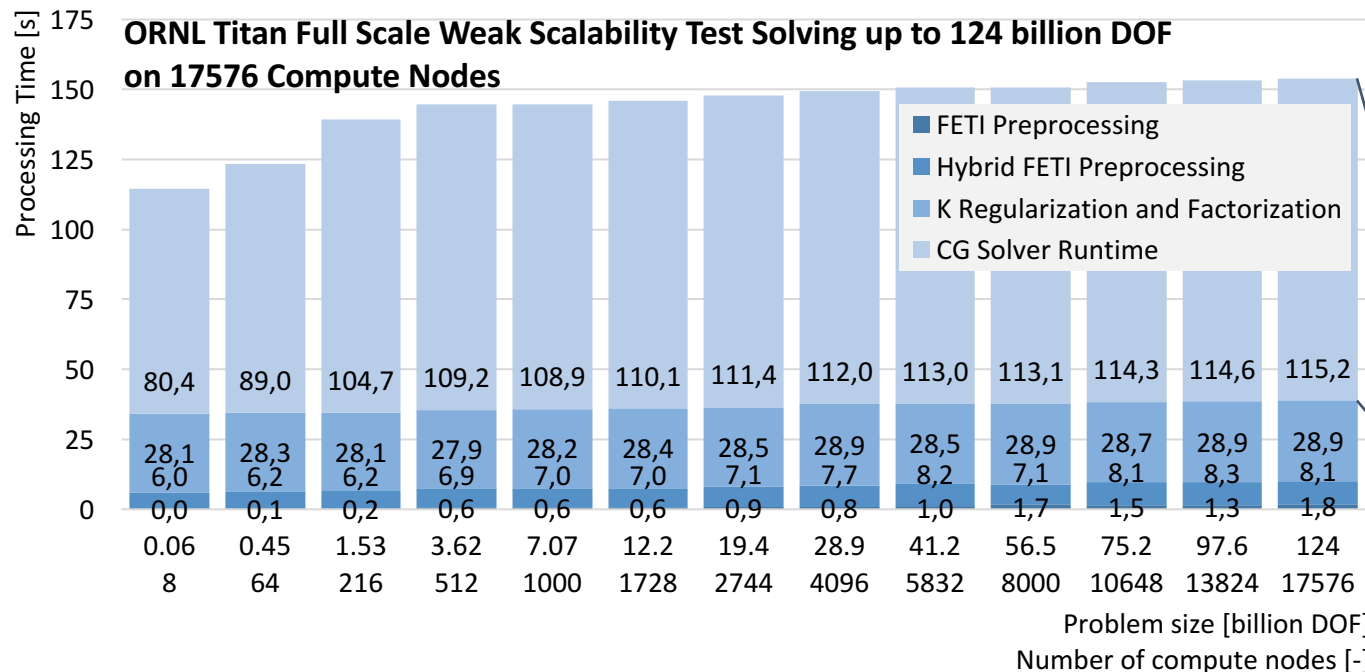

- Salomon Supercomputer - 1008 computational nodes of which 576 are regular compute nodes and 432 accelerated nodes (two Intel Xeon Phi 7120P co-processors per node)

## Selected Applications from Parallel Algorithms Research Lab at IT4Innovations

- **IPC Center at IT4Innovations** – acceleration of the **ESPRESO** library and its integration into ELMER and OpenFOAM community codes (<http://ipcc.it4i.cz>, <http://espresso.it4i.cz>)
- **BEM4I** – development of efficient Boundary Element Method (BEM) assembler for multi- and many-core architectures with wide SIMD units (<http://bem4i.it4i.cz>)
- **Blender** – parallelization of the Blender Cycles rendering engine by MPI and OpenMP with support for Intel Xeon Phi accelerators (<http://ipcc.it4i.cz/files/Download/>)

# ESPRESSO Massively Parallel Sparse Linear Solver

- FETI domain decomposition based solver supporting massively parallel solution of structural mechanics problems
- Supports acceleration by Intel Xeon Phi processors – KNC and KNL
- Provides general API for other open source libraries (ELMER, OpenFOAM)



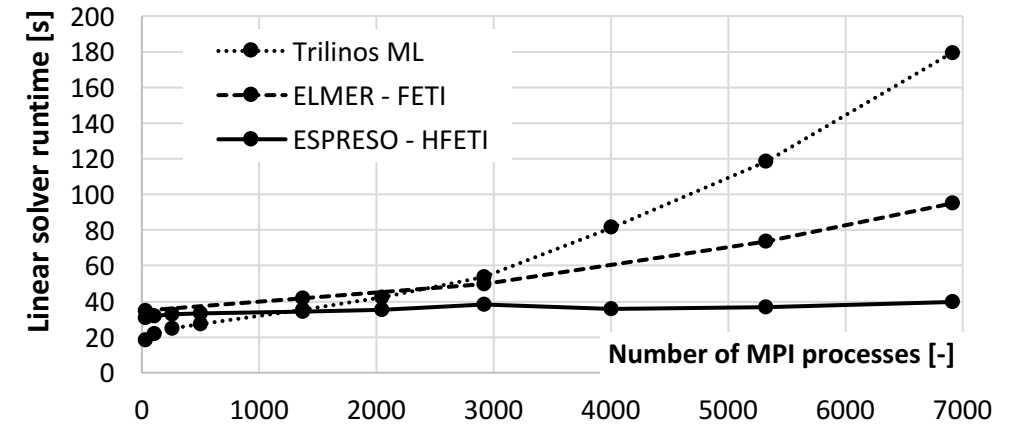
This part of the solver is accelerated by the Intel Xeon Phi co-processor

Designed to solve real world problems generated by open-source (ELMER and OpenFOAM) and commercial tools (Ansys Workbench)

# ESPRESO Advanced Features


- Within the IPCC project at IT4Innovations the library was accelerated by offloading to Intel Xeon Phi (KNC) co-processors
- Algorithmic changes have been made to replace actions with sparse FEM data structures with dense objects to fully utilize the coprocessor
- Acceleration of two key routines: (1) system matrix processing and (2) preconditioner application
- Supports offload to multiple coprocessors and dynamic CPU/co-processor load balancing
- Scalability tested on the Salomon supercomputer with up to 864 Xeon Phi 7120P co-processors
- Improved scalability of community codes (ELMER)

Weak Scalability Evaluation of Selected Linear Solvers in ELMER




Compute nodes [-];	2;	5;	11;	21;	58;	86;	122;	167;	222;	288;
MPI ranks [-];	32;	108;	256;	500;	1372;	2048;	2916;	4000;	5324;	6912;
Problem size [million DOF]	3.4	11.6	27.6	53.9	147	220	314	431	574	745

Dynamic load balancing in ESPRESO




# Offloading Approaches in ESPRESO




## Matrixpack offload

- one thread on host offloads all matrices within one array (MatrixPack)
- One parallel OpenMP region on device



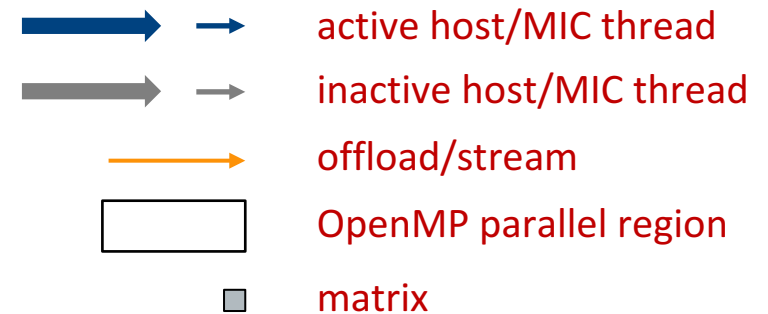
## Individual offload

- Each thread on host offloads its matrices individually
- No OpenMP region on device




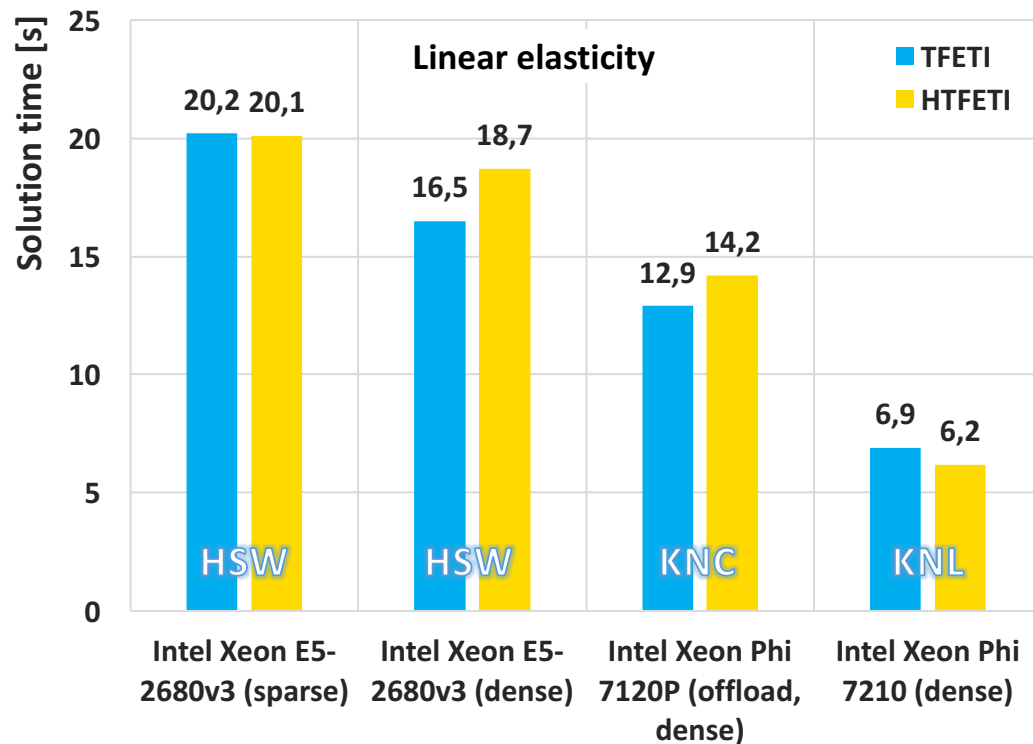
## Matrix semi-pack offload

- Each host thread processes one matrix pack
- Multiple OpenMP regions on device



# The Future of ESPRESO – Porting to KNL

- Tests were performed on Intel Xeon Phi 7210 processors at Intel Endeavor cluster
- Slower preprocessing phase, faster solve – suitable for problems requiring large number of iterations



**Note:** Comparison of iterative solver runtime on different Xeon architectures using sparse/dense matrix structures



# BEM4I – Xeon Phi Accelerated Library

## Boundary element method

- alternative to FEM
- discretization of boundary only
- singular surface integrals
- dense matrices (classic BEM)

## BEM4I ([bem4i.it4i.cz](http://bem4i.it4i.cz))

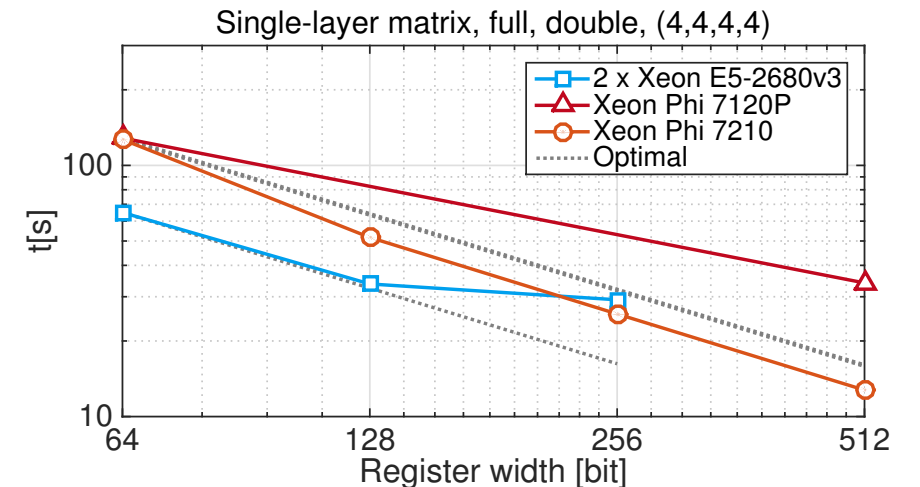
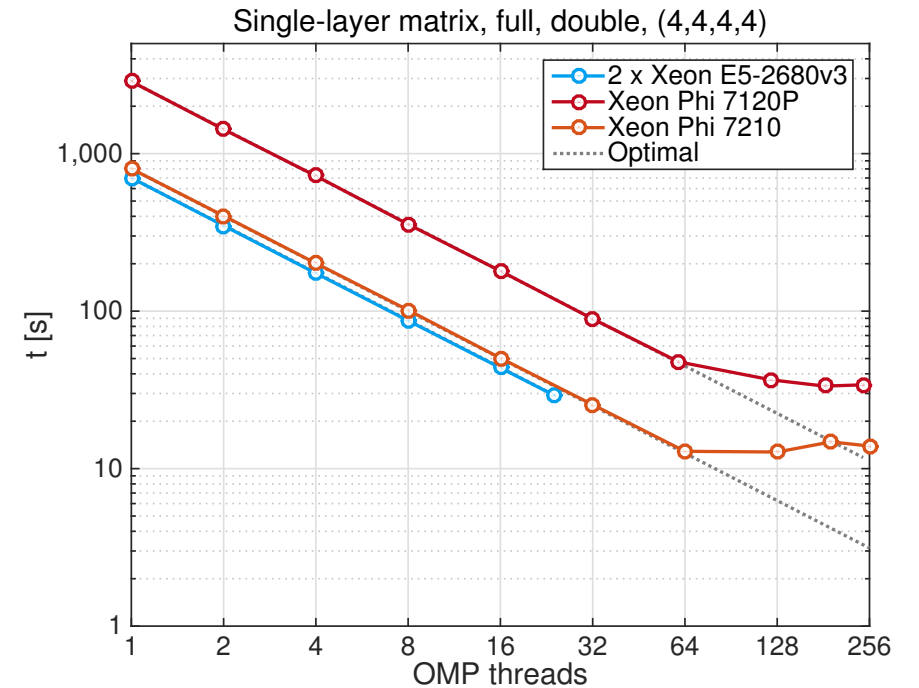
- solver developed at IT4Innovations
- optimized for Intel Xeon and Xeon Phi processors

## Performance (threading)


- 2 x Xeon E5-2680v3 (Haswell) processor with 24 cores
  - speedup 24 w.r.t. single threads
- Xeon Phi 7120P (Knights Corner) co-processor with 61 cores
  - speedup 61+
- Xeon Phi 7210, (Knights Landing) processor with 64 cores
  - speedup 64+

## Performance (SIMD)

- speedup 8+ with AVX-512 w.r.t. scalar code



# BEM4I – Xeon Phi Accelerated Library



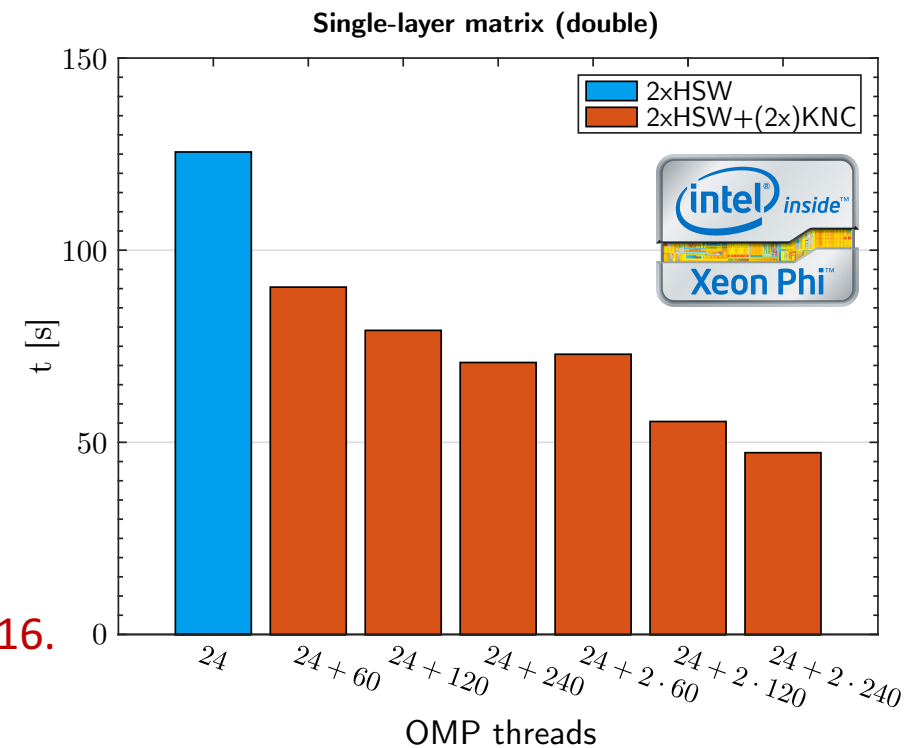
## Possible offload to Intel Xeon Phi (KNC) co-processor

- Intel LEO pragmas,
- one OpenMP thread per MIC,
- asynchronous assembly on host and MIC.

## Performance

- offload to 2 x Xeon Phi 7120P (KNC) co-processor
  - speedup up to 2.7.



Journal publication: Merta, M.; Zapletal, J.; Jaros, J. Many Core Acceleration of the Boundary Element Method. LNCS, Springer, 2016.





# Massively parallel version: BETI = BEM4I + ESPRESO

- Boundary Element Tearing and Interconnecting (alternative to FETI originally implemented in ESPRESO)
- massively parallel – SIMD + OpenMP + MPI + offload to Intel Xeon Phi (KNC) co-processor



Number of MPI processes – 2 MPI processes per node with OpenMP  
 Problem size in surface elements [in billions]  
 Problem size in volume elements [in billions]

# Acceleration of Blender Cycles Rendering Engine


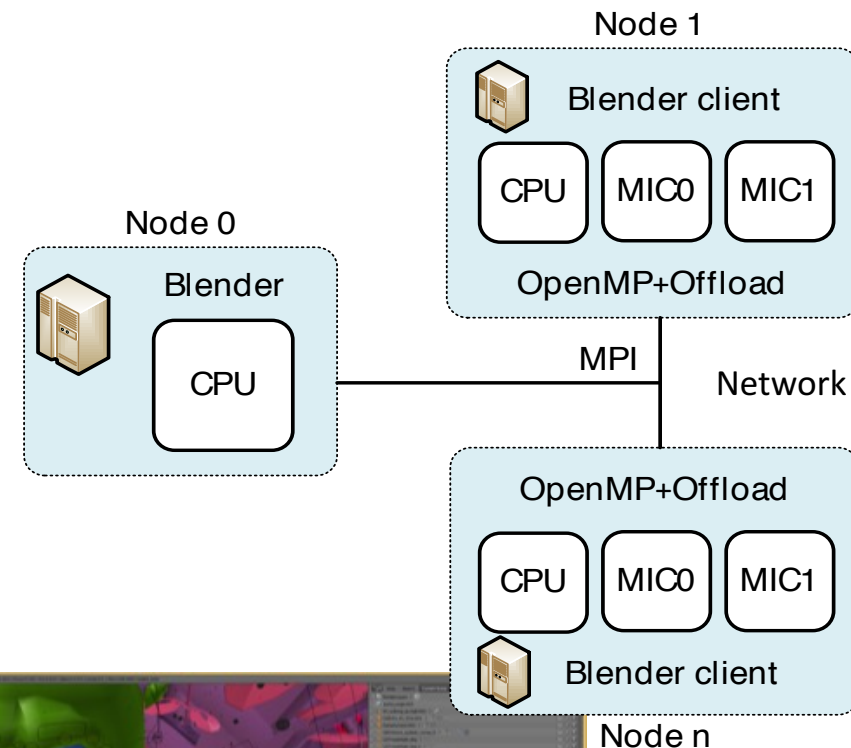
IT4Innovations  
national  
supercomputing  
center



CC-BY-SA, Blender Institute, based on Agent 327 comics (c) Martin Lodewijk

# CyclesPhi – Accelerated Blender Cycles

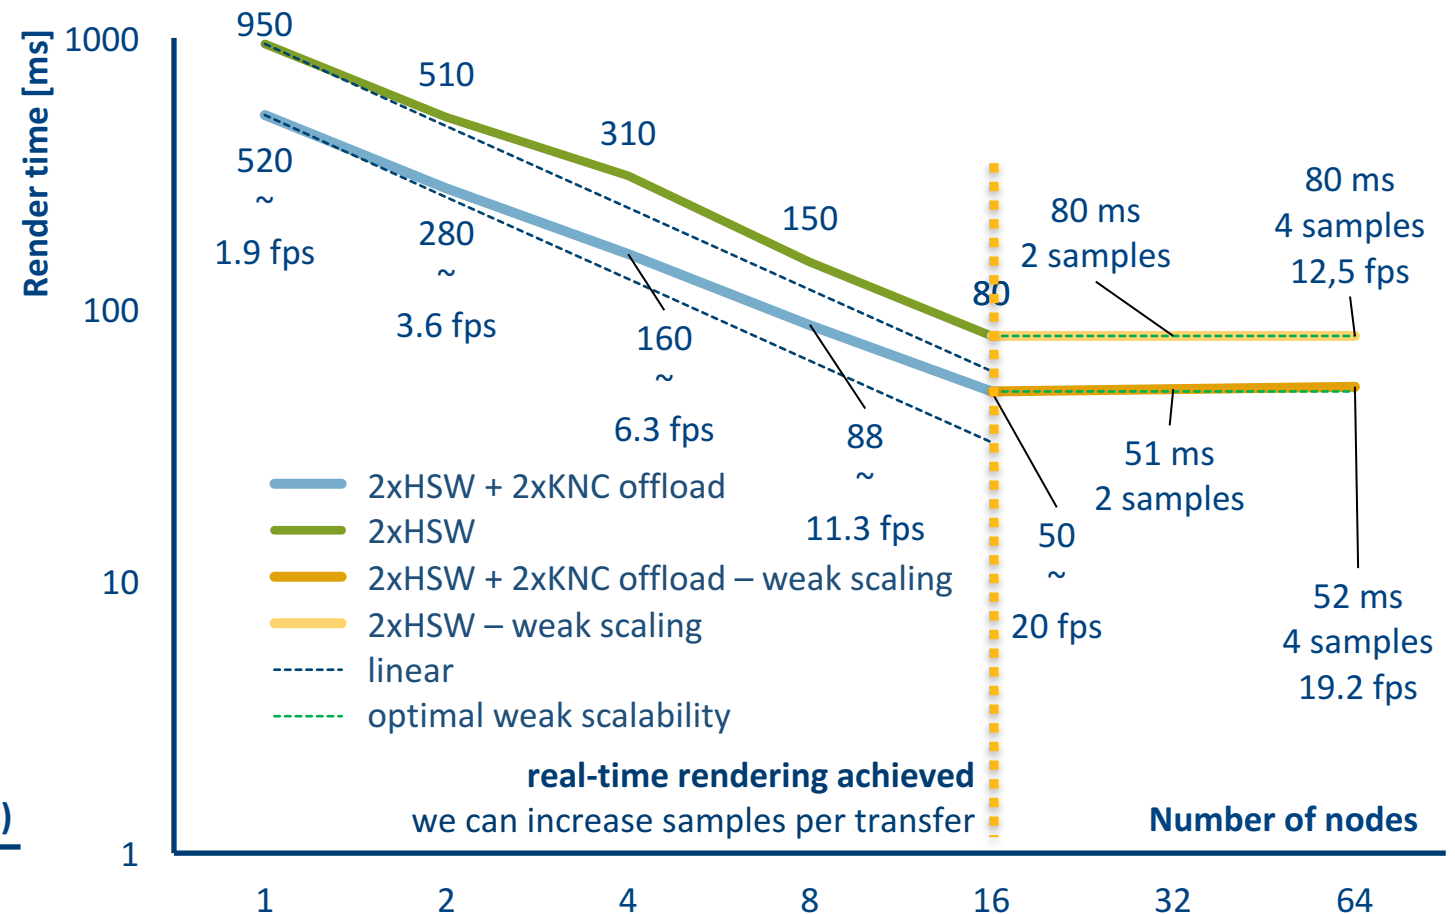
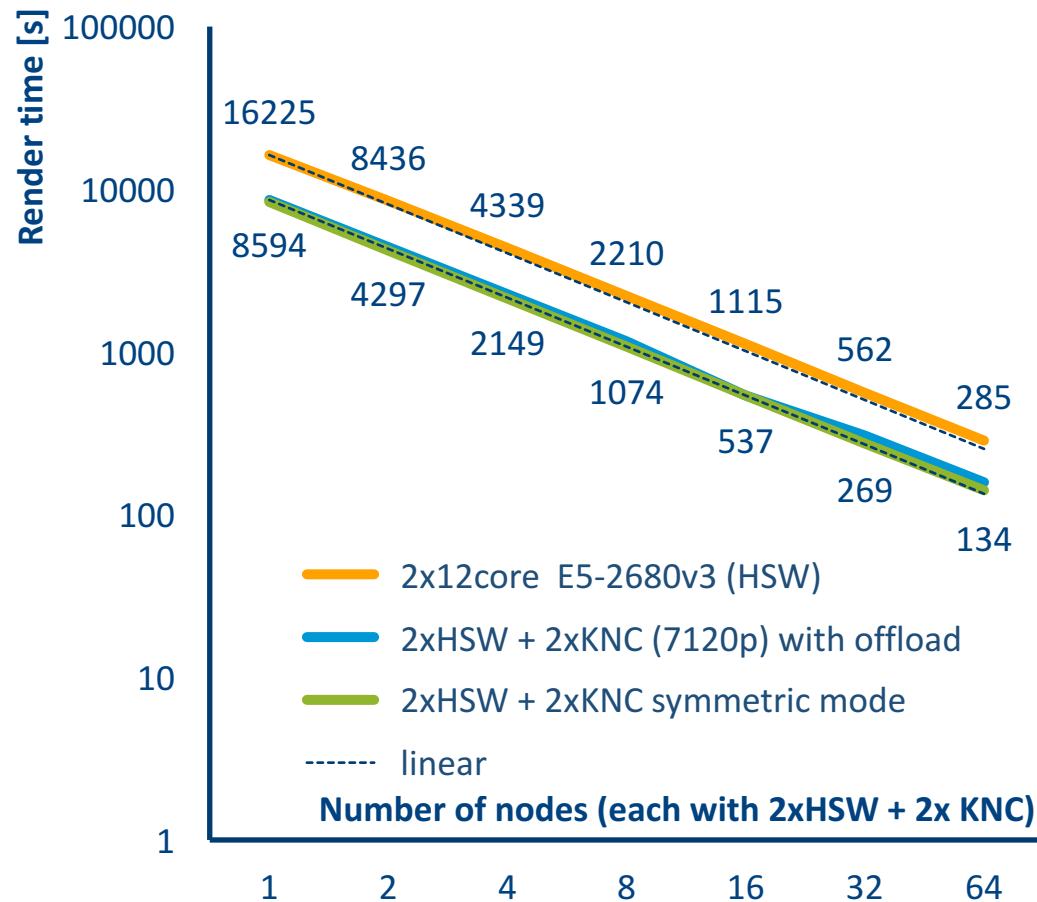
- Blender is an open source 3D creation suite. It has two render engines: Blender Internal and Cycles.
- Cycles is a raytracing based render engine with support for interactive rendering, shading node system, and texture workflow.
- We have modified the kernel of the Blender Cycles rendering engine and then extended its capabilities to support the HPC environment. We call this version the CyclesPhi and it supports following technologies:
  - OpenMP
  - MPI
  - Intel Xeon Phi (KNC) processor with Offload
  - Intel Xeon Phi (KNC) processor with Symmetric mode
  - And their combinations




# MPI Parallelization of Blender Cycles

Supports both interactive and offline rendering – with almost linear strong scalability

Tests on Salomon supercomputer: 2x E5-2680v3 (HSW) and 2x7120p (KNC) processors per node




# Acceleration of Blender Cycles by Xeon Phi Processor



Preprocessing time [s]	Classroom	Fishy cat	Koro	Pabellon Barcelona	Victor
Scene (see figures)					
1x Xeon Phi 7250 (KNL)	18 s	110 s	120 s	11 s	1680 s
1x Xeon Phi 7120p (KNC)**	3 s	12 s	14 s	4 s	120 s
2x HSW + 2x Xeon Phi 7120p (KNC)**	6 s	14 s	19 s	4 s	150 s
2x 12core E5-2680v3 (HSW)	1 s	8 s	10 s	1 s	90 s

\*\* preprocessing is done by the 2x 12core E5-2680v3 (HSW) and includes data transfer over PCIe bus

- Salomon supercomputer: 2x Intel Xeon E5-2680v3 CPUs + 2x Intel Xeon Phi 7120P (KNL) processors
- Intel Xeon Phi 7250 - HLRN-III Cray System





IT4Innovations  
national  
supercomputing  
center



THANK YOU