ComPat

Computing Patterns for High Performance Multiscale Computing

www.compat-project.eu

12 May 2016, Prague

Tomasz Piontek
Poznan Supercomputing and Networking Center

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 671564.

ComPat Project

Computing Patterns for High Performance Multiscale Computing

- Horizon 2020
- Call: H2020-FETHPC-2014
- Duration: 36 months
- Start: October 2015
- Current Status: 8M
- Coordinator: Prof. Alfons Hoekstra, UvA

ComPat Consortium

COMPAT
Computing Patterns for High Performance Multiscale Computing

University of Amsterdam

University College London

Poznan Supercomputing and Networking Centre

Leibniz Supercomputing Centre

Max-Planck-Institut für Plasmaphysik

University Leiden

The Hartree Centre/STFC

Allinea Software

CBK Sci Con Limited

ITMO University

World is multi-scale

All the studied complex phenomena consist of many sub-processes on **disparate length and time scales** that interact in strong and non-linear ways.

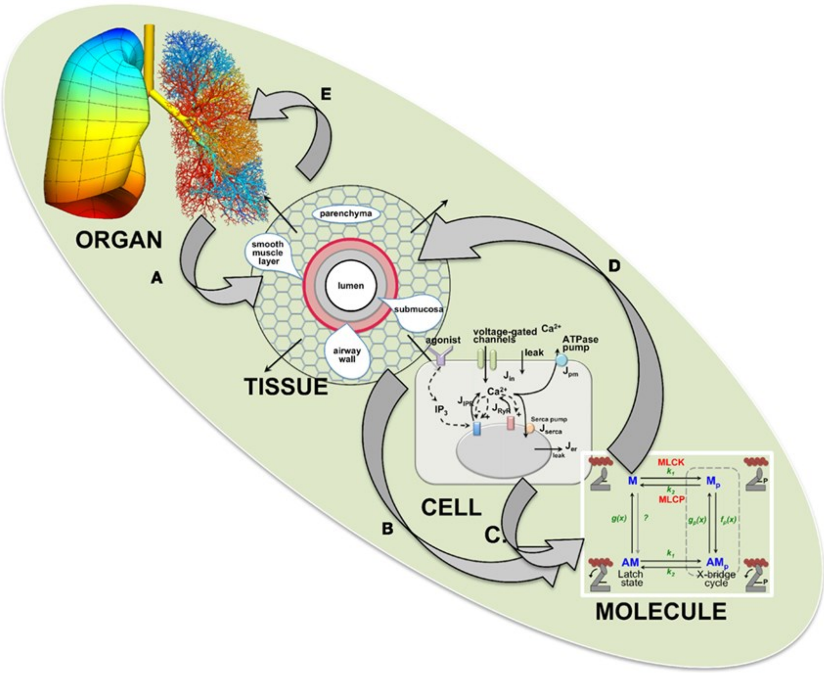

Multi-scale approach

In a multiscale simulation, each relevant scale needs its own type of solver. Accordingly, multiscale model is defined as a **collection of coupled single scale models** that can be computed reliably with a dedicated, so-called “monolithic” solver.

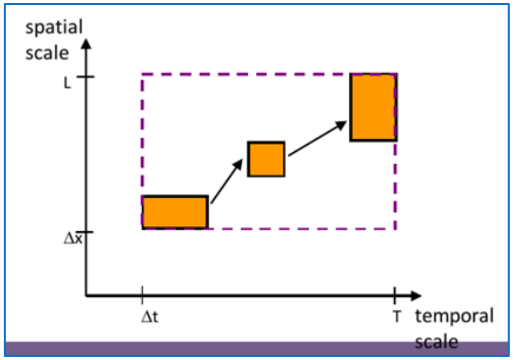

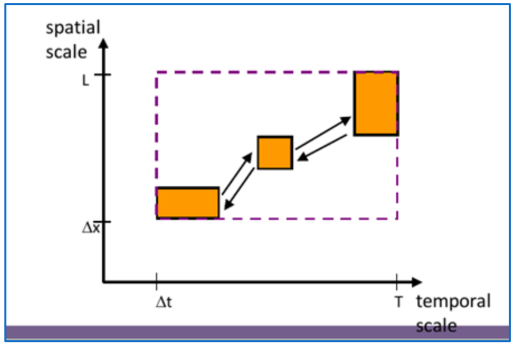

ComPat Objectives

Project objective is to develop generic and reusable **High Performance Multiscale Computing Patterns (schemas of processing)** that will address the exascale challenges posed by heterogeneous architectures and will enable to run **multiscale applications** with extreme data requirements while achieving scalability, robustness, resiliency, and energy efficiency.

ComPat Objectives

Our **ambition** is to establish new standards for multiscale computing at exascale, and provision a robust and reliable software technology stack that empowers multiscale modellers to transform computer simulations into predictive science.

HPMC Patterns

We have proposed and formalised three **multiscale computing patterns for multi-scale applications**, incorporating customized algorithms for load balancing, data handling, fault tolerance and energy consumption under generic exascale application scenarios.

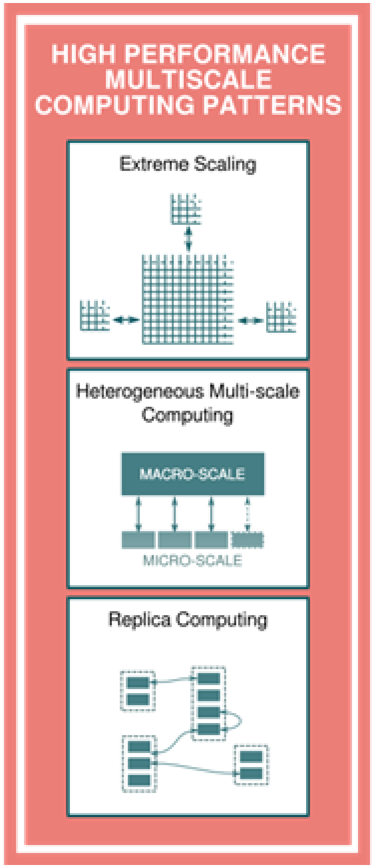

HPMC Patterns

- **Extreme Scaling** - one (or perhaps a few) of the single scale models in the overall multiscale model dominates all others by far, in terms of required computing power.
- **Heterogeneous Multiscale Computing** – coupling of a macroscopic model to a large and dynamic number of microscopic models. A database stores previously calculated values that can be used to interpolate input for macroscale model.
- **Replica Computing** - combines a potentially very large number of independent simulations ('replicas') to explore the parameter space.

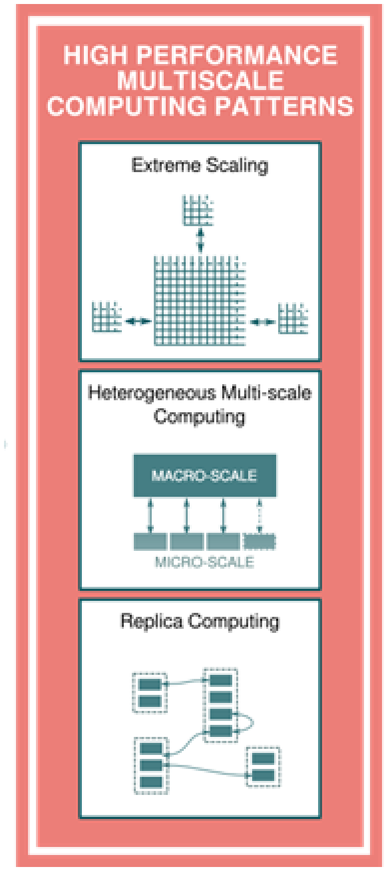

- [Hybrid Approach]– combination of basic patterns

Multiscale modeling and simulation framework

Conceptual Framework | Computational Framework

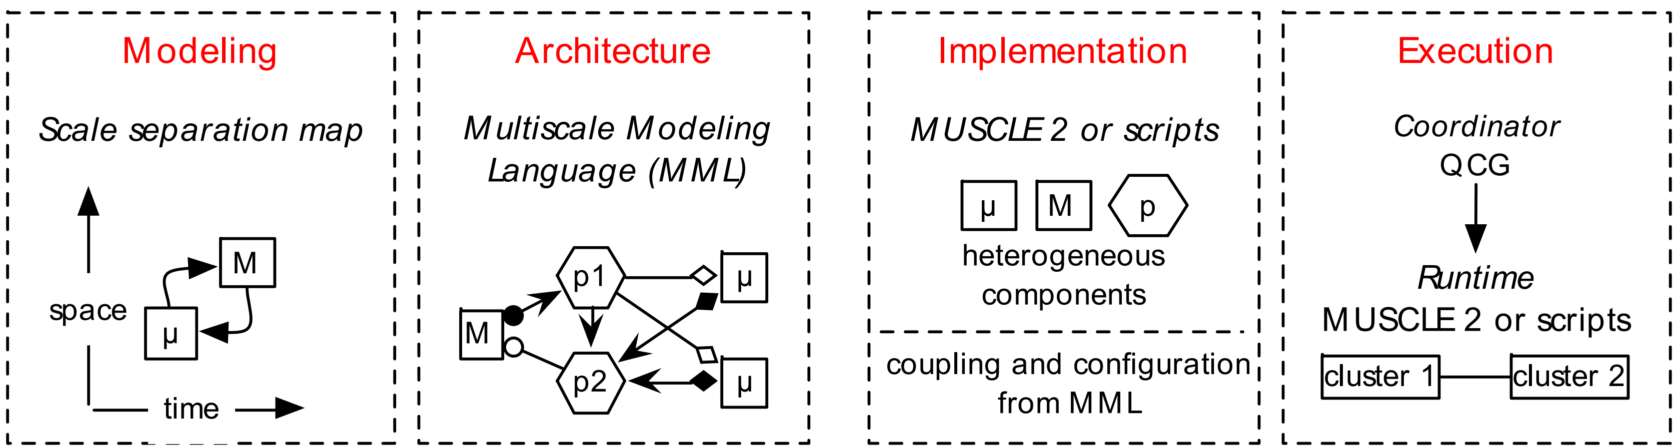

1. Modelling of the phenomena by identifying relevant processes (single scale models) and their relevant scales, as well as their mutual couplings.
2. The single scale models and their coupling are specified with the Multiscale Modelling Language (MML) thereby forming the architecture of a multiscale model.
3. A coupling library like MUSCLE ensures communication between heterogeneous components is possible, with minimal and local changes to the single scale code.
4. Sub-models are executed on a computing infrastructure. Each sub-model may require different computing resources.

ComPat stack

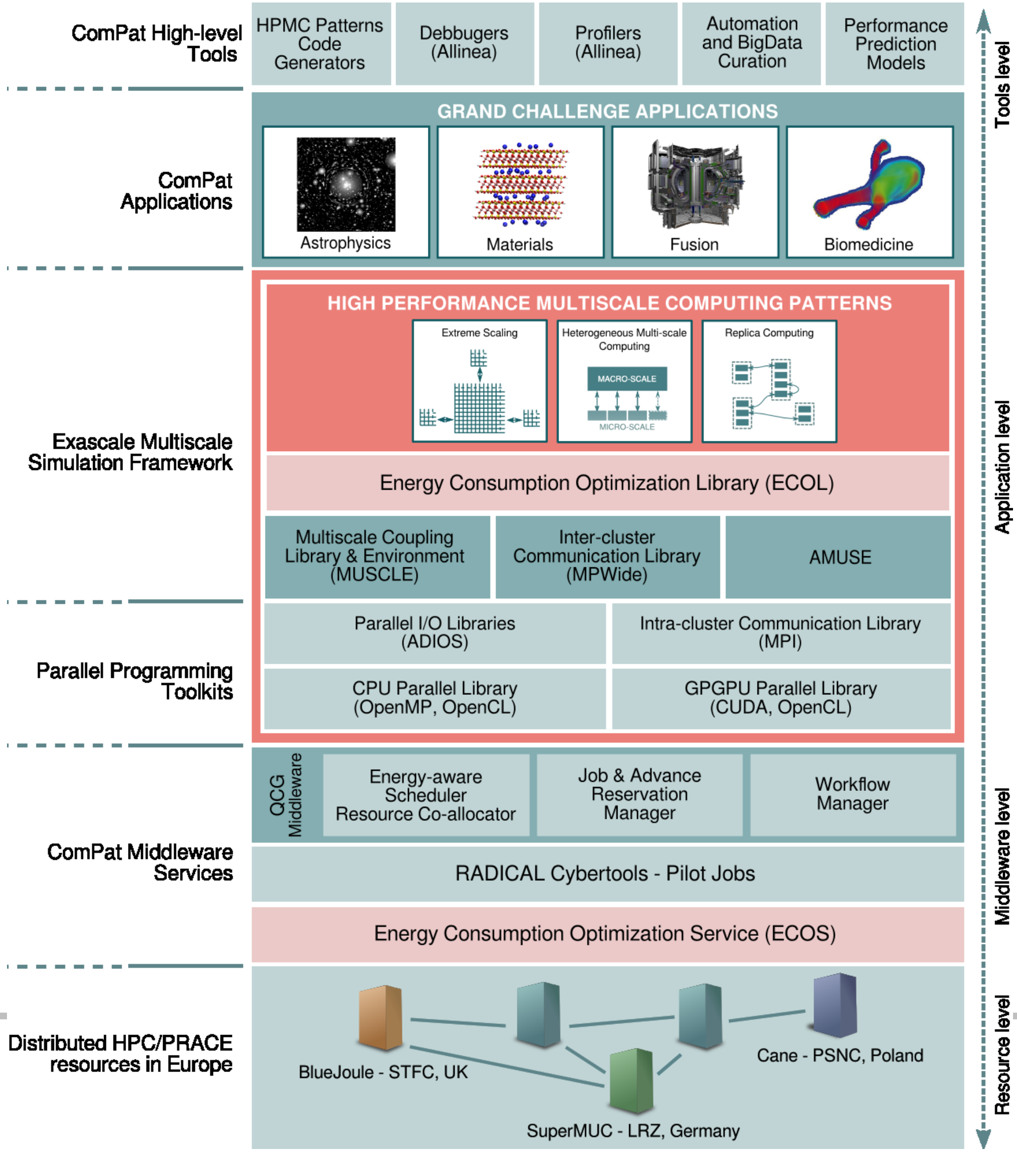

Multiscale software development cycle

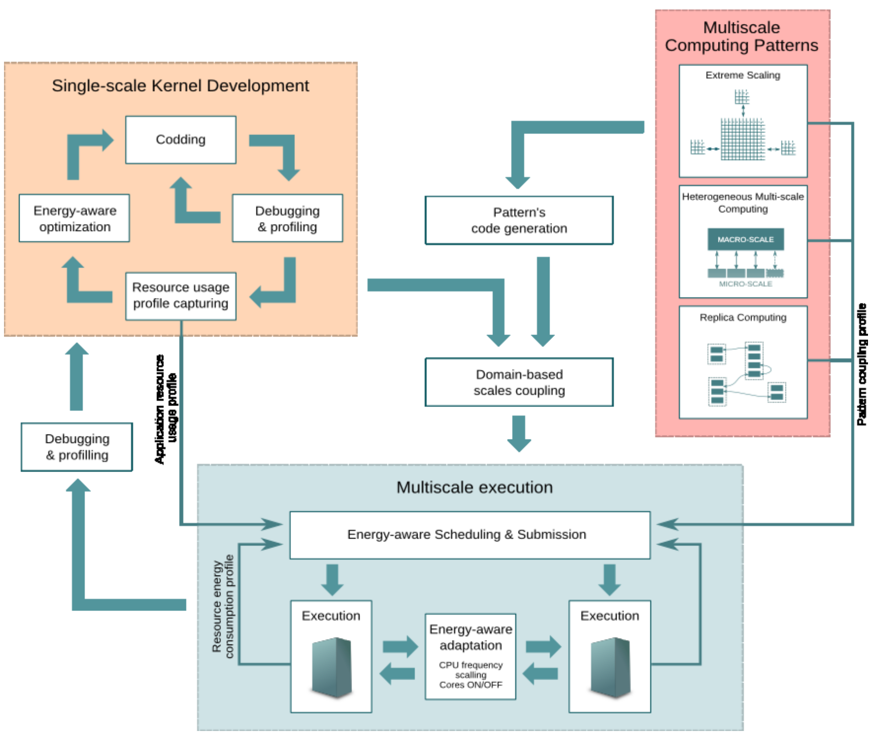

1. Independent design, implementation and optimisation of every single-scale application kernel
2. Identification of the HPMC pattern
3. Generation of the template / skeleton for the pattern based on formalized description
4. Embedding of the single scale models into the generated multiscale application skeleton (coupling)
5. Execution of the application on the infrastructure with taking into account energy efficiency (dynamic adaptation of resource properties)

Applications details

	Extreme Scaling	Heterogeneous Multiscale Computing	Replica Computing
Fusion (MPG-IPP)	global turbulence simulation	flux-tube chain	-
Biomedicine (UvA)	RBC and platelet transport	blood rheology	-
Biomedicine (IMTO + UvA)	In-stent restenosis	-	In-stent restenosis (*)
Biomedicine (UCL)	aneurysm flow dynamics	aneurysm flow dynamics(*)	
Material Science (UCL)		"on-the-fly" coarse-graining	phase behaviour (*)
Astrophysics (UL)	Milky-Way Galaxy simulation	Milky-Way Galaxy simulation (*)	-

	Core count (state-of-the-art)	Core count (desired)
Extreme Scaling		
Fusion (MPG-IPP)	400,000	4,000,000
Biomedicine (UvA)	45,000	4,000,000
Biomedicine (ITMO + UvA)	4,000	4,000,000
Biomedicine (UCL)	49,000	600,000
Astrophysics (UL)	500,000	10,000,000
Heterogeneous Multiscale Computing		
Fusion (MPG-IPP)	16,000	120,000
Biomedicine (UvA)	45.000	4,000,000
Biomedicine (UCL)	49,000	750,000
Material Science (UCL)	294.000	2,000,000
Astrophysics (UL)	1,000	100,000
Replica Computing		
Biomedicine (ITMO + UvA)	4,000	400,000
Material Science (UCL)	294,000	3,000,000

ComPat vs EsD

ComPat fully support the idea of cross-project integration and technology uptake by industry.

ComPat already follows the EsD guidelines

ComPat – example of the Single Project EsD

- Co-design (technology driven by applications)
- Technology Readiness Level – QCG, MUSCLE (9)
- Structure of the Consortium (TP, AO, RP)
- Phase A + B

ComPat in EsD Projects

- Technology Providers
  - HPMC Patterns (ComPat)
  - QCG Middleware (PSNC)
  - MUSCLE Coupling Library (UvA, PSNC)
  - Energy Consumption Optimization Service and Library (PSNC)
- Application owners
  - 9 grand challenge applications
- HPC Centers
  - LRZ - Leibniz Supercomputing Centre
  - PSNC – Poznan Supercomputing and Networking Center
  - STFC - Science and Technology Facilities Council

Tomasz Piontek
piontek@man.poznan.pl

www.compat-project.eu

This project has received funding from the European Union's Horizon 2020
research and innovation programme under grant agreement No 671564.