UNIVERSITY OF ICELAND
SCHOOL OF ENGINEERING AND NATURAL SCIENCES
FACULTY OF INDUSTRIAL ENGINEERING, MECHANICAL ENGINEERING AND COMPUTER SCIENCE

RAISE Center of Excellence

HPC Systems Engineering with the Interaction Room

Helmut Neukirchen, Matthias Book, Morris Riedel

University of Iceland

{helmut, book, morris}@hi.is

Interaction Room Might Work Well for Business Applications… …but HPC Is Not a Business Application

- In fact, similar considerations apply as well, e.g.:
  different implicit assumptions/knowledge between stakeholders:
- HPC projects often start with a domain expert writing HPC code.
  - E.g. Science/Engineering PhD student with no formal education in HPC nor Computer Science.
- The same problems at different occasions:
  - Code survives often academic scientists with temporary contracts:
    Code needs to be handed over to new scientists.

  ⇒ **Implicit assumptions and knowledge lost**.
  - Domain expert sooner or later struggles with HPC problems: code needs to be handed to HPC experts, e.g. at a supercomputing centre.

  ⇒ **Implicit assumptions and knowledge need to be transferred**.

Same within the CoE RAISE: Experts with different knowledge working together.

Interaction Room for HPC?

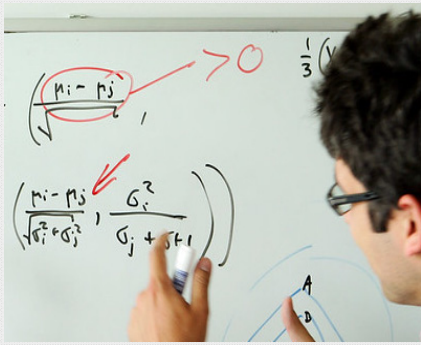

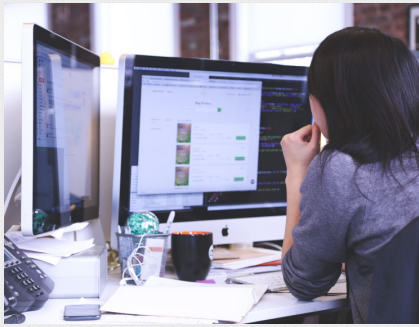

- Goal: **Facilitate collaboration** of experts from
  - the **natural science/engineering domain**,
  - the **HPC domain**,
  - the **AI/ML domain**,
  - the **computer science/software engineering domain**.

- **Adapt proven Interaction Room concepts**:
  - But: **canvases** needed that are **specific to HPC/AI/ML needs**.

Interaction Room for HPC!

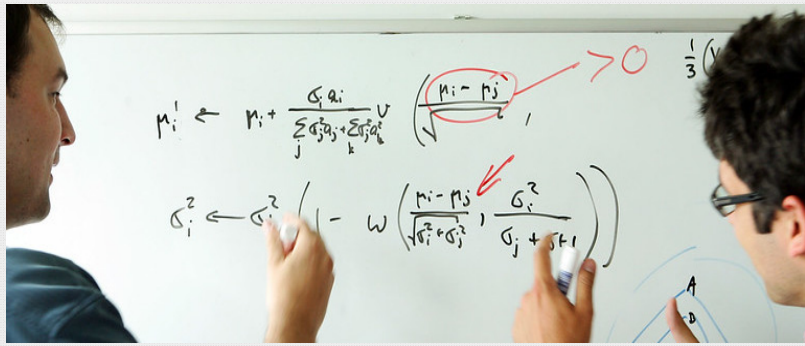

- HPC/AI-specific canvases that address:
  - Crucial **interdisciplinary discussion points**,
  - **Typical HPC/AI software development phases**.
- **Different domains** needing different canvas types, e.g.:
  - **Simulation sciences** (“classic” HPC),
  - **Data sciences** (HPDA) / **AI/ML**.
  - So far: Interaction Room canvases developed for HPC simulation sciences…
  - …to foster **collaboration within CoE RAISE**:
    go beyond ⇒ canvases for **HPC combined with AI**.

Book, Riedel, Neukirchen, Goetz: *Facilitating Collaboration in High-Performance Computing Projects with an Interaction Room*. 4th ACM SIGPLAN International Workshop on Software Engineering for Parallel Systems (SEPS 2017)

Crucial Interdisciplinary Communication Points in HPC Simulation Science Projects

RAISE Center of Excellence

- Science/Engineering domain experts need to help HPC experts **understand**, e.g.:
  - Research question we trying to answer?
  - What assumptions/simplifications are made?
  - What parameters and variables are there?
    - How do they affect each other? How do they evolve over time?
- HPC experts need to **validate** technical decisions with domain experts, e.g.:
  - Domain decomposition (How to map the problem most efficiently to the HPC cluster?),
  - HPC cluster HW architecture? (CPU or GPU? Memory-intensive or compute-intensive?),
  - Memory model (Distributed memory? Shared memory? Hybrid?),
  - Communication patterns (Choice of communication type? Ghosts/halos?).

Typical HPC Simulation Science Software Process to be Supported by Interaction Room

1. Understand the problem domain:
   a) In general (goal and scope),
   b) In detail. (How would serial problem domain code look like?)
2. Perform appropriate domain decomposition, choose appropriate communicators, etc.
3. Implement communication between processes/threads; mix into code of problem domain.
4. Test and validate simulation model and code.
5. Optimize accuracy, tune performance (specific to HW architecture).

Domain concepts to be supported by IR

Domain concepts to be supported by IR

HPC concepts to be supported by IR

HPC concepts to be supported by IR

Interaction Room Canvases for HPC Simulation Science Projects

- **Problem canvas:**
  - Goal and scope of research question (=the science domain).
- **Real-world canvas:**
  - Description of the pertinent aspects of the science domain.
- **Decomposition canvas:**
  - Breakdown of scientific model into parallelizable units.
- **Architecture canvas:**
  - Implementation of simulation on suitable HPC technology/HW architecture.

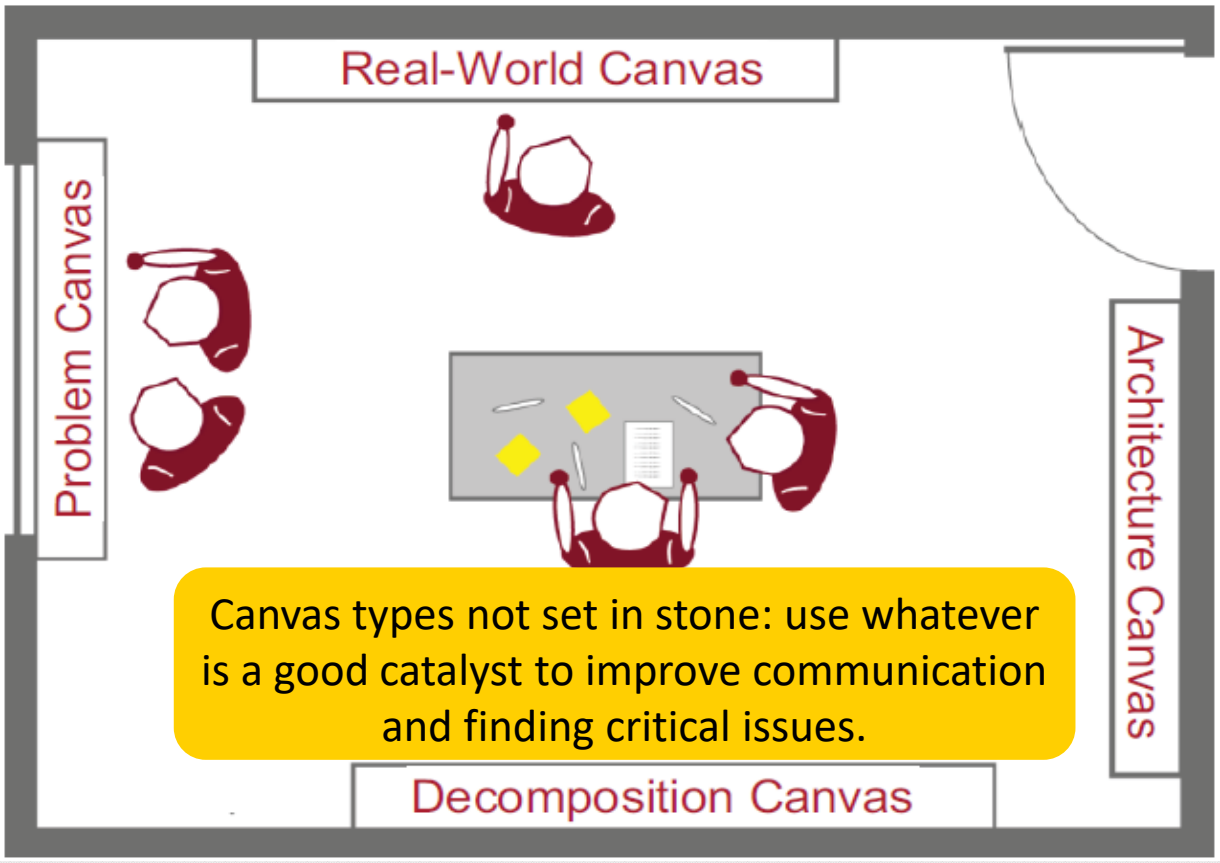

- Not necessary sequential flow, but iterative refinement of canvas contents.

Problem Canvas
Goal and scope of research question about the domain

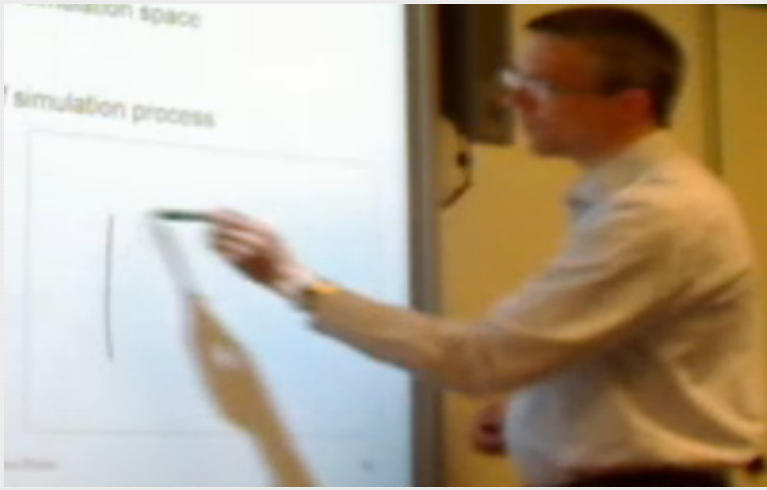

- Domain experts collect:
  - Research question,
  - Boundary conditions,
  - Assumptions,
  - Abstractions,
  - Quality requirements.
- Example: Heat dissipation problem
  - **Research Question**: “What will the temperature in a room after running an air conditioner on one side and a heater on the other for several hours?”
  - **Boundary conditions**: Starting temperature, A/C and heater setting.
  - **Abstractions**: Consider heat transfer by air flow / convection only, not by radiation.
  - **Assumptions**: No moving objects in the room, no windows/doors.
  - **Quality requirements**: Temperature must be determined as double precision float.

Real-World Canvas
Description of the pertinent aspects of the domain

- Domain experts sketch **static properties** of the simulation space:
  - Physical laws,
  - Spatial setup.
- Domain experts sketch **dynamic properties** of simulation process:
  - Forces,
  - Events.

- Example: Heat dissipation problem
  - Room geometry, heater & A/C location,
  - Working of convection forces/air flows,
  - Appropriate formulae for physical laws.

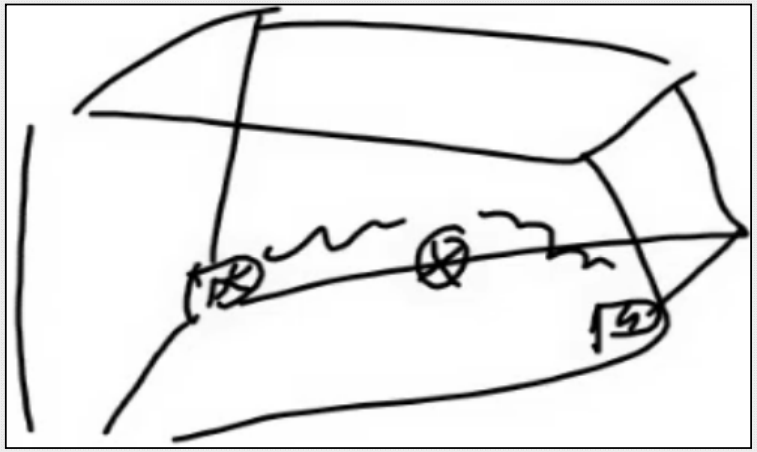

Decomposition Canvas
Breakdown of scientific model into parallelizable units

- HPC experts sketch parallel model reflecting real-world model:
  - **Static HPC aspects**: Domain decomposition, data structures.
  - **Dynamic HPC aspects**: Communication patterns, halos/ghosts, adaptive mesh refinements, iterative numerical methods.

- Example: Heat dissipation problem
  - Adaptive mesh refinement,
  - Halo creation & communication strategy,
  - Cartesian communicator,
  - Iterative method to solve formula of physical heat transfer law.

Architecture Canvas
Implementation of simulation on suitable HPC technology

- HPC experts sketch mapping of parallel model to actual cluster HW:
  - **Static aspects**: Cluster architecture, memory model, tools.
  - **Dynamic aspects**: Communication protocols, I/O operations, checkpointing.

- Example: Heat dissipation problem
  - 1024 CPU cores,
  - Hybrid MPI/OpenMP,
  - Jacobi solver,
  - Parallel I/O,
  - Checkpoints every 2.000 iterations.

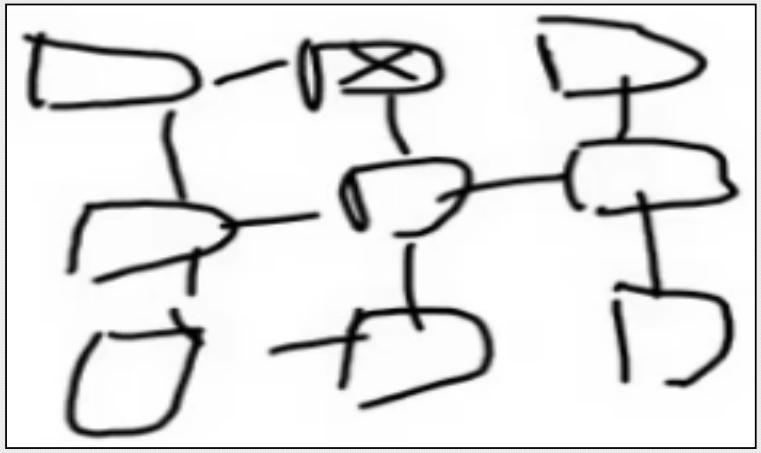

Summary

- Interaction Room **facilitates collaboration of**
  - **domain experts (engineers/ scientists),**
  - **HPC experts,**
  - **AI/ML experts,**
  - **Computer science/software engineering experts**.
- Canvas topic types not set in stone. (Add what is a good catalyst.)
- Not only a tool within CoE RAISE WP2,
  - but also outcome of CoE RAISE itself,
  - reaching beyond CoE RAISE.

Outlook

- To be developed during the CoE RAISE journey:
  - **Canvases for AI/ML**.
    - Including: combining HPC & AI, suitable annotations.

- Virtual (due to COVID) Interaction Room sessions with individual RAISE CoE WP2 partners:
  - Collaborative whiteboard software,
  - Caveat: lots of interaction lost in video meeting, e.g. body language,
  - Participants need to talk.
    - If no one talks, the best Interaction Room does not help – but canvases are a good catalyst.

drive. enable. innovate.

Book, Riedel, Neukirchen, Goetz:
Facilitating Collaboration in High-Performance Computing Projects with an Interaction Room.
4th ACM SIGPLAN International Workshop on Software Engineering for Parallel Systems (SEPS 2017)
https://doi.org/10.1145/3141865.3142467

RAISE
Center of Excellence

The CoE RAISE project receives funding from
the European Union's Horizon 2020 –
Research and Innovation Framework Programme
H2020-INFRAEDI-2019-1 under grant agreement no. 951733

Typical Pitfalls in HPC Simulation Science Projects

- Choosing appropriate solvers vs. reinventing the wheel
- Inefficient domain decomposition; load imbalance
- Dealing with differences between & unique strengths of individual architectures
- Dealing with different schedulers and their job scripts
- Debugging costs high amount of (possibly expensive) computing time
- Approximation of real world, insufficient validation data
- Integrating different physical models/processes with each other (multi-physics)
- Constant change of hardware, software, modus operandi
  - Constant need for porting; always an early adopter; changing code ownership

- Many of these revealed only in late (i.e. expensive to fix) project stages

Conceptual Levels in HPC Simulation Science Projects

RAISE Center of Excellence

- **Problem level:** Statement of research question / project goal and scope
  - Goal, context, scope: Research question, boundary conditions, assumptions, abstractions
  - Quality requirements: Accuracy, generalizability, performance
- **Scientific level:** Description of the pertinent aspects of the domain
  - Static aspects: Coordinates, variables, sources of influence, points of interest, physical laws
  - Dynamic aspects: Forces, interactions, events, timing, discontinuities
- **Distribution level:** Breakdown of the scientific model into parallelizable units
  - Static aspects: Domain decomposition, data structure, initial conditions
  - Dynamic aspects: Communication patterns, stencils, halos, ghosts, adaptive mesh refinements, iterative numerical methods
- **Technical level:** Implementation of distribution model on particular architecture
  - Static aspects: Cluster architecture, (parallel) file system, memory model, interconnect
  - Dynamic aspects: Communication protocols, I/O operations, available libraries, solvers