



APPLYING THE HPC SKILL TREE

Early experiences with course classification and design

December 14, 2021 | Marc-André Hermanns

This presentation contains work in progress and should be regarded as a basis for discussion.

- For users:
 - Orientation in knowledge levels of different HPC topics
 - Self-assessment of personal knowledge level (after course participation)
 - Independently-assessed of personal knowledge level (after exam)

- For users:
 - Orientation in knowledge levels of different HPC topics
 - Self-assessment of personal knowledge level (after course participation)
 - Independently-assessed of personal knowledge level (after exam)
- For HPC centers:
 - Better assessment of users' knowledge (→ “HPC Führerschein”)
 - Strategically improve users' knowledge on HPC topics
 - Towards “standardization” of HPC curricula

- Extending an existing MPI course
 - 2-day MPI introduction as part of PPCES @ RWTH Aachen University
 - Basis for HPC.NRW online MPI tutorial (in development)
 - Extend to broad coverage of MPI

- Extending an existing MPI course
 - 2-day MPI introduction as part of PPCES @ RWTH Aachen University
 - Basis for HPC.NRW online MPI tutorial (in development)
 - Extend to broad coverage of MPI
- Initial assessment:
 - Classify the skill level of course content
 - Which skills does the current course build up?

- Extending an existing MPI course
 - 2-day MPI introduction as part of PPCES @ RWTH Aachen University
 - Basis for HPC.NRW online MPI tutorial (in development)
 - Extend to broad coverage of MPI
- Initial assessment:
 - Classify the skill level of course content
 - Which skills does the current course build up?
 - Classify the skill level of requirements
 - Which skills does the participant need to follow the course?

- Extending an existing MPI course
 - 2-day MPI introduction as part of PPCES @ RWTH Aachen University
 - Basis for HPC.NRW online MPI tutorial (in development)
 - Extend to broad coverage of MPI
- Initial assessment:
 - Classify the skill level of course content
 - Which skills does the current course build up?
 - Classify the skill level of requirements
 - Which skills does the participant need to follow the course?
 - Group content according to skill level
 - Which content is needed for “basic”, “intermediate”, and “expert” level?

- SD – Software Development
 - SD 1 – Programming Concepts for HPC
 - SD 1.2 – Parallel Programming
 - SD 1.2.3 – Programming Message Passing Systems
- K – HPC Knowledge
 - K 3 – Program Parallelization
 - K 3.1 – Level of Parallelization
 - K 3.2 – Parallelization Overheads
 - K 3.3 – Domain Decomposition

- SD – Software Development
 - SD 1 – Programming Concepts for HPC
 - SD 1.2 – Parallel Programming
 - SD 1.2.3 – Programming Message Passing Systems
- K – HPC Knowledge
 - K 3 – Program Parallelization
 - K 3.1 – Level of Parallelization
 - K 3.2 – Parallelization Overheads
 - K 3.3 – Domain Decomposition

- SD – Software Development
 - SD 1 – Programming Concepts for HPC
 - SD 1.2 – Parallel Programming
 - SD 1.2.3 – Programming Message Passing Systems ←
- K – HPC Knowledge
 - K 3 – Program Parallelization
 - K 3.1 – Level of Parallelization
 - K 3.2 – Parallelization Overheads
 - K 3.3 – Domain Decomposition

- Blocking vs. Nonblocking

SD 1.2.3 PROGRAMMING MESSAGE PASSING SYSTEMS

Basic Knowledge

- Blocking vs. Nonblocking ✓
- Point-to-Point vs. Collective

- Blocking vs. Nonblocking ✓
- Point-to-Point vs. Collective ✓
- Detect Race Conditions & Deadlocks?

- Blocking vs. Nonblocking ✓
- Point-to-Point vs. Collective ✓
- ~~Detect~~ Understand Race Conditions & Deadlocks?

- Blocking vs. Nonblocking ✓
- Point-to-Point vs. Collective ✓
- ~~Detect~~ Understand Race Conditions & Deadlocks?
- Assess the impact of communication and synchronization on program performance

- Blocking vs. Nonblocking ✓
- Point-to-Point vs. Collective ✓
- ~~Detect~~ Understand Race Conditions & Deadlocks?
- ~~Assess~~ Understand the impact of communication and synchronization on program performance

- Blocking vs. Nonblocking ✓
 - Point-to-Point vs. Collective ✓
 - ~~Detect~~ Understand Race Conditions & Deadlocks?
 - ~~Assess~~ Understand the impact of communication and synchronization on program performance
- Detection and assessment are topics of “performance engineering” and “correctness checking”

SD 1.2.3 PROGRAMMING MESSAGE PASSING SYSTEMS

Intermediate Knowledge

- ~~Assess~~ Understand the impact of communication and synchronization on program performance

SD 1.2.3 PROGRAMMING MESSAGE PASSING SYSTEMS

Intermediate Knowledge

- ~~Assess~~ Understand the impact of communication and synchronization on program performance
- What else?
 - (see classification proposal)

- ~~Assess~~ Understand the impact of communication and synchronization on program performance
- Apply the concept of overlay networks

- ~~Assess~~ Understand the impact of communication and synchronization on program performance
- Apply the concept of overlay networks
 - What is meant here? Virtual Topologies? Why not earlier?

Classification of a large API

Marc-André Hermanns

- Semantic Terms
 - Local vs. nonlocal
 - Blocking vs. nonblocking
 - Collective vs. noncollective

- Semantic Terms [Basic]
 - Local vs. nonlocal
 - Blocking vs. nonblocking
 - Collective vs. noncollective

- Semantic Terms [Basic]
- Point-to-Point Communication
 - Blocking
 - Nonblocking
 - Persistent
 - Partitioned

- Semantic Terms [Basic]
- Point-to-Point Communication
 - Blocking [Basic]
 - Basic understanding of routine
 - Nonblocking [Basic]
 - Initial understanding of nonblocking semantics
 - Persistent [Intermediate]
 - Persistence as intermediate concept (→ tuning opportunities)
 - Partitioned [Expert]
 - Partitioned communication requires understanding of persistence

- Semantic Terms [Basic]
- Point-to-Point Communication
 - Blocking [Basic]
 - Nonblocking [Basic]
 - Persistent [Intermediate]
 - Partitioned [Expert]
- Datatypes

- Semantic Terms [Basic]
- Point-to-Point Communication
 - Blocking [Basic]
 - Nonblocking [Basic]
 - Persistent [Intermediate]
 - Partitioned [Expert]
- Datatypes [Basic]
 - Understanding of type map & type signature

- Collective Communication
 - Blocking
 - Nonblocking
 - Persistent
 - Neighborhood

- Collective Communication
 - Blocking [Basic]
 - Basic understanding of collective concept
 - Nonblocking [Basic]
 - Extend the nonblocking concept to collective concept
 - Persistent
 - Neighborhood

- Collective Communication
 - Blocking [Basic]
 - Basic understanding of collective concept
 - Nonblocking [~~Basic~~ Intermediate]
 - Extend the nonblocking concept to collective concept
 - Time constraints may prohibit this topic to be in “Basic” level
 - Persistent
 - Neighborhood

- Collective Communication
 - Blocking [Basic]
 - Nonblocking [~~Basic~~ Intermediate]
 - Persistent [Expert]
 - Persistence as intermediate concept (→ tuning opportunities)
 - Neighborhood [Expert]
 - Relying on virtual topologies as intermediate concept

- Collective Communication
 - Blocking [Basic]
 - Nonblocking [~~Basic~~ Intermediate]
 - Persistent [Expert]
 - Neighborhood [Expert]
- Groups & Communicators
 - Intra-communicators
 - Inter-communicators

- Collective Communication
 - Blocking [Basic]
 - Nonblocking [~~Basic~~ Intermediate]
 - Persistent [Expert]
 - Neighborhood [Expert]
- Groups & Communicators
 - Intra-communicators [Basic]
 - Inter-communicators [Expert]

- Collective Communication
 - Blocking [Basic]
 - Nonblocking [~~Basic~~ Intermediate]
 - Persistent [Expert]
 - Neighborhood [Expert]
- Groups & Communicators
 - Intra-communicators [Basic]
 - Inter-communicators [Expert]
- Process Topologies

- Collective Communication
 - Blocking [Basic]
 - Nonblocking [~~Basic~~ Intermediate]
 - Persistent [Expert]
 - Neighborhood [Expert]
- Groups & Communicators
 - Intra-communicators [Basic]
 - Inter-communicators [Expert]
- Process Topologies [Intermediate]
 - Relying on virtual topologies as intermediate concept

- Collective Communication
 - Blocking [Basic]
 - Nonblocking [~~Basic~~ Intermediate]
 - Persistent [Expert]
 - Neighborhood [Expert]
- Groups & Communicators
 - Intra-communicators [Basic]
 - Inter-communicators [Expert]
- Process Topologies [Intermediate]
- Info Object

- Collective Communication
 - Blocking [Basic]
 - Nonblocking [~~Basic~~ Intermediate]
 - Persistent [Expert]
 - Neighborhood [Expert]
- Groups & Communicators
 - Intra-communicators [Basic]
 - Inter-communicators [Expert]
- Process Topologies [Intermediate]
- Info Object [Intermediate]
 - User-Library-Interaction as intermediate concept

A first proposal

- Process Initialization
 - World Model
 - Sessions Model

- Process Initialization
 - World Model [Basic]
 - Single-threaded initialization [Basic]
 - Hybrid initialization [Basic/Intermediate?]
 - Sessions Model [Intermediate]
 - Tune process startup
 - Isolate libraries

A first proposal

- Process Initialization
 - World Model [Basic]
 - Sessions Model [Intermediate]
- One-sided Communication

- Process Initialization
 - World Model [Basic]
 - Sessions Model [Intermediate]
- One-sided Communication [Intermediate]
 - Relies on understanding nonblocking communication
 - Complex semantics

A first proposal

- Process Initialization
 - World Model [Basic]
 - Sessions Model [Intermediate]
- One-sided Communication [Intermediate]
- Generalized Requests

A first proposal

- Process Initialization
 - World Model [Basic]
 - Sessions Model [Intermediate]
- One-sided Communication [Intermediate]
- Generalized Requests [Expert?]
 - Chicken-egg problem? (mostly unknown ↔ mostly unused)

A first proposal

- Process Initialization
 - World Model [Basic]
 - Sessions Model [Intermediate]
- One-sided Communication [Intermediate]
- Generalized Requests [Expert?]
- I/O

A first proposal

- Process Initialization
 - World Model [Basic]
 - Sessions Model [Intermediate]
- One-sided Communication [Intermediate]
- Generalized Requests [Expert?]
- I/O [Intermediate]
 - Important scalability concept

A first proposal

- Process Initialization
 - World Model [Basic]
 - Sessions Model [Intermediate]
- One-sided Communication [Intermediate]
- Generalized Requests [Expert?]
- I/O [Intermediate]
- Tool Support

- Process Initialization
 - World Model [Basic]
 - Sessions Model [Intermediate]
- One-sided Communication [Intermediate]
- Generalized Requests [Expert?]
- I/O [Intermediate]
- Tool Support [Expert]
 - Tool experts should have an expert understanding of MPI

A proposed curriculum

Marc-André Hermanns

- Semantiv terms
- Process Initialization (World Model)
- Blocking & nonblocking point-to-point communication
- Datatypes
- Blocking collective communication
- Nonblocking collective communication
- Groups & Communicators (Intra-communicators)

- Process Initialization (Sessions Model)
- Persistent point-to-point communication
- Nonblocking collective communication
- One-sided communication
- I/O
- Process Topologies

- Process Initialization (Sessions Model)
- Persistent point-to-point communication
- Nonblocking collective communication
- One-sided communication
- I/O
- Process Topologies

- Increased understanding of performance implications
 - Require basic cost awareness (PE 1-B)
 - Require basic understanding of performance profiling (PE 2.2-B)
- Best practices on communication design

- Partitioned point-to-point communication
- Persistent collective communication
- Neighborhood collective communication
- Groups & Communicators (Intra-communicators)
- Generalized requests
- Tool Support

- Partitioned point-to-point communication
- Persistent collective communication
- Neighborhood collective communication
- Groups & Communicators (Intra-communicators)
- Generalized requests
- Tool Support

- Further detailed understanding of performance implications
 - Require intermediate knowledge level of related skills?
 - Which skills specifically?

HOW TO HANDLE RELATED SKILLS

A first proposal

- Domain Decomposition (K 3.3)
- Load Balancing (SD 1.2.4)
- Performance Engineering (PE)
- I/O Programming (SD 1.2.5)

- Domain Decomposition (K 3.3)
- Load Balancing (SD 1.2.4)
- Performance Engineering (PE)
- I/O Programming (SD 1.2.5)

- Not (direct) part of SD 1.2.3 course
 - Use in exercises?

A first proposal

- Domain Decomposition (K 3.3)
- Load Balancing (SD 1.2.4)
- Performance Engineering (PE)
- I/O Programming (SD 1.2.5)

- Not (direct) part of SD 1.2.3 course
 - Use in exercises?
- Requirements for specific levels of SD 1.2.3
 - K 3.3-B should be a requirement for SD 1.2.3-I/E
 - K 3.3-B is a good motivation for virtual topologies
 - ...

- Requirements for specific levels of SD 1.2.3
 - ...
 - SD 1.2.3-B or SD 1.2.2-B requirement for K 3.3-I
 - At least one parallelization paradigm should be known

- Requirements for specific levels of SD 1.2.3
 - ...
 - SD 1.2.3-B or SD 1.2.2-B requirement for K 3.3-I
 - At least one parallelization paradigm should be known
 - SD 1.2.5-B should be a requirement for SD 1.2.3-I
 - Information to fundamentals of filesystems are needed to motivate MPI I/O

- Requirements for specific levels of SD 1.2.3
 - ...
 - SD 1.2.3-B or SD 1.2.2-B requirement for K 3.3-I
 - At least one parallelization paradigm should be known
 - SD 1.2.5-B should be a requirement for SD 1.2.3-I
 - Information to fundamentals of filesystems are needed to motivate MPI I/O
- Focus on functionality first and performance later?
 - Enable participants to gain experience with the interface
 - Participants may get into bad practice without proper guidance?

- Skill tree is still a work in progress (e.g., SD 1.2.3)
- Skill definition needs to be precise enough to full coverage in a block
 - A “course” with a skill label should cover the full skill
 - Participants can assume full coverage for the exam

- Skill tree is still a work in progress (e.g., SD 1.2.3)
- Skill definition needs to be precise enough to full coverage in a block
 - A “course” with a skill label should cover the full skill
 - Participants can assume full coverage for the exam
- Large interfaces not trivial to break down (due to cross-cutting concepts)
 - If non-blocking point-to-point and blocking collective communication is covered, why not also cover nonblocking collective communication?
 - Time constraint

- Skill tree is still a work in progress (e.g., SD 1.2.3)
- Skill definition needs to be precise enough to full coverage in a block
 - A “course” with a skill label should cover the full skill
 - Participants can assume full coverage for the exam
- Large interfaces not trivial to break down (due to cross-cutting concepts)
 - If non-blocking point-to-point and blocking collective communication is covered, why not also cover nonblocking collective communication?
 - Time constraint
 - Unwanted splitting of similar concepts

- Skill tree is still a work in progress (e.g., SD 1.2.3)
- Skill definition needs to be precise enough to full coverage in a block
 - A “course” with a skill label should cover the full skill
 - Participants can assume full coverage for the exam
- Large interfaces not trivial to break down (due to cross-cutting concepts)
 - If non-blocking point-to-point and blocking collective communication is covered, why not also cover nonblocking collective communication?
 - Time constraint
 - Unwanted splitting of similar concepts
 - Do we need further subdivisioning: SD 1.2.3.X?
 - Where to place MPI I/O?

- Complex relationships among skill tree nodes
 - How to ensure requirements for a course are known and met?
 - How to encode optional knowledge?
 - How to build paths from basic to expert level?

- Complex relationships among skill tree nodes
 - How to ensure requirements for a course are known and met?
 - How to encode optional knowledge?
 - How to build paths from basic to expert level?
- Do we need working groups to generate specific progress in different parts of the skill tree?

Thank you.