



Seminar *Performance* Portable Programming of HPC Applications – Kick-off

Erwin Laure

Max Planck Computing and Data Facility & Chair of Computer Architecture and Parallel Systems Faculty of Informatics Technical University of Munich

with Material from Alexis Engelke, David Hildenbrand, Michael Petter, and Josef Weidendorfer

26.04.2022





Organization



- Kick-off meeting (today)
- Literature research + derive article structure
- Discuss structure with advisor
- Write draft paper deadline 2022-05-31



Organization



- Kick-off meeting (today)
- Literature research + derive article structure
- Discuss structure with advisor
- Write draft paper deadline 2022-05-31
- Peer-review two other papers deadline 2022-06-17



Organization



- Kick-off meeting (today)
- Literature research + derive article structure
- Discuss structure with advisor
- Write draft paper deadline 2022-05-31
- Peer-review two other papers deadline 2022-06-17
- Incorporate feedback from peers and advisor
- ► Final submission + presentation July 5, 2022



Background





FROM SIMPLE VON NEUMANN ARCHITECTURES TO MODERN HPC SYSTEMS







AND THEN WE ALSO ADD ACCELERATORS (GPUS)



Background

TITLE | DD.MM.YYYY 3



Background





AND USE MANY, REALLY MANY OF THESE NODES

- Summit Supercomputer @ ORNL:
 - 4.608 nodes
 - · 200 PF performance
 - · 13 MW power consumption



MAX PLANCK COMPUTING AND DATA FACILITY | NAME



Background





(SOME) CHALLENGES IN PROGRAMMING THESE SYSTEMS

- Level of parallelism
 - O(10⁹) FPUs
- Hardware heterogeneity
 - · CPUs, GPUs, other
 - · HBM, SSD, object store
- Programming/Performance Portability
- Novel numerical/methodological approaches

TITLE | DD.MM.YYYY









THE GOOD OLD TIMES

- Programms written in Fortran (or C/C++)
- MPI (Message Passing Interface) for moving data across distributed memory
- OpenMP for expressing parallelism on shared memory

MAX PLANCK COMPUTING AND DATA FACILITY | NAME

TITLE | DD.MM.YYYY



Background



CRAY

EPIGRAMHS







Goals & Topics



- Investigate different programming approaches particularly wrt portability and performance portability
 - Focus on heterogeneous architectures (GPUs, shared/distributed memory)
- Topics
 - Kokkos
 - RAJA
 - Alpaka
 - HiP
 - SYCL
 - MPI/CUDA (synchronization semantics in CUDA-integrated MPI communication)
 - ADIOS
 - Just-in-time compilation (JAX, Numba, Julia ...)



Goals & Topics



- Investigate different programming approaches particularly wrt portability and performance portability
 - Focus on heterogeneous architectures (GPUs, shared/distributed memory)
- Topics
 - Kokkos
 - RAJA
 - Alpaka
 - HiP
 - SYCL
 - MPI/CUDA (synchronization semantics in CUDA-integrated MPI communication)
 - ADIOS
 - Just-in-time compilation (JAX, Numba, Julia ...)
 - Bring your own topic





- Find 3-4 suitable publications & suitable websites
- Describe briefly how the mentioned approach works and helps with portability
- Does the mentioned approach enable a code to be portable across architectures?
- How good is the performance of a code portable?
- Discuss advantages and drawbacks





Literature and sources

- Finding literature and citable sources/references
- Writing a seminar paper
 - Structure, style, citing
- Presentation techniques
 - Structure, slide design, presentation style



Citable Literature

ТШ

Good to use

- Books, book chapters
- Papers (conf./journal)
- Published articles
- Manuals



Citable Literature

ТШ

Good to use

- Books, book chapters
- Papers (conf./journal)
- Published articles
- Manuals
- Websites with identifiable author

(cite with URL+access date)



Citable Literature

Good to use

- Books, book chapters
- Papers (conf./journal)
- Published articles
- Manuals
- Websites with identifiable author

(cite with URL+access date)

Try to avoid

- Wikipedia
- Facebook, etc.
- Advertisements
- Lecture slides
- Source code









- Starting points: IEEExplore, ACM DL, Google Scholar, arxiv.org, ...
 - Select appropriate keywords
 - Many papers/books accessible freely via the library





- Starting points: IEEExplore, ACM DL, Google Scholar, arxiv.org, ...
 - Select appropriate keywords
 - Many papers/books accessible freely via the library
- Other starting point: your advisor



- Starting points: IEEExplore, ACM DL, Google Scholar, arxiv.org, ...
 - Select appropriate keywords
 - Many papers/books accessible freely via the library
- Other starting point: your advisor
- Graph algorithms
 - Publications of the same author(s)
 - Publications at the same venue
 - Cites . . . (listed references)
 - Cited by ...









- Abstract: Brief summary of area, problem, approach, result
- Introduction: introduce area, problem, key results, contributions, outline
- Background: if needed, describe prerequisites





- Abstract: Brief summary of area, problem, approach, result
- Introduction: introduce area, problem, key results, contributions, outline
- Background: if needed, describe prerequisites
- Main part (approach, evaluation, discussion, etc.)





- Abstract: Brief summary of area, problem, approach, result
- Introduction: introduce area, problem, key results, contributions, outline
- Background: if needed, describe prerequisites
- Main part (approach, evaluation, discussion, etc.)
- (In a paper: Related Work)
- Summary & outlook



Writing style





Writing style



► Factual, precise, focused

- Stay on topic, no story telling, ...
- Limit to important and necessary topics
- Don't omit necessary prerequisites





► Factual, precise, focused

- Stay on topic, no story telling, ...
- Limit to important and necessary topics
- Don't omit necessary prerequisites
- Avoid forward references
- Avoid I, prefer we (or passive voice)
- We only describes the authors, not the reader







All work that is not yours must be cited

- Clearly describe source
- But: no wrong/inaccurate attributions
- Citing styles:
 - Literal (direct) quote
 - indirect quote (rephrase)

 $\leftarrow \mathsf{strongly} \ \mathsf{preferred}$

 Exception: foundations can be assumed (generally first few Bachelor semesters)









The x86 architecture defines the register CR2 [1].

The x86 architecture defines the register CR2 c (intel2019man).





The x86 architecture defines the register CR2 [1].

The x86 architecture defines the register CR2. It can be used with the instruction MOV. [1] The x86 architecture defines the register CR2 c (intel2019man).

The x86 architecture defines the register CR2. It can be used with the instruction MOV.~\cite{intel2019man} (Absatz)





The x86 architecture defines the register CR2 [1].

The x86 architecture defines the register CR2. It can be used with the instruction MOV. [1]

Valgrind [1] is a tool for run-time instrumentation.

The x86 architecture defines the register CR2 c (intel2019man).

The x86 architecture defines the register CR2. It can be used with the instruction MOV.~\cite{intel2019man} (Absatz)

Valgrind~\cite{nethercote2007} is a tool for run-time instrumentation.





The x86 architecture defines the register CR2 [1].

The x86 architecture defines the register CR2. It can be used with the instruction MOV. [1]

Valgrind [1] is a tool for run-time instrumentation.

Other approaches [1,2,3] ...

The x86 architecture defines the register CR2 $^$ cite{intel2019man}.

The x86 architecture defines the register CR2. It can be used with the instruction MOV.~\cite{intel2019man} (Absatz)

Valgrind~\cite{nethercote2007} is a tool for run-time instrumentation.



Presentation: Content Selection





Presentation: Content Selection



Presentation for the audience!



Presentation: Content Selection



Presentation for the audience!

- What do you want the audience to take away? (Not: what can I talk about!)
- What are the key points?
- How much content fits into the time slot?







Motivation

- Why is the topic relevant?
- Background
 - Consider referencing information from previous talks
- Concept
- Evaluation
 - How good is the described concept?
- Conclusions and outlook





- Important: topics build upon each other! (avoid forward references)
- Only give important details
- Use good/helpful examples
- Critical discussion of the topic





ТШ

- Slides (Beamer)
 - For use during the talk
 - Good to prepare
 - Backup-Folien as preparation for questions
- Whiteboard, blackboard
 - Permanently needed information
 - Answering questions
- Hardware, demonstrators, etc.
- Check possibilities in advance



Before the Talk

- Prepare slides, etc.
- Do a dry-run
 - Always recommended
 - Helps with uncertainity and time estimation
- Prepare on-site
 - Laptop, Beamer, laser pointer, clock, etc.



Talking Style



- Speak freely
- Don't go too fast/slow
- Stay in contact with the audience
 - Eye contact, position, etc.
- Usually at least 1 minute per slide
- Stay in time limit
 - Optional slides can fill time
 - Regularly consult a watch

Stay calm





- One topic per slide
- Avoid text
 - \leq 8 lines
- Prefer graphics/illustrations



- One topic per slide
- Avoid text
 - \leq 8 lines
- Prefer graphics/illustrations
- No unused points
 - Cover everything on the slides in your talk





Title page

▶ Title, name, institution, date, location





Title page

▶ Title, name, institution, date, location

On every other slide: number and title



ТШ

Title page

► Title, name, institution, date, location

- On every other slide: number and title
- Conclusion
 - Alle important points on one slide



Slides: Colors



Black on white

Literature



Slides: Colors



- Black on white
- Black on white



Slides: Colors



- Black on white
- Black on white
- Sufficient contrast
- Use colors sparingly, but systematically
- Be careful with gradients
- No annoying backgrounds (wave textures, etc.)
- Animations only with sufficiently added value



- Double-check text for typos, etc.
- Use a readable, sans-serif font
- Prefer vector graphics (or images with a high resolution)
- Avoid screenshots/scans
- Important: citations
- Listings only with a sufficiently large value



Negative Example

Alignments in brief	
>108.4 rolpothe heigh 57 Petersian Charael Pothe ALIWER AGANT LANVILLARSY LAVER CAPGAQLIT TYRALWIN SPECTIVE OF OLIVIER OF CAPGAQLIT GPTSFGL/TAALATWFVGREQ	>TOT, C. meigenbild heigh-223 Detestion Channel IGDV MEDIFARELORS NAALAWAYNYN ECTMOL NGEVIN RAYANDILAU HIAADAYNK AN KOSODWAG YNKKTIATE PAAPAGLAADROILLAGI GELEFELYRCL APPLRILLING OFFICIAL ADROILLAGI GELEFELYRCL APPLRILLING DEFENSION REVENDAU WAAR VERTHYNG I GOLWAATPIGKY I GIAMILTGISALTLLIGTYNNMPORILY
Query= 1bl8_A mol:protein length:97 Potassium Channel Protein (97 letters)	
 >lorq_C mol:protein length:223 Potassium Channel Length = 223 Score = 58.5 bits (140), Expect = 4e-14 Identities = 26/72 (36%), Positives = 43/72 (59%) 	Scoring matrix Algorithm to optimize score
Query: 21 GSYIAVLAERGAPGACLITYPRALWWSVETATTVGYGDLYPVTLWGRCVAVVVMVAGITS 80 G++	
Query: 81 FGLVTAALATWF 92 L+ ++ F Sbjct: 207 LTLLIGTVSNMF 218	



Negative Example







Negative Example



\begin{frame}
\frametitle{Die Anti-Folie}
\frametitle{Die Anti-Folie}
\centering
\includegraphics[width=0.95\textwidth]{pictures/antifolie.jpg}
\caption{Werbe-Folie. Foto von Flickr-Benutzer niallkennedy
(https://www.flickr.com/photos/niallkennedy/58697220/sizes/l/)}
\label{fig:gliederung}
\end{figure}
\end{frame}

Abbildung: Screenshot of code with insufficient resultion









End of Presentation



Summary slide with main take-away points



End of Presentation



- Summary slide with main take-away points
- NO Questions slide!







- Bring your point to the audience written or spoken
- Good literature as starting point
- Logical structure for paper and presentation
- Presentation: good preparation is important







- Bring your point to the audience written or spoken
- Good literature as starting point
- Logical structure for paper and presentation
- Presentation: good preparation is important
- Chance to learn ご







- Paper 6-8 pages
- Presentation about 15 Minutes plus questions (5 minutes)







- Paper 6-8 pages
- Presentation about 15 Minutes plus questions (5 minutes)
- Grading
 - ► 40% paper
 - 40% presentation
 - 20% review
 - all needs to be positive
- start literature search now and get in contact with your tutor



Questions







Questions











- Paper 6-8 pages
- Presentation about 15 Minutes plus questions (5 minutes)
- Grading
 - ► 40% paper
 - 40% presentation
 - 20% review
 - all needs to be positive
- start literature search now and get in contact with your tutor