MAX PLANCK INSTITUTE
FOR DYNAMICS OF COMPLEX
TECHNICAL SYSTEMS
MAGDEBURG

25 1998 2023

CSC COMPUTATIONAL METHODS IN SYSTEMS AND CONTROL THEORY

Introduction to focus group topics

The f(A)bulous workshop on matrix functions and exponential integrators

Stéphane Gaudreault, Kathryn Lund, and Marcel Schweitzer

25 September 2023

Supported by:

DFG

Outline

1. Background
2. Knowledge transfer
3. High-performance and energy-aware computing
4. Benchmark problems and FAIR comparisons
5. Instructions for focus groups

Why $f(A)b$?

- From a mathematical point of view...
  - More challenges due to nonlinearity
  - Rich theories: complex analysis, functional analysis, PDEs, etc.
- From a computer science point of view...
  - Adaptation of tools and methods for $Ax = b$
  - Nontrivial implementations required for robust solvers
- From an applications point of view...
  - Lattice quantum chromodynamics
  - Evolutionary advection-diffusion-reaction equations
  - Computational fluid dynamics
  - Chemical master equation
  - Gaussian processes
  - Stiff matrix differential equations
  - Network analysis

Outline

1. Background
2. Knowledge transfer
3. High-performance and energy-aware computing
4. Benchmark problems and FAIR comparisons
5. Instructions for focus groups

One problem, many angles

- General (dense) matrix functions: $f(A)$
- Matrix functions times a vector: $f(A)\boldsymbol{b}$
- Exponential integrators: $\exp(tA)\boldsymbol{b}$
- Related problems:
  - $\operatorname{trace}(f(A))$
  - $\boldsymbol{b}^T f(A)\boldsymbol{b}$
  - $f\{A, B\}(C)$
    - Lyapunov, Sylvester, and Stein equations
    - Fréchet derivatives
    - $\kappa(f(A))$

https://commons.wikimedia.org/wiki/Category:Elephants

Key questions

- What do these problems have in common? What are the underlying "kernels"?
- What techniques have been tested for one problem, but not yet another?
- How do these problems differ significantly, i.e., where should we be careful not to over simplify?
- Which functions f aren't getting enough attention?
- What role could / should machine learning play here?
- How do we measure success / performance / effort / accuracy / efficiency / stability?
- How can we improve communication between diverse members of this community?

Outline

1. Background
2. Knowledge transfer
3. High-performance and energy-aware computing
4. Benchmark problems and FAIR comparisons
5. Instructions for focus groups

General challenges

MareNostrum-4 in Barcelona

https://commons.wikimedia.org/wiki/File:2017_BSC_Superordenador_MareNostrum-4_Barcelona-Supercomputing-Center.jpg

- Machine access
- Scaling
- Parallelization
- Fault tolerance
- Hardware limits– CPU speed, bandwidth, latency, RAM, harddrive storage
- Operating costs
- CO2 emissions

Key questions

- What strategies can transfer from linear systems? What doesn't transfer?
- What existing tools and libraries can be adapted, built upon, or extended?
- How to optimize the dense problem?
- What should computational models for $f(A)\boldsymbol{b}$ look like?
- What challenges do exponential integrators or other specific f face in particular?
- Is $f(A)\boldsymbol{b}$ ready for exascale?
- What role could / should machine learning play here?
- How do we measure success / performance / effort / accuracy / efficiency / stability?
- How can we improve communication between diverse members of this community?

Outline

1. Background
2. Knowledge transfer
3. High-performance and energy-aware computing
4. Benchmark problems and FAIR comparisons
5. Instructions for focus groups

Research data management

- FAIR principles: Findability, Accessibility, Interoperability, and Reproducibility
- Benchmark problem collections
- Algorithm comparison workflows
- Bare minimum: publish your code and data with an open license!

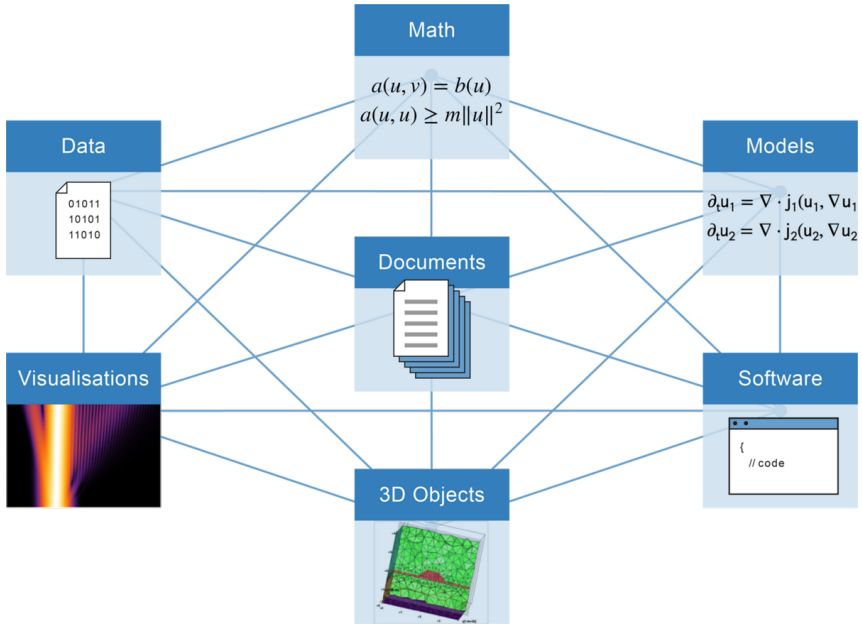

Borrowed from the MAthematical Research Data Initiative (MaRDI)

`www.mardi4nfdi.de`

Key questions

- What should benchmark problems look like for $f(A)\boldsymbol{b}$ and exponential integrators? Are there existing collections that we can adapt? What types of problems would help you in your own research?
- What problems arise when you try to share your data and code? What problems arise when you try to use someone else's data and code?
- How do we incentivize community contributions?
- When developing new algorithms, how do we establish what is state-of-the-art?
- How do we define "algorithm isotopes"?
- What role could / should machine learning play here?
- How can we improve communication between diverse members of this community?

Outline

1. Background
2. Knowledge transfer
3. High-performance and energy-aware computing
4. Benchmark problems and FAIR comparisons
5. Instructions for focus groups

Choose your topic...

- ...find your group leader[a]
  - Marcel: Knowledge transfer
  - Stéphane: HPC
  - Kathryn: FAIR Benchmarking
- ...find somewhere nice to work
  - Prigogine (here)
  - Wiener (next door)
  - Technikum (3rd floor with a foosball table)
- ...appoint 1-2 people to take notes
- ...grab some sugar and caffeine
- ...discuss!

[a] Try to evenly distribute!

There's no homework, but...

- Plan to present your results tomorrow afternoon
  - 30 min per group
  - Any format welcome:
    - Flipchart
    - Typed notes
    - Powerpoint/Beamer
    - Chalkboard
    - Interpretive dance
- Be creative, controversial, counterintuitive... but courteous!
- Extra motivation: submissions to ETNA special issue due 21 January 2024

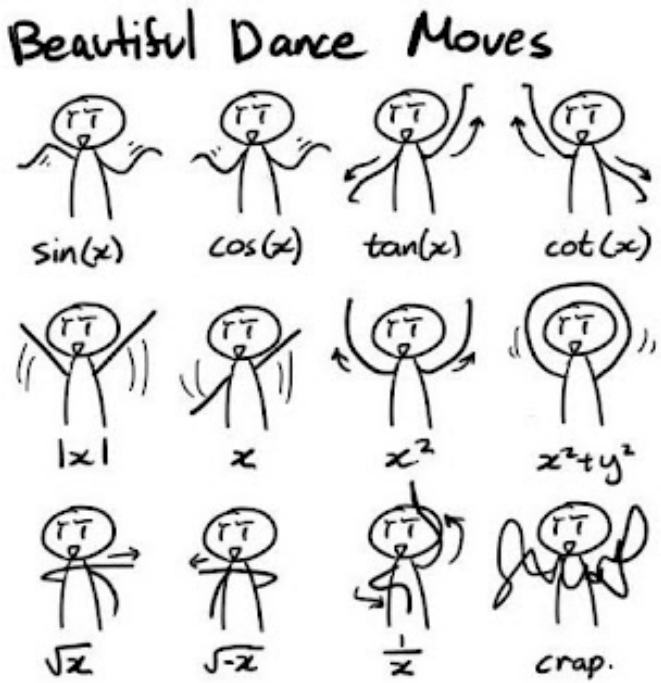

Where to find these slides

https://indico3.mpi-magdeburg.mpg.de/e/fabulous2023