

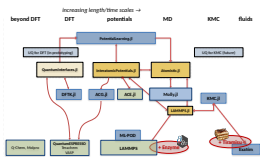
Fostering interdisciplinary research by composable **julia** software

Michael F. Herbst

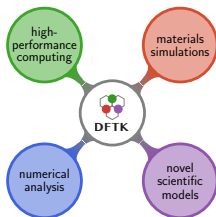
Mathematics for Materials Modelling (matmat.org), EPFL

20 February 2024

Slides: https://michael-herbst.com/talks/2024.02.20_elstruct_code_workshop.pdf



Real-world multi-physics
software stack for materials modelling



Stress =

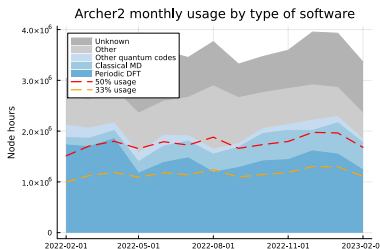
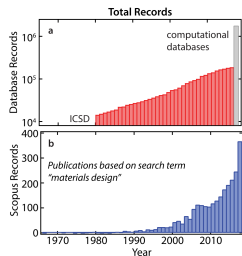
$$\frac{1}{\det(\mathbf{L})} \left. \frac{\partial E[P_*, (\mathbf{I} + \mathbf{M}) \mathbf{L}]}{\partial \mathbf{M}} \right|_{\mathbf{M}=0}$$

```
# Run SCF, get P*
scfres = self_consistent_field(basis)
L = basis.model.lattice
stress = 1/det(L) * gradient(
M -> recompute_energy(
    scfres, (I + M) * L),
zero(L)
)
```

julia vision: Math \equiv code

Tackling 21st century challenges

- 21st century challenges:
 - Renewable energy, green chemistry, health care ...
- Current solutions limited by properties of available materials
 - ⇒ Innovation driven by **discovering new materials**
- Crucial tool: **Computational materials discovery**
 - Systematic simulations on $\simeq 10^4 - 10^6$ compounds
 - Complemented by data-driven approaches
 - **Noteworthy share** of world's supercomputing resources



Tackling 21st century challenges

- 21st century challenges:
 - Renewable energy, green chemistry, health care ...
- Current solutions limited by properties of available materials
 - ⇒ Innovation driven by **discovering new materials**
- Crucial tool: **Computational materials discovery**
 - Systematic simulations on $\simeq 10^4 - 10^6$ compounds
 - Complemented by data-driven approaches
 - **Noteworthy share** of world's supercomputing resources
- Multi-disciplinary effort: **Software** takes a key role
 - E.g. growing list of data / workflow management tools
 - Challenges of combining efforts & integrating communities



AFLOW
Automates FLOW for Materials Discovery



AiiDA



A focus on robust materials simulations

- Goal in ~~Mat~~Mat group:

- Obtain **reliable & efficient** simulations
- Develop and employ **mathematical analysis** of error
- Transform **empirical wisdom** to built-in **convergence guarantees**

⇒ Understand **where and how** to spend efforts best

- Practical error indicators:

- Automatic & robust verification
- Multi-fidelity statistical surrogates
- Active learning of missing physics

- Leverage inexactness:

- Error balancing: Optimal adaptive parameter selection
- Adaptive tolerances & selective precision

⇒ Multidisciplinary expertise required

A focus on robust materials simulations

- Goal in ~~Mat~~Mat group:

- Obtain **reliable & efficient** simulations
- Develop and employ **mathematical analysis** of error
- Transform **empirical wisdom** to built-in **convergence guarantees**

⇒ Understand **where and how** to spend efforts best

- **Practical error indicators:**

- Automatic & robust verification
- Multi-fidelity statistical surrogates
- Active learning of missing physics

- **Leverage inexactness:**

- Error balancing: Optimal adaptive parameter selection
- Adaptive tolerances & selective precision

⇒ Multidisciplinary expertise required

(Exaggerative) state of codes in this field

Mathematical research

- **Goal:** Numerical experiments
- **Scope:** Reduced models
- High-level **language:**
Matlab, python, ...
- **Lifetime:** 1 paper
- **Size:** < 1k lines
- Does not care about performance

Application research

- **Goal:** Modelling physics
- **Scope:** All relevant systems
- Mix of **languages:**
C, FORTRAN, python, ...
- **Lifetime:** 100 manyears
- **Size:** 100k – 1M lines
- Obligated to write performant code

- Working with these codes requires different skillsets
⇒ **Orthogonal** developer & user **communities**
- Obstacle for knowledge transfer:
 - Mathematical methods **never tried in practical setting**
(and may well not work well in the real world)
 - **Some issues cannot be studied** with mathematical codes
(and mathematicians may never get to know of them)
- What about emerging hardware, accelerators, performance?
 - Should be the regime of Computer Science (yet another community)

Difficulties of interdisciplinary research

- Community conventions (e.g. publication culture)
- Language barriers and context-sensitive terms
- Speed of research (development of model vs. its analysis)
- A social problem . . .
 - (Communication, convention, compromises, . . .)
- . . . that is **cemented in software**:
 - **Priorities differ** ⇒ What is considered “a good code” differs
 - Insurmountable obstacles to integrate codes
 - Collaborations can stop before they begin . . .
- **Hypothesis:** **People compose if software composes**



high-performance computing

materials simulations



DFTK

numerical analysis

novel scientific models



- **Julia**-based DFT code in 7500 lines
- Cross-community: Mathematical research & applications
- Allows restriction to **relevant model problems**,
- **and scale-up** to application regime (1000 electrons)

● Integrated with high-throughput:  

● Lessons learned:

- Software integration is **hard work**
 - **Unexpected catalytic effects** from integration discussions
 - Parties understand their role, change of viewpoint
- ⇒ As software composes, communities compose




● **Central:** How can we **lower the barrier to integrate?**

We already want a lot from good software . . .


- Integration across communities (users, developers, scientists)
- Maintainability
- Reproducibility
- Documentation / Accessibility
- Portability (future technologies & hardware)
- Performance
- . . .






- Can we the get the best in each category?
 - Probably not . . .

- To **maximise integration**: Where should we compromise?

- 1 Composability aspects of  **julia**
- 2  **DFTK** and related  **julia** efforts

Separating the what from the how

- Why is this separation so important ...
 - ... for composable software?
 - ... for multidisciplinary research?
- Consider the **goal**: Modelling a physical system
- **Traditionally** users code in detail **how** the computation should proceed (Imperative programming)
 - How = architecture
 - How = algorithm
 - How = memory layout
 - How = discretisation
 - ...
- But all this has **nothing to do with physics!**
- Can the **how** be abstracted away?
 - such that CS / Math can deal with it *independently*
- Let's see some  **julia** developments

Accelerators	Shared Mem	Distributed
 CUDA.jl	 AMDGPU.jl	 OneAPI.jl
JuliaGPU/ GPUArrays.jl		 JuliaParallel/ Dagger.jl
Reusable array functionality for Julia's various GPU backends. https://github.com • JuliaGPU • KernelAbstractions		
KernelAbstractions.jl - Heterogeneous programming in Julia Heterogeneous programming in Julia. Contribute to JuliaGPU/KernelAbstractions.jl development by creating an ... JuliaGPU / KernelAbstractions.jl Public.		JuliaParallel/MPI.jl MPI wrappers for Julia
Julia Folds/FLoops.jl Fast sequential, threaded, and distributed for-loops for Julia-fold for humans™ Announcing composable multi-threaded parallelism in Julia 23 July 2017 Jeff Bezanson (Julia Computing), Aravind Narayana (Julia Computing), Kimon Pappas (Intel)		JuliaParallel/ Dagger.jl A framework for out-of-core and parallel execution Standard Library / Distributed Computing
Base / Multi-Threading Multi-Threading		Base Threads.@threads – Macro
Julia Atomics Manifesto This proposal aims to define the memory model of Julia and to provide certain guarantees to the consumers of data races, both by default and through providing hints to allow the user to specify the level of guarantees required. This should allow native implementations to Julia to employ system primitives (like mutexes), integrate with native hardware capabilities, and allow to give generally repeatable behaviors without incurring significant performance cost. Additionally, it should be the general audience and get clear about the user to identify opportunities with respect to ensuring that an atomic-type field is accessed with proper care for synchronization.		 GitHub - vth-cscc/impl3d@beatzj · github.com

```
function power_method(A, x; niter=100)
    for i = 1:niter
        x = A * x
        x ./= norm(x)
    end
    x
end
```

$$A = \text{rand}(10, 10); A = A + A' + 10I; x = \text{rand}(10)$$

```
using LinearMaps, IterativeSolvers
itinv(A) = LinearMap(x -> cg(A, x), size(A)...)

using CUDA
power_method(itinv(CuArray(A)), CuArray(x))

using AMDGPU
power_method(itinv(ROCArray(A)), ROCArray(x))
```

Code reinterpretation & self-implementing features

```
using OrdinaryDiffEq, Plots

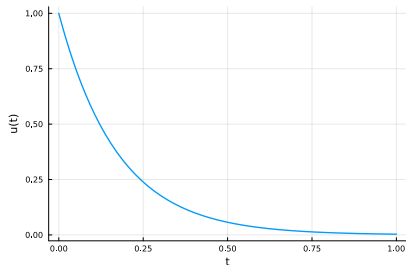
# Half-life of Carbon-14 is 5730 years.
c = 5.730

# Setup
u0 = 1.0
tspan = (0.0, 1.0)

# Define the problem
radioactivedecay(u, p, t) = -c*u

# Pass to solver
prob = ODEProblem(radioactivedecay, u0, tspan)
sol = solve(prob, Tsit5();
            reltol=1e-8, abstol=1e-8)

plot(sol.t, sol.u;
     ylabel="u(t)", xlabel="t", lw=2, legend=false)
```



Code reinterpretation & self-implementing features

```
using OrdinaryDiffEq, Measurements, Plots

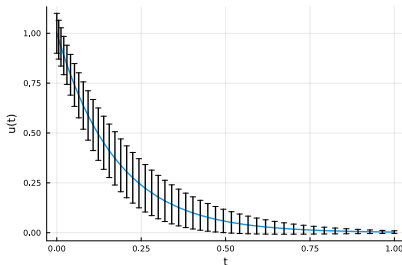
# Half-life of Carbon-14 is 5730 years.
c = 5.730 ± 2

# Setup
u0 = 1.0 ± 0.1
tspan = (0.0, 1.0)

# Define the problem
radioactivedecay(u, p, t) = -c*u



# Pass to solver
prob = ODEProblem(radioactivedecay, u0, tspan)
sol = solve(prob, Tsit5();
            reltol=1e-8, abstol=1e-8)

plot(sol.t, sol.u;
     ylabel="u(t)", xlabel="t", lw=2, legend=false)
```



- User says: I want to track measurement error
- Numerics adapts, plotting adapts
 - **No prior discussion** with/amongst package maintainers to “make this happen”
- `Measurement.jl` reinterprets floating-point operations
 - In some sense this feature “implemented itself”

Aside: package manager and binary dependencies

-  makes no compromises in reproducibility
 - Package environments automatically tracked in plain-text files
(Can be committed along code)
 - Includes **python & foreign-code binaries**
- Difficulty: Integration with HPC clusters:
 - E.g. making use of vendor-specific MPI / BLAS libraries
 -  solution: *trampoline* libraries
 - BLAS & MPI libraries can be switched *at runtime*

⇒ Sane defaults for laptops & flexibility
- Some pain points remain:
 - Default binaries cannot make full use of hardware
 - Automatic detection of vendor libraries




- Magic of Julia:
 - Painless generics and abstractions
 - Enables unusual code *reinterpretation*
(Algorithmic differentiation, symbolics, cross-platform compilation)

⇒ Separation of **what** and **how**:

- Hardware & architecture (Computer Science)
- Algorithms (Mathematics)
- Model building (Physics)
- Interactive scripting (Application scientists)

⇒ Cross-disciplinary **expertise can compose** in one code

- Modelling and algorithm code stays high-level
 - Appropriate **specialisations unlock performance**
 - We can **add them gradually** as needed (Iterative optimisation)

- 1 Composability aspects of  **julia**
- 2  **DFTK** and related  **julia** efforts



high-performance computing



materials simulations

DFTK

numerical analysis

novel scientific models



- **julia** code for **cross-disciplinary research**:
 - Allows restriction to **relevant model problems**,
 - **and scale-up** to application regime (1000 electrons)
 - **Sizeable feature set** in **7500 lines** of code
 - Including some unique features (Self-adapting algorithms)
 - Integrated with high-throughput:  
- **Fully composable** due to **julia** abstractions:
 - Arbitrary precision (32bit, >64bit, ...)
 - Algorithmic differentiation (AD)
 - HPC tools: GPU acceleration, MPI parallelisation
- Accessible **high-productivity** research framework:
 - Key contributions by undergrads (AD, GPU, Pseudos, ...)
 - Over 30 contributors in 5 years (Maths, physics, CS, ...)





DFTK design: Keeping code concise & accessible

Stress =

$$\frac{1}{\det(\mathbf{L})} \left. \frac{\partial E[P_*, (\mathbf{I} + \mathbf{M}) \mathbf{L}]}{\partial \mathbf{M}} \right|_{\mathbf{M}=0}$$

```
# Run SCF, get P*
scfres = self_consistent_field(basis)
L = basis.model.lattice
stress = 1/det(L) * gradient(
    M -> recompute_energy(
        scfres, (I + M) * L),
    zero(L)
)
```

- Stress computation (Definition vs.  code)¹
- Post-processing step \Rightarrow Not performance critical
- Comparison of implementation complexity:
 -  DFTK: 20 lines¹ (forward-mode algorithmic differentiation)
 - Quantum-Espresso: 1700 lines²
 - \simeq 10-week GSoC project

\Rightarrow No performance impact & accessible code

¹<https://github.com/JuliaMolSim/DFTK.jl/blob/master/src/postprocess/stresses.jl>

²<https://github.com/QEF/q-e/blob/develop/PW/src>

New features from generic code: Sensitivity analysis

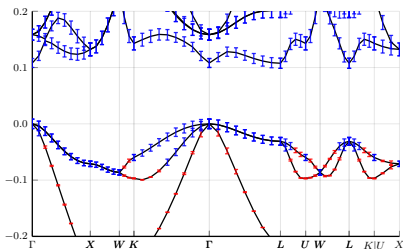
```
function dft_energy(a,  $\theta$ )  
    model = Model(a, PbeExchange( $\theta$ ), ...)  
    scf(model).energies.total  
end  
optimise_lattice( $\theta$ ) = optimise(a -> dft_energy(a,  $\theta$ ))  
  
sensitivities =  
    ForwardDiff.gradient(optimise_lattice,  $\theta$ )
```

$$a_* = \arg \min_a \mathcal{E}(a, \theta)$$

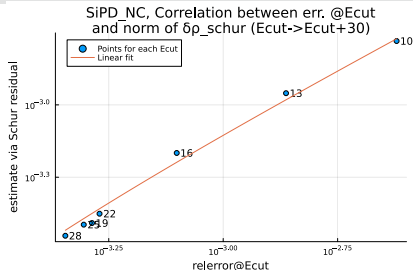
$$\text{sensitivities} = \frac{da_*}{d\theta}$$

- Arbitrary, user-desired derivatives in **one line of code**
 - Breaks “one PhD student per derivative” paradigm
 - ⇒ New properties/derivatives by **non-DFT experts!**
- Avoids combinatorial explosion
 - Unusual derivatives equally supported
- ⇒ Setting the scene for new approaches:
 - Sensitivity analysis & UQ
 - Combined analytical and statistical error estimation

Support of *a posteriori* error analysis



Band structure with guaranteed errors¹



Estimation of basis set error in ρ

- Momentum towards **error estimators** for DFT
 - Focus on basis set error (some also tackle floating-point, SCF convergence)
 - Estimate numerical error for modelled system (e.g. for **density** and **forces**²)
- Results promising, but many challenges & caveats remain
 - Crucial to play with simplifications / numerics etc.

⇒  **DFTK** is **major research tool** for this development¹⁻⁴

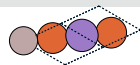
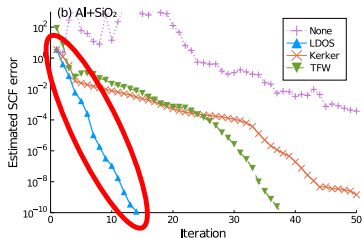
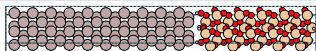
¹MFH, A. Levitt, E. Cancès. Faraday Discus. **223**, 227 (2020).

²E. Cancès, G. Dusson, G. Kemplin et. al. SIAM J. Sci. Comp., **44**, B1312 (2022).

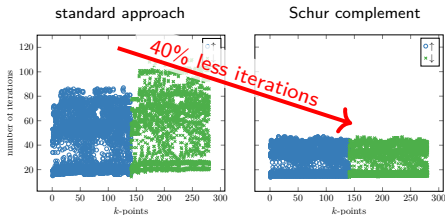
³E. Cancès, G. Kemplin, A. Levitt. J. Matrix Anal. Appl., **42**, 243 (2021).

⁴E. Cancès, G. Kemplin, A. Levitt. J. Sci. Comput., **98**, 25 (2024)

Robust & efficient algorithms



Fe₂MnAl Heusler alloy



- LDOS mixing for inhomogeneous systems¹ (surfaces, clusters, ...)
- ca. 50% less iterations
- Automatic & system-adapted selection of damping²

- First-principle properties of metals
- Schur-complement approach to perturbation theory²
- ca. 40% less iterations

⇒ Maths / physics collaboration:


Exchange of ideas between simplified & practical settings crucial

¹MFH, A. Levitt. J. Phys. Condens. Matter **33**, 085503 (2021).

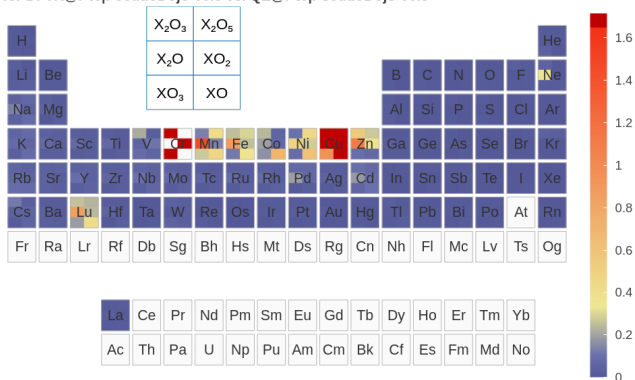
²MFH, A. Levitt. J. Comput. Phys. **459**, 111127 (2022).


³E. Cancès, MFH, G. Kemlin, et. al. Lett. Math. Phys. **113**, 21 (2023).

Integration with AiiDA


- Integration with  AiiDA high-throughput workflow manager
 - <https://github.com/aiidaplugins/aiida-dftk>
- Used in **automated verification tests**: ($r_{\text{cut}} = \infty$)

v for DFTK@PW|PseudoDojo-v0.5 vs. QE@PW|PseudoDojo-v0.5

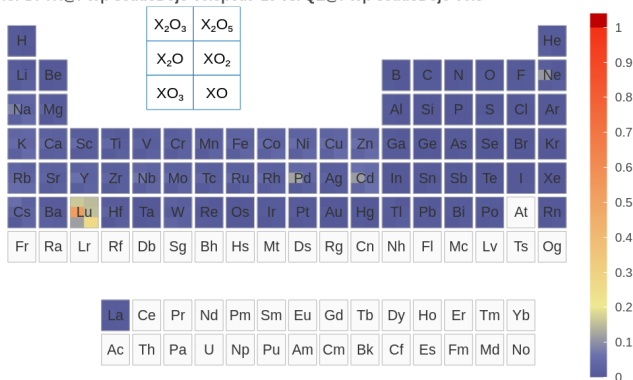



⇒ Excellent agreement Quantum-Espresso vs.  DFTK

Integration with AiiDA



- Integration with  AiiDA high-throughput workflow manager
 - <https://github.com/aiidaplugins/aiida-dftk>
- Used in **automated verification tests**: ($r_{\text{cut}} = 10\text{a.u.}$)


ϵ for DFTK@PW|PseudoDojo-v0.5|rcut=10 vs. QE@PW|PseudoDojo-v0.5



⇒ Excellent agreement Quantum-Espresso vs.  DFTK

Quick overview of materials codes

- Some  materials science codes:
 - `github.com/ACEsuit`: **Atomic Cluster Expansion** (ML potential)
 - `JuliaMolSim/Molly.jl`: **Molecular dynamics**
 - `qiaojunfeng/Wannier.jl`: **Wannierisation**
 - `JuliaMolSim/DFTK.jl`: **Density-functional theory**
- Community desire for **common interfaces** across  materials ecosystem
 - `AtomsBase.jl` & `AtomsCalculators.jl`
 - E.g. compatibility of **structural representations** *across codes*

⇒ Link *within*  ecosystem, but also *link to external codes* (e.g. ASE, Quantum Espresso, LAMMPS)
- ⇒ Generic & re-usable utility packages:
 - `AtomsIO.jl`: **File parsing**
 - `AtomsView.jl`: **Structure viewing**
 - Many just slim bindings to existing foreign-language codes ...
- Overview talk: **Julia for Materials Modelling**:
<https://michael-herbst.com/julia-for-materials> (youtube recording)

Summary

- People compose if software composes
 - Key ingredient: Separating **what** and **how**
 - ⇒ Better collaboration by separation of concern
- What makes **julia** codes so composable?
 - **Specialisation**: Performance & hardware specifics
 - **Abstraction**: Code becomes the math
 - **Multiple dispatch**: Repurpose existing code (e.g. AD)
- **julia**-based materials codes: Bridging communities
 - Multiple cross-disciplinary projects: Maths ↔ applications
 - Community emphasis on **composable interfaces**
(e.g. `AtomsBase.jl` & `AtomsCalculators.jl`)
 - Details: michael-herbst.com/julia-for-materials

Acknowledgements

- Alan Edelman (MIT)
- Valentin Churavy (MIT)
- Antoine Levitt (Université Paris-Saclay)
- All  DFTK contributors
- Joe Greener (Cambridge)
- Rachel Kurchin (CMU)
- Christoph Ortner (UBC)
- Spencer Wyatt (MIT)
- Pablo Zubieta (Chicago)

EPFL


Mt Mat




Summer of code





Questions?


 <https://matmat.org>

 mfherbst

 michael.herbst@epfl.ch

 https://michael-herbst.com/talks/2024.02.20_elstruct_code_workshop.pdf

 DFTK <https://dftk.org>

 <https://michael-herbst.com/julia-for-materials>