

Loki v0.1.5: Freely Programmable Source-to-Source Translation for IFS and beyond



Balthasar Reuter, Ahmad Nawab, Michael Staneker, Michael Lange

🏠 Research Department, ECMWF, Bonn (Germany) ✉️ {firstname}.{lastname}@ecmwf.int

Motivation

- **Performance portability** of Numerical Weather Prediction (NWP) codes across a broad range of HPC architectures, including **accelerators** (such as GPUs), from a **single code base**
- **Static code analysis/linting** of source code to aid development

Challenges

- Different **programming paradigms** and environments
- **Hardware-specific** optimisation (loop order, memory layout, ...)
- Handling a **large and complex Fortran code base**
- **Compatibility** with operational requirements and scientific changes

Methodology

- **Source-to-source (S2S) translation tool** to inspect/transform code:
 - **Static code analysis** using internal representation
 - **Build-time transformation** of source code using bespoke recipes

Open development on Github

Source code & bug tracker



📄/ecmwf-ifs/loki

Documentation



sites.ecmwf.int/docs/loki

Jupyter Notebook Tutorials



📄/ecmwf-ifs/loki/tree/main/example

Loki: overview and internal representation

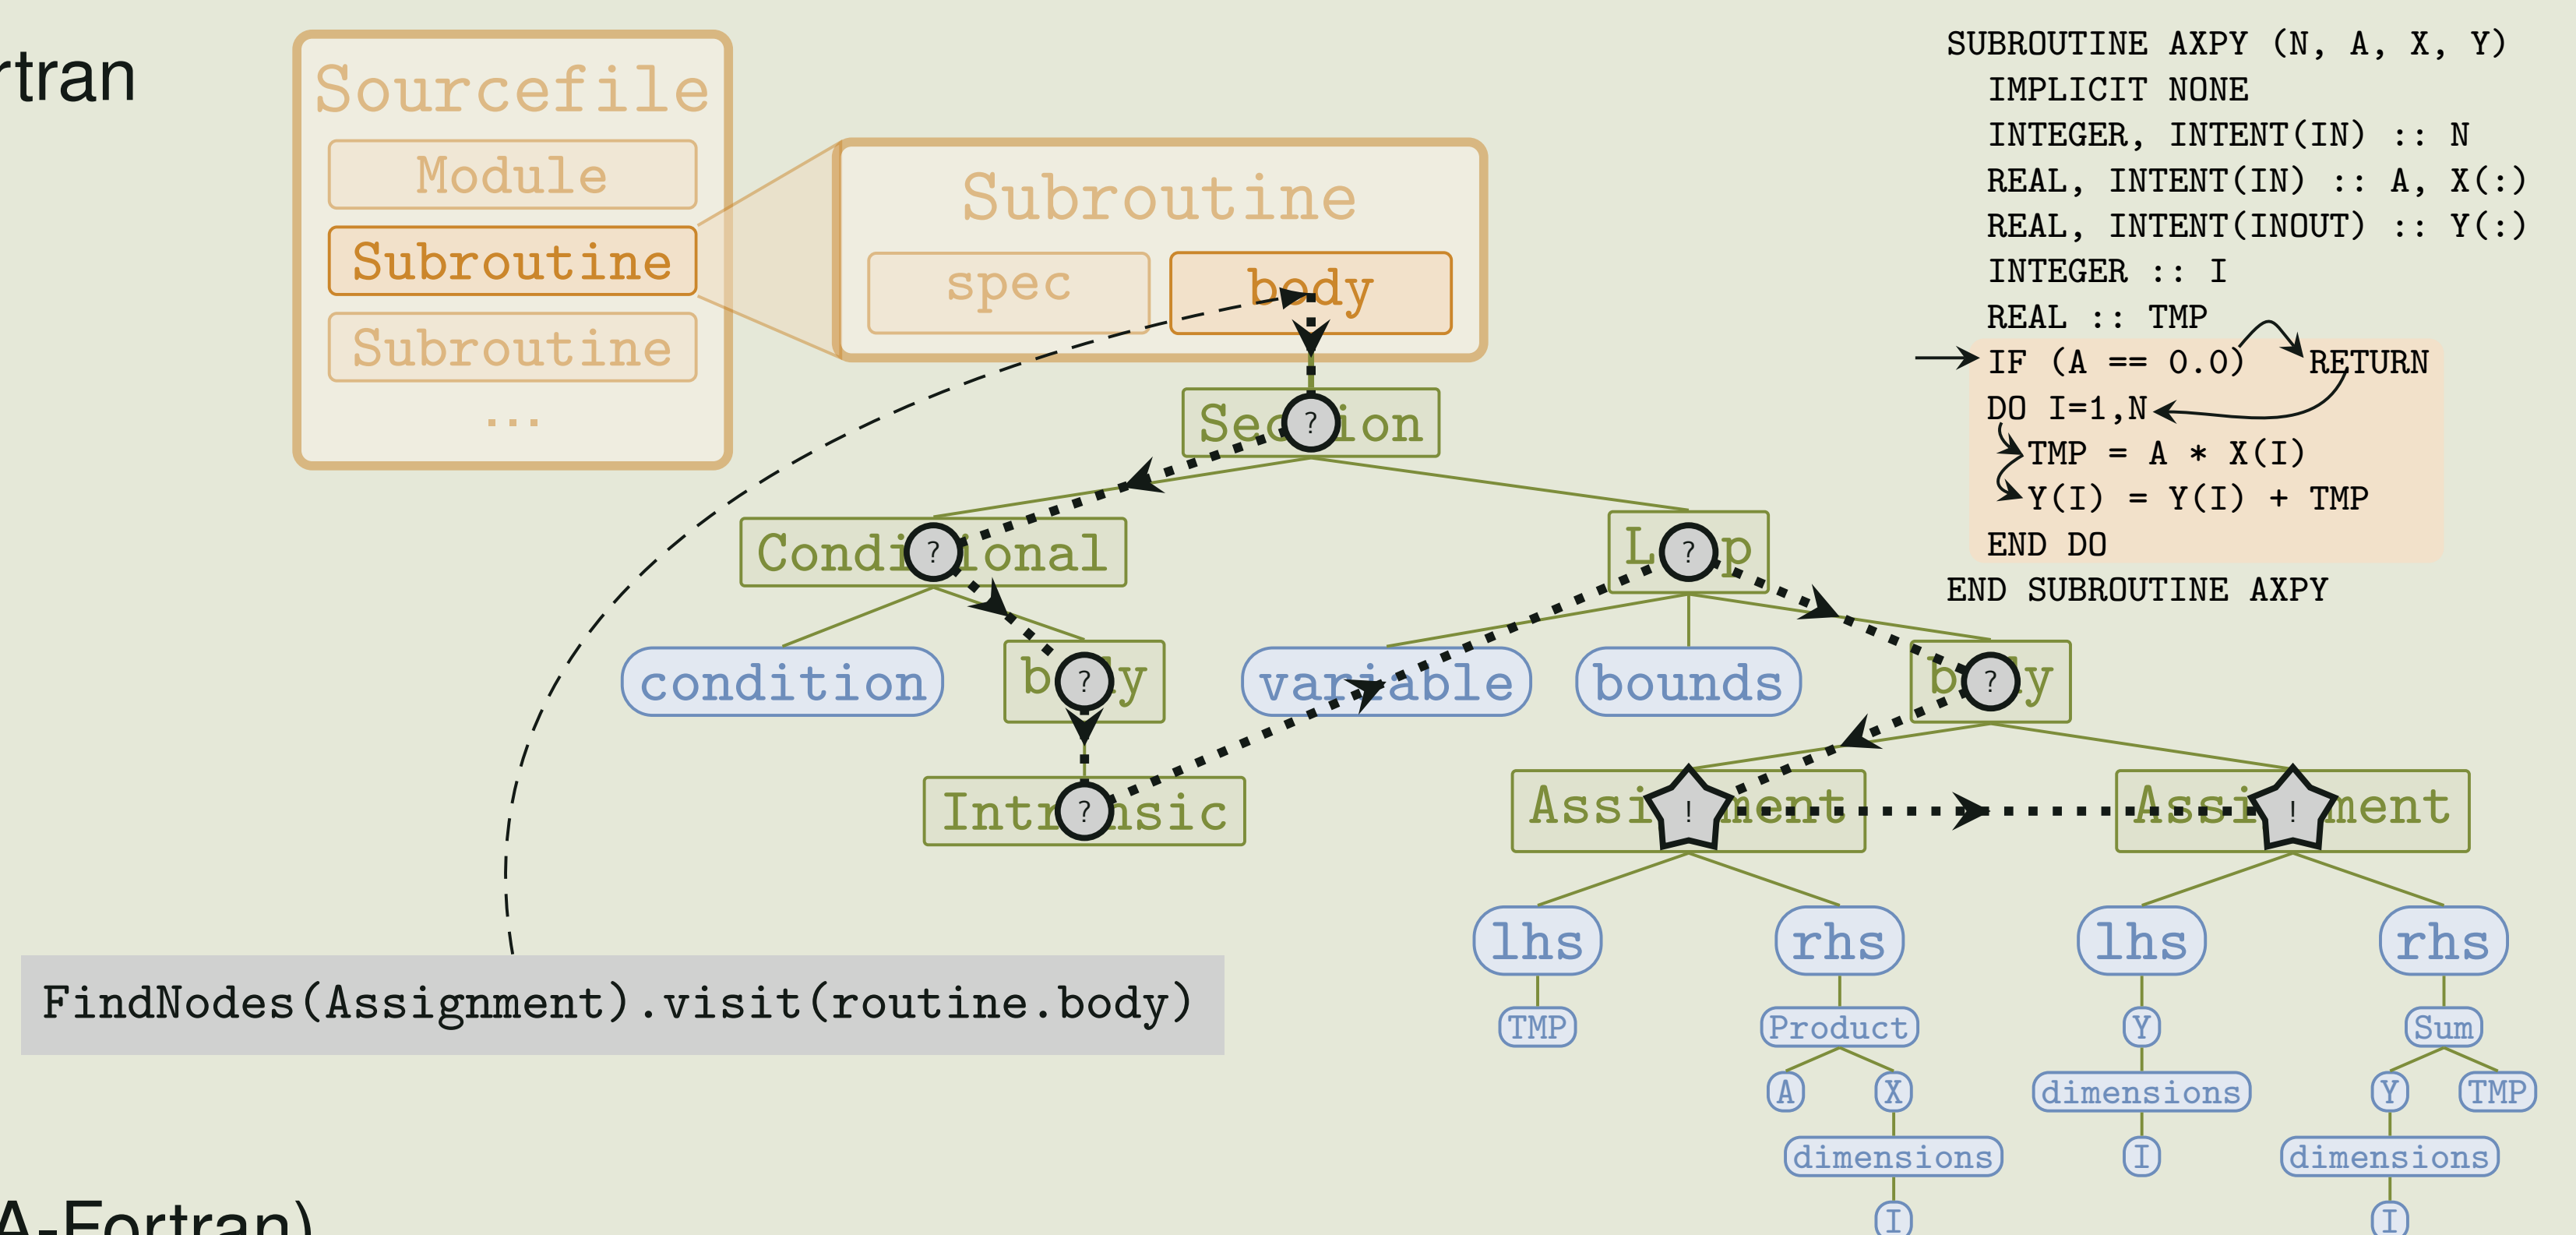
Loki is a **Python package** to encode S2S translation recipes for Fortran

Core library: Internal representation (IR) and API to **encode custom transformations** or **analysis/linting pipelines**

- *Fparser2*¹ is used to generate parse tree of Fortran source
- The parse tree is converted into Loki's **two-level IR**, separating (Fortran-tinted) **control-flow** from **expression tree**

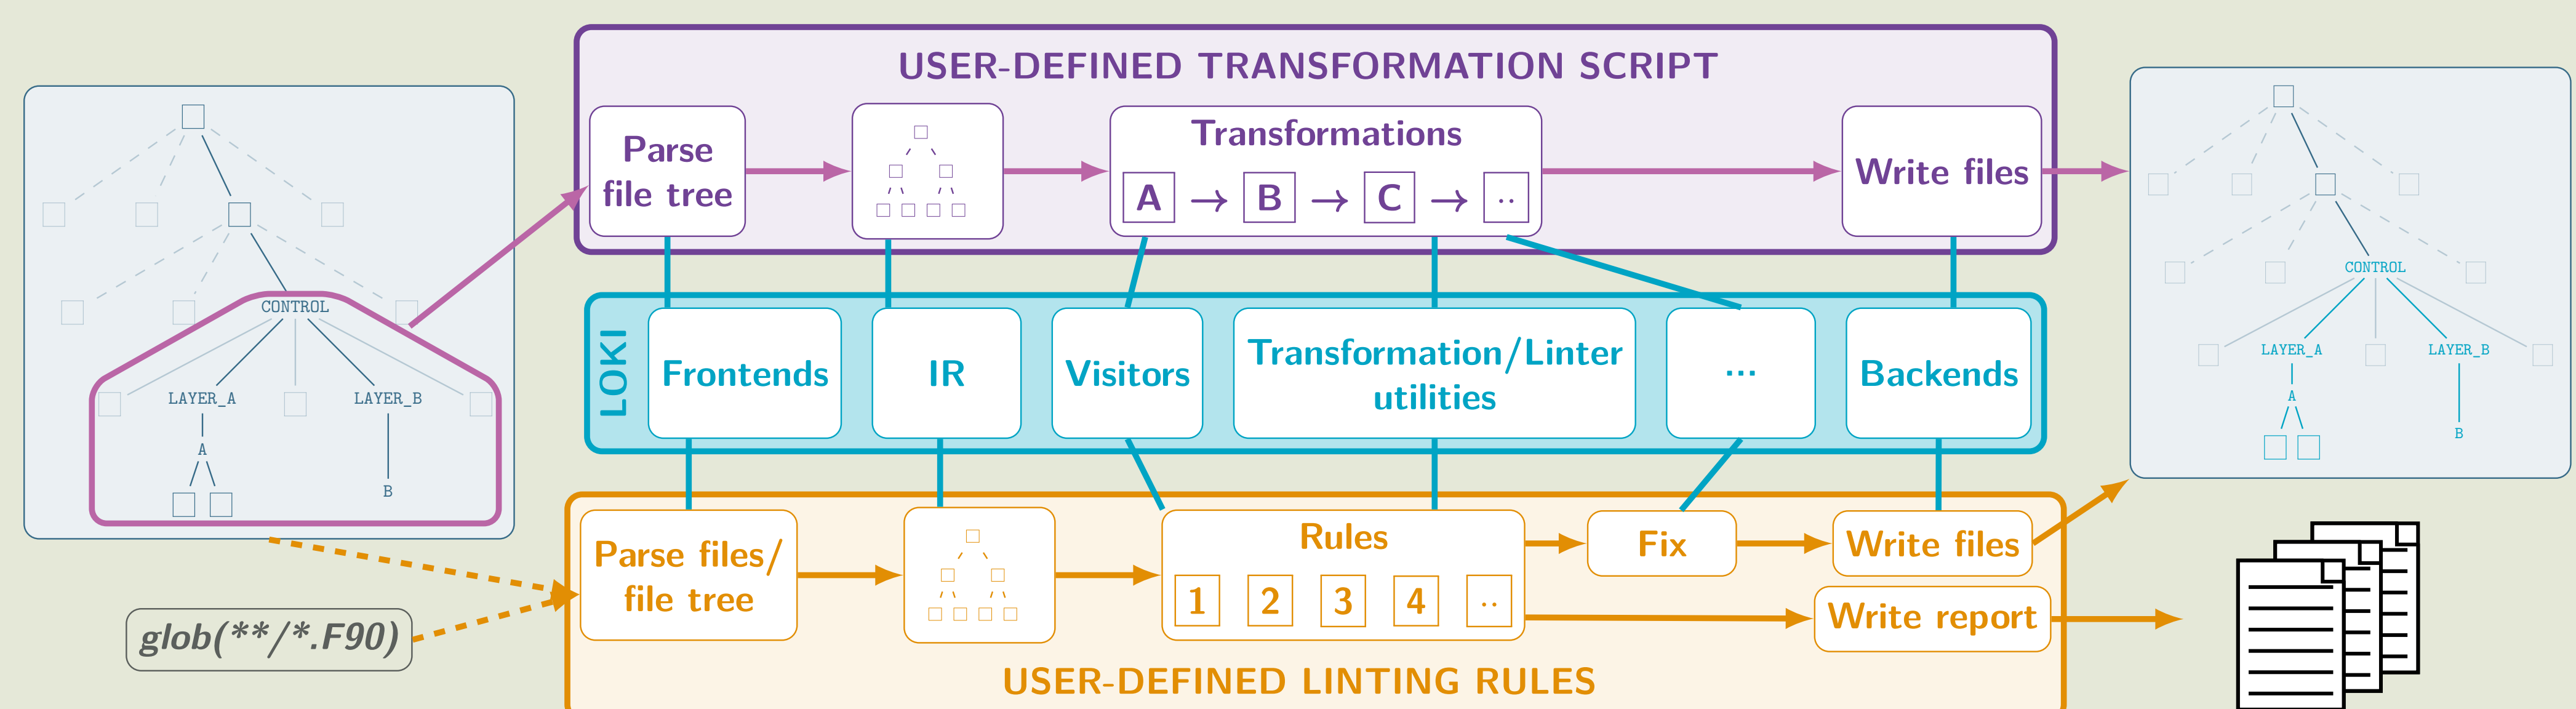
Features:

- **Visitors** are used to traverse and transform the IR
- Scope-aware symbol tables **manage type information**
- **Scheduler** builds a dependency graph for call trees across source files and allows for **inter-procedural analysis**
- **Backends** to generate Fortran (experimental: C, Python, or CUDA-Fortran)



Bulk transformation and analysis of source code

- Typical S2S translation recipes consist of **multiple bespoke transformation steps**
- **User-defined** pipeline of transformation steps can be built using **core library** utilities
- **Scheduler** applies transformations in the order of the dependency graph
- **CMake integration** automatically updates dependencies of build system targets
- Same infrastructure unlocks **custom static code analysis** and experimental fixing of coding rule violations



Loki integration into standalone mini-apps extracted from IFS

CLOUDSC (cloud microphysics) serves as a **proxy for single-column algorithms** to develop transformation recipes for IFS physics

ecWAM is the **operational IFS wave model**, consisting of dynamical core and physics

SCC: Single Column Coalesced GPU transformation


- Swap horizontal/vertical loop, demote arrays
- **HOIST**: pre-allocate temporaries in driver
- **STACK**: pool-allocator for temporaries
- Comparison: AMD EPYC 7742 vs. NVIDIA A100 40GB



📄/ecmwf-ifs/dwarf-p-cloudsc




📄/ecmwf-ifs/ecwam



Growing user base

Loki is used for the GPU adaptation of NWP models at **ECMWF**, **Météo-France**, in the **ACCORD** consortium, and as a key component for the digital twins in **Destination Earth**.



stories.ecmwf.int/destination-earth

Outlook and Plans

Core library:

- Add **parallel processing** support in the Scheduler
- Expand **data flow analysis** to unlock advanced transformation steps
- **Fortran-to-C** translation for kernel languages (CUDA, HIP, SYCL)

IFS transformation recipes:

- Automate advanced transformation recipes (e.g., k-caching)
- Expand offload support from OpenACC to OpenMP

Static code analysis:

- Automatic checking and integration into PR review process

1. Science and Technology Facilities Council. *fparser*. <https://github.com/stfc/fparser>.