



BERKELEY LAB

Bringing Science Solutions to the World



# Please, No More Loops (Than Necessary):

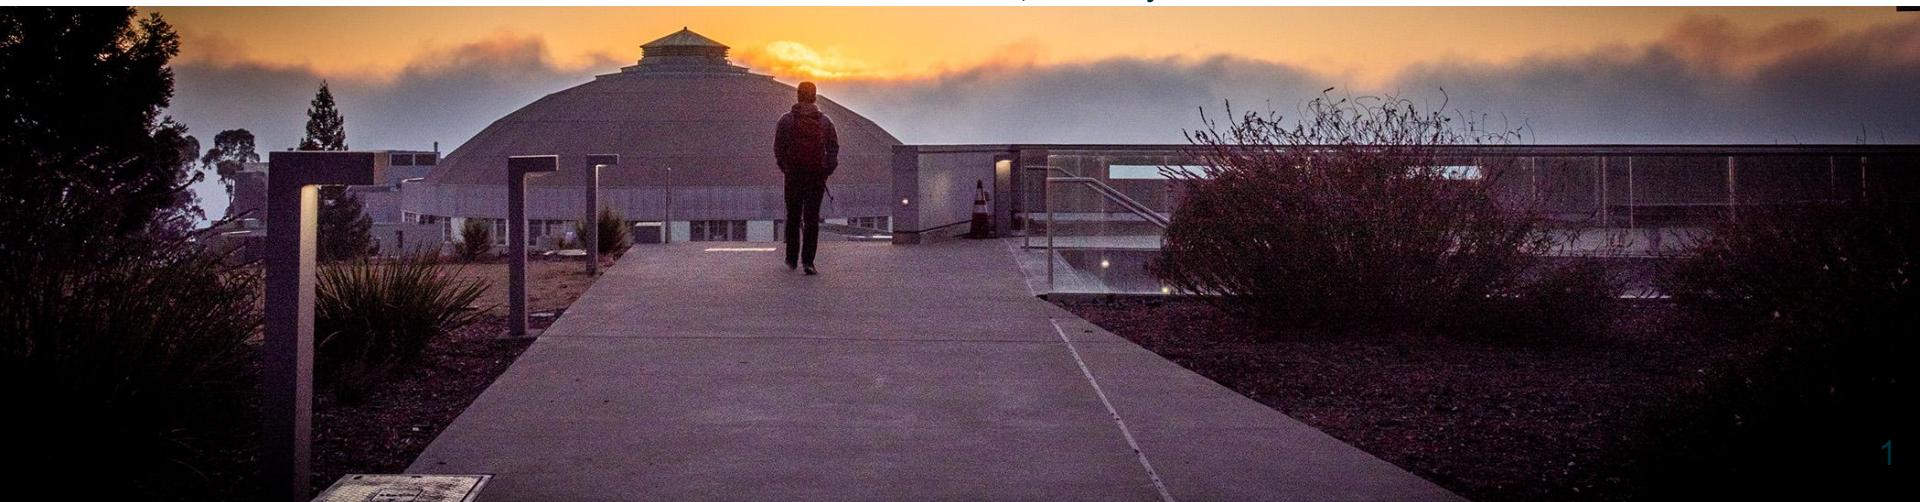
New Patterns in Fortran 2023

Damian Rouson

Computer Languages and Systems Software (CLaSS) Group

<https://go.lbl.gov/fortran>

HPC Best Practices Webinar, 21 January 2026



# Acknowledgements

## The Berkeley Lab Fortran Team

Dan Bonachea, Brandon Cook, Paul Hargrove, Hugh Kadhem, Kate Rasmussen

## Collaborators

**Julienne:** Katherine Rasmussen, Dan Bonachea

**Fiats:** Zhe Bai, Jeremiah Bailey, Baboucarr Dibba, Dan Bonachea, Ethan Gutmann, Brad Richardson, Sameer Shende, David Torres, Katherine Rasmussen, Kareem Jabbar Weaver, Jordan Welsman, Yunhao Zhang

**Formal, MOLE:** Dan Bonachea, Jose Castillo, Johnny Corbino, Joseph Hellmers

**Flang, Caffeine, PRIF:** Dan Bonachea, Kareem Ergawy, Jeff Hammond, Michael Klemm, Jean-Didier Pailleux, Etienne Renault, Katherine Rasmussen, Brad Richardson

**Matcha:** Dan Bonachea, David Torres, Dominick Martinez

## Sponsors

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, and Office of Nuclear Physics. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231. This research was supported by the Lawrence Berkeley National Laboratory Laboratory Research and Development (LDRD) program.

# Overview

---

01

**Background:**  
Deep origins

02

**Features & Paradigms:**  
A walking tour from Fortran 90 to 202Y

03

**Use Cases:**  
Nooks, Crannies, and Pastures

04

**Conclusions**



# Overview

---

01

**Background:**  
Deep origins

02

**Features & Paradigms:**  
A walking tour from Fortran 90 to 202Y

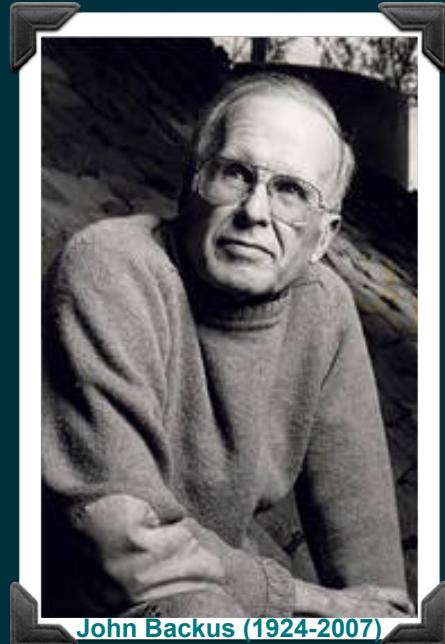
03

**Use Cases:**  
Nooks, Crannies, and Pastures

04

**Conclusions**



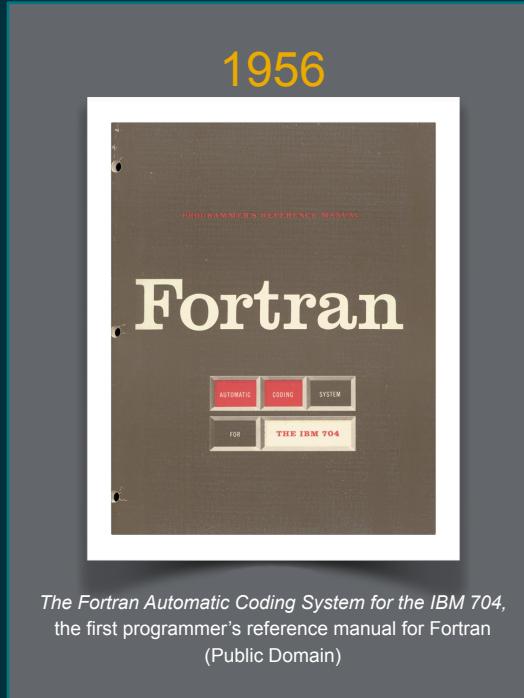


**John Backus (1924-2007)**

Pioneers in Science and Technology Series: John Backus, 1984

© City of Oak Ridge, Oak Ridge, TN 3783 (Public Domain)

<https://cdm16107.contentdm.oclc.org/digital/collection/p15388coll1/>





Fortran Manual (1956).pdf  
View Page 8 of 54

Sort By: **Search Rank** **Page Order**

Inspector Zoom Share Highlight Rotate Markup Form Filling Search

Fortran Manual (1956).pdf

CHAPTER 1. GENERAL PROPERTIES  
OF A FORTRAN SOURCE PROGRAM

A FORTRAN source program consists of a sequence of FORTRAN *statements*. There are 32 different types of statement, which are described in detail in the chapters which follow.

**Example of a Fortran Program**

The following brief program will serve to illustrate the general appearance and some of the properties of a FORTRAN program. It is shown as coded on a standard FORTRAN coding sheet.

FORTRAN STATEMENT	CONTINUATION
1. PROGRAM FOR FINDING THE LARGEST VALUE ATTAINED BY A SET OF NUMBERS	2. 3.
2. BIGA = A(1)	
3. DO 20 I = 2,N	
4. IF (BIGA .LT. A(I)) 10, 20, 20	
5. 10. BIGA = A(I)	
6. 20. CONTINUE	

This program examines the set of  $n$  numbers  $a_i$  ( $i=1, \dots, n$ ) and sets  $biga$  (after a continuation card) to the largest value attained in the set. It begins by setting  $biga$  equal to the first number in the set. The DO statement causes the succeeding statements to be included in the loop and carried out repeatedly, first with  $i=2$ , then with  $i=3$ , etc., until finally with  $i=n$ . During each repetition of this loop the IF statement compares  $biga$  with  $a_i$ ; if  $biga$  is less than  $a_i$ , statement 10, which replaces  $biga$  by  $a_i$ , is executed before continuing.

Each statement is punched on a separate card. If a statement is too long to fit on a single card it can be continued over as many as 9 additional continuation cards. For each statement the initial card must contain either a zero or a blank in column 6; on continuation cards column 6 must not contain a zero or a blank, and it should be used to number the continuation cards consecutively from 1 to 9. If a statement is too long to fit on a single line of the coding form, the programmer can signal to the keypuncher that he has continued on to the next line by placing a mark in the column labeled CONTINUATION.

7

“... the DO statement causes the succeeding statements to be carried out repeatedly...”

# Overview

---

01

Background:  
Deep origins

02

**Features &  
Paradigms:**

A walking tour  
from Fortran 90  
to 202Y+

03

Use Cases:  
Nooks,  
Crannies, and  
Pastures

04

Conclusions



# New Features



90



Pointers (required for dynamically allocating object components)

Defined operations and defined assignments

Recursive procedures

Allocatable arrays

Array statements, structure constructors, and intrinsic functions

Modules

Interface blocks & bodies, including generic interfaces

Structured branching and looping

Kind parameters and related intrinsic functions

Free-form source (.f90)

# Paradigms



Object-Based Programming



Functional programming



Array programming



Modular programming



Structured programming

# New Features



pure procedures  
elemental procedures  
where construct  
forall construct (obsolescent)

# Paradigms

- } Functional programming
- } Array programming
- } Parallel/vector programming

# New Features



Type extension, type-bound procedures, polymorphism, and final subroutines

Generic bindings, including type-bound operators and assignments

Type/source allocation & automatic (re-)allocation via intrinsic assignment

Allocatable scalars and components

Parameterized derived types

Vector subscripts

Abstract types, abstract interfaces, and deferred bindings

**Associate** construct

`iso_c_binding` module

`Iso_fortran_env` module: `input_unit`, `error_unit`, ...

# Paradigms

}

Object-Oriented Programming (OOP)

}

Generic programming

}

Array programming

}

Object-Oriented Design Patterns (OOD)

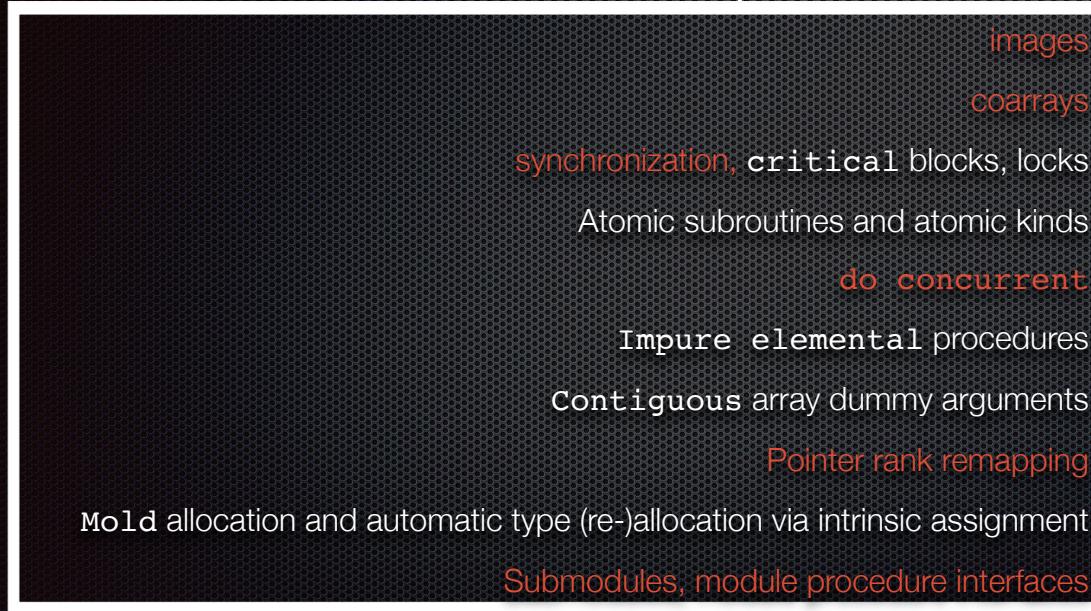
}

Functional programming

}

(C Interoperability)

# New Features



# Paradigms



- Parallel/GPU programming:
  - SPMD
  - PGAS
  - Shared or distributed memory



- Array programming



- Modular programming

# New Features



90



95



2003



2008



2018



Assumed-type (`type(*)`) and assumed-rank (`a(...)`) dummy arguments and rank guarding (select rank)

`ISO_Fortran_binding.h:`

**Data structures:** `CFI_cdesc_t`, `CFI_dim_t`, `CFI_rank_t`, `CFI_type_t`, `CFI_attribute_t`

**Functions:** `CFI_address()`, `CFI_allocate()`, `CFI_deallocate()`, `CFI_establish()`,  
`CFI_is_contiguous()`, `CFI_section()`

**Collective subroutines:** `co_sum`, `co_broadcast`, `co_min`, `co_max`, `co_reduce`

**Events:** `event_type`, `event_query()`, `event_wait`, `event_post`

**Locality specifiers** for `do concurrent`: `local`, `shared`, `default(none)`

Teams

Failed images

# Paradigms



(C-Interoperability)



Parallel/vector programming

# New Features



```
do concurrent reduce locality
put notifications: notify_type, notify_wait
enumeration types
maximum statement length: 1 million characters
```

} Parallel/vector programming

# New Features



90



95



2003



2008



2018



2023



202Y

# Paradigms

Type-safe templates

Generic subroutines

Enhanced collective subroutines: team collectives,  
asynchronous collectives

} Generic programming

} Parallel programming

# Nooks & Crannies

www.merriam-webster.com

M NOOK Definition & Meaning - Merriam-Webster

Merriam-Webster

nook

Definition Synonyms Example Sentences Word History Phra

## nook noun

'nūk (ə)

Synonyms of *nook* >

1 chiefly Scotland : a right-angled corner

2 a : an interior angle formed by two meeting walls  
b : a secluded or sheltered place or part  
| searched every *nook* and *cranny*

c : a small often recessed section of a larger room  
| a *breakfast nook*

www.marthastewart.com/r

martha stewart

## 22 Reading Nook Ideas for Turning Any Space Into a Cozy Escape

Get inspired to find the perfect place to forget about your cares and slip away into a good book with these reading nook ideas.

By Heather Bien | Published on July 3, 2024



Credit: Barr Joinery & Lucy Walters

www.merriam-webster.com

M CRANNY Definition & Meaning - Merriam-Webster

Merriam-Webster

cranny

Definition Synonyms Example Sentences Word History Phra

## cranny noun

cran·ny (kra-nē)

plural *crannies*

Synonyms of *cranny* >

1 : a small break or slit : *CREVISE*

2 : an obscure nook or corner

**crannied** (kra-nēd) adjective

# Fortran 2018 Intrinsic Functions

2018-08-28

Table 16.

Procedure
ALLOCATED
ANY
ASIN
ASINH
ASSOCIATED
ATAN
ATAN2
ATANH
ATOMIC_ADD
ATOMIC_AND
ATOMIC_CAS
ATOMIC_DEFINE
ATOMIC_FETCH_<sub>ADD</sub>
ATOMIC_FETCH_<sub>AND</sub>
ATOMIC_FETCH_<sub>OR</sub>
ATOMIC_FETCH_<sub>XOR</sub>
ATOMIC_REF
BESSEL_10
BESSEL_J1
BESSEL_JN
BESSEL_N0
BESSEL_Y1
BESSEL_YN
BIGE
BGT
BIT_SIZE
BLE
BIT
BTTEST
CEILING
CHAR
CMPXL
CO_BROADCAST
CO_MAX
CO_MIN
CO_REDUCE
CO_SUM
COMMAND_ARGUMENT_COUNT
CONJG
COS
COSH
COSHAPEx
COUNT
COPY_TIME
CSHIFT
DATE_AND_TIME
DBLE
DIGITS
DIM
DOF_PRODUCT
DPROD
DSHIFTL
DSHIFTTR
EOSHIFT
EPSILON
ERF
ERFC
ERFC_SCALED
EVENT_QUERY
EXTRACT_COM- MAND_LINE
EXP
EXPONENT
EXTENDS_TYPE_<sub>OF</sub>
FAILED_IMAGES
FINLOC
FLOOR
FRACTION
GAMMA
GET_COMMAND_<sub>ARGUMENT</sub>
GET_ENVIRON- MENT_VARIABLE
LOG10
LOGICAL
MASKL
MASKR
MATMUL
MAX
MAXEXPONENT
MAXLOC

Table 16.

Procedure
CEPHEID
HUGE
HYPOT
LACHAR
DALL
LAND
IANY
IBCLR
IBITS
ISRCT
ICAHM
IEOR
IMAGE_INDEX
IMAGE_STATUS
INDEX
INT
IOR
IPARTY
ISHFT
ISHFTC
IS_CONTIGUOUS
IS_EQUIV_END
IS_RSTAT_FOR
KIND
LBOUND
LCBOUND
LEADZ
LEN
LEN_TRIM
LGE
LGT
LLE
LLT
LOG
LOG_GAMMA
RANDOM
RANDOM_GAMMA
RANDOM_NUMBER
RANDOM_SEED
RANGE
RANK
REAL
REDUCE

334

2018-08-28

Table 16.

Procedure
MAXVAL
MERGE
MERGE_BITS
MIN
MINEXPONENT
MINLOC
MINVAL
MOD
MODULO
MOVE_ALLOC
MVBITS
NEAREST
NEW_LINE
NINT
NORM2
NOT
NULL
NUM_IMAGES
OUT_OF_RANGE
PACK
PARTY
POPCNT
POPPAR
PRECISION
PRESENT
PRODUCT
RADIX
RANDOM_INIT
RANDOM_NUMBER
RANDOM_SEED
RANGE
RANK
REAL
REDUCE

2018-08-28

Table 16.

Procedure
MAXVAL
MERGE
MERGE_BITS
MIN
MINEXPONENT
MINLOC
MINVAL
MOD
MODULO
MOVE_ALLOC
MVBITS
NEAREST
NEW_LINE
NINT
NORM2
NOT
NULL
NUM_IMAGES
OUT_OF_RANGE
PACK
PARTY
POPCNT
POPPAR
PRECISION
PRESENT
PRODUCT
RADIX
RANDOM_INIT
RANDOM_NUMBER
RANDOM_SEED
RANGE
RANK
REAL
REDUCE

336

WD 1539-1

J3/18-007r1

Procedure	Arguments	Class	Description
REPEAT	(STRING, NCOPIES)	T	Repetitive string concatenation.
RESHAPE	(SOURCE, SHAPE [, PAD, ORDER])	T	Arbitrary shape array construction.
RRSPACING	(X)	E	Reciprocal of relative spacing of model numbers.
SAME_TYPE_AS	(A, B)	I	Dynamic type equality test.
SCALE	(X, B)	E	Real number scaled by radix power.
SCAN	(STRING, SET [, BACK, KIND])	E	Character set membership search.
SELECTED_CHAR_<sub>NAME</sub>	(NAME)	T	Character kind selection.
KIND		T	Integer kind selection.
SELECTED_INT_<sub>NAME</sub>	(R)	T	Real kind selection.
KIND		T	Real kind selection.
SELECTED_REAL_<sub>NAME</sub>	(P, R, RADIX)	T	Real kind selection.
KIND		T	Real kind selection.
SET_EXPONENT	(X, I)	E	Real power with specified exponent.
SHRINK	(SOURCE [, KIND])	I	Truncate of an array or a scalar.
SHIPTA	(I, SHIFT)	E	Right shift with fill.
SHIPTL	(I, SHIFT)	E	Left shift.
SHITR	(I, SHIFT)	E	Right shift.
SIGN	(A, B)	E	Magnitude of A with the sign of B.
SINH	(X)	E	Hyperbolic sine function.
SINH	(X)	E	Hyperbolic sine function.
SIZE	(ARRAY [, DIM, KIND])	I	Size of an array or one extent.
SPREAD	(ARRAY [, DIM, NCOPIES])	T	Splicing of model numbers.
SQRT	(X)	E	Square root.
STOPPED_IMAGES	(TEAM, KIND)	I	Indices of stopped images.
STORAGE_SIZE	(A, KIND)	I	Storage size in bits.
SUM	(ARRAY [, DIM [, MASK]]) or (ARRAY [, MASK])	T	Array reduced by addition.
SYSTEM_CLOCK	(COUNT, COUNT_RATE, COUNT_MAX)	S	Query system clock.
TAN	(X)	E	Tangent function.
TANH	(X)	E	Hyperbolic tangent function.
TEAM_NUMBER	(TEAM)	T	Team number.
THIS_IMAGE	(TEAM)	T	Index of the invoking image.
THIS_IMAGE	(COARRAY [, TEAM]) or (COARRAY [, DIM [, TEAM]])	T	Conscript(s) for the image.
TINY	(I)	I	Smallest positive model number.
TRAILZ	(SOURCE, MOLD [, SIZE])	E	Number of trailing zeros.
TRANSPOSE	(MATRIX)	T	Transpose of an array of rank two.
TRIM	(STRING)	T	String without trailing blanks.
UBOUND	(ARRAY [, DIM, KIND])	I	Upper bound(s).
UCOUNT	(COARRAY [, TEAM [, KIND])	I	Upper bound(s) of a coarray.
UNPACK	(VECTOR, MASK, FIELD)	T	Vector unpacked into an array.
VERIFY	(STRING, SET [, BACK, KIND])	E	Character set non-membership search.

<sup>1</sup> The effect of calling EXECUTE\_COMMAND\_LINE on any image other than image 1 in the initial team is processor dependent.

J3/18-007r1

337

# Expansive Pastures



Davos, Switzerland, on the way to PASC23

# Fortran 2023 Multi-Image Features

- Statements
  - Synchronization
    - Explicit: `SYNC ALL`, `SYNC IMAGES`, `SYNC MEMORY`, `SYNC TEAM`
    - Implicit: `ALLOCATE`, `DEALLOCATE`, `STOP`, `END`, `MOVE_ALLOC`
  - Events: `EVENT POST`, `EVENT WAIT`
  - Notify: `NOTIFY WAIT`
  - Error termination: `ERROR STOP`
  - Locks: `LOCK`, `UNLOCK`
  - Failed images: `FAIL IMAGE`
  - Teams: `FORM TEAM`, `CHANGE TEAM`
  - Critical sections: `CRITICAL`, `END CRITICAL`
- Coarray Accesses (`[...]`)
- Intrinsic functions: `NUM_IMAGES`, `THIS_IMAGE`, `LCOBOUND`, `UCOBOUND`, `TEAM_NUMBER`, `GET_TEAM`, `FAILED_IMAGES`, `STOPPED_IMAGES`, `IMAGE_STATUS`, `COSHAPE`, `IMAGE_INDEX`

- Intrinsic subroutines
  - Collective subroutines: `CO_SUM`, `CO_MAX`, `CO_MIN`, `CO_REDUCE`, `CO_BROADCAST`
  - Atomic subroutines: `ATOMIC_ADD`, `ATOMIC_AND`, `ATOMIC_CAS`, `ATOMIC_DEFINE`, `ATOMIC_FETCH_ADD`, `ATOMIC_FETCH_AND`, `ATOMIC_FETCH_OR`, `ATOMIC_FETCH_XOR`, `ATOMIC_OR`, `ATOMIC_REF`, `ATOMIC_XOR`
  - Other subroutines: `EVENT_QUERY`
- Types, kind type parameters, and values
  - Intrinsic derived types: `EVENT_TYPE`, `TEAM_TYPE`, `LOCK_TYPE`, `NOTIFY_TYPE`
  - Atomic kind type parameters: `ATOMIC_INT_KIND` and `ATOMIC_LOGICAL_KIND`
  - Values: `STAT_FAILED_IMAGE`, `STAT_LOCKED`, `STAT_LOCKED_OTHER_IMAGE`, `STAT_STOPPED_IMAGE`, `STAT_UNLOCKED`, `STAT_UNLOCKED_FAILED_IMAGE`

# Overview

---

01

**Background:**  
Deep origins

02

**Features & Paradigms:**  
A walking tour from Fortran 90 to 2023 & beyond

03

**Use Cases:**  
Nooks, Crannies, and Pastures

04

Conclusions



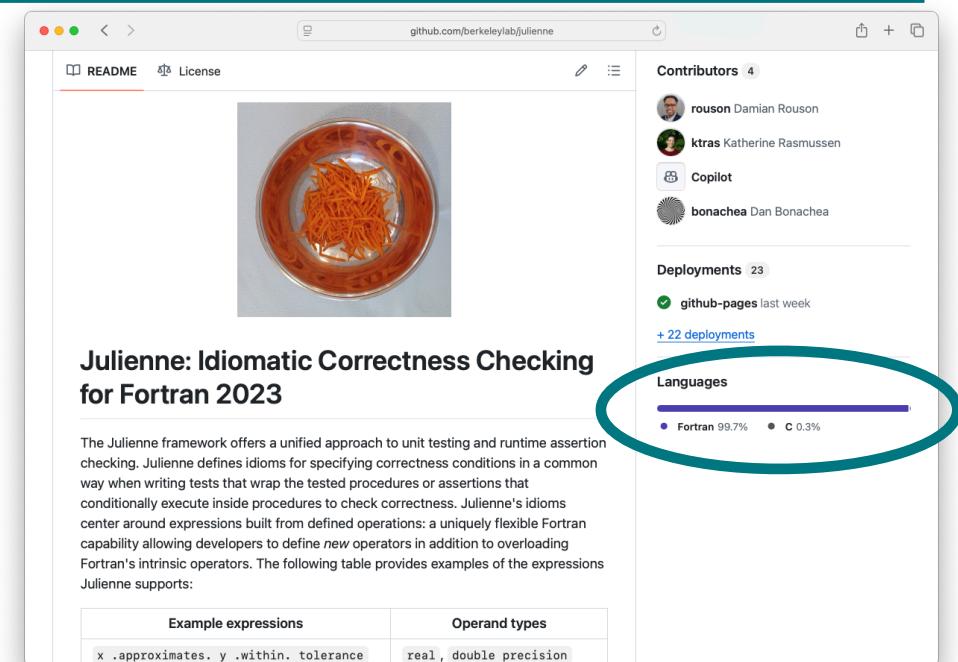
# Correctness-Checking with Julienne

## Unified Idioms for writing

- Unit tests
- Assertions

## Support for Fortran 2023 parallelism

- Multi-image testing: a collective reduction detects failure on a subset of images
- Assertions are pure procedures as required for invocation inside a `do concurrent` construct.



<https://go.lbl.gov/julienne>

Rouson, Bonachea, & Rasmussen, "[Idiomatic Correctness-Checking via Julienne in Fortran 2023](#)", *Proceedings of the US Research Software Engineering Conference*, October 2025.  
DOI: [10.25344/S4BG65](https://doi.org/10.25344/S4BG65)

Please, No More Loops (Than Necessary)



**BERKELEY LAB**

Bringing Science Solutions to the World

# Julienne Idioms

Row	Example expressions	Supported operand types of the bold operator
1	<code>x .approximates. y .within. tolerance</code>	real, double precision
2	<code>x .approximates. y .withinFraction. tolerance</code>	real, double precision
3	<code>x .approximates. y .withinPercentage. tolerance</code>	real, double precision
4	<code>.all. ([i,j] .lessThan. k)</code>	test_diagnosis_t
5	<code>.all. ([i,j] .lessThan. [k,m])</code>	test_diagnosis_t
6	<code>.all. (i .lessThan. [k,m])</code>	test_diagnosis_t
7	<code>(i .lessThan. j) .also. (k .equalsExpected. m)</code>	test_diagnosis_t
8	<code>x .lessThan. y</code>	integer, real, double precision
9	<code>x .greaterThan. y</code>	integer, real, double precision
10	<code>i .equalsExpected. j</code>	integer, character, type(c_ptr)
11	<code>i .isAtLeast. j</code>	integer, real, double precision
12	<code>i .isAtMost. j</code>	integer, real, double precision
13	<code>s .isBefore. t</code>	character
14	<code>s .isAfter. t</code>	character
15	<code>(.expect. allocated(A)) // '(expected an allocated array "A")'</code>	logical

## Elemental Operators:

- Defined as pure functions
- Binary operators accept conformable operands:
  - Same-shaped arrays
  - Scalar/array combinations



Any 1 elemental function defining a binary operator accepts 46 combinations of operands:

- 1 for scalar operands
- 15 for array operands of rank 1-15
- 30 for scalar/array combinations in either order



# Writing PDE Solvers with Formal

## Fortran mimetic abstraction language

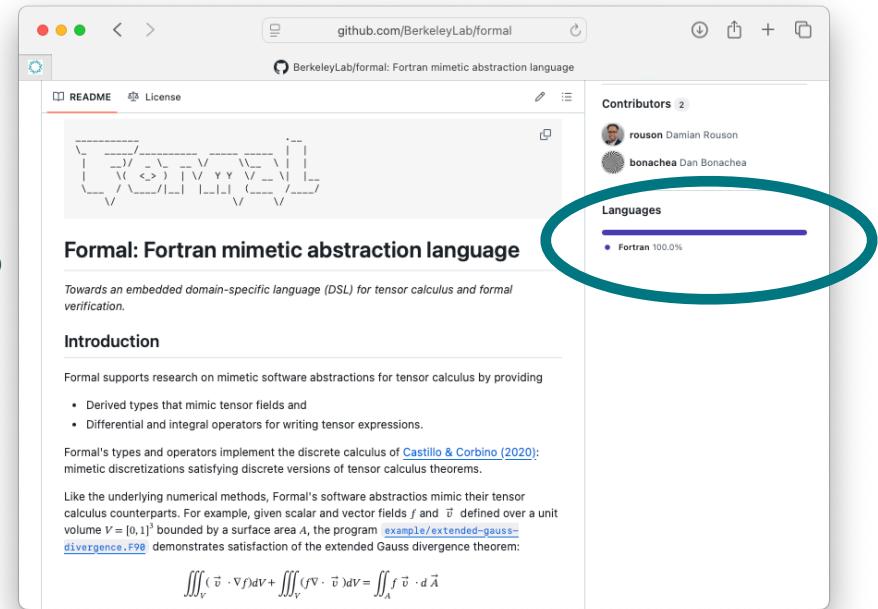
- Derived types that mimic tensor fields
- Differential and integral operators
- Supported by a discrete calculus based on the mimetic discretizations of Corbino & Castillo (2020).

## Future work

- An embedded domain-specific language (DSL) for tensor calculus
- Formal verification leveraging problem-specific proof by testing
- Tensor contractions for machine learning

Corbino, J., & Castillo, J. E. (2020). "High-order mimetic finite-difference operators satisfying the extended Gauss divergence theorem". *Journal of Computational and Applied Mathematics*, 364, 112326.

Please, No More Loops (Than Necessary)

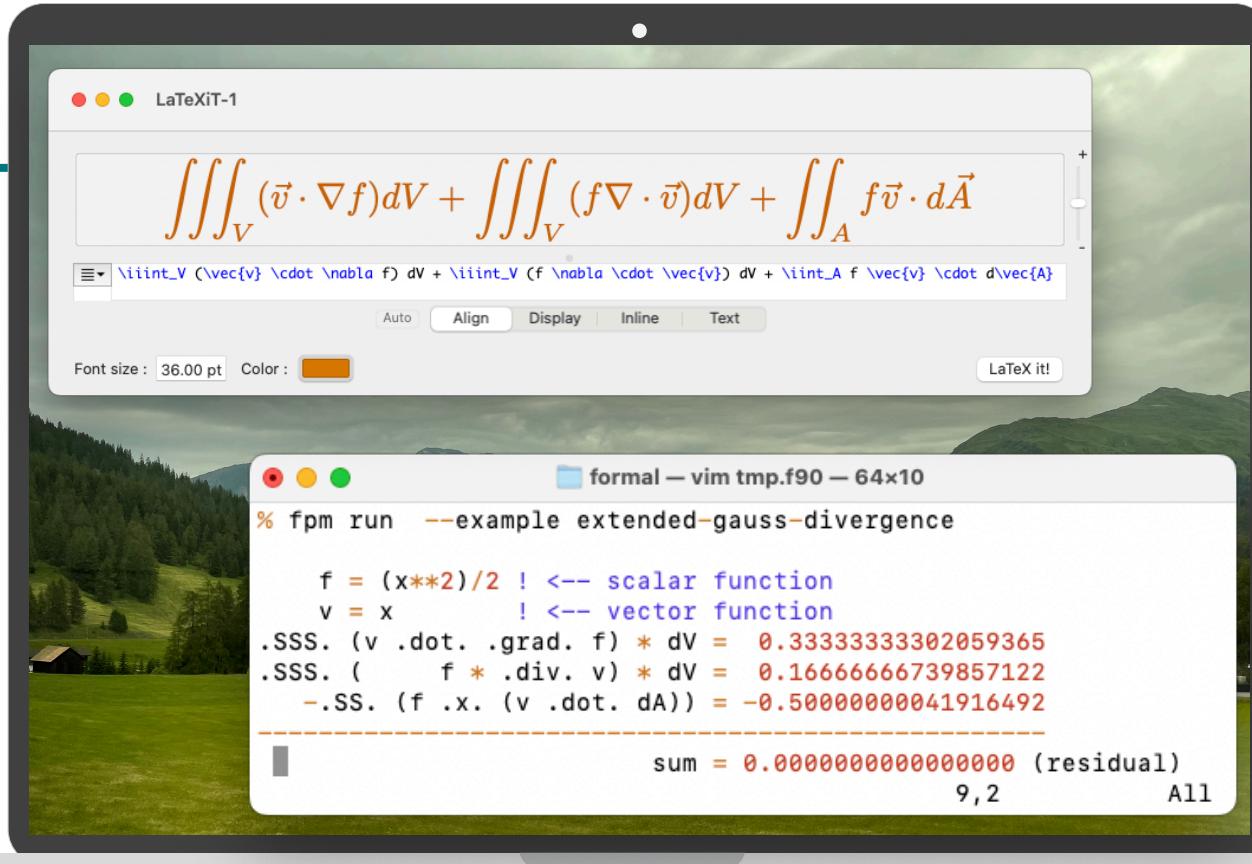


<https://go.lbl.gov/formal>



BERKELEY LAB

Bringing Science Solutions to the World



Please, No More Loops (Than Necessary)



BERKELEY LAB

Bringing Science Solutions to the World

“When writing type-safe templates in Fortran, you can consider the requirements as defining a DSL for the template body. Such DSLs are extremely cheap to define, just a collection of derived declarations, and have no runtime overhead.”

Prof. Magne Haveraaen  
Bergen Language Design Laboratory  
University of Bergen

# REQUIREMENT Construct

```
REQUIREMENT binop(op, T, U, V)
  DEFERRED TYPE :: T, U, V
  DEFERRED INTERFACE
    FUNCTION op(x,y) RESULT(z)
      TYPE(T), INTENT(IN) :: x
      TYPE(U), INTENT(IN) :: y
      TYPE(V) :: z
    END FUNCTION
  END INTERFACE
END REQUIREMENT
```

```
TEMPLATE my_temp1(T, U, plus, times)
  USE requirements_mod, only: binop

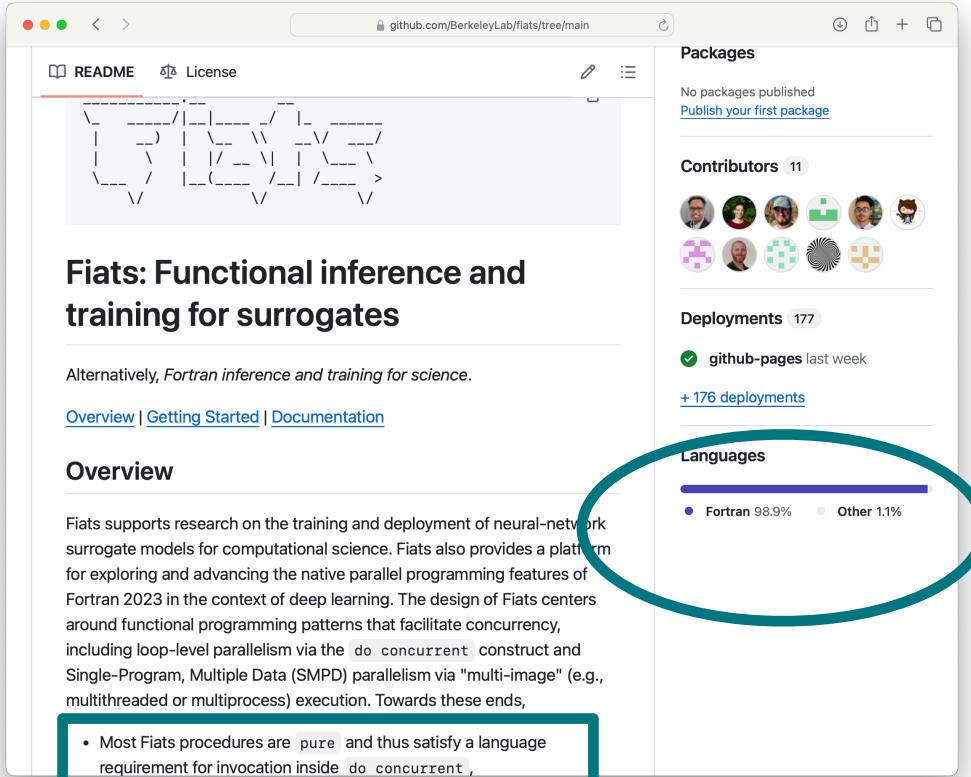
  REQUIRE binop(plus, T, U, U) ! Real+complex -> complex
  REQUIRE binop(times, T, U, U) ! Real*complex -> complex
  ...
END TEMPLATE
```

A requirement encapsulates a reusable relationship among deferred arguments.

The **REQUIRE** statement enforces a REQUIREMENT

- Mismatch at template instantiation is compile-time error
- Transitively declares its arguments
- Can appear in template specification and requirement construct

# Deep Learning with Fiats



The screenshot shows the GitHub repository for Fiats. The README file contains a logo consisting of a grid of characters. Below the logo, the text reads: "Fiats: Functional inference and training for surrogates". A subtitle "Alternatively, Fortran inference and training for science." is present. Navigation links for "Overview", "Getting Started", and "Documentation" are shown. The repository statistics on the right include: "Packages" (0), "Contributors" (11), "Deployments" (177, with one recent deployment from "github-pages" last week and 176 more), and "Languages" (98.9% Fortran, 1.1% Other). A teal oval highlights the "Languages" section.

Alternatively, *Fortran inference and training for science.*

[Overview](#) | [Getting Started](#) | [Documentation](#)

## Overview

Fiats supports research on the training and deployment of neural-network surrogate models for computational science. Fiats also provides a platform for exploring and advancing the native parallel programming features of Fortran 2023 in the context of deep learning. The design of Fiats centers around functional programming patterns that facilitate concurrency, including loop-level parallelism via the `do concurrent` construct and Single-Program, Multiple Data (SMPD) parallelism via "multi-image" (e.g., multithreaded or multiprocess) execution. Towards these ends,

- Most Fiats procedures are `pure` and thus satisfy a language requirement for invocation inside `do concurrent`,

Please, No More Loops (Than Necessary)

<https://go.lbl.gov/fiats>



BERKELEY LAB

Bringing Science Solutions to the World

# Fiats: Inference

# Fiats: Inference

```
example — vim concurrent-inferences.f90 — 90x9
129      !$omp parallel do default(None) shared(neural_network,inputs,outputs) collapse(3)
130      do j=1,lon
131          do k=1,lev
132              do i=1,lat
133                  outputs(i,k,j) = neural_network%infer(inputs(i,k,j))
134              end do
135          end do
136      end do
```

# Fiats: Inference

```
example — vim concurrent-inferences.f90 — 90x9
129      !$omp parallel do default(None) shared(neural_network,inputs,outputs) collapse(3)
130      do j=1,lon
131          do k=1,lev
132              do i=1,lat
133                  outputs(i,k,j) = neural_network%infer(inputs(i,k,j))
134              end do
135          end do
136      end do
```

```
fiats — rouson@login38:~/fiats — login38 — vim example/concurrent-inferences.f90 — 103x5
122      do concurrent(i=1:lat, k=1:lev, j=1:lon) default(None) shared(outputs, neural_network, inputs)
123          outputs(i,k,j) = neural_network%infer(inputs(i,k,j))
124      end do
:
```

# Fiats: Inference

```
example — vim concurrent-inferences.f90 — 90x9
129      !$omp parallel do default(None) shared(neural_network,inputs,outputs) collapse(3)
130      do j=1,lon
131          do k=1,lev
132              do i=1,lat
133                  outputs(i,k,j) = neural_network%infer(inputs(i,k,j))
134              end do
135          end do
136      end do
```

```
fiats — rouson@login38:~/fiats — login38 — vim example/concurrent-inferences.f90 — 103x5
122      do concurrent(i=1:lat, k=1:lev, j=1:lon) default(None) shared(outputs, neural_network, inputs)
123          outputs(i,k,j) = neural_network%infer(inputs(i,k,j))
124      end do
:
```

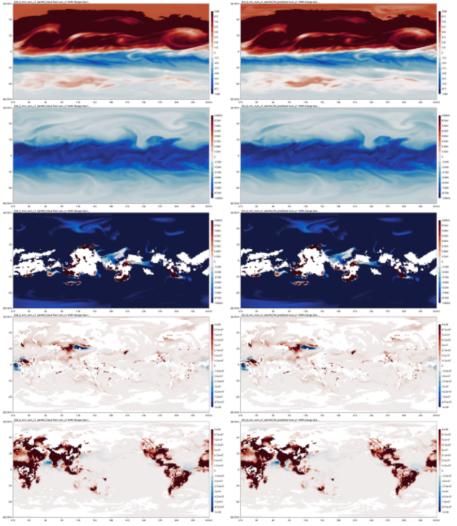
```
example — vim concurrent-inferences.f90 — 50x5
73      !$omp workshare
74          outputs = neural_network%infer(inputs)
75      !$omp end workshare
```

# Automatic Parallelization on Perlmutter CPU<sup>28</sup>

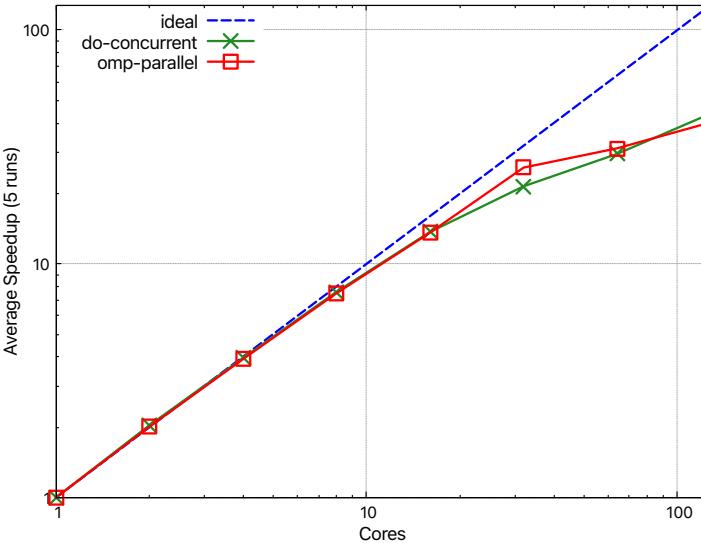
Automatically parallelizing batch inference on deep neural networks using Fiats and Fortran 2023 “do concurrent”

Damian Rouson<sup>1</sup> , Zhe Bai<sup>1</sup> , Dan Bonachea<sup>1</sup> , Kareem Ergawy<sup>2</sup> , Ethan Gutmann<sup>3</sup> , Michael Klemm<sup>2</sup> , Katherine Rasmussen<sup>1</sup> , Brad Richardson<sup>1</sup> , Sameer Shende<sup>4</sup> , David Torres<sup>5</sup> , and Yunhao Zhang<sup>1</sup> 

<sup>1</sup> Lawrence Berkeley National Laboratory, Berkeley CA 94720, USA



1. Fiats and Fortran 2023 “do concurrent” applications have evaluated or adopted deep neural networks as surrogate models. At least two categories of solutions have emerged to satisfy the inference and training needs of Fortran applications: (1) application programming interfaces (APIs) that expose functionality provided by software packages written



```
OMP_NUM_THREADS=128 fpm run \
--example concurrent-inferences \
--runner "srun --cpu_bind=cores -c 128 -n 1" \
--network model.json
```

Rouson, Bai, Bonachea, Ergawy, Gutmann, Klemm, Rasmussen, Richardson, Shende, Torres, and Zhang (2025). Automatically parallelizing batch inference on deep neural networks using Fiats and Fortran 2023 “do concurrent”. In *5th International Workshop on Computational Aspects of Deep Learning*, Hamburg, Germany.



**BERKELEY LAB**

Bringing Science Solutions to the World

# Fiats: Training

```
 907 #if F2023_LOCALITY
 908         iterate_through_batch: &
 909         do concurrent (pair = 1:mini_batch_size) local(a,z,delta) reduce(+: dcdb, dc当地
 910
 911 #elif F2018_LOCALITY
 912
 913         reduce_gradients: &
 914         block
 915             real reduce_dcdb(size(dcdb,1),size(dcdb,2),mini_batch_size)
 916             real reduce_dc当地
 917             reduce_dcdb = 0.
 918             reduce_dc当地 = 0.
 919
 920         iterate_through_batch: &
 921         do concurrent (pair = 1:mini_batch_size) local(a,z,delta)
 922
 923 #else
 924
 925         reduce_gradients: &
 926         block
 927             real reduce_dcdb(size(dcdb,1),size(dcdb,2),mini_batch_size)
 928             real reduce_dc当地
 929             reduce_dcdb = 0.
 930             reduce_dc当地 = 0.
 931
 932         iterate_through_batch: &
 933         do concurrent (pair = 1:mini_batch_size)
 934
 935         iteration: &
 936         block
 937
 938             real a(maxval(self%nodes_), input_layer:output_layer) ! Activations
 939             real z(size(b,1),size(b,2)), delta(size(b,1),size(b,2))
 940 #endif
 941
```

# Stochastic Gradient Descent + Adam Optimizer

~96 statements in which nearly every statement implicitly exposes parallelism, e.g., multidimensional array statements inside do concurrent constructs

# Deep Learning with Fiats

```
flats - vim neural_network_m.f90 - 90x14
21  type neural_network_t(k)
22    !! Encapsulate the information needed to perform inference
23    integer, kind :: k = default_real
24    type(tensor_map_t(k)), private :: input_map_, output_map_
25    type(metadata_t), private :: metadata_
26    real(k), allocatable, private :: weights_(:,:,:,:), biases_(:,:,:)
27    integer, allocatable, private :: nodes_()
28    type(activation_t), private :: activation_
29  contains
30    generic :: infer => default_real_infer, double_precision_infer
31    procedure, private, non_overridable :: default_real_infer, double_precision_infer
32    generic :: learn => default_real_learn
33    procedure, private, non_overridable :: default_real_learn
```

Kind type parameter allows us to set an object's precision in its declaration *without* recompiling.

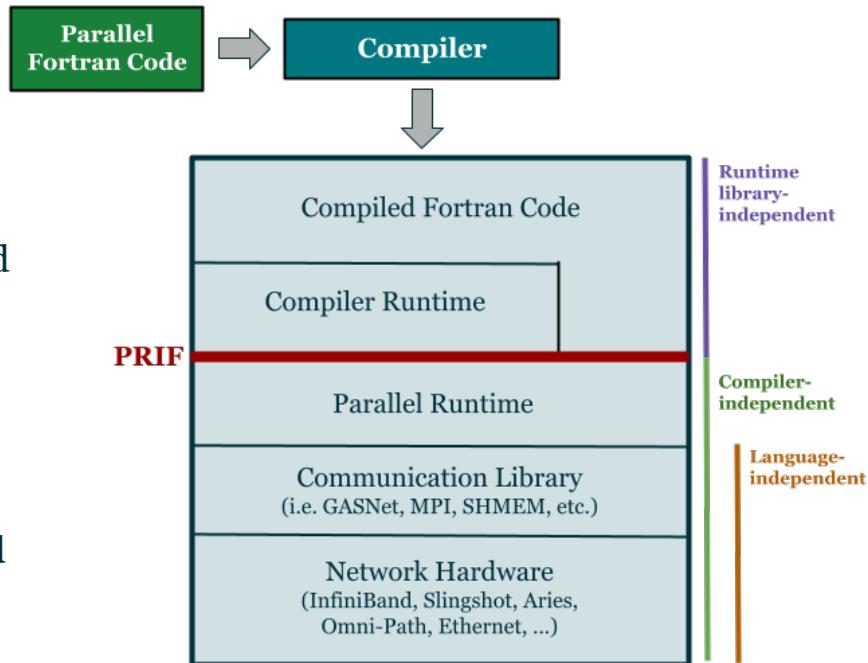
Non\_overridable attribute prevents dynamic dispatch, thereby facilitating *future* GPU execution.

<https://doi.org/10.25344/S4VG6T>



# Parallel Runtime Interface for Fortran (PRIF)<sup>32</sup>

- Compiler- and runtime-agnostic interface to support multi-image parallel Fortran features
- A runtime interface written *in Fortran*: `prif` module
- Tight correspondence between PRIF procedures and Fortran's multi-image parallel features, e.g.,
  - `num_images` → `prif_num_images`
  - `real x(N)[*]` → `prif_allocate_coarray`
- For more information, please see [go.lbl.gov/prif](http://go.lbl.gov/prif) and [fortran.lbl.gov](http://fortran.lbl.gov).



D. Bonachea, K. Rasmussen, B. Richardson, D. Rouson, "Parallel Runtime Interface for Fortran (PRIF): A Multi-Image Solution for LLVM Flang", *Tenth Workshop on the LLVM Compiler Infrastructure in HPC (LLVM-HPC2024)*, Nov. 2024. [doi:10.25344/S4No17](https://doi.org/10.25344/S4No17).



# LLVM-HPC Workshop at SC25 Paper

- Paper highlights the increased LLVM Flang compiler support for Fortran's multi-image features, a subset of which has now been upstreamed, thanks to the support of the NERSC/CLaSS collaboration
- Perlmutter runs in distributed memory show LLVM Flang is comparable with Cray's long extant multi-image Fortran support
- Cray ftn compiler bug encountered while compiling a coarray benchmark on Perlmutter: [NERSC ticket INC0241058](#)

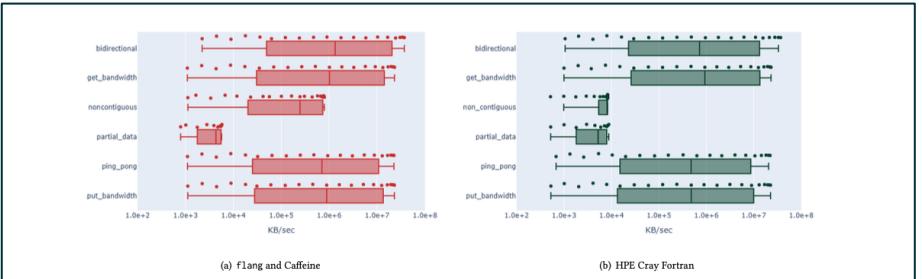


Figure 5: Results from running Caf-testsuite in distributed memory across two nodes of the NERSC Perlmutter supercomputer, connected over an HPE Slingshot 11 network

Dan Bonachea, Katherine Rasmussen, Damian Rouson, Jean-Didier Pailleux, Etienne Renault, Brad Richardson. "Lowering and Runtime Support for Fortran's Multi-Image Parallel Features using LLVM Flang, PRIF, and Caffeine", Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC Workshops '25), November 16–21, 2025, St Louis, MO, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.25344/S4G883>

The Eleventh Annual Workshop on the LLVM Compiler Infrastructure in HPC (LLVM-HPC '25)

## Lowering and Runtime Support for Fortran's Multi-Image Parallel Features using LLVM Flang, PRIF, and Caffeine

Dan Bonachea Katherine Rasmussen Damian Rouson Jean-Didier Pailleux Etienne Renault Brad Richardson   
Dan Bonachea, Katherine Rasmussen, Damian Rouson, Jean-Didier Pailleux, Etienne Renault, and Brad Richardson. Lowering and Runtime Support for Fortran's Multi-Image Parallel Features using LLVM Flang, PRIF, and Caffeine. In Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC Workshops '25), November 16–21, 2025, St Louis, MO, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.25344/S4G883>

Brad Richardson   
Brad.richardson@nasa.gov   
Ametum Services, Inc.   
Houston, Texas, USA

### ABSTRACT

This paper provides an overview of the multi-image parallel features in Fortran 2023 and their implementation in the LLVM Flang compiler and the Caffeine parallel runtime library. The features of interest support a Single-Program, Multiple-Data (SPMD) programming model based on creating multiple "images", each with its own memory space. The features also support a shared Global Address Space (PGAS) in the form of "coarray" distributed data structures. The paper discusses the lowering of multi-image features to the Fortran 2023 standard (via PRIF) and the implementation of PRIF in the Caffeine parallel runtime library. This paper also provides an early view into the design of a new multi-image dialect of the LLVM Multi-Level Intermediate Representation (MLIR). We describe validation and testing of the resulting software, and demonstrate that performance is comparable, if not better, than the Cray compiler and runtime library, GNU Compiler Collection (GCC) `ftn90` and OpenCoarrays, respectively. Performance of LLVM Flang in the multi-image dialect is still missing in flang.

### CCS CONCEPTS

• Software and its engineering → Runtime environments; Parallel programming languages; Computing methodologies → Parallel programming languages.

### KEYWORDS

Fortran, Parallel programming, HPC, PGAS, RMA, LLVM Flang, Exascale Computing, Runtime Libraries, Caffeine, GASNet-EX

### ACM Reference Format:

Dan Bonachea, Katherine Rasmussen, Damian Rouson, Jean-Didier Pailleux, Etienne Renault, and Brad Richardson. 2025. Lowering and Runtime Support for Fortran's Multi-Image Parallel Features using LLVM Flang, PRIF, and Caffeine. In Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC Workshops '25), November 16–21, 2025, St Louis, MO, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.25344/S4G883>

LLVM-HPC '25, November 17, 2025, St Louis, MO, USA  
© 2025. Copyright held by the author or author's institution.

This is the author's version of the work. It is posted here for your personal use. Not for distribution outside of the author's institution. The final version of record was published in Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC Workshops '25), November 16–21, 2025, St Louis, MO, USA. ACM, https://doi.org/10.25344/S4G883

Use the term multi-image in this paper to manipulate coarray features (allocation, deallocation, access... – as well as all other multi-image parallel features (synchronization, collective reductions, atomic operations, image teams, etc.)

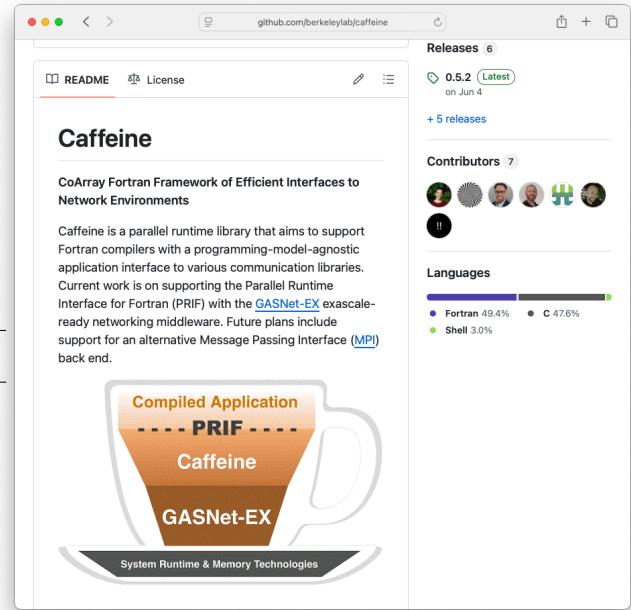
©2025 LBNL. doi: 10.25344/S4G883

1

# Caffeine: Co-Array Fortran Framework of Efficient Interfaces to Network Environments

- Caffeine is written mostly in (serial) Fortran
- Invokes GASNet-EX communication library
- PRIF implementation status:  
[go.lbl.gov/caffeine-status](http://go.lbl.gov/caffeine-status)

Multi-image Fortran Feature	Status
Program startup and shutdown (incl. normal and error termination): STOP, ERROR STOP, END PROGRAM statements	yes
Collective subroutines: CO_{BROADCAST, SUM, MIN, MAX, REDUCE}	yes
Image queries: THIS_IMAGE, NUM_IMAGES, etc, intrinsic functions	yes
Synchronization: SYNC {ALL, IMAGES, MEMORY, TEAM} statements	yes
Storage management: Coarray allocation, deallocation and coarray aliases	yes
Coarray Queries: LCOBOUND, UCOBOUND, COSHAPE, etc.	yes
Contiguous and strided coarray access: Coarray puts and gets	yes
Teams: TEAM_TYPE intrinsic type and {FORM, CHANGE, END} TEAM statements	yes
Events: EVENT_TYPE intrinsic type, EVENT_QUERY subroutine and EVENT {POST, WAIT} statements	yes
Notifications: NOTIFY_TYPE intrinsic type and NOTIFY WAIT statement	yes
Atomics: ATOMIC_{INT, LOGICAL}_KIND kind parameters and ATOMIC_{DEFINE, REF, ...} subroutines	yes
Critical construct: CRITICAL and END CRITICAL	no
Locks: LOCK and UNLOCK statements	no
FAIL IMAGE statement	no



[go.lbl.gov/caffeine](http://go.lbl.gov/caffeine)



# Fortran Package Manager (fpm)



fpm  
build system  
(actual size)



# Overview

---

01

**Background:**  
Deep origins

02

**Features & Paradigms:**  
A walking tour from Fortran 90 to 202Y+

03

**Use Cases:**  
Nooks, Crannies, and Pastures

04

**Conclusions**





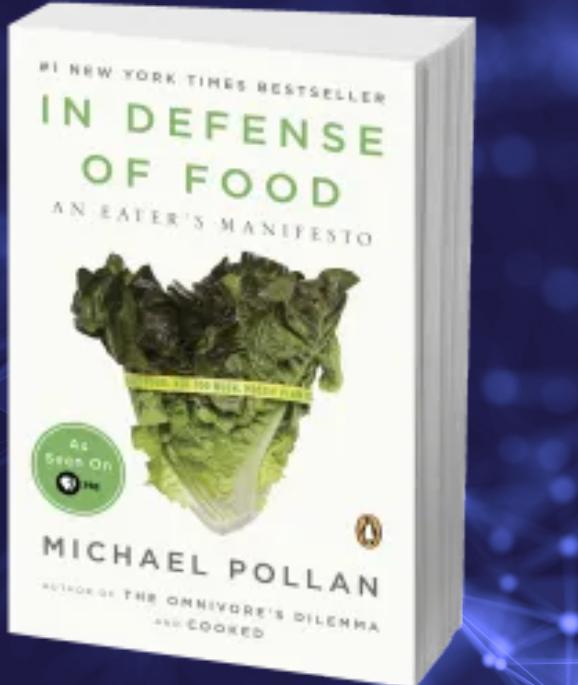
1961



“Fortran is a new and exciting language used by programmers to communicate with computers. It is exciting as it is the wave of the future.”

Character of Dorothy Vaughan,  
a NASA mathematician and programmer,  
as played by Octavia Spencer in  
*Hidden Figures* (20th Century Fox, 2016).

# In Defense of Food: An Eater's Manifesto



*Eat food.*

*Not too much.*

*Mostly plants.*

A dark blue background featuring a glowing, semi-transparent network of white dots and lines, representing a complex system or data structure.

# In Defense of Software: A Developer's Manifesto

*Write software.*

*Not too much.*

*Mostly pure functions.*

# Conclusions

“Fortran [2023] is a new and exciting language used by programmers to communicate with [each other].”

Several underutilized feature sets facilitate writing

- **Compact code:**
  - 46-fold savings in supporting binary operators with elemental functions
  - State-of-the-art neural network training in fewer than 100 lines of code
- **Parallel programs:**
  - Multi-image execution for SPMD/PGAS programming
  - Automatic loop-level multithreading or offloading to a GPU
- **Functional programming patterns:**
  - Pure procedures
  - Immutable state: associate construct
- **Expressive abstractions:**
  - Natural language idioms
  - Textbook forms of partial differential equations

# Thank You!

<https://go.lbl.gov/fortran>