

IRPF90 : a Fortran code generator for HPC

Anthony Scemama¹ <scemama@irsamc.ups-tlse.fr>

François Colonna²

¹ Laboratoire de Chimie et Physique Quantiques
IRSAMC (Toulouse)

² Laboratoire de Chimie Théorique (Paris VI)



Introduction

- Scientific codes need *speed* -> Fortran/C
- Low level language -> difficult to maintain
- High-level features of Fortran 95 or C++ can kill the efficiency (pointers, array syntax, objects, STL, etc) -> not a good solution for HPC

We need to hide the code complexity and keep the code efficient :

1. Implicit Reference to Parameters programming strategy
2. IRPF90 : Facilitates programming with IRP in Fortran

What is a scientific code?

A program is a function of its input data:

```
output = program (input)
```

A program can be represented as a **production tree** where

- The root is the output
- The leaves are the input data
- The nodes are the intermediate variables
- The edges represent the relation *needs/needed by*

Example:

```
u(x,y) = x + y + 1
v(x,y) = x + y + 2
w(x)   = x + 3
t(x,y) = x + y + 4
```


What is the production tree of $t(u(d_1, d_2), v(u(d_3, d_4), w(d_5)))$?

$$u(x, y) = x + y + 1$$

$$v(x, y) = x + y + 2$$

$$w(x) = x + 3$$

$$t(x, y) = x + y + 4$$



Traditional Fortran implementation

```
program example_t
implicit none
integer :: d1, d2, d3, d4, d5
integer :: u1, u2, v, w, t

call read_data(d1,d2,d3,d4,d5)          !      t
                                              !      /
                                              !      u1      v
call compute_u(d1,d2,u1)                 !      /      |
call compute_u(d3,d4,u2)                 !      /      |      /
call compute_w(d5,w)                     !      d1      d2      u2      w
call compute_v(u2,w,v)                   !      /      \      \
call compute_t(u1,v,t)                   !      d3      d4      d5

write(*,*), "t=", t
end program
```

Difficulties

The subroutines need to be called in **the correct order**:

- The programmers needs have the **global knowledge** of the production tree : Production trees are usually too complex to be handled by humans
- Programmers may not be sure that their modification did not break some other part
- Collaborative work is difficult : any user can alter the production tree

Using the functional paradigm

```
program compute_t
    implicit none
    integer :: d1, d2, d3, d4, d5
    integer :: u, v, w, t
    call read_data(d1,d2,d3,d4,d5)
    write(*,*), "t=", t( u(d1,d2), v( u(d3,d4), w(d5) ) )
end program
```

```
!      t
!      /   \
!      u1   v
!      /   /
!      d1   d2   u2   w
!           /   \
!           d3   d4   d5
```

- Instead of telling the machine **what to do**, we express **what we want**
- The production tree is now explored from the root to the leaves.
- The programmer doesn't handle the execution sequence

From global to local knowledge

For each node, we can express the needed entities

```
* t -- needs --> u1 and v
* u1 -- needs --> d1 and d2
* v -- needs --> u2 and w
* u2 -- needs --> d3 and d4
* w -- needs --> d5
```

In this way, all the knowledge is **local**, and much easier to handle by the programmer.

Let's write our program in this way:

```
program compute_t
  implicit none
  integer, external :: t
  write(*,*), "t=", t()
end program
```

```
integer function t()
  implicit none
  integer, external :: u1, v
  t = u1() + v() + 4
end
```

```

integer function v()
    implicit none
    integer, external :: u2, w
    v = u2() + w() + 2
end

integer function w()
    implicit none
    integer :: d1,d2,d3,d4,d5
    call read_data(d1,d2,d3,d4,d5)
    w = d5+3
end

integer function f_u(x,y)
    implicit none
    integer, intent(in) :: x,y
    f_u = x+y+1
end

```

```

integer function u1()
    implicit none
    integer :: d1,d2,d3,d4,d5
    integer, external :: f_u
    call read_data(d1,d2,d3,d4,d5)
    u1 = f_u(d1,d2)
end

integer function u2()
    implicit none
    integer :: d1,d2,d3,d4,d5
    integer, external :: f_u
    call read_data(d1,d2,d3,d4,d5)
    u2 = f_u(d3,d4)
end

```

- Problem : The same data will be recomputed multiple times.
- Solution : memo functions

Implicit Reference to Parameters programming strategy

1. Each entity has only one builder : a subroutine that builds a *valid* value of an entity

```
subroutine build_t(x,y,result)
    implicit none
    integer, intent(in) :: x, y
    integer, intent(out) :: result
    result = x + y + 4
end subroutine build_t

subroutine build_w(x,result)
    implicit none
    integer, intent(in) :: x
    integer, intent(out) :: result
```

```
    result = x + 3
end subroutine build_w

subroutine build_v(x,y,result)
    implicit none
    integer, intent(in) :: x, y
    integer, intent(out) :: result
    result = x + y + 2
end subroutine build_v

subroutine build_u(x,y,result)
    implicit none
    integer, intent(in) :: x, y
    integer, intent(out) :: result
    result = x + y + 1
end subroutine build_u

subroutine build_d(d1,d2,d3,d4,d5)
```

```
implicit none
integer, intent(out) :: d1,d2,d3,d4,d5
read(*,*) d1,d2,d3,d4,d5
end
```

2. Each entity has only one provider : a subroutine with no input arguments whose role is to prepare a *valid* value of an entity.

```
module nodes

! Nodes
integer :: u1
logical :: u1_is_built = .False.

integer :: u2
logical :: u2_is_built = .False.

integer :: v
```

```
logical :: v_is_built = .False.  
  
integer :: w  
logical :: w_is_built = .False.  
  
integer :: t  
logical :: t_is_built = .False.  
  
! Leaves  
integer :: d1, d2, d3, d4, d5  
logical :: d_is_built = .False.  
  
end module
```

```
subroutine provide_t
  use nodes
  implicit none
  if (.not.t_is_built) then
    call provide_u1
    call provide_v
    call build_t(u1,v,t)
    t_is_built = .True.
  endif
end subroutine provide_t
```

```
subroutine provide_w
  use nodes
  implicit none
  if (.not. w_is_built) then
    call provide_d
    call build_w(d5,w)
    w_is_built = .True.
```

```
    endif
end subroutine provide_w

subroutine provide_v
    use nodes
    implicit none
    if (.not. v_is_built) then
        call provide_u2
        call provide_w
        call build_v(u2,w,v)
        v_is_built = .True.
    endif
end subroutine provide_v

subroutine provide_u1
    use nodes
    implicit none
    if (.not. u1_is_built) then
```

```
    call provide_d
    call build_u(d1,d2,u1)
    u1_is_built = .True.
  endif
end subroutine provide_u1

subroutine provide_u2
  use nodes
  implicit none
  if (.not. u2_is_built) then
    call provide_d
    call build_u(d3,d4,u2)
  endif
end subroutine provide_u2

subroutine provide_d
  use nodes
  implicit none
```

```
if (.not. d_is_built) then
    call build_d(d1,d2,d3,d4,d5)
    d_is_built = .True.
endif
end
```

3. Calling a provider always *guarantees* that the entity of interest is **valid** after the provider has been called

The main program is simply:

```
program test
use nodes
implicit none
call provide_t
print *, "t=", t
end program
```

Summary

With the IRP method:

- Code is easy to develop for a new developer : Adding a new feature only requires to know the names of the needed entities
- If one developer changes the dependence tree, the others will not be affected : collaborative work is simple
- Forces to write clear code : one builder builds only one thing
- Forces to write efficient code : temporal locality is good, as in cache oblivious algorithms

But in real life:

- A lot of typing is required
- Programmers are lazy

IRPF90

- Code generator that will write all the IRP glue code for you
- Fortran with additional keywords
- Extends fortran to add very useful features :
 - Automatic makefile generation
 - Text editor integration
 - Some Introspection
 - Meta programming
 - Many more interesting things
- No problem using external libraries (MKL, MPI, etc)
- Generated code is **very** efficient : sustained 960 Tflops/s on Curie in 2011 with QMC=Chem (12 GFlops/s / core)

```
BEGIN_PROVIDER [ integer , t ]
```

```
    t = u1+v+4
```

```
END_PROVIDER
```

```
BEGIN_PROVIDER [ integer , w ]
```

```
    w = d5+3
```

```
END_PROVIDER
```

```
BEGIN_PROVIDER [ integer , v ]
```

```
    v = u2+w+2
```

```
END_PROVIDER
```

```
BEGIN_PROVIDER [ integer , u1 ]
```

```
    integer :: fu
```

```
    u1 = fu(d1,d2)
```

```
END_PROVIDER
```

```
BEGIN_PROVIDER [ integer , u2 ]
```

```
integer :: fu
u2 = fu(d3,d4)
END_PROVIDER
```

```
integer function fu(x,y)
integer, intent(in) :: x,y
fu = x+y+1
end function
```

```
program irp_example
print *, 't=' , t
end
```

When you write a provider for *x*, you **only** have to focus on

- How do I build *x*?
- What are the variables that I need to build *x*?
- Am I sure that *x* is built correctly when I exit the provider?

Features

Declaring variables

Variables can be declared anywhere. It allows this:

```
integer :: i
do i=1,N
  a(i) = b(i) + c(i)
end do
```

```
integer :: j
do j=1,M
  c(j) += d(j)
end do
```

Arrays

```
BEGIN_PROVIDER [ double precision, A, (dim1, 3) ]  
    ...  
END_PROVIDER
```

- Allocation of IRP arrays done automatically
- Dimensioning variables can be IRP entities, provided before the memory allocation
- FREE keyword to force to free memory. Invalidates the entity.

Example:

```
BEGIN_PROVIDER [ integer, fact_max ]
    fact_max = 10
END_PROVIDER

BEGIN_PROVIDER [ double precision, fact, (0:fact_max) ]
    implicit none
    integer :: i
    fact(0) = 1.d0
    do i=1,fact_max
        fact(i) = fact(i-1)*dbl(e(i))
    end do
END_PROVIDER
```

Assertions

Assertion

Boolean expression that **must** be true at run time

Example:

```
BEGIN_PROVIDER [ integer, u2 ]
    call compute_u(d3,d4,u2)
    ASSERT (u2 < d3)
END_PROVIDER
```

```
Stack trace:          0
-----
provide_t
provide_v
provide_u2
u2
-----
u2: Assert failed:
file: uvwt.irp.f, line: 23
(u2 < d3)
u2 =          8
d3 =          3
STOP 1
```

Documentation

Every subroutine/function/provider should have a documentation section:

```
BEGIN_PROVIDER [ double precision, Fock_matrix_beta_mo, (mo_tot_num_align,mo_tot_num) ]
  implicit none
  BEGIN_DOC
    ! Fock matrix on the MO basis
  END_DOC
  ...
END_PROVIDER
```

```
$ irpman fock_matrix_beta_mo
```

IRPF90 entities(1)	fock_matrix_beta_mo	IRPF90 entities(1)
Declaration		
double precision, allocatable :: fock_matrix_beta_mo (mo_tot_num_align,mo_tot_num)		
Description		
Fock matrix on the MO basis		
File		
Fock_matrix.irp.f		
Needs		
ao_num fock_matrix_alpha_ao mo_coef mo_tot_num mo_tot_num_align		
Needed by		
fock_matrix_mo		
IRPF90 entities	fock_matrix_beta_mo	IRPF90 entities(1)

Templates

Avoids duplicating similar code (but not exactly the same)

```
BEGIN_TEMPLATE
```

```
BEGIN_PROVIDER [ $type , $name ]
  call find_in_input( '$name' , $name )
END_PROVIDER
```

```
logical function $name_is_zero()
  $name_is_zero = ($name == 0)
end function
```

```
SUBST [ type , name ]
```

```
integer      ;    size_tab1  ;;
integer      ;    size_tab2  ;;
real        ;    distance   ;;
```

```
real      ; x      ; ;  
real      ; y      ; ;  
real      ; z      ; ;
```

```
END_TEMPLATE
```

Augmented assignment operators

Operators : `+=`, `-=`, `*=`

```
my_very_explicit_name(dim1, dim2, dim3) = &
    my_very_explicit_name(dim1, dim2, dim3) &
    + b*c - d

my_very_explicit_name(dim1, dim2, dim3) += b*c - d
```

- More robust to typos
- More robust to changing `dim1`...`dim3`
- More readable

Embedding scripts

- Info at compile time
- Specific formulas (ex: fast power functions)

```
BEGIN_SHELL [ /bin/bash ]
    echo print *, `Compiled by `whoami` on `date`\`"
END_SHELL
```

```
BEGIN_SHELL [ /usr/bin/python ]
for i in range(100):
    print """
        double precision function times_%d(x)
            double precision, intent(in) :: x
            times_%d = x**%d
        end
    """%locals()
END_SHELL
```

Conditional compilation

It is not possible to use the C preprocessor with IRPF90 files. Instead, we can do

```
IRP_IF HAS_MPI
```

```
  subroutine abort(message)
    call mpi_abort(message)
  end
```


```
IRP_ELSE
```

```
  subroutine abort(message)
    print *, message
    stop
  end
```


```
IRP_ENDIF
```

Iterative processes


Iterative processes may involve cyclic dependencies:




TOUCH A : A is valid, but everything that needs A is invalidated



(a)



(b)



(c)

Debugging

Display tree

The debug option -d activates the printing of the tree:

```
$ ./irp_example1
0 : -> provide_t
0 :   -> provide_u1
0 :     -> provide_d1
0 :       -> d1
d1
1
d2
2
d3
3
d4
```

```
4
d5
5
0 :      <- d1  6.88999999999999E-004
0 :      <- provide_d1  8.07000000000000E-004
0 :      -> u1
0 :      -> fu
0 :      <- fu  9.99999999996990E-007
0 :      <- u1  2.9000000000038E-005
0 :      <- provide_u1  8.99999999999998E-004
0 :      -> provide_v
0 :      -> provide_w
0 :      -> w
0 :      <- w  1.00000000000133E-006
0 :      <- provide_w  2.70000000000011E-005
0 :      -> provide_u2
0 :      -> u2
0 :      -> fu
```

```
0 :      <- fu  1.00000000000133E-006
0 :      <- u2  2.01500000000000E-003
0 :      <- provide_u2  2.02600000000000E-003
0 :      -> v
0 :      <- v  0.00000000000000E+000
0 :      <- provide_v  2.08900000000000E-003
0 :      -> t
0 :      <- t  0.00000000000000E+000
0 :      <- provide_t  3.03300000000000E-003
0 :      -> irp_example1
t =          26
0 : <- irp_example1  0.00000000000000E+000
```

Fortran compiler errors

All the generated code has a comment to point to the irpf90 file:

```
$ ifort -I IRPF90_temp/ -O2 -c IRPF90_temp/uvwt.irp.F90 -o IRPF90_temp/uvwt.irp.o
IRPF90_temp/uvwt.irp.F90(71): error #5082: Syntax error, found END-OF-STATEMENT
when expecting one of: ( * ) :: , . % + - [ : . ** / // .LT. < .LE. <= .EQ. ==
...
    u2 = fu(d3,d4)                                ! uvwt.irp.f: 40
-----
                                         ^
compilation aborted for IRPF90_temp/uvwt.irp.F90 (code 1)
make: *** [IRPF90_temp/uvwt.irp.o] Error 1
```

Getting the name of the current subroutine

All subroutines / functions / providers have a character variable `irp_here`:

```
print *, irp_here//': a = ', a
```

Tracing memory allocations

-m command line option.

```
...
    10 Allocating ci_electronic_energy(N_states_diag)
    10 Allocating ci_eigenvectors(N_det,N_states_diag)
    10 Allocating ci_eigenvectors_s2(N_states_diag)
128260 Allocating psi_det(N_int,2,psi_det_size)
...
Deallocating ci_eigenvectors
    30600 Allocating ci_eigenvectors(N_det,N_states_diag)
    6120 Allocating det_connections(N_con_int,N_det)
...
```

Debugging an embedded script

--preprocess command line option.

```
program test
  BEGIN_SHELL [ /bin/bash ]
cat << EOF | sed 's/\\(.*)\\)/echo "\1"/g'
    print *, 'Compiled by `whoami` on `date`'
    print *, '$PWD'
    print *, '$(hostname)'
EOF
  END_SHELL
end
```

```
$ irpf90 --preprocess test.irp.f
program irp_program
call test
call irp_finalize_1624498827()
end program
subroutine test
character*(4) :: irp_here = 'test'
    print *, 'Compiled by scemama on Mon Mar  9 18:41:50 CET 2015'
    print *, '/dev/shm/tmp'
    print *, 'lpqlx139'
end
```

Debugging TOUCH statements

-t command line option displays the entities that will be invalidated by a touch.

```
$ irpf90 -t psi_coef
Touching psi_coef invalidates the following entities:
- ci_electronic_energy
- ci_energy
- coef_hf_selector
- exc_degree_per_selectors
- n_det_generators
- n_det_selectors
- one_body_dm_mo
- psi_average_norm_contrib
- psi_det_sorted
- psi_generators
- psi_selectors
- s2_values
```

IRPF90 for HPC

In this section, it is recommended to use the Intel Fortran compiler (ifort).

Array alignment

- Vector instructions (ADD/MUL/LOAD/STORE/ etc) operate on **aligned** data.
- SSE : 16 bytes, AVX/AVX2 : 32 bytes, AVX512 : 64 bytes.
- If we can easily align data -> performance gain
 - Array : !DIR\$ ATTRIBUTES ALIGN : 32 :: A
 - Loop : !DIR\$ VECTOR ALIGNED
- For an *aligned* multi-dimensional array, all columns are aligned *if* the LDA is a multiple of the alignment

Using the --align <n> option, IRPF90 can introduce compiler directives for ifort such that *all* the IRP arrays are *n*-byte aligned. The \$IRP_ALIGN variable corresponds *n*.

```
integer function align_double(i)
  integer, intent(in) :: i
  integer :: j
  j = mod(i,max($IRP_ALIGN,4)/4)
  if (j==0) then
    align_double = i
  else
    align_double = i+4-j
  endif
end
```

```
BEGIN_PROVIDER [ integer, n ]
&BEGIN_PROVIDER [ integer, n_aligned ]
  integer :: align_double
  n = 19
  n_aligned = align_double(19)
END_PROVIDER
```

```
BEGIN_PROVIDER [ double precision, Matrix, (n_aligned,n) ]
    Matrix = 0.d0
END_PROVIDER
```

- All IRP entities are aligned
- All columns of array Matrix are aligned
- -> We can happily use !DIR\$ Vector aligned

Variable substitutions

Create a binary targeted for a given input :

```
if (choice1) then
  !DIR$ VECTOR ALIGNED
  do i=1,lmax
    call do_stuff
  enddo
else
  !DIR$ VECTOR ALIGNED
  do i=1,nmax
    call do_something_else
  enddo
endif
```

```
irpf90 --align=32 -s lmax:100 -s nmax:48 -s choice1:.True.
```

```
if (.True.) then
  !DIR$ VECTOR ALIGNED
  do i=1,100
    call do_stuff
  enddo
else
  !DIR$ VECTOR ALIGNED
  do i=1,48
    call do_something_else
  enddo
endif
```

!

! Compiler knows

! what is the best

! optimization

!

!

! Dead code

! removed by

! the compiler

Profiling

--profile command line option.

Time measures using hardware counter rdtsc:

	N.Calls	Tot Cycles	Avg Cycles	Tot Secs	Avg Secs

...					
ci_energy	6.	1662765.	277127.+/- 59043.	0.00072451	0.00012+/- 0.00003
coef_hf_selector	7.	13009101.	1858443.+/- 605660.	0.00566841	0.00081+/- 0.00026
davidson_criterion	1.	1736.	1736.+/- 0.	0.00000076	0.00000+/- 0.00000
davidson_sze_max	1.	18.	18.+/- 0.	0.00000001	0.00000+/- 0.00000
det_connections	1.	6945057.	6945057.+/- 0.	0.00302614	0.00303+/- 0.00000
diag_algorithm	6.	15253.	2542.+/- 246.	0.00000665	0.00000+/- 0.00000
do_pt2_end	1.	233928.	233928.+/- 0.	0.00010193	0.00010+/- 0.00000
elec_alpha_num	1.	751170.	751170.+/- 0.	0.00032730	0.00033+/- 0.00000
exc_degree_per_selectors	7.	209402.	29915.+/- 10827.	0.00009124	0.00001+/- 0.00000
expected_s2	1.	240961.	240961.+/- 0.	0.00010499	0.00010+/- 0.00000
ezfio_filename	1.	386883.	386883.+/- 0.	0.00016858	0.00017+/- 0.00000

Codelet generation

--codelet command line option.

Creates a new program to time a given provider.

Example:

```
$ irpf90 --codelet v:t:100000
```

Generates a codelet to time v. t has to be provided first and the measurement will be done 100000 times.

OpenMP

- OpenMP can be used as in Fortran.
- Use `--openmp` to make providers thread-safe

CoArray Fortran

`--coarray` : all providers are co-arrays

Example:

```
program caf_test
implicit none

PROVIDE X
if (image_id == 1) then
  print *, 'This image:'
  print *, X
  print *, 'Image 2:'
  print *, X[2]
```

endif

end

More info

Source on GitHub

<https://github.com/scemama/irpf90>

GitBook

<http://scemama.gitbooks.io/irpf90/>

Web page

<http://irpf90.ups-tlse.fr>

Quantum Package : *Quantum Chemistry (OpenMP)*

https://github.com/LCPQ/quantum_package

QMC=Chem : *Quantum Monte Carlo (ZeroMQ)*

<http://qmcchem.ups-tlse.fr>

EPLF : *Electron pair localization function (MPI)*

<http://eplf.sourceforge.net>

EZFIO : *Easy Fortran I/O library generator*

<https://github.com/scemama/ezfio>