

41391 High performance computing: Miscellaneous parallel programmes in Fortran

Nilas Mandrup Hansen, Ask Hjorth Larsen

January 19, 2010

1 Introduction

This document concerns the implementation of a Fortran programme capable of calculating digits of pi, one which can do matrix-vector multiplication, and one which calculates the Mandelbrot set. The programmes are parallelized using the automatic parallelization (which the `f90` compiler supports through the `-xautopar` option) and/or explicitly with OpenMP.

2 Calculating digits of pi

It is known that

$$\pi = \int_0^1 \frac{4}{1+x^2} dx \approx \frac{4}{N} \sum_{i=1}^{\infty} \left[1 + \left(\frac{i - \frac{1}{2}}{N} \right)^2 \right]^{-1}. \quad (1)$$

The terms of this sum are independent of each other, and the calculation therefore embarrassingly parallel.

With OpenMP we implement this calculation using the following loop:

```
!$omp parallel do shared(n) private(i) reduction(+: pi)
  do i=1, n
    pi = pi + 1.0 / (1.0 + ((i - 0.5) / n)**2)
  end do
!$omp end parallel do
```

The exact same code is used with automatic parallelization, with the exception of the parameters `-xautopar` `-xreduction` instead of `-xopenmp` in the makefile.

Figure 1 shows the speedup as a function of processor count for the π calculation code using OpenMP as well as automatic parallelization, which seem to differ little in terms of performance. With a small number of iterations (2^{20}) the parallelization is, unsurprisingly, less efficient than with a larger number (2^{24})

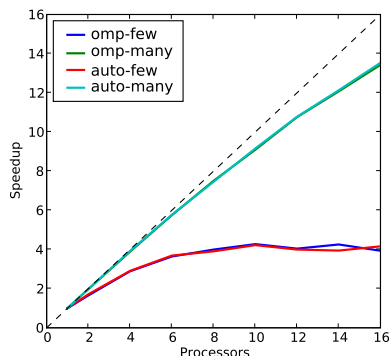


Figure 1: Speedup for parallel calculation of π with automatic parallelization and OpenMP when calculating with few or many iterations. The total time is on the order of milliseconds when doing only few iterations, which means process startup overhead prevents perfect scaling of this otherwise embarrassingly parallel operation.

(both iteration counts are rather large given that the calculated value is only a double precision number, but the high iteration counts make timings more reliable).

The source files can be found in the appendix.

3 Matrix times vector multiplication

The product y of a matrix A and a vector x has the elements

$$y_i = \sum_j A_{ij}x_j. \quad (2)$$

This we implement in terms of an outer loop over i and an inner loop over j , where we parallelize the outer loop.

Since Fortran uses column-major ordering, and since we don't want to worry too much about transposes and definitions, we use the indexing scheme $A(i*N + j)$ rather than actual two-dimensional arrays.

With OpenMP and automatic parallelization we use the same loop:

```
!$omp parallel do shared(A, x, y, M, N) private(tmp, i, j)
do i = 1, M
  tmp = 0.0
  do j = 1, N
    tmp = tmp + A(j + N * i) * x(j)
  end do
  y(i) = tmp
```

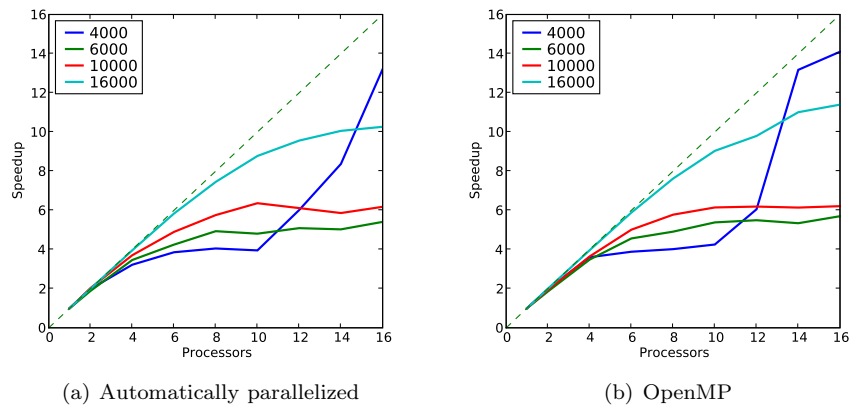


Figure 2: Speedup of matrix times vector function using `-xautopar` (a) and explicit OpenMP (b).

```

end do
!$omp end parallel do

```

Each processor writes to different areas of the destination matrix, so no synchronization is required.

Figure 2 shows the speedup of the matrix times vector function with `-xautopar` and OpenMP for different matrix sizes and processor counts. With a 4000×4000 matrix, most or all of the matrix fits into the L2 cache when using more than 10 CPUs (7.6 MB for 16 CPUs with a cache size of 8 MB), which results in a dramatic speedup due to fewer cache misses.

The larger matrices will not fit into the cache (the 6000×6000 uses 17 MB on 16 CPUs), and thus behave in a more usual manner: Larger matrices parallelize better as the CPUs have more work compared to the (constant) amount of synchronization. We think the imperfect scaling observed for very large matrices is caused by having too many processes read data from the main memory.

We note that smaller matrices (of size around 500-2000) tend to scale superlinearly as they start fitting better into L1 and L2 caches (not shown in the figure).

Source files can be found in the appendix.

4 Mandelbrot

For this exercise, the Mandelbrot set is computed and visualized by using the provided `mandelbrot.zip` file, which contains the following FORTRAN files;

- `main.f90`: The main program for Mandelbrot.

- `mandel.f90`: Contains the Mandelbrot calculations.
- `timestamp.f90`: Simply prints the current YMDHMS date as a time stamp.

Also provided is a `Makefile` as well as several other files which are not important to mention.

For the first part of the exercise, the Mandelbrot program is simply to be compiled by the `gmake` command in order to produce a serial version of Mandelbrot. The output of the executable is shown in fig. 3




Figure 3: The Mandelbrot set

4.1 Parallelizing the Code

For the second part of the exercise, the `Mandelbrot` executable should be generated as a parallel version by using the OpenMP worksharing constructs. Most of the resources are used in the triple do-loop calculating the Mandelbrot set. In order to parallelize the code, a team of threads is created by using the command `!$parallel do`. Also the loop counters are set to private, while all other variables are set to shared (see source code in app. C).

When compiling the code with the mentioned modifications, one obtains a parallel version of the executable `mandelbrot`. The runtimes for a different number of threads is shown in Fig. 4.

It can be seen from the figure that the code does not scale well.




Figure 4: Speed up vs. number of threads. Using static workload distribution.

4.2 Modifications the to the Parallel version of the Mandelbrot

From the preceding section it was seen that the code did not scale well. In order to optimize the parallel version, we profiled the program by utilizing the `collect` command. The result was visualized in the Analyzer program. The result showed that approximately 95% of the workload was distributed to only two threads even though four threads were available.

In order to overcome the problem a dynamic distribution of the workload was specified in the code by specifying the `clause schedule(dynamic,5)` in the parallel `do` environment (see app. D). The result is presented in Fig. 5

From the figure it can be seen that the speed-up factor has improved considerably when using dynamic workload distribution.

4.3 Using Orphaning

As a last part of this exercise, the code should utilize orphaning. This is a useful feature of the OpenMP, since it allows the user to declare workshare or synchronization directives which are not located within a parallel region. This means that the programmer has the possibility to run the same subroutine with or without parallelization.

An example is shown in app. E. Here a flag has been inserted in the main document which allows the user to run the subroutine `mandel` in either serial or parallel computation.

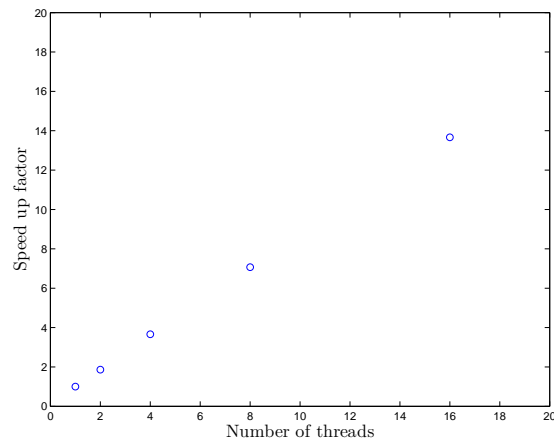


Figure 5: Speed up vs. number of threads. Using dynamic workload distribution.

A Pi calculation source files

A.1 pi/omp/pi.f90

```

1 module picalc
2
3 contains
4
5   subroutine findpi(n, pi)
6     implicit none
7
8     integer, intent(in) :: n
9     double precision, intent(out) :: pi
10    call findpi_ompreduce(n, pi)
11    !call findpi_ourreduce(n, pi)
12    pi = pi * 4.0 / n
13  end subroutine findpi
14
15  subroutine findpi_ompreduce(n, pi)
16    implicit none
17
18    integer, intent(in) :: n
19    double precision, intent(out) :: pi
20
21    integer :: i
22
23    pi = 0.0
24
25    !$omp parallel do shared(n) private(i) reduction(+: pi)
26      do i=1, n
27        pi = pi + 1.0 / (1.0 + ((i - 0.5) / n)**2)
28      end do
29    !$omp end parallel do
30  end subroutine findpi_ompreduce
31
32  subroutine findpi_ourreduce(n, pi)
33    implicit none
34
35    integer, intent(in) :: n

```

```

36     double precision, intent(out) :: pi
37
38     integer :: i
39     double precision :: local_pi
40
41     pi = 0.0
42 !$omp parallel shared(n, pi) private(local_pi)
43     local_pi = 0.0
44 !$omp do private(i)
45     do i=1, n
46         local_pi = local_pi + 1.0 / (1.0 + ((i - 0.5) / n)**2)
47     end do
48 !$omp end do
49
50 !$omp critical (sum)
51     pi = pi + local_pi
52 !$omp end critical (sum)
53 !$omp end parallel
54     end subroutine findpi_ourreduce
55
56     subroutine iterate(i, n, tmp)
57         implicit none
58
59         integer, intent(in) :: i, n
60         double precision, intent(out) :: tmp
61
62         tmp = 1.0 / (1.0 + ((i - 0.5) / n)**2)
63     end subroutine iterate
64
65 end module picalc

```

A.2 pi/omp/main.f90

```

1  program calculate_pi
2  use picalc, only: findpi
3  use omp_lib
4
5  integer :: n! = 2**27
6  double precision :: pi, t1, t2
7
8  read*, n
9
10 t1 = omp_get_wtime()
11 call findpi(n, pi)
12 t2 = omp_get_wtime()
13
14 print*, omp_get_max_threads(), t2 - t1
15
16 end program calculate_pi

```

A.3 pi/omp/makefile

```

1  ARGS=-O3 -openmp
2  AUTOPAR_ARGS = -xautopar -xloopinfo -xreduction
3
4  findpi: pi.o main.f90
5  f90 $(ARGS) pi.o main.f90 -o findpi
6
7  pi.o:
8  f90 -c $(ARGS) pi.f90
9
10 clean:
11 rm -f pi.o picalc.mod

```

B Matrix times vector source files

B.1 mxv/omp/main.f90

```
1 program matrix
2   use m_mxv
3   use omp_lib
4
5   integer :: M, N
6
7   double precision, dimension(:), allocatable :: A
8   double precision, dimension(:), allocatable :: y
9   double precision, dimension(:), allocatable :: x
10  double precision :: tmp
11  integer :: i, j
12  double precision :: time1, time2
13
14  integer :: ops_per_mxv, approx_opcount, number_of_runs, opcount, iter
15
16  read*, M
17
18  N = M
19
20  ops_per_mxv = M * N
21  approx_opcount = 8000 * 8000 * 4 ! probably reasonable in terms of real-
22    time
23  number_of_runs = (approx_opcount / ops_per_mxv) * omp_get_max_threads()
24  opcount = number_of_runs * ops_per_mxv
25
26  allocate(A(N * M), y(M), x(N))
27
28  write(unit=0, *) 'size:', M, '; runs:', number_of_runs, '; ops:', opcount,
29    &
30    '; cpus:', omp_get_max_threads()
31
32  !$omp parallel
33  !$omp do
34    do i=1, M
35      do j=1, N
36        A(j + M * i) = 1.0
37      enddo
38    enddo
39  !$omp end do
40  !$omp do
41    do i=1, N
42      x(i) = 1.0
43    enddo
44  !$omp end do
45  !$omp end parallel
46
47  time1 = omp_get_wtime()
48  do iter=1, number_of_runs
49    call mxv(A, M, N, x, y)
50  enddo
51  time2 = omp_get_wtime()
52  print*, omp_get_max_threads(), ((time2 - time1) / real(number_of_runs))
53 end program matrix
```

B.2 mxv/omp/mxv.f90

```
1 module m_mxv
2
3 contains
4
5   subroutine mxv(A, M, N, x, y)
6     integer, intent(in) :: M, N
```



```

7      double precision, dimension(N * M), intent(in) :: A
8      double precision, dimension(N), intent(in) :: x
9      double precision, dimension(M), intent(out) :: y
10
11     integer :: i, j
12
13     double precision :: tmp
14 !$omp parallel do shared(A, x, y, M, N) private(tmp, i, j)
15     do i = 1, M
16         tmp = 0.0
17         do j = 1, N
18             tmp = tmp + A(j + N * i) * x(j)
19         end do
20         y(i) = tmp
21     end do
22 !$omp end parallel do
23
24     end subroutine mxv
25
26 end module m_mxv

```

B.3 mxv/omp/makefile

```

1  ARGS = -g -fast -xopenmp -xloopinfo
2
3  run: mxv.o main.f90
4      f90 $(ARGS) main.f90 mxv.o -o run
5
6  mxv.o:
7      f90 -c $(ARGS) mxv.f90
8
9  clean:
10     rm -f run mxv.o m_mxv.mod

```

C Source code - mandel.f90

```

1      subroutine mandel(n, image, max_iter)
2
3      integer ( kind = 4 ) :: n
4      integer ( kind = 4 ) :: image(n,n)
5      integer ( kind = 4 ) :: max_iter
6      integer ( kind = 4 ) :: i, j, k
7      real ( kind = 8 ) :: x, x1, x2
8      real ( kind = 8 ) :: y, y1, y2
9      real ( kind = 8 ) :: x_max = 1.25D+00
10     real ( kind = 8 ) :: x_min = - 2.25D+00
11     real ( kind = 8 ) :: y_max = 1.75D+00
12     real ( kind = 8 ) :: y_min = - 1.75D+00
13
14     image = 0.0
15
16 !$omp parallel default(none) shared(n,image,max_iter,x,y,x_min,x_max,y_min,
17     y_max,x1,y1,x2,y2) private(i,j,k)
18 !$omp do
19     do i = 1, n
20         do j = 1, n
21             x = ( real ( j - 1, kind = 8 ) * x_max &
22                 + real ( n - j, kind = 8 ) * x_min ) &
23                 / real ( n - 1, kind = 8 )
24
25             y = ( real ( i - 1, kind = 8 ) * y_max &
26                 + real ( n - i, kind = 8 ) * y_min ) &
27                 / real ( n - 1, kind = 8 )
28

```

```

29      !      image(i,j) = 0
30
31      x1 = x
32      y1 = y
33
34      do k = 1, max_iter
35
36          x2 = x1 * x1 - y1 * y1 + x
37          y2 = 2 * x1 * y1 + y
38
39          if ( x2 < -2.0D+00 .or. &
40              2.0D+00 < x2 .or. &
41              y2 < -2.0D+00 .or. &
42              2.0D+00 < y2 ) then
43
44              image(i,j) = k
45              exit
46
47          end if
48
49          x1 = x2
50          y1 = y2
51
52      end do
53
54  end do
55  end do
56  !$omp end do
57  !$omp end parallel
58  end subroutine

```

D Source code - mandel.f90

```

1      subroutine mandel(n, image, max_iter)
2
3      integer ( kind = 4 ) :: n
4      integer ( kind = 4 ) :: image(n,n)
5      integer ( kind = 4 ) :: max_iter
6      integer ( kind = 4 ) :: i, j, k
7      real ( kind = 8 ) :: x, x1, x2
8      real ( kind = 8 ) :: y, y1, y2
9      real ( kind = 8 ) :: x_max = 1.25D+00
10     real ( kind = 8 ) :: x_min = - 2.25D+00
11     real ( kind = 8 ) :: y_max = 1.75D+00
12     real ( kind = 8 ) :: y_min = - 1.75D+00
13
14     image = 0.0
15
16     !$omp parallel default(none) shared(n,image,max_iter,x,y,x_min,x_max,y_min,
17     y_max,x1,y1,x2,y2) private(i,j,k)
18     !$omp do schedule(dynamic,5)
19     do i = 1, n
20         do j = 1, n
21
22             x = ( real ( j - 1, kind = 8 ) * x_max &
23                 + real ( n - j, kind = 8 ) * x_min ) &
24                 / real ( n - 1, kind = 8 )
25
26             y = ( real ( i - 1, kind = 8 ) * y_max &
27                 + real ( n - i, kind = 8 ) * y_min ) &
28                 / real ( n - 1, kind = 8 )
29
30             !      image(i,j) = 0
31
32             x1 = x
33             y1 = y

```

```

33
34         do k = 1, max_iter
35
36             x2 = x1 * x1 - y1 * y1 + x
37             y2 = 2 * x1 * y1 + y
38
39             if ( x2 < -2.0D+00 .or. &
40                 2.0D+00 < x2 .or. &
41                 y2 < -2.0D+00 .or. &
42                 2.0D+00 < y2 ) then
43
44                 image(i,j) = k
45                 exit
46
47             end if
48
49             x1 = x2
50             y1 = y2
51
52         end do
53
54     end do
55 end do
56 !$omp end do
57 !$omp end parallel
58 end subroutine

```

E Source code - main.f90 and mandel.f90

```

1 program main
2
3 !*****80
4 !
5 ! MAIN is the main program for MANDELBROT.
6 !
7 ! Discussion:
8 !
9 !     MANDELBROT computes an image of the Mandelbrot set.
10 !
11 ! Licensing:
12 !
13 !     This code is distributed under the GNU LGPL license.
14 !
15 ! Modified:
16 !
17 !     08 August 2009
18 !
19 ! Author:
20 !
21 !     John Burkardt
22 !
23 ! Modified by:
24 !     Bernd Dammann
25 !     Boyan Lazarov
26 !
27 ! Local Parameters:
28 !
29 !     Local, integer COUNT_MAX, the maximum number of iterations taken
30 !     for a particular pixel.
31 !
32 implicit none
33
34 integer ( kind = 4 ) :: n = 2501
35 integer ( kind = 4 ) :: count_max = 800
36

```

```

37  integer    ( kind = 4 ) :: c
38  real      ( kind = 8 ) :: c_max, c_max_inv
39  integer    ( kind = 4 ), dimension(:,:), allocatable :: image
40  character  ( len = 255 ) :: filename
41
42  real      ( kind = 8 ) :: x_max = 1.25D+00
43  real      ( kind = 8 ) :: x_min = - 2.25D+00
44  real      ( kind = 8 ) :: y_max = 1.75D+00
45  real      ( kind = 8 ) :: y_min = - 1.75D+00
46  real:: flag
47  flag=10.1
48
49  allocate(image(n,n))
50
51  write ( *, '(a)' ) ' '
52  write ( *, '(a)' ) 'MANDELBROT'
53  write ( *, '(a)' ) ' FORTRAN90 version'
54  write ( *, '(a)' ) ' '
55  write ( *, '(a)' ) ' Create an PNG image of the Mandelbrot set.'
56  write ( *, '(a)' ) ' '
57  write ( *, '(a)' ) ' For each point C = X + i*Y'
58  write ( *, '(a,g14.6,a,g14.6,a)' ) ' with X range [', x_min, ', ', x_max,
59  write ( *, '(a,g14.6,a,g14.6,a)' ) ' and Y range [', y_min, ', ', y_max,
60  write ( *, '(a,i8,a)' ) ' carry out ', count_max, ' iterations of the map'
61  write ( *, '(a)' ) ' Z(n+1) = Z(n)^2 + C.'
62  write ( *, '(a)' ) ' If the iterates stay bounded (norm less than 2)'
63  write ( *, '(a)' ) ' then C is taken to be a member of the set.'
64  write ( *, '(a)' ) ' '
65  write ( *, '(a)' ) ' A PNG image of the set is created using'
66  write ( *, '(a,i8,a)' ) ' N = ', n, ' pixels in the X direction and'
67  write ( *, '(a,i8,a)' ) ' N = ', n, ' pixels in the Y direction.'
68  write ( *, '(a)' ) ' '
69  !
70  ! Carry out the iteration for each pixel, determining COUNT.
71
72  call timestamp ( )
73
74  !! private(i,j,k)
75  if(flag.GT.1) then
76  !$omp parallel default(none) shared(n,image,count_max)
77  call mandel(n,image,count_max)
78  !$omp end parallel
79  else
80  call mandel(n,image,count_max)
81  endif
82
83  write ( *, '(a)' ) ' '
84  write ( *, '(a)' ) ' Calculation of the image finished. '
85  call timestamp ( )
86
87  ! uncomment the following line, if you don't need the PNG output
88  ! stop
89  !
90  ! call writepng to save a PNG image in filename
91
92  filename = "mandelbrot.png"//CHAR(0)
93  call writepng(filename, image, n, n)
94  deallocate(image)
95
96  write ( *, '(a)' ) ' '
97  write ( *, '(a)' ) &
98  ' PNG image data stored in "' // trim ( filename ) // "'. '
99  write ( *, '(a)' ) ' '
100 write ( *, '(a)' ) 'MANDELBROT'
101 write ( *, '(a)' ) ' Normal end of execution.'
102 write ( *, '(a)' ) ' '

```

```

103     call timestamp ( )
104
105     stop
106 end

1     subroutine mandel(n, image, max_iter)
2
3     integer ( kind = 4 ) :: n
4     integer ( kind = 4 ) :: image(n,n)
5     integer ( kind = 4 ) :: max_iter
6     integer ( kind = 4 ) :: i, j, k
7     real ( kind = 8 ) :: x, x1, x2
8     real ( kind = 8 ) :: y, y1, y2
9     real ( kind = 8 ) :: x_max = 1.25D+00
10    real ( kind = 8 ) :: x_min = - 2.25D+00
11    real ( kind = 8 ) :: y_max = 1.75D+00
12    real ( kind = 8 ) :: y_min = - 1.75D+00
13
14    image = 0.0
15
16
17    !$omp do schedule(dynamic,5)
18    do i = 1, n
19    do j = 1, n
20
21    x = ( real ( j - 1, kind = 8 ) * x_max &
22    + real ( n - j, kind = 8 ) * x_min ) &
23    / real ( n - 1, kind = 8 )
24
25    y = ( real ( i - 1, kind = 8 ) * y_max &
26    + real ( n - i, kind = 8 ) * y_min ) &
27    / real ( n - 1, kind = 8 )
28
29    ! image(i,j) = 0
30
31    x1 = x
32    y1 = y
33
34    do k = 1, max_iter
35
36    x2 = x1 * x1 - y1 * y1 + x
37    y2 = 2 * x1 * y1 + y
38
39    if ( x2 < -2.0D+00 .or. &
40    2.0D+00 < x2 .or. &
41    y2 < -2.0D+00 .or. &
42    2.0D+00 < y2 ) then
43
44    image(i,j) = k
45    exit
46
47    end if
48
49    x1 = x2
50    y1 = y2
51
52    end do
53
54    end do
55 end do
56 !$omp end do
57
58 end subroutine

```