

Porting and Performance of PCM on the Cray X1

Vince Wayland, NCAR



NCAR

Acknowledgements

- ORNL consulting staff, esp. Mark Fahey
- Cray Consultants John Levesque, Ilene Carpenter
- NCAR Consultant Jeff Keuhn
- CCP Colleague Gary Strand



Outline

- Description and History of PCM
- Porting Process
- Preliminary Performance on X1



About PCM



NCAR

PCM

- PCM – Parallel Climate Model
 - Joint DOE -NCAR effort to develop a parallel climate model, 1996 - 1997
 - Major Collaborators:
 - NCAR
 - LANL
 - USNPGS
 - ORNL
 - others

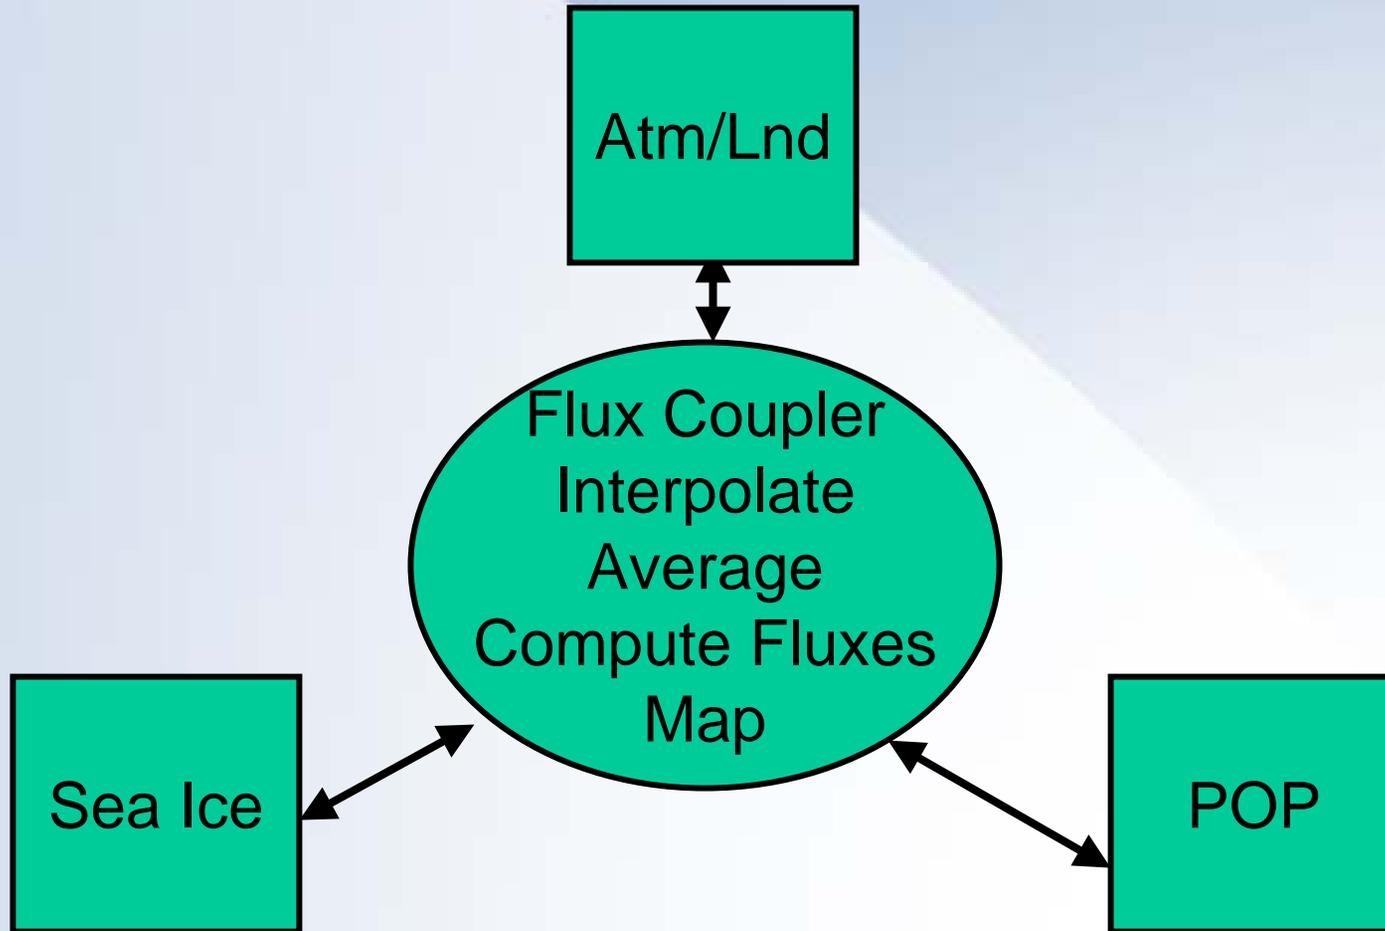


PCM Components

- Atmosphere/Land
 - NCAR CCM3, T42 Resolution
- Ocean
 - LANL POP, 2/3 degree resolution, 1/2 at equator
- Ice
 - USNPGS Sea Ice Model, 27 km resolution
- NCAR Flux Coupler



PCM Schematic



PCM History

- CRI T3D and SGI O2K Original targets
- Ported to:
 - Cray T3E
 - IBM Power 2,3,4
 - HP, nee Compaq, nee DEC Alpha Cluster
 - HP Precision Cluster
- Components Upgraded, RTM added, SHMEM replaced by MPI



PCM Usage

- ~185 cases/experiments run
- ~21000 simulated climate years
- ~100 TB data archived



PORTING



NCAR

Porting PCM to Cray X1

- PCM Source
 - Mainly Fortran 77, some Fortran 90
 - Lots of Previous Ports, Has been Flint-ed
 - A few C routines
 - Max of 64 Processors, due to T42 resolution
 - Lots of Cray CPP macros – artifact from CCM3 C90 days



Porting PCM to Cray X1

- Getting clean compile
 - Couldn't find suggested compiler option strings for typical ports in documentation
 - Convert "CRAY" CPP macro to "CRAY90"
 - Use __crayx1 CPP macro
 - Find X1 acceptable subset of I/O keywords
 - Convert system calls
 - Fortran versions of ceil2.c and pair.c



Porting PCM to Cray X1

- Attempt to execute PCM
 - Seg Fault
 - No output file
 - No traceback
 - Learn about `call flush (101)`
 - Insert lots of “I am here” print statements
 - Learn about `ftn -D` option



Porting PCM to Cray X1

– -D option has fatal interactions with passed character variable lengths

- **SPR 727900**

- Temporary Workaround

Split dependency lists into char and non-char handling lists for each model



Porting PCM to Cray X1

- Finally 1 uninitialized variable found
- Finally 1 call statement with mismatched array dimension and array name arguments found



Porting PCM to Cray X1

- PCM has multiple namelists in a single physical file
 - Depends on closing and reopening of file to not change location of file pointer
 - Didn't work on X1
 - Workaround
 - Separate namelists into individual files



Porting PCM to Cray X1

- Attempt to replace local lenchr function with len_trim intrinsic
- Test case on next slide



```
program stest100
implicit none
integer lenchr, l1, l2, l3
character*18 dummy
character*22 a
character*20 b
character*18 c
lenchr(dummy) = len_trim(dummy)
a = '12345678901234567890'
b = '12345678901234567890'
c = '123456789012345678'
l1 = lenchr(a)
l2 = lenchr(b)
l3 = lenchr(c)
end
```



Porting PCM to Cray X1

- Cray X1 results:
- 40 phoenix:/dfs/home/wayland/test>a.out
3*18

Not quite what I expected, but....

IBM gives same result



Porting PCM to Cray X1

- Second try:
- Change length of ASF dummy argument
- Test case on next slide



```
program stest101
implicit none
integer lenchr, l1, l2, l3
character*20 dummy
character*22 a
character*20 b
character*18 c
lenchr(dummy) = len_trim(dummy)
a = '12345678901234567890'
b = '12345678901234567890'
c = '123456789012345678'
l1 = lenchr(a)
l2 = lenchr(b)
l3 = lenchr(c)
end
```



Porting PCM to Cray X1

- Results in the following compile time diagnostic:

```
43 phoenix:/dfs/home/wayland/test> ftn stest101.F
```

```
l3 = lenchr(c)
      ^
```

```
ftn-848 ftn: ERROR STEST100, File = stest101.F, Line = 14,
Column = 21
```

```
Dummy argument "DUMMY" is longer than the actual argument.
```

- Not quite what I expected
- IBM gives: 20, 20, 18 **What I expected**

Porting PCM to Cray X1

- coreinfo dump analysis utility not much help
- Totalview hard to set up, hard to get X-Window tunnel view through to remote display
- TV did find location of next bug:



Porting PCM to Cray X1

- In ccm subroutine infil.F, loop with many exp expressions mis-compiled, resulting in runtime errors.
- Solution: Convert expression arguments into temporary variables, insert if tests for out of range value and print statements.
- `if (c04 .gt. 709.) print *, 'infil: k,c04',k,',',',c04`
- Note: Print statement required!
- Will submit PSR for simplified case someday.



```

do k = 1, npt
  qsurf = qmelt(k) + qsde(w(k)
  v = bch(k)*smps(k) / (0.5*dzsoi(1,k)*1000.)
  s = h2osoi(1,k)/watsat(k)
  c1 = qsurf - hksat(k)*(1.-v)
  c2 = hksat(k)*v
  sq = max( min(c1/c2, 1.), 0. )
  c3 = max( qrain(k), 1.e-20 )
  c4 = exp(-1./s)
  c5 = exp(-sq/s)
  c6 = exp( min(80., c1/c3 ) )      !upper limit prevents overflow
  c7 = exp( min(80.,-sq/s-sq*c2/c3 ) ) !upper limit prevents overflow
  if (qsurf .le. hksat(k)) then
    flx1 = qrain(k)*c4*exp( (qsurf-hksat(k))/c3 )
  else
    flx1 = (qrain(k)+qsurf-hksat(k)) * c4
  end if
  flx2 = (qrain(k)+c1)*(c5-c4) + c2*(c4*(1.+s)-c5*(sq+s))
  flx3 = qrain(k)**2 / (qrain(k)+c2*s) * c6 * (1.-c7)
  if (ist(k). eq. 1) then          !soil
    qover(k) = flx1+flx2+flx3
  else                             !glaciers, lakes, wetlands
    qover(k) = 0.
  end if
  qinfl(k) = qrain(k) + qsurf - qover(k)
end do

```



Porting PCM to Cray X1

- Junk produced by pat_hwpc
 - `ioctl PIOCWSTOP failed [No such file or directory]`
 - `ioctl PIOCSETPERFCONTROL failed [No such file or directory]`
 - `[1] Exit 1 pat_hwpc -f mpirun -n 4 -t 700 pcm.B00.x1.4 < /dev/null`
- Junk produced by pat_report SPR 726582
 - `pat_report -i pcm.B00.x1.4_inst pcm.B00.x1.4_inst+134061pdt.xf > pcm.4.prof`
 - `[/opt/open/open/bin/pwd: C\2034\2031\2032^A^O\2033Rh^N: not found`
 - `/opt/open/open/bin/pwd: syntax error at line 1: `)' unexpected`
 -
- Fixed in PE 5.2

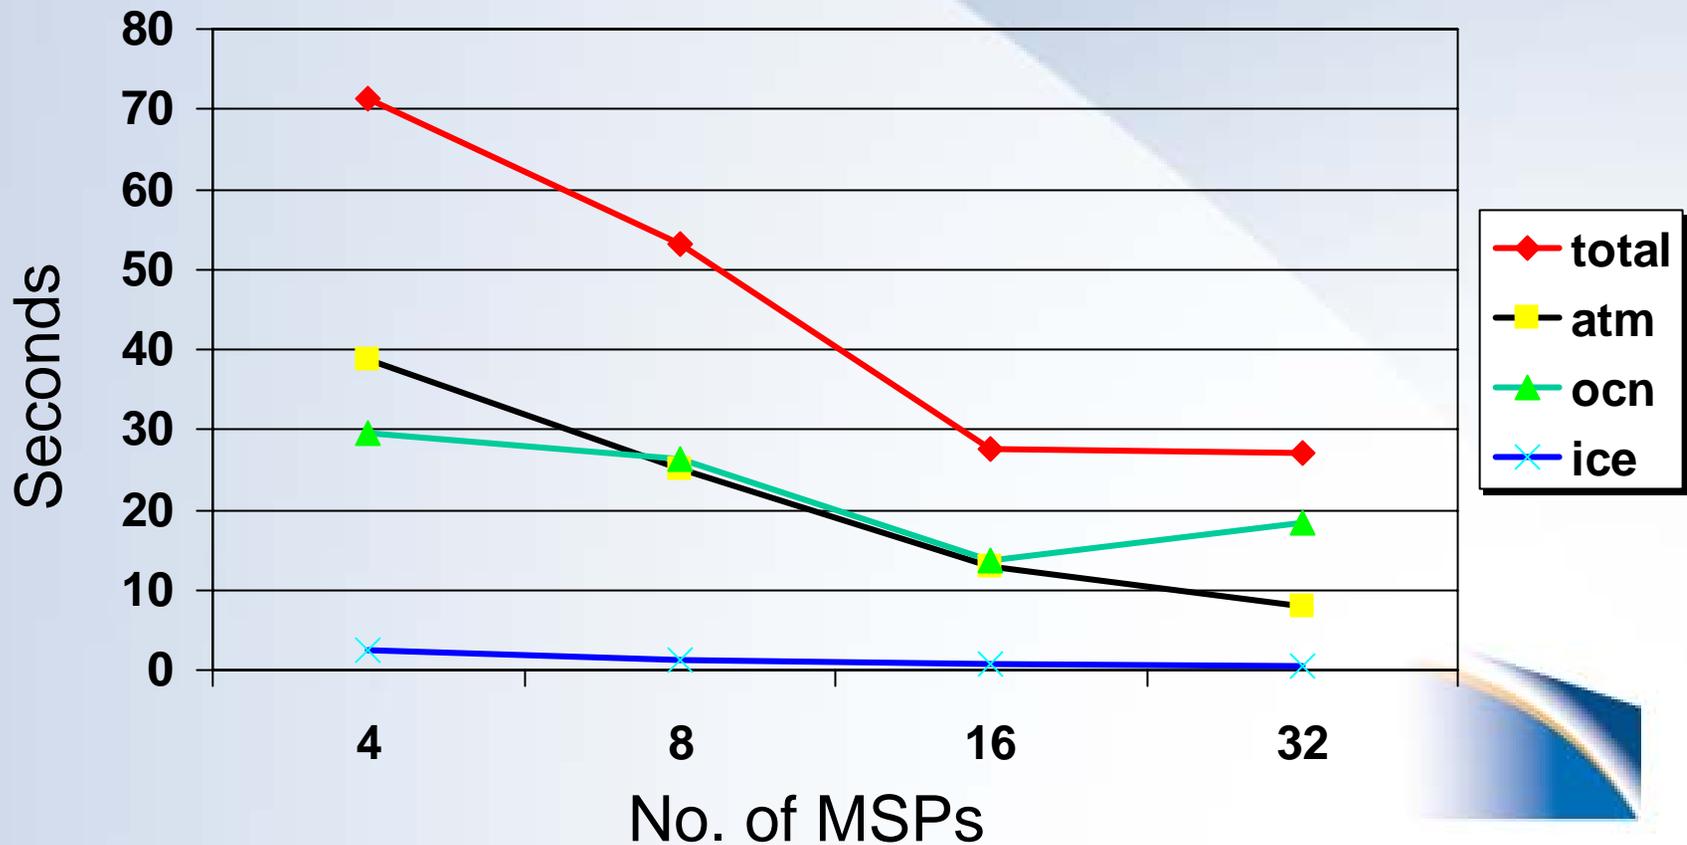


CURRENT PERFORMANCE



NCAR

Time per simulated day



Performance Profile

6.3% | 56.3% | 99610 | pe.1

0.7%	50.7%	11202	ninept_
0.7%	51.4%	10773	fivept_
0.6%	52.0%	10302	fcd_a2o_
0.3%	52.3%	4464	global_real_sum_
0.2%	52.6%	3666	MPI_CRAY_send_wait
0.2%	52.8%	3253	grcalc_
0.2%	52.9%	2667	__bcopy_wordstrm
0.2%	53.1%	2586	__bcopy_strm
0.1%	53.2%	2181	state_
0.1%	53.4%	2101	sw_4pt_
0.1%	53.5%	2065	__bcopy_prv
0.1%	53.6%	1931	fcd_o2a_
0.1%	53.7%	1775	grad_4pt_
0.1%	53.8%	1723	MPI_CRAY_recv_wait



Performance cont'd.

- 9 Point Operator

X X X

X O X

X X X

5 Point Operator

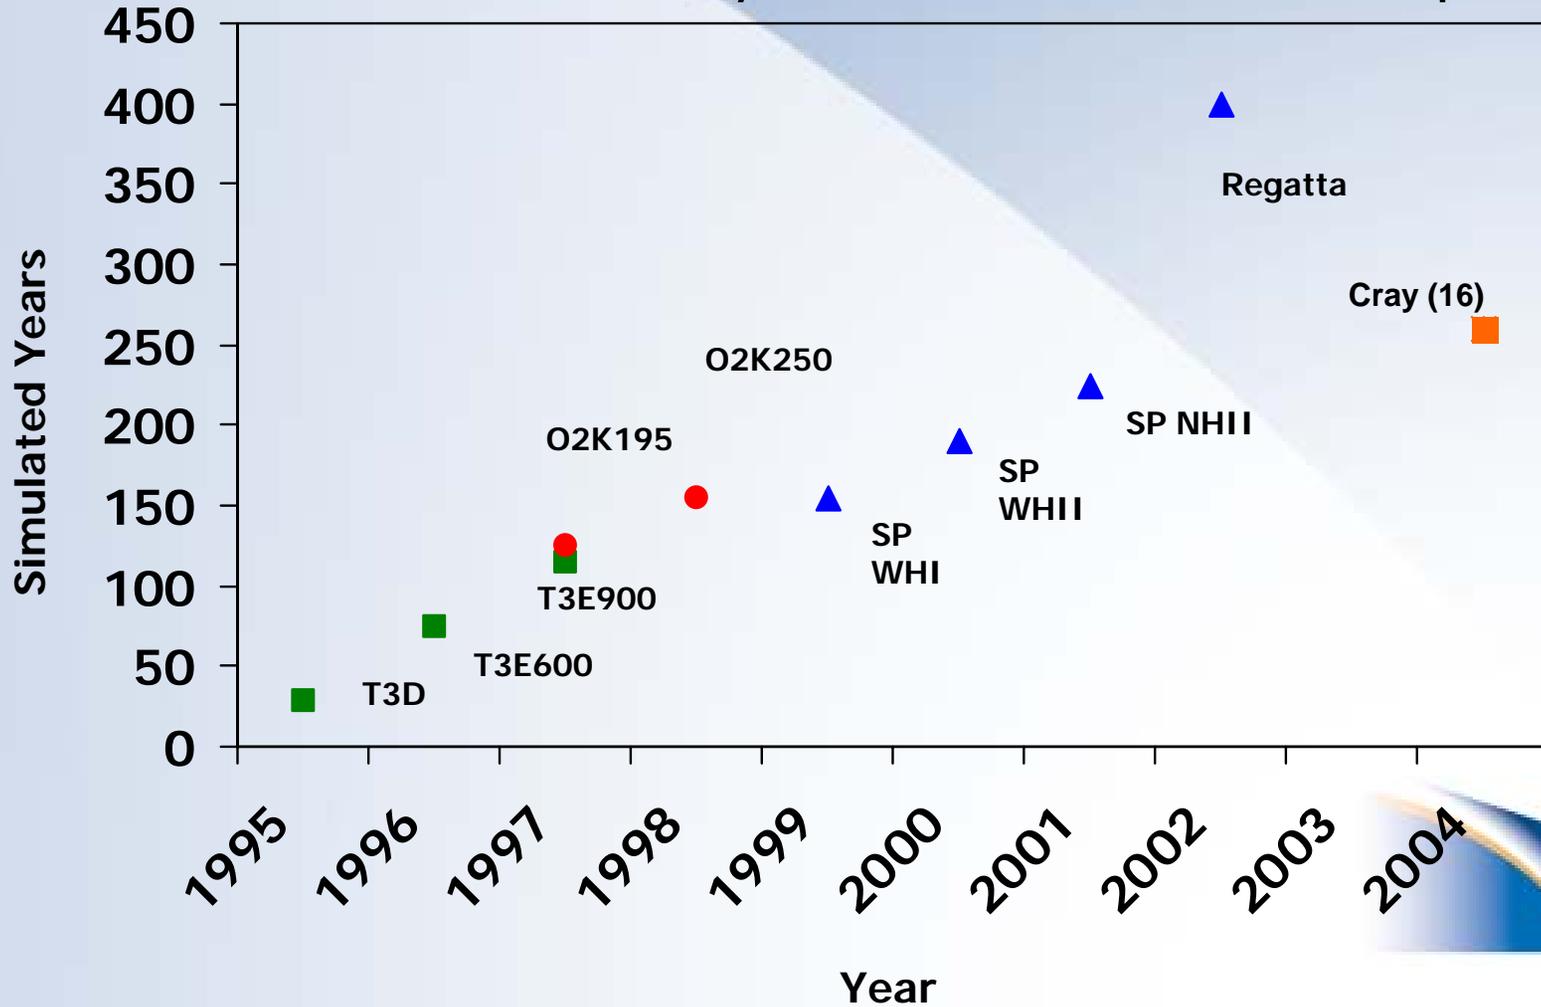
X X X

X O X

X X X

PCM Version 1

Simulated Years per Wallclock Month 64pes



The End



NCAR