

The Dr. HOOK instrumentation kit

Olivier Marsden

European Centre for Medium-Range Weather Forecasts



What is Dr.HOOK?

- A Fortran & C-callable instrumentation library
 - that tracks call tree for every MPI task and OpenMP thread
 - enabled by calling `DR_HOOK()` at entry and exit to every¹ routine
 - written by Sami Saarinen
- This allows Dr HOOK to :
 - trap run-time problems
 - print full traceback at crash site
 - gather performance profile info per subroutine :
 - wall clock or CPU time
 - memory usage
- behaviour is controlled

¹with caveat of overhead

Instrumentation example

```
SUBROUTINE SLOW_ROUTINE
```

```
USE YOMHOOK, ONLY : JPHOOK, LHOOK, DR_HOOK  
IMPLICIT NONE
```

```
REAL(KIND=JPHOOK) :: ZHOOK_HANDLE ! Must be a local (stack) variable
```

```
!-The very first statement in the subroutine
```

```
IF (LHOOK) CALL DR_HOOK('SLOW_ROUTINE',0,ZHOOK_HANDLE)
```

```
!---Body of the routine goes here---
```

```
IF (LHOOK) CALL DR_HOOK('SLOW_ROUTINE',1,ZHOOK_HANDLE)  
END SUBROUTINE SLOW_ROUTINE
```

Controlling Dr. HOOK

Environment variables (a few examples)

- `DRHOOK=1`
- `DR_HOOK_OPT=profile / memory / cputime`
- `DR_HOOK_CATCH_SIGNALS=1`
- `DR_HOOK_IGNORE_SIGNALS=8`

Example output

Dr Hook results fileName of the executable :

/ec/fws5/sb/work/rd/dipg/hui7/LWDA/2019080100/gfc/tmp.hui7_fc_main_fcgroup.model.5.201613/ifsMASTER.SP

Number of MPI-tasks : 62

Number of OpenMP-threads : 4

Wall-times over all MPI-tasks (secs) : Min=25.630, Max=25.840, Avg=25.731, StDev=0.052

Routines whose total time (i.e.sum) > 0.000 secs will be included in the listing

| Avg-% | Avg.time | Min.time | Max.time | Incl.ave | Incl.min | Incl.max | St.dev | Imbal-% | # of calls | Name of the routine |
|-------------------|----------|----------|----------|----------|----------|----------|--------|---------|------------|---------------------|
| 13.86% | 3.559 | 3.521 | 3.736 | 3.559 | 3.522 | 3.736 | 0.030 | 5.75% | 1550 | UPDECAEC |
| 0.06% | 0.016 | 1.007 | 1.007 | 0.016 | 1.007 | 1.007 | 0.128 | 0.00% | 1866 | GRIB_API:IGR |
| 2.90% | 0.745 | 0.736 | 0.753 | 0.745 | 0.736 | 0.753 | 0.004 | 2.26% | 62 | YOEAERC:SU_A |
| 2.03% | 0.522 | 0.461 | 0.572 | 0.626 | 0.537 | 0.691 | 0.023 | 19.37% | 20150 | radiation_mc |
| 1.46% | 0.375 | 0.354 | 0.397 | 0.506 | 0.503 | 0.507 | 0.021 | 10.81% | 62 | SUECAEOR |
| 1.08% | 0.278 | 0.261 | 0.298 | 0.312 | 0.295 | 0.331 | 0.008 | 12.38% | 60450 | CLOUDSC |
| 1.17% | 0.301 | 0.238 | 0.363 | 0.420 | 0.323 | 0.502 | 0.026 | 34.34% | 20150 | radiation_mc |
| 0.74% | 0.190 | 0.189 | 0.192 | 25.390 | 25.388 | 25.394 | 0.001 | 1.56% | 62 | MASTER |
| 0.01% | 0.003 | 0.174 | 0.174 | 0.003 | 0.174 | 0.174 | 0.022 | 0.01% | 78 | GRIB_API:IGR |
| 0.63% | 0.161 | 0.158 | 0.164 | 0.162 | 0.159 | 0.165 | 0.003 | 3.64% | 62 | SUAFN1 |
| 0.59% | 0.152 | 0.145 | 0.159 | 0.152 | 0.145 | 0.159 | 0.003 | 8.76% | 730236 | VERINT_DGEMM |
| 0.59% | 0.153 | 0.143 | 0.161 | 0.153 | 0.143 | 0.161 | 0.004 | 11.12% | 120900 | LATRI_WENO |
| 0.54% | 0.138 | 0.129 | 0.144 | 0.138 | 0.129 | 0.144 | 0.004 | 10.35% | 302250 | LARCHE |
| 0.50% | 0.128 | 0.126 | 0.131 | 0.129 | 0.127 | 0.131 | 0.001 | 3.80% | 107880 | SINPUT_ARD |
| 0.56% | 0.143 | 0.124 | 0.159 | 0.143 | 0.124 | 0.159 | 0.009 | 21.88% | 20150 | radiation_ae |
| 0.48% | 0.123 | 0.117 | 0.128 | 0.141 | 0.135 | 0.146 | 0.002 | 8.53% | 302250 | LASCAW |
| 0.46% | 0.117 | 0.117 | 0.117 | 0.117 | 0.117 | 0.117 | 0.000 | 0.01% | 62 | SUECAEBC:PAR |
| 0.01% | 0.002 | 0.117 | 0.117 | 0.002 | 0.117 | 0.117 | 0.015 | 0.01% | 6 | GRIB_API:IGR |
| 0.45% | 0.115 | 0.114 | 0.115 | 0.242 | 0.241 | 0.243 | 0.000 | 0.87% | 62 | SUECAEBC |
| 0.52% | 0.134 | 0.109 | 0.156 | 2.215 | 2.124 | 2.273 | 0.011 | 29.94% | 1550 | RADINTG |
| 0.37% | 0.096 | 0.094 | 0.099 | 0.129 | 0.126 | 0.132 | 0.001 | 5.01% | 60450 | CLOUD_SATADJ |
| 17.86% | 4.586 | 0.092 | 8.793 | 4.586 | 0.092 | 8.793 | 4.306 | 98.94% | 4774 | >MPL- |
| TRMTOL_COMMS(807) | | | | | | | | | | |
| 0.39% | 0.100 | 0.092 | 0.105 | 0.100 | 0.092 | 0.105 | 0.003 | 12.27% | 60450 | CUANCAPE2 |