

Parallel NONMEM, Initial experience

Mark Sale M.D. Next Level Solutions LLC
Andy Sherman, ASPEED

The problem:

- The expectation is that the field will continue to move forward
 - ◆ More detailed biology models
 - ◆ Larger, longer duration data sets
 - ◆ More computationally intensive methods
 - ◆ Shorter timelines
- Intel has not released a faster CPU since mid 2004.

Modern programming:

- Parallel execution approaches
 - ◆ Multiple threads – shared memory
 - ◆ Multiple processes – no shared memory

Windows Task Manager

File Options View Shut Down Help

Applications Processes Performance Networking Users

Image Name	CPU	CPU Time	Mem Usage	Base Pri	Threads
nonmem.exe	99	0:00:37	3,144 K	Normal	1
EXCEL.EXE	00	0:00:00	15,484 K	Normal	5
WZQKPICK.EXE	00	0:00:00	412 K	Normal	1
RAMASST.exe	00	0:00:00	496 K	Normal	2
TOSCDSPD.exe	00	0:00:00	384 K	Normal	1
PDVDServ.exe	00	0:00:00	448 K	Normal	2
CFSServ.exe	00	0:00:14	3,424 K	Normal	14
ctfmon.exe	00	0:00:00	836 K	Normal	1
qtask.exe	00	0:00:00	300 K	Normal	2
TPSBattM.exe	00	0:00:00	364 K	Normal	1
oasclnt.exe	00	0:00:00	508 K	Normal	2
mcvshld.exe	00	0:00:00	812 K	Normal	2
pingr.exe	00	0:00:00	1,004 K	Normal	1
SmoothView.exe	00	0:00:00	356 K	Normal	1
PadExe.exe	00	0:00:03	4,592 K	Normal	8
TPSMain.exe	00	0:00:05	1,800 K	Normal	1
TFncky.exe	00	0:00:00	864 K	Normal	1
DLACTRLW.EXE	00	0:00:00	2,040 K	Normal	9
anrsmmsn.exe	00	0:00:00	404 K	Normal	2

Show processes from all users

End Process

Processes: 58 CPU Usage: 100% Commit Charge: 308M / 1056M

History:

- Demonstration project initially presented at AAPS last Nov 2005.
- 2.6 fold increase in speed on a 4 processor server

Collaboration between Next Level Solutions /ASPEED/Globomax

- Method shared with Aspeed, to commercialize product
- Globomax to provide technical support.
- Aspeed will actually develop/sell/support the product
- To run on network, rather than SMP
- Could use 100's of processors

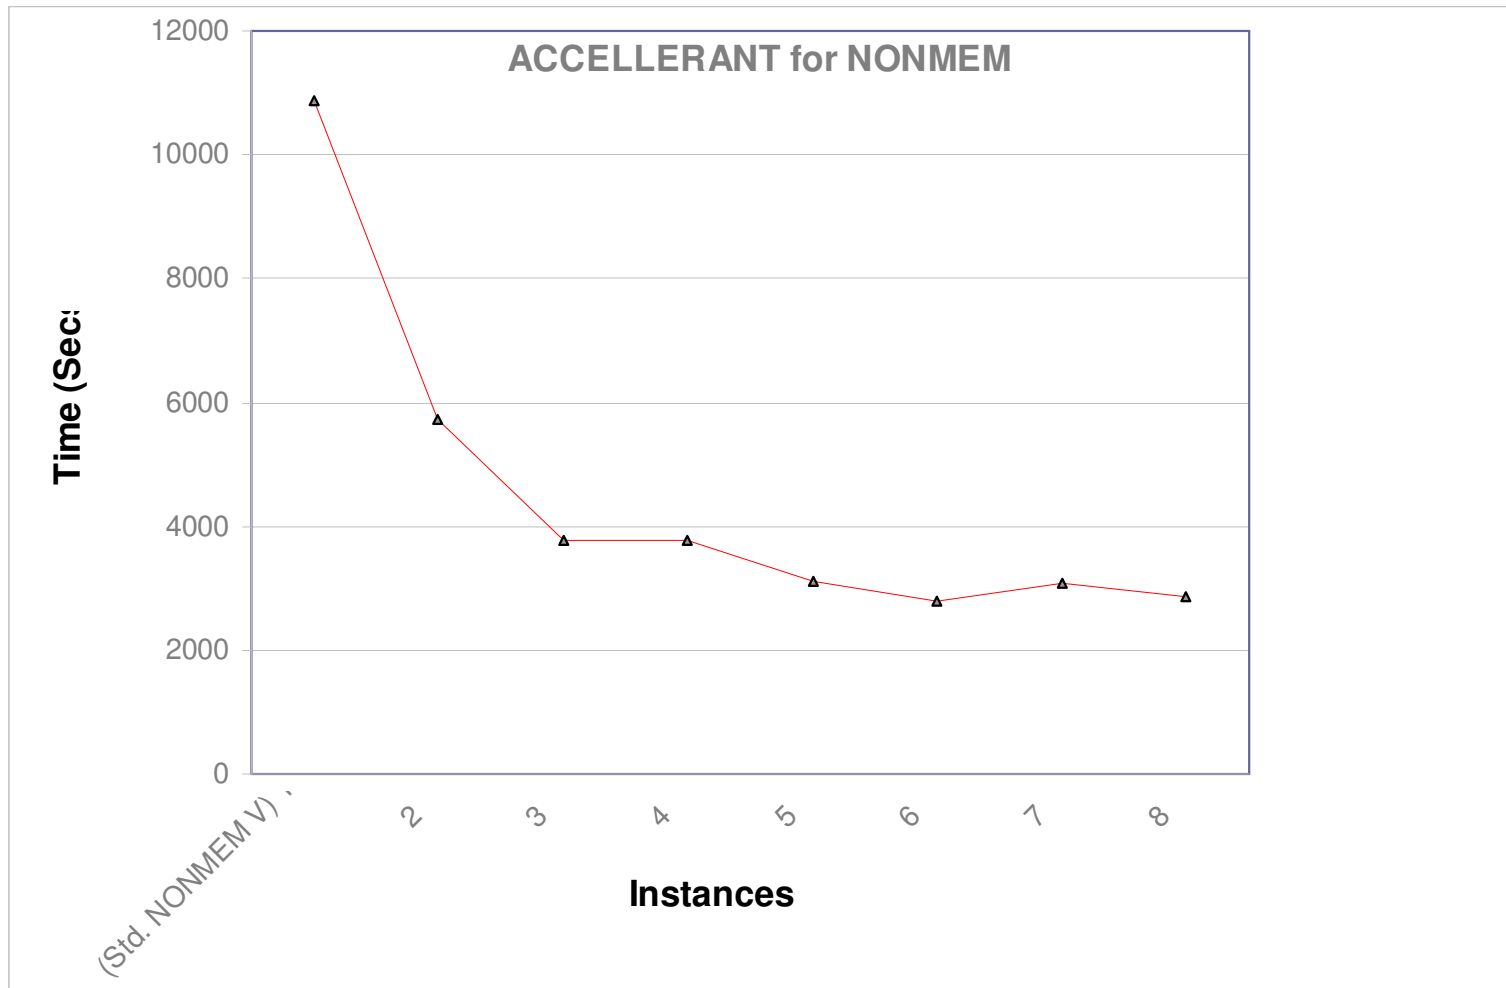
Requirements:

- All ADVANs
- FO and FOCE with/without INTER (center may not be available)
- +/- parallelization of covariance step
- Same (or better) answer than NONMEM

Benchmark:

- Actual data, NONMEM V, ADVAN6, somewhat stiff. (but refused to run in ADVAN8/9)
- Native NONMEM execution time: 10861 seconds (Linux, 2.8 Ghz Xeon machine, 32 bit).

Two processor, Dual core machine



Other stats:

	Total NONMEM Function Evaluations	Wall Clock Time (Secs)	Acceleration Factor
(Std. NONMEM V) 1		10861	
2	1459	5732	1.9
3	1459	3754	2.9
4	1459	3762	2.9
5	1459	3120	3.5
6	1459	2783	3.9
7	1459	3064	3.5
8	1459	2858	3.8

Comparison of results:

- Identical THETA/OMEGA/SIGMA (to 3 digits)
 - ◆ Differences in gradient only.

Parallel NONMEM

- Networked computers (still needs some work)

Network

2 x 2 (1)

9406

2 x 2 (2)

7757

How is this done???

- Stuart has said (at least twice) in public that NONMEM cannot be parallelized
- But, method to parallelize non-linear regression is published in standard textbooks on high performance computing.

How is this done???(2)

- For hard problems (ODE solvers, ADVAN6, 8, 9) NONMEM spends nearly all of its time calculating predictions (integrating ODEs).
- The method of getting the prediction is defined in PRED, which in turn calls ADVAN? to get the predictions.
- Subroutine PRED return prediction which is then used to calculate the likelihood for each predicted value.

Predicted values are dependent within an individual

- But, each individual gets independently calculated predictions, and then the likelihood for each subject is summed
- Then the likelihood for the entire data set (at current values for THETA/OMEGA/SIGMA) is calculated.

Current status:

- In previous version ADVAN6 was tested for FOCE INTER
- CENTER did not work.
- FO was not examined.
- Not even close to a push button (nmfe6, WFN like) solution

Current version

- Not really tested much
- ADVAN6 only so far.
- But, achieved nearly linear speed increase (3.8 fold increase on 4 processor machine).
- Old version wrote disk file to be read by “slave” processes.
- ASPEED version “streams” data to slave processes.

Better answer???

- When the contributions of each individual to the log likelihood is summed, they should be summed in decreasing absolute value, to reduce rounding error.

Better answer???

Terminated due to rounding errors?

- Assume 5 digits of precision:

- ◆ $1000.0 + 0.04 = 1000.0$

- Assume 3 digits of precision

$$9.19 + 1.14 + 1.11 = 11.4$$

$$(9.19 + 1.14 = 10.3)$$

$$(10.3 + 1.11 = 11.4)$$

$$1.11 + 1.14 + 9.19 = 11.5$$

$$(1.11 + 1.14) = 2.25 ;$$

$$(2.25 + 9.19 = 11.5)$$

Microsoft Excel - nmsequence

Type a question for help

File Edit View Insert Format Tools Data Window Help

Reply with Changes... Egd Review...

K11 5

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S
1	OBJ	THETA(1)	THETA(2)	THETA(3)	THETA(4)	THETA(5)	THETA(6)	THETA(7)	OMEGA(1)	OMEGA(2)	OMEGA(3)	OMEGA(4)	OMEGA(5)	OMEGA(6)	OMEGA(7)	SIGMA(1)	MINIMIZED?	COVARIANCE?
2	-856.624	4.23	1.78	2.08	2.18	2.59	2.05	1.2	0.977	7.84	5	1.06	19	8.83	11.4	1.01	FALSE	FALSE
3	-856.624	4.23	1.78	2.08	2.18	2.59	2.05	1.2	0.977	7.84	5	1.06	19	8.83	11.4	1.01	TRUE	FALSE
4	-856.624	4.23	1.78	2.08	2.18	2.59	2.05	1.2	0.977	7.84	5.01	1.06	19	8.83	11.4	1.01	TRUE	FALSE
5	-856.624	4.23	1.78	2.08	2.18	2.59	2.05	1.2	0.977	7.84	5	1.06	19	8.83	11.4	1.01	TRUE	FALSE
6	-856.624	4.23	1.78	2.08	2.18	2.59	2.05	1.2	0.977	7.84	5	1.06	19	8.83	11.4	1.01	FALSE	FALSE
7	-856.624	4.23	1.78	2.08	2.18	2.59	2.05	1.2	0.977	7.84	5	1.06	19	8.83	11.4	1.01	FALSE	FALSE
8	-856.624	4.23	1.78	2.08	2.18	2.59	2.05	1.2	0.977	7.84	5	1.06	19	8.83	11.4	1.01	TRUE	FALSE
9	-856.624	4.23	1.78	2.08	2.18	2.59	2.05	1.2	0.977	7.84	5	1.06	19	8.83	11.4	1.01	TRUE	FALSE
10	-856.624	4.23	1.78	2.08	2.18	2.59	2.05	1.2	0.977	7.84	5	1.06	19	8.83	11.4	1.01	TRUE	FALSE
11	-856.624	4.23	1.78	2.08	2.18	2.59	2.05	1.2	0.977	7.84	5	1.06	19	8.83	11.4	1.01	FALSE	FALSE
12	-856.624	4.23	1.78	2.08	2.18	2.59	2.05	1.2	0.977	7.84	5	1.06	19	8.83	11.4	1.01	TRUE	FALSE
13	-856.624	4.23	1.78	2.08	2.18	2.59	2.05	1.2	0.977	7.84	5	1.06	19	8.83	11.4	1.01	TRUE	FALSE
14	-856.624	4.23	1.78	2.08	2.18	2.59	2.05	1.2	0.977	7.84	5	1.06	19	8.83	11.4	1.01	TRUE	FALSE
15	-856.624	4.23	1.78	2.08	2.18	2.59	2.05	1.2	0.977	7.84	5	1.06	19	8.83	11.4	1.01	TRUE	FALSE
16	-856.624	4.23	1.78	2.08	2.18	2.59	2.05	1.2	0.977	7.84	5	1.06	19	8.83	11.4	1.01	FALSE	FALSE
17	-856.624	4.23	1.78	2.08	2.18	2.59	2.05	1.2	0.977	7.84	5	1.06	19	8.83	11.4	1.01	TRUE	FALSE
18	-856.624	4.23	1.78	2.08	2.18	2.59	2.05	1.2	0.977	7.84	5	1.06	19	8.83	11.4	1.01	FALSE	FALSE
19	-856.624	4.23	1.78	2.08	2.18	2.59	2.05	1.2	0.977	7.84	5	1.06	19	8.83	11.4	1.01	FALSE	FALSE
20	-856.624	4.23	1.78	2.08	2.18	2.59	2.05	1.2	0.977	7.84	5	1.06	19	8.83	11.4	1.01	FALSE	FALSE
21	-856.624	4.23	1.78	2.08	2.18	2.59	2.05	1.2	0.977	7.84	5	1.06	19	8.83	11.4	1.01	TRUE	FALSE
22	-856.624	4.23	1.78	2.08	2.18	2.59	2.05	1.2	0.977	7.84	5	1.06	19	8.83	11.4	1.01	TRUE	FALSE

Ready

Does this happen?

```
0ITERATION NO.: 3  OBJECTIVE VALUE: -0.84904E+03  NO. OF FUNC. EVALS.:12
CUMULATIVE NO. OF FUNC. EVALS.: 44
PARAMETER: 0.1011E+00 0.5584E+00 0.1007E+00 0.7456E-01 0.2069E+00
            0.1525E+00 0.5456E-01 0.8555E-01 0.1052E+00 0.7900E-01
            0.1363E+00 0.1527E+00 0.9231E-01 0.1232E+00 0.1589E+00
GRADIENT: -0.6751E+00 0.2870E+02 0.1354E+02 -0.2137E+01 -0.1794E+02
           0.3482E+02 -0.1013E+02 -0.2693E+00 0.4964E+01 0.2505E+01
           -0.7018E+00 -0.1148E+02 0.3564E+01 0.3308E+01 0.1175E+03
SEQUENTIAL SUM
SUM = -728.053425913415
SORTED SUM
SUM = -728.053425913415

SEQUENTIAL SUM
SUM = -828.966701280901
SORTED SUM
SUM = -828.966701280900
```

So, we've fixed this?

- No yet, put in sorting, then summing - Didn't help.
- But, it is a little more complicated. NONMEM calls OBJ in a very complex way that we haven't fully figured out yet – and has changed since NONMEM V.
- We think it will probably (modestly) improve both convergence and covariance stability.

New performance benchmark

- Intel, which next week is expected to announce plans to move to a new processor architecture, is switching to a new yardstick to measure processor performance: performance per watt*.

*<http://www.eweek.com/article2/0,1895,1848394,00.asp>

Future: from Intels web site

- Intel will also deliver a quad-core (4 full execution cores) processor to the high-end desktop based upon this new microarchitecture, codenamed Kentsfield. Kentsfield is targeted for introduction in the first quarter of 2007*.

*<http://www.intel.com/technology/architecture/coremicro/>

Webinar on availability Aug 2, 2006

- www.aspeed.com/NONMEMReg
- I'm told launch for version 5 early September, version 6 mid September.