

Chemora!

Computational **H**ierarchy for
Engineering **M**odel-**O**riented
Re-adjustable **A**pplications [tm]

Chemora Hands-On

Steven R. Brandt¹

David M. Koppelman²

Yue Hu²

Louisiana State University
Baton Rouge, LA U.S.A.

¹ Division of Computer Science

² Division of Electrical & Computer Engineering
Center for Computation and Technology

Einstein Toolkit Workshop, 14 August 2015, Stockholm, Sweden

Chemora! Goals

Someday in the near future

Write a differential equation description of some system ...

... enter a simple command ...

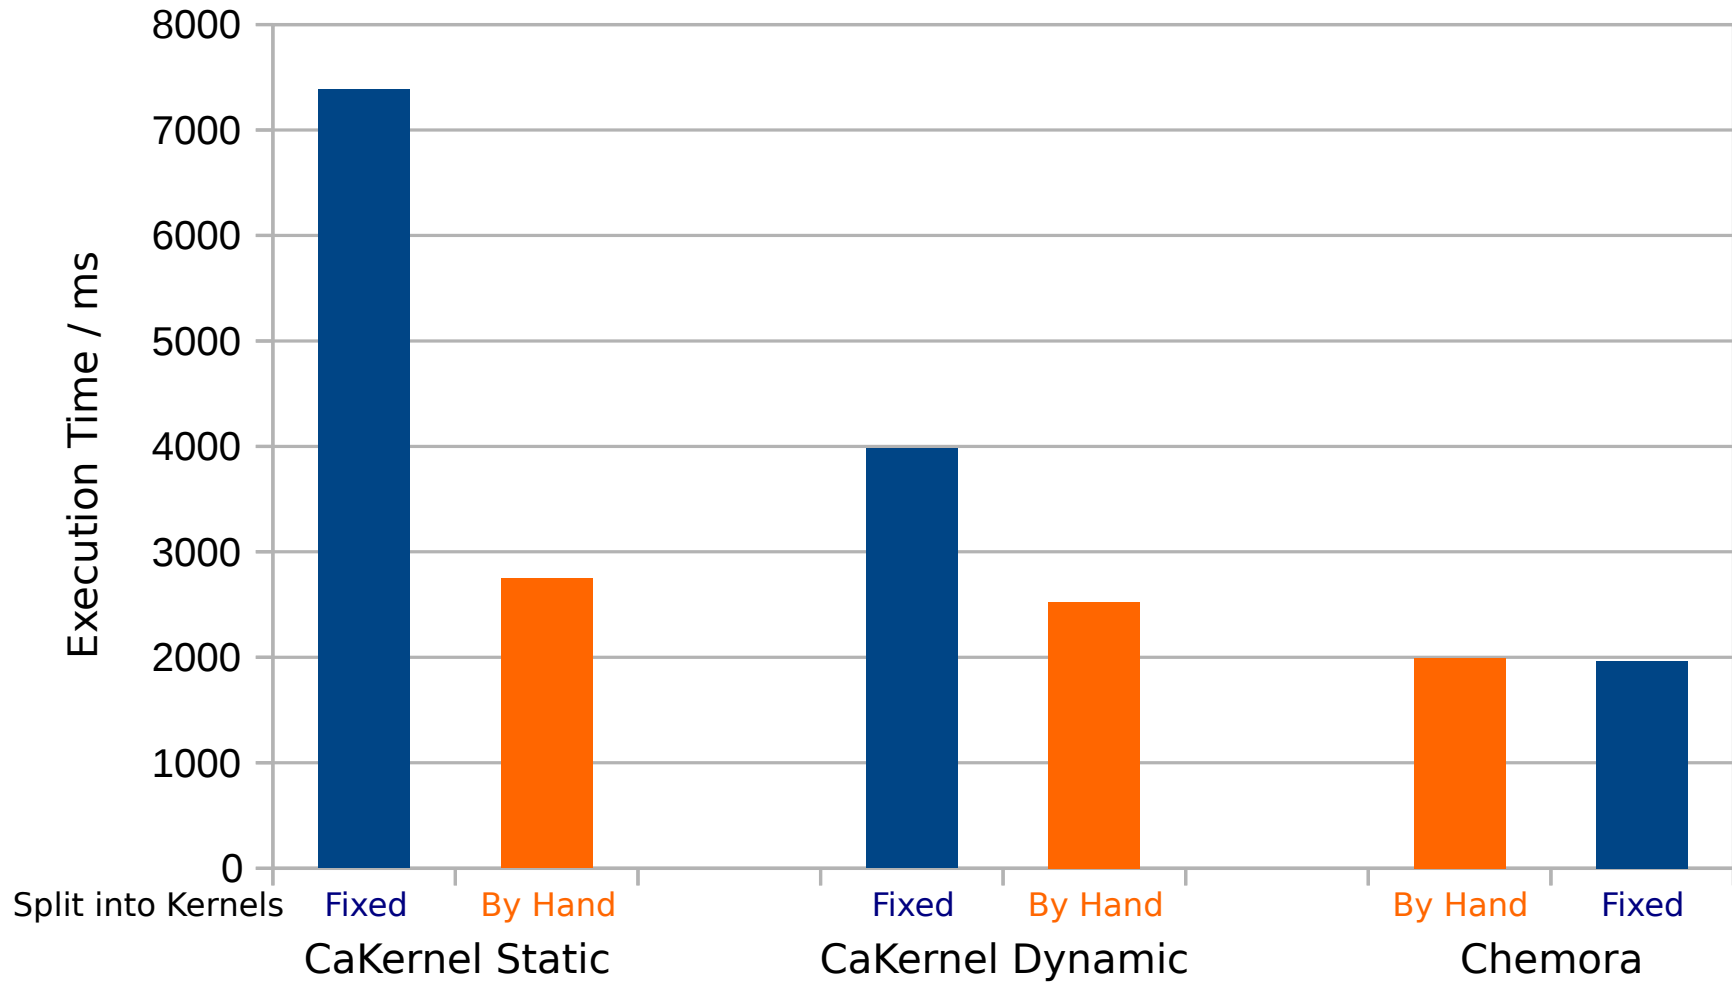
... the Chemora system will generate highly tuned code...

... tailored to your system and simulation parameters.

You won't need to think about what kind of accelerator you have ...

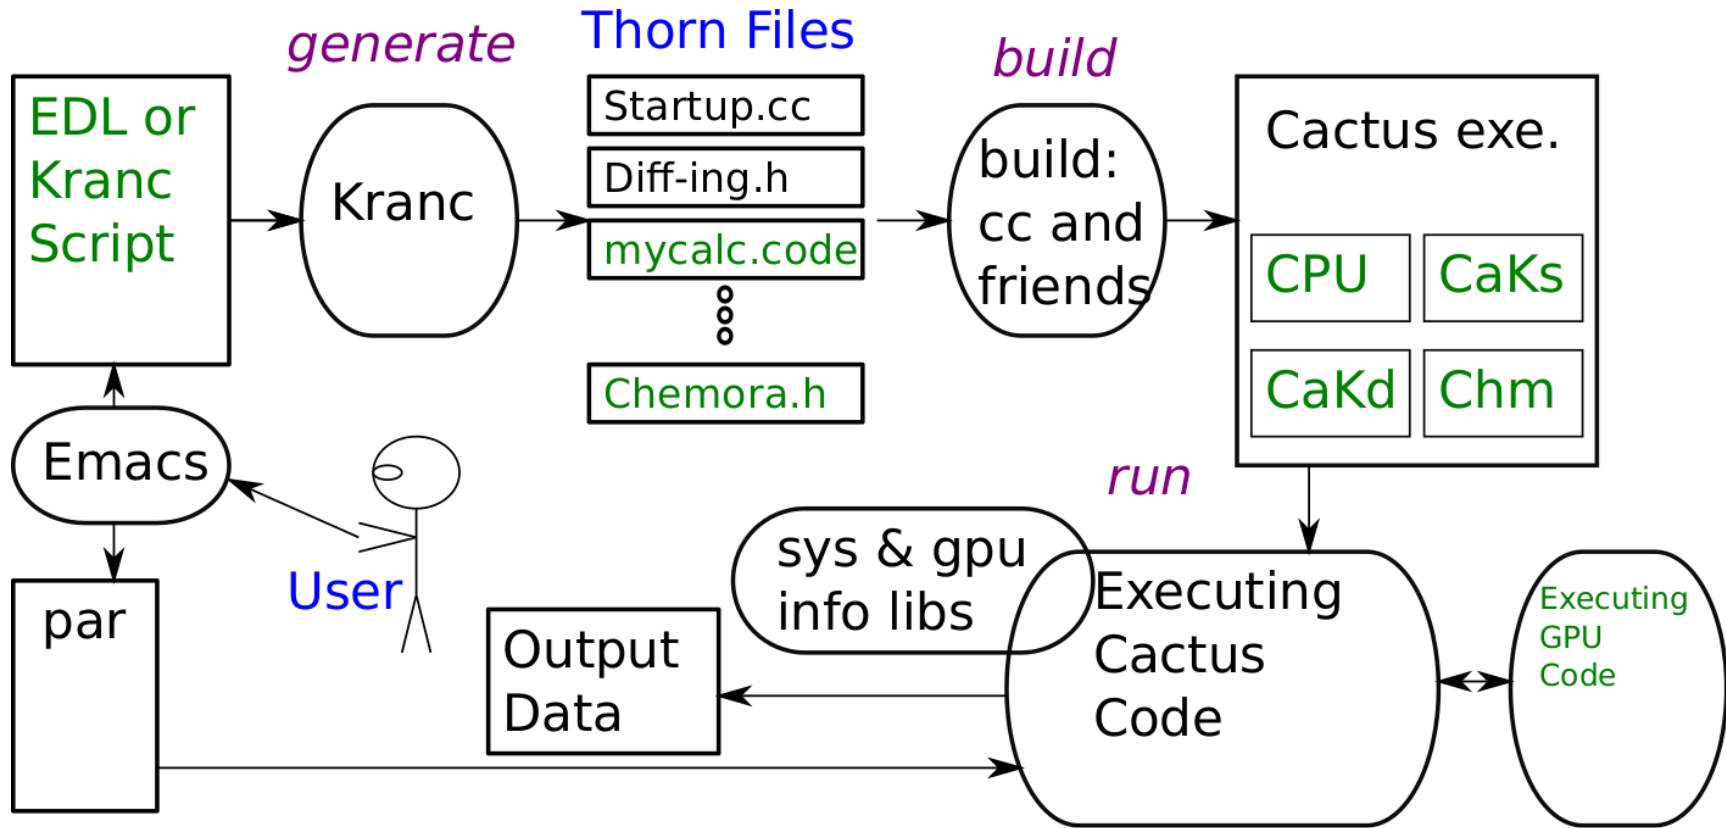
... let alone how to tune for it.

Speedup



Overview

From EDL to GPU Code



From EDL to GPU Code

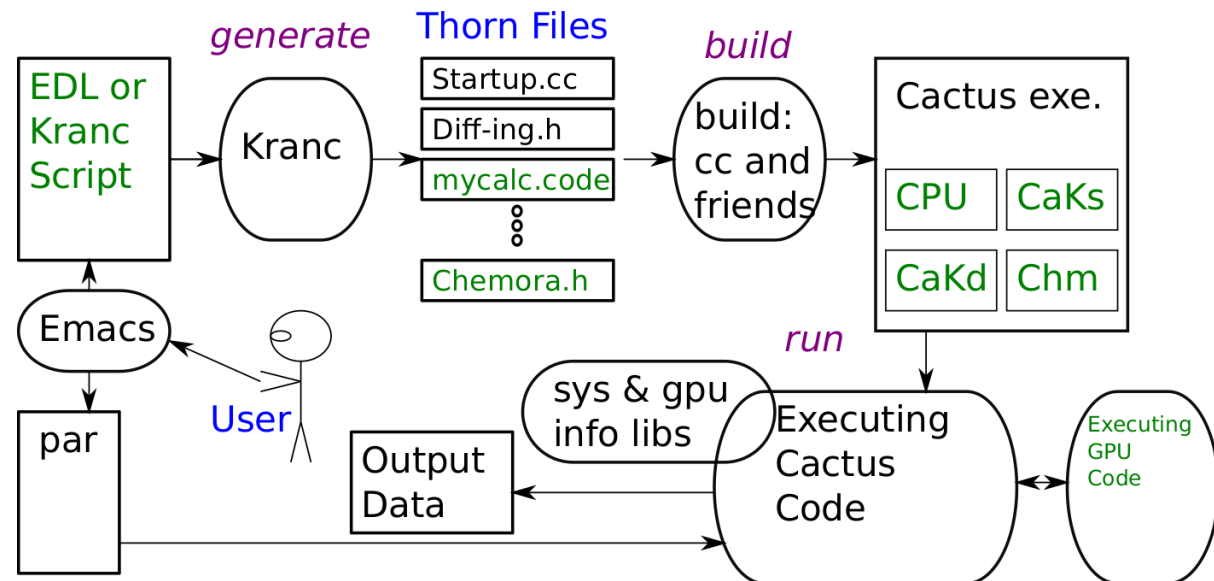
User writes EDL or Kranc script.

User writes parameter file describing simulation.

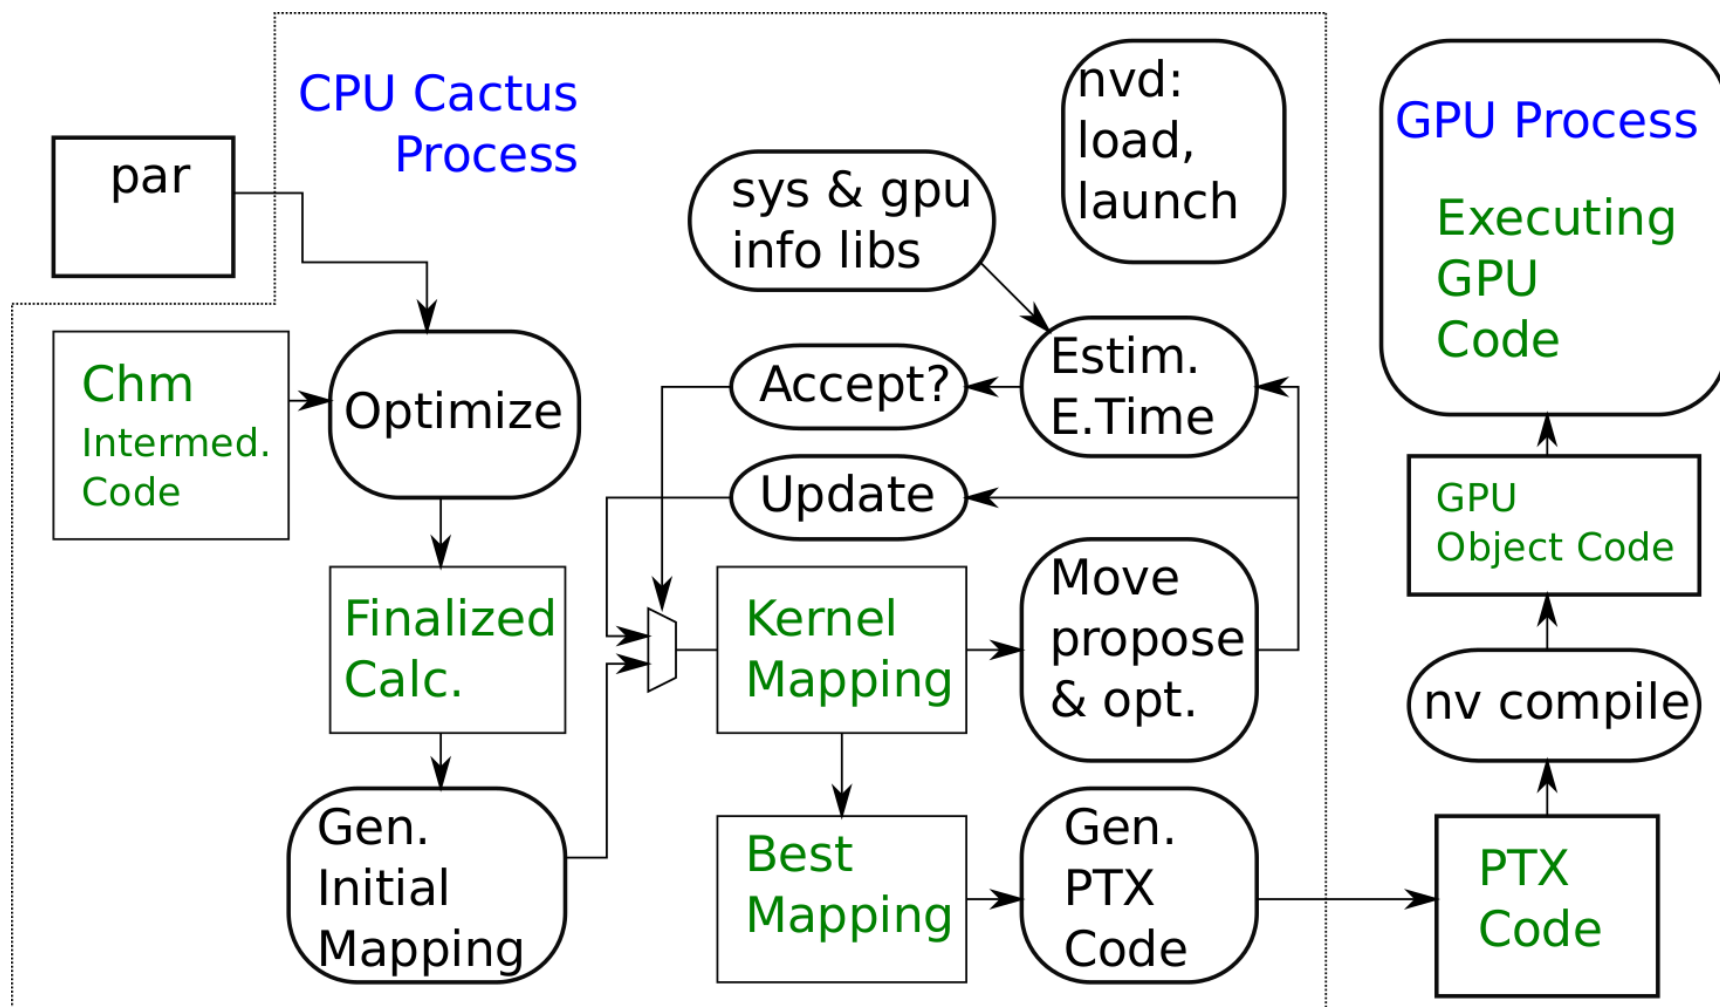
User runs Kranc to *generate* a Cactus thorn.

User *builds* Cactus to generate a Cactus executable.

User *runs* Cactus executable.



Chemora! Overview



Steps

For Each Calculation

Assign *nodes* (operations such as $+$, \times , GF access) ...
... to one or more *kernels*.

For Each Kernel

Based on stencil shape, ...
... choose a *buffering* strategy.

Based on buffering strategy and number of nodes ...
... determine number of threads and tile shape.

Based on all this, estimate kernel execution time.

Compute calculation execution time.

Go to step 1, maybe.

Buffering Strategies

Problem: *Reuse*: Multiple accesses to same grid point.

For example, $a_{i-1,j,k} + a_{i,j,k}$.

The same element accessed by grid points (i, j, k) and $(i + 1, j, k)$.

If not handled properly reuse can multiply amount of off-chip data access.

Buffering Strategies

G: No Offset (Stencil) Accesses

No buffering needed, direct access to global memory.

|: All offsets only along y axis.

Iterate along y axis.

Buffer values in registers, no shared memory used.

This is the most favourable case.

Strategy also used for the z axis.

x: All offsets only along x axis.

This is like $|$ accesses, but requires a different optimization ...
... because of mandatory spatial locality along x axis.

Currently not optimized.

+: If an offset along y , then x and z are zero.

Example of **+**: $a_{0,1,0} + a_{1,0,1} + a_{0,0,-1}$.

Example of **NOT +**: $a_{0,1,0} + a_{1,0,1} + a_{0,1,-1}$.

Iteration along y axis.

Shared memory used for xz plane at offset 0 from y .

Time needed to access shared memory is overhead.

+ also includes the pattern for the z access.

B: Offsets not following the pattern above.

Not encountered in BSSN.

Chemora Tuning Output

Purposes

Developers: Chemora development purposes.

Everyone: Track down possible performance problems.

Users: Feel good about all the work you don't need to do.

Tuning Tables

Information Common to All Tables

Chosen tile size.

Shows important factors contributing to et.

Shows **estimated execution time (et)** for each kernel.

Example: Two kernels from alpha3_2.

:	Tile	Iter	B	FP/Insn	IK	Elts	Issue	Data	Lat	ET	OV	H/I
0:	70, 1,14	10	y !	8/ 9	0	2.8	1.2	0.9	0.9	1.9	0	1.80
1:	70, 1,14	10	y !	8/ 9	0	2.8	1.2	0.9	0.9	1.9	0	1.80
	-----	-----								-----		
	Tile	Iter								Estim.		
	Dims	Direction								Exec.		
	(Threads)	and buffering.								Time		

Example, continued.

	Tile	Iter	B	FP/Insn	IK	Elts	Issue	Data	Lat	ET	OV	H/I
0:	70, 1,14	10	y !	8/ 9	0	2.8	1.2	0.9	0.9	1.9	0	1.80
1:	70, 1,14	10	y !	8/ 9	0	2.8	1.2	0.9	0.9	1.9	0	1.80

Execution Time
Limiters

Ideal: $ET = \max\{ \text{Issue}, \text{Data}, \text{Lat} \}$

Used: $ET = \max\{ \text{Issue}, \text{Data}, \text{Lat} \} + f(\text{Issue}, \text{Data}, \text{Lat})$

Issue: Limit based on number of FP units, etc.

Data: Limit based on off-chip bandwidth.

Lat: Limit based on operation latency.

Example, continued.

:	Tile	Iter	B	FP/Insn	IK	Elts	Issue	Data	Lat	ET	OV	H/I
0:	70, 1,14	10	y !	8/ 9	0	2.8	1.2	0.9	0.9	1.9	0	1.80
1:	70, 1,14	10	y !	8/ 9	0	2.8	1.2	0.9	0.9	1.9	0	1.80

Number of
Insns.
per grid
point

Reals
per
grid
point

Halo (ghost zone)
Overhead.
(No stencil = 1)
f(# of gp/kernel)

Example, continued.

:	Tile	Iter	B	FP/Insn	IK	Elts	Issue	Data	Lat	ET	OV	H/I
1:	70,14,	1	10 z G	0/ 1	0	1.0	0.0	0.3	0.0	0.4	0	1.00
2:	70,14,	1	10 z G	613/ 99	41	140.0	82.5	46.5	30.1	109.3	235	1.00

Interkernel
data transfer.

Spill
Regs.

Good and Bad

:	Tile	Iter	B	FP/Insn	IK	Elts	Issue	Data	Lat	ET	OV	H/I
1:	70,14, 1	10	z G	0/ 1	0	1.0	0.0	0.3	0.0	0.4	0	1.00
2:	70,14, 1	10	z G	613/ 99	41	140.0	82.5	46.5	30.1	109.3	235	1.00
	-----				--						---	
	Higher				Low						Very	
											Low	

Post-Run Table

	Tile	Iter	L-FP	L-Dat	Issue	Data	Lat	ET	SET	ET/ms	Ratio	
23:	70,	1,14	10y	2.0	1.9	2.3	1.9	1.0	3.4	3.8	17.357	1.13
2:	70,	1,14	10y	2.0	1.9	2.3	1.9	1.0	3.4	3.8	17.360	1.13
27:	70,	1,14	10y	2.0	1.9	2.3	1.9	1.0	3.4	3.8	17.390	1.13
14:	70,	1,14	10y	2.0	1.9	2.3	1.9	1.0	3.4	3.8	17.409	1.13
21:	70,	1,14	10y	2.0	1.9	2.3	1.9	1.0	3.4	3.8	17.435	1.13
0:	70,	1,14	10y	1.0	0.9	1.2	0.9	0.9	1.9	2.1	9.657	1.14

ET/ms: Measured execution time.

SET: Measured execution time scaled to ET units.

Ratio: SET/ET: Ratio of predicted to measured time.

Ideal value is 1.

For Chemora it needs to be close to the same value for a calculation.

How much faster?

	Tile	Iter	L-FP	L-Dat	Issue	Data	Lat	ET	SET	ET/ms	Ratio	
23:	70,	1,14	10y!	2.0	1.9	2.3	1.9	1.0	3.4	3.8	17.357	1.13
2:	70,	1,14	10y!	2.0	1.9	2.3	1.9	1.0	3.4	3.8	17.360	1.13
27:	70,	1,14	10y!	2.0	1.9	2.3	1.9	1.0	3.4	3.8	17.390	1.13
14:	70,	1,14	10y!	2.0	1.9	2.3	1.9	1.0	3.4	3.8	17.409	1.13
21:	70,	1,14	10y!	2.0	1.9	2.3	1.9	1.0	3.4	3.8	17.435	1.13
0:	70,	1,14	10y!	1.0	0.9	1.2	0.9	0.9	1.9	2.1	9.657	1.14

Fundamental	Current
Limiters	Limiters

(Fundamental) L-FP omits all non-FP instructions.

L-Data omits inter-kernel data.

Specialized Tables

Pre-Compilation.

Results of Compilation.

Number of registers, and *spill* storage.

Maximum number of threads.

Results of Run.

Measured execution time.

Sequence of Tables

Several table printed per calculation.

Code Generation for a one-kernel mapping.

Code Generation for a fixed n -kernel mapping.

Code Generation for last (not best) mapping in optimization.

Code Generation for best mapping.

In addition optimization progress shown:

```
0 -> 2 Sc    -0.4 = ( 0,1546,1059)   20.1 ( 47, 27)  - 18 - 18 + 18 Ac
0 -> 27 Sc   -1.4 = ( 0,1599,1060)   18.7 ( 29, 43)  -  2 -  2 +  2 Ac
-----
ET
```

One- and n -Kernel Comparison

```
INFO (CaCUDALib): Chemora Code Generation Report for DEVICE_ML_BSSN_RHS_Dalpha3_2_etc
:      Tile  Iter B   FP/Insn  IK   Elts  Issue  Data  Lat   ET  OV  H/I
0:   70, 1,14  10 y   88/ 99   0  30.8  15.9  10.2  29.3  38.5  308 1.80
Total of 1 ( 0)   88/ 99   0  30.8  15.9  10.2  29.3  38.5
```

-- Used 0 MiB of storage here, total 2 MiB.

```
INFO (CaCUDALib): Chemora Code Generation Report for DEVICE_ML_BSSN_RHS_Dalpha3_2_etc
:      Tile  Iter B   FP/Insn  IK   Elts  Issue  Data  Lat   ET  OV  H/I
0:   70, 1,14  10 y   8/  9   0   2.8   1.2   0.9   0.9   1.9   0  1.80
1:   70, 1,14  10 y   8/  9   0   2.8   1.2   0.9   0.9   1.9   0  1.80
2:   70, 1,14  10 y   8/  9   0   2.8   1.2   0.9   0.9   1.9   0  1.80
3:   70, 1,14  10 y   8/  9   0   2.8   1.2   0.9   0.9   1.9   0  1.80
4:   70, 1,14  10 y   8/  9   0   2.8   1.2   0.9   0.9   1.9   0  1.80
5:   70, 1,14  10 y   8/  9   0   2.8   1.2   0.9   0.9   1.9   0  1.80
6:   70, 1,14  10 y   8/  9   0   2.8   1.2   0.9   0.9   1.9   0  1.80
7:   70, 1,14  10 y   8/  9   0   2.8   1.2   0.9   0.9   1.9   0  1.80
8:   70, 1,14  10 y   8/  9   0   2.8   1.2   0.9   0.9   1.9   0  1.80
9:   70, 1,14  10 y   8/  9   0   2.8   1.2   0.9   0.9   1.9   0  1.80
10:  70, 1,14  10 y   8/  9   0   2.8   1.2   0.9   0.9   1.9   0  1.80
Total of 11 ( 0)  88/ 99   0  30.8  12.7  10.2   9.8  20.4
```

Optimization Progress

0 -> 2 Sc	-0.4 = (0,1546,1059)	20.1 (47, 27)	- 18 - 18 + 18 Ac
0 -> 27 Sc	-1.4 = (0,1599,1060)	18.7 (29, 43)	- 2 - 2 + 2 Ac

Last Mapping and Best Mapping

```
INFO (CaCUDALib): Chemora Code Generation Report for DEVICE_ML_BSSN_RHS_Dalpha3_2_etc
:      Tile  Iter B    FP/Insn  IK    Elts  Issue  Data  Lat    ET  OV  H/I
0:   70, 1,14  10 y    51/ 71  27    46.0   10.6  15.3  27.7   36.8  252 1.80
2:   70,14, 1  10 z G     0/  1  1     2.0    0.0   0.7   0.8    1.1   0 1.00
6:   70,14, 1  10 z G     0/  1  1     2.0    0.0   0.7   0.8    1.1   0 1.00
8:   70, 1,14  10 y     7/  4  9    12.8    1.1   4.3   2.5    5.6   0 1.40
13:  70, 1,14  10 y     8/ 10  1     4.8    1.2   1.6   1.7    2.7   0 1.80
22:  70, 1,14  10 y     7/  3  8    11.3    1.1   3.8   2.5    5.1   0 1.30
27:  70, 1,14  10 y    15/  9 11    17.4    2.4   5.8   3.8    8.0  11 1.70
Total of 7 ( 2935)    88/ 99  58    96.3   16.4  32.0  39.9   60.5
```

```
INFO (CaCUDALib): Chemora Code Generation Report for DEVICE_ML_BSSN_RHS_Dalpha3_2_etc
:      Tile  Iter B    FP/Insn  IK    Elts  Issue  Data  Lat    ET  OV  H/I
0:   70, 1,14  10 y     8/  9  0     2.8    1.2   0.9   0.9    1.9   0 1.80
2:   70, 1,14  10 y    16/ 18  0     5.6    2.3   1.9   1.0    3.4   0 1.80
14:  70, 1,14  10 y    16/ 18  0     5.6    2.3   1.9   1.0    3.4   0 1.80
21:  70, 1,14  10 y    16/ 18  0     5.6    2.3   1.9   1.0    3.4   0 1.80
23:  70, 1,14  10 y    16/ 18  0     5.6    2.3   1.9   1.0    3.4   0 1.80
27:  70, 1,14  10 y    16/ 18  0     5.6    2.3   1.9   1.0    3.4   0 1.80
Total of 6 (   0)    88/ 99  0    30.8   12.7  10.2  5.7   18.7
```


Best Mapping, and Post-Run Table

```
INFO (CaCUDALib): Chemora Code Generation Report for DEVICE_ML_BSSN_RHS_Dalpha3_2_etc
:      Tile  Iter B      FP/Insn  IK    Elts  Issue  Data  Lat    ET  OV  H/I
0:    70, 1,14 10 y      8/   9   0    2.8   1.2  0.9  0.9   1.9  0 1.80
2:    70, 1,14 10 y     16/  18   0    5.6   2.3  1.9  1.0   3.4  0 1.80
14:   70, 1,14 10 y     16/  18   0    5.6   2.3  1.9  1.0   3.4  0 1.80
21:   70, 1,14 10 y     16/  18   0    5.6   2.3  1.9  1.0   3.4  0 1.80
23:   70, 1,14 10 y     16/  18   0    5.6   2.3  1.9  1.0   3.4  0 1.80
27:   70, 1,14 10 y     16/  18   0    5.6   2.3  1.9  1.0   3.4  0 1.80
Total of 6 ( 0)      88/ 99   0   30.8  12.7  10.2  5.7   18.7
```

```
INFO (CaCUDALib): Chemora Code Execution Report for DEVICE_ML_BSSN_RHS_Dalpha3_2_etc
:      Tile Iter  L-FP L-Dat  Issue  Data  Lat    ET    SET    ET/ms Ra-
tio
23:   70, 1,14 10y    2.0  1.9    2.3  1.9  1.0    3.4    3.8   17.357  1.13
2:    70, 1,14 10y    2.0  1.9    2.3  1.9  1.0    3.4    3.8   17.360  1.13
27:   70, 1,14 10y    2.0  1.9    2.3  1.9  1.0    3.4    3.8   17.390  1.13
14:   70, 1,14 10y    2.0  1.9    2.3  1.9  1.0    3.4    3.8   17.409  1.13
21:   70, 1,14 10y    2.0  1.9    2.3  1.9  1.0    3.4    3.8   17.435  1.13
0:    70, 1,14 10y    1.0  0.9    1.2  0.9  0.9    1.9    2.1    9.657  1.14
Total of 6 ( 30)  11.0  10.2  12.7  10.2  5.7    18.7  21.2   96.609  1.13
```