esiwace2
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Accelerating Chemistry Modules in Atmospheric Models using GPUs

Christian Guzman Ruiz, Mario C. Acosta, Matthew Dawson*, Oriol Jorba, Carlos Pérez García-Pando, Kim Serradell

Earth Sciences Department
***Currently at National Center for Atmospheric Research (NCAR)**

24/03/2020

NVIDIA GTC 2020

Outline

- Introduction:
  - BSC and the Computational Earth Science group
  - Motivation
  - Tools: MONARCH & CAMP
- Implementations:
  - GPU
  - Multi-cells & GPU
- Conclusions and future work

> Preliminary work on exploiting GPU capacity <

Motivation

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

BSC Departments

To influence the way machines are built, programmed and used: programming models, performance tools, Big Data, Artificial Intelligence , computer architecture, energy efficiency

To develop and implement global and regional state-of-the-art models for short-term air quality forecast and long-term climate applications

To understand living organisms by means of theoretical and computational methods (molecular modeling, genomics, proteomics)

To develop scientific and engineering software to efficiently exploit super-computing capabilities (biomedical, geophysics, atmospheric, energy, social and economic simulations)

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Earth Sciences

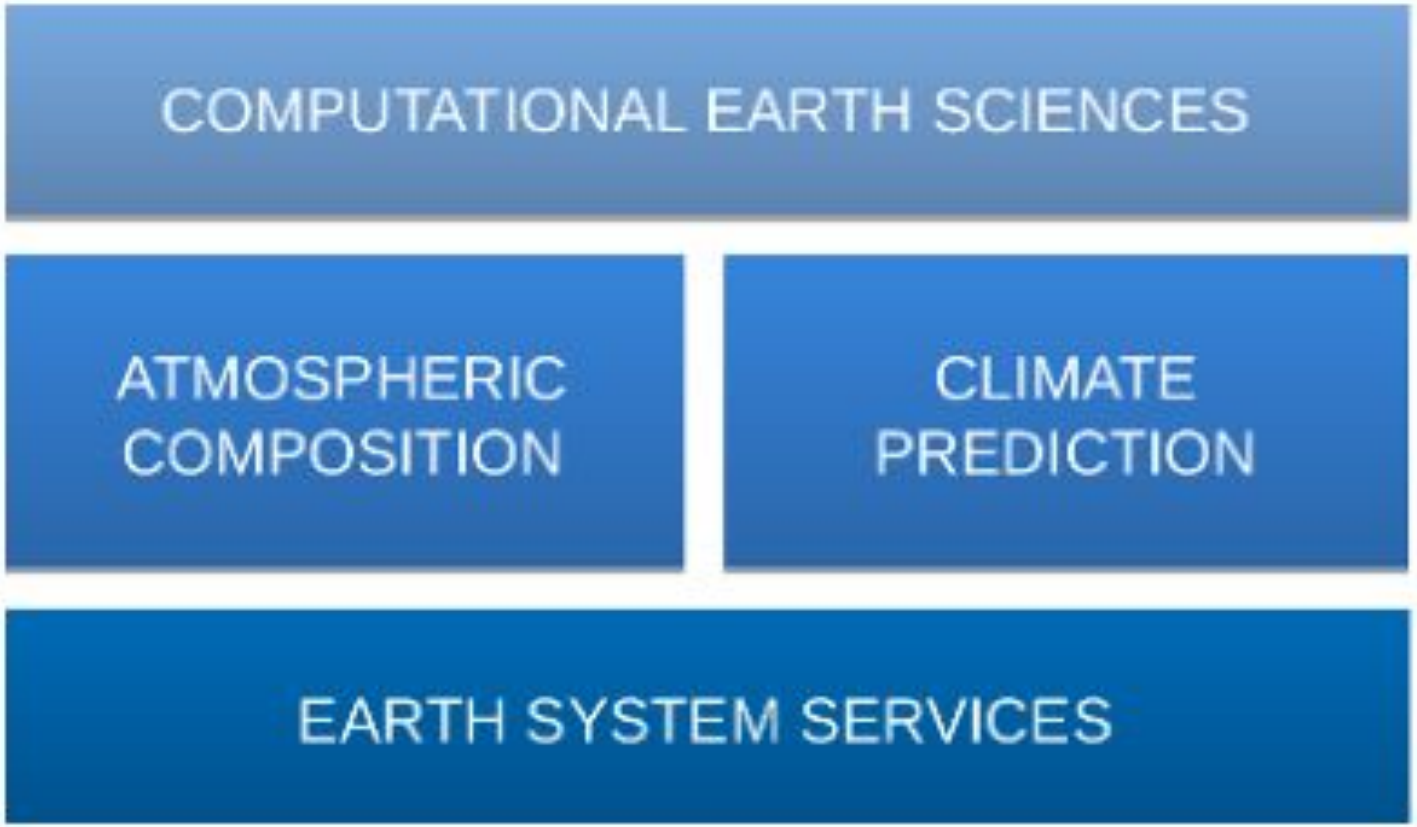

Computational Earth Science

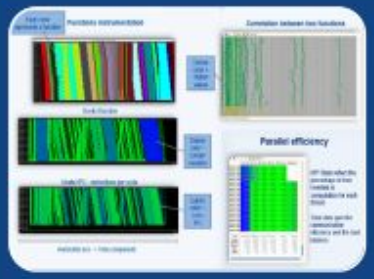

Performance Team

- Provide HPC Services (profiling, code audit, ...) to find main bottlenecks of our operational models
- Research and apply new computational methods for current and new platforms

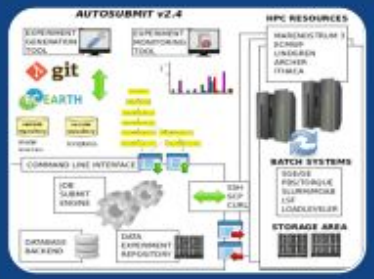

Models and Workflows Team

- Development of HPC user-friendly software framework
- Support the development of atmospheric research software

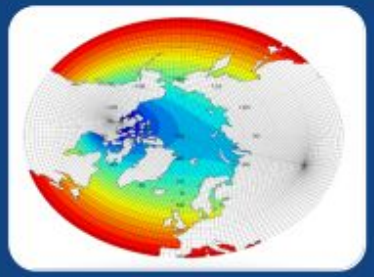

Data and Diagnostics Team

- Big Data in Earth Sciences
- Provision of data services
- Visualization

Performance Team

- Knowledge of the mathematical and computational aspects of Earth System Applications
- Knowledge of the specific HPC needs of Earth Systems Applications
- Research of HPC methods specifically designed for Earth Systems Applications

Atmospheric models

Atmospheric models are a **mathematical representation** of atmospheric water, gas, and aerosol cycles.

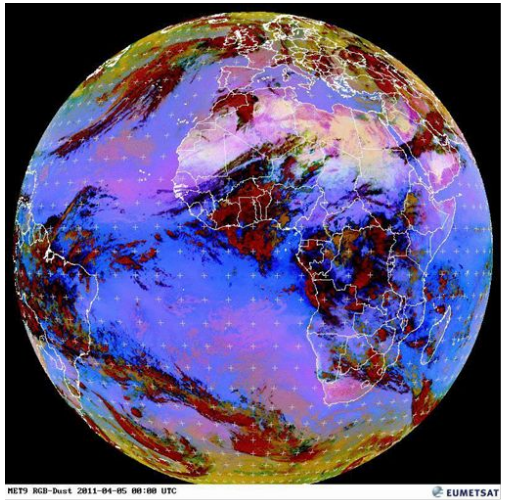

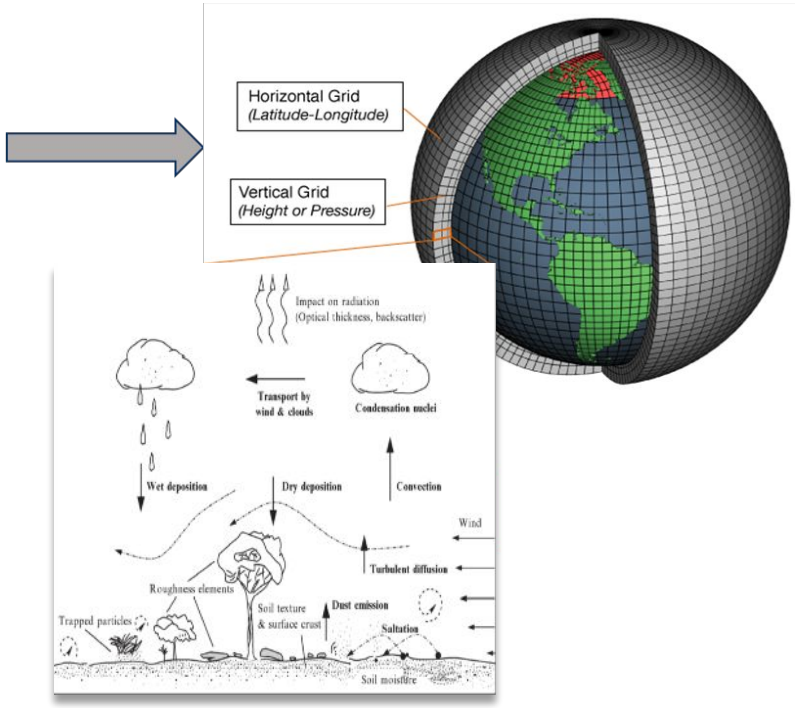

Chemical mechanism

Reaction Number	Reaction	Rate Cons
	Inorganic Chemistry	
(1)	$NO_2 + h\nu \xrightarrow{O_2} NO + O(^3P)$	J_{NO_2}
(2)	$NO_3 + h\nu \longrightarrow 0.89NO_2 + 0.89O(^3P) + 0.11NO$	J_{NO_3}
(3)	$HNO_2 + h\nu \longrightarrow OH + NO$	J_{HNO_2}
(4)	$HNO_3 + h\nu \longrightarrow OH + NO_2$	J_{HNO_3}
(5)	$HNO_4 + h\nu \longrightarrow HO_2 + NO_2$	J_{HNO_4}
(6)	$O_3 + h\nu \longrightarrow O(^3P)$	J_{O3a}
(7)	$O_3 + h\nu \longrightarrow O(^1D)$	J_{O3b}
(8)	$H_2O_2 + h\nu \longrightarrow 2OH$	$J_{H_2O_2}$
(9)	$O(^1D) + O_2 \longrightarrow O(^3P) + O_2$	3.2×10^{-11} exp (7(
(10)	$O(^1D) + N_2 \longrightarrow O(^3P) + N_2$	1.8×10^{-11} exp (11
(11)	$O(^1D) + H_2O \longrightarrow 2OH$	2.2×10^{-10}
(12)	$O(^3P) + O_2 \xrightarrow{M} O_3$	$F(6.0(-34), 2.3, 0.0$
(13)	$O(^3P) + O_3 \longrightarrow O_2 + O_2$	8.0×10^{-12} exp (−
(14)	$O(^3P) + NO_2 \longrightarrow NO$	6.5×10^{-12} exp (−
(15)	$O(^3P) + NO_2 \xrightarrow{M} NO_3$	$F(9.0(-32), 2.0, 2.2$
(16)	$O(^3P) + NO \xrightarrow{M} NO_2$	$F(9.0(-32), 1.5, 3.0$
(17)	$O_3 + NO \longrightarrow NO_2$	2.0×10^{-12} exp (−
(18)	$O_3 + NO_2 \longrightarrow NO_3$	1.2×10^{-13} exp (−
(19)	$O_3 + OH \longrightarrow HO_2$	1.6×10^{-12} exp (−
(20)	$O_3 + HO_2 \longrightarrow OH$	1.1×10^{-14} exp (−
(21)	$OH + H_2 \longrightarrow HO_2 + H_2O$	5.5×10^{-12} exp (−
(22)	$OH + NO \xrightarrow{M} HNO_2$	$F(7.0(-31), 2.6, 3.6$
(23)	$OH + NO_2 \xrightarrow{M} HNO_3$	$F(2.5(-30), 4.4, 1.6$
(24)	$OH + NO_3 \longrightarrow HO_2 + NO_2$	2.2×10^{-11}
(25)	$OH + HNO_3 \longrightarrow NO_3$	1.8×10^{-11} exp (−
(26)	$OH + HNO_3 \xrightarrow{M} NO_3$	$k_a = [M]k_b/(1 + [M$ $k_a = 7.2 \times 10^{-15}$ e $k_b = 1.9 \times 10^{-33}$ e $k_c = 4.1 \times 10^{-16}$ e
(27)	$OH + HNO_4 \longrightarrow NO_2$	1.3×10^{-12} exp (38
(28)	$OH + HO_2 \longrightarrow H_2O + O_2$	4.8×10^{-11} exp (28
(29)	$OH + H_2O_2 \longrightarrow HO_2$	2.9×10^{-12} exp (−
(30)	$HO_2 + HO_2 \xrightarrow{M} H_2O_2$	$(k_d + [M]k_e)$ $k_d = 2.3 \times 10^{-13}$ e $k_e = 1.7 \times 10^{-33}$ e
(31)	$HO_2 + HO_2 + H_2O \xrightarrow{M} H_2O_2$	$k_{30} \times 1.4 \times 10^{-21}$
(32)	$HO_2 + NO \longrightarrow OH + NO_2$	3.5×10^{-12} exp (28
(33)	$HO_2 + NO_2 \xrightarrow{M} HNO_4$	$F(1.8(-31), 3.2, 4.7$
(34)	$HO_2 + NO_2 \longrightarrow HNO_2$	5.0×10^{-16}
(35)	$HNO_4 \xrightarrow{M} HO_2 + NO_2$	$k_{33} \times 4.76 \times 10^{26}$
(36)	$NO_3 + NO \longrightarrow 2NO_2$	1.5×10^{-11} exp (13
(37)	$NO_3 + NO_2 \longrightarrow NO + NO_2$	4.5×10^{-14} exp (−
(38)	$NO_3 + NO_2 \xrightarrow{M} N_2O_5$	$F(2.2(-30), 3.9, 1.5$
(39)	$NO_3 + NO_3 \longrightarrow 2NO_2 + O_2$	8.5×10^{-13} exp (−
(40)	$NO_3 + HO_2 \longrightarrow .3HNO_3 + .7NO_2 + .7OH$	3.5×10^{-12}
(41)	$N_2O_5 + H_2O \longrightarrow 2HNO_3$	2.0×10^{-21}
(42)	$N_2O_5 \xrightarrow{M} NO_3 + NO_2$	$k_{38} \times 3.7 \times 10^{26}$ e
(43)	$NO + NO + O_2 \xrightarrow{O_2} 2NO_2$	3.3×10^{-39} exp (53
(44)	$CO + OH \xrightarrow{O_2} HO_2$	$1.5 \times 10^{-13}(1 + .6$
(45)	$SO_2 + OH \longrightarrow H_2SO_4 + HO_2$	$F(3.0(-31), 3.3, 1.5$

Reaction Number	Reaction	Rate Constant, $k^\dagger$	Note
	Carbonyl Chemistry		
(50)	$HCHO + h\nu \xrightarrow{O_2} 2HO_2 + CO$	J_{HCHOa}	13,18
(51)	$HCHO + h\nu \longrightarrow CO$	J_{HCHOb}	13,18
(52)	$HCHO + OH \xrightarrow{O_2} HO_2 + CO$	1.0×10^{-11}	
(53)	$HCHO + NO_3 \xrightarrow{O_2} HNO_3 + HO_2 + CO$	3.4×10^{-13} exp (−	
(54)	$ALD2 + h\nu \xrightarrow{2O_2} CH_3O_2 + HO_2 + CO$	J_{ALD2}	
(55)	$ALD2 + OH \longrightarrow C_2O_3$	5.6×10^{-12} exp (2	
(56)	$ALD2 + NO_3 \xrightarrow{O_2} C_2O_3 + HNO_3$	1.4×10^{-12} exp (−	
(57)	$AONE + h\nu \xrightarrow{2O_2} C_2O_3 + CH_3O_2$	J_{AONE}	
(58)	$AONE + OH \longrightarrow ANO2$	$T^2 5.3 \times 10^{-18}$ exp	
(59)	$MGLY + h\nu \longrightarrow C_2O_3 + CO + HO_2$	$9.64 \times J_{HCHOa}$	
(60)	$MGLY + OH \longrightarrow XO_2 + C_2O_3$	1.7×10^{-11}	
(61)	$MGLY + NO_3 \longrightarrow HNO_3 + C_2O_3 + CO$	1.4×10^{-12} exp (−	
	Olefin chemistry		
(62)	$ETH + O_3 \longrightarrow HCHO + 0.22HO_2 + 0.12OH + 0.24CO + 0.24CH_2O_2 + 0.52HCOOH$	1.2×10^{-14} exp (−	
(63)	$ETH + OH \longrightarrow XO_2 + 1.56HCHO + HO_2 + 0.22ALD2$	$F(1.0(-28), 0.8, 8.$	
(64)	$OLET + O_3 \longrightarrow 0.57HCHO + 0.47ALD2 + 0.33OH + 0.26HO_2 + 0.08H_2 + 0.07CH_3O_2 + 0.06ETHP + 0.03RO_2 + 0.13C_2O_3 + 0.04MGLY + 0.03CH_3OH + 0.06CH_4 + 0.01C_2H_6 + 0.31CO + 0.22CO_2 + 0.22HCOOH + 0.09RCOOH - 1.06PAR$	4.2×10^{-15} exp (−	
(65)	$OLEI + O_3 \longrightarrow 1.03ALD2 + 0.07AONE + 0.60OH + 0.22HO_2 + 0.10CH_3O_2 + 0.05ETHP + 0.09RO_2 + 0.11ANO_2 + 0.19C_2O_3 + 0.07MGLY + 0.04CH_3OH + 0.08CH_4 + 0.01C_2H_6 + 0.30CO + 0.18CO_2 + 0.16RCOOH - 2.26PAR$	8.9×10^{-16} exp (−	
(66)	$OLET + OH \longrightarrow XO_2 + HO_2 + HCHO + ALD2 - PAR$	5.8×10^{-12} exp (4	
(67)	$OLEI + OH \longrightarrow XO_2 + HO_2 + 0.23AONE + 1.77ALD2 - 2.23PAR$	2.9×10^{-11} exp (2	
(68)	$OLET + NO_3 \longrightarrow NAP$	3.1×10^{-13} exp (−	
(69)	$OLEI + NO_3 \longrightarrow NAP$	2.5×10^{-12}	
	Aromatic Chemistry		
(70)	$TOL + OH \longrightarrow 0.08XO_2 + 0.2HO_2 + 0.12CRES + 0.8TO_2$	2.1×10^{-12} exp (3	
(71)	$XYL + OH \longrightarrow 0.5XO_2 + 0.55HO_2 + 0.8MGLY + 1.1PAR + 0.45TO_2 + 0.05CRES$	1.7×10^{-11} exp (1	
(72)	$TO_2 + NO \longrightarrow 0.95(NO_2 + OPEN + HO_2) + 0.05ONIT$	8.1×10^{-12}	
(73)	$CRES + OH \longrightarrow 0.4CRO + 0.6XO_2 + 0.6HO_2 + 0.3OPEN$	4.1×10^{-11}	
(74)	$CRES + NO_3 \longrightarrow CRO + HNO_3$	2.2×10^{-11}	
(75)	$CRO + NO_2 \longrightarrow ONIT$	1.4×10^{-11}	
(76)	$OPEN + OH \longrightarrow XO_2 + C_2O_3 + 2CO + 2HO_2 + HCHO$	3.0×10^{-11}	
(77)	$OPEN + h\nu \longrightarrow C_2O_3 + CO + HO_2$	$9.04 \times J_{HCHOa}$	
(78)	$OPEN + O_3 \longrightarrow 0.03ALD2 + 0.62C_2O_3 + 0.7HCHO + 0.69CO + 0.08OH + 0.03XO_2 + 0.76HO_2 + 0.2MGLY$	5.4×10^{-17} exp (−	
	Isoprene Chemistry		
(79)	$ISOP + OH \longrightarrow ISOPP + 0.08XO2$	2.55×10^{-11} exp (	

Reaction Number	Reaction	Rate Constant, $k^\dagger$	Note
	Organic Hydroperoxides		
(86)	$CH_3OOH + h\nu \xrightarrow{O_2} HCHO + HO_2 + OH$	J_{CH_3OOH}	11,18
(87)	$ETHOOH + h\nu \longrightarrow ALD2 + HO_2 + OH$	same as reaction (86)	9,11
(88)	$ROOH + h\nu \longrightarrow OH + 0.4XO_2 + 0.74AONE + 0.3ALD2 + 0.1ETHP + 0.9HO_2 - 1.98PAR$	same as reaction (86)	9,11
(89)	$CH_3OOH + OH \longrightarrow 0.7CH_3O_2 + 0.3HCHO + 0.3OH$	3.8×10^{-12} exp (200/T)	1,11
(90)	$ETHOOH + OH \longrightarrow 0.7ETHP + 0.3ALD2 + 0.3OH$	3.8×10^{-12} exp (200/T)	9,11
(91)	$ROOH + OH \longrightarrow 0.77RO_2 + 0.19MGLY + 0.04ALD2 + 0.23OH - 0.42PAR$	3.8×10^{-12} exp (200/T)	9,11
	Organic Nitrates		
(92)	$ONIT + OH \longrightarrow NAP$	1.6×10^{-11} exp (−540/T)	11,12
(93)	$ONIT + h\nu \longrightarrow NO_2 + 0.41XO_2 + 0.74AONE + 0.3ALD2 + 0.1ETHP + 0.9HO_2 - 1.98PAR$	J_{ONIT}	11,18
(94)	$C_2O_3 + NO_2 \longrightarrow PAN$	$F(9.7(-29), 5.6, 9.3(-12), 1.5)$	1,13
(95)	$PAN \longrightarrow C_2O_3 + NO_2$	$k_{94} 1.1 \times 10^{28}$ exp (−14000/T)	1,13
	Alkyl and Acyl Peroxy Radical Chemistry		
(96)	$CH_3O_2 + NO \longrightarrow HCHO + HO_2 + NO_2$	3.0×10^{-12} exp (280/T)	1,11
(97)	$ETHP + NO \longrightarrow ALD2 + HO_2 + NO_2$	2.6×10^{-12} exp (365/T)	1,11
(98)	$RO_2 + NO \longrightarrow 0.16ONIT + 0.84NO_2 + 0.34XO_2 + 0.62AONE + 0.25ALD2 + 0.08ETHP + 0.76HO_2 - 1.68PAR$	4.0×10^{-12}	8,11
(99)	$C_2O_3 + NO \xrightarrow{O_2} CH_3O_2 + NO_2 + CO_2$	5.3×10^{-12} exp (360/T)	1,10
(100)	$ANO2 + NO \longrightarrow NO_2 + C_2O_3 + HCHO$	4.0×10^{-12}	8,11
(101)	$NAP + NO \longrightarrow 1.5NO_2 + 0.5HCHO + 0.5ALD2 + 0.5ONIT + 0.5HO_2 - PAR$	4.0×10^{-12}	8,11
(102)	$ISOPP + NO \longrightarrow 0.09ONIT + 0.91NO_2 + 0.91HO_2 + 0.63 HCHO + 0.91ISOPRD + 0.18PAR$	4.0×10^{-12}	8,15
(103)	$ISOPN + NO \longrightarrow NO_2 + 0.8ALD2 + 0.8ONIT + 0.8HO_2 + 0.2ISOPRD + 0.2NO_2 + 1.6PAR$	4.0×10^{-12}	8,15
(104)	$ISOPO_2 + NO \longrightarrow NO_2 + HO_2 + 0.59CO + 0.55ALD2 + 0.25HCHO + 0.34MGLY + 0.63AONE$	4.0×10^{-12}	8,15
(105)	$XO_2 + NO \longrightarrow NO_2$	4.0×10^{-12}	8,13
(106)	$CH_3O_2 + NO_3 \longrightarrow HCHO + HO_2 + NO_2$	1.1×10^{-12}	7,11
(107)	$ETHP + NO_3 \longrightarrow ALD2 + HO_2 + NO_2$	2.5×10^{-12}	7,11
(108)	$RO_2 + NO_3 \longrightarrow NO_2 + 0.4XO_2 + 0.74AONE + 0.3ALD2 + 0.1ETHP + 0.9HO_2 - 1.98PAR$	2.5×10^{-12}	7,11
(109)	$C_2O_3 + NO_3 \longrightarrow CH_3O_2 + NO_2$	4.0×10^{-12}	8,11
(110)	$ANO2 + NO_3 \longrightarrow NO_2 + C_2O_3 + HCHO$	1.2×10^{-12}	8,11
(111)	$NAP + NO_3 \longrightarrow 1.5NO_2 + 0.5HCHO + 0.5ALD2 + 0.5ONIT + 0.5HO_2 - PAR$	4.0×10^{-12}	8,11
(112)	$XO_2 + NO_3 \longrightarrow NO_2$	2.5×10^{-12}	7,11
(113)	$CH_3O_2 + HO_2 \longrightarrow CH_3OOH$	3.8×10^{-13} exp (800/T)	1,11
	OH	7.5×10^{-13} exp (700/T)	1,11
		1.7×10^{-13} exp (1300/T)	8,11
	)OH + O_3)	4.5×10^{-13} exp (1000/T)	1,10
		1.2×10^{-13} exp (1300/T)	8,11
		1.7×10^{-13} exp (1300/T)	8,11
		1.7×10^{-13} exp (1300/T)	8,15
	+ 2PAR	1.7×10^{-14} exp (1300/T)	8,15
	H	1.7×10^{-13} exp (1300/T)	8,15
		1.7×10^{-13} exp (1300/T)	8,11
	Parameterized Permutation Reactions		
	$0.32HO_2 + 0.34CH_3OH$	$k_i^{(1)}$, $i = CH_3O_2$	11,16
	$HO_2 + 0.2C_2H_6$	$k_i^{(1)}$, $i = ETHP$	11,16
	$LD2 + 0.57AONE + 0.06ETHP$	$k_i^{(1)}$, $i = RO_2$	11,16
			11,16
		$k_i^{(1)}$, $i = C_2O_3$	11,16
(127)	$ANO2 \longrightarrow 0.7(C_2O_3 + HCHO) + 0.15(MGLY + AONE)$	$k_i^{(1)}$, $i = ANO2$	11,16
(128)	$NAP \longrightarrow 0.5(NO_2 + HCHO + ALD2 + ONIT) - PAR$	$k_i^{(1)}$, $i = NAP$	11,16

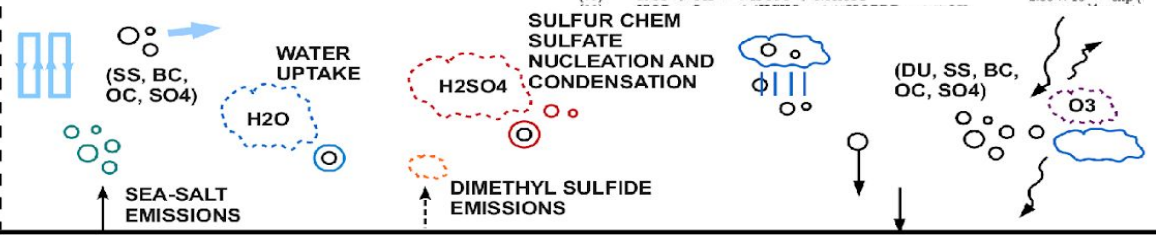

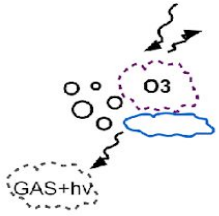

Computational design

- Programming languages:
  - FORTRAN, C
- Parallelization strategy:
  - Domain decomposition
  - MPI, MPI+OpenMP
  - New approaches: MPI + OpenACC | CUDA

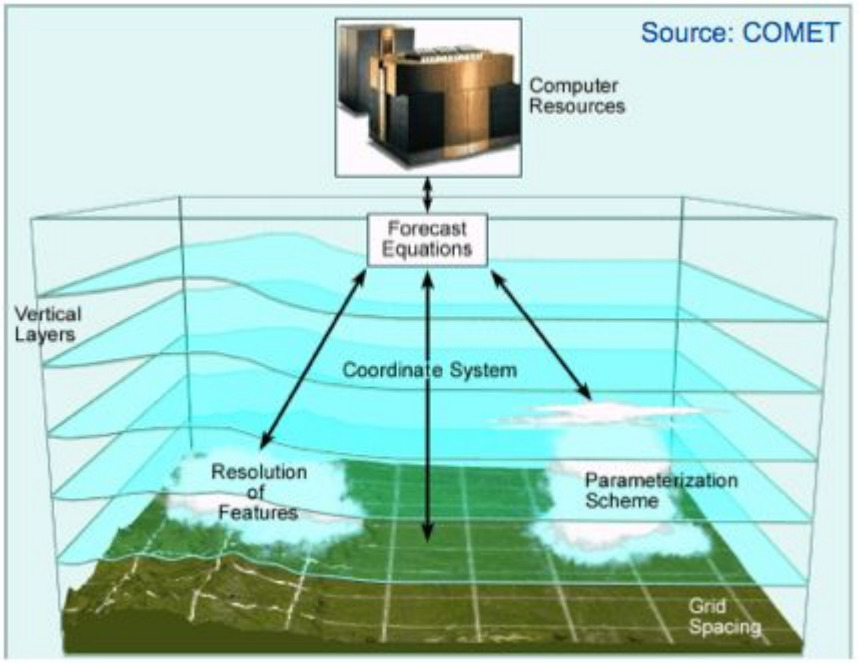

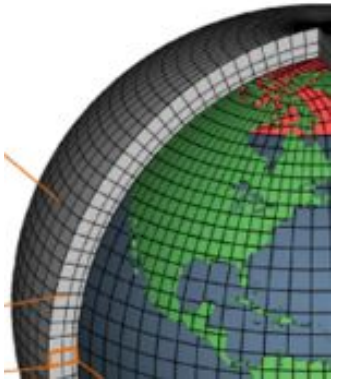

Parallelization techniques

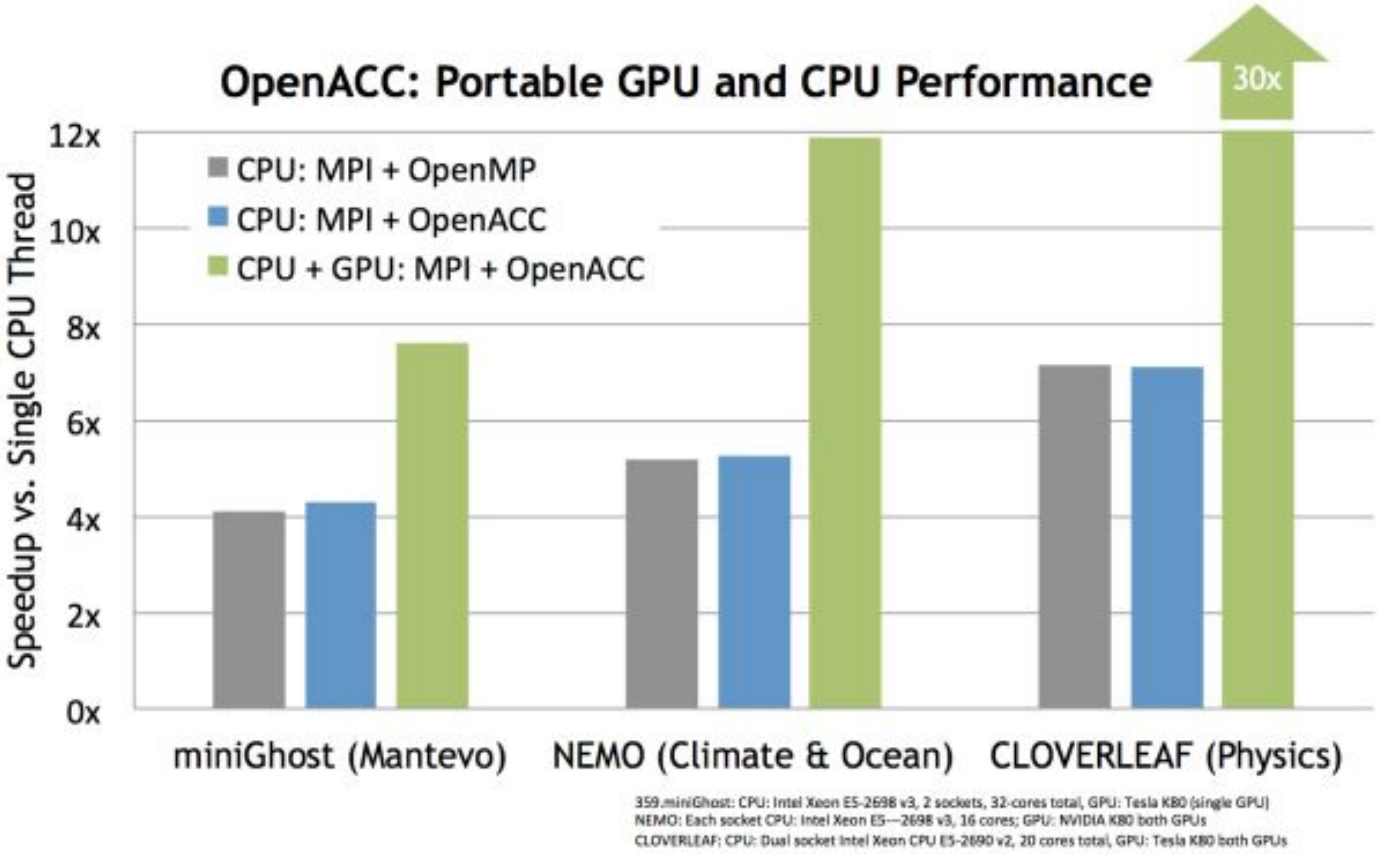

Performance Portability from GPUs to CPUs with OpenACC, Nvidia Developer Blog

Chemistry in the GPU: CUDA

Configuration	Median CPU exec time (s)	Median accelerated exec time (s)	Performance over CPU
Intel Xeon X5650 + M2070	4.502	0.999	4.50×
Intel Xeon E5-2680 v3 + K80	1.476	0.283	5.21×
IBM POWER8 + P100	3.040	0.149	20.40×

Configuration	MPI Processes	CPU exec time (s)	Accelerated exec time (s)	Performance over CPU
2 × 6-core Intel Xeon X5650 + 2 × NVIDIA M2070	2 MPI processes	5199	2358	2.27 ×
	12 MPI processes	1388	1368	1.01 ×
2 × 12-core Intel E5-2680 v3 + 2 × NVIDIA K80	4 MPI processes	7362	3384	2.17 ×
	24 MPI processes	1756	1473	1.19 ×
2 × 10-core IBM POWER8 + 4 × NVIDIA P100	4 MPI processes	2294	918	2.50 ×
	20 MPI Processes	814	437	1.86 ×

Michail Alvanos and Theodoros Christoudia, GPU-accelerated atmospheric chemical kinetics in the ECHAM/MESSy (EMAC) Earth system model , 2017

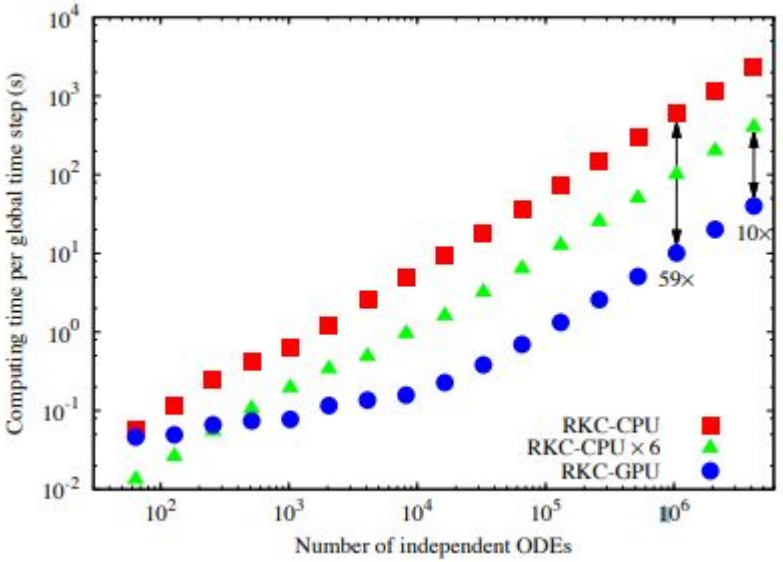

Kyle E. Niemeyera,b,1, Chih-Jen Sungb, Accelerating moderately stiff chemical kinetics in reactive-flow simulations using GPUs, 2018

...and more

Our goal

- **Challenges** Addressed
  - Siloed treatment of physical/chemical processes
  - Huge heterogeneous codebase
  - Efficient solving of complex physical/chemical systems

- **How** we do it
  - Integrated stand-alone chemistry solver
  - Standardized description of physical/chemical processes
  - Porting high-cost functions to GPUs
  - Simultaneous solving of multiple grid-cells

Tools: MONARCH & CAMP

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

MONARCH: Multiscale On-line Atmosphere Chemistry Model

MONARCH

~20%

NCEP/NMMB

→ *Janjic and Gall (NCAR/TN 2012)*
→ *Janjic and Vasic (EGU2012)*
→ *Janjic et al. (MWR 2011)*
→ *(...)*

~80%

BSC/Chemistry

AEROSOLS

→ *Pérez et al. (ACP 2011)*
→ *Haustein et al. (ACP 2012)*
→ *Spada et al. (ACP 2013)*
→ *Spada et al. (AE 2014)*
→ *Spada (2015)*
→ *DiTomaso et al. (GMD 2017)*

VOLCANIC ASH

→ *Martí et al. (ACP 2016)*

GAS-PHASE CHEMISTRY

→ *Jorba et al. (JGR 2012)*
→ *Badia and Jorba (AE 2014)*
→ *Badia et al. (GMD 2017)*

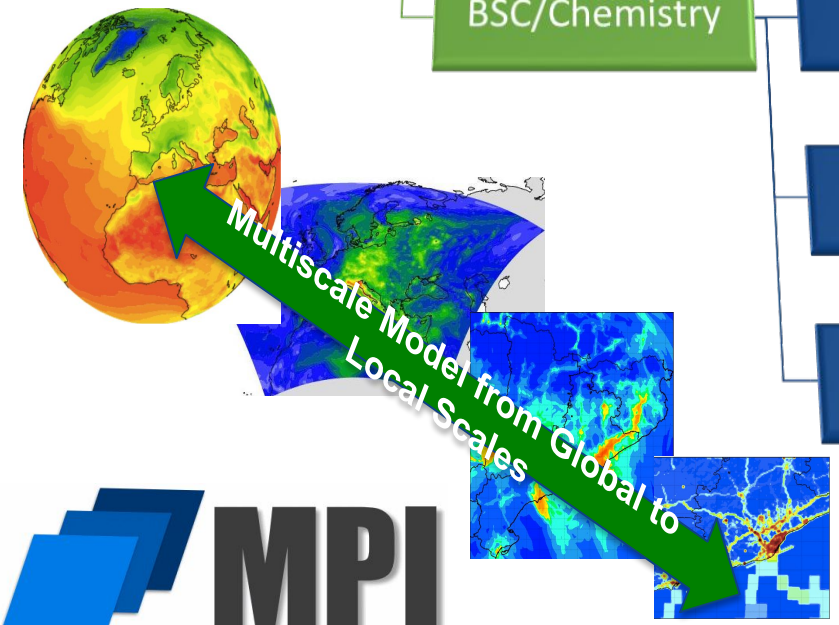

MPI

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Atmospheric chemistry - Classical approach

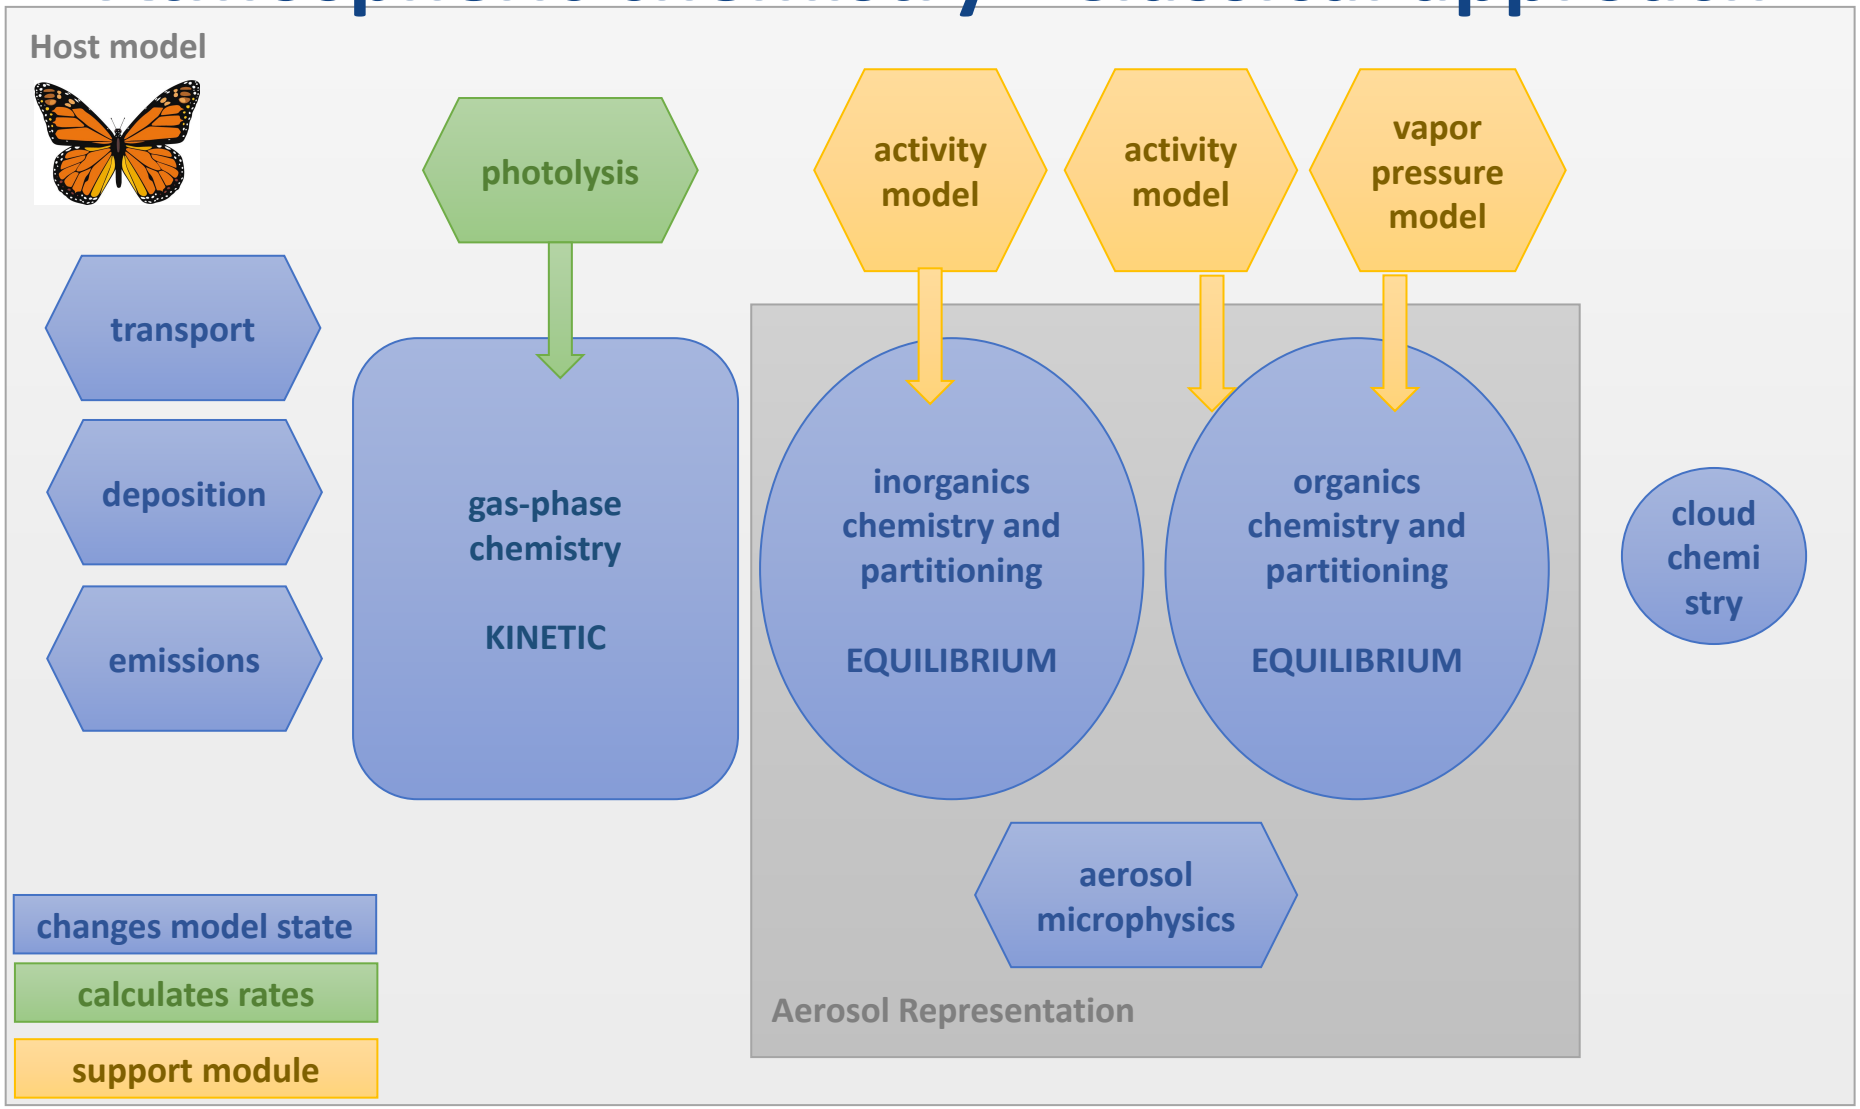

CAMP: Chemistry Across Multiple Phases

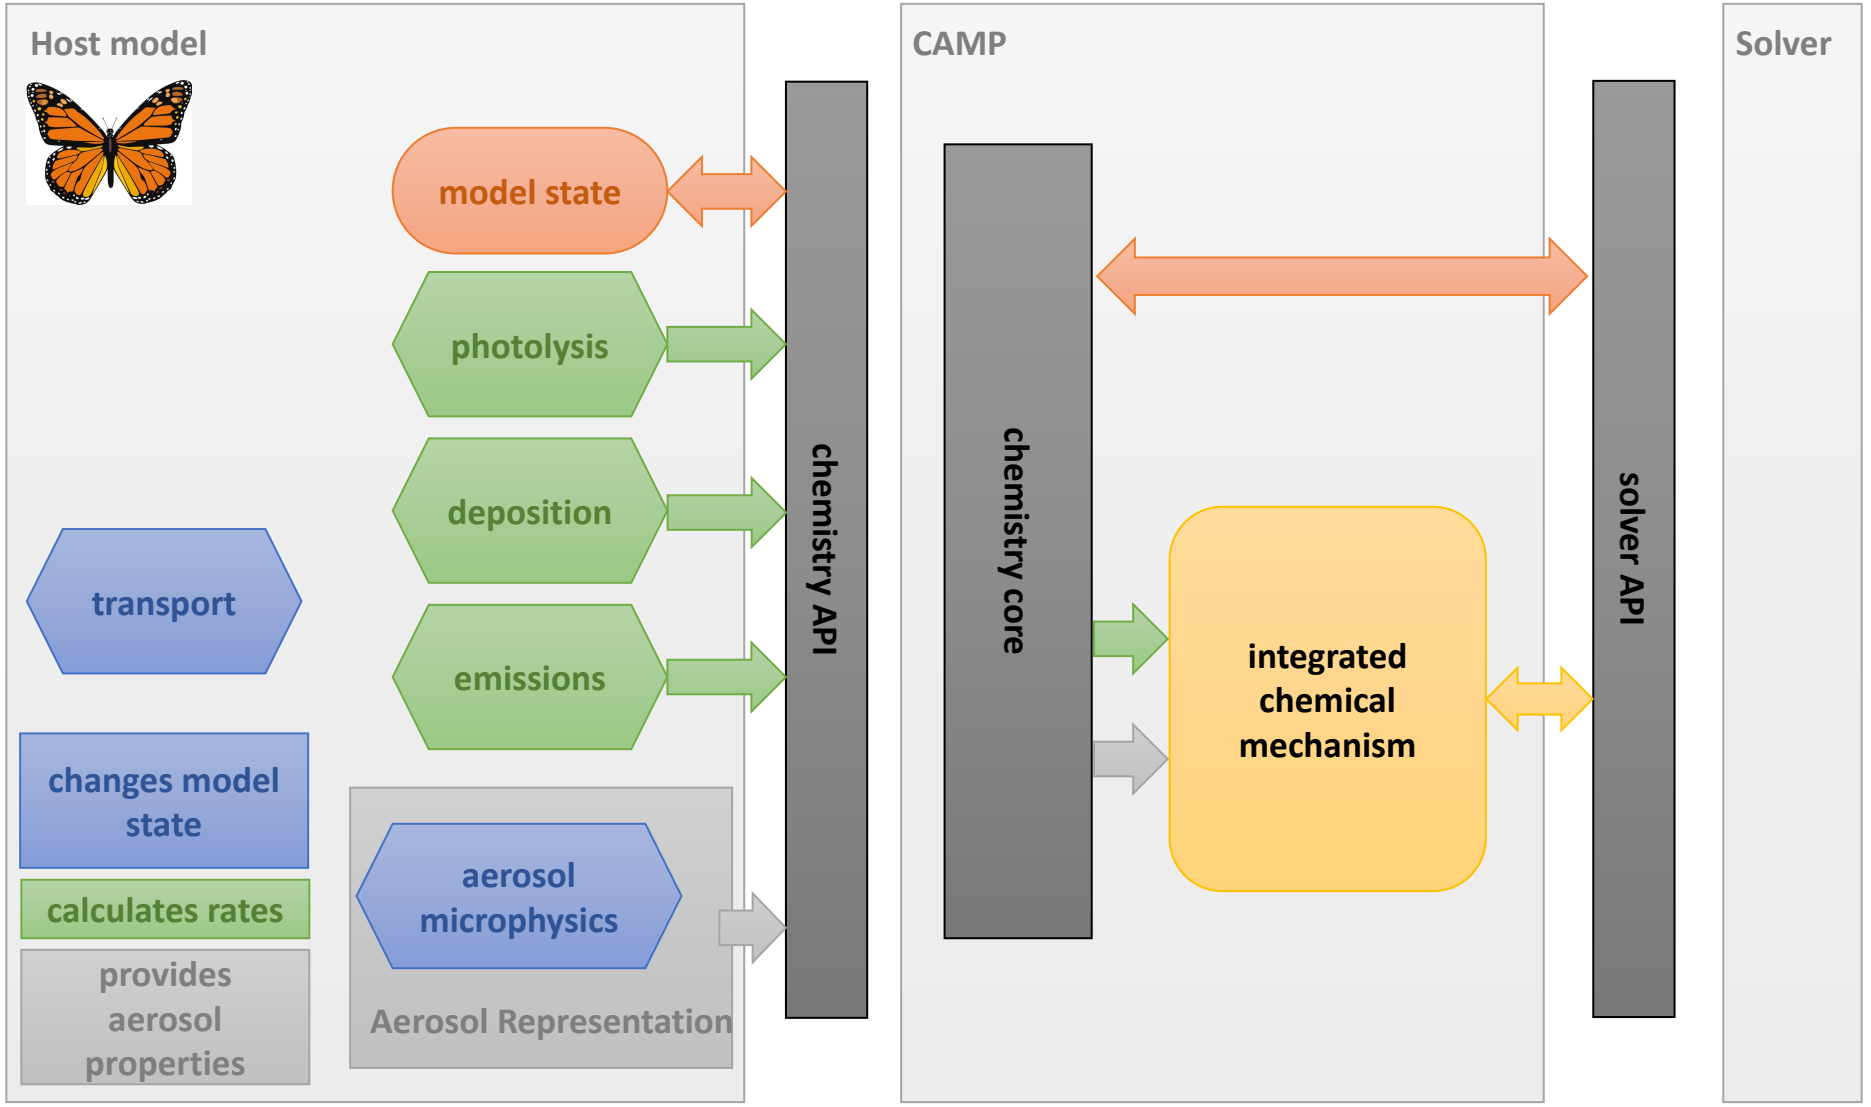

CAMP: Chemistry Across Multiple Phases

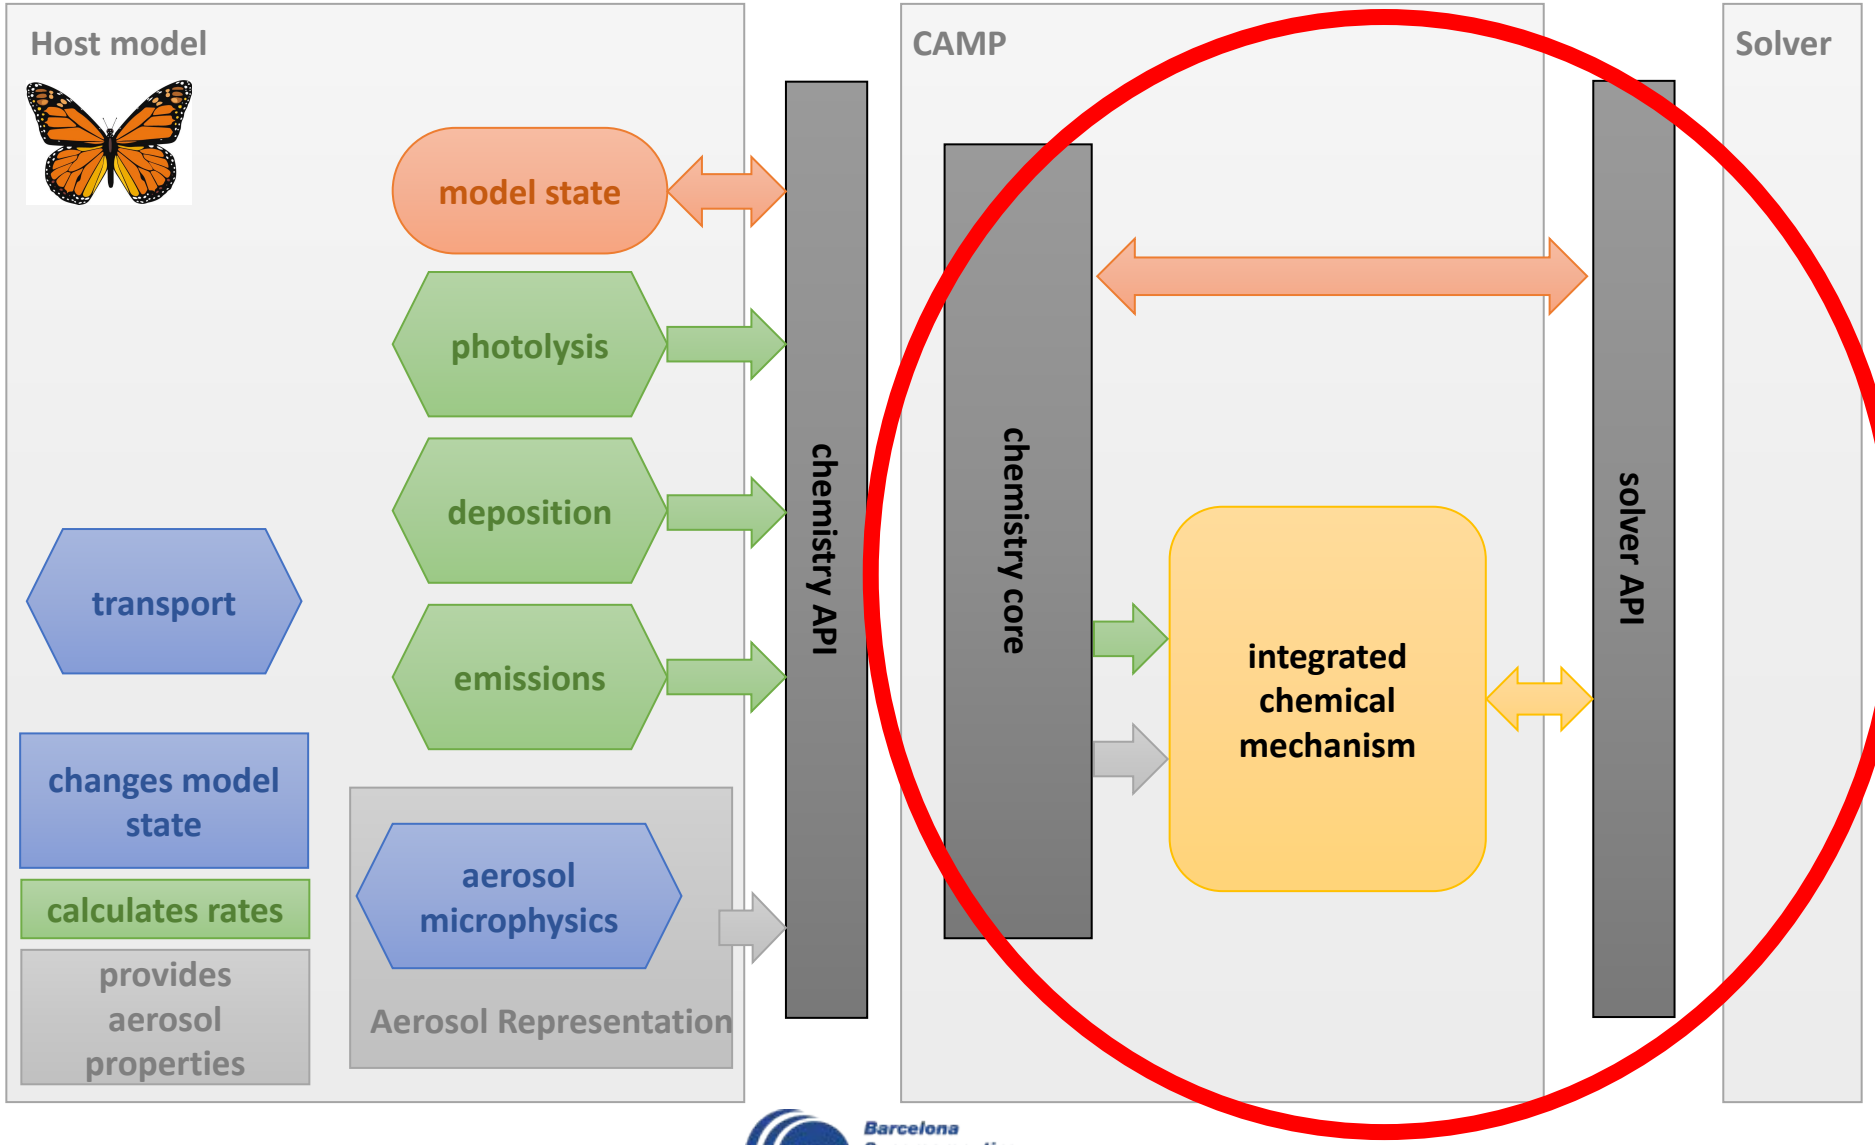

ODE Solver

- **Purpose**: Iteratively solves **y**'=**f**(t,**y**) using a Backward Differentiation Formula (CVODE) and the SuiteSparse KLU Linear Solver

- **Needs**: **f**(**y**) and **J** = ∂**f** /∂**y** (*Derivative* & *Jacobian*)

ODE Solver

- **Purpose**: Iteratively solves **y**'=**f**(t,**y**) using a Backward Differentiation Formula (CVODE) and the SuiteSparse KLU Linear Solver ~70%

- **Needs**: **f**(**y**) and **J** = ∂**f** /∂**y** (*Derivative* & *Jacobian*) ~30%

  ~20% ~10%

ODE Solver

- **Purpose**: Iteratively solves **y**'=**f**(t,**y**) using a Backward Differentiation Formula (CVODE) and the SuiteSparse KLU Linear Solver ~70%

- **Needs**: **f**(**y**) and **J** = ∂**f** /∂**y** (*Derivative* & *Jacobian*) ~30%

  ~20% ~10%

CAMP: Optimization strategy

- Profiling: Identify the most-time consuming functions
- GPU-based derivative function: Porting to GPUs and performance analysis
- Multi-cells & GPU: Improve performance using a multiple grid-cell solving strategy

CAMP: Optimization strategy

- ~~Profiling: Identify the most-time consuming functions~~ -> **Derivative (~20%)**

- GPU-based derivative function: Porting to GPUs and performance analysis

- Multi-cells & GPU: Improve performance using a multiple grid-cell solving strategy

GPU-based Derivative Function

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

CAMP workflow in MONARCH

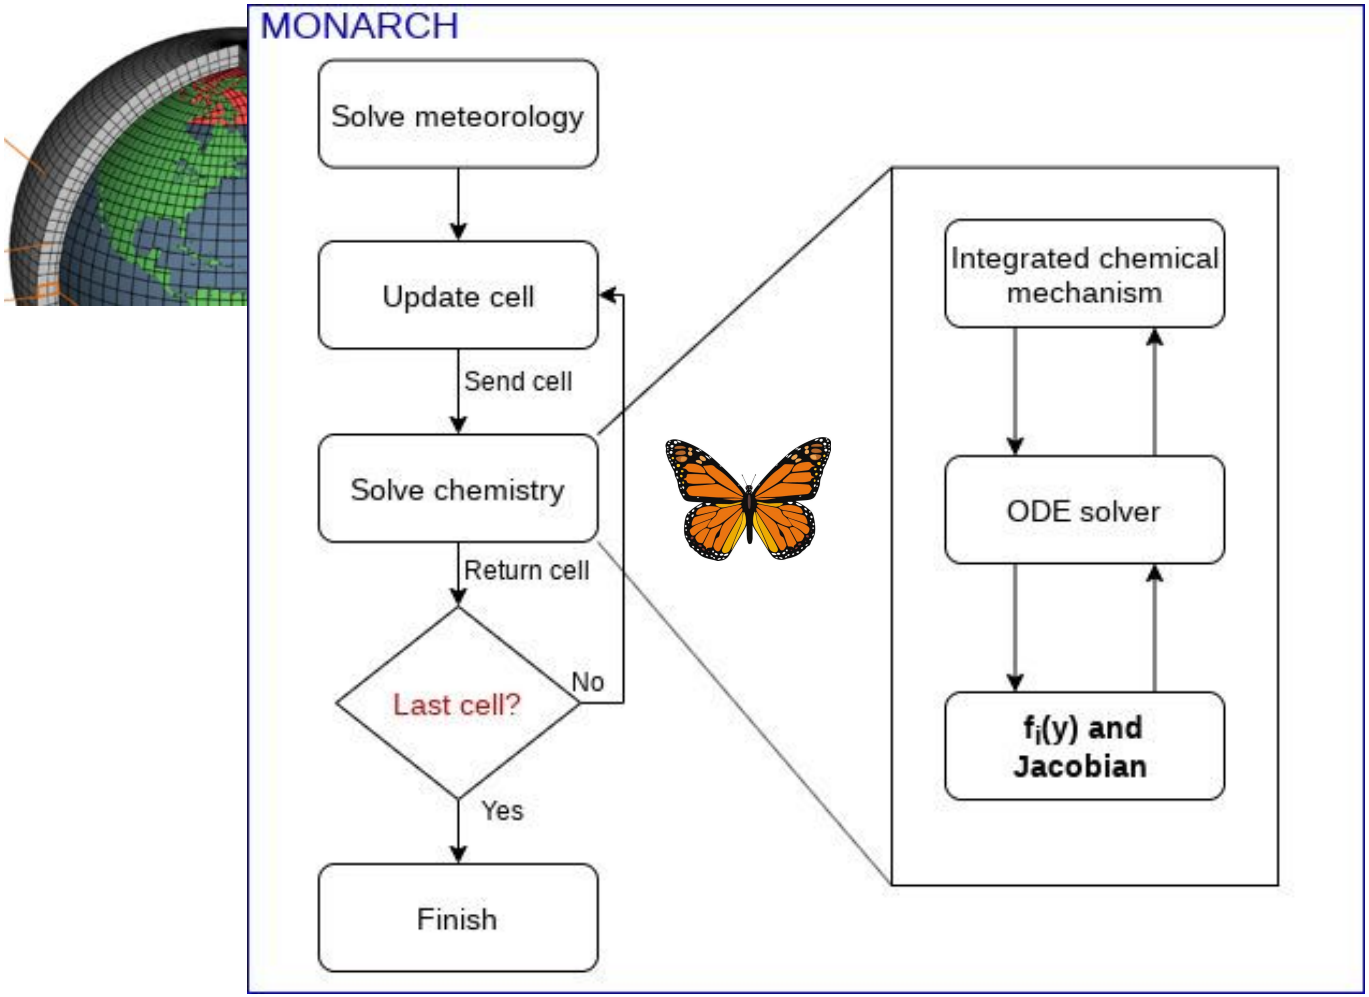

Derivative: f(y)

- A reaction:

$$c_1 y_1 + c_2 y_2 + \cdots + c_m y_m \leftrightarrow c_{m+1} y_{m+1} + c_{m+2} y_{m+2} + \cdots + c_n y_n,$$

- Derivative:

$$\left(\frac{dy_i}{dt}\right)_j = \begin{cases} -c_i r_j(\mathbf{y}, T, P, \ldots) & \text{for } i \le m \\ c_i r_j(\mathbf{y}, T, P, \ldots) & \text{for } m < i \le n \end{cases}$$

$$f_i \equiv \frac{dy_i}{dt} = \sum_j \left(\frac{dy_i}{dt}\right)_j$$

c = stoichiometric coefficient
t = time
r = rate
j = reaction
i = species
y_i = concentration of species i
m = number of reactants
n = number of products
p = number of reactions
w = number of species

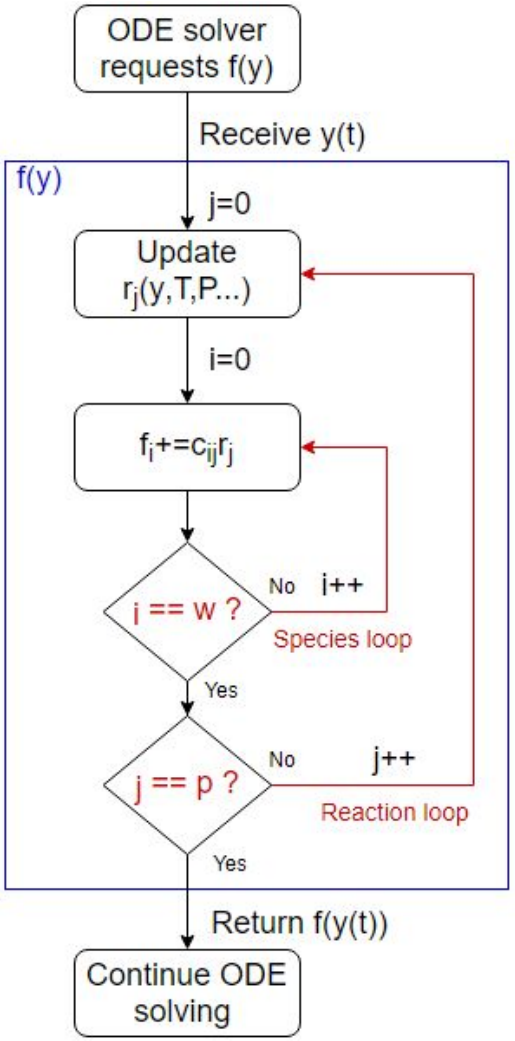

Derivative GPU

- Parallelize reactions loop
- Add data transfer
- Atomic operations

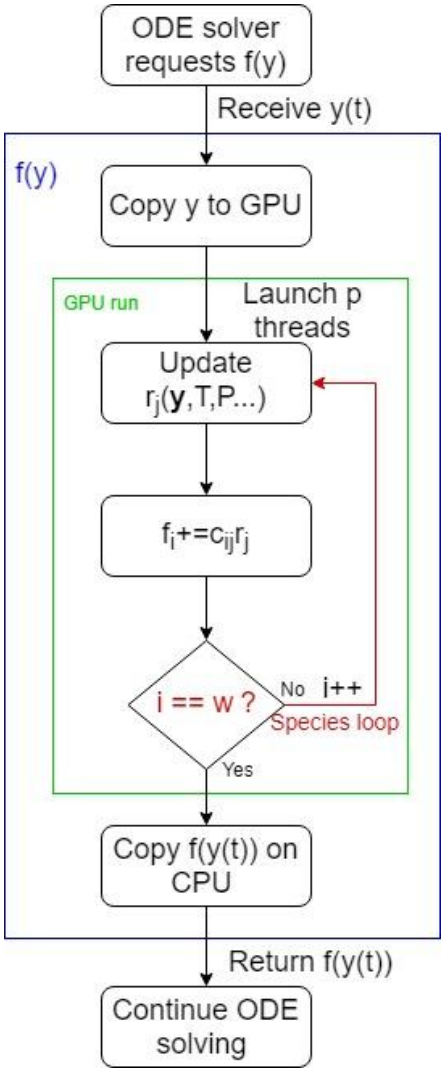

Test environment

- **Plaftorm:** CTE-POWER cluster, each node with:
  - 2 x IBM Power9 8335-GTH @ 2.4GHz
  - 4 x GPU NVIDIA V100 (Volta)
  - GCC version 6.4.0 and NVCC version 9.1

- **Configuration**: CB05

Mechanism	Reactions	Reactants	Cells	GPUs	MPI processes
CB05 (CPU)	186	72	1	0	1
CB05 (GPU)	186	72	1	1	1

Derivative GPU

- Test impact of GPU reaction calculations
- Scale by repeating CB05 mechanism
- We can still improve memory access

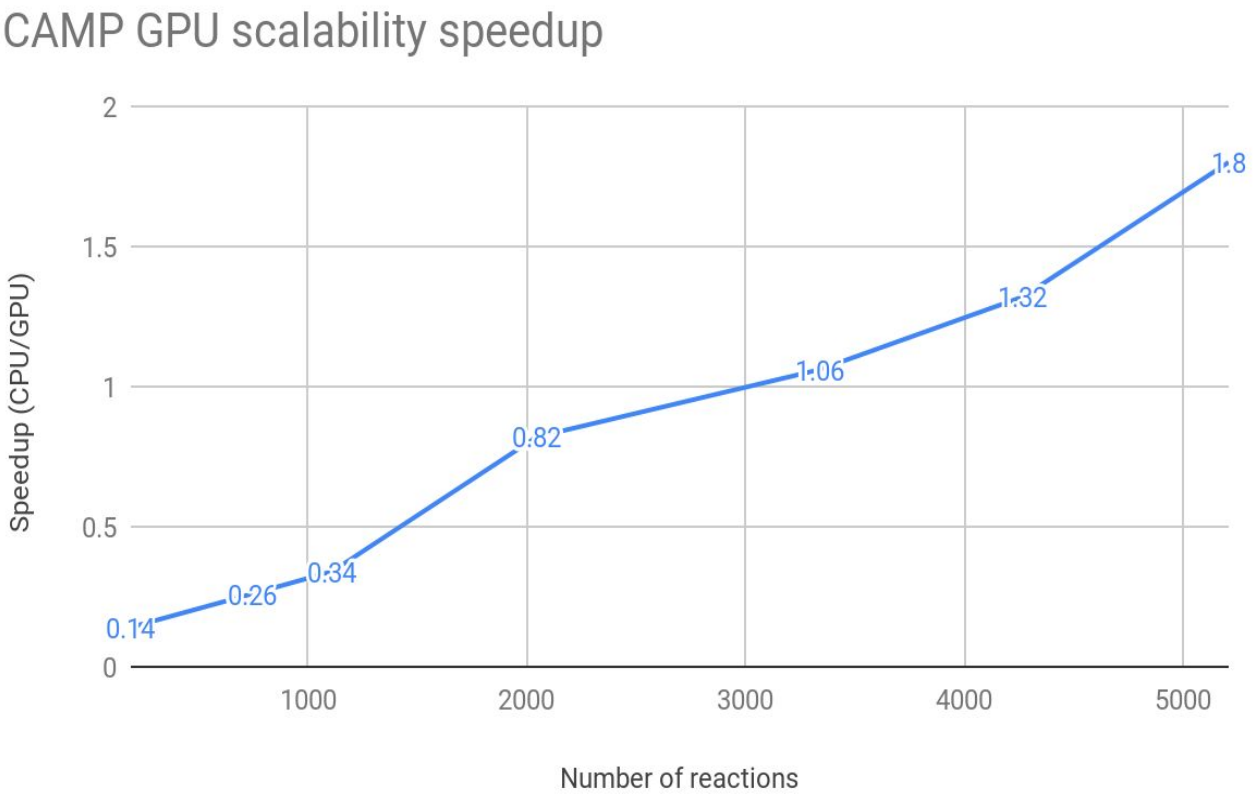

Reaction data structure

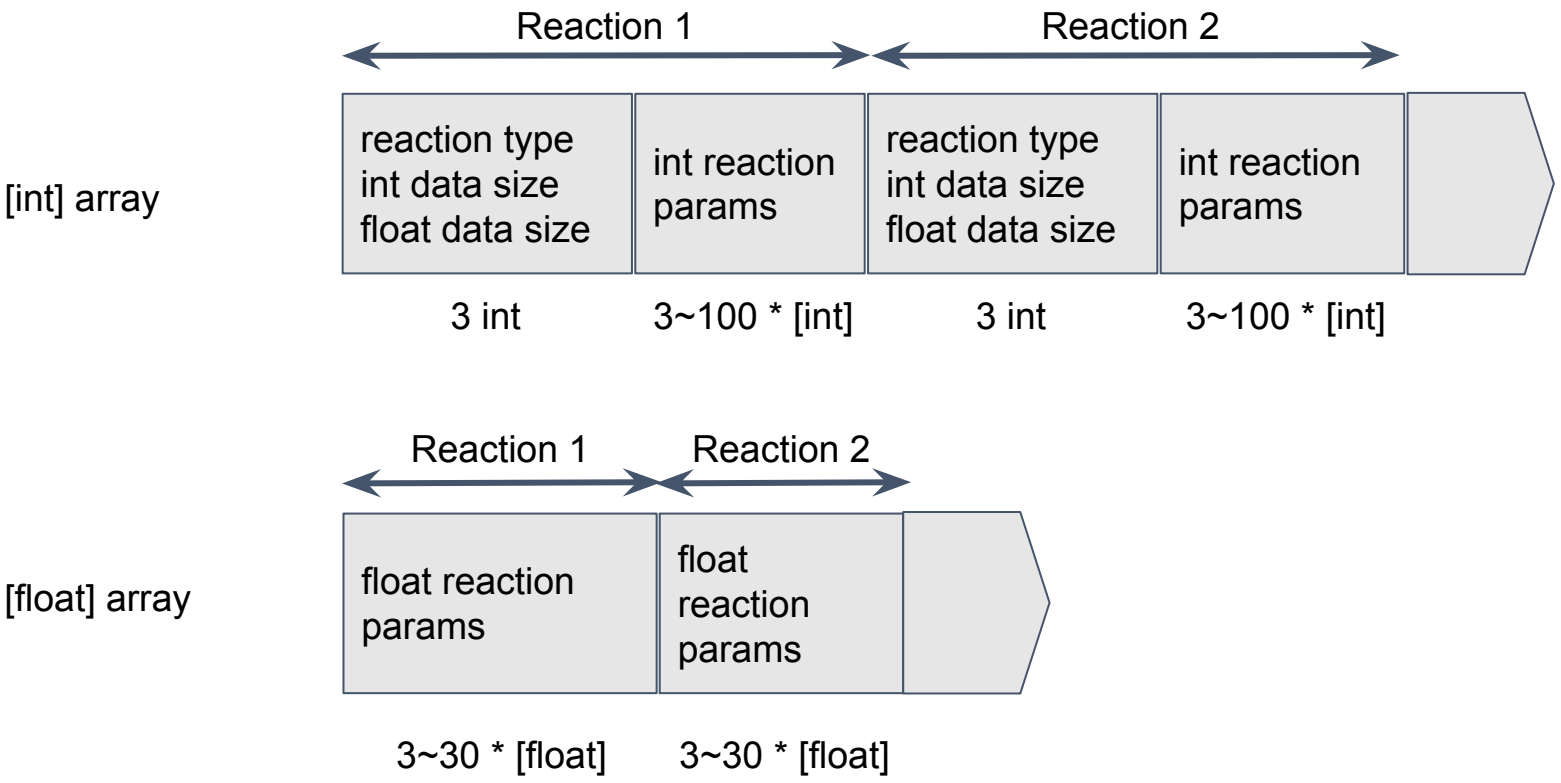

Derivative GPU: Inverted Data Structure

j = Reaction
p = Parameter
Value # = GPU access order and arrangement in memory

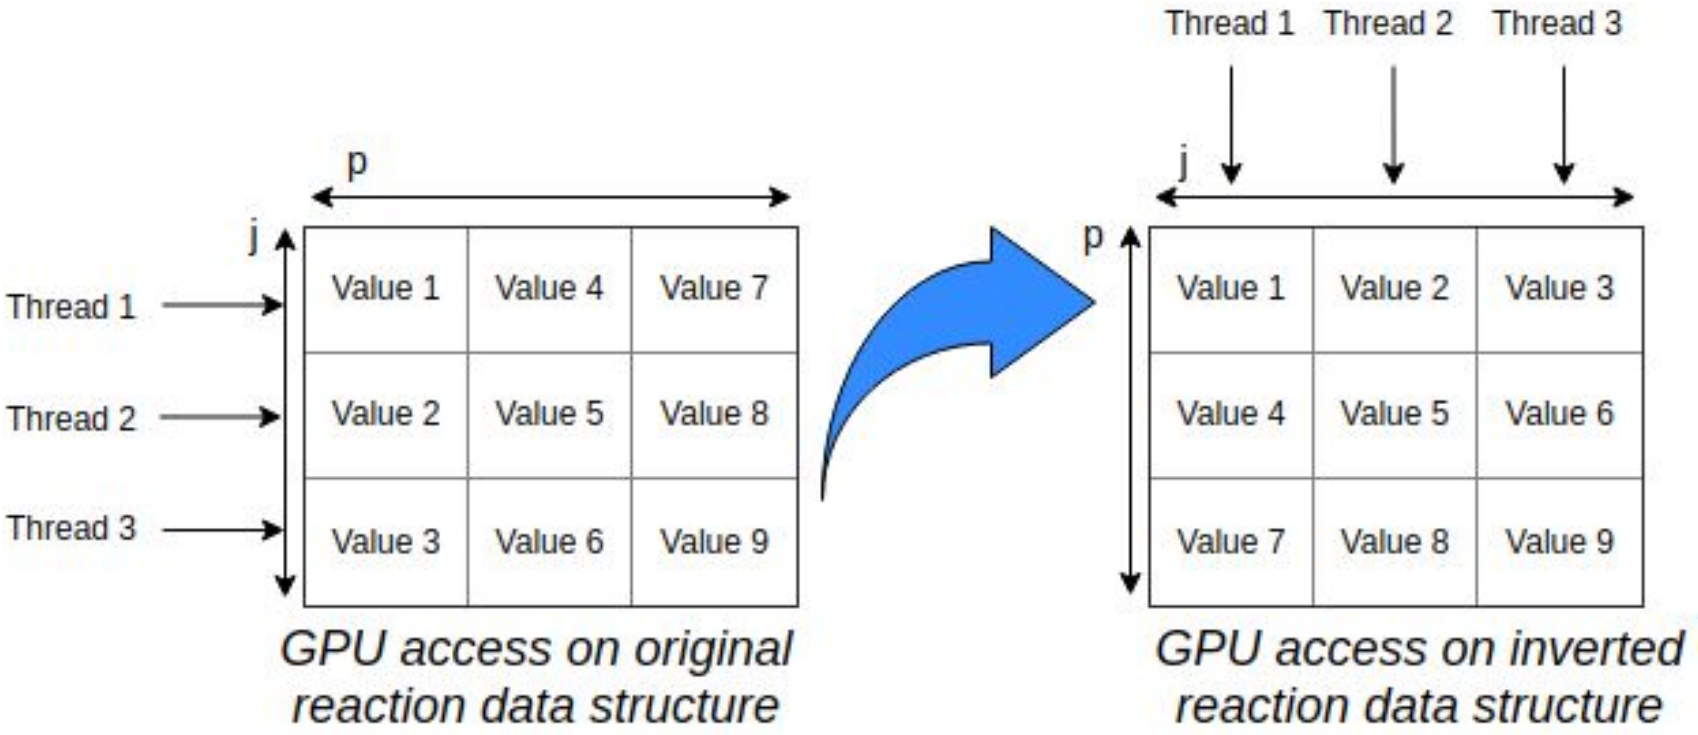

Derivative GPU: Inverted Data Structure

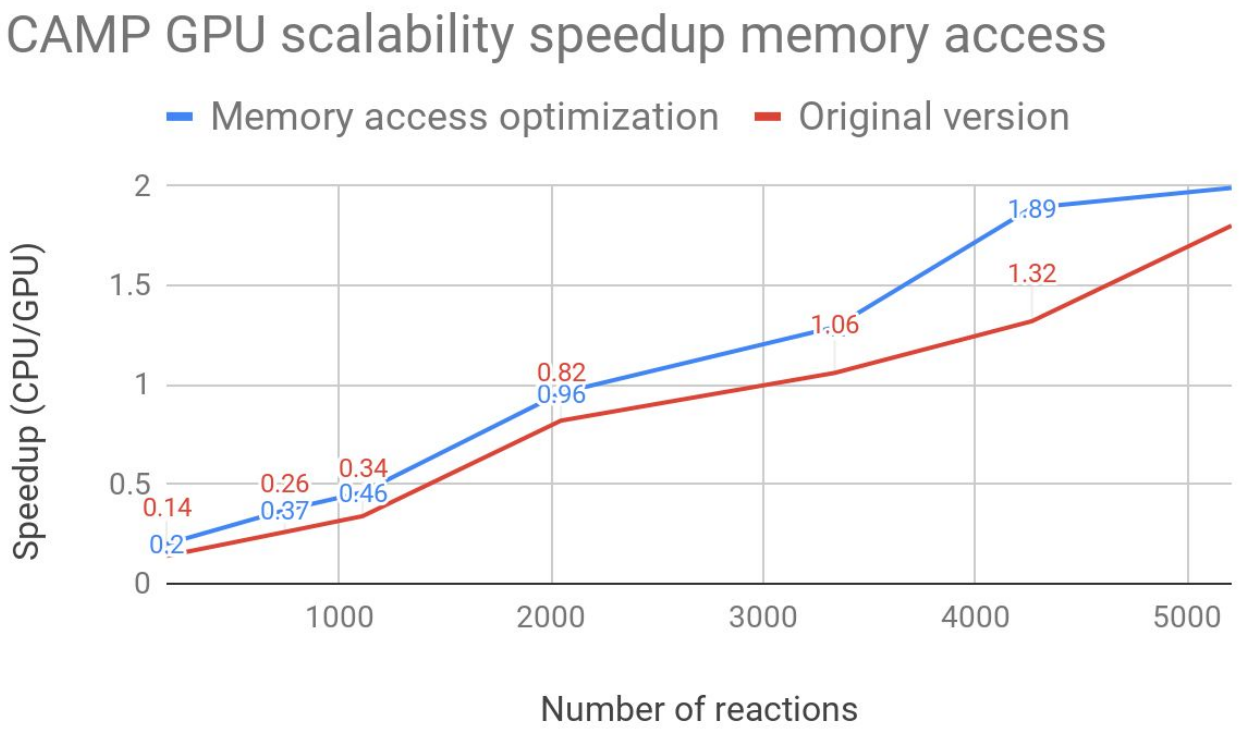

CAMP: Optimization strategy

- ~~Profiling: Identify the most-time consuming functions~~ -> **Derivative (~20%)**

- ~~GPU-based derivative function: Porting to GPUs and performance analysis~~ -> **GPU initialization is costly for small amounts of data**

- Multi-cells & GPU: Improve performance with a multiple grid-cell solving strategy

Multi-Cells & GPU

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

CAMP: Multi-cells

- **System Features:**
  - Cells are **not interdependent** w.r.t. chemistry
  - GPUs need sufficient work to offset initialization
  - **MPI domain decomposition** results in multiple cells per node

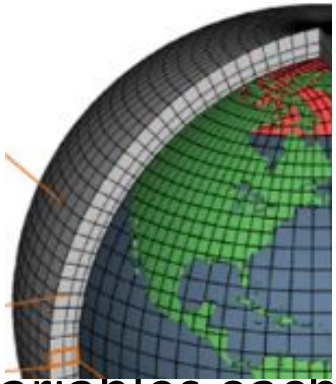

- **Goals:**
  - Avoid resetting variables each iteration
  - Reduce cache misses
  - Reduce ODE iterations

CAMP: Multi-cells

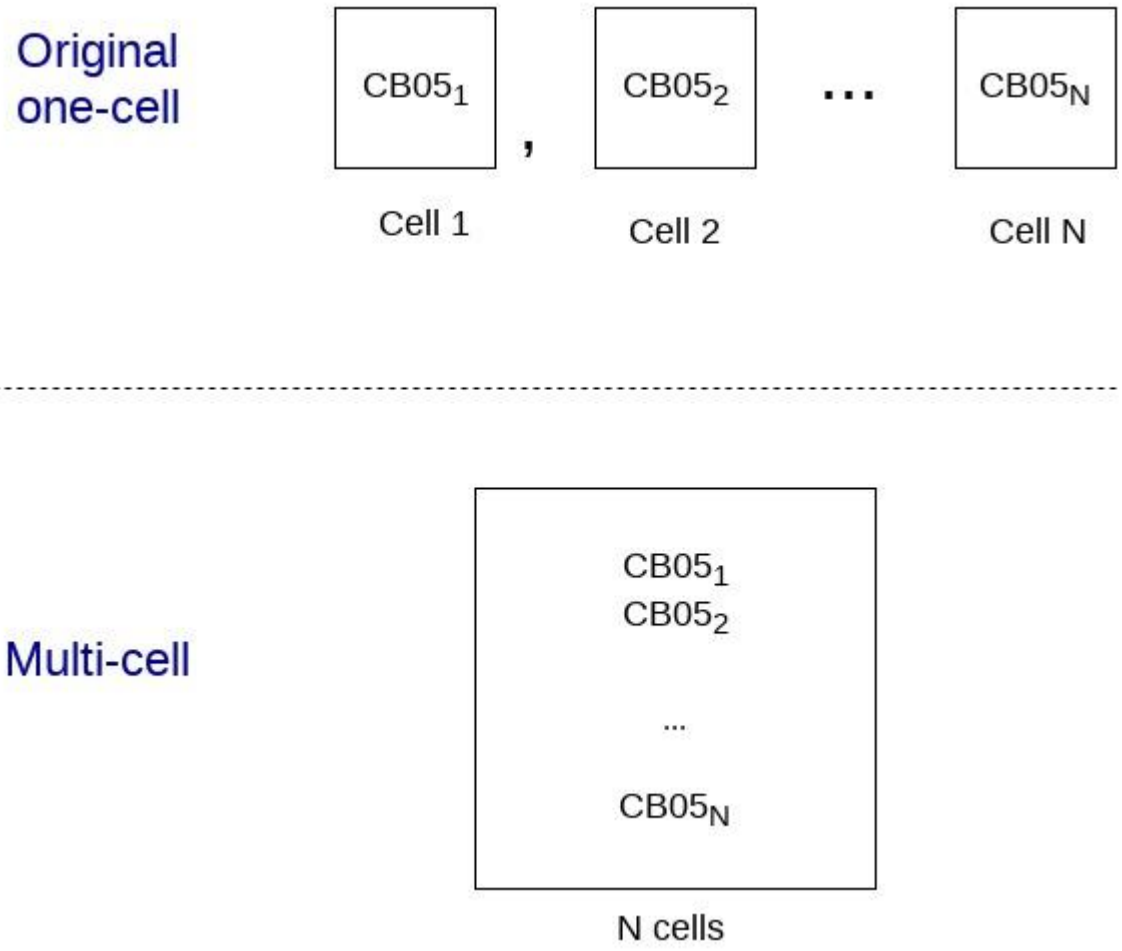

CAMP: Multi-cells

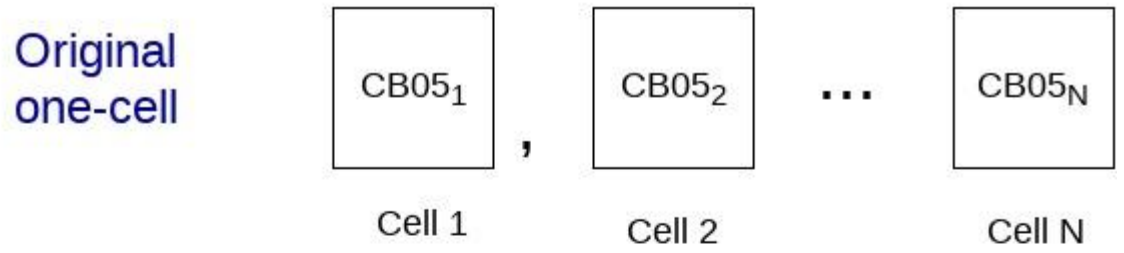

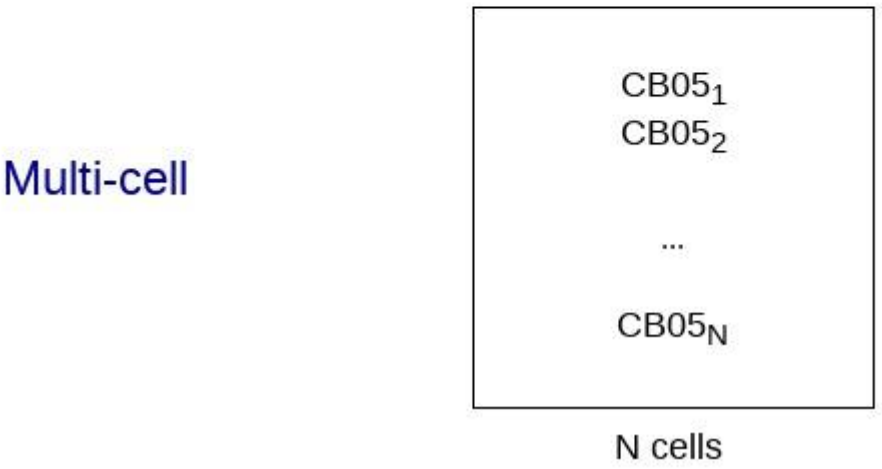

- Species "replication": $O3_1$, $O3_2$.. O_N
- Common ODE solver parameters

CAMP: Multi-cells

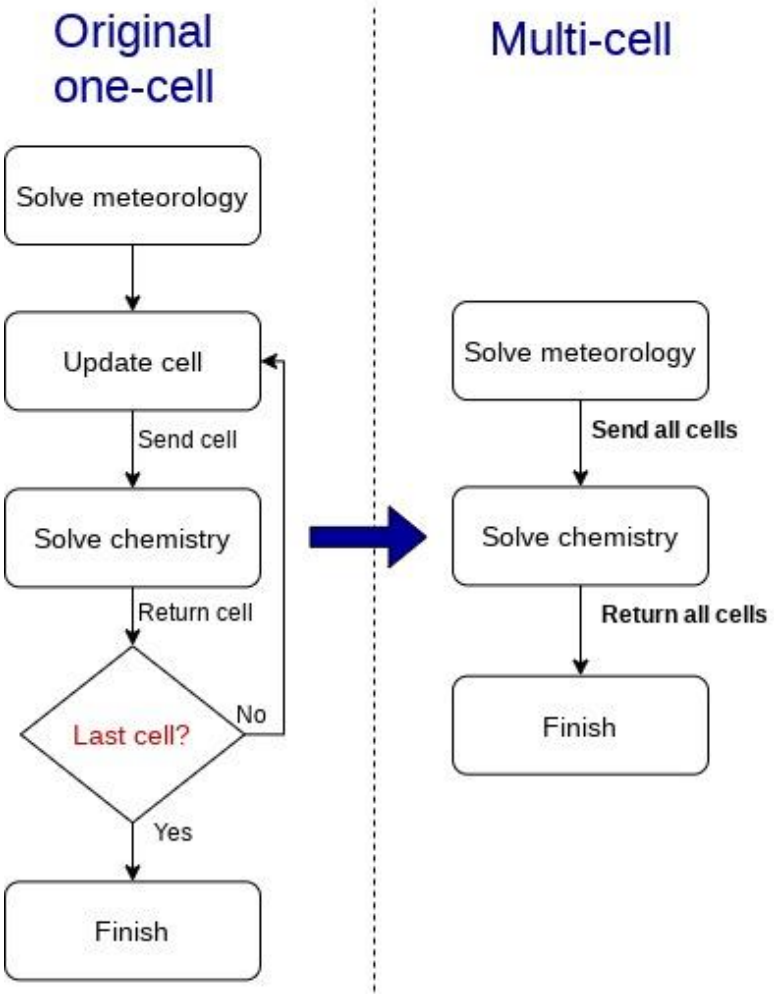

CAMP: Multi-cells

$$f_i \equiv \frac{dy_{ik}}{dt} = \sum_j \left(\frac{dy_{ik}}{dt} \right)_j$$

f_i = derivative
t = time
j = reaction
i = species
k = cell
y_{ik} = concentration of species i in cell k
p = number of reactions
q = number of cells

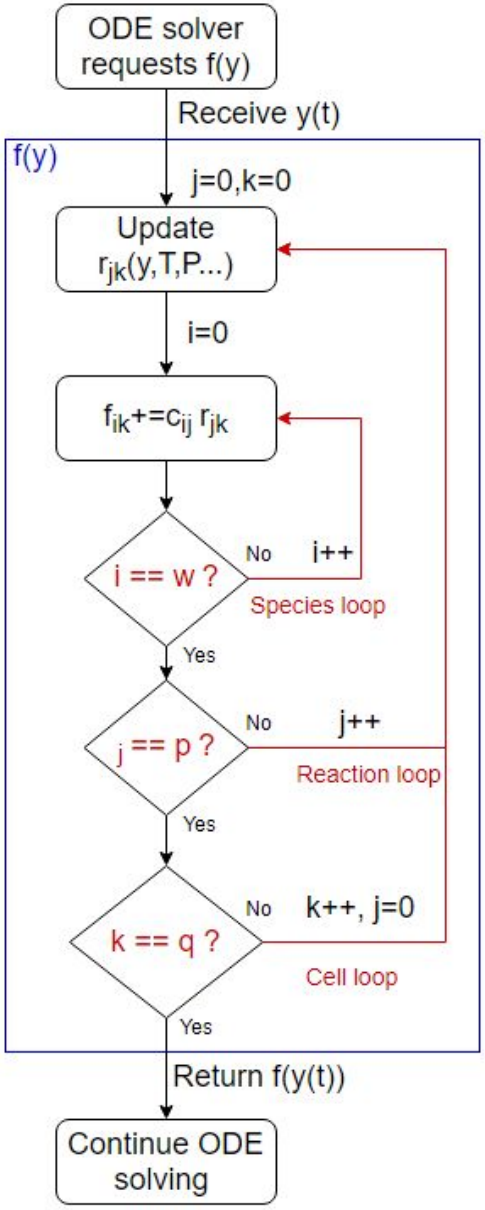

Test environment

- **Plaftorm:** CTE-POWER cluster, each node with:
  - 2 x IBM Power9 8335-GTH @ 2.4GHz
  - 4 x GPU NVIDIA V100 (Volta)
  - GCC version 6.4.0 and NVCC version 9.1

- **Configuration**: Basic

Mechanism	Reactions	Reactants	Cells*	GPUs	MPI processes
Basic (One-cell)	2	3	100 - 10,800	0	1
Basic (Multi-cell)	2	3	100 - 10,800	0	1

**10,800 cells is the common configuration in MONARCH*

CPU Multi-cells: Results

- Reduced ODE solver iterations

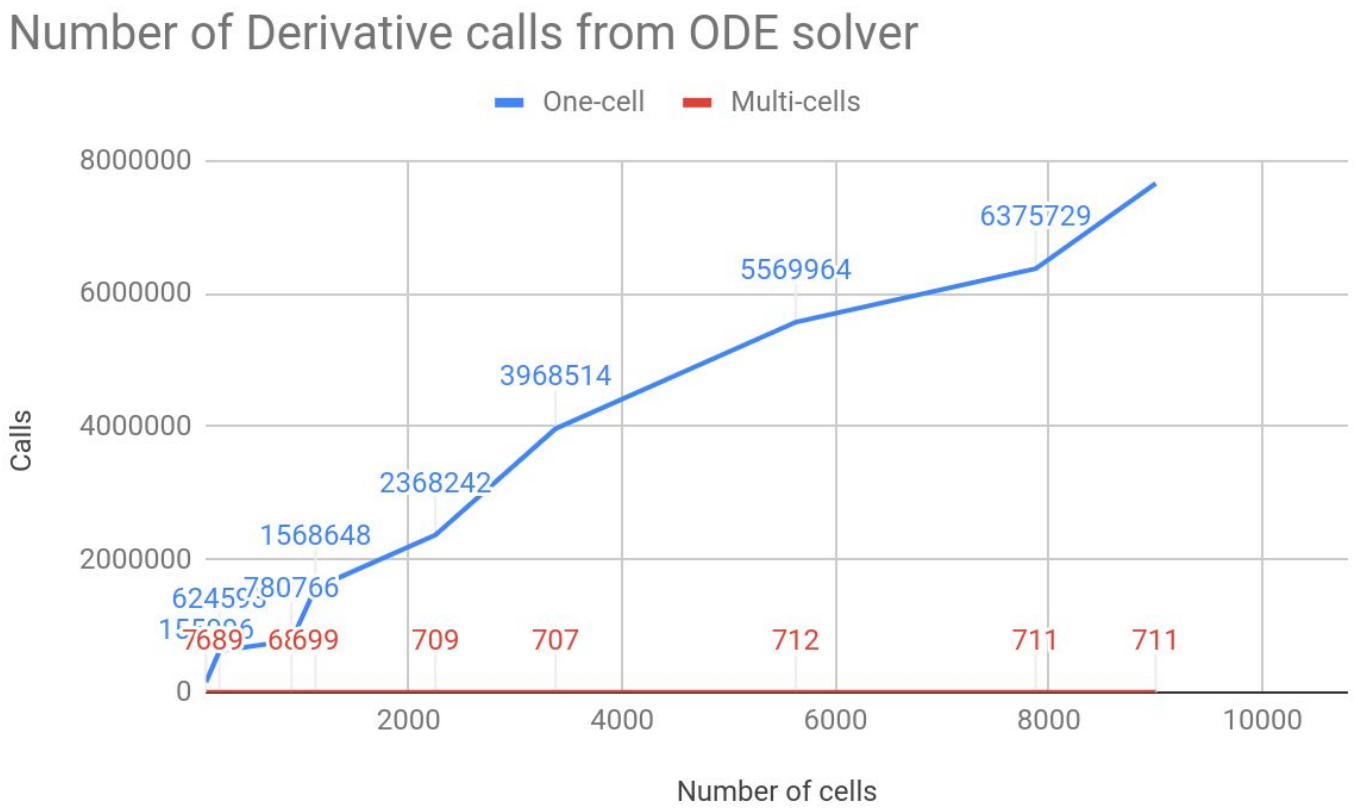

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

CPU Multi-cells: Results

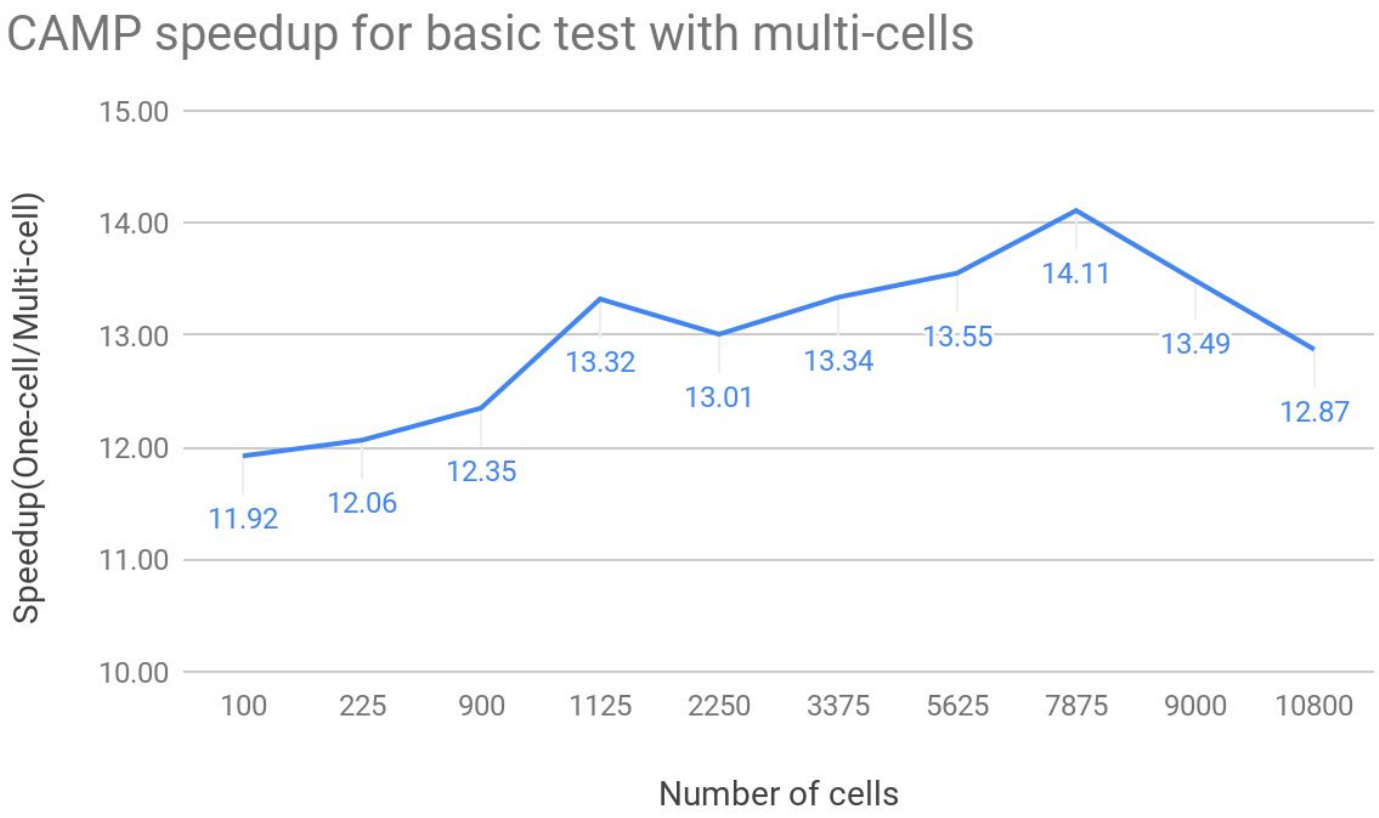

GPU Multi-cells

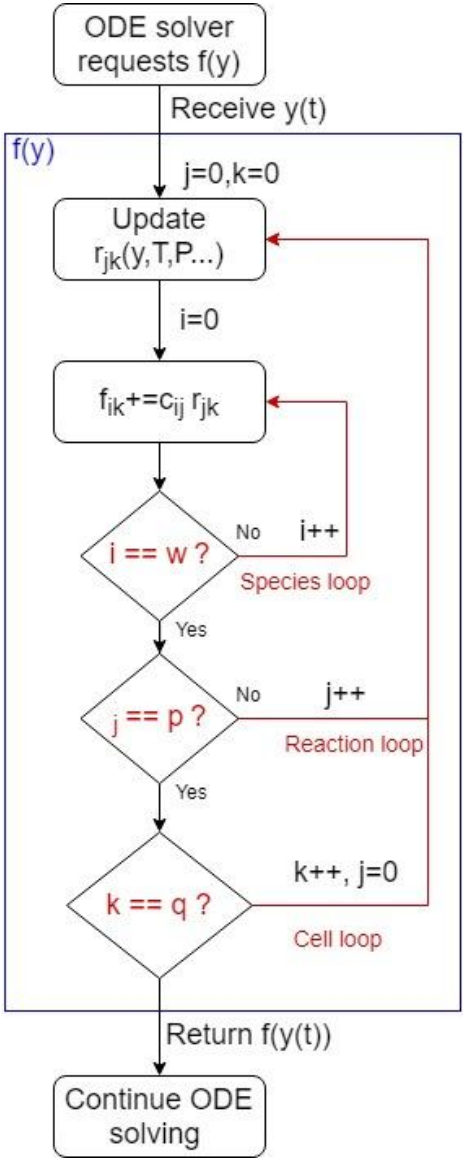

Reaction & Cell parallelization

f_i = *derivative*
t = time
r = rate
j = reaction
i = species
k = cell
y_{ik} = *concentration of species i in cell k*
p = number of reactions
q = number of cells

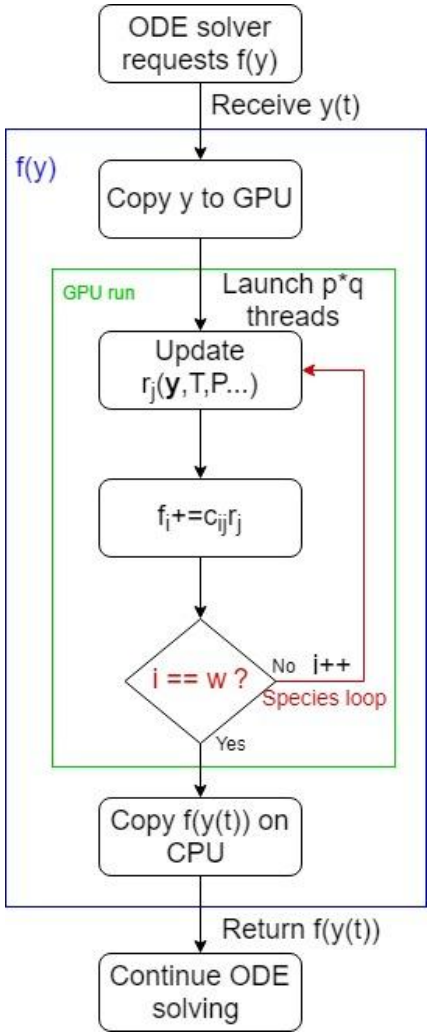

Test environment

- **Plaftorm:** CTE-POWER cluster, each node with:
  - 2 x IBM Power9 8335-GTH @ 2.4GHz
  - 4 x GPU NVIDIA V100 (Volta)
  - GCC version 6.4.0 and NVCC version 9.1

- **Configuration**: Basic GPU

Mechanism	Reactions	Reactants	Cells*	GPUs	MPI processes
Basic (GPU)	2	3	100-10,800	1	1

**10,800 cells is the common configuration in MONARCH*

GPU Multi-cells: Results

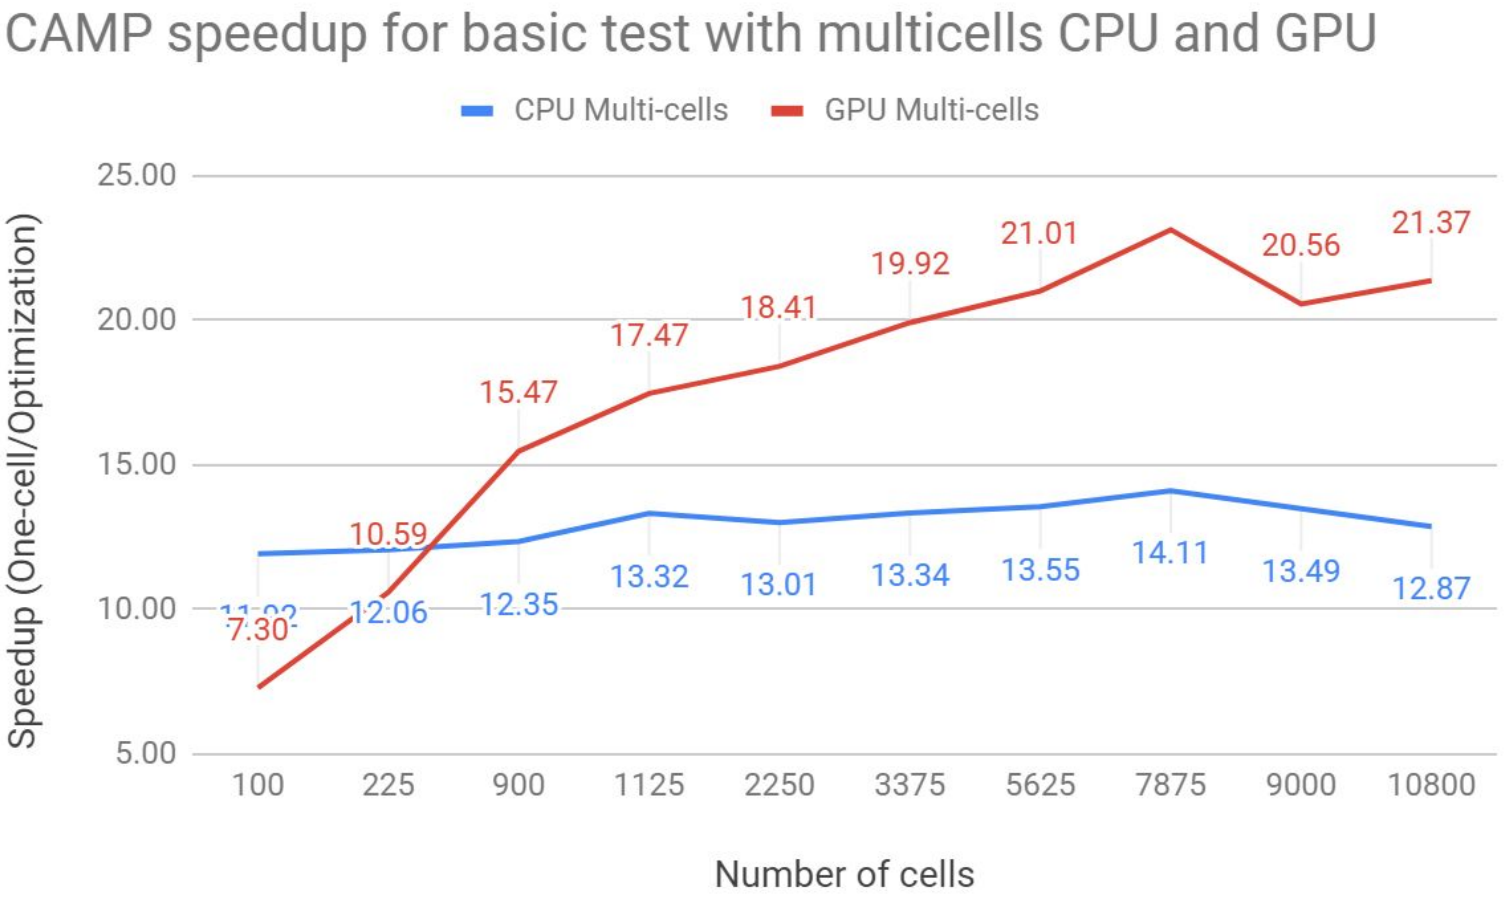

GPU Multi-cells: Data & Compute

Mechanism	Reactions	Reactants	Cells	GPUs	MPI processes
Basic (GPU)	2	3	10,800	1	1

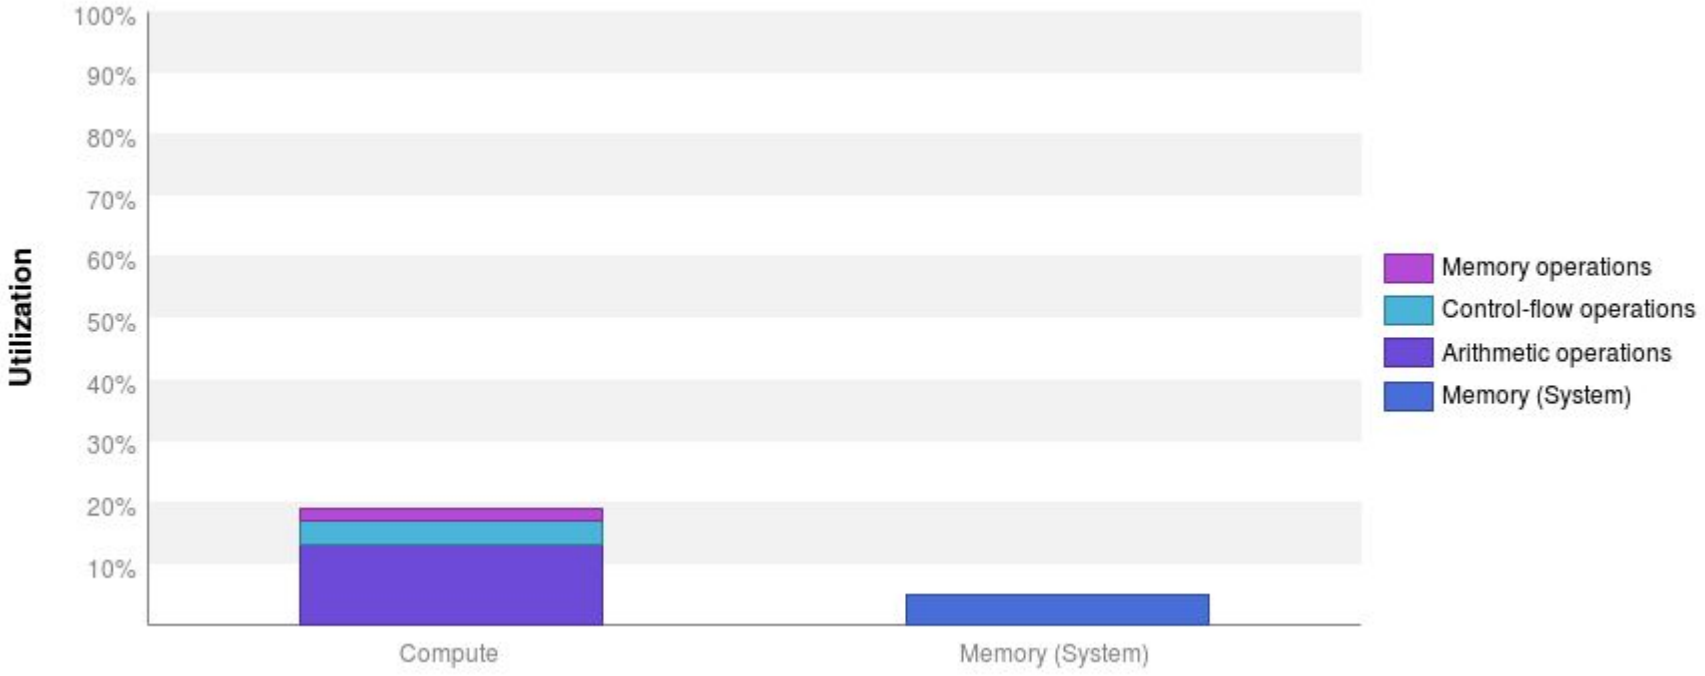

GPU Multi-cells: Data & Compute

Mechanism	Reactants	Cells	GPUs	Processes
Basic (GPU)	3	131072	1	1
Basic (MPI)	3	131072	0	40

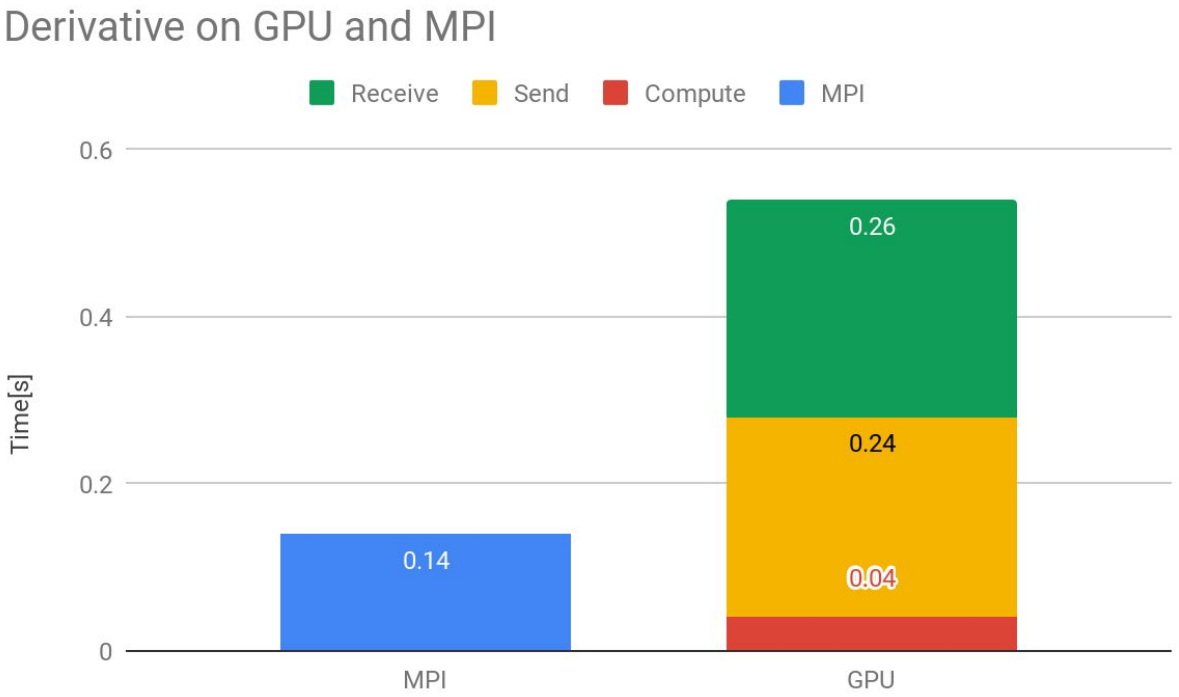

CAMP: Optimization strategy

- ~~Profiling: Identify the most-time consuming functions~~ -> **Derivative (~20%)**
- ~~GPU-based derivative function: Porting to GPUs and performance analysis~~ ->**GPU initialization is costly for small amounts of data**
- ~~Multi-cells & GPU: Improve performance using a multiple grid-cell solving strategy~~ ->**GPUs can handle even more load** ... For upcoming work!

Conclusions

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Conclusions

- GPU-based derivative function improves systems with >2000 reactions

  -> GPUs significantly speed up solving large mechanisms

- Optimizing GPU memory access by reconfiguring data structures improves execution up to 30%

  -> Ensuring optimal memory access always improves efficiency

Conclusions

- Multi-cell approach makes solving 12–14 times faster

  -> Simultaneously solving cells reduces solver iterations

- Porting solver functions to GPUs coupled with multi-cell treatment improves chemistry solving by 7–21×

  -> Maximizing parallelization improves GPU functions

- Data movement accounts for most multi-cell GPU computation time for large numbers of grid-cells

  -> Search for alternatives (async & more computation)

Future work

- Porting all solver functions to GPUs will reduce data movement and improve efficiency

- Load balancing GPU & CPU + asynchronous communication

- Evaluate GPU-based chemistry solving in MONARCH

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Thank you

christian.guzman@bsc.es

Appendix

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Initial CPU-based CAMP

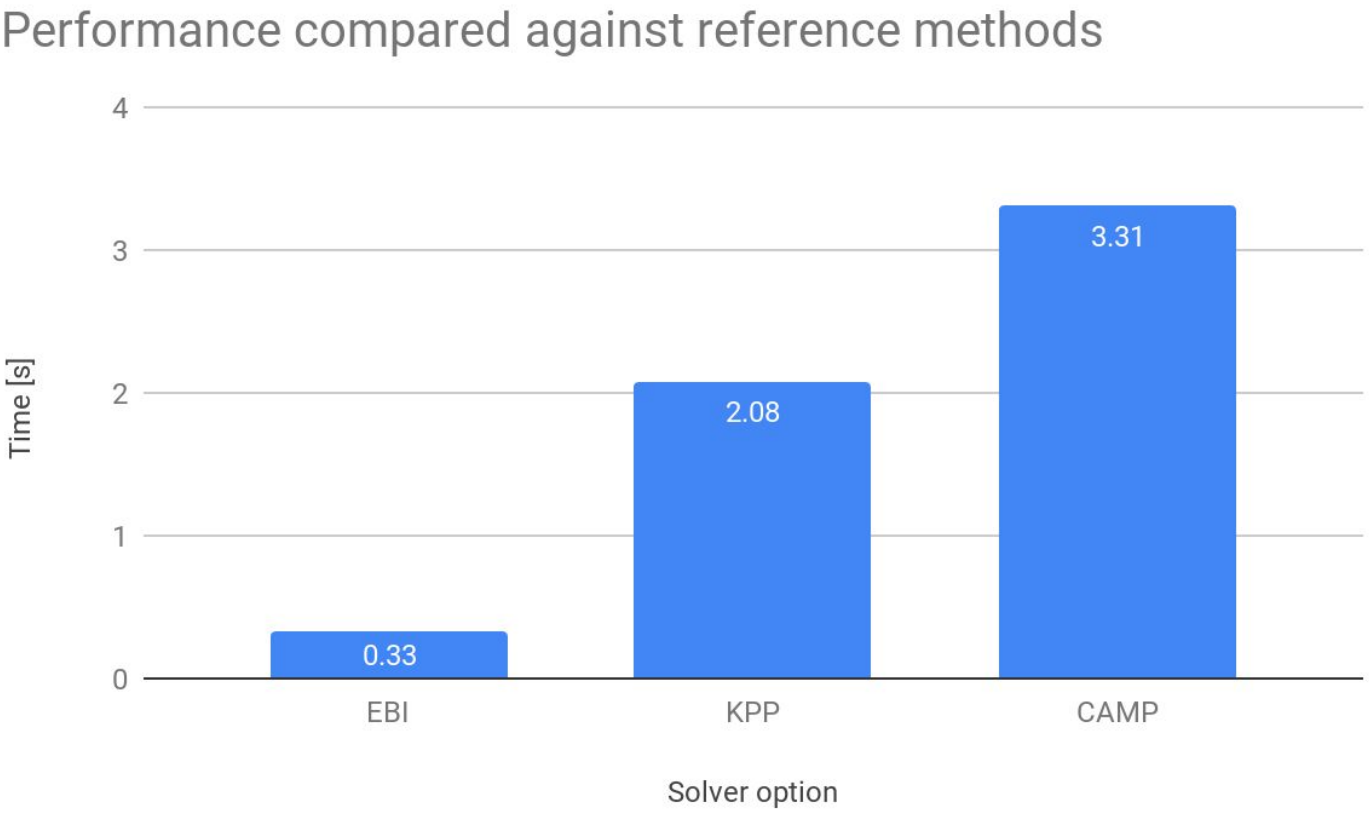

GPU Multi-cells: Memory & Compute

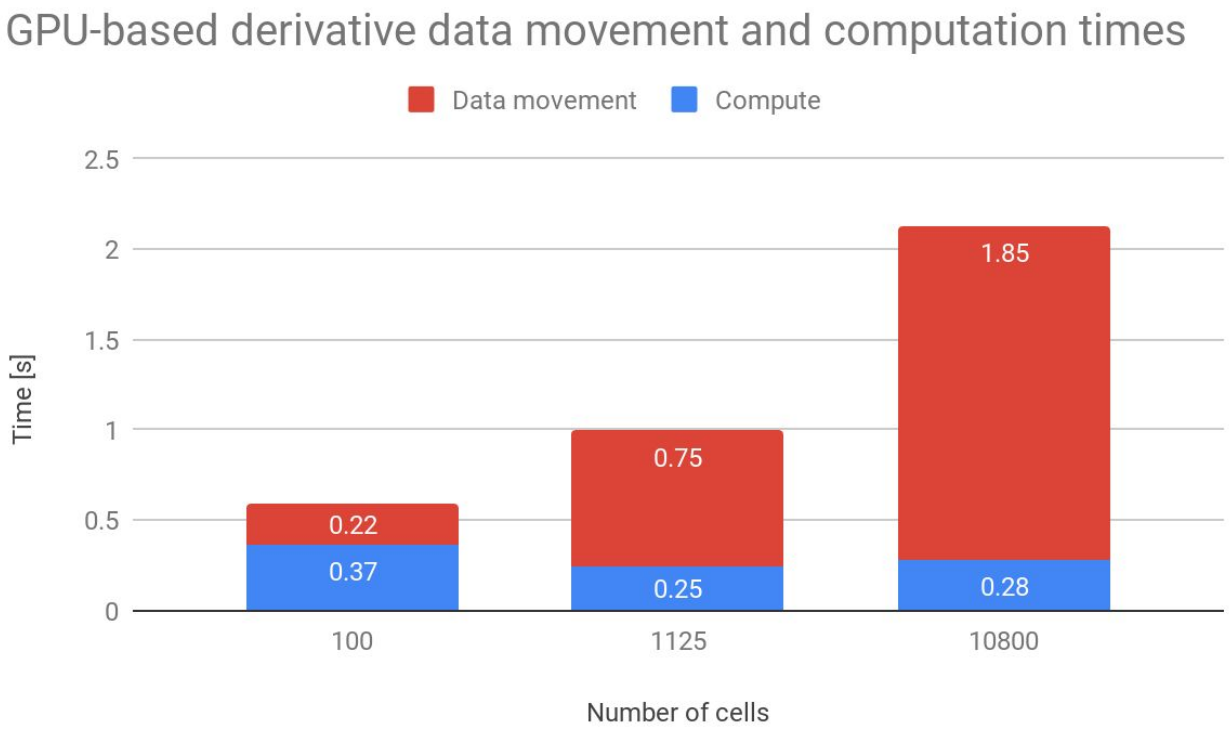

CAMP GPU vs KPP GPU

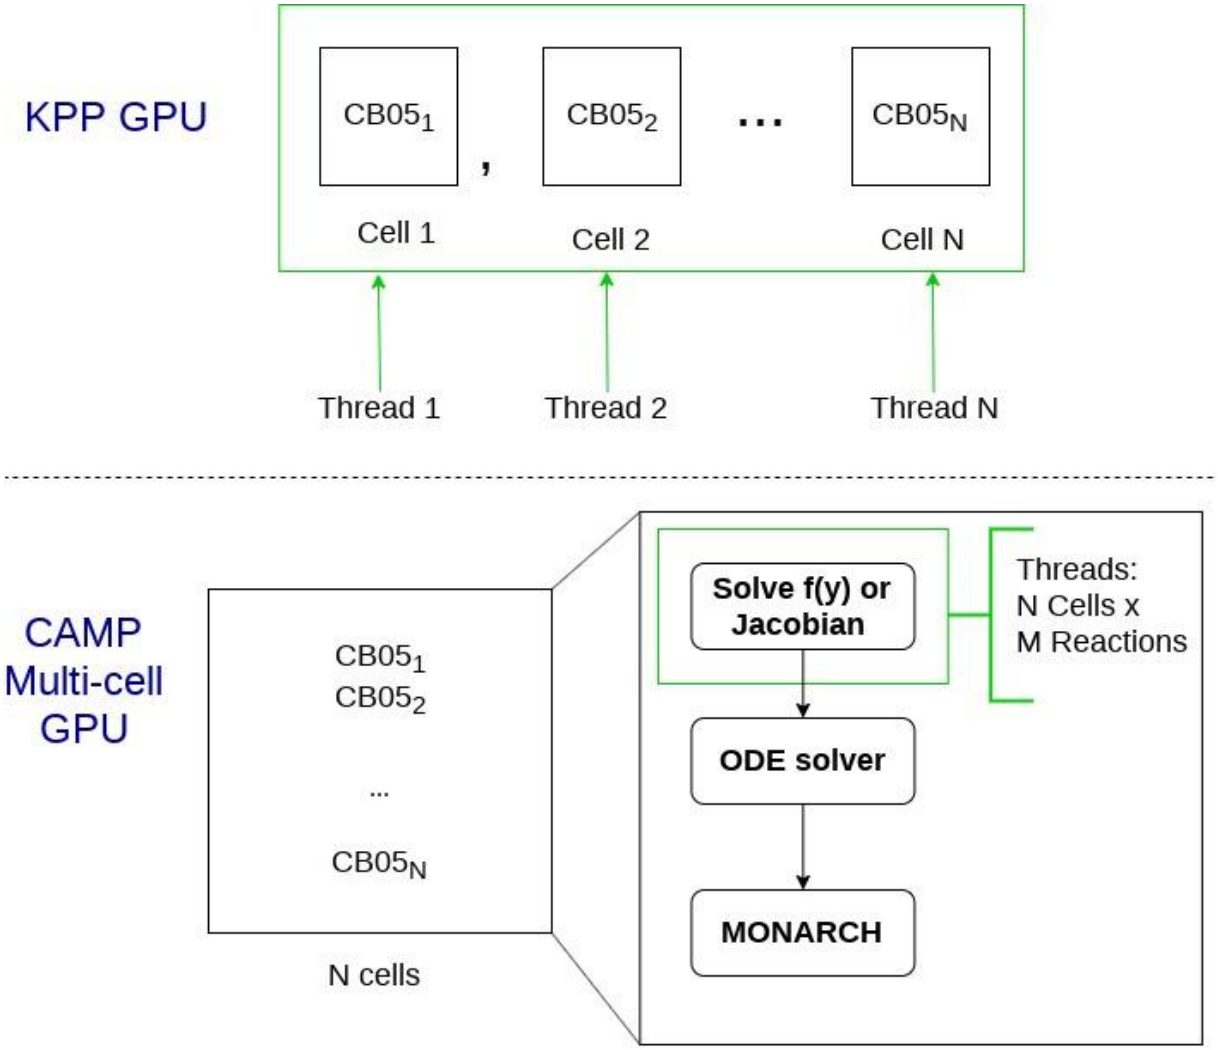

CAMP GPU (including future work) vs KPP GPU

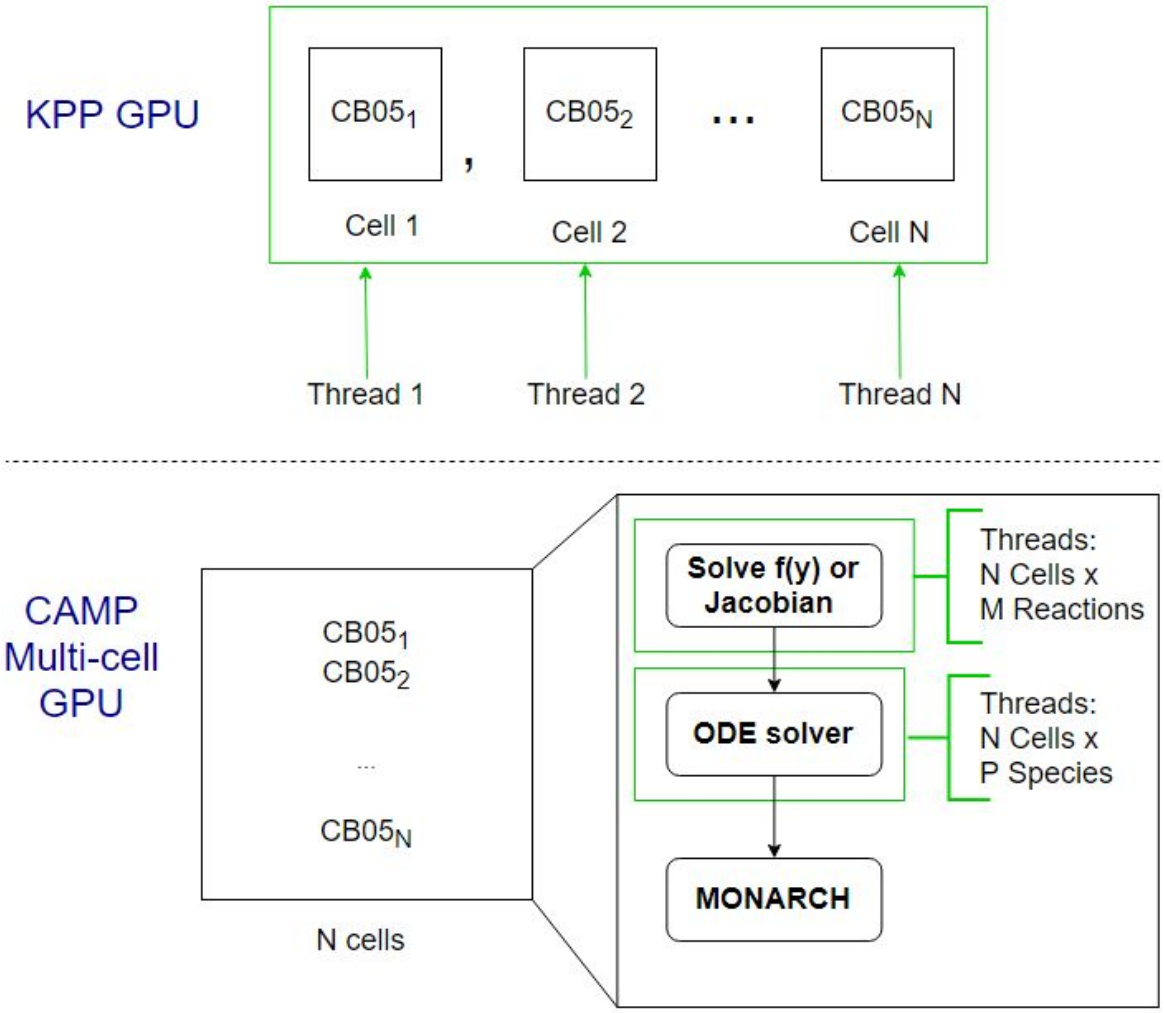

GPU Multi-cells: Block processing

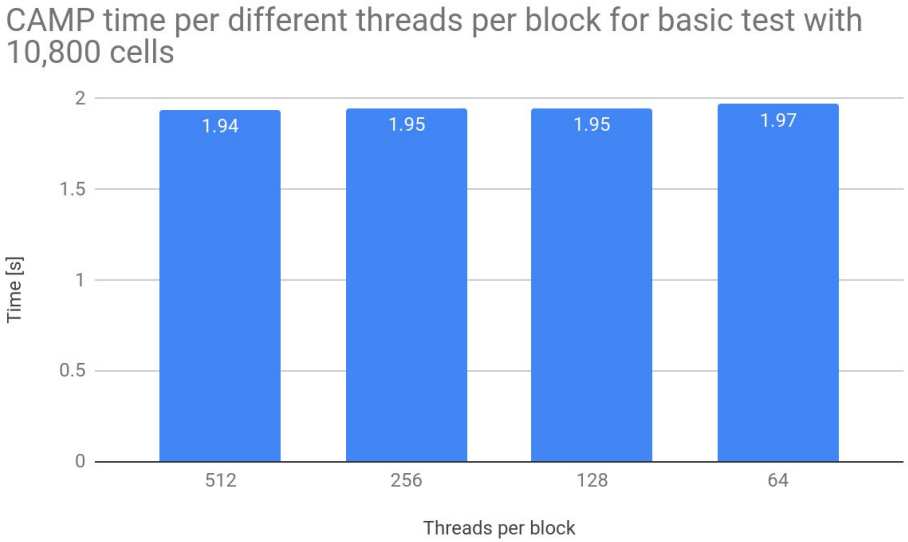

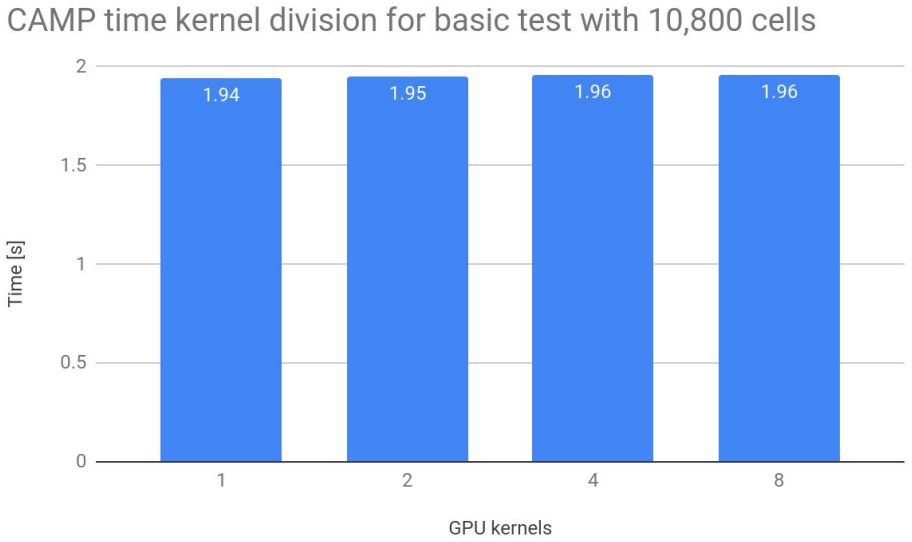

Platform

CTE-POWER:

2 login nodes and 52 compute nodes, each of them:

- 2 x IBM Power9 8335-GTH @ 2.4GHz (3.0GHz on turbo, 20 cores and 4 threads/core, total 160 threads per node)
- 512GB of main memory distributed in 16 dimms x 32GB @ 2666MHz
- 2 x SSD 1.9TB as local storage
- **4 x GPU NVIDIA V100 (Volta) with 16GB HBM2.**
- GPFS via one fiber link 10 GBit
- **Compilers: GCC version 6.4.0 and NVCC version 9.1**

Reaction data