CAP-387(2016) – Tópicos Especiais em Computação Aplicada: Construção de Aplicações Massivamente Paralelas

Aula 21: Supercomputing'16

Celso L. Mendes, Stephan Stephany

LAC / INPE

Emails: celso.mendes@inpe.br, stephan.stephany@inpe.br

Supercomputing (SC)

- Maior conferência mundial em supercomputação
  - Tipicamente, ~11.000 ou mais participantes
    - ~80% EUA, ~20% resto do mundo
  - Representantes de: indústria, universidades, labs nacionais
  - Combinação de: artigos técnicos, palestras, exposição, workshops, tutoriais, posters, BoF's, etc., etc.
    - Divulgação da lista de Novembro do Top500 (lista de Junho divulgada na Europa – conferência ISC)
  - Sempre realizada na terceira semana de Novembro, nos EUA
  - Organização: IEEE e ACM
  - Estudantes podem se candidatar para serviços voluntários!
  - URL: http://www.supercomputing.org

Supercomputing’16 (SC16)

- Edição de 2016:
  - 29ª edição da conferência
  - Local: Salt Lake City, Utah
  - Data: 13~18/Novembro/2016
  - Maior área de exposição de vendedores até hoje!
  - URL: http://sc16.supercomputing.org/
- Edição de 2017:
  - Local: Denver, Colorado
  - Data: 12~17/Novembro/2017

SC16 – Top500

- Top500 Novembro/2016:
  - #1: Sunway TaihuLight @93 PF (pico 125 PF), China – SW26010
  - #2: Tianhe-2 @34 PF (pico 55 PF), China – Intel KNC
  - #3: Titan (Cray XK) @18 PF (pico 27 PF), EUA – AMD+GPU
  - #4: Sequoia (IBM-BG) @17 PF (pico 20 PF), EUA - PowerBQC
  - **#5: Cori (Cray XC) @14 PF (pico 28 PF), EUA – Intel KNL**
  - **#6: Oakforest (Fujitsu) @13 PF (pico 25 PF), Japão – Intel KNL**
  - #7 (ex #5): K Computer (Fujitsu) @10PF (pico 11 PF), Japão – Sparc
  - **#8: PizDaint (Cray) @9,8PF (pico 16 PF), Suiça – GPU-P100**
  - #9 (ex #6): Mira (IBM) @9 PF (pico 10 PF), EUA – PowerBQC
  - #10 (ex #7): Trinity (Cray) @8 PF (pico 11 PF), EUA – Intel-Xeon

SC16 – Top500 (cont.)

- Top500 – Composição por Fabricantes

(Número de Sistemas)

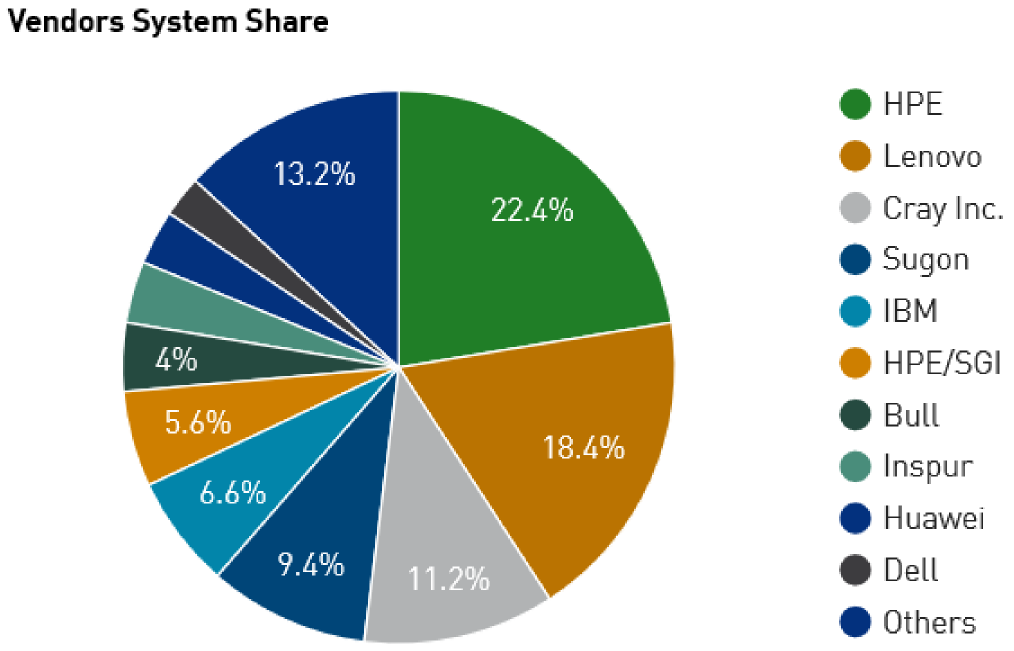

(Desempenho Agregado)

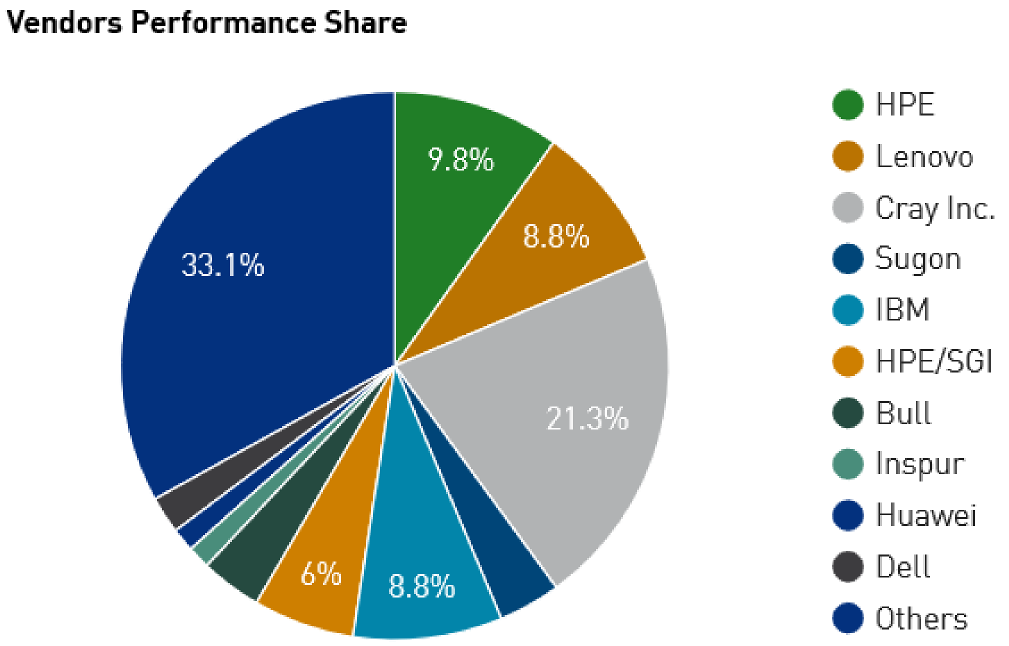

SC16 – Top500 (cont.)

- Top500 – Composição por Países

(Número de Sistemas)

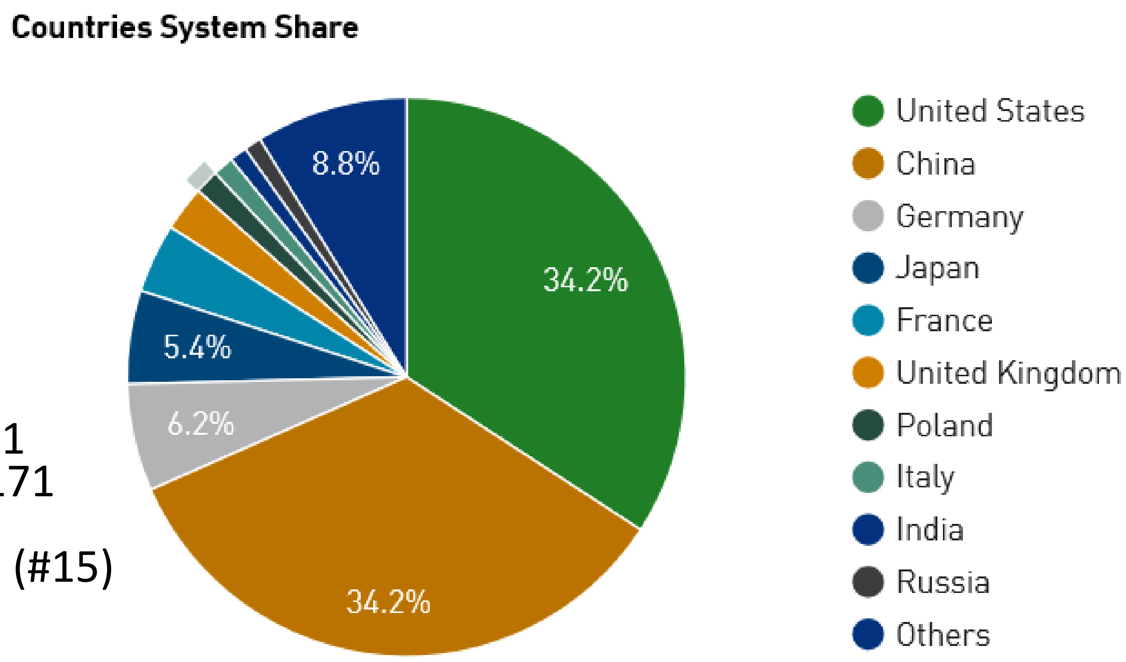

EUA: 171
China: 171
...
Brasil: 3 (#15)

(Desempenho Agregado)

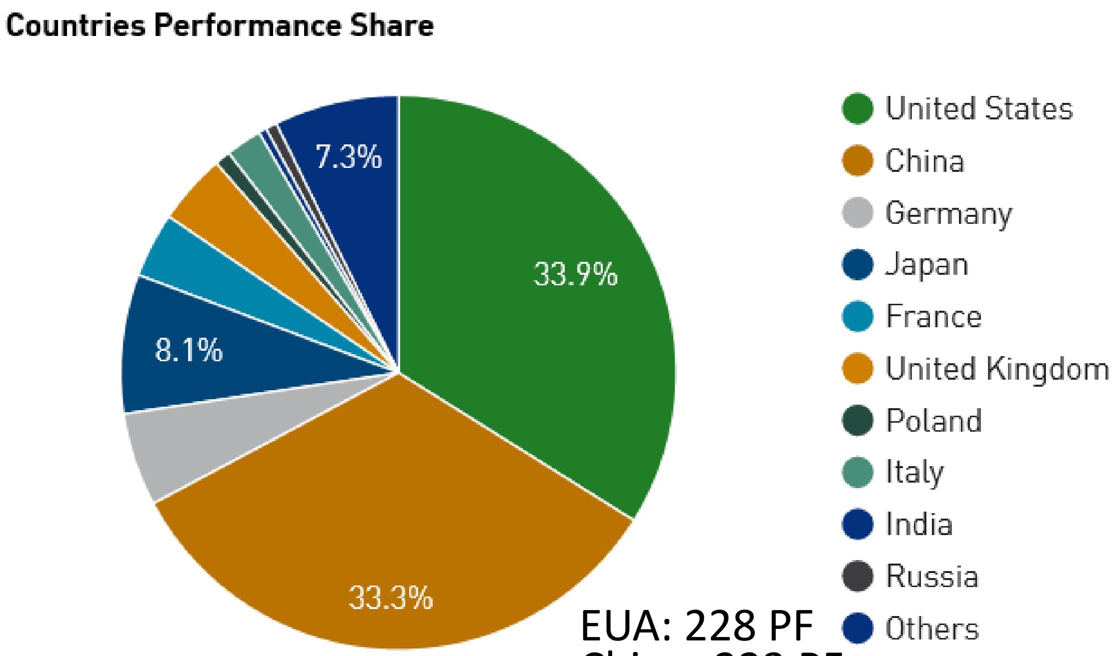

EUA: 228 PF
China: 223 PF
...
Brasil: 1.2 PF

SC16 – Top500

- Top500 Novembro/2016 – Outras Observações
  - Brasil: 3 sistemas (eram 4 em Junho/2016)
    - #364: S.Dumont/Bull, GPU   (ex #265)
    - #433: Cimatec/SGI   (ex #323)
    - #476: S.Dumont/Bull, Intel-KNC (ex #364)
  - Desempenho agregado (500 sistemas) cresceu 60% em 1 ano
    - Mais que a Lei de Moore (~41%)
  - Sistema #500: @349 TF (pico 510 TF)
    - Em Junho/2016: #500 @286 TF (pico 402 TF)

SC16 – Outros Destaques

- Nvidia: nova GPU P100
  - Sistema próprio, #28 no Top500, #1 Green500: 9.5 GFlop/Watt
  - Sistema #8 no Top500 (Suíça), #2 Green500: 7.5 Gflop/Watt
  - Alto engajamento empresarial em *Deep-Learning* (Intelig.Artif.)
- Fujitsu - K computer
  - processador Sparc será substituído por ARM no future
- Artigos Técnicos
  - Disponíveis nas bibliotecas digitais – IEEE e ACM

SC16 – Outros Destaques

- Gordon-Bell Award (Desempenho)
  - Artigo “10M-Core Scalable Fully-Implicit Solver for Nonhydrostatic Atmospheric Dynamics” (China)
    - Solver para dinâmica atmosférica
    - Executação: 10,5 milhões de núcleos da máquina #1 do Top500
    - Desempenho sustentado: 7,95 PF
      - ~6% do pico, “apenas”

10M-Core Scalable Fully-Implicit Solver for Nonhydrostatic Atmospheric Dynamics

Chao Yang*§, Wei Xue†‡**, Haohuan Fu‡**, Hongtao You¶, Xinliang Wang†‡, Yulong Ao*§, Fangfang Liu*§, Lin Gan†‡**, Ping Xu†‡, Lanning Wang||, Guangwen Yang†‡**, Weimin Zheng†
*Institute of Software and State Key Laboratory of Computer Science, Chinese Academy of Sciences, China
†Department of Computer Science and Technology, Tsinghua University, China
‡MOE Key Lab for Earth System Modeling, and Center for Earth System Science, Tsinghua University, China
§University of Chinese Academy of Sciences, China
¶National Research Center of Parallel Computer Engineering and Technology, China
||College of Global Change and Earth System Science, Beijing Normal University, China
**National Supercomputing Center in Wuxi, China

***Abstract*—An ultra-scalable fully-implicit solver is developed for stiff time-dependent problems arising from the hyperbolic conservation laws in nonhydrostatic atmospheric dynamics. In the solver, we propose a highly efficient hybrid domain-decomposed multigrid preconditioner that can greatly accelerate the convergence rate at the extreme scale. For solving the overlapped subdomain problems, a geometry-based pipelined incomplete LU factorization method is designed to further exploit the on-chip fine-grained concurrency. We perform systematic optimizations on different hardware levels to achieve best utilization of the heterogeneous computing units and substantial reduction of data movement cost. The fully-implicit solver successfully scales to the entire system of the Sunway TaihuLight supercomputer with over 10.5M heterogeneous cores, sustaining an aggregate performance of 7.95 PFLOPS in double-precision, and enables fast and accurate atmospheric simulations at the 488-m horizontal resolution (over 770 billion unknowns) with 0.07 simulated-years-per-day. This is, to our knowledge, the largest fully-implicit simulation to date.**

***Index Terms*—atmospheric modeling; fully implicit solver; Sunway TaihuLight supercomputer; heterogeneous many-core architecture.**

I. Justification for ACM Gordon Bell Prize

An important attempt is made to design an ultra-scalable fully-implicit solver for nonhydrostatic atmospheric simulations. With both algorithmic and optimization innovations, the solver scales to 10.5-million heterogeneous cores on Sunway TaihuLight at an unprecedented 488-m resolution with 770-billion unknowns, sustaining 7.95 PFLOPS performance in double-precision with 0.07 simulated-years-per-day (SYPD).

Performance Attributes	Content
Category of achievement	*Time-to-solution*
Type of method used	*Fully implicit*
Results reported on basis of	*Whole application including I/O*
Precision reported	*Double precision*
System scale	*Measured on full-scale system*
Measurement mechanism	*Timers*

II. Simulation of Atmospheric Dynamics

Every year, extreme weather/climate events may bring economic loss in hundreds of billion dollars [1] and sometimes cause catastrophic disasters to the living condition of human beings [2]. Ever since the ENIAC system in 1950s, generations of scientists have been continuously working on improving the simulation and prediction capability of atmosphere models by developing innovative numerical algorithms on state-of-the-art computing platforms [2]. With six decades passed, the continuous advance of the scientific understanding about the climate system, the computing methods, and the computing capabilities have finally pushed us to the edge of seamless weather-climate simulations/predictions at the resolution of the km-level and beyond.

On the road to the seamless weather-climate prediction, a major obstacle is the difficulty of dealing with various spatial and temporal scales [3]. The atmosphere contains time-dependent multi-scale dynamics that support a variety of wave motions. For example, the seasonal Asian summer monsoon usually comes at the planetary length scale of the earth with the order of 10^3-10^4 km, but thunderstorms and tornadoes often develop in minutes with an horizontal scale range of 10 km to

SC16; Salt Lake City, Utah, USA; November 2016
978-1-4673-8815-3/16/$31.00 ©2016 IEEE