



Driving Industry-University Cooperation Programs
to leverage Parallel & Distributed Computing
- Early experiences at SPRACE -

Rogério L. Iope
SPRACE/UNESP

INFIERI 7th Workshop
12-15 April 2016 - Lisboa, Portugal



- São Paulo State:
 - Area
 - ▣ ~United Kingdom
 - Population
 - ▣ ~Spain
 - Economy GDP
 - ▣ 2 X Switzerland
 - Investment in S & E
 - ▣ 1.66% GDP (~Argentina)



- Second largest among all Brazilian universities
- A successful model of a multicampus university

UNESP in Numbers

- Campuses: 24
- Institutes: 41
- Budget: US\$ 1 Billion
- Faculty: 3,900
- Alumni: 146,000
- Students: 51,000
 - Undergraduate: 38,000
 - Graduate: 13,000
- Programs
 - Undergraduate: 181
 - Graduate: 231
- Facilities
 - 6 Associated Hospitals
 - 30 Libraries
 - 3 Veterinary Hospitals
 - 5 Teaching and Research Farms
- Research groups: 600
- Scientific impact
 - 8% of Brazilian papers
- QS University Rankings 2015
 - Brazil: 5
 - Latin America: 17
 - BRICS: 27
 - Under 50 y: 71-80
 - World: 481-490
 - ❑ Dentistry: 31
 - ❑ Veterinary: 45
 - ❑ Agriculture-Forestry: 51
 - ❑ Pharmacology: 101
 - ❑ Environment, Engineering Physics, : 201
 - ❑ Biology, Medicine: 251

The SPRACE Project

- São Paulo Research and Analysis Center
 - High Energy Physics Group
 - ❑ Participation in the CMS Experiment of LHC at CERN
 - ❑ Funded by FAPESP (Thematic Project)
 - ❑ Started in 2004
- Physics Analysis
 - Physics Beyond the Standard Model
 - ❑ Extra-dimensions, Dark Matter
 - Heavy-Ion Collisions
 - ❑ Strong interaction at high density and temperature



The SPRACE Project

- BR-SP-SPRACE

- First Official WLCG Tier-2 in LA

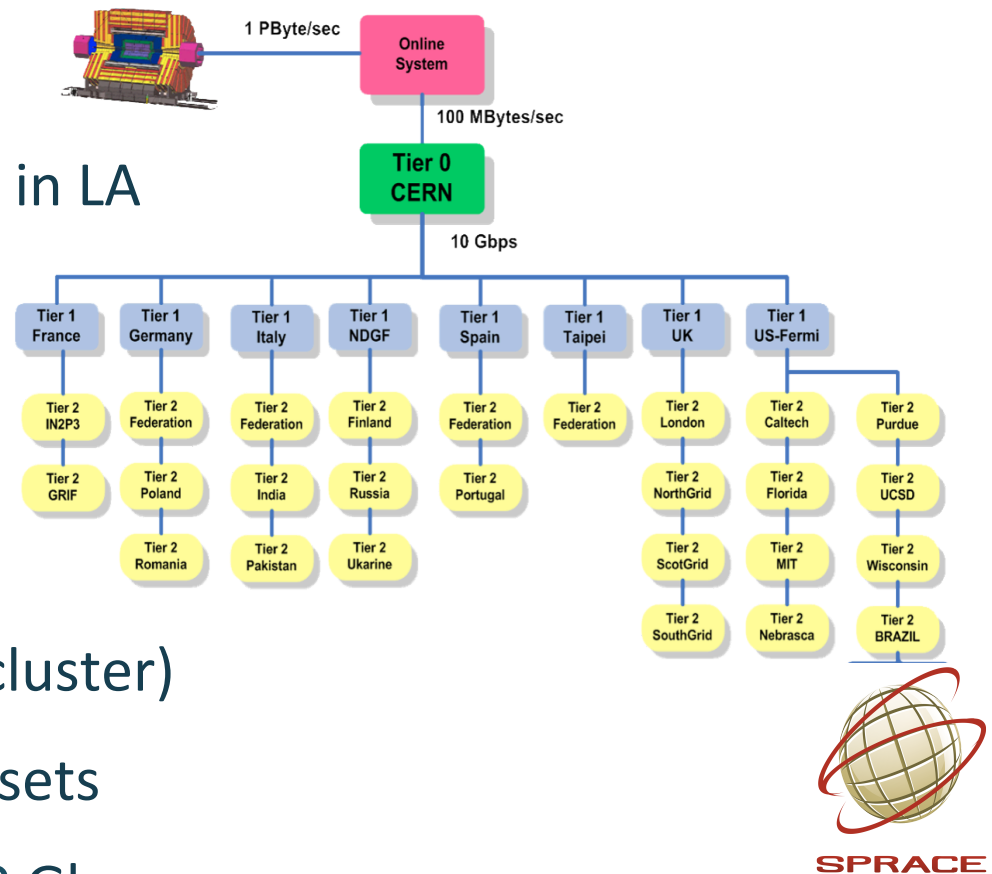
- MoU signed in April/2009
(FAPESP & CERN)

- Physics analysis and MC simulation for CMS

- Computing facility (HPC cluster)

- ~ 1PB of storage for datasets

- Network connection: 100 Gbps



One of the 'Top 10' WLCG Tier-2 sites



Tier-2 Availability and Reliability Report

CMS

January 2014

Federation Summary - Sorted by Reliability

Colour coding :

N/A

<30%

<60%

<90%

>=90%

Availability Algorithm: (OSG-CE + CE + CREAM-CE + ARC-CE) * (SRMv2 + OSG-SRMv2)

Federation	Availability	Reliability
BR-SP-SPRACE	100 %	100 %
FI-HIP-T2	100 %	100 %
PT-LIP-LCG-Tier2	100 %	100 %
CERN-PROD	100 %	100 %
GR-Ioannina-HEP	100 %	100 %
UK-London-Tier2	99 %	100 %
HU-HGCC-T2	99 %	99 %
PL-TIER2-WLCG	99 %	99 %
EE-NICPB	99 %	99 %
T2_US_MIT	99 %	99 %
CH-CHIPP-CSCS	98 %	98 %
FR-IN2P3-IPHC	98 %	98 %
DE-DESY-RWTH-CMS-T2	98 %	98 %
T2_US_Florida	98 %	98 %
IN-INDIACMS-TIFR	98 %	98 %
DE-DESY-ATLAS-T2	98 %	98 %
ES-CMS-T2	98 %	98 %
IT-INFN-T2	98 %	98 %
T2_US_Wisconsin	98 %	98 %



Federation	Availability	Reliability
CN-IHEP	97 %	97 %
T2_US_Caltech	97 %	97 %
T2_US_Purdue	97 %	98 %
T2_US_Nebraska	96 %	96 %
T2_US_UCSD	96 %	100 %
FR-GRIF	95 %	99 %
KR-KNU-T2	94 %	95 %
UK-SouthGrid	94 %	95 %
BE-TIER2	93 %	93 %
UA-Tier2-Federation	93 %	99 %
RU-RDIG	92 %	94 %
FR-IN2P3-CC-T2	91 %	91 %
AT-HEPHY-VIENNA-UIBK	89 %	89 %
TR-Tier2-federation	88 %	88 %
PK-CMS-T2	80 %	80 %
TW-FTT-T2	70 %	70 %
Taiwan-LCG2	70 %	70 %
T2-LATINAMERICA	31 %	37 %

GridUNESP Project: a SPRACE spinoff

- First Campus Grid in Latin-America

- Provides Scientific Computing to

- ❑ 60+ Scientific Projects

- ❑ 350+ Researchers

- Partnership with US Open Science Grid

- ❑ The only VO outside US

- ANSP Grid CA

- ❑ Grid Certificate Authority for the State of São Paulo

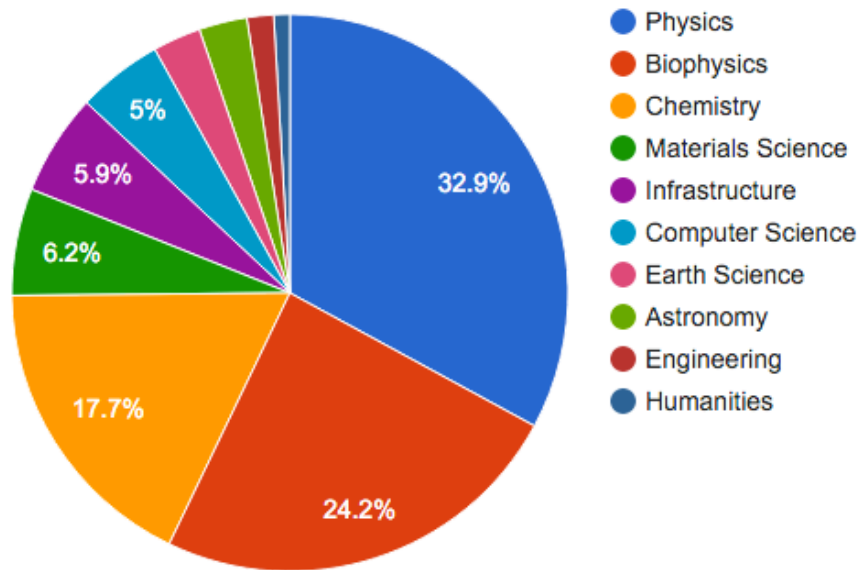


- UNESP “Center for Scientific Computing” (total investment: over USD 7 million)




GridUNESP: research projects and users

User distribution by research field



General-purpose infrastructure
serves distinct research fields

Evolution of subscribed users and projects



64 projects, 360 users
(as of April 2016)

Driving (small) Industry-University Cooperation Programs

Intel / Unesp first joint activities

- Cloud computing development (2012)
 - Cloud infrastructure (IaaS) focused on educational activities
 - Partnership with São Paulo Secretary of Education
 - ❑ Automated, customized load balancing system for deploying on-demand VMs for high-school teachers to use during classes
 - Two OpenStack clouds deployed at
 - ❑ Secretary of Education datacenter
 - ❑ UNESP CSC datacenter
- Manycore Testing Lab (2013)
 - An experimental platform for testing, validation and scaling of parallel algorithms and workloads for the academic community
 - ❑ primarily for courseware delivery
 - UNESP was the first Brazilian academic institution to receive a Xeon Phi card from Intel (Jan 2013)


Intel / Unesp Manycore Testing Lab

- One of the first manycore labs outside U.S.
 - Server donated by Intel
 - ❑ Intel Xeon Phi coprocessors
 - ❑ Suite of Intel's software development tools
 - Users
 - ❑ test the performance of their codes in a highly parallel system
 - ❑ are registered as guests into Unesp CSC user database (LDAP)
 - Authentication / authorization
 - ❑ can be controlled by digital certificates issued by ANSP Grid CA
- First results: hands-on activities at
 - INFIERI Summer School 2013 (University of Oxford)
 - Intel Software Conferences 2013 (UNESP/SP and COPPE/RJ)
 - SBAC/PAD 2013 Parallel Programming Marathon
 - INFIERI Summer School 2014 (Université Paris-Diderot)



intel
Xeon Phi™ Coprocessor

Delta
BE740



intel
Xeon Phi™ Coprocessor

900

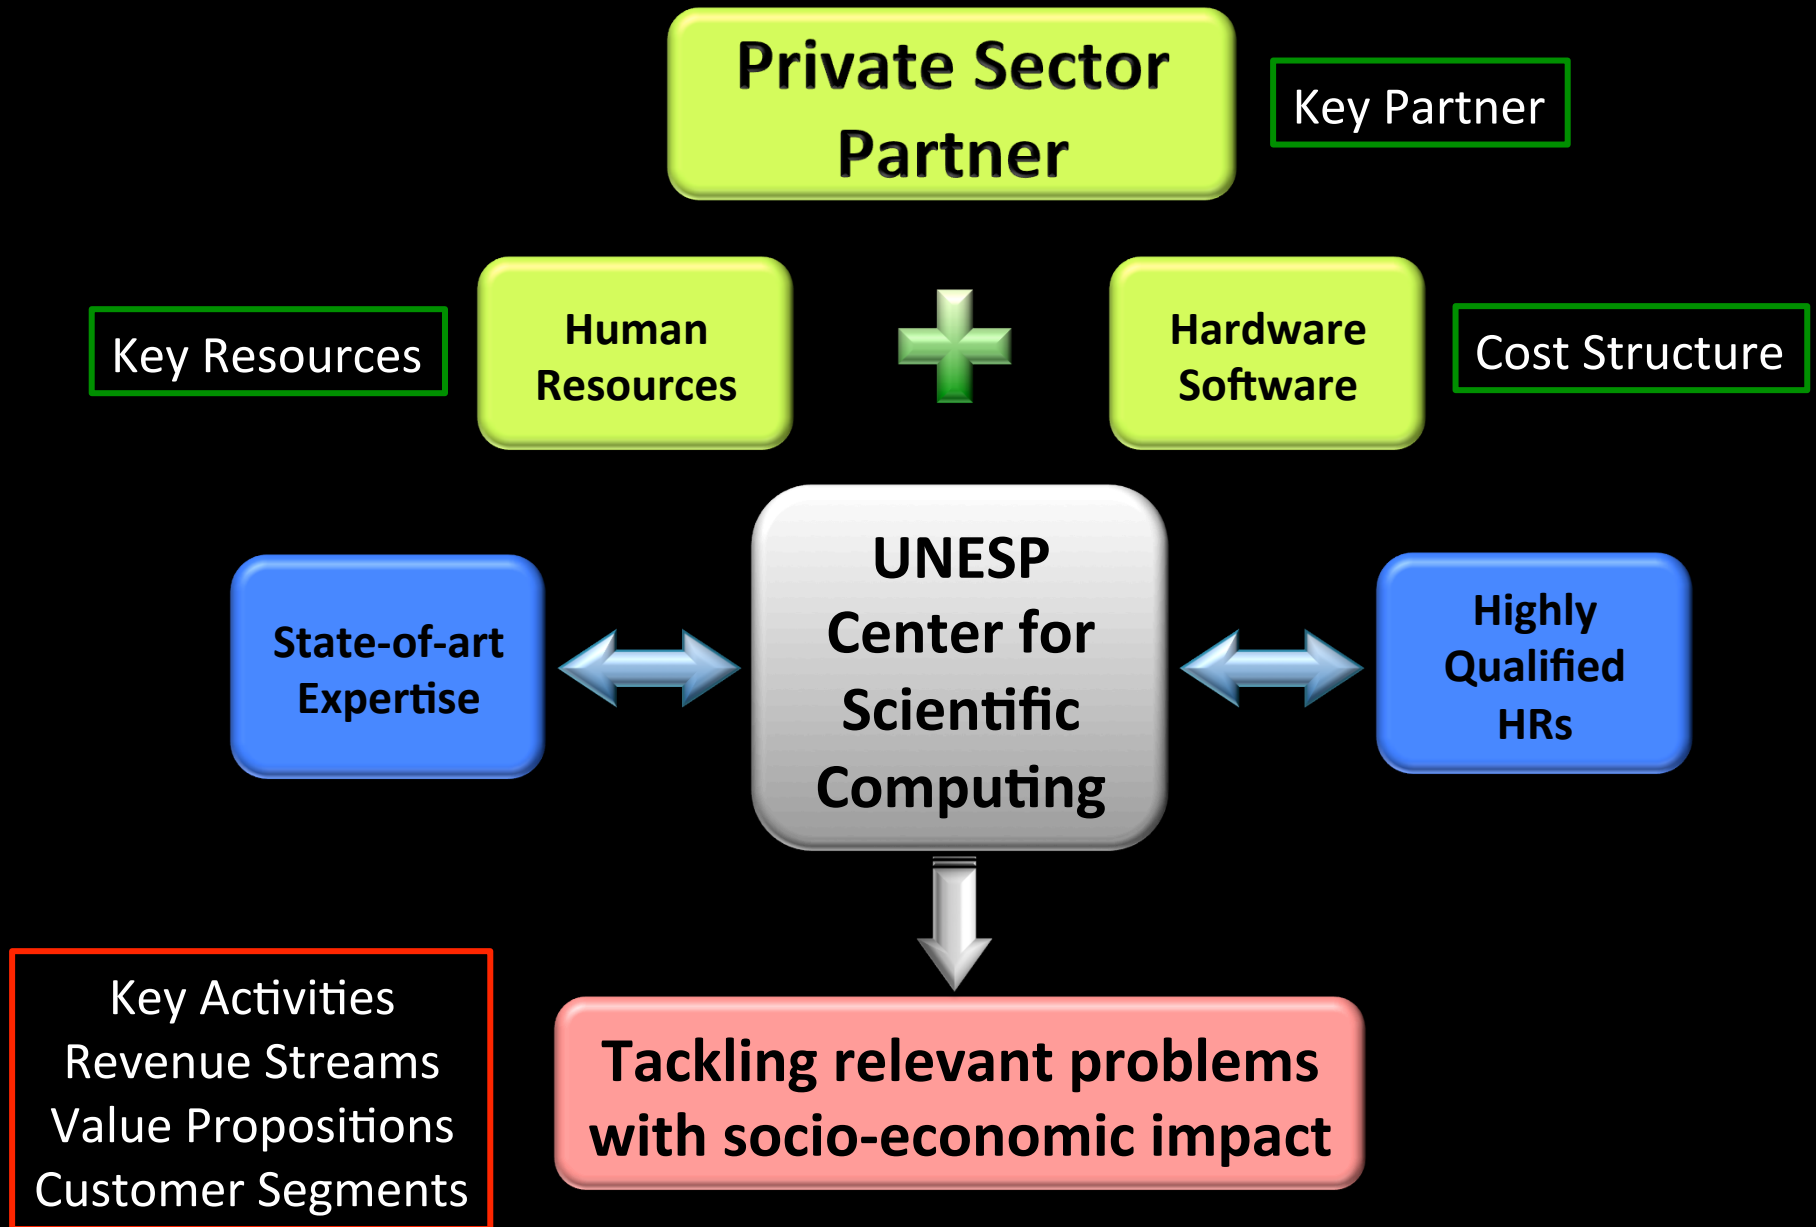
Important
Intel Xeon Phi
COPROCESSOR

FLAVORON

ASUS
M2N68-MLE
M2N68-MLE

"BUSINESS MODEL"

(Ref.: <http://www.businessmodelgeneration.com/>)



1. Intel® Parallel Computing Centers



- A Parallel Computing Center @UNESP
 - Project started in November 2014
 - Part of a select group of 64 leading R&D institutions from all over the world
 - R&D efforts to adapt HEP software tools and explore manycore architectures
 - Parallelization/vectorization of GEANT
 - Simulation of radiation-matter interaction
 - Broad Impact
 - Extremely important tool for HEP ...
 - ... and for other relevant S&E fields of socio-economic impact (e.g. medical applications)

- Partnership
 - CERN Geant-V development team
 - Fermilab Computing Division

Intel® PCC at UNESP

PARALLELIZATION OF HEP SOFTWARE TOOLS IN THE MANYCORE ERA

GUILHERME AMADIO¹, CALEBE BIANCHINI¹, ROGÉRIO IOPE¹, ANDREI GHEATA², SOFIA VALLECORSA² AND SANDRO WENZEL²

¹NCC/UNESP, ²CERN
for the GeantV Project (geant.cern.ch)

OVERVIEW

Geant4 (GEometry ANd Tracking) is the toolkit of choice for the simulation of particle interaction with matter. Its development started almost twenty years ago and has evolved since then with the effort of an international collaboration of physicists and computer scientists from many institutions. Developers interact constantly with users in a combined effort to validate the results for application in high energy physics (HEP) experiments, space projects and medical studies.

Intel® Parallel Computing Centers [1] aim to modernize applications to increase parallelism and scalability through optimizations that leverage cores, caches, threads, and vector capabilities of state-of-art microprocessors and coprocessors, such as UNESP-IPCC and CERN-IPCC. The Center for Scientific Computing at São Paulo State University (NCC/UNESP) is mainly involved in R&D efforts to adapt HEP software tools, including Geant, in order to exploit modern computing architectures that support multi-threading and other parallel processing techniques to make data processing more cost-effective. Geant requires long calculation times and it is ideal for compute-bound workloads that are suited for execution on Intel Xeon Phi coprocessors. The plan for code performance improvements at UNESP-IPCC is focused on testing vector-coprocessor prototypes in hybrid parallel computing systems and analyzing the performance of the next generation of Intel Xeon Phi coprocessor (KNL), evaluating various redesign strategies for this new platform. Ongoing activities are mainly related to the development of GeantV, the next generation of the Geant simulation engine, which will include massive parallelism natively at the track level.

- Intel funding: **USD 250,000.00** for a 2-year program (2 postdoc-level fellowships + extra budget for travel expenses, hardware, etc)
- Budget covers travel expenses to Fermilab and CERN (3-month periods, twice a year at each place, for each fellow)

2. Intel® Modern Code Partner Program

Training and support on the exploitation of multithreading and vectorization for modern multicore and manycore architectures

North America



India



Europe

BAYNCORE

Korea

DAEHAN CONSULTING GROUP




China

inspur 浪潮

PARATERA 并行

Latin America



- Project started in May 2015
- International events
 - INFIERI Summer Schools
 - University of Oxford
 - Paris Diderot University
 - University of Hamburg
 - IEEE/ACM CCGrid 2016 (Colombia)
 - VECPAR 2016 (Portugal)
- Brazilian HPC Regional Schools
 - ERAD-SP, -NE, -RJ, -RS
- Intel special events:
 - Software Days
 - Xeon Phi Coprocessor Workshops
 - Permanent contact with Intel USA
- Regular events (near future):
 - São Paulo State Technological Colleges (FATECs)
 - National Center for HPC in São Paulo (CENAPAD-SP)

<http://modern-code.ncc.unesp.br/events>

<https://software.intel.com/en-us/modern-code/live-workshops>

Modern Code: Training activities

2015			
Event name	Planned attendees	Actual attendees (Hands-on)	Actual attendees (Lectures)
Code Modernization Workshop @ UNESP	90	20	61
Intel Software Day	70	8	120
Engineering Week - UNESP Ilha Solteira Campus	20	15	200
Code Modernization Workshop @ ERAD-SP	20	15	-
Code Modernization Workshop @ ERAD-RJ	41	26	100
INFIERI Summer School @ Hamburg	20	20	-
Code Modernization Workshop @ ERAD-NE	100	30	75
Code Modernization Workshop @ CINTEC	60	13	50
Workshop for graduate students at Poli-USP	35	35	-
Total	261	182	481

2016			
Event name	Planned attendees	Actual attendees (Hands-on)	Actual attendees (Lectures)
Code Modernization Workshop @ UNESP	50	16	44
Code Modernization Workshop - TOTVS training	20	18	18
Code Modernization Workshop @ CENAPAD-SP	44	12	15
Code Modernization Workshop @ UFRN	140	27	54
Code Modernization Workshop @ ERAD-RS 2016			
Code Modernization Workshop @ CCGRID 2016			
Code Modernization Workshop @ VECPAR 2016			
Code Modernization Workshop @ FATEC-Santos			
Code Modernization Workshop @ UFES			
Total	254	73	131



Workshop on Parallel Programming and Optimization for Intel® Architecture @ CENAPAD-SP

14-15 March 2016

CENAPAD-SP - Centro Nacional de Processamento de Alto Desempenho em São Paulo

America/Sao_Paulo timezone

Overview

Programme

Agenda

Venue

Registration

Registration Form

Participant List

Contact Information

✉ events@ncc.unesp.br

☎ +55 11 3393-7780

The Intel Xeon Phi coprocessor, the first product of Intel's Many Integrated Core (MIC) Architecture, is a new accelerator technology developed by Intel to enable performance gains for highly parallel computing workloads. It possesses several interesting and appealing features, including the ability to use familiar programming models such as OpenMP and MPI.

The series of workshops on Parallel Programming and Optimization, offered by *Universidade Estadual Paulista (UNESP)* in partnership with *Intel Software do Brasil*, aims to provide a comprehensive, practical introduction to parallel programming and optimization techniques based on open standards and frameworks in order to fully utilize the scaling capabilities of Intel Xeon processor-based systems. They have been conceived with a special focus on the active participation of the attendees.

The first day provides a general introduction to the Intel Xeon Phi coprocessor. Participants will learn about the architecture, software infrastructure, supported programming models, and OpenMP and MPI programming and analysis.

The second day builds on information learned during the first day and provides practical coverage with hands-on activities. Participants will work on predefined sets of exercises that address a wide range of aspects aimed to help them get more familiar with the Intel Xeon Phi coprocessor architecture.

This event is sponsored by:



<https://indico.ncc.unesp.br/> → Intel / Unesp Modern Code events

For information on upcoming events please visit: <http://modern-code.ncc.unesp.br/events>

Modern Code: Milestones & Funding

- Project milestones / deliverables:
 - Regular parallel programming and code modernization training workshops using Intel hardware and Intel software development tools
 - Training initiatives on Data Science using Intel DAAL (Data Analytics Acceleration Library) and similar tools - exploring confluences of HPC and Big Data
 - Technical support to local developer community
 - Operations and maintenance on hardware assets
- Intel funding: **USD 150,000.00**
 - USD 95,000.00 for a 2-year program (1 full-time + 1 half-time postdoc-level fellowships + extra budget for covering instructors' travel expenses)
 - USD 55,000.00 for hardware and software:
 - ❑ 2 top-level servers with 4 Xeon Phi coprocessors (on each server)
 - ❑ Licenses for all Intel software tools
 - Extras
 - ❑ “Training the trainers” program (regular meetings & workshops w/ Intel experts)
 - ❑ Ongoing: we are waiting for a new server with 2nd generation Intel Xeon Phi (Knights Landing) - see next slide

Developer Access Program for KNL

Ninja Developer Platform Pedestal



Ninja Developer Platform Rack



- Developer Edition of Intel® Xeon Phi™ Processor: 16GB MCDRAM, 6 Channels of DDR4, AVX 512
- Liquid cooled
- MEMORY: 6x DIMM slots
- EXPANSION: 2x PCIe 3.0 x16, 1x PCIe 3.0 x4 (in a x8 mechanical slot)
- LAN: 2x Intel® i350 Gigabit Ethernet
- STORAGE: 8x SATA ports, 2x SATADOM support
- POWER SUPPLY: 1x 750W 80 Plus Gold
- CentOS 7.2
- Intel® Parallel Studio XE Professional Edition Named User 1 year license

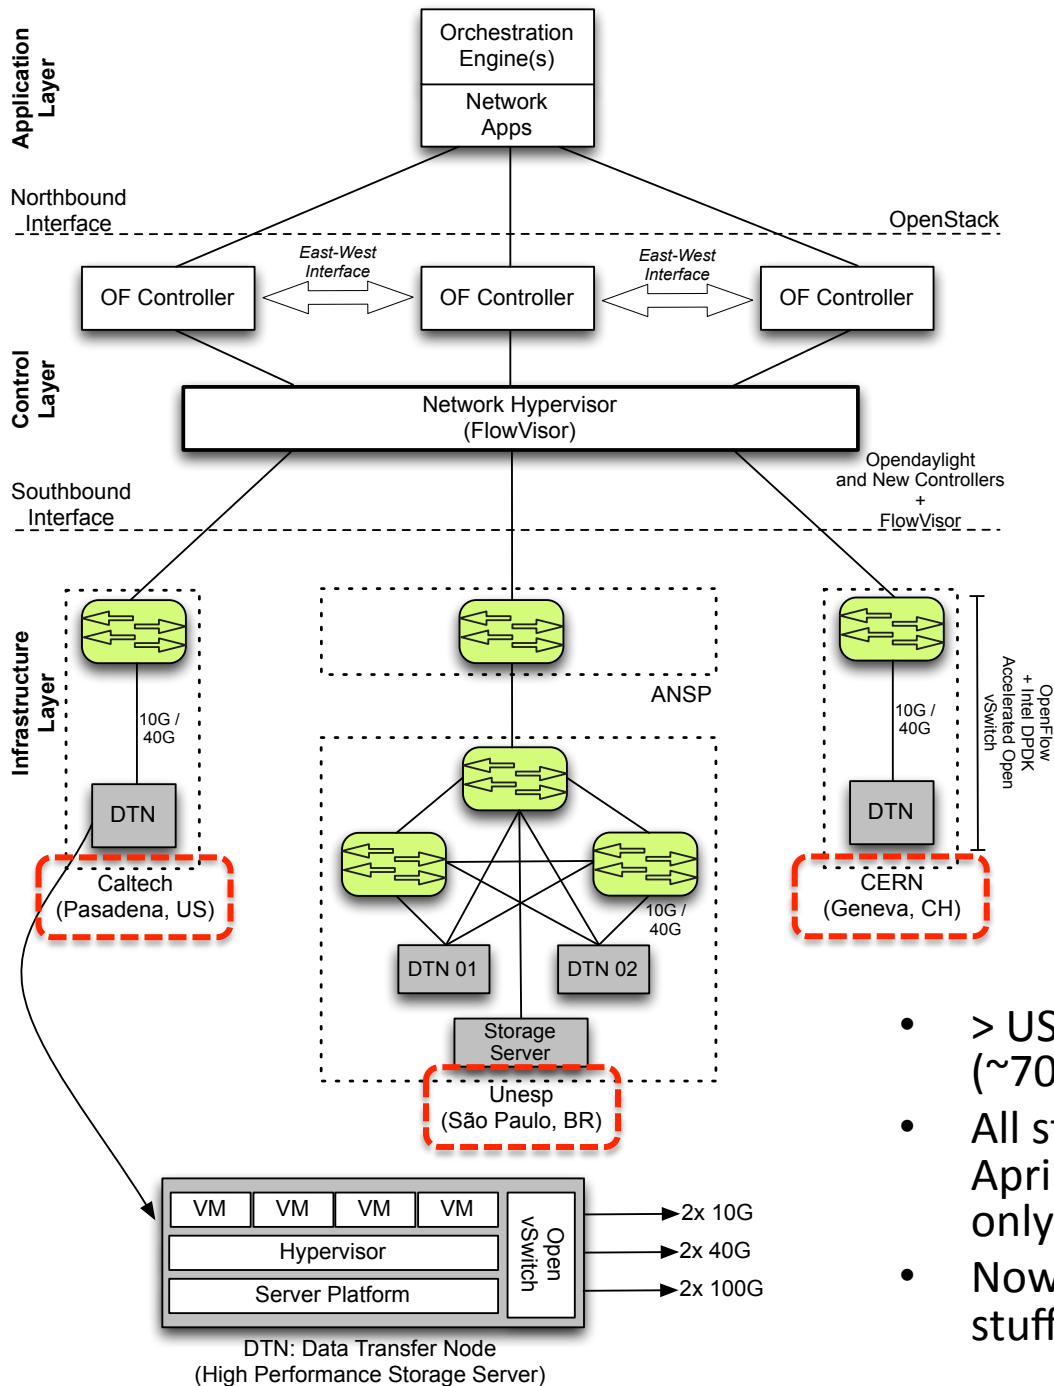
- Developer Edition of Intel® Xeon Phi™ Processor: 16GB MCDRAM, 6 Channels of DDR4, AVX 512
- 2U 4x Hot-Swap Nodes
- MEMORY: 6x DIMM slots / Node
- EXPANSION: Riser 1: 1x PCIe 3.0 x16, Riser 2: 1x PCIe Gen3 x 20 (x16 or x4) / Node
- LAN: 2x Intel® i210 Gigabit Ethernet / Node
- STORAGE: 12x 3.5" Hot-Swap Drives
- POWER SUPPLY: 2x 2130W Common Redundant 80 Plus Platinum
- CentOS 7.2
- Intel® Parallel Studio XE Cluster Edition Named User 1 year license

Ref.: <http://dap.xeonphi.com/>

3. Huawei: R&D on SDN



- Huawei Technologies Co. Ltd.
 - Leading global ICT provider
 - Largest telecom equipment manufacturer in the world
 - Over 170K employees (more than 45% engaged in R&D)
- R&D on Software-Defined Networking (SDN) over WAN for Data-Intensive Science
 - Project started in January 2016
 - Development of SDN for Global Scale Science with Caltech & CERN
 - International OpenFlow + OpenStack testbed
 - ❑ 3 “islands”: Unesp (Brazil), Caltech (USA), CERN (Europe)
 - Deployment of a high-end data transfer experimental system
 - ❑ São Paulo – Miami at 100 Gbps
 - Demonstrations at annual Supercomputing Conference in U.S.



- > USD 280K in equipment (~700 Kg)
- All stuff arrived at UNESP on April 1st (import process took only 2 months)
- Now we need to export some stuff to Caltech and CERN

R&D in SDN: Milestones & Funding

- Project milestones / deliverables
 - System design and testbed deployment
 - Deployment of web portal and monitoring tools
 - Deployment of cloud-based tools on each island
 - 100G “alien wavelength” prototype system (DWDM)
 - Development of control tools for WAN network orchestration
 - Development of an open-source SDN controller
- Huawei funding: **USD 500,000.00** per year (up to 3 years)
 - 20-40% in hardware/software each year
 - 4-6 full-time fellowships (Master / PhD level)
 - Extra budget for covering 3rd party services and travel expenses

Lessons learned (so far)

- Project management for R&D is definitely far from trivial
 - the “research” part of the project is the main issue: the work is subject to unexpected developments and results
 - managers have to be capable of continuously adjusting to new situations
- Systems engineering complements project management
 - an experienced engineer that assumes direct responsibilities for the development and control of activities at a much deeper level compared to a non-technical project manager is key to success
- PI and management staff (PM, SE) need to handle a huge amount of non-technical stuff
 - a permanent interaction with lawyers and experts from the University Office of Technology Transfer (TTO) is mandatory
- The most fruitful partnerships take time and effort to bear fruit
 - Unesp team has been interacting with Intel Brazil representatives for more than 10 years

Engineering Team

Unesp

Marcio A. Costa
Sidney T. Santos
Jadir M. da Silva
Allan Szu

FundUnesp

Raphael M. O. Cobe
Rogério L. Iope
Beraldo C. Leal
Thiago C. Paiva
Gabriel A. von Winckler

Fapesp

Artur Baruchi

Intel

Guilherme Amadio
Calebe P. Bianchini
André L. V. Pereira
Silvio Stanzani

Huawei

Diego R. Oliveira
André T. de Carvalho



Thank you

Rogério L. Iope
SPRACE/UNESP

rogerio@ncc.unesp.br