High Performance Computing and High Performance humans in Computing



Maria Grazia Pia INFN Genova

mariagrazia.pia@ge.infn.it maria.grazia.pia@cern.ch

IAEA (Virtual) Workshop on Computational Nuclear Science and Engineering 12-16 July 2021

MGP

Does not like writing bios

- Physicist at INFN (Istituto Nazionale di Fisica Nucleare), Section of Genova, Italy
 - Pre-covid: large fraction of time at CERN, looking forward to post-covid era
- Associate Editor of IEEE Transactions on Nuclear Science
- Moderator of arXiv Computational Physics (physics.comp-ph)
- High Energy Physics background
 - CERN, FNAL, SLAC
- Monte Carlo development (Geant4) and applications, physics data libraries validation
- Statistical data analysis, scientometrics, epistemology
- No Facebook, Twitter, LinkedIn, Instagram... but google me, and you will find track of my research activity

What is HPC?







HPC systems

- HPC computers are networks of processors
 - Very fast memories
 - Low-latency, high-bandwidth communication systems
 - between the processors
 - between the processors and the associated memories
- Homogeneous: only CPUs
- Hybrid: CPUs + GPUs
 - GPUs can handle millions of threads simultaneously, are more energy efficient, have faster memories, require less data transfer
 - **CPU**s oversee the computation

Key features of **HPC architecture**

speed of operation

- **parallelism** to perform multiple operations
- efficient use of critical components
- electrical power that it consumes
- reliability
- how easy to program

Top supercomputers

	Rmax PFLOPS	Location	Manufacturer		
Fugaku	442.010	RIKEN, Japan	Fujitsu	Performance Country PELOPS	
Summit	148.600	ORNL, USA	IBM		
Sierra	94.640	LLNL, USA	IBM		669
Sunway TaihuLight	93.015	Natl. Supercomp. Centre, China	NRCPC	Japan	594
Perlmutter	64.590	NERSC, USA	Nvidia	China	564

FLOPS: floating-point operations per second



Commodity clusters

A group of integrated computer systems

- Standalone components (COTS), capable of independent operation
- Integration network is separately developed
- Off the shelf mass storage
- Interfaces adhere to industry standards
- Very successful
 - First one in 1997, ~50% of Top 500 systems in 2005, ~ 85% today

Parallel programming modalities

- **Throughput computing:** efficiently run a large number of jobs that are independent or require minimal communication
- **Message-passing:** requires a significant amount of communication and coordination within the application
 - communicating sequential processes model, exemplified by the MPI
- Shared-memory multiple-thread applications

exemplified by the OpenMP (open multiprocessing) programming model
Maria Grazia Pia, INFN Genova

LHC

WLCG

LHC data are currently handled by the Worldwide LHC Computing Grid

- >170 sites
- 42 countries
- ~10⁶ CPU cores
- 1 exabyte storage
- 2 million tasks per day
- global transfer rates > 60 GB/s
- >12000 physicists around the world

2028

2024

WLCG is a massive **distributed computing** infrastructure



accelerated event reconstruction and simulation

HPC requires suitable software to benefit from the hardware

Parallel algorithms

Physical
SIMD: single-instruction multiple data parallelism
parallelism
MIMD: multiple-instruction multiple data

- shared memory parallelism
- distributed memory parallelism
- Multiple parts of the workload are performed concurrently to reduce the time to achieve the solution
- Several parallel algorithms are used in scientific computing
 - fork–join
 - divide and conquer
 - halo exchange
 - permutation
 - embarrassingly parallel
 - manager–worker
 - task dataflow

Emerging numerical methods in supercomputing applications:

- graph traversal
- finite state machines
- combinational logic
- statistical machine learning

Some algorithms are better suited for one kind of physical parallelism versus another

Maria Grazia Pia, INFN Genova

Chances are that you would deal with HPC and parallel computing in

Monte Carlo simulation in particle/nuclear physics

Embarrassingly parallel

- Parallelism with essentially no inter-task communication
- Highly partitionable workload with minimal overhead
- Concurrency is trivially extracted from the workflow
- Often require gathering the results at the end into a manager process
- Monte Carlo transport is naturally suitable to event-level parallelism
- Methods for embarrassing parallel simulation documented in most popular Monte Carlo particle transport codes
 - Threading: OpenMP threading on a single multicore computer or on a single node of a server or cluster
 - Message-passing between nodes on a cluster using MPI environment
 - Used separately or together

Maria Grazia Pia, INFN Genova

Multi-threaded simulations

Ability to exploit hardware multi-threading capabilities

Goal: **reduce** the **memory** footprint of parallel applications, while preserving the **linear speedup** as a function of the number of physical cores

Based on a master-worker model





- master thread prepares geometry and physics setups
- worker threads compete for the next events to be simulated

an application example

Min Cheol Han *et al*, **Multi-threading performance of Geant4, MCNP6, and PHITS** Monte Carlo codes for tetrahedral-mesh geometry, 2018 *Phys. Med. Biol.* 63 09NT02 *Maria Grazia Pia, INFN Genova* **Quantum computing** is a computing paradigm that exploits quantum mechanical properties (superposition, entanglement, interference...) of matter to do calculations

Quantum computing

Quantum supremacy using a programmable superconducting processor

Nature volume 574, pp. 505-510 (2019)

https://doi.org/10.1038/s41586-019-1666-5 Received: 22 July 2019 Accepted: 20 September 2019 Published online: 23 October 2019



Frank Arute¹, Kunal Arya¹, Ryan Babbush¹, Dave Bacon¹, Joseph C. Bardin^{1,2}, Rami Barends¹, Rupak Biswas³, Sergio Boixo¹, Fernando G. S. L. Brandao^{1,4}, David A. Buell¹, Brian Burkett¹, Yu Chen¹, Zijun Chen¹, Ben Chiaro⁵, Roberto Collins¹, William Courtney¹, Andrew Dunsworth¹, Edward Farhi¹, Brooks Foxen^{1,5}, Austin Fowler¹, Craig Gidney¹, Marissa Giustina¹, Rob Graff¹, Keith Guerin¹, Steve Habegger¹, Matthew P. Harrigan¹, Michael J. Hartmann^{1,6}, Alan Ho¹, Markus Hoffmann¹, Trent Huang¹, Travis S. Humble⁷, Sergei V. Isakov¹, Evan Jeffrey¹, Zhang Jiang¹, Dvir Kafri¹, Kostyantyn Kechedzhi¹, Julian Kelly¹, Paul V. Klimov¹, Sergey Knysh¹, Alexander Korotkov^{1,8}, Fedor Kostritsa¹, David Landhuis¹, Mike Lindmark¹, Erik Lucero¹, Dmitry Lyakh⁹, Salvatore Mandrà^{3,10}, Jarrod R. McClean¹, Matthew McEwen⁵, Anthony Megrant¹, Xiao Mi¹, Kristel Michielsen^{11,12}, Masoud Mohseni¹, Josh Mutus¹, Ofer Naaman¹, Matthew Neeley¹, Charles Neill¹, Murphy Yuezhen Niu¹, Eric Ostby¹, Andre Petukhov¹, John C. Platt¹, Chris Quintana¹, Eleanor G. Rieffel³, Pedram Roushan¹, Nicholas C. Rubin¹, Daniel Sank¹, Kevin J. Satzinger¹, Vadim Smelyanski¹, Kevin J. Sung¹¹³, Matthew D. Trevithick¹, Amit Vainsencher¹, Benjamin Villalonga^{11,4}, Theodore White¹, Z. Jamie Yao⁰, Ping Yeh¹, Adam Zalcman¹, Hartmut Neven¹ & John M. Martinis^{15,4}

The Sycamore processor can run a test computation in 200 seconds that would take the world's biggest supercomputers 10000 years to complete

IBM researchers claim Google's challenge would take a classical computer just two and half days

In quantum circuits:

- data = qubits
- operations = quantum gates
- results = measurements

NewScientist

China beats Google to claim the world's most powerful quantum computer

5 July 2021

TECHNOLOGY 5 July 2021



High Performance humans in Computing









Maria Grazia Pia, INFN Genova



Producing results

Measurements over 10 years Publicly accessible data Average productivity



[...] collaboration

Fred P. Brooks,

"No Silver Bullet - Essence and Accidents of Software Engineering"

IEEE Computer, vol. 20, no. 4, pp.10-19, April 1987

As we look to the horizon of a decade hence, we see **no silver bullet**. There is no single development, in either technology or in management technique, that by itself promises even one order-of-magnitude improvement in productivity, in reliability, in simplicity.

Although we see no startling breakthroughs - and indeed, I believe such to be inconsistent with the nature of software - many encouraging innovations are under way. A **disciplined**, **consistent effort** to develop, propagate, and exploit these innovations should indeed yield an order-of-magnitude improvement. There is no royal road, but there is a road.

Introduction to concepts and methods

- 📌 Pills of wisdom
- Food for thought

📌 Curiosity

Background for further learning



Software development **methods** and **techniques** are seldom part of academic programs for physics and engineering degrees

Cowboy programming



Emphasis on ingenious artistry

- Galloping off on one's own without a prior plan
- Brute-force programming
- Uncertain design requirements, code rewrite
- Quick and dirty: code and fix later
- Lack of comments, documentation, reviews
- Reinventing the wheel

The results are often spotty and difficult to duplicate



Inexperienced developers are unfamiliar with **technologies** and **methodologies** that support producing quality software effectively



These complex **disciplines** include **activities**, generate

products and involve responsibilities in various roles

articulated over the software **life-cycle**:

get-go, elaboration, construction, use, maintenance...

Built on **best practices** derived from experience

Software development methodologies are conceptual frameworks to structure, plan and control the process of developing the software



Old, risky... and most common



Risk of discovering problems at a late stage of the project

Variant of waterfall: V-model



Emphasis on testing at all levels of software development

Each development phase is associated with a testing phase

Grasp the nettle: non-linear view of the software life cycle

Spiral development

basis of modern methodologies



loop in the spiral = phase of software development

Barry W. Boehm, A Spiral Model of Software Development and Enhancement, IEEE Computer, vol. 21 no. 5, pp. 62-72, 1988

Unified Process (UP, USDP, RUP)

Iterative, incremental process, with emphasis on modeling



Complex, but highly **customizable**

Kent Beck et al. (2001)

Manifesto for Agile Software Development

https://agilemanifesto.org/

We are uncovering better ways of developing software by doing it and helping others do it. Through this work we have come to value:

Individuals and interactions over processes and tools Working software over comprehensive documentation Customer collaboration over contract negotiation Responding to change over following a plan

> That is, while there is value in the items on the right, we value the items on the left more.

Emphasis on

- Effective communication among all stakeholders
- Adaptive response to change
- Rapid, incremental delivery of software

Agile Methods



B. Boehm, "Get Ready for Agile Methods, With Care", IEEE Computer, 2002, http://dx.doi.org/10.1109/2.976920
B. Meyer, Agile!: The Good, the Hype and the Ugly, Springer, 2014
R. C. Martin, Clean Agile: Back to Basics, Prentice Hall, 2019



- is a facilitator
- measures progress against the backlog
- communicates with customers and management

K. Schwaber and M. Beedle, Agile Software Development with Scrum, Prentice Hall, 2001 S. Ockerman and S. Reindl, Mastering Professional Scrum: A Practitioners Guide to Overcoming Challenges and Maximizing the Benefits of Agility, Addison-Wesley, 2019



Many different approaches are possible

- Positive and negative sides in any of them
- Good or bad often depends on the context
- Small/large scale project, short/long lifetime etc.
- Process frameworks may (should) be adapted and extended
 - A good software process is **tailored to the project**
- Grain of salt
 - Commercial vs. scientific environment
 - ...let's not forget that often we are not only the developers, but also the customers!

How to improve the way we develop software?

Improvement requires **measurement**: quantify *before/after*



Helpful guidance towards adopting good practices

Capability Immaturity Model: A. Finkelstein, ACM SIGSOFT Software Eng. Notes 17 (4) 1992 pp 22–23 0 Negligent -1. Obstructive -2. Contemptuous -3. Undermining <u>https://dl.acm.org/doi/10.1145/141874.141878</u>

For singles

What if I work at a project where I am the only software developer?

The benefits of sound methodologies are not restricted to large scale projects or sizeable teams





+

https://doi.org/10.1184/R1/6585197.v1

PSP

A Self-Improvement Process for Software Engineers



Watts S. Humphrey











THE UNIFIED SOFTWARE

DEVELOPMENT

PROCESS

JAMES RUMBAUGH

The complete guide to the Unified Process from the original designers

Maria Grazia Pia, INFN Genova



Get a mentor!

Technology

No time for an extensive overview and in-depth analysis

A few highlights on key technologies and tools

ModellingTestingToolsDealing with legacy code



Basic techniques of programming hygiene *"agile software craftmanship"*

"If you have been a programmer for more than two or three years, you have probably been significantly slowed down by someone else's messy code." *R. C. Martin, Clean Code [Your own?]*

Clean Co	D de Craftsmanship	
	i di kana seria seria Interna seria s	iys ble id c
Foreword by James O. Coplien	Robert C. Martin	

Design rules

Keep configurable data at high levels

Prefer polymorphism to if/else or switch/ca

Separate multi-threading code

Prevent over-configurability

Use dependency injection

Follow Law of Demeter A class should know its direct dependencies

Understandability

Be consistent If you do something a certain do all similar things in the same way

Use explanatory variables

Encapsulate boundary conditions. Boundar conditions are hard to keep track of Put the processing for them in one place

Prefer dedicated value objects to primitive

Avoid logical dependency. Don't write met which works correctly depending on somet else in the same class. Maria Grazia Pia, I

Avoid	negative	conditionals	
	nobacito	oomancionaio	

Best practices for clean code

	Names	Use as explanation of intent	Small number of instance variables	
	Choose descriptive and unambiguous names	Use as clarification of code	Base class should know nothing about their derivatives	
better	Make meaningful distinction	Use as warning of consequences	Better to have many functions than to pass some code into a function to select a behavior	
leaner than	Use pronounceable names	Source code structure structure	Prefer non-static methods to static methods	
he root:	Use searchable names	Separate concepts vertically	Tests	
	Replace magic numbers with named constants	Related code should appear vertically dense	One assert per test	
	Avoid encodings. Don't append prefixes or type information	Declare variables close to their usage	Readable	
h/case	Functions	Dependent functions should be close	Fast	
	Small	Similar functions should be close	Independent	
	Do one thing	Place functions in the downward direction	Repeatable	
	Use descriptive names	Keep lines short	Code smells	
now only	Prefer fewer arguments	Don't use horizontal alignment	Rigidity. The software is difficult to change. A small change causes a cascade of subsequent changes	
	Have no side effects	Use white space to associate related things and disassociate weakly related	Fragility. The software breaks in many places due to a single change	
у	Don't use flag arguments. Split method into several independent methods that can be called from the client without the flag	Don't break indentation	International Series C. Martin Series S	
tain way,	Comments	Object and data structure	The Clean Coder	
	Always try to explain yourself in code	Hide internal structure	A Code of Conduct for Professional Programmers	
dary the	Don't be redundant	Prefer data structures	0	
ive type	Don't add obvious noise	Avoid hybrids structures (half object and half data)		
nethods mething	Don't use closing brace comments	Should be small	they want	
INFN Ge	enova			
	Don't comment out code Just remove	Do one thing	and a second	

Robert C. Martin

Forward by Hatthey Houses, Software Praces Arguabit,

[Visual] Modeling

Enormously helpful

- Think before typing implementation code on the keyboard
- Bird's-eye view of the software (and of the hardware)
- Visualize relationships among the players, which may not be easy to catch just by looking at source code

• UML (Unified Modelling Language)

- Standard, lingua franca for the communication of models
- Plenty of educational material and tools, from simple to very powerful
- Umbrello (<u>https://umbrello.kde.org/</u>) Sparx Enterprise Architect™
- Diagrams: class, package, component, deployment, object, state, activity, sequence, interaction, use case...

Design patterns

- elements of reusable object-oriented software

Maria Grazia Pia, INFN Genova



management of cross sections data for electron-photon interactions

> Atomic data libraries are used by all major Monte Carlo codes for particle transport



(from Compare libraries)





ArchiMate diagram of the computational environment of a software system for the verification test of atomic data libraries

How do you trust the software you use?



"I'm just doing what the other ones are doing"

"Testing shows the presence, not the absence, of bugs." In other words, a program can be proven incorrect by a test, but it cannot be proven correct. All that tests can do, after sufficient testing effort, is allow us to deem a program to be correct enough for our purposes.

				Regression testing		
lest	 Levels of testing Unit Integration 			Performance testin		
No time to				Stress testingConfiguration testingSecurity testing		
discipline	• Syst	System				
in depth!	<i>n depth!</i> • Acceptance					
Verification test Functiona		Functional	/non-functional testing			
Validation test		Black/white-box testing				
Test coverage		Test harness			Test automation	
Test cases		Test planning			Test frameworks	
🚺 🔞 unit						
pyte:	st _{va}	JUnit	6		googletestetc. Google C++ Testing Framework 35	

Test-driven development (TDD)



Facilitates regression testing

Discover problems early during the software development

Limited to unit testing,

still need system testing, performance, reliability testing etc.

D. Astels, Test Driven Development: A Practical Guide, Prentice Hall, 2003

conventional wisdom

"If it ain't broken, don't fix it"

A piece of software can be broken in many ways

Functional

it no longer delivers the function it is designed to perform

Maintenance

it can no longer be maintained

- Obsolete or no documentation
- Missing tests
- Original developers or users have left
- Inside knowledge about the system has disappeared
- Limited understanding of the entire system
- Too long to turn things over to **production**
- Too much time to make simple changes
- Need for constant bug fixes
- Big build times
- Difficulties separating products
- Duplicated code
- Code smells

Warnings you are heading into trouble

usually do not occur isolated



Take a loan on your software \Rightarrow pay back via reengineering

Investment for the future

 \Rightarrow paid back during maintenance





Refactoring is a disciplined technique for improving the design of an existing code

In the ideal world there would be hardly any need for refactoring In the real world most software needs to be refactored By learning refactoring you also learn writing code that minimizes the need to be refactored



Further learning







Technology Methodology



Master them, so that you can choose what is most appropriate to your research problem

This 40' introduction only scratches the surface, more to follow next year

Feel free to contact me for further information and for suggestions for next year