30th NENE 2021
September 6 - 9
BLED-SLOVENIA
International Conference Nuclear Energy for New Europe

Upgrading of HPCC at Reactor Center JSI in the Last Decade

Sandi Cimerman, Bojan Žefran
"Jozef Stefan" Institute
Jamova cesta 39
SI-1000 Ljubljana
sandi.cimerman@ijs.si,bojan.zefran@ijs.si

ABSTRACT

There has been great progress in the development of computers and computational science in recent history. For the last four decades or so this history has been also reflecting in computational history of the Reactor Center in Podgorica. Especially for the last decade, four compute clusters that have been installed on premise, have established RCP as one of the few HPC centers in Slovenia.

HPC clusters become essential in the nuclear energy science, as they are used in the fields of thermal hydraulics, material science, structural integrity and neutronics. The Reactor Engineering Department (R4) and the Reactor Physics Department (F8), have jointly used their HPCC's for fourteen years. Both of the departments have, with the support of the Slovenian Research Agency (ARRS), heavily invested in their HPCC systems during this time. HPCs represent critical equipment for their research activities. They improve research quality and help increase nuclear safety [1]. Students and young researchers also regularly use them in education.

1 INTRODUCTION: BRIEF NUCLEAR COMPUTATIONAL HISTORY

Computational science is a discipline concerned with the design, implementation and use of mathematical models to analyse and solve scientific problems. Analog computers have existed for centuries, but technological advancement modern era of the late 19 century has made them widely available in science. Electro-mechanical computers that has been developed in late 1930s, has paved the way for vacuum tubes and digital electronics computers, that have replaced them in the early '40s.

Nuclear and computing were intertwined from the beginning. »Electronic Numerical Integrator and Computer«, or more commonly known as ENIAC, was the first programmable, electronic, general-purpose digital computer. It was commissioned in December 1945, initially for calculating artillery firing tables for the United States Army, but it was soon used by Edward Teller, who used the machine for calculating nuclear reactions.

First computers provided an alternative to the expensive and complex or simply impossible physical experiments. Initially as the supplement to the experiment, but with time became a viable replacement.

In the 1950s and 1960s, computers were rare and only available to the American government and research agencies. The first working programmable, fully automatic digital computer Zuse was built in 1941, but due to missing conditional branching it was not fully Turing-complete. Jožef Stefan Institute purchased one of the Zuse Z3 computers in 1962, and nuclear physicists were the first ones to use it for solving the neutron transport equations.

The CDC 6600, developed by »father of supercomputing« Seymour Cray, was released in 1964 by a company named Control Data Corporation, and it is considered the first (successful) supercomputer. In 1976, four years after leaving CDC, Cray delivered Cray-1, the most successful supercomputer in history. In 1985 Cray-2, a supercomputer with 4 processors, was the first supercomputer that breached one gigaflop (a measure of computer performance) limit.

For simple calculations the commodity computers could be eventually used, something that was impossible before without rare, large and expensive computers. That progressed pushed that even further, when first Beowulf supercomputer was built in 1994 by NASA. They merged commodity and of the shelf systems, connected together with a high performance local area Ethernet network. The cluster ran on top of one of the newest operating system – Linux, the operating system that was first written less than three years before that.

2 HISTORY OF POWERFUL COMPUTING AT THE REACTOR CENTER IN PODGORICA

Since the late 70s there have been multiple powerful computers at the RCP. In the late 70s there has been CDC Cyber 172 [2], a 16-bit minicomputer, first manufactured by Control Data Corporation in 1966. It was used for simulations of neutron scattering.

After CDC, there were several powerful computers used in the Reactor Center in the following years. Programmed Data Processor - PDP 11/34, first manufactured by Digital Equipment Corporation (DEC) in 1970. Then several VAX machines, VAX-750 (32-bit minicomputer, first produced in 1977), MicroVAX2 (32-bit, first introduced in 1983), and VAX-8650 (first manufactured in 1985) [3].

In the mid 80s the researchers of RCP were working on rewriting the entire programs (like RELAP5/MOD1, COBRA-IV and FRAPCON-2) from CDC into VAX environment. Like to GPGPUs of the modern era, VAX had also computing vector coprocessors - CSPI MAP6420. Researchers on RCP have rewritten some of the test programs, and ran them on them coprocessor.

In 1998, 23 years before commissioning of the first Slovenian national supercomputer Vega (today ranked 106 on top 500), there was a prototype mini compute cluster by the same name on the F8 Reactor physics. It was Beowulf type of compute cluster, four individual single processor PC's, running on top of Red Hat Linux, connected over 10/100 ethernet network. Vega was a testing predecessor of later compute clusters purchased by R4 and F8. It was successfully used as parallel multi-node computation cluster for nuclear program calculations (MCNP, Tracos and Triga) [4].

3 EXAMPLES OF MODERN HPC USE IN THE NUCLEAR FIELD

3.1 At the Reactor Engineering Department (R4)

Department is involved in a great number of research areas, and a large fraction of them are using on-premise HPC systems. In the field of thermal hydraulic safety analyses

- for the validation and uncertainty of system and CFD code,
- counter-current gas-liquid flow modelling using interface tracking methods,
- turbulent flows DNS, LES (Water, Liquid Sodium),
- fast transients in two-phase flows (Water hammer),
- two-phase flows with two-fluid models and interface tracking schemes,
- transitional two phase flows,
- modelling of flows in complex geometries (fuel bundles, multiple jets

- subcooled flow boiling and critical heat flux),
- CFD analysis of conjugate heat transfer and design assessment of a divertor cooling finger for ITER.

In the fields of severe accidents

- interaction between containment safety systems and atmosphere, including management of hydrogen,
- fuel-coolant interactions (water, liquid sodium).

For the integrity and ageing of

- multiscale models of intergranular/intragranular damage [5],
- models of cracks in complex conditions, including dissimilar welds,

Multiphysics modelling computational research of

- models of thermal fatigue (computational fluid dynamics and structural mechanics),
- coupling of neutronics and thermal hydraulics.

Researchers of R4 are using programs and tools developed within the group, along with a number of reliable and modern commercial programs, and many freely available and open software. The tools used for solving CFD problems are Ansys CFX, Ansys Fluent, OpenFOAM and IcemCFD. Modeling of structural mechanics (FEA/FEM) is run on top of SIMULIA ABAQUS and Neper. For nuclear engineering modeling specialized tools are used, like

- MC3D, code for simulation of fuel-coolant interactions (by IRSN, France),
- MELCOR, for simulating severe nuclear power plant accidents (by Sandia National Laboratories),
- ASTEC system codes for severe accidents,
- WAHA for simulations of 1D water hammer in two-phase systems,
- RELAP5 for thermal hydraulics of nuclear power plants,
- TRACE for deterministic safety analysis, and
- TRIPOLI-4 for general purpose radiation transport code (using Monte Carlo method) to simulate neutron and photon behaviour in three-dimensional geometries.

3.2 At the the Reactor Physics Department (F8)

Main parallel applications are multi-purpose three-dimensional continuous-energy Monte Carlo particle transport calculations. Codes used for HPC computing are

- MCNP - developed by Los Alamos National Laboratory [6]
- Serpent - developed at VTT Technical Research Centre of Finland, Ltd. [7]

They are used for

- reactor physics applications
- criticality calculations
- fuel cycle studies
- research reactor modelling
- validation of deterministic transport codes
- neutron and photon transport simulations for radiation dose rate calculations, shielding, fusion research and medical physics

For the calculations are used evaluated nuclear data libraries for nuclear science and technology, for instance ENDF/B (VIII.0, VII.1, VII.0), JEFF (3.3, 3.2), FENDL, IRDFF II and TENDL.

4 WHAT DOES IT MEAN TO IMPLEMENT HPC?

High performance compute cluster is an assembly of software, hardware and the infrastructure that is housing the hardware.

The **hardware** of compute clusters on RCP all consists of one or more management servers, high capacity high reliability permanent storage pool (sometimes storage server), high performance/high bandwidth/low latency network, and independent multiple core/multiple sockets servers with high memory capacity.

The **software** of compute clusters is stacked on top of enterprise grade operating system. The compiled software on the system is either open source, closed source (some with limited access) or commercial software. Special care is needed when granting access to code and data, as we are bound to follow the regulations and limitations of software access. Some of the programs, like MCNP, are limited per users within the organizational unit. Licensing access, for the most of the commercial software, allows the usage of programs to the specific organization, for the educational or research purposes only.

HPC infrastructure has evolved substantially in the last decade. Purchase of Krn in 2010 and following upgrades demanded the expansion of the server room. We increased the floor space by demolishing wall and start using entire basement room. Today, a new building with the necessary infrastructure is in the planning.

The output of the compute clusters is data and heat. The clusters could consume up to 35 kW power in 2020. All that energy converts directly into heat. Lifetime of the servers depends on constant maintaining of »clean« electrical power (without short term electrical outages and without voltage or frequency variations), providing clean and vibrant free server room, and having a powerful and reliable cooling system.

Figure 1: 80 kW modular UPS (left), (middle) 30 kW in-row cooling system and Skuta first two purchases (two racks on the right).

Today, HPC part of the RCP server room, has 4 independent cooling systems. Two backup support split AC systems, and two main cooling systems. A new 30 kW split system was added in 2019, next to the reliable and cost effective 20 kW water-cooled cabinet, that has been installed in 2010.

For the power redundancy, there are two transformers installed on the location, providing mains electrical power for the entire reactor centre. From there, two separate power lines connect the transformer site and the HPC server room, and so provide the option of having electrical power in case of major maintenance on one of the main buildings. All the main HPC hardware, the network switches, storage pools and the management servers have a redundant (dual) power supply, in case one of the PSU units fails.

In event of mains power outage, a rack of 80 kW modular uninterruptable power supply (UPS), provide electricity for the entire server room from the cabinet of lead batteries. Together they can keep all clusters and support hardware running at peak power consumption (up to 40 kW) for several minutes. During that time, the RCP generator has the capability of automatically starting, and providing the power for the entire reactor for several hours, while completely unmaintained.

The facilities also consists of advanced fire alarm, that can trigger fire suppression system by releasing auto-extinguishing gas, that doesn't damage the electronics. Remote night-vision camera provides additional safety of the server room. A dedicated temperature sensing system provides an automatic alerting system via a cellular network in case of emergency (overheating, leakage warning, cooling water loss).

Figure 2: Flask of automatic fire protection system (left), new main electrical distribution board (middle), and (right) half of the modular UPS system – rack of lead batteries providing several minutes of power for entire cluster.

Server room had been last upgraded in 2019, when the entire electrical wiring have been redone, and the new redundant secondary mains power line has been connected to the server room. In same year, a large modular UPS system has replaced older ones, and a new In-Row cooling unit has been added.

5 REVIEW OF GROWTH OF R4 AND F8 ACQUISITIONS FROM MANGRT TO SKUTA

During the last 15 years, purchases and updates of multiple compute clusters were carried out. Systems, named after the highest Slovenian mountains, are (listed by the order of purchase) Mangrt, Krn, Razor and Skuta.

5.1 Mangrt

In 2007, departments R4 and F8 jointly bought the Mangrt computing cluster. The cluster contained all the elements of modern and all the following compute clusters. The cluster was stacked in a single cabinet connected via dedicated gigabit network switch. The servers had motherboards with dual-socket dual-core processors, and when upgraded the following year they increased to the quad-core. The cluster contained a total of 118 physical cores. The system has been shut down in 2019.

5.2 Krn

In 2010, departments took into account the shortcomings of the Mangrt cluster, and with a new tender bought calculation cluster Krn. State-of-the-art 40 gigabit communications network called Infiniband was connecting 11 compute nodes. They were also based on dual Intel Xeon X56xx x86 processors, each containing 6 compute cores. In the following year, the cluster was upgraded with 22 new nodes and a year after that with 15 more. The last upgrade added additional Nvidia Tesla general-purpose graphics processing units (GPGPU). That made Krn one of the first Slovenian compute clusters with usable GPGPUs, and at that moment it consisted of 49 computing nodes with a total of 600 physical cores. The system is today still partially operating.

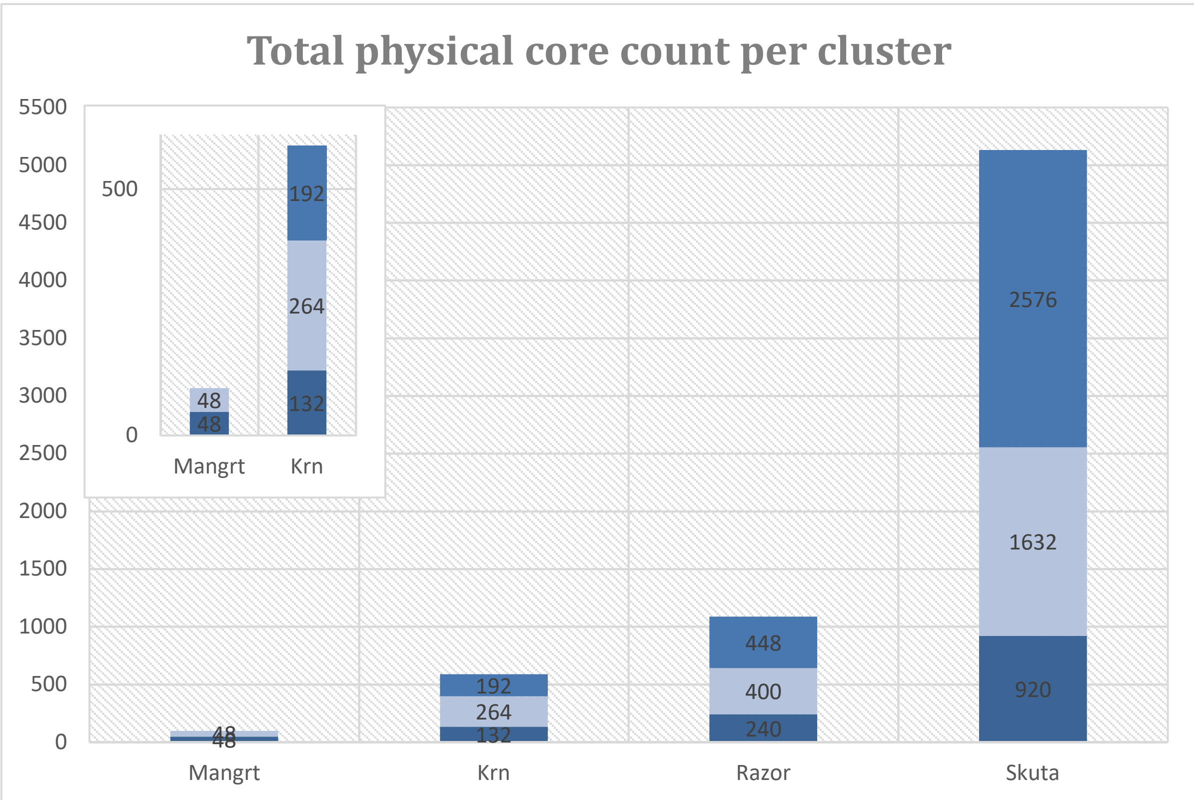

Figure 3: Growth of the RCP HPC systems between 2007 and 2021 – total number of physical processors per cluster on y-axis, and physical core count per upgrade in the individual sub segment. Cluster acquisition and subsequent upgrades are marked with different shades of blue.

5.3 Razor

With accumulated HPC and public tender purchase knowledge, departments made a new EU tender, and purchased Razor in 2014. The initial configuration was upgraded in just after 4 months, almost tripling the total node count. In the following year, the second upgrade was carried out - it was also the last one for Razor. The entire cluster is connected via Infiniband FDR. Computing nodes are dual-socket Intel x86 Xeon E5 26xx, with each processor having between 10 and 14 physical computing cores. After 2015, 50 computing nodes has a total of 1136 physical cores. It has also a separate 90 TB storage file server.

5.4 Skuta

In 2019 again both departments jointly invested in another, even larger, EU procurement tender. It enabled the purchase of Skuta compute cluster. 25 nodes were connected with a new type of a communications network, Omni-Path. The compute nodes are two socket Intel x86 Gold processors, each processor contains 20 physical computing cores and 192 GB of memory. In 2020 the cluster was upgraded for the first time, with 28 new nodes, each with two 24-physical core processors. At that point, cluster had 56 computing nodes with a total of 2584 physical compute cores.

Currently, we are in on-going process of the second upgrade of Skuta. This is the last upgrade of this cluster, and the largest upgrade in RCP HPC history. It will add additional 45 dual socket nodes with 28-core Intel processors. By the end of 2021, Skuta will have 104 computing nodes with a total of 5048 physical cores and a dedicated fast scratch file server. The entire cluster is interconnected over 100 Gbps low latency Omni-Path network.

Figure 4: Last (still on-going) upgrade of cluster Skuta (two racks on the left), cluster Razor (two racks in the middle-right) and (far right) part of the cluster Krn.

6 FUTURE AND PERSPECTIVE

The well-known Moore's law, states that the number of transistors doubles about every two years. The CPU power dissipation limits clock rates and instruction-level parallelism of single core processors, and also single threaded code execution. Therefore, modern processors contain multi-core chip designs, and the software must be written in a multi-threaded form in order to use it in entirety. The future will most likely continue to use muti-node multi-socket multi-core systems.

The increasing computational power can hardly follow the computational need. There are several reasons for increase of demand, and just some of most common are: increase of available programs, higher software accuracy and probably most important one - ease of software usage. The growing need for higher precision of calculations is also obtainable only by providing a sufficient amount of available compute resources. For the last few years, we notice that this trend is only increasing.

Today there is also a wide option of multiple hardware acceleration available, like FPGAs or GPGPUs. Currently, their usage is either limited for the rare types of compute subroutines, or the complex implementation of the hardware and software outweigh the benefits, therefore the conventional CPUs are still preferred. Since the future might bring a solution to either of those, technology remains to be of great interest to us.

On the software side, the future is even more encouraging. Recently there were new provisioning solutions published, providing instructions and software stack for deploying the entire cluster. Future promises more simple deployment for even more programs and environments as well as entire operating environments. This will greatly simplify the deployment of new clusters, or refurbishing the old ones. Lightweight virtualization has been providing a new way of giving the users more freedom in software deployment, while seamlessly using all the underlying hardware resources. All these solutions provide a better end-user experience, but at the price of adding to the complexity to the systems, and increasing level of system administration.

HPC systems will continue to be used for the near future by both of the departments as well as the rest of the nuclear sector.

ACKNOWLEDGMENTS

This project has been partially supported by The Slovenian Research Agency (ARRS).

REFERENCES

[1] I. Simonovski, S. Clements, B. Žefran, S. Cimerman, The role of High Performance Computing in the Nuclear Energy Sector, Proc. Int. Conf. Conference Nuclear Energy for New Europe 2015, Portorož, Slovenia, September 14-17, Nuclear Society of Slovenia, 2015

[2] M. Zeljko, M. Gregorič. G. Peterlin, Analiza delovanja uparjalnika jedrske elektrarne s programom RELAP5, IJS 1983

[3] Poročilo o uporabi jedrske energije v letu 1988, M. Gregorič M.F. Levstek M. Trampuž, RUJV 1989

[4] Using Vega Linux Cluster at Reactor Physics Department, Bojan Žefran, Robert Jeraj, Jure Skvarč, Bogdan Glumac, Proc. Int. Conf. Nuclear Energy in Central Europe 1999, Portorož, Slovenia, September 6-9, Nuclear Society of Slovenia, 1999

[5] I. Simonovski, L. Cizelj, Grain-Scale Modeling Approaches for Polycrystalline Aggregates, Polycrystalline materials - theoretical and practical aspects, pp 49-74 (2011)

[6] T. Goorley, et al., "Initial MCNP6 Release Overview", Nuclear Technology, 180, pp 298-315 (Dec 2012)

[7] Leppänen, J., et al. (2015) "The Serpent Monte Carlo code: Status, development and applications in 2013", Ann. Nucl. Energy, 82 (2015) 142-150.