Handbook for Getting Started with HPC

For industry and academia

Computer-based simulations have become standard tools for industry and academia. While the power of standard PCs or workstations with multiprocessors/cores is tremendous nowadays, they still often don't fulfill the requirements for raw computational speed, storage and main memory that repeated detailed simulations may pose. Here, **High-Performance Computing (HPC) can come to the rescue.**

Objectives of this Handbook

HPC Competency Centres operate large installations of HPC systems and have acquired considerable expertise about HPC Technologies. The centres of the SESAME network *offer services ranging from compute time over HPC software development to HPC technology consulting.*

This handbook will explain the ***economic benefits*** of using HPC for people at managerial positions and will provide a path for attaining this benefits through collaboration with the HPC Competency Centres in SESAME Net. Apart from the purely economic reasons we consider the ***improvements in innovation and marketing*** and the competitive advantages stemming from the use of HPC to provide justification for taking the managerial decision to explore the use of HPC.

Furthermore, the technical aspects of HPC are explained in a simplified way, oriented towards newcomers in this field, offering a first introduction of common technical terms and formal aspects in HPC as well as a short *practical examples of how to connect to and start an application* on one of the HPC resources belonging to a centre of the SESAME Net.

At the end, the reader will find some representative industrial use cases.

What's NOT in this handbook?

The guide will not cover the software development of applications on HPC systems. You can find more in-depth information about programming for HPC here:

- https://www.mcs.anl.gov/~itf/dbpp/text/book.html, Designing and building parallel programs,
- http://www-users.cs.umn.edu/~karypis/parbook/, Introduction to Parallel Computing,
- https://computing.llnl.gov/tutorials/parallel_comp/, WWW – parallel computing: Blaise Barney.

Recent developments like modern accelerator technologies, for example based on graphics processing units or FPGAs, have been largely omitted for the sake of clarity.

What is HPC?

The term HPC is occasionally used as a synonym for supercomputing, although usually *only the fastest HPC systems currently operated world-wide are referred to as Supercomputers*.

The term High Performance Computing (HPC) usually refers to the application of state-of-the-art computing systems for efficiently running advanced applications such as computer simulations, often for engineering or natural sciences purposes. While in the past the term was mostly used for the work with very specialized and often highly expensive machines, the evolution of computer and server systems has led to a commoditization of HPC systems and to a mainstreaming of its concepts and technologies. Nowadays desktop computers and workstations potentially offer such high capacities, while at the same time HPC systems have converged so much with off-the-shelf hardware, that the term also includes achieving the most performance on mainstream systems.

The term HPC is used for nearly all aspects of managing and using high-performance systems and the software running on them. Major topics in the area include activities like:

- **Performance** assessment of existing applications,
- **Optimization** of applications for high performance,
- Specialized software **development** for HPC systems, and of course
- **Actual execution** of applications, with eventual automation of workflows or interactive visualization of large result sets.

Why use HPC?

The vast increase of computational power in the last decades has created exciting opportunities. High-Performance Computing (HPC) has become **indispensable tool for industry and academia** to innovate in such domains as Computer aided engineering, simulations, renewable energy, financial services, satellite, Earth observation, advanced image analysis, data science and precision agriculture. But it does not stop there: **supercomputing can help your business too.** Supercomputing can help spark your business innovations. It can improve design turnaround time for new products, reduce time to market and increase your overall competitiveness.

'97% of companies that have adopted the technology said they could not compete or survive without it' (IDC, 2014).

The most obvious motivation for using HPC is, as mentioned before, a lack of computing power and storage on your enterprise infrastructure to run simulations or other demanding computer applications. Broadly speaking, HPC technology may be useful in all cases where large amounts of computations are necessary or large amounts of data need to be processed. And even when compute infrastructure can easily be extended, it will be difficult to reach significant performance capacity for your applications such as provided by HPC centres.

However, most of the advantages of HPC are related to the kinds of computations that cannot be accomplished on regular infrastructure. Due to the advancements in applied sciences, accurate simulations and modelling of new products and process increasingly require HPC-type equipment that is usually not available on-site, **even in medium or large scale enterprise.**

The size of the **HPC market** is projected to grow from USD 28.08 Billion in 2015 to USD 36.62 Billion by 2020, at a Compound Annual Growth Rate (CAGR)[1] of 5.45%. The use of HPC technologies is a testament of the innovative character of the SMEs and brings substantial Return On Investment (ROI). Various studies have been performed on the ROI for HPC. One pilot study from IDC indicates that one can expect $356.5 on average in revenue per dollar of HPC invested and $38.7 on average of profits (or cost savings) per dollar of HPC invested[2]. The European project EESI concluded that HPC has clearly added great economic value across Europe and the ROI results in an average of $867 revenue dollars generated for each dollar invested in HPC in this study and $69 dollars in profits (or in cost savings) for each dollar invested in HPC[3]. Obviously, the actual results will depend on the type of industry and type of use of HPC technology, but the potential is obviously interesting. It is important to understand that there are synergetic effects for a country's economy from increased use of HPC.

Another reason to investigate HPC could be a pressure to realize shorter development cycles, e.g., because of competition on the market. Typically, new products are nowadays developed with a combination of physical experiments and simulations. Having large computing power for large and complex simulations opens possibilities for new ideas, innovation and new products in pharmacy, production, automotive components, etc. By affording to run multiple simulations with different input parameters, HPC systems for example offer the possibility to adapt the simulation closer to the experiments or to optimize products through simulations.

Large industrial enterprises, e.g. from the automotive industry, are said to operate their own HPC clusters since a decade. HPC centres like the ones collaborating in SESAME Net also enable small and medium enterprises to access such powerful resources.

HPC can add ***tremendous value*** to any SME that develops products:

- Creation of **large, high-fidelity models** that yield accurate and detailed insight into a design's performance.
- **Extremely accurate simulations** to predict real-world conditions to convince a customer of your product.
- In engineering HPC enables **high mesh densities** for improved accuracy, numerous geometric details, or sophisticated treatment of physical phenomena.

[1] http://www.marketsandmarkets.com/Market-Reports/Quantum-High-Performance-Computing-Market-631.html

[2] https://hpcuserforum.com/roi.html

[3] http://www.eesi-project.eu/wp-content/uploads/2015/05/EESI2_D7.4_Final-report-on-HPC-Return-on-Investment.pdf

Aspects and benefits
of using HPC technology,

specifically by SMEs

Many studies establish how the use of HPC **opens new markets** or allows for **expansion of the existing markets**, through improvements in innovation and enabling of new technologies and products.

- Economic reasons
- Innovative reasons
- Marketing reasons
- Competitive advantage reasons

Since all the decisions taken in an enterprise are intended to have a **measurable impact in their economic balance**, in order to maximize profit: marketing, research, innovation reasons should at last have an economic impact in the enterprise balance, but in this section we are going to focus in the economic reasons (the ones related more directly to economic balance) to integrate HPC as a key aspect in the enterprise value chain.

Economic reasons:

explaining how HPC will reduce costs and/or improve benefits.

For example, reducing designing and prototyping costs; detecting design errors in early stages of product development. It's important to always focus on the CEO-like argument, so in detecting design errors earlier it should be mentioned the positive economic impacts that this would have: early detected design errors are cheaper to overcome. Another example, if prototyping has a major role in the SME's chain value then HPC saves the costs of building many expensive physical prototypes which now can be built virtually, minimizing the number of physical prototypes to be built.

The three key economic reasons to consider to integrate HPC in SMEs (and also in big enterprises) are:

- **Cost reduction**: Prototyping (wind tunnel, physical testing) and design costs;
- Reducing design times, reducing time to market, (translating this **time reduction** in money);
- **New product/service** development that is not possible to develop without HPC.

Each of these aspects will be analyzed and introduced with more detail.

Cost reduction

HPC can help the SME to save costs, for example in **prototyping and design procedures**, as well as in **operational costs**. Physical prototyping is to create a physical prototype ("An early sample, model, or release of a product build to test a concept or process or to act as a thing to be replicated or learned from"[4]). This process is very time and cost consuming, because of the time needed to produce an item outside of an assembly line and the cost of prime matters, and this only gets worse when multiple prototypes are needed, especially in case of destructive physical testing when prototypes can only be used once. Physical prototypes can be substituted by virtual prototypes that are cheaper to develop and can be used again and again without destroying them.

Virtual prototypes once developed can be used to simulate the physical behavior or interactions under many circumstances: aero and hydrodynamics, substituting the need of complexity and expensiveness of wind tunnels or fluid tests, useful in: Aerospatiale, automotive, naval and constructive; Physical tests, virtual prototypes can be used to perform multiple virtual test using simulation, no matter if these tests simulate physical damage or full destruction of the prototype.

Design costs can be reduced using HPC, since HPC brings us the power to design prototypes of products without the need of any physical prototype and begin to get knowledge about it in the very early stages of design. This also improves the design process allowing detecting design errors or misconceptions in these early stages of design, where they can be corrected easily and cheaper than in later stages. Operational costs of the SMEs can also be saved by accurate and extensive simulation, modelling and optimization of certain kind of operations. For example, HPC can be used in optimization of public transport or logistics. It is widely used in the search of oil and gas or in energy production problems.

Reducing design times

In some sectors like aeronautics, engineering or energy, where the interaction of fluids with products like: aircrafts, shipbuilding, turbines, wind or offshore energy generators is a key factor of efficiency, HPC can be a key to success because it can **reduce design times drastically**.

The times needed to construct a virtual prototype and discover the full interactions with their surrounding or inner contained fluids are much shorter

[4] Blackwell, A. H.; Manar, E., eds. (2015). "Prototype". UXL Encyclopedia of Science (3rd ed.).

than the time needed to get this knowledge from physical ones. This is an uncontestable advantage that can make the difference between a successful and a ruinous project, as well as enable a project that otherwise is never started because HPC possibilities were not on mind or simply were unknown when tight schedules were present and standard design times are not possible.

As an example of the prior two reasons, under the umbrella of the EU funded project **FORTISSIMO**, the Italian SME EnginSoft reduced the time needed for full design and optimization of a single pump turbine **from 2-3 years to 6 months** and this fact allowed to recover the full SME investment in HPC Cloud based simulation and external expertise in less than six months[5].

Products impossible to produce without HPC

Another example of economic benefit from HPC is the possibility to **produce or study products that with standard technology limitations are impossible or too costly to design and produce**.

The study of hazardous substances (explosive, toxic or mutagenic) has heavy limitations due to their inherent properties and security measures needed to manipulate them. The physical measures needed to overcome these limitations are prohibitive in terms of cost, so the physical experimentation with these substances will have a bad ROI in most cases and involve high risks in all of them. HPC can overcome these limitations and simulate the behavior of these substances with virtual models getting knowledge about them in a non-risky environment.

Limiting the number of physical tests only to experiments with high percentages of success

Another FORTISSIMO experiment shows how a Swiss SME, Lonza from the chemical sector, suppliers to the Pharma&Biotech and specially ingredients market, explored with success the possibility to determinate physical properties of compounds with desired precision using HPC. The use of HPC clusters reduced the amount of time needed to perform calculations **from a six-months-long calculation in a 16-core cluster, to less than one day on an HPC system.**

[5] HPC-Cloud-based design of centrifugal pumps: http://www.fortissimo-project.eu/experiments/418

Economic reasons to use HPC Center services:

- Reducing costs using a Cloud-based HPC system on a pay-per-use basis at approximately half the price of owning and maintaining a sufficiently powerful in-house system.
- Running big processes that need very powerful HPC systems not affordable for SMEs.

The SESAME Net members possess **state-of-the-art HPC facilities** and even more-importantly, **substantial expertise** in development and use of HPC applications for solving diverse practical problems. The use of resources from SESAME Net provides **opportunity for the SMEs to evaluate new technology at low financial risk**, so that they can build the business case for larger investment in HPC. They will have the **flexibility to buy application services** instead of licenses and computing systems.

Some SMEs that are the business of provision of software and services may also expand significantly their portfolio if they incorporate the use of HPC technology as part of their offering. By co-designing with SESAME Net partners, SMEs can address new business cases and markets.

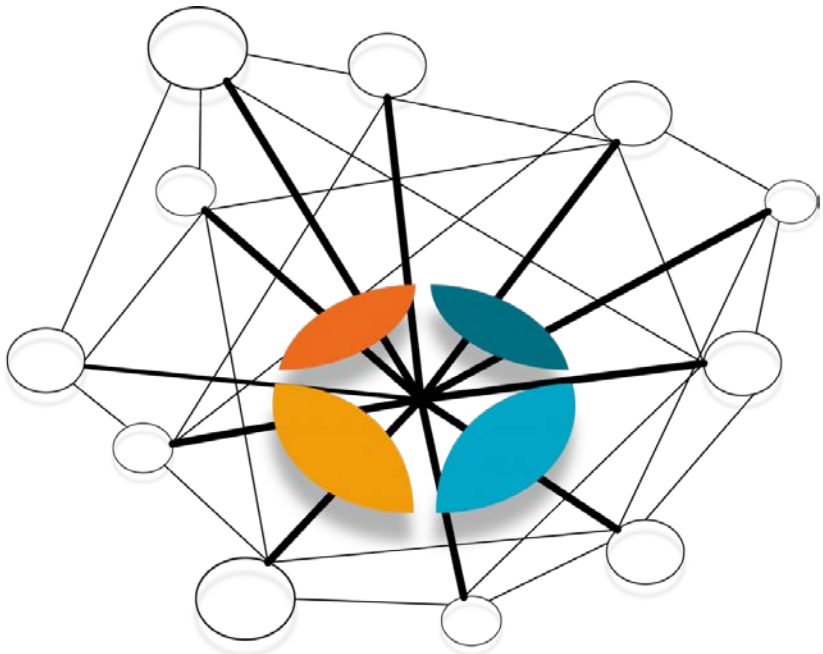

Most of the advantages of using HPC technology can be obtained by SMEs that have substantial Research & Development department, since the access to HPC resources can improve the possibilities to develop innovative products and services. By collaborating with an established HPC Competency Centre, the **SME will not necessitate costly on-premise infrastructure**. The costs of owning and operating such infrastructure are a heavy burden for most SMEs. By using flexible contracts SMEs can have a low-risk, low-effort, pay-as-you go access to HPC and leverage state-of-the-art software and tools to speed-up their development process.

Innovative reasons:

explaining how HPC will improve possibilities to develop innovative products and services.

Apart from the cost-related reasons, there are some main advantages of using HPC:

- Obtaining results much **faster** by employing high amount of compute power;
- Achieving much more **accurate** representation of the product and its working environment;
- Exploring much **wider search space** for the parameters of the envisaged product.

By collaborating with a leading HPC Competency Centre, the personnel at the SME's Research & Development department acquire **precious expertise** that is also useful in other, more standard situations. For example, the increased use of multi-core devices in all kinds of products makes knowledge of parallel computing paradigms an important asset. The need to achieve a given results with as little use of energy as possible stresses the importance of using optimal algorithms and scalable programming approaches.

Overall, the use of HPC for completing simulations that would normally run for days or weeks on a desktop workstation, in a matter of hours enables much faster time-to-market and thus first-mover advantage for the SME.

It should not be underestimated that being in close contact with the academia and HPC experts in industry, the SME may have advanced knowledge of new developments in HPC technology or applied sciences that will enable cross-fertilization and rapid incorporate of the new technologies in their product portfolio.

Best product build with best technology:

In our society, most of the products get an improved perception by clients when they are built with the latest technology. Even if the use of this technology only involves a very low percentage of improvement, there are clients that always want to own the "state-of-art" product with the last improvements, while others are more concerned about environment and want to have "green" products, with more efficiency and low power consumption. HPC technology can deliver all these properties, since it is an innovative technology with relatively low use by SMEs because it is not generalized yet. Innovative technologies can lead to innovative products, and for SMEs willing to improve their products it is a competitive advantage. HPC can be applied successfully to improve products in order to make them more "green". Under the umbrella of FORTISSIMO there are some examples of products optimized using HPC that gained special characteristics that can be valuable for some products or differentiate them from the product´s competence. In the aerodynamics sector, Pipistrel, a Slovenian SME, used HPC to improve the design of their aircraft and KEW optimize the aircraft wiring design not only to reduce costs saving prime matters but less prime matters is less weight, and less weight is less fuel consumption and over the 30 to 50 years of an aircraft life this has at last economic relevance. In the automotive sector AVL used HPC in the cloud to perform studies about CO2 emissions reduction. In the civil construction sector, the Scottish SME IES performs building energy efficiency studies but making the jump to HPC let them make scale their studies from buildings to cities.

An important marketing reason for the use of HPC is to present the enterprise as a cutting-edge technological firm. This has a positive effect in its market position and on the perception of its products and services by customers. And of course, it results in the subsequent positive economic impacts, e.g. more sales, price premiums, and so forth.

Marketing reasons:

present your SME as a cutting-edge technological firm.

More specifically, the marketing reasons to use HPC can be listed as:

Leading brands partnerships are a key:

When selling a product and you want to reach a market where a "brand" or "organization" is well known and with good reputation, it is a good idea to associate it with this leading "brand" or "organization". So your product gets an improved perception over other products that do not own this advantage. SESAME Net is composed of very known and well considered partners, so the collaboration with SESAME HPC partners can lead to a win-win situation, where the new product is associated with a well-known brand and on the other hand SESAME Net gets an improved visibility.

Be a leading brand in their sector:

HPC can be the difference between one leading brand and their competitors that don't use this technology. To associate your brand with innovation and differentiate against competitors' products, the use of HPC for product improvement is a good point that highlighted correctly can increase product brand´s sells. Final customers always want to buy best cost-value products and include HPC in the chain value of the SME can be a technology to achieve these added value products. This consideration is even more visible in provisioning of services, where the use of powerful HPC equipment can lead to increased consumer trust.

In all those cases it is important to be able to substantiate the marketing claims with scientific and technical arguments.

"High Performance Computing (HPC) plus data science allows public and private organizations get actionable, valuable **intelligence from massive volumes of data** and use **predictive and prescriptive analytics** to make better decisions and create game-changing strategies. The integration of computing resources, software, networking, data storage, information management, and data scientists using machine learning and algorithms is the secret sauce to achieving the fundamental goal of creating **durable competitive advantage**."[5]

Competitive advantage reasons:

Technological capable firms have inherent advantages compared to other ones.

Being able to bring new products to the market, thanks to the HPC edge on research and development process, allows the SME to gain a competitive advantage over their competitors. This *first-mover advantage*, thanks to the technological leadership, has several positive impacts in the firm. For example brand recognition, this is the automatic association in costumers' mind of the new type of product to its initial manufacturer, e.g. the copying machines and Xerox, the tissue paper and Kleenex. Another example of first-mover advantage is the preemption of scarce assets, which allows the first-mover firm which has superior information to purchase assets at market prices below those that will prevail later in the evolution of the market.[6]

Additionally, HPC capable firms are able to **reduce the *product cycle time***, this is the period for a manufacturer to complete development and production of a new, or modified, product. It does not only reduce costs, but an **improvement in the *time to market* response** allows for a bigger maneuverability to adapt to the changing conditions of modern markets.

The use of HPC enables the design of a more accomplished product, with less defects or covering wider specter of consumer needs.

In the today's economy it is increasingly more important to find the consumer niche and to engage the potential clients. With HPC technology the SME can employ advanced algorithms and processes that will increase their market share and reach wider consumer audience.

[5] http://www.datasciencecentral.com/profiles/blogs/high-performance-computing-data-science-competitive-advantage

[6] https://en.wikipedia.org/wiki/First-mover_advantage

By establishing collaboration, possibly through a pilot project, not only with HPC resource providers, but also with Independent Software Vendors and research software providers, application experts and consultants, the SME is immersed in a network with high innovation potential and thus will have increased flexibility and high potential for sustainable development of innovative products.

To gain an overview of exemplary problems that have been successfully solved using HPC technology, it might be useful to browse through the collection of use cases documented at the SESAME Net website[7]. If you are working in similar fields or face similar challenges, chances are high that HPC could provide real benefits for your business.

Benefits of accessing HPC through **SESAME Net**

The HPC centres in the SESAME Net network offer consulting services, which includes individual **advice on specific scenarios** and help to find out if the utilization of HPC might be beneficial in your situation. That is why SMEs are encouraged to contact an HPC competency centre or to make use of the resources the SESAME Net partners have made available to take an informed decision as to the usefulness of HPC for a particular problem.

One particular option to get in touch is to fill the online survey at the SESAME Net website[7]. By answering a couple of basic questions, the SME provides enough information so that SESAME Net partners may give a preliminary evaluation whether the SME can benefit from the use of HPC (see Figure 1). This survey is available in 12 languages.

Figure 1: SESAME Net online survey for SMEs on the official webpage.

[7] https://sesamenet.eu/survey/index.php/265291/link

A second option is to directly contact a HPC competence centre. The SESAME Net network currently consists of 15 members from various European countries and consequently offers the opportunity to support and consult in various languages. A short description, an overview of available HPC recourses, a service portfolio and contact data of all SESAME Net members can be found at the project website[8].

Even if the initial problem that was considered is not found to be a clear case for HPC, SMEs are encouraged to join the online forum[9] and to explore all services that are offered by SESAME Net partners. For example, SMEs can benefit from free training events that are organized throughout Europe, even if they are not strictly doing HPC.

As mentioned before, there are different possibilities to enter into collaboration with a SESAME Net HPC centre and take advantage of HPC.

SESAME Net partners offer you:

- **Pure resource usage agreements**, being accounted by hardware resource usage, e.g. compute cycles, data volume, etc.
- **Consulting service contracts** to benefit from HPC expert knowledge, possibly as a supplement to a resource usage contract.
- **Application service provision** / software-as-a-service, accounting for usage time of a pre-deployed and configured software product such as a simulation suite including the necessary resources (hardware, license …).
- **Setting up a joint R & D project**
  - based on contract research;
  - with bilateral benefits, e.g. with the goal to develop a joint project sharing the usage rights or distributing results between the partners.
- **Collaborative projects based on public funding**. The European Union regularly provides funds for collaborative projects to strengthen SMEs and encourage the adoption of new technologies. There might also be national funding opportunities for a collaborative project. SESAME Net tries to provide updated information about funding opportunities e.g. through the Webpage and the Newsletter. There are also other possibilities for collaborative projects, such as the SHAPE (SME HPC Adoption Programme in Europe) pan-European initiative supported by PRACE (Partnership for Advanced Computing in Europe) and the collaborative Fortissimo project that is focused on enabling Manufacturing SMEs to benefit from HPC-based Simulations.

[8] https://sesamenet.eu/centres/
[9] http://forum.sesamenet.eu/

Although rather standard and not particularly outstanding for HPC projects, we don't want to disregard some additional legal aspects that need to be taken care of, including:

- Terms and conditions, including the security, access policies, service level agreements (SLA), accounting and billing procedures;
- Non-disclosure agreements (NDA);
- Intellectual property (IP) rights and agreements.

This chapter briefly describes some of the most used terms in the general area of HPC. It is not intended to be an exhaustive list of terms, but should rather allow a rough orientation. The list will be extended in the future.

HPC terms and technologies

Clusters

A computer cluster is a single logical unit consisting of a set of connected servers, working together as a single powerful machine. The individual servers (aka nodes) of a computer cluster are usually connected through very fast network connections (called the interconnect). The nodes of a cluster often have different purposes, e.g. the cluster is often composed of compute, login, storage and other specialised nodes e.g. for visualisation or with extra-large memory.

A computer cluster can accommodate much more processors and provide much more storage capacity than a single computer and therefore can provide much faster processing speed. The higher speed is mainly achieved by processing numerous tasks in parallel. It means that a number of processors cooperatively solves a problem at the same time (parallel computing).

An example for a computer cluster is the HPC system “Avitohol” which forms the core of the computing infrastructure in the Institute of Information and Communication Technologies (IICT) in Bulgaria, is presented in Figure 2. It consists of 150 HP Cluster Platform SL250S GEN8 servers, each with 2 Intel Xeon E2650 v2 CPUs and 2 Intel Xeon 7120P coprocessors.

Figure 2: Avitohol HPC system at IICT

HPC terms and technologies

Supercomputers

Currently the fastest systems worldwide are often called Supercomputers. There are different constructions and architectures for Supercomputers existing. While in 2000 diverse architectures for supercomputers were popular, today **most supercomputers are built as computer clusters** consisting of mostly standard server, often equipped with faster network connections.

Another interesting development concerns the number of processors a Supercomputer provides. While around 1970 the supercomputers used only a couple of processors, in the 1990 the first machines with thousands of processors appeared. By the end of the 20th century, massive supercomputers with tens of thousands of processors were the norm. Currently the fastest system at the time of writing, "Tianhe-2" of the National Super Computer Centre in Guangzhou, is reported to employ a combined number of 3,120,000 compute cores (a term which also includes specialized processors, c.f. 5.4 Hybrid Systems).

As a **measure of the computation power of a system, the number of floating point operations it can perform in one second (FLOPS)** is often used. The theoretical maximum that the system can achieve in the best case based on its hardware limits is called the (theoretical) peak performance and can simply be calculated using the technical specifications of the vendor. However, for practical measurements and especially for the intercomparison of the performance of these systems, the High-Performance Linpack Benchmark (HPL) has prevailed. This benchmark measures the performance of the system solving a dense system of linear equations.

A list of the highest results achieved using this HPL benchmark is regularly compiled as the "Top500 list of Supercomputers".[10] From this list one can obtain lots of information regarding types of supercomputing systems, their hardware, software, etc. For example, at the time of writing, the 81 fastest supercomputers documented in the list work at more than one Petaflops or 10^{15} floating-point operations per second.

However, problems that arise in industry rarely require such high amounts of computational power. An HPC Cluster of medium size, even though it has never been in the TOP500 list, may be entirely appropriate to handle most problems.

[10] https://www.top500.org/

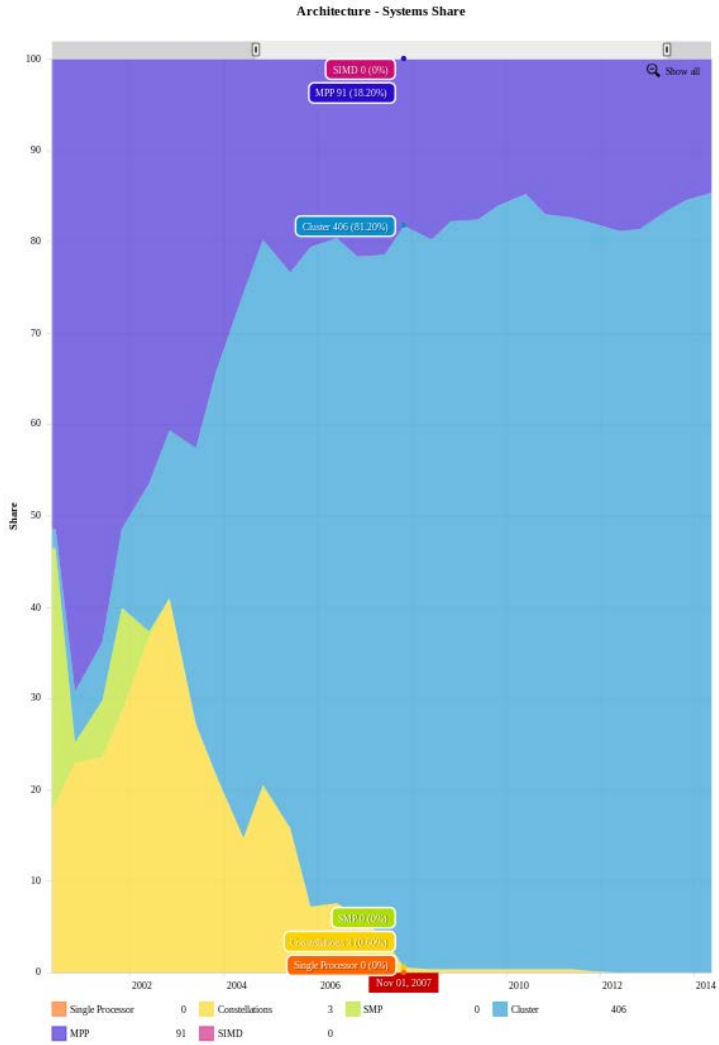

Figure 3: The share of clusters in the Top500 list of supercomputers.[11]

From the above Figure 3 you will see that **85, 2% of the Supercomputers today are clusters** – clusters of commodity hardware, which means they largely utilize the same hardware you will find in your servers. This has had a tremendous impact on HPC and technology providers – maybe also SMEs – making them able to design systems for various applications. Software, in the past one of the major barriers to adopt HPC, now starts to be standardized, which will allow companies with limited IT resources to take advantage of HPC Technology as well.

[11] Image courtesy of Top500 list of supercomputers, source http://www.top500.org/statistics/overtime/

Often modern supercomputer architectures are differentiated by the connection of the processors to memory, because it influences heavily the methods of programming those machines.

HPC terms and technologies

Different types of coupled systems – **Shared memory**

When a system is said to have shared memory, this means that although the system may have hundreds of processors, the memory space is not divided, but is considered shared between these processors. This means that changes made by one processor are visible to the others and that explicit communication is not required. Nevertheless, programming for shared memory systems requires some way of synchronization to be implemented. Nowadays, even desktops and smartphones can be considered as shared memory parallel machines (see Figure 4). The most used application programming interface (API) is the OpenMP® API specification for parallel programming.

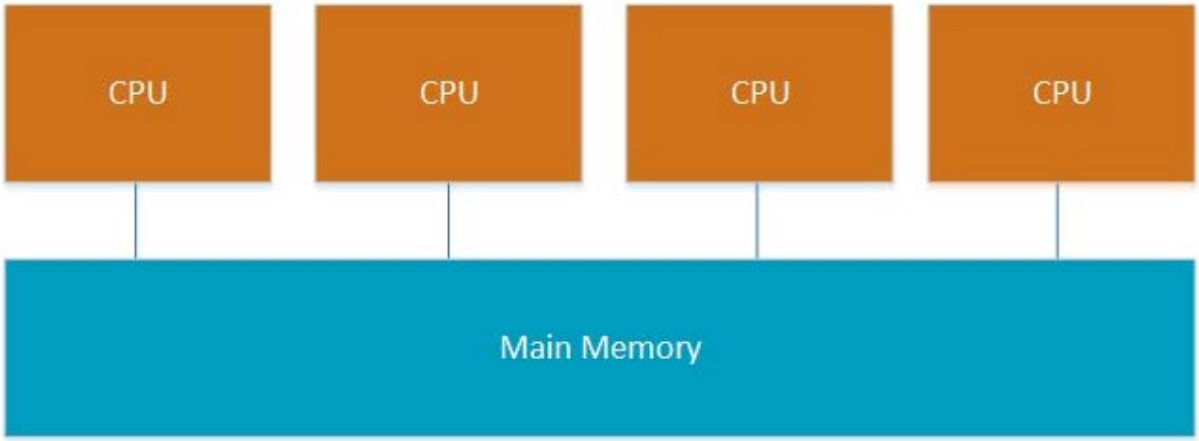

Figure 4: Shared memory systems

HPC terms and technologies

Different types of coupled systems – **Distributed memory**

In systems with distributed memory (see Figure 5) the memory is distributed between the processors. Each processor can only see his own part. Changes made by one processor need to be communicated explicitly if other processors need to see them. The Message Passing Interface (MPI) consists of standardized libraries for message-passing and allows applications to explicitly transfer data between processors and synchronizes the computations.

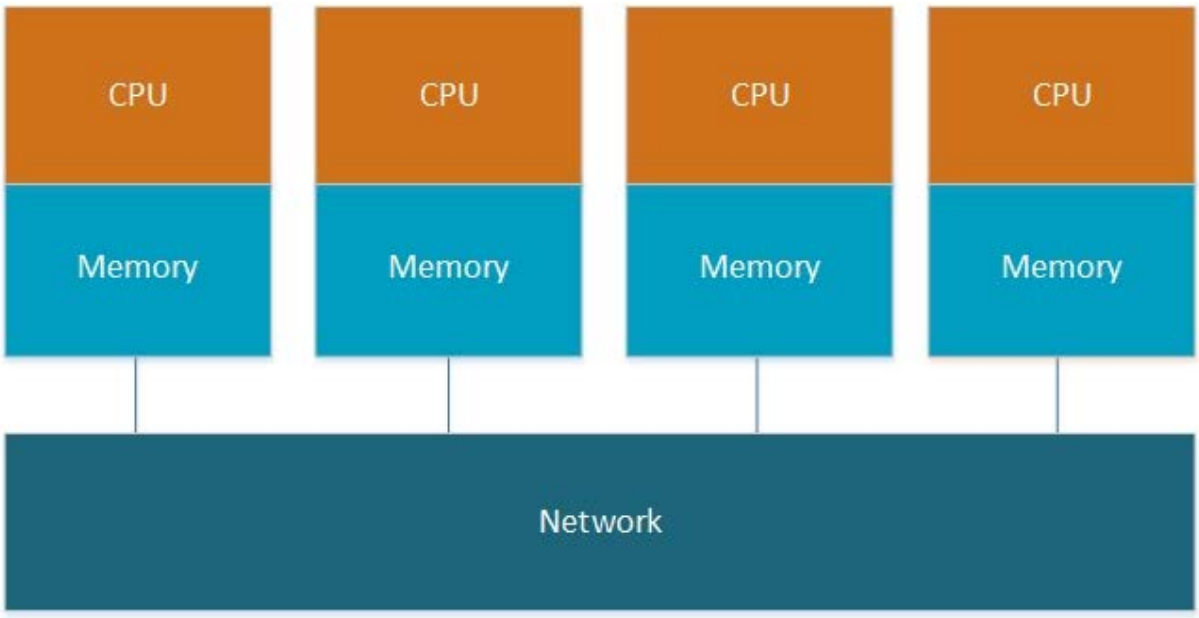

Figure 5: Distributed memory systems

HPC terms and technologies

Different types of coupled systems – **Hybrid Systems**

Most parallel computers are neither only a shared-memory nor a purely distributed-memory type but a mixture or hybrid of both. To get as much as possible performance out of a single compute node they are sometimes combined with accelerators.

Examples of such accelerators are graphics cards or Graphics Processing Units (GPUs as opposed to the general CPUs), general floating-point accelerators or Field Programmable Gate Arrays (FPGAs). Even CPUs for general purposes have vector instructions that are advantageous for HPC. Software that exploits vector processing offers much higher efficiency.

The most popular concepts for developing or porting applications to the above architectures described previously are OpenMP and MPI.

HPC terms and technologies

Parallel Programming

OpenMP

Open Multi-Processing (OpenMP) is a set of compiler directives, library routines and environment variables. OpenMP gives programmers a simple and flexible interface for developing parallel applications for platforms ranging from the standard desktop computer to the supercomputer. OpenMP is used on shared memory machines. On high performance clusters OpenMP is often combined with MPI. This is called hybrid programming.

MPI

Message Passing Interface (MPI) is a standard to allow communication between the processes by exchanging messages on and between the compute nodes. It works on a variety of architectures often also called parallel computers. MPI is generally considered to be the industry standard and forms the basis for most communication interfaces adopted by HPC programmers. The standard defines the syntax and semantics of a core of library routines.

This Chapter provides a practical example of the usage of an HPC resource managed by one of the partners of SESAME Net. It demonstrates how simple command-line applications can be easily executed on one or more systems, and how simple C programs can make use of more than one CPU core and even multiple nodes at the same time.

HPC terms and technologies

Usage example

Connecting to the HPC centre

While not all use cases actually require the SME to manually perform computations on the HPC system, here we describe a common way of doing so.

Typically, the HPC centre will provide access to the system at user-level, meaning without any kind of administrative rights, and through an interactive command-line interface for general purpose access.

The most typical scenario, when we have access via MS Windows, would proceed as follows:

- An account request is submitted to the HPC centre via a web interface or similar;
- The account is validated and confirmed. A username and password or alternative login credentials are provided to the user.
- The user can then login to a server which is connected to the Cluster (login node) via a program implementing the Secure Shell (SSH) protocol, also called an SSH client. SSH client software is available for all popular operating systems, including mobile systems such as Android.

The SSH protocol also accommodates file transfers, so data can easily be transferred to and from the HPC system. Software to provide more comfortable file transfers through SSH exists for several operating systems as well, e.g. free WinSCP software for Windows.

Optionally, to use programs with a graphical user interface on the HPC system, the client computer, i.e. your desktop PC, will need the installation of additional software. The HPC systems typically use the “X Window System”, which requires a counterpart on the client PC to function over network connections. The method of displaying a GUI window of an application running on the remote server on your PC is called “X-forwarding”, and

typically needs to be enabled in the SSH client. Usually Linux PCs provide this functionality out-of-the-box, and there are free applications that provide such functionality on MS Windows.

Once connected to the HPC system, the user can interactively discover what kind of software packages and libraries are installed. If the particular software intended to be used is not readily available on the HPC system, it can be deployed by the HPC centre. Typically, the centre provides web-based documentation, as well as a support helpdesk for these kind of questions and general help.

Because the HPC systems may support numerous users at the same time, they usually provide a way to "enqueue" the execution of programs in a kind of waiting line system. A single call of an application in such a waiting system is then called a "(compute) job" or old fashioned "batch job". A practical example of this mechanism is given in the next section.

Running jobs

Usually the computational workloads on an HPC system are divided into jobs. One job takes a particular subset of the available processors (or processor cores) and executes an application on them for a period of time, until completion or until its allotted time expires. Most of the jobs on an HPC system are parallel, taking much more than one CPU core for their execution. Jobs are sent to queues, where they are executed following some prioritization strategy.

Every user is given information about the suggested queues. A job is usually executed without need for user interaction, but if required, users can also request an interactive type of job. Interactive jobs are usually used for development and testing. The status of a job can be queried and usually e-mail notifications may be requested to be notified of status changes.

Some HPC centres are able to provide higher levels of security, by creating virtual (or real) private sub-clusters, for example created for exclusive use by a company. Because of the exclusive access, the company may then organize the execution of workloads differently.

Access to the HPC cluster at IICT and execution of the tasks (examples)

Here we demonstrate some examples how the given user (a student who has account on the HPC cluster at IICT) can *login in* and *run* training MPI and OpenMP tasks (jobs).

To access the cluster enter the following command:

```
ssh trainee<number>@ gw.ipp.acad.bg
```

Enter your account name (student<number>) and your account password.

```
cp -r /opt/exp_software/documentation/Training/* /home/trainee<number>/
```

In our first example we will submit a simple job which will print Hello World in its output. Navigate to:

```
/home/traineet<number>/Examples/Hello
```

and enter the following commands:

```
cat hello.sh
#!/bin/bash
#PBS -q lifesci
echo "Hello World!"
```

To submit the shell script to the cluster use the following command:

```
qsub hello.sh
```

This command will print the following output:

```
<Job_ID>.torq.hpcg
```

Your job will be completed almost instantly and the standard output and the standard error output will be in ***hello.sh.o<Job_ID> and hello.sh.e<Job_ID>*** respectively.

Enter the following commands:

```
cat hello.sh.o<Job_ID>
cat hello.sh.e<Job_ID>
```

Our second example will print the resource granted to your job. Navigate to:

```
/home/trainee<number>/Examples/Resources
```

Print the shell script ***res.sh*** (use the command ***cat res.sh***).

```
#!/bin/bash
#PBS -q lifesci
#PBS -l nodes=2:ppn=4
cat $PBS_NODEFILE
```

The ***#PBS -l nodes=2:ppn=4*** denotes that our job will run on 2 nodes and will use 4 logical cores from each node. This can be modified by altering the ***res.sh*** or by using the ***-l*** option of the ***qsub*** command. Submit the script using following command:

```
qsub -l nodes=1:ppn=16 res.sh
```

Open the output:

```
cat res.sh.o<Job_ID>
```

The third example will familiarize you with the commands used to monitor and terminate your job. Navigate to

```
/home/traineet<number>/Examples/Loop
```

Here you will see the script ***loop.sh***:

```
cat loop.sh
#!/bin/bash
#PBS -q lifesci
#PBS -l nodes=1:ppn=1
while(true)
do
a=1+1
done
```

Now let us submit ***loop.sh***. Again we will use the ***qsub*** command, but we will specify an e-mail on which we will receive status updates on our job:

```
qsub -M <your_e_mail_address> -m abe loop.sh
```

Such options to ***qsub*** can be specified on the command line, but alternatively they can be specified in the shell script. Use the following commands to view the information for your job:

```
qstat <Job_ID>
qstat -f <Job_ID>
```

qstat -n <Job_ID>

The last command will print a list of the nodes and cores granted to your job. The first entry in the list (***wn<number>.hpcg***) is the host node. To navigate to it:

ssh wn<number>

Notice that you don't have to specify password. Once you are logged to the execution node enter the ***top*** command. You will see a list of all the running processes sorted by their CPU usage. You can notice that your process is on top. To close the ***top*** application press ***q*** button on your keyboard. Now return to ***wn02*** node by entering the ***exit*** command.

You can also get information about all jobs of a given user using the ***qstat*** command. Enter the following command:

qstat -u <user_name>

Using ***qstat*** you can also get information about the queues. Enter the following commands and consider their outputs:

qstat -q will print information about all the queues.

qstat -q lifesci will print the information about the lifesci queue.

qstat -Q will print more detailed information about all the queues.

Now it is time to terminate our job:

qdel <Job_ID>

If you don't remember your ***Job_ID*** you can either get a list of all your current jobs using the command:

qstat –u <user_name>

or use the command:

qdel ALL

which will try to delete all jobs, but will delete only yours and will produce a lot of error messages (caused by your unauthorized deletion requests).

Fourth example: Execution of MPI and OpenMP jobs.

Before we continue with MPI we must choose an MPI compiler. Enter the following command:

mpi-selector-menu

From the list of options choose ***openmpi_gcc-1.3.3***. Select this setting to be per-user (u) and overwrite any existing settings ***(y)***. Next exit the ***mpi-selector-menu*** (press ***Q***). Close the terminal and reopen it again (you will have to log on to ***gw.ipp.acad.bg*** using ***ssh*** again).

Next navigate to

```
/home/trainee<number>/Examples/MPI
```

In the directory there are two files – ***helloMPI.c*** and ***helloMPI.sh***. View them by using the following commands:

```
cat helloMPI.c
```

Enter the following command:

```
mpicc helloMPI.c -o helloMPI
```

This will build your program. To run it on the login node (only useful during testing):

```
mpirun -np 10 ./helloMPI
```

This will run 10 instances of your program. Alternatively, you can submit your job using the ***qsub*** command. Print the shell script using:

```
cat helloMPI.sh
```

Submit your job using the following command:

```
qsub helloMPI.sh
```

Check your job's status or enter the ***ls*** command to check if your results are present.

```
cat helloMPI.sh.o<Job_ID>
```

Next navigate to:

```
/home/trainee<number>/Examples/OMP
```

Again, there are two files in the directory –***helloOMP.c*** and ***helloOMP.sh***. View them using:

```
cat helloOMP.c
cat helloOMP.sh
```

Build your program using the command:

```
gcc -fopenmp helloOMP.c -o helloOMP
```

To run it on the login node (only useful during testing), enter:

```
./helloOMP <number_of_threads>
```

To submit it enter:

```
qsub helloOMP.sh
```

Navigate to:

```
/home/trainee<number>/Examples/MPInOMP
```

There are two files in the directory –***helloMPInOMP.c*** and ***helloMPInOMP.sh***. Build the program using the following command:

```
mpicc -fopenmp helloMPInOMP.c -o helloMPInOMP
```

To run in on your machine use:

```
mpirun -np 10 ./helloMPInOMP 3
```

This will run the 10 instances of the application with 3 threads each.

Print the ***helloMPInOMP.sh***. Submit it using the ***qsub*** command and review its output.

Based on the excellent collaboration between HPC competence centers and SMEs, SESAME Net partners developed several examples (use cases) illustrating how real companies benefit from HPC technologies in different areas of business.

Here we present some selected success stories. The full list of the use cases can be found here:

https://sesamenet.eu/resources/success-stories/ .

Use Cases:

Use Case-1: Biomedicine

HPC-Competence Centre

The Institute of Information and Communication Technologies (IICT) at the Bulgarian Academy of Sciences hosts the Bulgarian largest HPC centre (http://www.hpc.acad.bg/) and provides knowledge, expertise and services to businesses and researchers to develop innovative applications and products.

Enterprise

Established in 1995 AMET Ltd. is a company dedicated to development, modern manufacturing and distribution of electronic medical equipment and modules, mechanical parts and units for incorporation. With ~100 employees AMET Ltd achieved a position in a highly competitive market and needs scientific expertise in order to develop products with enhanced power and precision.

How HPC makes the difference

AMET Ltd., (see http://amet-bg.com/en/), started to develop medical device for radio-frequency ablation of hepatic tumours. It needed precise mathematical modelling and computer simulation of the heat transfer process in order to optimize the parameters of this low-invasive therapy technique. Adequate representation of the problem was achieved by FEM discretization for a time-dependent partial differential equation of a parabolic type, generated based on segmented medical image. The 3D simulations (see Figure 6) were performed using a supercomputer and allowed precise estimation of the ablation parameters and an increase in power efficiency.

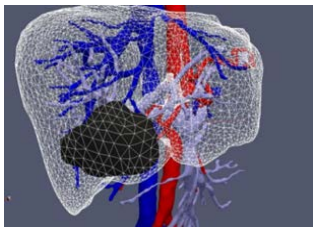

Figure 6: 3D FEM discretization of the liver

Figure 7: HFiT-T system produced by AMET, commercially available in Europe

Significant industrial impact

As a result of the successful project "Innovative technology solutions for radio-frequency thermal ablation" (2012-2014), implemented in collaboration with IICT, AMET Ltd. developed a new device (see Figure 7) for high frequency thermal ablation HFIT-T and started its production. The device ensures accurate determination of the relevant parameters of the process of thermal ablation in accordance with the tumour size, needle size, and distance from the walls of the organs at risk (major blood tract, bile ducts, etc.) duration of the procedure and the radio-frequency power.

Use Cases:

Use Case-2: Manufacturing and Materials

HPC-Competence Centre

The Fraunhofer Institute for Algorithms and Scientific Computing SCAI offers the application of sophisticated in-house and/or commercial optimization and robust design methods in combination with numerical simulations for complex flow problems.

Enterprise

The participating industrial company was Hennecke Polyurethane Technology. Hennecke designs and constructs high-performance machine and plant technology for polyurethane processing, enabling their customers to achieve high-quality and efficient production results.

How HPC makes the difference

The cooperation between SCAI and Hennecke featured the optimization of two different subdomains of polyurethane processing. The first case examined an injector for mixing the polyurethane components Polyol and Isocyanate. Based on a sensitivity analysis varying several geometric parameters, an optimization process was launched, which yielded a set of optimal parameters for a required range of operating mass flow. In the second case, slab stock foam plants used for the production of foam blocks have been optimized, to achieve a foam of high quality. In both tasks the fluid solver STAR-CCM+ from cdadapco, the optimization software DesParO from Fraunhofer SCAI and a self-developed process chain environment was applied to optimize processes and simulate them in acceptable computation times on SCAI's HPC Cluster.

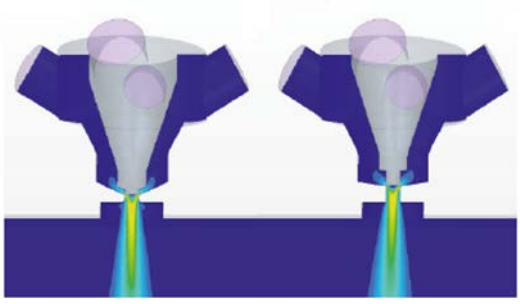

Figure 8: Mixing injectors

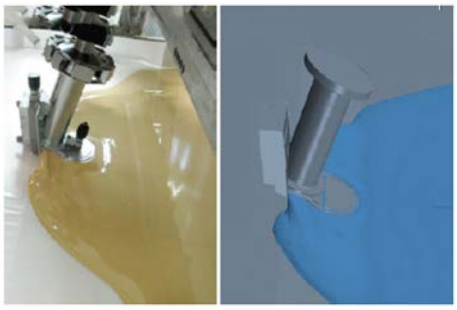

Figure 9: Slab stock foam (left); Simulated process (right).

Significant industrial impact

- Optimization of plants and components for the polyurethane production;
- Cost reduction;
- Shorter computation time at simulations.

Use Cases:

Use Case-3: Virtual Prototyping

HPC-Competence Centre

IT4Innovations National Supercomputing Centre is an important part of the e-Infrastructure of the Czech Republic focused on HPC research and services. The Centre operates the most advanced HPC technologies and services and makes them available to Czech and foreign research teams from both academia and industry. One of the key functions of the IT4Innovations National Supercomputing Centre is to support industry in the Czech Republic.

Enterprise

The company BORCAD cz was established in 1990 as a construction and development studio. Today BORCAD cz is a leading European producer of railway and medical technologies. The company employs 190 people, exports to more than 80 countries worldwide and thanks to its unique design and original construction solutions is one of the most innovative companies in this field.

How HPC makes the difference

Numerical modelling and simulations are commonly used in the research and development of new products because they significantly shorten the time necessary to bring a new product to the market and also save a lot of money. The old fashioned way of product development by trial and error, where usually quite costly prototypes have to be built and tested, is increasingly being replaced by numerical modelling of virtual prototypes.

IT4Innovations has cooperated with BORCAD since June 2013. The company searched for a solution of a problem they had with entering the UK market with passenger seats for regional and long-distance rail transport. BORCAD needed certification that their passenger seat complied with regulations. The certification process involves a crash test which checks not only seat integrity but also bio-mechanical criteria measured on crash test dummies. Since physical tests are very expensive, BORCAD decided to use numerical modelling and simulation to perform a virtual crash test on new designs and perform only two physical tests. The cooperation helped BORCAD to obtain the certificate after successfully passing the crash tests with much less resources and time invested.

Figure 10: Simulation of crash test of railway passenger seat.

Significant industrial impact

- Shortening the design cycle;
- Reducing costs;
- Adding value to existing products;
- Creating new products.

Summary

This handbook tried to facilitate the entrance into the broad topic of HPC, to demonstrate how SMEs can add tremendous value to their products with HPC and to motivate SMEs to try it.

Even when you don't need to know much about the technology, because you are only interested to use an HPC system with HPC - ready applications - commercial codes from independent software vendors (ISVs) or open source - you may find our first short introduction to the common used technical terms interesting and helpful.

For those SME managers that were interested by the benefits of the HPC technologies, it is important to understand that all SESAME Net partners are open for contacts and collaboration and have established processes and procedures.

As demonstrated by our short example of starting an application on a Supercomputer of one of the SESAME Net partners, there is no magic with HPC. Even when the HPC centres are more engaged in the research world, our partners are familiar with SLAs and contracts. Therefore, we also addressed it.

The real life examples - you will find more at www.sesamenet.eu - and the information about the HPC competence centres and their services give you possibilities for further reading, establishing contacts and taking the first steps towards using HPC.

As often mentioned in the media, supercomputers are no longer limited to large academic, government and industrial research facilities. The HPC competence centres of the SESAME Net are part of the "democratization of HPC".

Get in touch with us. The staff of the SESAME Net HPC Centres are always available for a personal conversation. You will be able to learn more about HPC and how you can benefit from the technology. The HPC centres are also informed about public funding options to support SME activities in Europe.