



INSTITUTE OF INFORMATION
AND COMMUNICATION
TECHNOLOGIES
BULGARIAN ACADEMY OF
SCIENCE



Avitohol supercomputer: features and applications

Assoc. Prof. Emanouil Atanasov

Head of Grid Technologies and Applications Department

Institute of Information and Communication Technologies – BAS

Email: emanouil@parallel.bas.bg



Academic HPC resources in Bulgaria

AVITOHOL at IICT-BAS

150x HP ProLiant SL250s Gen8 each with
2x Intel Xeon E5-2650 v2 (8C/16T),
64 GB DDR3-1866 RAM and
2x Intel Xeon Phi 7120P
6x HP ProLiant DL380p Gen8 nodes with
2x Intel Xeon E5-2650v2 (8C/16T),
64 GB DDR3-1866 RAM
Infiniband 56 Gb/s FDR
Storage system with 96 TB



Total Performance:
RPeak: 412.3 TFlop/s
RMax: 264.2 TFlop/s
Top 500 position: 389

HPCG cluster at IICT-BAS

36 blades BL 280c(2x Intel X5560(4C/8T); 24GB DDR3);
8 management nodes HP DL 380 G6(2x Intel
X5560(4C/8T); 32GB DDR3);
2 HP ProLiant SL390s G7(2x Intel E5649(6C/12T); 96GB
DDR3)
8x nVidia TESLA M2090 per server;
2 HP SL270s Gen8 (2x Intel Xeon E5-2650 v2(8C/16T);
128GB DDR3)
Total number of Xeon Phi 5110P coprocessors: 9
Total 132TBs of system storage

TOTAL PERFORMANCE:
RPEAK: 22.94 TFlop/s



PHYSON at Sofia University

53 Intel Xeon x86_64 processors
524Gibs of system memory
6.5TBs of system storage
2x nVidia Tesla M2090 graphics processors

TOTAL PERFORMANCE:
RPEAK: 3.57 TFlop/s
RMAX: 3.22 TFlop/s



NCSA IBM Blue Gene/P

8192 PowerPC 450 processors
4TBs of system memory
12TBs of system storage
IBM proprietary interconnect with
2.5 μ s latency and 10GBps bandwidth

TOTAL PERFORMANCE:
RPEAK: 27.85 TFlop/s
RMAX: 23.45 TFlop/s



MADARA at IOCCP-BAS

54 Primergy RX200 S5 servers with
2 Intel Xeon E5520(4C/8T) each
and a total of 800GB DDR3 1066MHz
20Gb/s DDR Infiniband
108TB System Storage by Fujitsu FibreCat SX100



The Avitohol supercomputer



Hardware features

Servers	150 x Dual CPU HP ProLiant SL250s Gen8
CPUs	Intel Xeon E5-2650v2 2.6GHz – 8 cores /16 HT
RAM	64 GB per node
Coprocessors	300 x Intel Xeon Phi 7120P(x86) - 61 cores / 244 HT
Total CPUs	2400 cores/4800 HT + 18300 cores/ 73200 HT
Total RAM	14400 GB (9600 + 4800)
Disk Storage	100 TB
Interconnect	Non-blocking FDR Infiniband
Latency	1.1 μ s
Bandwidth	56 Gb/s

Environment and

Peak Performance CPU	50 Tflop/s
Pear Performance Accelerators	362 Tflop/s
Total Peak Performance	412 Tflop/s
Real Measured Performance	264 Tflop/s
Max Power	250 kW

The system consists of 4 dual racks of type HP MCS 200. Each one provides power and cooling for up to 50 kW of equipment, cooled by water. Last on Top500 List on 389 place (Nov 2015)

<http://www.top500.org/system/178609>



Xeon Phi accelerators

- x86 Architecture
- 61 cores
- 244 threads
- 512 bit SIMD
- 16 GB RAM
- 352 GB/s bandwidth
- Works synergistically with Intel Xeon processors
- Give about 90% of the theoretical performance



The Avitohol supercomputer

Architecture of one node



Low-Latency, High-Bandwidth
Network - Infiniband FDR 56
Gbps

- RedHat Enterprise for HPC
- Intel Cluster Studio (compilers, MPI, MKL)
- GNU Compiler Collection
- Torque/moab for resource management with web and command line interface
- Monitoring and accounting tools
- Services running on virtual machines
- Rich set of applications and libraries like:
 - Boost, mumps, octave, OpenFOAM, WRF, WRF-Chem, NetCDF, Gromacs
- ...



Connected resources - Heterogeneous HPC Cluster

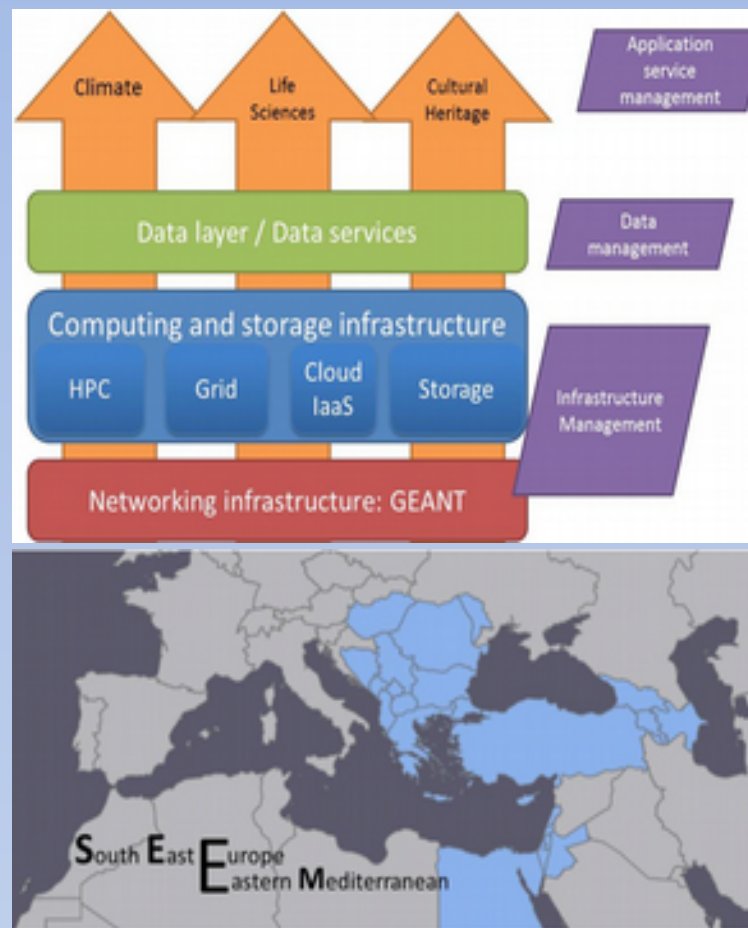
- 36 blade servers 2 Xeon X5560 @ 2.80GHz, 24 GB RAM – 576 cores total
- 132 TB - three SAN systems with Fibre Channel interconnects
- 2 servers with 16 NVIDIA Tesla M2090 graphic cards
- 2 server with up to 8 Intel Xeon Phi 5110P Coprocessors
- Interconnected with non-blocking InfiniBand
- Peak performance 22.93 Tflop/s



- IICT-BAS operated different types of clusters and supercomputers like the Origin 2000
- IICT-BAS coordinates the Bulgarian participation in the European Grid Initiative, first through the EGEE series of projects, then *EGI-InSPIRE* and *EGI-Engage*
- IICT-BAS is one of the leading partners in a series of projects that spread the expertise of advanced countries in the region to the less developed countries – from SEEGRID to HP-SEE to VI-SEEM
- Innovative scientific applications are developed and run on the available HPC equipment at IICT-BAS by teams of scientists from institutes of the Bulgarian Academy of Sciences and leading universities in Bulgaria.

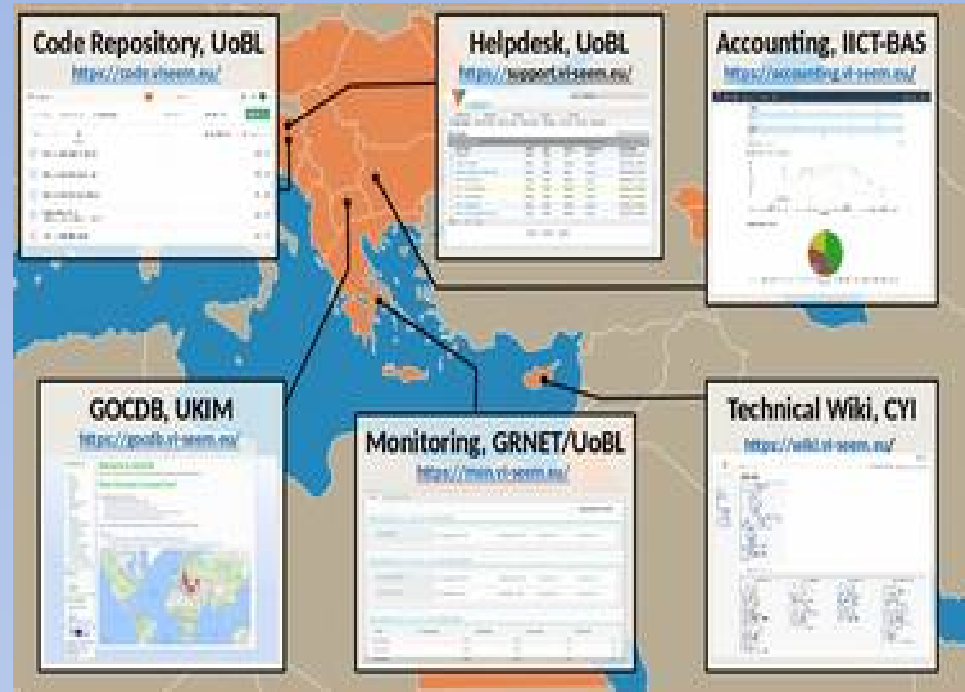


- The project VI-SEEM aims to integrate all kinds of electronic infrastructure in the region – Grid, Cloud, HPC, BigData, in order to provide single point of access for scientists in the region of South Eastern Europe and the Mediterranean.
- Bulgaria is the leader in the HPC area.
- Provides user-friendly integrated e-Infrastructure platform for Scientific Communities in:
 - Climatology,
 - Life Sciences,
 - Cultural Heritage
 for the SEEM region by linking compute, data, and visualization resources, as well as services, software and tools.
- VRE Portal is available at <http://vre.vi-seem.eu/>

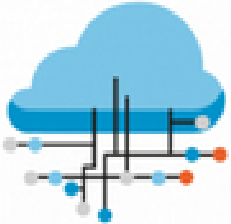


VI-SEEM Regional Services

- Code Repository – GitLab
- GOCDB - topology and resources information.
- Monitoring – backend based on Nagios with custom frontend developed by GRNet
- Accounting portal – developed by IICT-BAS



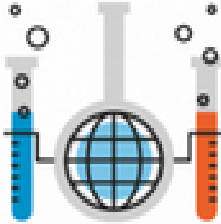
- VRE Portal is available <http://vre.vi-seem.eu/> Continuously populated with new content



Virtual Communities

Climate

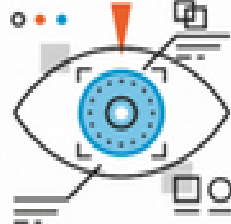
- Climate Applications
- MMS
- OPENFOAM
- WRF
- WRF-CHEM



Virtual Communities

Life Sciences

- Life Sciences Applications
- VMD
- NAMD
- JVM
- GROMACS



Virtual Communities

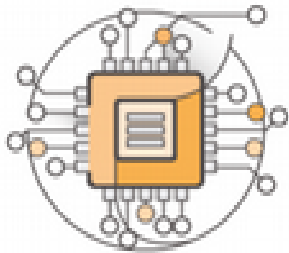
Digital Cultural Heritage

- Digital Cultural Heritage Applications
- SOL
- Clowder
- CH-CBIR
- AutoGR


Quick Links

[How to upload code](#)


[How to upload a dataset](#)



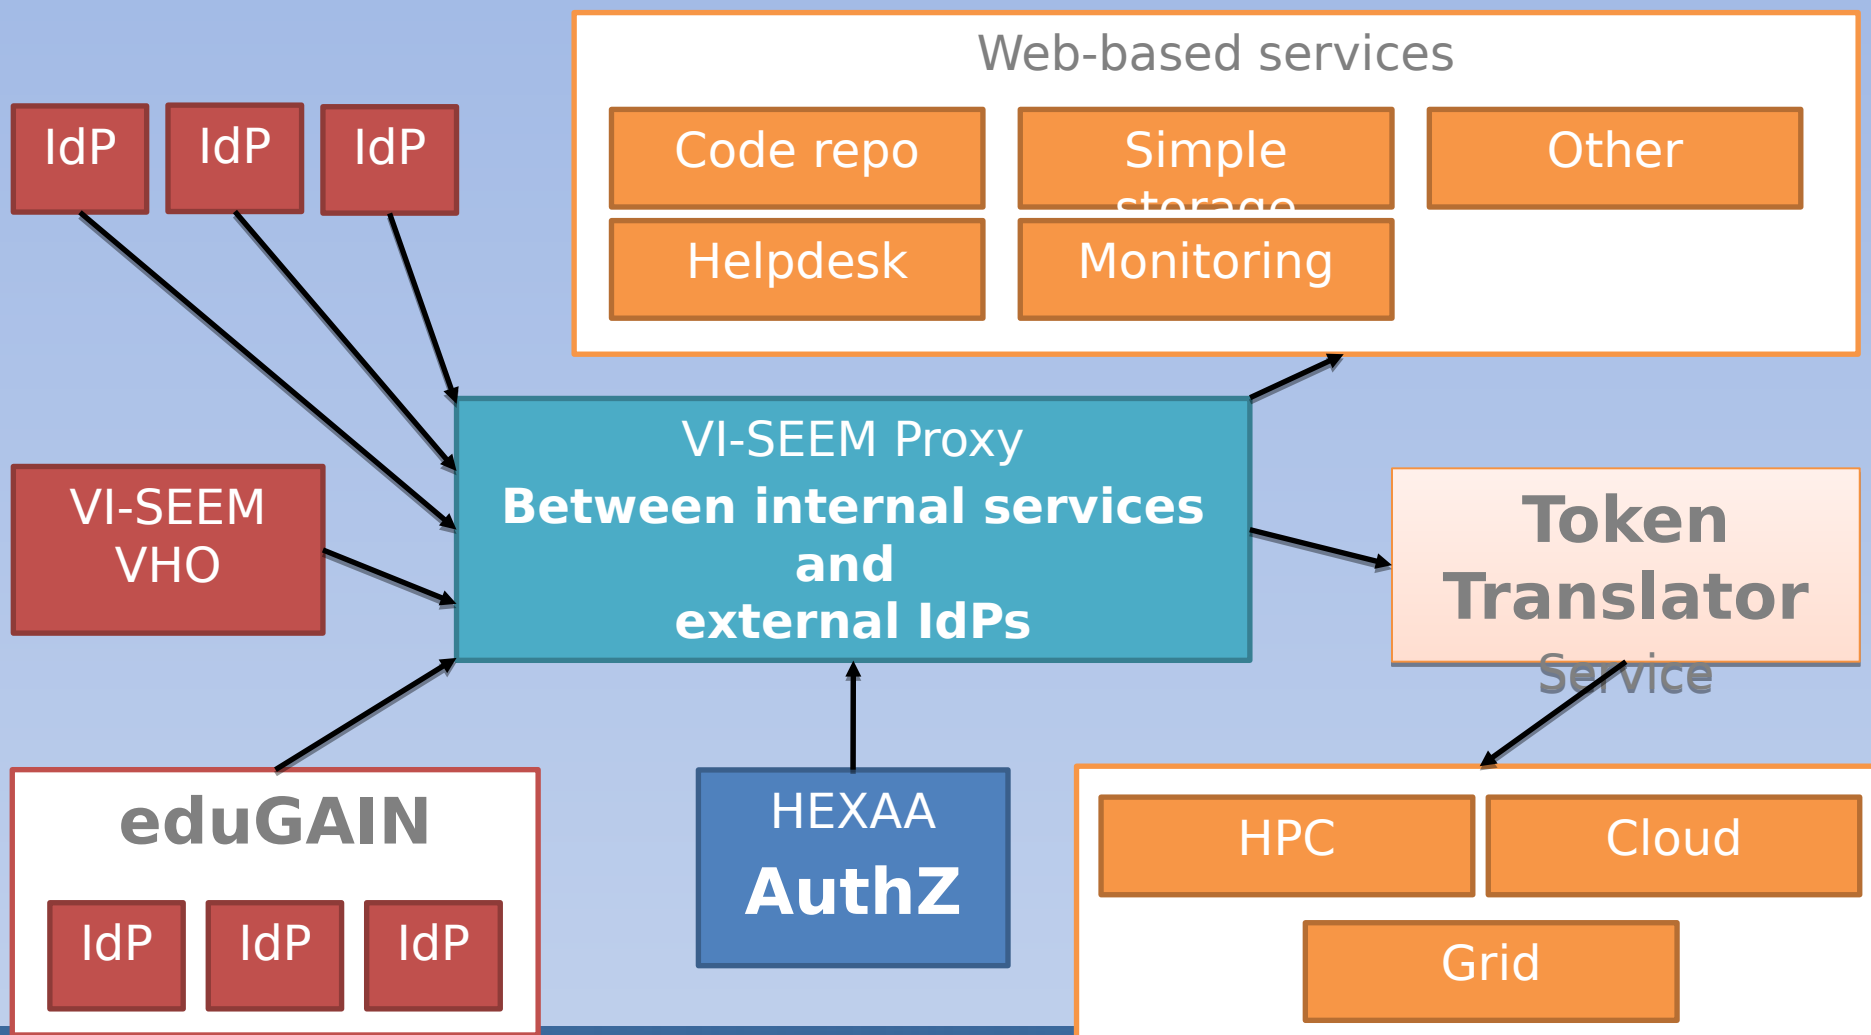
Access to Compute Resources



Access to Data Resources



Access to VI-SEEM Training Portal

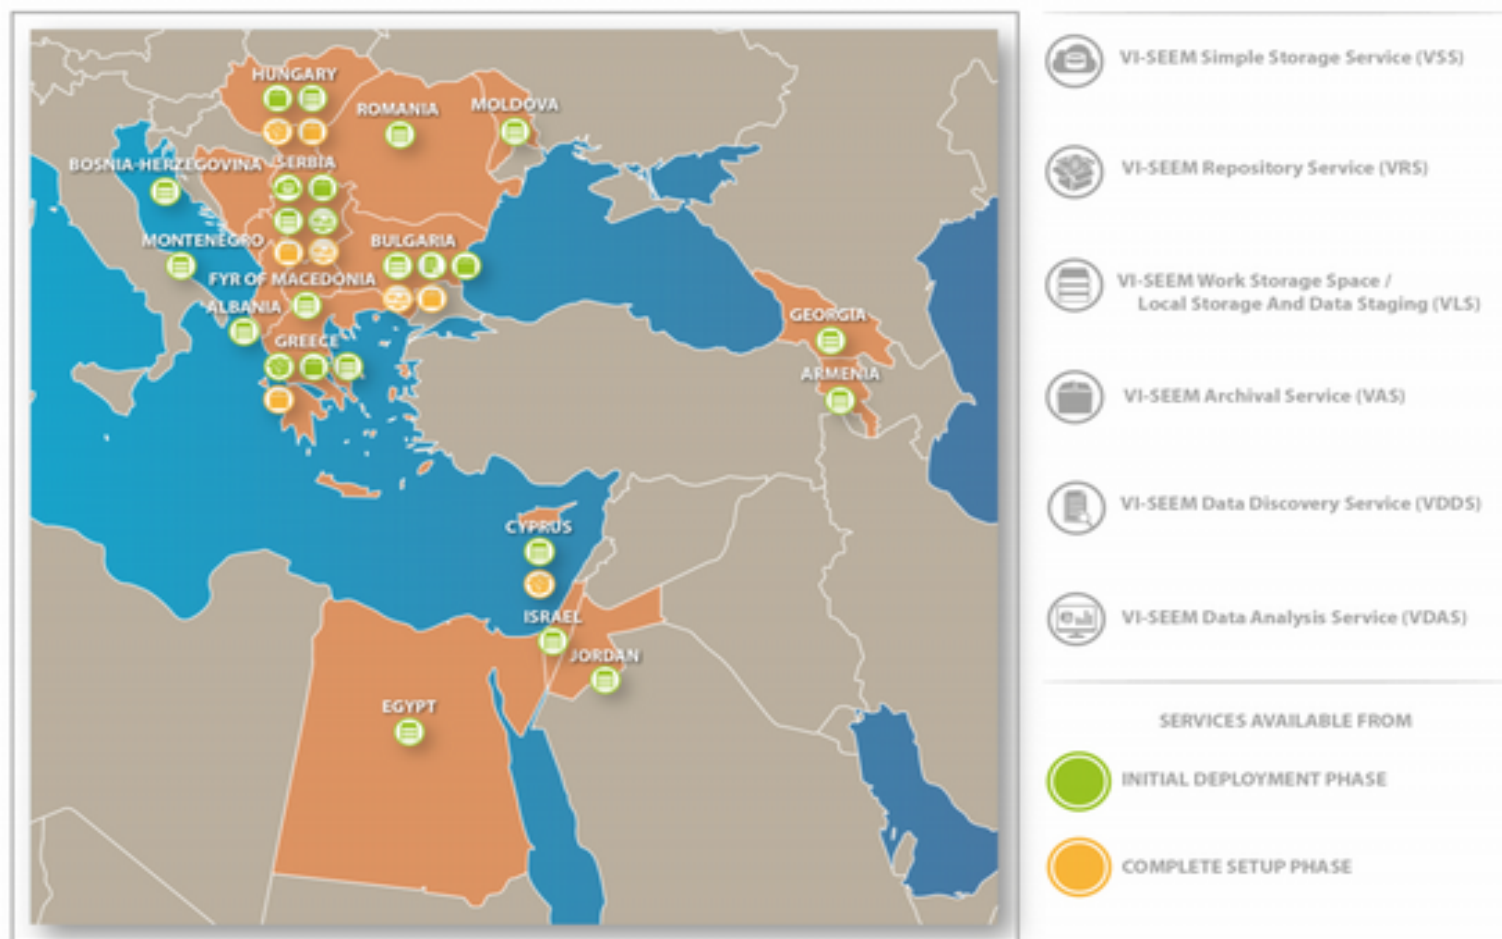


Data management services

- Functions allowing for data management for selected Scientific Communities, engage the full data management lifecycle
 - VSS – Simple Storage Service (simplestorage.vi-seem.eu)
 - VRS – Repository Service (repo.vi-seem.eu); integrated with PID service
 - VAS – Archival Service (deployed at 6 sites – GRNET, IPB, IICT-BAS, NIIF, IUCC, BA)
 - VLS – work storage space / local storage and data staging (at 12 sites)
 - VDDS – Data Discovery Service (search.vi-seem.eu)
 - VDAS – Data Analysis Service (hadoop.ipb.ac.rs)
 - PIDs (handle.grnet.gr)



Data management services - spread



VRE portal



- All services integrated through the user-facing VRE portal
- <https://vre.vi-seem.eu/>
- Organized per Scientific Community
 - Climate SC
 - Life Sciences SC
 - Digital Cultural Heritage SC
- Access to VI-SEEM services and resources: Compute, Data, Domain-specific, Training
- Guidelines on how to contribute to
 - Applications
 - Workflows/codes
 - Datasets
 - Domain-specific services
- Domain-specific services integrated in the portal in a series of phases



Domain-specific services

- ❑ VRE Scientific Application Environment
 - ❑ Optimized applications and libraries
 - ❑ Virtual Machine (VM) images
 - ❑ Codes from the three scientific communities

- ❑ Workflow, software tools repository

- ❑ Regional community datasets

- ❑ Application level services



- ❑ Climate
 - ❑ Live Access Server



- ❑ Digital Cultural Heritage
 - ❑ VI-SEEM Clowder
 - ❑ 3DINV
 - ❑ AUTOGR



- ❑ Life Sciences
 - ❑ ChemBioServer
 - ❑ AFMM
 - ❑ NANO-Crystal
 - ❑ Subtract



Application-level service flagships



□ Climate

- Live Access Server, <http://las.vi-seem.eu/las>

A web server providing flexible access to geo-referenced scientific data, offering visualization & post-processing capabilities for climate data



□ Digital Cultural Heritage

- VI-SEEM Clowder, <http://dchrepo.vi-seem.eu/>

A Digital Culture Heritage repository which also offers integrated interactive visualization tools



□ Life Sciences

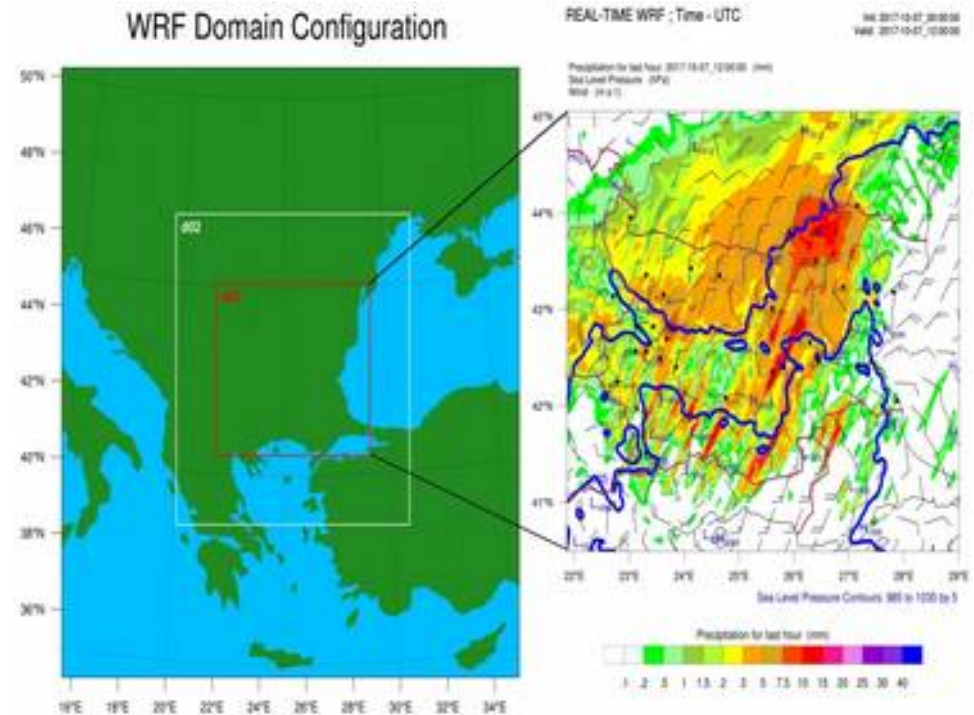
- ChemBioServer, <http://bioserver-3.bioacademy.gr/Bioserver/ChemBioServer/>

A web-based pipeline for filtering, clustering and visualization of chemical compounds used in drug discovery

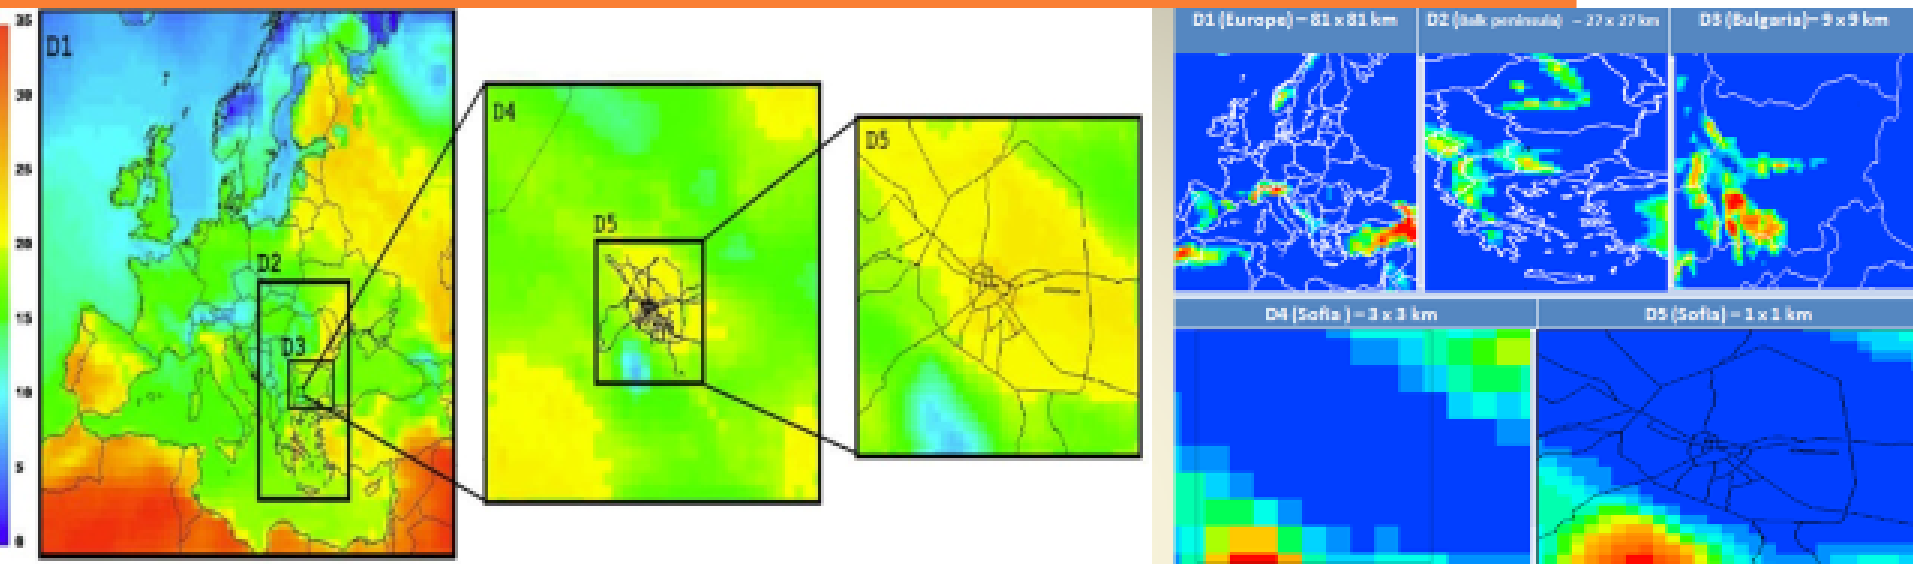


Weather Research and Forecasting (WRF) Model

- numerical weather prediction (NWP) system designed to serve both atmospheric research and operational forecasting needs
- 3D model that takes into account landscape peculiarities
- Used for both operational and research purposes



Applications



Study Atmospheric Composition Impact on Quality of Life and Human Health in Sofia

- WRF and CMAQ nesting capabilities are applied for downscaling the simulations to a 1 km step for the innermost domain.
- The simulations is performed day by day for a period of 7 years (2008-2014). The surface concentrations are calculated for the four seasons and annually by averaging the typical fields.

Development of NIGGG-BAS and support of IICT



National projects supported

Project Name	Institution	PI Name	Grant
Development and Investigation of quasi-Monte Carlo Algorithms for Extreme Parallel Computer Systems, funded by Bulgarian NSF, 2014 - 2016	IICT-BAS	Assoc. Prof. T. Gurov	#DFNI-I02/8
Estimate of brown bear population in Bulgaria on the basis of mathematical, statistical and biological analysis of monitoring data", funded by PUDOOS, 2013 - 2015	IICT-BAS	Assoc. Prof. T. Gurov	#D9190
Parallel and Distributed Computing Practices	IICT-BAS	Assoc. Prof. I. Lirkov	
Numerical methods for coupled systems and computer modelling of biomedical and environmental problems	IICT-BAS	Assoc. Prof. K. Georgiev	#DFNI-I01/5
Efficient Parallel Algorithms for Large-Scale Computational Problems, 2015 - 2016	IICT-BAS	Prof. I. Dimov	#DFNI-I02/20
Preparation of operational numerical physical model of the Black Sea	Faculty of Physics-SU	Assoc. Prof. E. Peneva	
Dynamics , mechanisms and sequence of bonding of the proteins responsible for DNA repair in living cells	IMB-BAS	Assoc. Prof. S. Stoykov	#DFNI-B02/16
Modeling of catalytic systems	Faculty of Chemistry-SU	Prof. G. Vaysilov	
Efficient numerical methods and parallel algorithms for dynamical analysis of plate structures, 2016 - 2017	IICT-BAS	Assist. Prof. PhD S. Stoykov	229/18.02.2016
New mathematical methods of machine learning with applications in modern technology for genetic sequencing	IMI-BAS	Prof. E. Stoimenova	#DFNI-I02/19
Fundamental research on stochastic branching processes - borderline behavior, statistical inferences and applications	FMI-SU	Prof. M. Bojkova	#DFNI-I02



International projects supported

Acronym	Project Name	PI Name	Grant
VI-SEEM	VRE for regional Interdisciplinary communities in Southeast Europe and the Eastern Mediterranean, 2015 - 2018	Assoc. Prof. T. Gurov	#675121, EC H2020
MMAC	Centre of Excellence for Mathematical Modeling and Advanced Computing, 2015 - 2016	Corr. Mem. S. Margenov	#664406, EC H2020
SESAME-NET	Supercomputing Expertise for SmAll and Medium Enterprise Network, 2015 - 2017	Assoc. Prof. T. Gurov	#654416, EC H2020
EGI-Engage	Engaging the EGI Community towards an Open Science Commons, 2015 - 2017	Assoc. Prof. E. Atanassov	#654142, EC H2020
CRoNoS	Computationally-intensive methods for the robust analysis of non-standard data, 2015 - 2019	Prof. A. Karaivanov	COST Action: IC1408
PRACE 4 IP	PRACE Fourth Implementation Phase, 2015 - 2017	Prof. S. Markov	#653838, EC H2020
AComIn	Advanced Computing for Innovation, 2012 - 2015	Prof. G. Angelova	#316087, EC FP7
QTLeap	Quality Translation by Deep Language Engineering Approaches, 2013 - 2016	Assoc. Prof. K. Simov	#610516, EC FP7

- Processes modeling. Assistance using various simulation codes.
- Design and optimization of HPC and Big Data infrastructures for SMEs.
- HPC Consultancy: how can your company benefit from supercomputing.
- Compute time: buy compute cycles at internationally competitive rates with dedicated helpdesk included.
- User support: optimize your computation work flow with our assistance. Software installs done at your request.
- Code optimizing & Software development. Performance benchmarking.
- Training (MPI, OpenMP, CUDA, Intel MIC). Frequent training events open to all users.
- Joint research/development in mathematical modelling and applied mathematics
- Provider of HPC and cloud resources
- Provider of fast and reliable networking
- Data analytics on HPC systems




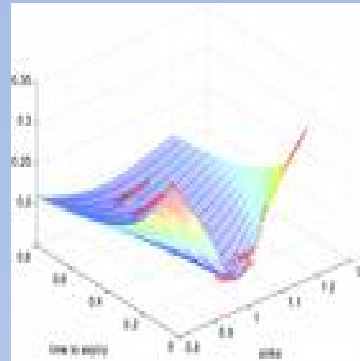
European Open Science Cloud services



Quantitative methods for problems in the field of financial mathematics, machine learning

Methods and algorithms for assessing the price of financial instruments and risk assessment. Calculate the parameters of stochastic models and model the movement of prices.

Development of IICT-BAS



Tempus – boosting finance technology performance and revolutionizing business operations through Machine Learning predictive models – project with Radioactive Zarko Asenov.

- **Network of HPC Centers, sharing expertise in support of SMEs**
- Bangor University, United Kingdom
- Irish High End Computing, National University of Galway, Ireland
- Fraunhofer Institute for Algorithms and Scientific Computing (SCAI), Germany
- Poznan Supercomputing and Networking Center IBCH (Inst. of Bioorganic Chemistry) PAS, Poland
- Greek Research and Technology Network S.A., Greece
- Centro de Supercomputación de Galicia CESGA, Spain
- PT Cloud, Portugal
- IT4Innovations National Supercomputing Center, VŠB – Techn. University of Ostrava, Czech Republic
- Yotta Advanced Computing d.o.o., Croatia
- West University of Timisoara UVT, Romania
- Inst. of Inform. and Comm. Technologies of BAS (IICT-BAS), Bulgaria
- Ruder Bošković Institute, Croatia
- Herculesstichting / Flemish Supercomputing Centre, Belgium
- Vilnius University – Faculty of Mathematics and Informatics, Lithuania





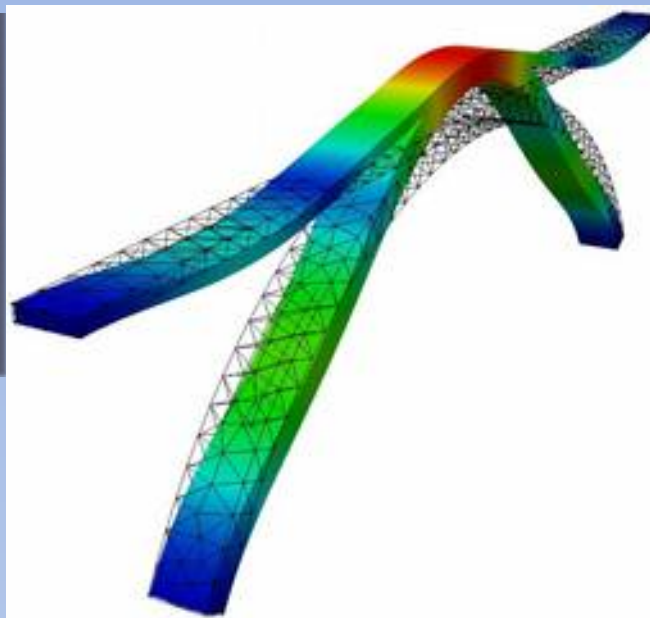



SESAME Net expands the use of the industry best practices, aligned with the European strategy for development of HPC infrastructure for use by SMEs.

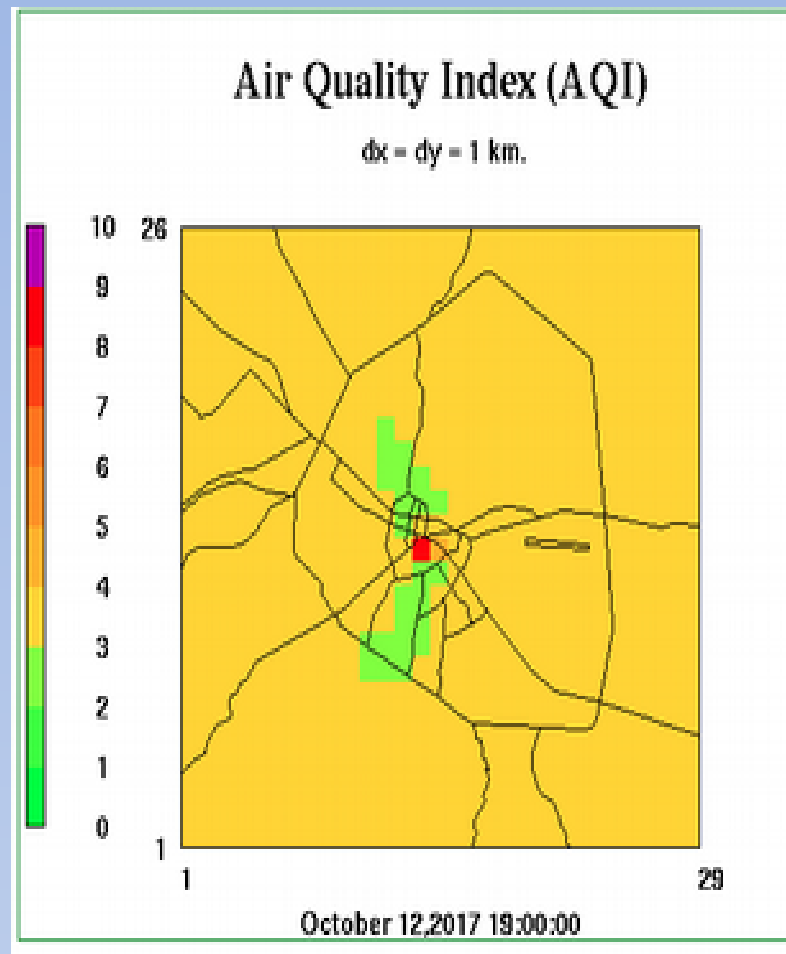
- IT4I – partner in SESAMENET
- Simulations:
- Crash Tests:
- Outcomes
 - Shortening the design cycle
 - Reducing costs
 - Adding value to existing products
 - Creating new products



IT4Innovations
national@IS#80
supercomputing
center@#01%101



Computations of the Air Quality Index

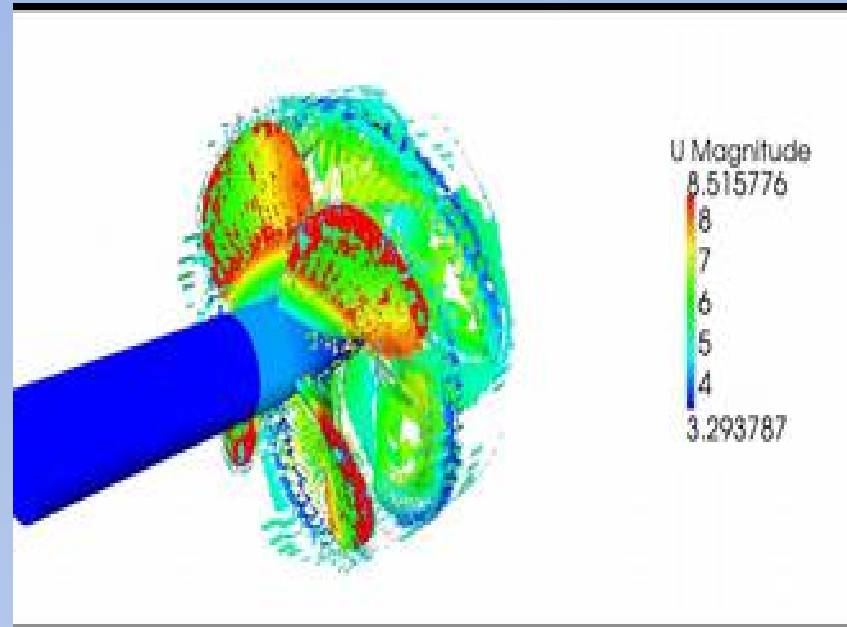


Air Quality Index produced with
the WRF Model

<http://niggg.bas.bg/cw3>

OpenFOAM



- C++ toolbox for the development of customized numerical solvers for solution of continuum mechanics problems, including computational fluid dynamics (CFD).
- Optimisation for Xeon Phi developed at IICT-BAS



underground reservoirs and oil and gas fields

New faster algorithms and software solutions for computer simulation of industrial flows.


Development of IICT-BAS



3D simulations for the study of liver tumor ablation

Presentation of the computing area and sampling using 3D high resolution medical imaging. Radio frequency exposure parameters are evaluated to maximize the reliability of the ablation.

*Joint development of IICT-BAS
and AMET LTD*



- Avitohol is only the first system of this caliber in Bulgaria
- Bulgaria joined the Euro HPC initiative
- Bulgaria is striving to maintain its place as leader in the region in the domain of HPC
- The Centre is actively participating in regional and European initiatives in the domain of electronic infrastructures
- IICT-BAS coordinates contract No: BG05M2OP001-1.001-0003-001 “CoE on Informatics and ICT” supported by the Operational Programme “Science and Education for Smart Growth” and is a partner in contract No BG05M2OP001-1.002-0002-C 01 „Digitalisation of the Economy in Big

