High Performance Computing I
(Fall 2020, Prof.dr. H. Wijshoff)

- Four parts:
  - Introduction to Parallel Programming and Parallel Architectures (just partly based on slides from Ananth Grama, Anshul Gupta, George Karypis, and Vipin Kumar accompanying ``Introduction to Parallel Computing'', Addison Wesley, 2003.)
  - Parallel Algorithms
    - Parallel Algorithm Design
    - Parallel Numerical Computing
    - Parallel Graph Computing
    - Parallel Sorting
  - Existing Programming Paradigms
  - New Programming Paradigms

Lab/Homework/Assignments

During the course of the semester, programming assignments.

These assignments make for 40% of the total load.

Open “book” final exam will make up for remainder 60%.

A Long History

- The advent of parallel computing dates back to the fifties of the last century
  - IBM introduced the 704 (full parallel floating point arithmetic) in 1954, through a project in which Gene Amdahl was one of the principal architects.
  - In April 1958, Stanley Gill (Ferranti Ltd, inventor of the subroutine) discussed parallel programming and the need for branching and waiting.
  - Also in 1958, IBM researchers John Cocke and Daniel Slotnick discussed the use of parallelism in numerical calculations for the first time.
- In 1969, US company Honeywell introduced its first Multics system, a symmetric multiprocessor system capable of running up to eight processors in parallel.
- The ILLIAC IV (1971) was one of the first attempts to build a massively parallel computer. One of a series of research machines (the ILLIACs from the University of Illinois), the ILLIAC IV design featured fairly high parallelism with up to 256 processors.

Milestones (the race to the bottom)

1972: First Supercomputer: CRAY 1

1 MFLOP = 1 000 000 operaties/sec.

1989: CRAY YMP

1 GFLOP = 1 000 000 000 operaties/sec.

1996: ASCI red (Intel based parallel processor)

1 TFLOP = 1 000 000 000 000 oper./sec.

2008: IBM Roadrunner

1 PFLOP = 1 000 000 000 000 000 oper./sec.

2020: Fugaku (Japan): 513 PFLOP (0.513 EXAFLOP)

The goal stays to realize more than 1 ExaFlop sustained! performance.

Liquid Cooled Cray 2

In December 2014 (http://top500.org/lists/2014/11)

RANK	SITE	SYSTEM	CORES	RMAX (TFLOP/S)	RPEAK (TFLOP/S)	POWER (KW)
1	National Super Computer Center in Guangzhou China	**Tianhe-2 (MilkyWay-2)** - TH-IVB-FEP Cluster, Intel Xeon E5-2692 12C 2.200GHz, TH Express-2, Intel Xeon Phi 31S1P NUDT	3,120,000	33,862.7	54,902.4	17,808
2	DOE/SC/Oak Ridge National Laboratory United States	**Titan** - Cray XK7 , Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x Cray Inc.	560,640	17,590.0	27,112.5	8,209
3	DOE/NNSA/LLNL United States	**Sequoia** - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom IBM	1,572,864	17,173.2	20,132.7	7,890
4	RIKEN Advanced Institute for Computational Science (AICS) Japan	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect Fujitsu	705,024	10,510.0	11,280.4	12,660
5	DOE/SC/Argonne National Laboratory United States	**Mira** - BlueGene/Q, Power BQC 16C 1.60GHz, Custom IBM	786,432	8,586.6	10,066.3	3,945
6	Swiss National Supercomputing Centre (CSCS) Switzerland	**Piz Daint** - Cray XC30, Xeon E5-2670 8C 2.600GHz, Aries interconnect , NVIDIA K20x Cray Inc.	115,984	6,271.0	7,788.9	2,325
7	Texas Advanced Computing Center/Univ. of Texas United States	**Stampede** - PowerEdge C8220, Xeon E5-2680 8C 2.700GHz, Infiniband FDR, Intel Xeon Phi SE10P Dell	462,462	5,168.1	8,520.1	4,510
8	Forschungszentrum Juelich (FZJ) Germany	**JUQUEEN** - BlueGene/Q, Power BQC 16C 1.600GHz, Custom Interconnect IBM	458,752	5,008.9	5,872.0	2,301

As of November 2015 (http://top500.org/lists/2015/11/)

RANK	SITE	SYSTEM	CORES	RMAX (TFLOP/S)	RPEAK (TFLOP/S)	POWER (KW)
1	National Super Computer Center in Guangzhou China	**Tianhe-2 (MilkyWay-2)** - TH-IVB-FEP Cluster, Intel Xeon E5-2692 12C 2.200GHz, TH Express-2, Intel Xeon Phi 31S1P NUDT	3,120,000	33,862.7	54,902.4	17,808
2	DOE/SC/Oak Ridge National Laboratory United States	**Titan** - Cray XK7 , Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x Cray Inc.	560,640	17,590.0	27,112.5	8,209
3	DOE/NNSA/LLNL United States	**Sequoia** - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom IBM	1,572,864	17,173.2	20,132.7	7,890
4	RIKEN Advanced Institute for Computational Science (AICS) Japan	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect Fujitsu	705,024	10,510.0	11,280.4	12,660
5	DOE/SC/Argonne National Laboratory United States	**Mira** - BlueGene/Q, Power BQC 16C 1.60GHz, Custom IBM	786,432	8,586.6	10,066.3	3,945
6	DOE/NNSA/LANL/SNL United States	**Trinity** - Cray XC40, Xeon E5-2698v3 16C 2.3GHz, Aries interconnect Cray Inc.	301,056	8,100.9	11,078.9	
7	Swiss National Supercomputing Centre (CSCS) Switzerland	**Piz Daint** - Cray XC30, Xeon E5-2670 8C 2.600GHz, Aries interconnect , NVIDIA K20x Cray Inc.	115,984	6,271.0	7,788.9	2,325
8	HLRS - Höchstleistungsrechenzentrum Stuttgart Germany	**Hazel Hen** - Cray XC40, Xeon E5-2680v3 12C 2.5GHz, Aries interconnect Cray Inc.	185,088	5,640.2	7,403.5	
9	King Abdullah University of Science and Technology Saudi Arabia	**Shaheen II** - Cray XC40, Xeon E5-2698v3 16C 2.3GHz, Aries interconnect Cray Inc.	196,608	5,537.0	7,235.2	2,834
10	Texas Advanced Computing Center/Univ. of Texas United States	**Stampede** - PowerEdge C8220, Xeon E5-2680 8C 2.700GHz, Infiniband FDR, Intel Xeon Phi SE10P Dell	462,462	5,168.1	8,520.1	4,510

As of June 2016 (http://top500.org/lists/2016/06/)

Rank	Site	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	National Supercomputing Center in Wuxi China	**Sunway TaihuLight** - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway NRCPC	10,649,600	93,014.6	125,435.9	15,371
2	National Super Computer Center in Guangzhou China	**Tianhe-2 (MilkyWay-2)** - TH-IVB-FEP Cluster, Intel Xeon E5-2692 12C 2.200GHz, TH Express-2, Intel Xeon Phi 31S1P NUDT	3,120,000	33,862.7	54,902.4	17,808
3	DOE/SC/Oak Ridge National Laboratory United States	**Titan** - Cray XK7 , Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x Cray Inc.	560,640	17,590.0	27,112.5	8,209
4	DOE/NNSA/LLNL United States	**Sequoia** - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom IBM	1,572,864	17,173.2	20,132.7	7,890
5	RIKEN Advanced Institute for Computational Science (AICS) Japan	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect Fujitsu	705,024	10,510.0	11,280.4	12,660
6	DOE/SC/Argonne National Laboratory United States	**Mira** - BlueGene/Q, Power BQC 16C 1.60GHz, Custom IBM	786,432	8,586.6	10,066.3	3,945
7	DOE/NNSA/LANL/SNL United States	**Trinity** - Cray XC40, Xeon E5-2698v3 16C 2.3GHz, Aries interconnect Cray Inc.	301,056	8,100.9	11,078.9	
8	Swiss National Supercomputing Centre (CSCS) Switzerland	**Piz Daint** - Cray XC30, Xeon E5-2670 8C 2.600GHz, Aries interconnect , NVIDIA K20x Cray Inc.	115,984	6,271.0	7,788.9	2,325
9	HLRS - Höchstleistungsrechenzentrum Stuttgart Germany	**Hazel Hen** - Cray XC40, Xeon E5-2680v3 12C 2.5GHz, Aries interconnect Cray Inc.	185,088	5,640.2	7,403.5	
10	King Abdullah University of Science and Technology Saudi Arabia	**Shaheen II** - Cray XC40, Xeon E5-2698v3 16C 2.3GHz, Aries interconnect Cray Inc.	196,608	5,537.0	7,235.2	2,834

Move Over, China: U.S. Is Again Home to World's Speediest Supercomputer

Summit, the world's fastest supercomputer, is made up of rows of black refrigerator-size units that weigh a total of 340 tons.

Shawn Poynter for The New York Times

Europe's Answer

Commission proposes to invest EUR 1 billion in world-class European supercomputers

Brussels, 11 January 2018

The European Commission unveiled today its plans to invest jointly with the Member States in building a world-class European supercomputers infrastructure.

Supercomputers are needed to process ever larger amounts of data and bring benefits tothe society in many areas from health care and renewable energy to car safety and cybersecurity.

Today's step is crucial for the EU's competitiveness and independence in the data economy. Today, European scientists and industry increasingly process their data outside the EU because their needs are not matched by the computation time or computer performance available in the EU. This lack of independence threatens privacy, data protection, commercial trade secrets, and ownership of data in particular for sensitive applications.

A new legal and funding structure – the EuroHPC Joint Undertaking – shall acquire, build and deploy across Europe a world-class High-Performance Computing (HPC) infrastructure. It will also support a research and innovation programme to develop the technologies and machines (hardware) as well as the applications (software) that would run on these supercomputers.

The EU's contribution in EuroHPC will be around EUR 486 million under the current Multiannual Financial Framework, matched by a similar amount from Member States and associated countries. Overall, around EUR 1 billion of public funding would be invested by 2020, and private members of the initiative would also add in kind contributions.

http://europa.eu/rapid/press-release_IP-18-64_en.htm

Current Status (2020)

The **EuroHPC Joint Undertaking** will support activities through procurement and open Calls in 2019 and 2020, and **will initially operate from 2019 to 2026**.
The EuroHPC JU foresees the **initial co-investment with Member States of about EUR 1 billion**, out of which EUR 486 million come from the actions already planned by the Commission in Horizon 2020 and Connecting Europe Facility (CEF) programmes in the current Multiannual Financial Framework (MFF). An additional EUR ~422 million will be contributed by private or industrial players in the form of in-kind contributions to the JU activities.
This level of funding will kick-start the activities to address the overall European HPC strategy, mainly: **to acquire two pre-exascale machines and several petascale systems by 2020**, and R&I actions covering the full HPC ecosystem launched in 2019 and 2020, including the support for HPC Competence Centres. A minimum of EUR 180 million will be dedicated to implement the R&I agenda with indirect actions.
Further funds in the next MFF (under discussion) would allow a full coverage of the HPC strategy, where we expect to have similar support from the Member States to **continuing the present EuroHPC JU in the 2021-2028 period, with the already agreed strategic objectives**.
These objectives include the acquisition in 2022-2023 of two exascale systems, at least one of them with European technology, one post-exascale system, networking and coordination of HPC Competence Centres, support for the first hybrid HPC / Quantum computing infrastructure in Europe, and coordination with the other digital priorities like Artificial Intelligence, Cybersecurity and digital skills. The Commission proposal foresees EUR 2.7 billion for HPC in the Digital Europe Programme, complemented with Research and Innovation actions for HPC and Big Data in the Horizon Europe programme.
The Commission has proposed to support EuroHPC in the next MFF with EUR 2.7 billion from the Digital Europe Programme (DEP) and with additional funds from Horizon Europe.

https://eurohpc-ju.europa.eu

As of June 2018 (http://top500.org/lists/2018/06/)

Rank	Site	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	DOE/SC/Oak Ridge National Laboratory United States	**Summit** - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband IBM	2,282,544	122,300.0	187,659.3	8,806
2	National Supercomputing Center in Wuxi China	**Sunway TaihuLight** - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway NRCPC	10,649,600	93,014.6	125,435.9	15,371
3	DOE/NNSA/LLNL United States	**Sierra** - IBM Power System S922LC, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband IBM	1,572,480	71,610.0	119,193.6	
4	National Super Computer Center in Guangzhou China	**Tianhe-2A** - TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000 NUDT	4,981,760	61,444.5	100,678.7	18,482

As of June 2019 (https://www.top500.org/lists/top500/2019/06/)

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	**Summit** - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148,600.0	200,794.9	10,096
2	**Sierra** - IBM Power System AC922, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM / NVIDIA / Mellanox DOE/NNSA/LLNL United States	1,572,480	94,640.0	125,712.0	7,438
3	**Sunway TaihuLight** - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway, NRCPC National Supercomputing Center in Wuxi China	10,649,600	93,014.6	125,435.9	15,371
4	**Tianhe-2A** - TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000, NUDT National Super Computer Center in Guangzhou China	4,981,760	61,444.5	100,678.7	18,482

148,600

94,640

93,014

As of June 2020 (https://www.top500.org/lists/top500/2020/06/)

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	**Supercomputer Fugaku** - Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D, Fujitsu RIKEN Center for Computational Science Japan	7,299,072	415,530.0	513,854.7	28,335
2	**Summit** - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148,600.0	200,794.9	10,096
3	**Sierra** - IBM Power System AC922, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM / NVIDIA / Mellanox DOE/NNSA/LLNL United States	1,572,480	94,640.0	125,712.0	7,438
4	**Sunway TaihuLight** - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway, NRCPC National Supercomputing Center in Wuxi China	10,649,600	93,014.6	125,435.9	15,371
5	**Tianhe-2A** - TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000, NUDT National Super Computer Center in Guangzhou China	4,981,760	61,444.5	100,678.7	18,482
6	**HPC5** - PowerEdge C4140, Xeon Gold 6252 24C 2.1GHz, NVIDIA Tesla V100, Mellanox HDR Infiniband, Dell EMC Eni S.p.A. Italy	669,760	35,450.0	51,720.8	2,252

Screenshot

Power Usage of Top 1 Supercomputer per TFLOP

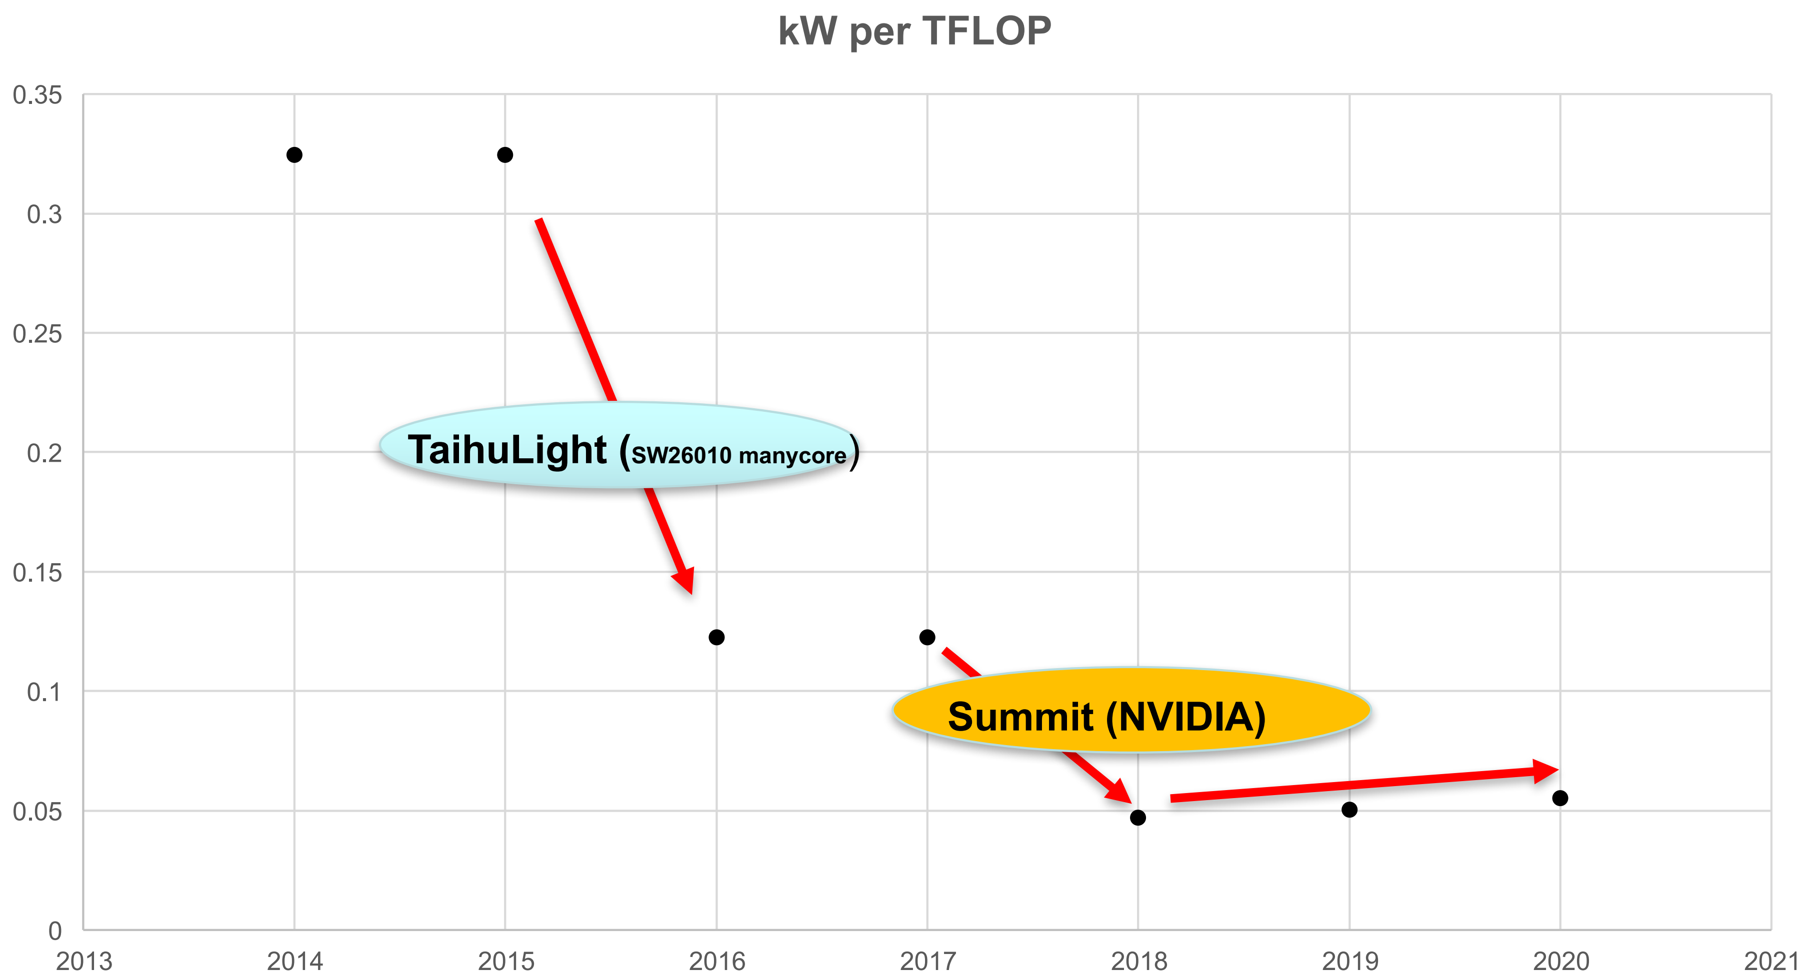

What does 18 MW mean?

Country	Population	Power per capita (W/p)*	18 MW equivalent in people
China	1.360.000.000	458	39.000
United States	318.000.000	1683	11.000
European Union	504.000.000	688	26.000
India	1.243.000.000	101	178.000
Netherlands	17.000.000	764	24.000
Syria	22.000.000	147	122.000
Afghanistan	30.000.000	1	18.000.000

*** Taken from https://en.wikipedia.org/wiki/List_of_countries_by_electricity_consumption**

What does 1 PFLOP mean?

Multiplying 2 numbers with 15 decimals

- Paper and Pencil: 1 per 4 minutes
- Calculator: 10 per minute based on 300 cpm (character per minute) (750 cpm world champion, Guinness, 2014)
- World Population of 7 000 000 000 ($7*10^9$), gives a speed of 70 000 000 000 per minute by calculators, which is equivalent to 70x60=4200 GFLOP

1 PFLOP is 1000000/4200 is more than 200 faster!

(And 1 ExaFLOP is more than 200000 faster!!!)

WHY do we need to compute at these rates?

Exponential growth of computational complexity

(Easy) example: CHESS

- Assume an average of 10 possible moves per turn
- Average chess match: 80 turns

So 10^80 different possible outcomes

With 1 PFLOP: 10^65 sec =

4x10^57 years =

4x10^54 centuries =

10^48 x the existence of the universe

WHY (II)

Large Scale of computations

Example: Weather Forecasting

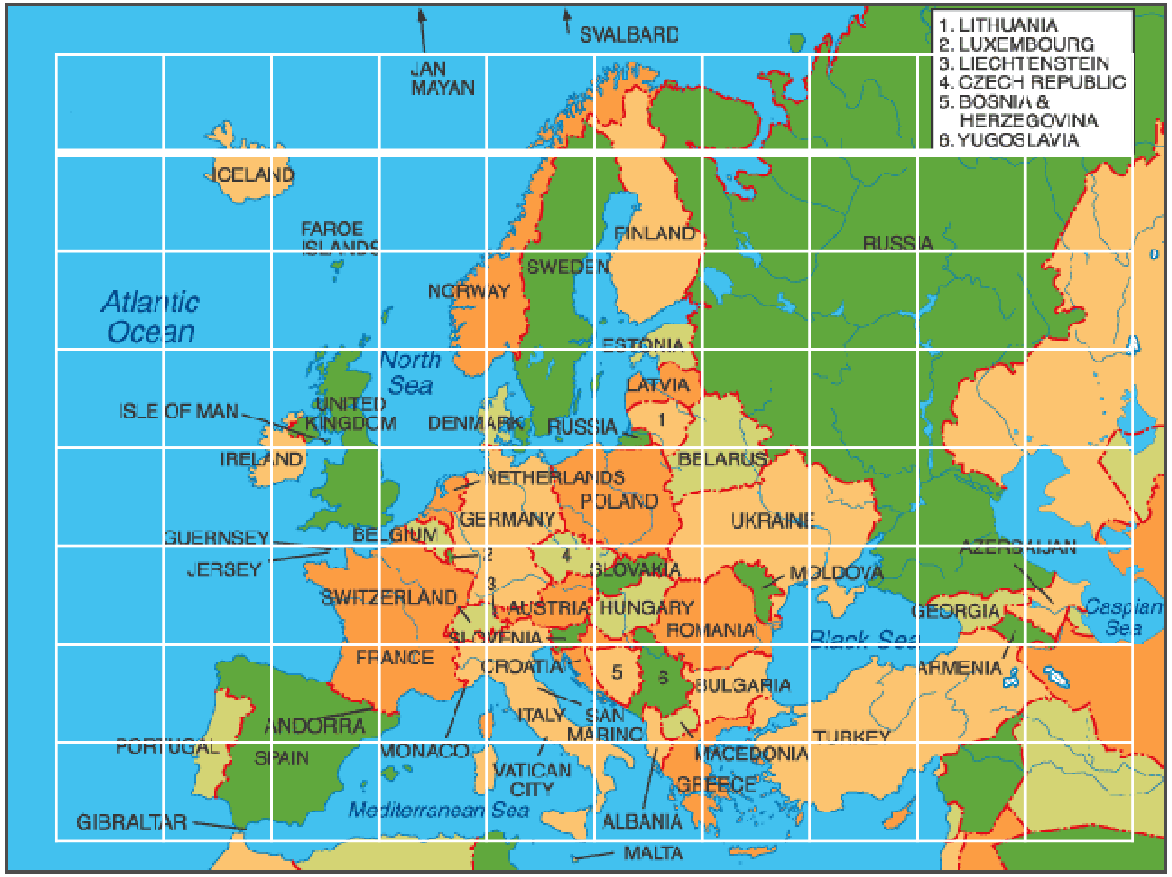

Computation (Simulation)

- For each grid point the interaction with its neighbor gridpoints are computed with respect to temperature, air pressure, moisture, etc
- Europe’s surface: 5 700 000 km^2
- Air heigth: 10 km
- With a 1m x 1m x 1m grid this results in:

  57 000 000 000 000 000 = 5.7 x 10^16 grid points

Computation (II)

- Several operations per grid point:

  Assume for each second and for 5 variables,

  then for a prediction of **12 hours**:

  5 x 12 x 60 x 60 = 216 000 operations per grid point
- With a 1 PFLOP computer this takes:

  5.7 x 10^16 X 216 x 10^3 / 10^15=

  12 x 10^21 / 10^15 = 12 x 10^6 sec. =

  3333 hours = **138 days** !!!!!!!!!

For a 1mm x 1mm x 1mm grid:

10^9 x 138 days = 380.000 centuries compute time required at 1 PFLOP rate for a 12 hours forecast

HPC Grand Challenges

"A Research and Development Strategy for High Performance Computing", Executive Office of the President, Office of Science and Technology Policy, November 20, 1987

- Prediction of weather, climate, and global change
- Challenges in materials sciences
- Semiconductor design
- Superconductivity
- Structural biology
- Design of pharmaceutical drugs
- Human genome
- Quantum chromodynamics
- Astronomy
- Challenges in Transportation
- Vehicle Signature
- Turbulence
- Vehicle dynamics
- Nuclear fusion
- Efficiency of combustion systems
- Enhanced oil and gas recovery
- Computational ocean sciences
- Speech
- Vision
- Undersea surveillance for anti-submarine warfare

Recently this list of applications was enlarged significantly

Next to applications in engineering and design, we have

- Scientific Applications: structural characterization of genes and proteins, new materials: understanding chemical pathways, bio-informatics and astrophysics, etc
- Commercial Applications: servers for large scale web servers (google, facebook, etc.), trading systems, etc.
- Applications in Computer Systems (the Internet itself): intrusion detection, cryptography, etc.
- Applications for social networks: online data mining,…
- Data Mining at large