Platform Parallel Netherlands

GPGPU-day, Amsterdam, 28 June 2012

Welcome to SARA

Dr. Peter Michielse

Deputy Director Policies & Strategy SARA

You are on historical ground !

CD Cyber 205 (1983)
1 CPU, 100 MFlop/s

- - - -

SGI Origin 3800 (2000)
1024 CPU's, 1 TFlop/s

SGI Altix 3700 (2003)
1400 CPU's, 3,2 TFlop/s

IBM Power5+ (2007)
1920 cores, 15 TFlop/s

IBM Power6 (2008)
3456 cores, 65 TFlop/s

factor 650.000
in 25 jaar
(Moore's Law)

The art of performance engineering

- In 1983, there was scalar and vector computing
- Scalar computing was standard, vector computing could significantly improve the performance through pipelining vector/array operations in the vector functional units

```
do i = 1,n
      y(i) = y(i) + a*x(i)
enddo
```

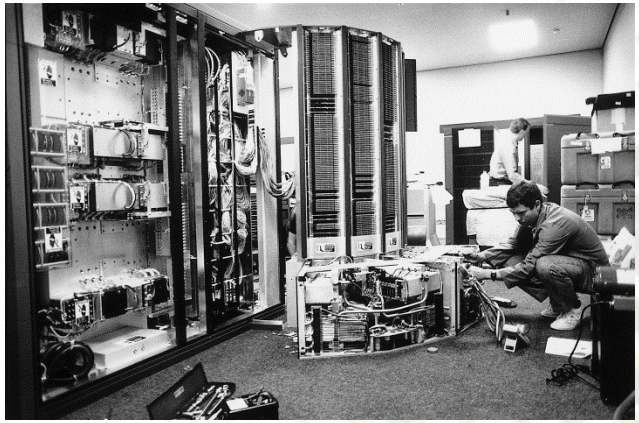

- In the early days, this did not come for free …..
- This may sound easy today, but specialists took care of the code and assisted the compiler in generating vector code

```
sdot = q8sdot(a(1;n),b(1;n))
```

The art of performance engineering

> After some time, one vector pipe was not enough anymore:

> More pipes and more processors, still shared memory

```
          subroutine cc_prod( a, b, n, c )
        real a(n), b(n), c, c_priv
cdir$ shared a(:block), b(:block), c, n
cdir$ master
        c = 0.0
cdir$ end master
        c_priv = 0.0
cdir$ do shared(i) on a(i)
        do i = 1, n
        c_priv = c_priv + a(i)*b
        end do
cdir$ atomic update
        c = c + c_priv
        return
        end
```

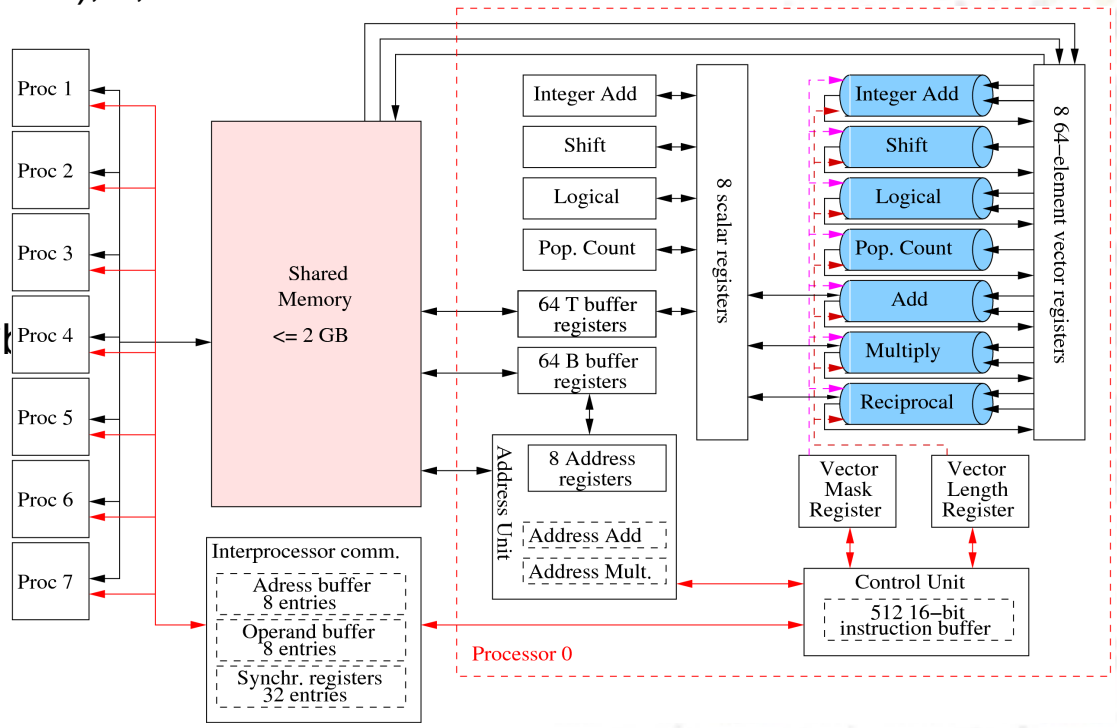

The art of performance engineering

- Massively parallel systems appeared, with 100's tot 1000's of processors (2000-now)
  - SMP, NUMA, ….
  - Both shared and distributed memory
  - Again the effort had to be made by the programmer with OpenMP and MPI

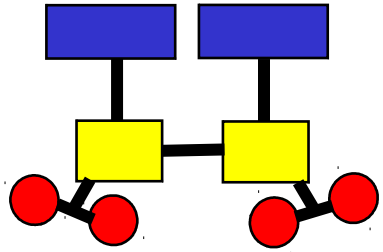

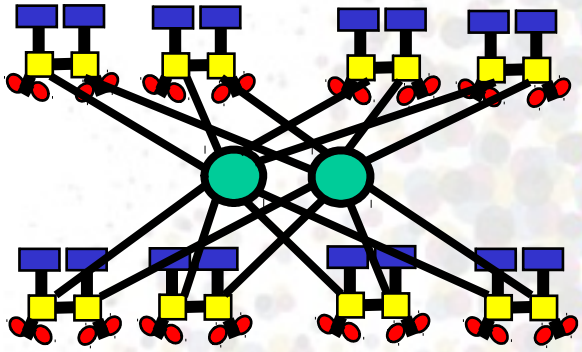

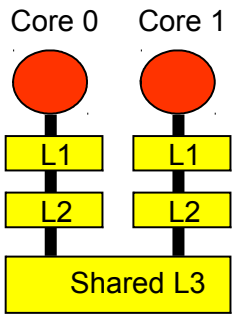

P6 dual-core chip (P6)

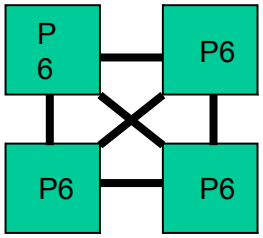

Quad Group (QG)

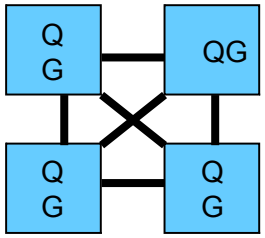

32-core Compute Node

Where are we now ?

- Facilities at SARA:
  - Capability HPC system (IBM Power6 HydroCluster)
  - Capacity system (Dell InfiniBand cluster)
  - HPC Cloud computing
  - Large-scale data processing
  - Hadoop services
  - Hybrid computing (GPU cluster)
  - Collaboratorium for advanced visualisation
- SARA experts and performance engineers are able to:
  - Translate the scientific wishes into implementations
  - Perform the actual programming work

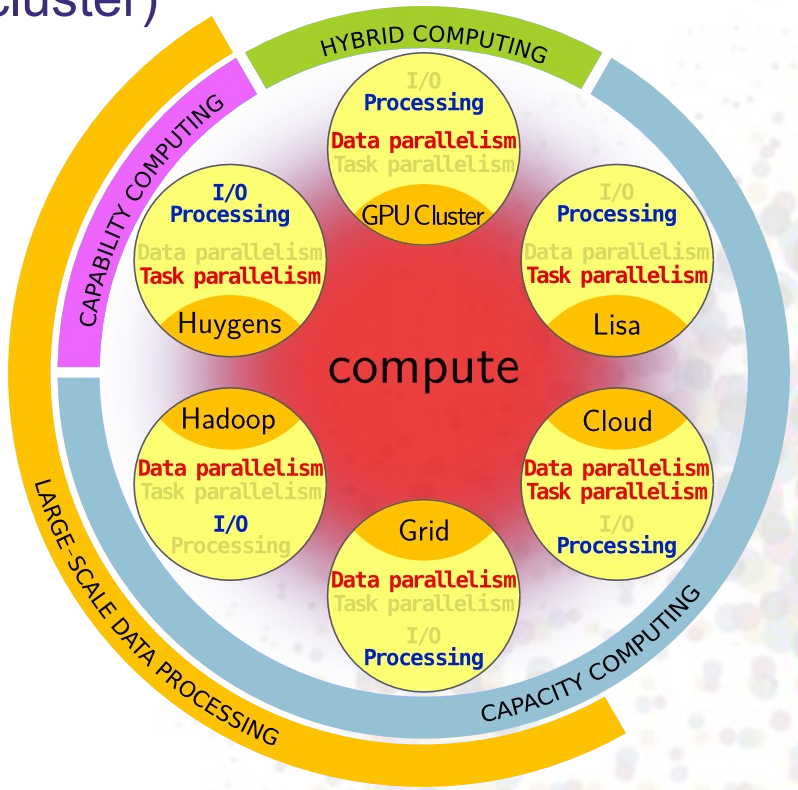

GPU cluster at SARA

- 8 HP dual quad core nodes, 2 NVIDIA Tesla GPUs/node
  - 16 Tflop/s
  - IB interconnect
  - CUDA environment
  - GPU code: C-like, CPU code: C+++
  - OpenCL
  - GPU code: own C version, CPU code: native C compiler (Intel / GCC)
  - (OpenGL / Assembler)
  - For pilot experiments and production runs
  - Size and growth with demands
  - Extensive support, training and guidelines
- Programming …..

Towards exascale computing …

Begin Full System Delivery (Yr)	2004	2007	2012	2015	2019
Design Parameters	**BG/L**	**BG/P**	**25PF**	**300PF**	**1200PF**
Cores / Node	2	4	8-24	32-64-128	96-128-500
Clock Speed (GHz)	0.7	0.85	1.6-4.1	2.3-4.8	2.8-6.0
Flops / Clock / Core	4	4	8-32	8-32	16-64
Nodes / Rack	1024	1024	100-1024	256-1024	256-1024
Racks / Full System Config	64	72	128-350	128-400	256-400
MB RAM/core	256	512	1024-4096	1024-4096	1024-4096
Total Power	2.5MW	4.8MW	8MW-20MW	20MW-50MW	30MW-80MW
Flops / Node (GF)	5.6	14	128-640	640-2000	2000-6000
Flops / Rack (TF)	5.7	14	200-400	400-1200	1600-4800
LB Concurrency	5.E+05	1.E+06	1M-2M	10M-100M	400M-1000M
Full System					
Total Cores (Millions)	0.13	0.3	.3M-1.2M	1M-10M	4M-30M
Total RAM (TB)	33.6	151	2,000-4,400	3,000-10,000	5,000-25,000
Total Racks	64	72	128-350	128-400	256-400
Peak Flops System (PF)	**0.37**	**1**	**25**	**300**	**1200**

Will accelerators be required to reach exascale within a reasonable power envelope ?

Will the programming models become “easy” again ?

Or better: what has changed since vector computing ?

The challenge is yours …. !

On behalf of SARA, I wish you a successful meeting !