

Empowered by Innovation **NEC**

NEC SX vectorization workshop

Jens-Olaf Beismann
Senior Benchmarking Analyst
NEC Deutschland GmbH



Objectives ...

```

***** Program Information *****
Real Time (sec)      : 274.312165
User Time (sec)     : 272.141663
System Time (sec)   : 1.473929
Vector Time (sec)   : 257.934811
Instruction Count    : 38211119460
Vector Instruction Count : 8688304757
Vector Element Count : 2057078964016
FLOP Count          : 890893604266
MOPS                : 7667.336768
MFLOPS              : 3273.639157
Average Vector Length : 236.764135
Vector Op. Ratio (%) : 98.585125
Memory size used (MB) : 1497.01942
MIPS                : 140.408929
Inst. Cache miss (sec) : 0.229672
Operand Cache miss (sec) : 2.678918
Bank Conflict Time (sec) : 9.299637

Start Time (date)   : 2004/08/12 05:25:08
End Time (date)    : 2004/08/12 05:29:46

```

Page 2

© NEC Deutschland GmbH 2015

Empowered by Innovation **NEC**

Agenda

DAY 1

- 9:30 – 11:00 Overview and basics on vectorization
11:00 – 11:30 Coffee break
- 11:30 – 13:00 Compilers, options and performance analysis
13:00 – 14:00 Lunch break
- 14:00 – 15:30 Optimization and vectorization examples,
indirect addressing
- 16:00 – 17:30 Run-time environment, Parallelization on SX

DAY 2

- 9:00 – 10:30 NEC SX-ACE: Architecture and new features
10:30 – 11:00 Coffee break
- 11:00 – 12:30 SX-ACE usage at HLRS
12:30 – 13:30 Lunch break
- 13:30 – 16:30 Hands-on session


Page 3

© NEC Deutschland GmbH 2015

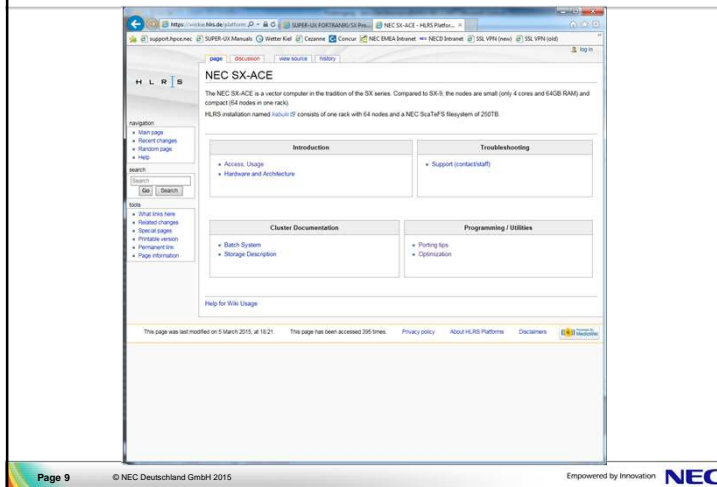
Empowered by Innovation

NEC

SX History and Technical Evolutions



SX Documentation



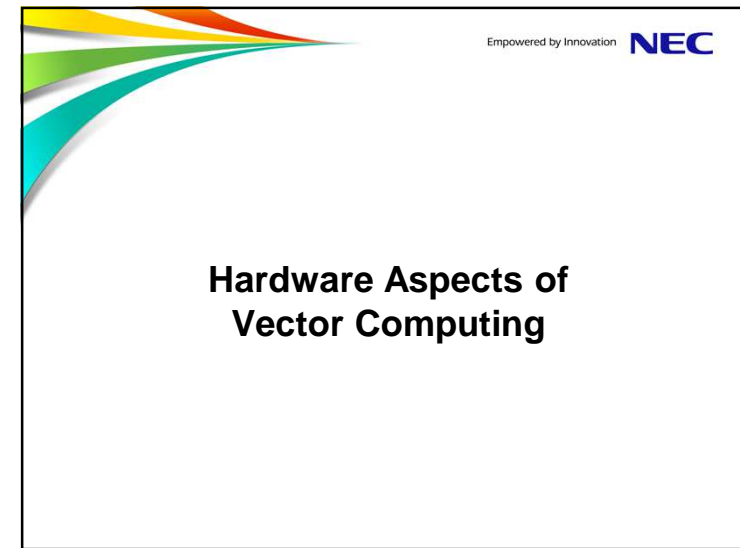
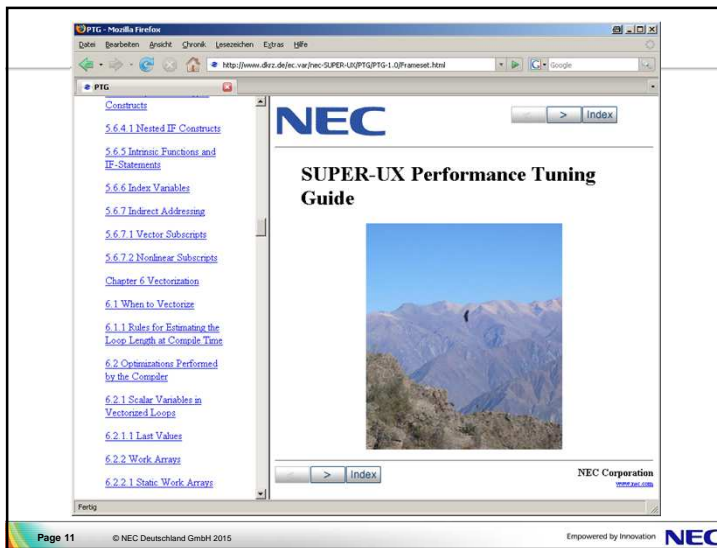
Programming languages

NEC

- FORTRAN90/SX Programmer's Guide
- FORTRAN90/SX Language Reference Manual
- C++/SX Programmer's Guide

Vectorization

- SUPER-UX Performance Tuning Guide
- FORTRAN90/SX Programmer's Guide



How do Vector Computers work?

Operation Segmentation:

Operations are decomposed into segments

Example: addition of floating point number

- compare exponent
- shift mantissa
- add mantissa
- select exponent and normalize

1.14e9 - 2.78e8
 1.14e9 - 0.278e9
 0.862e9
 8.62e8

Page 13

© NEC Deutschland GmbH 2015


Empowered by Innovation

NEC


How do Vector Computers work?

Pipelining:

- without Pipelining



- with Pipelining



Page 14

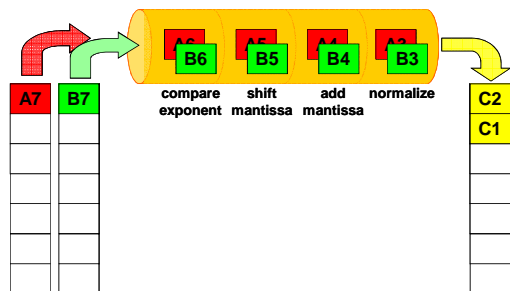
© NEC Deutschland GmbH 2015

Empowered by Innovation

NEC

How do Vector Computers work?

Vector "Add" Pipeline:



Page 15

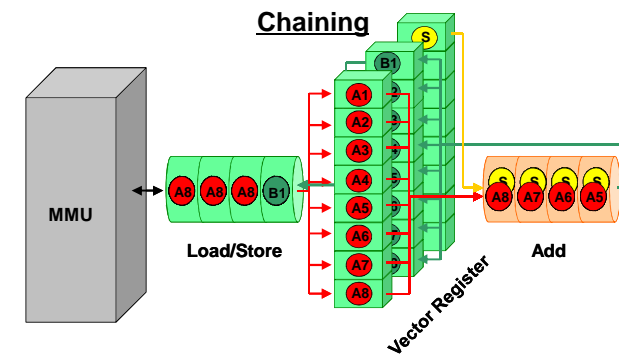
© NEC Deutschland GmbH 2015

Empowered by Innovation

NEC

How do Vector Computers work?

Chaining

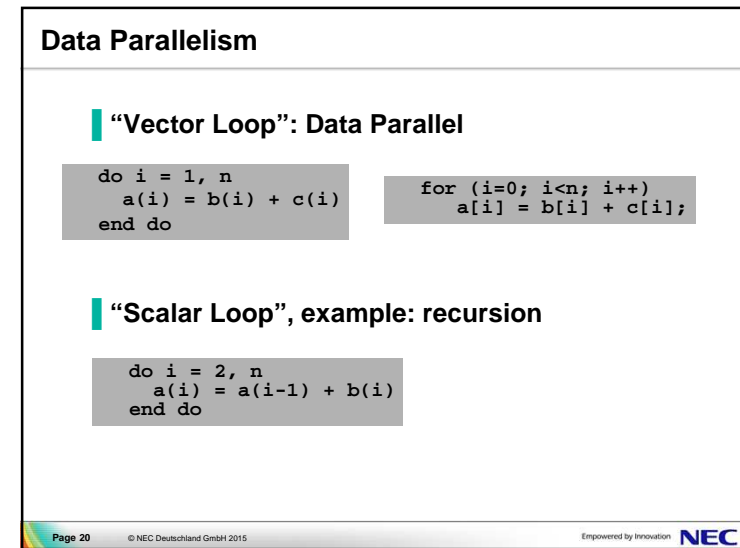
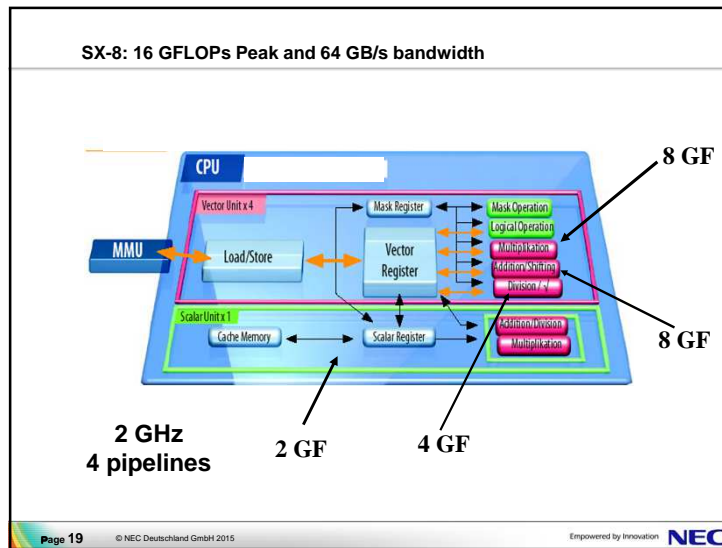
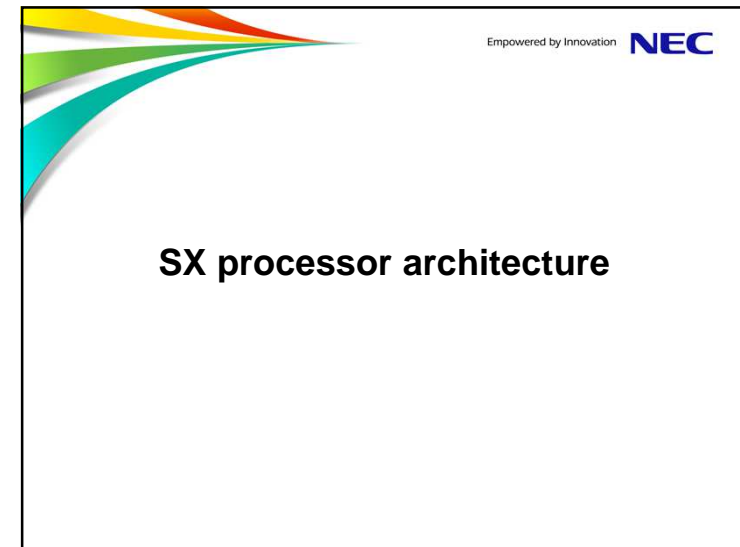
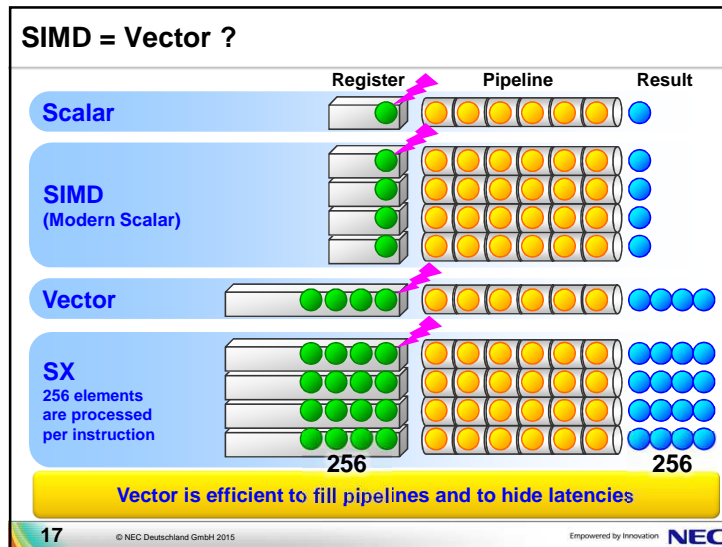


Page 16

© NEC Deutschland GmbH 2015

Empowered by Innovation

NEC



Levels of Parallelism


Segmentation



multiple pipes



parallel usage of functional units



parallel CPUs



parallel nodes



Levels of Parallelism

Segmentation

multiple pipes

parallel usage of functional units

parallel CPUs

parallel nodes


Efficiency:

$$\frac{\text{used units}}{\text{available units}}$$


To achieve high efficiency:

- Keep all units busy

Parallel programming on SX architecture



CPU level
• Vectorisation



Shared Memory level
• Microtasking
• Macrotasking
• Autotasking
• OpenMP
• MPI

Distributed Memory level
• MPI
• Co-array Fortran

Hybrid

Thinking Vector ?

'Scalar' thinking / coding:

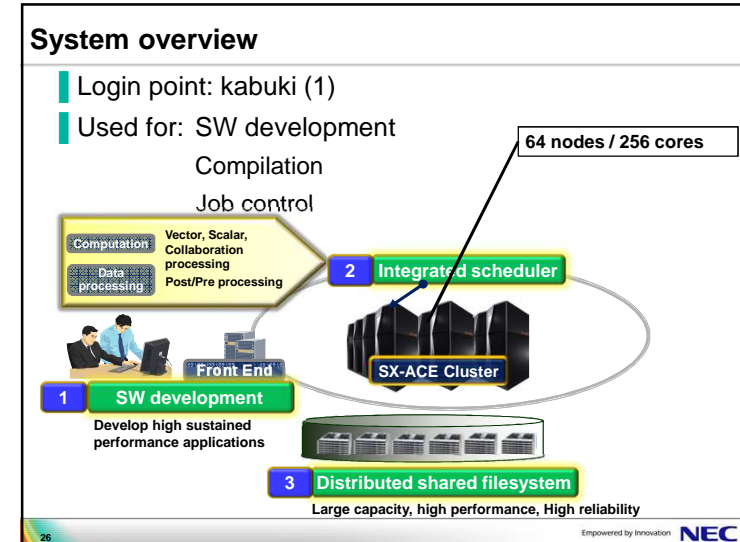
*There is a (grid-point,particle,equation,element),
what am I going to do with it?*

'Vector' thinking / coding:

*There is a certain action or operation, to which
(grid-points,particles,equations,elements)
am I going to apply it simultaneously?*

Empowered by Innovation **NEC**

First steps: compiling a code, submitting a job



Basic steps on the SX

Don't use the SX for compiling and linking !
Use cross compilers on frontend (**kabuki**)

<code>sxc++</code>	C++ and C Compiler
<code>sxf90</code>	F90 Compiler
<code>sxar</code>	ar archiver
<code>sxmpic++</code>	MPI C++/C Compiler
<code>sxmpif90</code>	MPI F90 Compiler
<code>sxftrace</code>	Analyzing Trace files

Page 27 © NEC Deutschland GmbH 2015 Empowered by Innovation **NEC**

FORTTRAN90/SX

Syntax :

```
sxf90 <command line options> \
    -Wf,<detailed optimisation options> \
    -Wf,<detailed vectorisation/parallelisation options> \
    -Wf,<miscellaneous options> \
    -Wl,<linker options> ...
```

Page 28 © NEC Deutschland GmbH 2015 Empowered by Innovation **NEC**

Basic set of compiler options

-R2	transformation and formatted listing
-Wf,-L summary	summary of compiler options
-pi	enable automatic inlining
-Chopt	highest optimisation level
-Wf,-pvctl fullmsg	full vectorisation diagnostics
-Wf,-pvctl noassume loopcnt=1000000	information about loop lengths

Command line options : optimization

F90 and C++

-C aopt	(only C++) aggressive vectorization and optimization
-C hopt	Full use of vectorization and optimization
-C vopt	Standard use of vectorization and optimization
-C sopt	Full use of optimization in scalar code
-C vsafe	Very safe use of vectorization and optimization
-C ssafe	Safe use of optimization in scalar code
-C noopt	(only C++) almost all optimizations and vectorization are suppressed
-C debug	No vectorization and optimization at all

Useful environment variables

```

F_PROGINF=[NO|YES|DETAIL]
    performance information from hardware counters

F_SETBUF=n
    set I/O buffers to n KB

F_FTRACE=[NO|FMT1|FMT2]
    flow trace performance information

F_FILEINF=[NO|YES|DETAIL]
    I/O statistics

Most environment variables also available as C_<varname> !

```

How to compile an MPI code

```

sxf90 code.f90 -o code
sxmpif90 mpicode.f90 -o mpicode

```

sxmpif90 does it all for you – no need to link or include anything !

NQSII : manipulating jobs, queue information, ...

On `kabuki` :

Some useful `NQSII` commands :


```

| qdel 1726           delete job with job ID 1726
| qstat -Qf smallque  get full information about settings of
                        queue "smallque"
| qcat -o 1726        see job output (log file contents) during
                        execution
| man qsub            access man pages for NQSII qsub
                        command

```

Job workflow

Arranging a set of jobs with dependencies...



NQSII workflow tools.

NQSII workflow tools

NQSII commands `wstart`, `wdel`, `wstat`
`qsub` (and options)
`qwait`, `qwait2`

Describe job workflow in a "workflow script"

```

#!/bin/bash

qsub -N TEST32n-1 ORCA025.L46-KBM01_inode.ksh
qsub -N TEST32n-2 ORCA025.L46-KBM01_inode.ksh --after TEST32n-1
qsub -N TEST32n-3 ORCA025.L46-KBM01_inode.ksh --after TEST32n-2
qwait2 TEST32n-3

```

NQSII workflow tools

```

nesh-fe : wstart wfl.sh >log.wfl 2>&1 &

nesh-fe : wstat
WFL-ID      Request Owner   Stat
-----
2_ace-ssiox 3 xeext098 RUN

nesh-fe : wstat 2
WFL-ID      RequestID RequestName Stat Exit
-----
2_ace-ssiox 78436.ace-ss TEST32n-1 RUN -
2_ace-ssiox 78437.ace-ss TEST32n-2 QUE -
2_ace-ssiox 78438.ace-ss TEST32n-3 QUE -

nesh-fe : qstat
RequestID ReqName Username Queue Pri STT S Memory CPU Elapse R H M Jobs
-----
78436.ace-ssiox TEST32n- xeext098 smallque 0 RUN - 1.03G 9.75 65 Y Y Y 32
78437.ace-ssiox TEST32n- xeext098 smallque 0 QUE - 0.00B 0.00 0 Y Y Y 32
78438.ace-ssiox TEST32n- xeext098 smallque 0 QUE - 0.00B 0.00 0 Y Y Y 32

```

NQSII workflow tools

Simple cases :

```
kabuki : qsub job1 job2 job3 job4 ... jobN
```

→ job chain (N consecutive jobs)

```
kabuki : qsub job1:job2:job3:job4:...:jobN
```

→ N parallel jobs running

```
kabuki : qsub --after <nqs-id> job1
```

→ start after job <nqs-id> has finished

Empowered by Innovation

NEC