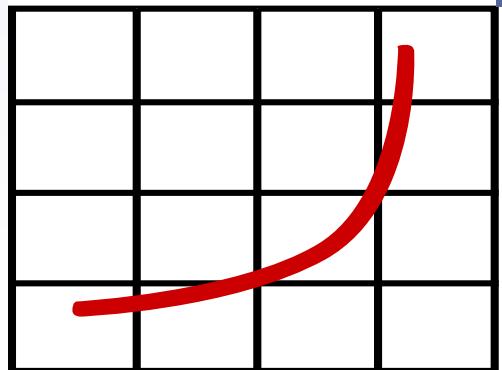


How to get, setup and run SPEC benchmarks

Swen Boehm, Mayara Gimenes, Robert Henschel, Veronica G.
Vergara Larrea, Junjie Li, Sandra Wienke



spec

<https://www.spec.org/hpg/publications>

<http://pages.iu.edu/~lijunj/sc19/>

Contents

- Cluster login 
- Overview of system requirements
- How to get SPEC benchmarks?
 - Benchmark acquisition & licensing
 - Download & unpacking
- How to setup SPEC benchmarks?
 - Installation 
- How to run SPEC benchmarks? 
 - Benchmark components & workloads
 - Runspec & run rules
 - Configuration files
 - From base to peak runs
 - Switch of compiler
- How to publish SPEC benchmark results?
 - Output files
 - Reportable runs
 - Process of publishing

Cluster login

- Read section “Login” in handout
- Follow instructions in section “Setup” in handout
- Then, stop.
- Follow along interactive demo

- Interactive demo time!
- We present SPEC Accel config files
- Opportunity to follow instructions interactively (also see handouts)
- Later: run benchmarks
Accel OpenACC on CPU and GPU
(Note: OpenMP and MPI runs take long – not covered here)
- Interpret results



If you have any problems, let us know immediately! We are happy to help you!

Contents

- Cluster login 
- Overview of system requirements
- How to get SPEC benchmarks?
 - Benchmark acquisition & licensing
 - Download & unpacking
- How to setup SPEC benchmarks?
 - Installation 

- How to run SPEC benchmarks? 
 - Benchmark components & workloads
 - Runspec & run rules
 - Configuration files
 - From base to peak runs
 - Switch of compiler
- How to publish SPEC benchmark results?
 - Output files
 - Reportable runs
 - Process of publishing

System Requirements

- Different benchmarks suites
→ different requirements
 - SPEC OMP2012, SPEC MPI2007, SPEC ACCEL
- Supported operating systems: AIX, Linux, MacOS, Solaris, Windows (except very old Windows)
 - Please do not use Windows/Unix compatibility products
- Compatible processors
 - CPU
 - GPU
 - APU
 - Xeon Phi

OpenMP: <http://spec.org/omp2012/Docs/system-requirements.html>
MPI: <http://spec.org/mpi2007/Docs/system-requirements.html>
ACCEL: <https://www.spec.org/accel/Docs/system-requirements.html>



System Requirements

- Memory requirements
 - OpenMP: 28GB for the whole system
 - MPI: 1GB/rank (medium size) and 2GB/rank (large size)
 - ACCEL: 4GB of host mem + 2GB of device mem
 - Otherwise, you are measuring your paging file, not your system
- Disk space requirements
 - OpenMP: 8GB
 - MPI: 10GB (medium), 17GB (large, big endian), 24GB (large, little endian)
 - ACCEL: 9GB
- Support of compilers
 - C99, C++98 and Fortran-95 compilers + MPI library for SPEC MPI 2007

OpenMP: <http://spec.org/omp2012/Docs/system-requirements.html>
MPI: <http://spec.org/mpi2007/Docs/system-requirements.html>
ACCEL: <https://www.spec.org/accel/Docs/system-requirements.html>

links 

Contents

- Cluster login 
- Overview of system requirements
- **How to get SPEC benchmarks?**
 - Benchmark acquisition & licensing
 - Download & unpacking
- How to setup SPEC benchmarks?
 - Installation 

- How to run SPEC benchmarks? 
 - Benchmark components & workloads
 - Runspec & run rules
 - Configuration files
 - From base to peak runs
 - Switch of compiler
- How to publish SPEC benchmark results?
 - Output files
 - Reportable runs
 - Process of publishing

Acquisition of SPEC Benchmarks

Standard Performance Evaluation Corporation

Home Benchmarks Tools Results Contact Site Map Search Help

[f](#) [in](#) [t](#) [g+](#)

Benchmarks

- Cloud
- CPU
- Graphics/Workstations
- ACCEL/MPI/OMP
- Java Client/Server
- Mail Servers
- Storage
- Power
- Virtualization
- Web Servers

Results Search

Submitting Results
Cloud/CPU/Java/Power
SFS/Virtualization
ACCEL/MPI/OMP
SPECcpu/SPECvleaper/SPECmpc

Tools

- SERT
- PTDaemon
- Chauffeur WDK

Order Benchmarks

- Order Form
- Downloads

SPEC

- About SPEC
- GWPG
- HPG
- OSG
- RG
- SC19

Purchase Current SPEC Benchmark Suites

SPEC benchmarking software is available via download through SWREG.

Benchmark Suite	Action	Price	Notes
ACCEL V1.2	Purchase (\$2000)	\$2000	Non-profit/educational organizations: request a free license
Chauffeur WDK V2.0.0	Purchase (\$50)	\$50	
Cloud IaaS 2016 V1.1	Purchase (\$2000)	\$2000	To purchase via download at the non-profit (\$500) pricing, contact the SPEC office for further information and to verify eligibility.
CPU2017 V1.0.2	Purchase (\$1000)	\$1000	To purchase via download at the upgrade (\$500) or non-profit (\$250) pricing, contact the SPEC office for further information and to verify eligibility.
CPU2006 V1.2	Purchase (\$800)	\$800	To purchase via download at the upgrade (\$400) or non-profit (\$200) pricing, contact the SPEC office for further information and to verify eligibility.
JBB2015 V1.01	Purchase (\$1500)	\$1500	To purchase via download at the non-profit (\$375) pricing, contact the SPEC office for further information and to verify eligibility.
jEnterprise2010 V1.03	Purchase (\$2000)	\$2000	To purchase via download at the non-profit (\$500) pricing, contact the SPEC office for further information and to verify eligibility.
MPI2007 V2.0.1	Purchase (\$2000)	\$2000	To purchase via download at the upgrade (\$250) pricing, contact the SPEC office for further information and to verify eligibility.
OMP2012 V1.0	Purchase (\$2000)	\$2000	Non-profit/educational organizations: request a free license
SEPT V2.0.1	Purchase (\$2000)	\$2000	Non-profit/educational organizations: request a free license

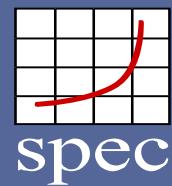
links 

<http://spec.org/order.html>
<https://www.spec.org/hpg/joining.html>

- Single SPEC suites
 - Commercial license
 - Non-profit license
 - SPEC membership
 - Receive benchmarks for free
- Payment Information
 - Upgrade Information
 - Non-profit/Educational Discounts
 - Tax/VAT
 - Site Licenses
 - Refunds
 - Purchasing Retired Benchmarks
 - SPECcpu Benchmark Requirements
 - MD5 checksums

Let's go shopping!

Non-profit organizations get 100% off



Press Release 2018: <https://www.spec.org/news/hpgnonprofitpricing.html>

Non-commercial download + definition:
<http://spec.org/hpgdownload.html>

links 

- Commercial license
 - Must be purchased via order form
 - Commercial enterprises (not being SPEC/HPG member) engaging in marketing, developing, testing, consulting for and/or selling computers, computer services, accelerator devices, drivers, software or other high performance computing systems or components in the computer marketplace
- Non-commercial license
 - Free of charge
 - Organizations that do not require a commercial license
 - Valid for the organization (not individual)
 - Institutional e-mail address required

Benchmark Suite	Non-Profit	Commercial
CPU2017 V1.0.2	\$250	\$1,000
ACCEL V1.3	free	\$2,000
OMP2012 V1.0	free	\$2,000
MPI2007 V2.0.1	free	\$2,000
SPECpower_ssj2008 V1.12	\$400	\$1,600

Download

Order Form

Currency: US Dollar ?

Locale: English ?

Product	Price	Qty	Delivery	Total
SPEC OMP2012 V1.0 - Retail download	US \$2380.00	1	Download ISO (1.6 GB)	US \$2380.00 including VAT (19%)
Add Extended Download Service for just US \$26.12 > What is this?				
			Subtotal	US \$2380.00 including US \$380.00 VAT (19%)
Coupon Code:			<input type="button" value="update cart"/>	

Billing Information
This is the address that your billing information is sent to.

Shipping Address is same as Billing Address

Email: *

First Name: *

Last Name: *

Payment Information

CREDIT CARD online now
 CREDIT CARD by PHONE
 PAYPAL
 WIRE TRANSFER with Proforma Invoice

Download as member

Index of /private/hpg/benchmarks/omp

Name	Last modified	Size
 Parent Directory		-
 omp2001-3.2.iso	27-Jul-2010 18:32	679M
 omp2001-3.2.iso.md5	06-Sep-2017 15:45	50
 omp2001-3.2.iso.sha256	06-Sep-2017 15:46	82
 omp2012-1.0.iso	17-Oct-2012 19:20	1.6G
 omp2012-1.0.iso.sha256	28-Feb-2018 16:16	82
 omp2012-1.0.iso.sha512	28-Feb-2018 16:24	146
 omp2012-1.0.iso.xz	17-Oct-2012 19:20	695M
 omp2012-1.0.iso.xz.md5	06-Sep-2017 15:46	53
 omp2012-1.0.iso.xz.sha256	06-Sep-2017 15:47	85

Apache/2.2.15 (CentOS) Server at pro.spec.org Port 443

Typically an ISO image

Unpacking (when you can mount ISO images)

```
$> md5sum -c accel-1.2.iso.xz.md5  
  
$> xz -d accel-1.2.iso.xz  
  
$> mkdir spec-iso  
  
$> mount -t iso9660 -o loop,ro accel-1.2.iso spec-iso
```

Check for correct download

Unpack archive

Mount ISO image in subdirectory

Useful hint:

Generate a subdirectory for every benchmark suite! Move the downloads there!

Unpacking (when you cannot mount ISO images)

- Use the tar.xz file (available to members and upon requests)

```
$> md5sum -c accel-1.2.tar.xz.md5
```

Check for correct download

```
$> tar xvJf accel-1.2.tar.xz
```

Unpack archive

- OR: Copy tar archive from the ISO

- Use tools such as `isoinfo` or `mc`

List files in iso image

```
$> isoinfo -J -l -i accel-1.2.iso
```

Extract md5 and tar ball from iso image
`-i`: iso image
`-x`: extracts to stdout, redirect needed

```
S> isoinfo -J -i accel-1.2.iso -x /install_archives/accel.tar.xz.md5 >
accel.tar.xz.md5
```

```
$> isoinfo -J -i accel-1.2.iso -x /install_archives/accel-1.2.tar.xz > accel-
1.2.tar.xz
```

Then: unpack tar ball (see above)

Contents

- Cluster login 
- Overview of system requirements
- How to get SPEC benchmarks?
 - Benchmark acquisition & licensing
 - Download & unpacking
- How to setup SPEC benchmarks?
- Installation 

- How to run SPEC benchmarks? 
- Benchmark components & workloads
- Runspec & run rules
- Configuration files
- From base to peak runs
- Switch of compiler
- How to publish SPEC benchmark results?
 - Output files
 - Reportable runs
 - Process of publishing

Installation

```
$> ./install.sh
```

Install SPEC suite
[-d <dest dir>]

SPEC ACCEL Installation



Top of the ACCEL tree is '/home/spec/accel-1.2'

Installing FROM /home/spec/accel-1.2

Installing TO /home/spec/accel-1.2

Is this correct? (Please enter 'yes' or 'no')
yes

Type yes and hit
enter

<http://spec.org/omp2012/Docs/install-guide-unix.html>

Installation

The following toolsets are expected to work on your platform. If the automatically installed one does not work, please re-run install.sh and exclude that toolset using the '-e' switch.

The toolset selected will not affect your benchmark scores.

linux-suse10-amd64

For 64-bit AMD64/EM64T Linux systems running SuSE Linux 10 or later, and other compatible Linux distributions, including some versions of RedHat Enterprise Linux and Oracle Linux Server.

Built on SuSE Linux 10 with
GCC v4.1.0 (SUSE Linux)

linux-redhat72-ia32

For x86, IA-64, EM64T, and AMD64-based Linux systems with GLIBC 2.2.4+.
Built on RedHat 7.2 (x86) with gcc 3.1.1

=====

Attempting to install the linux-suse10-amd64 toolset...

Attempt to
automatically
determine platform

<http://spec.org/omp2012/Docs/install-guide-unix.html>

links

Installation

```
=====
```

```
Attempting to install the linux-suse10-amd64 toolset...
```

```
Checking the integrity of your source tree...
```

```
Checksums are all okay.
```

```
Removing previous tools installation
```

```
Unpacking binary tools for linux-suse10-amd64...
```

```
Checking the integrity of your binary tools...
```

```
Checksums are all okay.
```

```
Testing the tools installation (this may take a minute)
```

Automatic
unpacking of files

Automatic testing of
installation

Loading SPEC tools

[..]

Installation successful. Source the shrc or cshrc in
/home/spec/accel-1.2
to set up your environment for the benchmark.

Hint how to proceed

```
>$ source ./shrc.sh
```

Source shrc or cshrc!
Without this nothing will work!!

Setup of environment variables
and paths for SPEC, e.g.,
\$SPEC to root path



<http://spec.org/omp2012/Docs/install-guide-unix.html>

Setup SPEC Accel on Summit

- Change the directory:
`cd $MEMBERWORK/trn001`
- Create a new directory to install the benchmark suite:
`mkdir accel && cd accel`
- Install the spec suite into your scratch space:
`/ccs/proj/trn001/accel-1.2/install.sh -d $PWD`
- Follow the prompt instructions

Setup SPEC Accel on Summit

- Config files provided in :

`/ccs/proj/trn001/sc19tut/config/*.cfg`

Copy into your install directories config directory

- Flag descriptions files:

`/ccs/proj/trn001/sc19tut/config/*.xml`

Copy into your install dir

- Job scripts:

`/ccs/proj/trn001/sc19tut/scripts/* .`

Copy into your install dir

We will run SPEC Accel while
I introduce the config file

Contents

- Cluster login 
- Overview of system requirements
- How to get SPEC benchmarks?
 - Benchmark acquisition & licensing
 - Download & unpacking
- How to setup SPEC benchmarks?
- Installation 

- How to run SPEC benchmarks? 
 - Benchmark components & workloads
 - Runspec & run rules
 - Configuration files
 - From base to peak runs
 - Switch of compiler
- How to publish SPEC benchmark results?
 - Output files
 - Reportable runs
 - Process of publishing

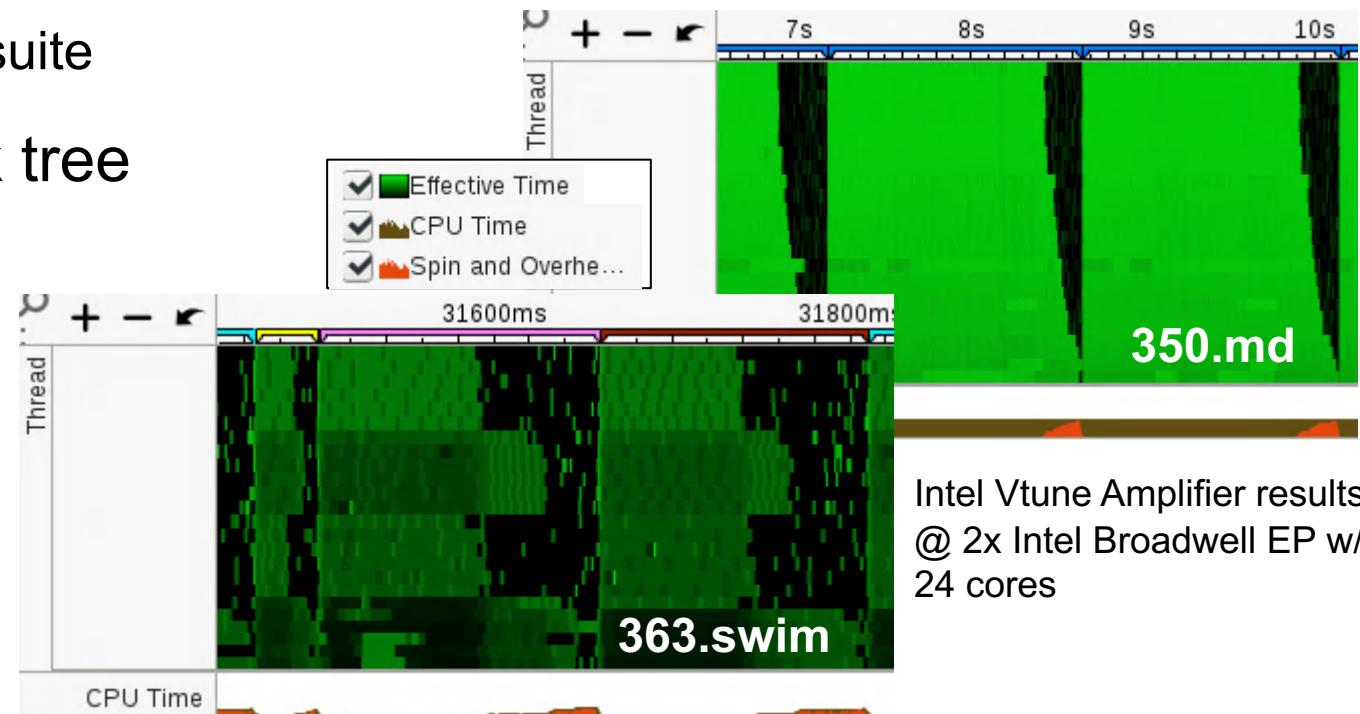
Benchmark components

- Coming from real-world applications
- May have completely different characteristics
- Identification: number and name
 - First number: affiliation to specific SPEC suite
- Documentation: website/ benchmark tree
 - Language, application domain,...

Description of benchmarks: <https://www.spec.org/omp2012/Docs/index.html>

links

Intel Vtune Amp.	350.md	363.swim
CPI rate	0.698	1.634
CPU utilization	69 %	23.6 %
Memory bound	11.8 %	48.5 %
% of packed FP instr	9.7 %	100 %



Workloads

Data set sizes

- **test**: data for a simple test that an executable is functional
 - **train**: data for feedback-directed optimization
 - **ref**: the real data set, required for all result reporting
- Runtime increases from a few seconds to tens of minutes per benchmark
(depending on your configuration)

SPEC Tools

Tools provided to ensure consistent operation of benchmarks across variety of platforms

- **specmake, specperl, specsha512sum, etc**
 - make, perl, sha512sum, etc.
- **runspec**
 - Primary tool in the suite
 - Used to build the benchmarks, run them, and report their results
 - Config file needed for usage (with detailed instructions on building/running the benchmarks)
- **rawformat**
 - Results formatter needed for publishing SPEC results
- And more

Tool overview: <http://spec.org/omp2012/Docs/tools-build.html>

runspec - Run SPEC benchmarks

```
$> . ./shrc
```

Reminder! Must be sourced before usage

```
$> runspec --config=openacc-pgi
```

Use specific config file

option: base & peak runs

--tune=base,peak 350

- Select benchmarks (suite or single ones)

- Specify config file

- Specify options

- E.g., data set size, iteration number, device number (for Accel), thread number (for OMP2012)

Run selected (single) benchmark: 350
Or: whole suite, i.e., all

- Specify actions
- E.g., compile single benchmarks, run them and validate them

Can overwrite parameters in config file

Runspec: <https://www.spec.org/accel/Docs/runspec.html>

Run rules

- Part of SPEC's philosophy
(see Part A "Overview of SPEC and SPEC HPG")
- Aim: fair and objective benchmarking
 - Published results are meaningful, comparable to other results, reproducible
- Public SPEC results must be compliant to these rules (license agreement)
 - or must be clearly described as *estimate*
- Estimates
 - May fail to provide one or more of the characteristics of public results
 - E.g., new chip design, prototype/ research compilers
 - Still encouraged to obey as many of the run rules as practical
 - Deviations from the rules must be clearly disclosed

OpenMP: <https://www.spec.org/omp2012/docs/runrules.html>
MPI: <https://www.spec.org/mpi/docs/runrules.html>
ACCEL: <https://www.spec.org/accel/docs/runrules.html>

links 

You can use SPEC results
as "estimates" in your
research environment
(also for publications)!

Run rules

Run rules basics:

- Same compiler flags for all in base run
- Individually tuned flags in peak run
- Three iterations needed for reportable run, medium is reported.
- Document all performance related software and hardware

Details see later

OpenMP: <https://www.spec.org/omp2012/docs/runrules.html>
MPI: <https://www.spec.org/mpi/docs/runrules.html>
ACCEL: <https://www.spec.org/accel/docs/runrules.html>
¹ For MPI2007: medium and large data sets

links 

Config Files

- Contain instructions for
 - Building benchmarks
 - Running them
 - Description of system under test



Key for reproducibility!

- How to write a config file?
 - Often start off using a config file that someone else has previously written¹
 - E.g. directory \$SPEC/config/
 - E.g., SPEC result submissions similar to your system²
 - Write your own³

Source: as of May 2018²

Test Sponsor	System Name	Base Threads	Processor				Results	
			Enabled Cores	Enabled Chips	Cores/Chip	Threads/Core	Base	Peak
Intel	Intel Server System R2208WFTZS (2 x Intel Xeon Platinum 8180, DDR4-2666, SMT ON Turbo ON)	112	56	2	28	2	21.1	25.5
Lenovo Global Technology	ThinkSystem SR950	224	112	4	28	2	40.2	Not Run
Oracle Corporation	SPARC T7-4	256	128	4	32	8	26.4	27.9
RWTH University Aachen	NEC HPC 1812Rg	24	24	2	12	2	7.22	Not Run
RWTH University Aachen	NEC HPC 1812Rg	48	24	2	12	2	7.65	Not Run
RWTH University Aachen	NEC HPC 1812Rg	14	144	8	18	2	33.6	Not Run
RWTH University Aachen	NEC HPC 1812Rg	72	144	8	18	2	19.5	Not Run
RWTH University Aachen	NEC HPC 1812Rg	36	144	8	18	2	9.97	Not Run
RWTH University Aachen	NEC HPC 1812Rg	18	144	8	18	2	5.23	Not Run

¹ https://www.spec.org/accel/docs/runspec.html#about_config

² <https://www.spec.org/omp2012/results/omp2012.html>

³ <https://www.spec.org/accel/docs/config.html>



Config Files

Use Case:

How to investigate single-GPU performance with SPEC ACCEL?

What we will investigate (as an outlook)...

- Basics on writing config files
- SPEC OpenACC with PGI compiler

Later: Have some (hands-on) time to run SPEC benchmarks

Run Accel OpenACC on GPU & CPU on Summit

Not covered here

- SPEC OpenCL
- SPEC OpenMP with target offloading

Structure of Config Files

Header section

- Usually runspec flags

```
#####
# what to do: build, validate = build+ runt+ check+ report
action      = validate
# Number of iterations of a test
iterations   = 1
# Tuning levels: base, peak
tune        = base
# Dataset size: test, train, ref
size         = test
# Environment variable will be set using "ENV_"
env_vars    = 1
# Output format: all, pdf, text, html and so on
output_format = text
flagsurl    = ${top}/../flagsfile/Intel.xml
teeout       = yes

# Run benchmarks according to your specific system config
# The variable "command" is the command used by spec
# submit = <system command> $command

#submit = numactl -p 1 $command
#submit = aprun -n 1 $command
```

Compiler Specification

```
#####
# Software information
#####
# Compilers. Using PGI compiler for example
default=default=default:
CC          = pgcc
CXX         = pgc++
FC          = pgfortran

# Environment variables at runtime
ENV_PGI_ACC_BUFFERSIZE = 8M
```

Compiler Flags

- Base & Peak (optional)
- Portability

```
#####
# Base
#####

openacc=base=default=default:
OPTIMIZE      = -fast -Mfprelaxed
FOPTIMIZE     = -acc -ta=tesla:cc35,cuda5.5
COPTIMIZE     = -acc -ta=tesla:cc35,cuda5.5

#####
# Portability flags for each benchmark
# Following flag should not have any impact on performance.
#####

359.miniGhost=default=default:
EXTRA_LDFLAGS += -Mno_main

#####
# Peak
#####

350.md=peak=default=default:
FOPTIMIZE = -acc -ta=tesla:cc35,cuda5.5,maxregcount:48

352.ep=peak=default=default:
FOPTIMIZE = -acc -ta=tesla:cc35,llvm

359.miniGhost=peak=default=default:
FOPTIMIZE = -acc -ta=tesla:cc35,cuda5.5,maxregcount:32
COPTIMIZE   = -acc -ta=tesla:cc35,cuda5.5,maxregcount:32

#####
# Hardware and software information for the machine under test.
# This information will be extracted for a reportable run.
# An example configuration can be copied from the website
# https://www.spec.org/accel/results/accel_acc.html
#####

company_name = SPEC Tutorial Company
test_sponsor = SPEC Tutorial Sponsor
tester       = SPEC Tutorial Tester
license_num  = SPEC Tutorial License
machine_name = SPEC Tutorial Machine
```

Host Information

```
#####
# HOST Hardware information
#####
# default=default=default=default:
hw_avail = Nov-2013
hw_cpu_name = Intel Core i7-3930K
hw_cpu_mhz = 3200
hw_cpu_max_mhz = 3800
hw_fpu = Integrated
hw_nchips = 1
hw_ncores = 6
hw_ncoresperchip = 6
hw_ntreadspercore = 2
hw_ncpuorder = 1 chip
hw_pcache = 32 KB I + 32 KB D on chip per core
hw_scache = 256 KB I+D on chip per core
```

Accelerator Information

```
#####
# Accelerator Hardware information
#####

hw_accel_model = Tesla K40c
hw_accel_vendor = NVIDIA
hw_accel_name = NVIDIA Tesla K40c
hw_accel_type = GPU
hw_accel_connect = PCIe 3.0 16x
hw_accel_ecc = Yes
hw_accel_desc00 = GPU Boost set to use a graphic clock frequency
hw_accel_desc001 = of 810 MHz. See notes below.
```

MD5 section

- Automatically-generated

```
#####
# MD5 section. It will be created by SPEC automatically.
# It is used by SPEC to check whether an executable if
# available is created using the current compiler and flags
# settings.
#####
```

SPEC Accel Config File (1/8)

config file: \$SPEC/config/openacc-pgi.cfg

```
#####
# The header section of the config file. Must appear
# before any instances of "default="
#####

# what to do: build, validate = build + run + check + report
action      = validate

# Number of iterations of a test
iterations   = 1

# Tuning levels: base, peak
tune        = base

# Dataset size: test, train, ref
size        = test
```

- **build**: compile benchmarks
- **validate**: benchmarks are built if necessary, run and reports are generated



- How many times to run each benchmarks
- e.g. for reportable run = 3

- **base**: flags for all benchmarks the same
- **peak**: set of optimizations individually selected for that benchmark

- Data set sizes (from small to big): test, train, ref
- e.g. test for debugging new set of compilation options

SPEC Accel Config File (2/8)

config file: \$SPEC/config/openacc-pgi.cfg

```
#####
# The header section of the config file. Must appear
# before any instances of "default="
#####
```

```
# which accelerator to use
device      = 0
```

- Use the 1st accelerator device (optional)

- Environment settings
- ENV_VAR = ...
- Apply to build phase → rebuild if any changes

```
# Environment variable will be set using "ENV_*", see the next section
```

```
env_vars     = 1
```

- Different output formats possible
- In \$SPEC/results

```
# Output format: all, pdf, text, html and so on
```

```
output_format = text
flagsurl      = ${top}/../flagsfile/Intel.xml
teeout        = yes
```

Displays the build commands to screen



hands-on

- Description of portability & tuning options (“Flags File”)
- Information on syntax of flags and their meanings
- Needed for valid reports
- Flags file: <http://spec.org/accel/Docs/flag-description.html>

SPEC Accel Config File (3/8)

config file: \$SPEC/config/openacc-pgi.cfg

```
# How to run the benchmarks according to your specific system configuration
# The variable "command" is the command used by spec
# submit = <system command> $command
# e.g., submit = aprun -n 1 $command
# e.g., submit = taskset -c 0-23 $command
# e.g., submit = numactl -n 1 $command
```

- **How to execute the benchmarks**
- Use \$command for SPEC command
- Preferred to assign work to processors
 - May place benchmarks on desired processors or benchmark memory on a desired memory unit
 - Especially needed for MPI runs
 - Example 1: run job on one node
 - Example 2: assign job to cores
- Can be used to change the run time environment
submit = export ENV_VAR=...; ...
→ no rebuild if any changes occur

SPEC Accel Config File (4/8)

config file: \$SPEC/config/openacc-pgi.cfg

```
#####
# Software information
#####

# Compilers. Using Intel compiler for example
default=default=default=default:
CC          = pgcc
CXX         = pgc++
FC          = pgfortran

#####
# Base
#####

openacc=base=default=default:
OPTIMIZE    = -fast -Mfprelaxed
FOPTIMIZE   = -acc -ta=tesla:cc70,cuda10.1
COPTIMIZE   = -acc -ta=tesla:cc70,cuda10.1
```

Named section

benchmark[, ...]=tuning[, ...]
=extension[, ...]=machine[, ...] :



hands-on

- Settings for base runs
- Compiler flags or environment variables

SPEC Accel Config File (5/8)

config file: \$SPEC/config/openacc-pgi.cfg

```
# Environment variables at runtime
ENV_PGI_ACC_BUFFERSIZE = 8M
#####
# Portability flags for each benchmark
# Following flag should not have any impact on performance.
#####
```

Environment variables

- Active if `env_vars` is set to 1 (see prior slides)
- Need to start with `ENV_`



SPEC Accel OpenACC current doesn't need portability flag with PGI.

Set portability flags if benchmark cannot be built/ execute correctly w/o these flags.

Other typical portability flags: language standards (`-std=c99`), Fortarn format (`-free`), memory models (`-mcmodel=medium`)

SPEC Config File: portability flags



hands-on

You will likely need different portability flags when you experiment with a different compiler.

Example here is portability flag of SPEC OMP2012.

```

# Environment variables at runtime
ENV_OMP_PROC_BIND = close
ENV_OMP_PLACES    = cores
ENV_OMP_NESTED    = FALSE
ENV_OMP_DYNAMIC   = FALSE

#####
####
# Portability flags for each benchmark
# Following flag should not have any impact on performance.
#####
####
350.md=default=default=default:
FPORTABILITY = -free

357.bt331=default=default=default:
PORTABILITY  = -mcmodel=medium

363.swim=default=default=default:
PORTABILITY  = -mcmodel=medium

367.imagick=default=default=default:
CPORTABILITY = -std=c99

```

Set portability flags if benchmark cannot be built/ execute correctly w/o these flags

SPEC Accel Config File (6/8)

config file: \$SPEC/config/openacc-pgi.cfg

```
#####
# Hardware and software information for the machine under test.
# This information will be extracted for a reportable run.
# An example configuration can be copied from the website
# https://www.spec.org/accel/results/accel_acc.html
#####
```

company_name	= SPEC Tutorial Company
test_sponsor	= SPEC Tutorial Sponsor
tester	= SPEC Tutorial Tester
license_num	= SPEC Tutorial License
machine_name	= SPEC Tutorial Machine

hw_vendor	= NEC
hw_avail	= NOV-2016
hw_cpu	= Intel Xeon E5-2650 v4
hw_cpu_mhz	= 2200
hw_cpu_max_mhz	= 2900

HW & SW description

- Needed only for reportable runs
- `runspec` tools captures information in submission file
- Very detailed information

Information on host configuration, e.g. CPU



hands-on

SPEC Accel Config File (7/8)

config file: \$SPEC/config/openacc-pgi.cfg

```
#####
# HOST Hardware information
#####
# default=default=default=default:
hw_avail = Nov-2013
hw_cpu_name = Intel Core i7-3930K
hw_cpu_mhz = 3200
hw_cpu_max_mhz = 3800
hw_fpu = Integrated
hw_nchips = 1
hw_ncores = 6
hw_ncoresperchip = 6
hw_ntreadspercore = 2
hw_ncpuorder = 1 chip
hw_pcache = 32 KB I + 32 KB D on chip per core
hw_scache = 256 KB I+D on chip per core
hw_tcache = 12 MB I+D on chip per chip
```

Performance critical aspects of Host
and Accelerator are documented

```
#####
# Accelerator Hardware information
#####

hw_accel_model = Tesla K40c
hw_accel_vendor = NVIDIA
hw_accel_name = NVIDIA Tesla K40c
hw_accel_type = GPU
hw_accel_connect = PCIe 3.0 16x
hw_accel_ecc = Yes
```

SPEC Accel Config File (8/8)

config file: \$SPEC/config/openacc-pgi.cfg

```
#####
# Software information
#####
default=default=default=default:
sw_avail = Feb-2014
sw_compiler = PGI Accelerator Server Complete, Release 14.2
sw_accel_driver = NVIDIA UNIX x86_64 Kernel Module 319.60
#####

# MD5 section. It will be created by SPEC automatically
# It is used by SPEC to check whether an executable
# the current compiler and flags settings.
#####
```

Information on software configuration, e.g. compilers

MD5 section

- Automatically generated by SPEC tools
- Used to check whether an executable is created using the current settings

Run Accel OpenACC: on GPU

batch script: accel-pgi.bsub

```
$> runspec --config=sc19tut-openacc-pgi --tune=base --size=ref 353 370
```

- Run in **base** mode
- Use ref data set
- Execute benchmarks 353.clvrleaf (Hydrodynamics) and 370.bt (PDE solver)

Name of the config file



hands-on

```
$> bsub accel-pgi.bsub
```

Submit to batch system.
Note: Machine is only available on day of tutorial

Timing on Ref

	353.clvrleaf	370.bt
V100 GPU	37s	10s



hands-on

Run Accel OpenACC: on CPU

batch script: accel-pgi-multicore.bsub

```
$> runspec --config=sc19tut-openacc-pgi-multicore --tune=base --size=ref 353 370
```

- Run in **base** mode
- Use ref data set
- Execute benchmarks 343.clvrleaf (Hydrodynamics) and 370.bt (PDE solver)

Name of the config file



hands-on

```
$> bsub accel-pgi-multicore.bsub
```

Submit to batch system.
Note: Machine is only available on day of tutorial

Timing on Ref



hands-on

	353.clvrleaf	370.bt
V100 GPU	37s	10s
Power 9 CPU	132s	110s

Run Accel OpenACC: published results

<https://www.spec.org/accel/results/res2018q3/accel-20180814-00114.html>

SPEC® ACCEL™ ACC Result	
Copyright 2015-2018 Standard Performance Evaluation Corporation	
IBM Corporation (Test Sponsor: NVIDIA Corporation)	SPECaccel_acc_base = 11.9
Tesla V100	SPECaccel_acc_peak = 11.9
IBM Power Systems AC922 for High Performance Computing (8335-GTH)	
ACCEL license: 019	Test date: Aug-2018
Test sponsor: NVIDIA Corporation	Hardware Availability: May-2018
Tested by: NVIDIA Corporation	Software Availability: Aug-2018

<https://www.spec.org/accel/results/res2018q3/accel-20180814-00113.html>

SPEC® ACCEL™ ACC Result	
Copyright 2015-2018 Standard Performance Evaluation Corporation	
IBM Corporation (Test Sponsor: NVIDIA Corporation)	SPECaccel_acc_base = 3.02
Power9	SPECaccel_acc_peak = 3.02
IBM Power Systems AC922 for High Performance Computing (8335-GTH)	
ACCEL license: 019	Test date: Aug-2018
Test sponsor: NVIDIA Corporation	Hardware Availability: May-2018
Tested by: NVIDIA Corporation	Software Availability: Aug-2018

Results available on
www.spec.org

Run Accel OpenACC: your full run on GPU

batch script: accel-pgi.bsub

```
$> runspec --config=sc19tut-openacc-pgi --tune=base --size=ref openacc
```

- Run in base mode
- Run ref data set
- Execute entire OpenACC suite (replace benchmark “353 370” by “**openacc**”)

Run the whole OpenACC suite



hands-on

```
$> bsub accel-pgi.bsub
```

Submit to batch system.
Note: Machine is only available on day of tutorial



hands-on

Flags - From base to peak runs

<https://www.spec.org/omp2012/docs/runrules.html#section2.2.4>

Portability flags

links 

- Allowed if benchmark cannot be built and execute correctly w/o these flags
- Must be performance neutral
- Requirements
 - Provided over PORTABILITY flag
 - Must be approved by SPEC HPG committee

Flags - From base to peak runs

Base optimization rules: <https://www.spec.org/omp2012/Docs/runrules.html#section2.3>

Base runs (recap from Part A)

links 

- Common set of optimizations & environment settings for all benchmarks
- “baseline”
 - single set of switches
 - single-pass make process
 - high degree of portability, safety, performance
- Must adhere to strict rules
 - e.g. same compiler for all modules of a given language
 - All flags, options must be the same, e.g., also the level of parallelism
 - Only portability switches allowed
 - More rules (base & peak) on names, library substitutions, data type sizes, source code changes

Flags - From base to peak runs

Peak optimization rules: <https://www.spec.org/omp2012/Docs/runrules.html#section2.4>
Published results OMP2012: <https://www.spec.org/omp2012/results/omp2012.html>

Peak runs

- Set of optimizations individually selected for each benchmark
 - e.g. different compilers, flags
- Called “aggressive compilation”

links 

Summary

- Many published results do not contain peak results (often coming from academic institutions)
- Results submitted by vendors often contain peak results

Flags - From base to peak runs

- Modifying the config file
 - Once you have a config file that runs on your system, it is easy to modify it
 - E.g. peak optimizations for better performance
 - SPEC Accel is relatively new, lack of peak run results (maybe you can produce one?)
 - Showing OMP2012 as example.
- Example:

SPEC® OMPG2012 Result	
Copyright 2012-2016 Standard Performance Evaluation Corporation	
Hewlett Packard Enterprise (Test Sponsor: HPE)	SPECompG_base2012 = 47.7
Integrity Superdome X (288 core, 2.50 GHz, Intel Xeon E7-8890 v3)	SPECompG_peak2012 = 55.3
OMP2012 license: 1 Test sponsor: HPE Tested by: HPE	Test date: Apr-2016 Hardware Availability: Oct-2015 Software Availability: Feb-2016

```

350.md: -O3 -openmp -ipo -xCORE-AVX2 -fno-alias -opt-malloc-options=1 -fp-model fast=2
         -no-prec-div -no-prec-sqrt -align array64byte

351.bwaves: -O3 -openmp -ipo -xCORE-AVX2 -fno-alias -fp-model fast=2 -no-prec-div -no-prec-sqr
              -align array64byte

357.bt331: Same as 351.bwaves

360.ilbdc: basepeak = yes

362.fma3d: -O3 -openmp -ipo -xCORE-AVX2 -fno-alias -align array64byte

363.swim: -O3 -openmp -ipo -xCORE-AVX2 -fno-alias -opt-streaming-stores always
           -opt-malloc-options=3 -align array64byte

370.mgrid331: -O3 -openmp -ipo -xCORE-AVX2 -fno-alias -opt-malloc-options=3 -fp-model strict

```

SPEC Peak run example: OMP Config File

Just show you how a peak run will look like:

```
#####
# Peak
#####
default=peak=default=default:
OPTIMIZE      = -O3 -qopenmp -ipo -xCORE-AVX2 -no-prec-div
COPTIMIZE     = -ansi-alias
CXXOPTIMIZE   = -ansi-alias
FOPTIMIZE     = -align
#
# [...] Environment variables
```

FP optimizations

```
350.md=peak=default=default:
OPTIMIZE=-O3 -qopenmp -ipo -xCORE-AVX2 -ansi-alias -qopt-malloc-options=1
FOPTIMIZE=-fp-model fast=2 -no-prec-div -no-prec-sqrt -align array64byte
```

```
363.swim=peak=default=default:
OPTIMIZE=-O3 -qopenmp -ipo -xCORE-AVX2 -ansi-alias -qopt-streaming-stores always
          -qopt-malloc-options=4
threads=24
```

- **-qopt-malloc-options: alternate algorithm for malloc**
- **-fp-model fast=2: aggressive optimization on FP computations**
- **-no-prec-sqrt: less precise square root computations/ more performance**
- **-align array64byte: align arrays to 64 Byte**
- **-qopt-streaming-stores always: use non-temporal stores (write through)**

memory optimizations

Summary

- **SPEC ACCEL for accelerators**
 - Including benchmarks implemented using OpenACC, OpenCL and OpenMP target
 - --device to choose device number
 - works with a wide range of acclerators (NVIDIA GPU, AMD GPU, Xeon Phi), also CPU.

Contents

- Cluster login 
- Overview of system requirements
- How to get SPEC benchmarks?
 - Benchmark acquisition & licensing
 - Download & unpacking
- How to setup SPEC benchmarks?
 - Installation 

- How to run SPEC benchmarks? 
 - Benchmark components & workloads
 - Runspec & run rules
 - Configuration files
 - From base to peak runs
 - Switch of compiler
- **How to publish SPEC benchmark results?**
 - Output files
 - Reportable runs
 - Process of publishing

Output files

- SPEC runs create results in `result` subdirectory
- Text files, “`.txt`”,
 - Preview of the result as it would look on the SPEC website
- Log files, “`.log`”, “`.log.debug`”
 - Verbose output of the benchmark run
- Raw files, “`.rsf`”,
 - Above the “line” are editable fields about the run such as system or software configuration
 - Below the “line” are the encoded results. Tampering with the results will corrupt the file.

Publish results on SPEC website

- Publishing SPEC HPG results helps to get a rich set of different HW, compilers, configurations, etc.
 - But it's **not** required
 - Note: non SPEC members pay publication fee
- Recap (Part A): Result Submission Process
 - Obtain and install the benchmark
 - Perform a valid run → Adhere to all **run rules** + create config file + **reportable run**
 - Supply hardware and software description → **Edit documentation** portion of results ((raw) file)
 - **Submit result for review (and publication) to SPEC HPG**
 - 2 week review process
 - (Define embargo period)
 - Use the result as you would like

<https://www.spec.org/omp2012/docs/runrules.html#section4.7>
https://www.spec.org/hpg/submitting_results.html



Run rules

A published result means

1. Performance observation → testing

- Generally no code modifications of provided sources allowed
- Tester supplies compiler, system, config files
- Tester provides description of performance-relevant conditions

2. Declaration of expected performance → reproducing

- Observed level of performance obtainable by others (e.g., used by vendors)
- Components (e.g., hardware, OS) obtainable by others

3. Claim about maturity of performance methods

- E.g., correct code generation & improved performance for a class of programs larger than the SPEC suite

OpenMP: <https://www.spec.org/omp2012/docs/runrules.html>

MPI: <https://www.spec.org/mpi/docs/runrules.html>

ACCEL: <https://www.spec.org/accel/docs/runrules.html>

Defs: <https://www.spec.org/omp2012/Docs/runrules.html#section4.2.1>

links 

Test Sponsor: entity sponsoring the testing (defaults to hardware vendor) → or can be your university

Tester: entity actually carrying out the tests (defaults to test sponsor) → or can be your name

Reportable runs

Create valid/ compliant result: **runspec --reportable [. .]**

- **--tune [base|all]**
- Entire SPEC suite (no single benchmarks)
- Workload: test, train, ref will be run → ref results are taken
 - Verification for all three data set sizes
- #iterations = 3 → median is taken
- #threads: one fixed number in base (variable per benchmark in peak)

Configuration disclosure (in config file or with rawformat – see next slide)

Reportable run: <https://www.spec.org/omp2012/Docs/runspec.html#section3.1.1>
<https://www.spec.org/omp2012/Docs/runspec.html#reportable>

Preparing a result for submission

- Flags and platform files
 - XML files containing detailed descriptions of the compiler flags and platform settings.
- Edit documentation portion of results: **rawformat**
 - Script used to format a raw file into text, html, Postscript, or PDF
 - Also performs a submission check to determine result is valid

Useful hint:
Make a backup copy of the rawfile before editing.

```
$> rawformat outfile.rsf  
$>  
$> rawformat -F path/to/flagsfile.xml
```

Runs offline verification of result (similar to submission), produces same output as online

Adds flags-file to the result

Rawformat: <https://www.spec.org/omp2012/docs/utility.html#rawformat>

Submitting results to SPEC

- Submission of SPEC results
 1. Process your rsf-file through rawformat to check for anything missing/ faulty
 2. Attach your rsf-file to an e-mail to, e.g., subaccel@spec.org
 3. Receive a reply with a sub-file attached
 4. For updates, modify the sub-file and attach to an e-mail to, e.g., resubaccel@spec.org
- Submitted results reviewed before publication by SPEC committee
 - Schedule: 2 weeks until reply (see (3))

Source: https://www.spec.org/accel/results/accel_acc.html

OpenACC (21):

Test Sponsor	System Name	Accelerator Name	Results		Energy	
			Base	Peak	Base	Peak
Indiana University	Lenovo NeXtScale nx360 M5	Intel Xeon E5-2680 v3	1.71	Not Run	--	--
Indiana University	HP Z820 Workstation	Intel Xeon E5-2640 v2	0.662	Not Run	1.10	--
Indiana University	Cray XC30	Intel Xeon E5-2697 v2	1.18	Not Run	--	--
Indiana University	Cray XK7	NVIDIA Tesla K20	1.71	Not Run	--	--
Indiana University	Cray XK7	NVIDIA Tesla K20	1.78	Not Run	--	--
Indiana University	Cray XK7	NVIDIA Tesla K20	2.00	Not Run	--	--
Indiana University	Cray XK7	NVIDIA Tesla K20	2.01	Not Run	--	--
Indiana University	Cray XK7	NVIDIA Tesla K20	2.07	Not Run	--	--
Lenovo Global Technology	ThinkSystem SR650	NVIDIA Tesla V100-PCIE-16GB	12.2	Not Run	--	--
Lenovo Global Technology	ThinkSystem SR670	Tesla V100-PCIE-16GB	12.0	Not Run	--	--
NVIDIA Corporation	SuperServer 1028GR-TR	Tesla K40m	2.56	2.56	--	--
NVIDIA Corporation	SuperServer 1028GR-TR	Intel Xeon E5-2698 v3	1.81	1.81	--	--
NVIDIA Corporation	SuperServer 1028GR-TR	Tesla P100-PCIE-16GB	8.02	8.02	--	--
Test Sponsor	System Name	Accelerator Name	Results		Energy	
			Base	Peak	Base	Peak
NVIDIA Corporation	SuperServer 1028GR-TR	Xeon E5-2698 v4	2.74	2.74	--	--
NVIDIA Corporation	A+ Server 1023US-TR4	EPYC 7451	2.59	2.59	--	--
NVIDIA Corporation	SuperServer 1029GQ-TRT	Xeon Gold 6148	3.77	3.77	--	--

<https://www.spec.org/accel/docs/runrules.html>

https://www.spec.org/hpg/submitting_results.html



Thank you!

Questions?

Contact

SPEC Headquarters:

Swen Boehm:

Mayara Gimenes:

Robert Henschel:

Verónica G. Vergara Larrea:

Junjie Li:

Sandra Wienke:

info@spec.org

boehms@ornl.gov

mayara@udel.edu

henschel@iu.edu

vergaravg@ornl.gov

lijunj@iu.edu

wienke@itc.rwth-aachen.de