Performance Evaluation of The APC Compiler

Mar 20, 2003

Performance Evaluation Group

✧ Tokuro ANZAKI (Hitachi, Ltd)
✧ Hiroki HONDA (The University of Electro-Communications)
✧ Kento AIDA (Tokyo Institute of Technology)

Agenda

1. Performance Evaluation Environment
2. Benchmarks and Run Rules
3. Evaluation Guideline
4. Benchmarks
5. Result of the Performance
6. Overall Results
7. Scalability
8. Conclusion

Evaluation Environments

Machine	SGI Origin2000	HP AlphaServer GS160 Model 6/731	IBM RS/6000	IBM pSeries690 Model 681
Processor	R10000 195MHz	Alpha 21264A 731MHz	PowerPC 604e 200MHz	Power4 1.1GHz
Processor Num	32 (2PE/node x 16)	8	8	16
Memory	11GB (704MB/node x 16)	4GB	1GB	8GB
Compiler	MIPSpro Fortran90 Version 7.30	Compaq Fortran V5.4 KAP Fortran V4.3	IBM XL Fortran for AIX Version 7.1.0	IBM XL Fortran for AIX Version 8.1

Benchmarks and Run Rule

- Select Standard benchmarks
  1. SPEC CFP2000:Most familiar S/E Benchmark Suite
  2. SPEC CFP95:
  3. NAS-PB 2.3-Serial
- Only FORTRAN77 programs!
- Use the SPEC CFP2000_base compiler option with vender specific automatic parallelizing options.
- Use the same compiler option for backend code generation.

Evaluation Guideline

1. **High level parallelism Benchmark**
   - *speedup ≥ half of #processors*
   - **Ex. speedup ≥ 4 on an 8 processor system**
   - **The APC compiler can not attain performance that is twice as good as original compilers, but it would attain some performance gain.**
2. **Low level parallelism Benchmark**
   - *speedup < half of #processors*
   - **Ex. speedup < 4 on an 8 processor system**
   - **The APC compiler potentially attains performance that is twice as good as original compilers.**
3. **Difficult to parallelize Benchmark**
   - *No performance gain on any number of processors*
   - **If the benchmarks had potential parallelism, we would be able to easily attain more than double performance!**

Benchmarks
(Only FORTRAN77)

Benchmark Suite	Benchmark	Description
SPEC CFP95	101.tomcatv	Vectorized mesh generation
	102.swim	Shallow water equations.
	103.su2cor	Monte-Carlo method
	104.hydro2d	Navier Stokes equations
	107.mgrid	3d potential field
	110.applu	Partial differential equations
	125.turb3d	Turbulence modeling
	141.apsi	Weather prediction
	145.fpppp	From Gaussian series of quantum chemistry benchmarks
	146.wave5	Maxwell's equations
SPEC CFP2000	168.wupwise	Quantum chromo-dynamics
	171.swim	Shallow water modeling
	172.mgrid	Multi-grid solver in 3D potential equations
	173.applu	Parabolic/elliptic partial differential equations

Benchmarks

(Only FORTRAN77) (cont.)

Benchmark Suite	Benchmark	Description
SPEC CFP2000 (cont.)	200.sixtrack	Particle accelerator model
	301.apsi	Solves problems regarding temperature, wind, distribution of pollutants
NAS-PB2.3 Serial	EP	Embarrassingly parallel
	MG	3D Multi-grid
	CG	Conjugate gradient
	FT	3D FFT partial differential equation
	LU	LU solver
	SP	Pentadiagonal solver
	BT	Block tridiagonal solver

Avoid Duplication of the same type of BMT(same color above)

- Duplicate benchmarks: Same Application, the code itself or problem sizes are different.
  - -102.swim, 171.swim
  - -107.mgrid, 172.mgrid, MG
  - -113.applu, 173.applu, LU
  - -141.apsi, 301.apsi

IBM RS/6000 (8 CPUs)

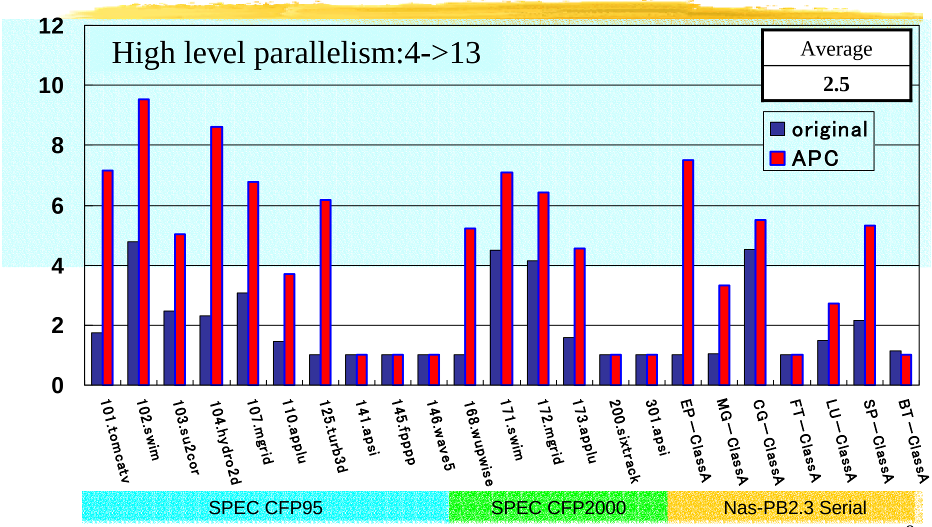

IBM pSeries690 (16 CPUs)

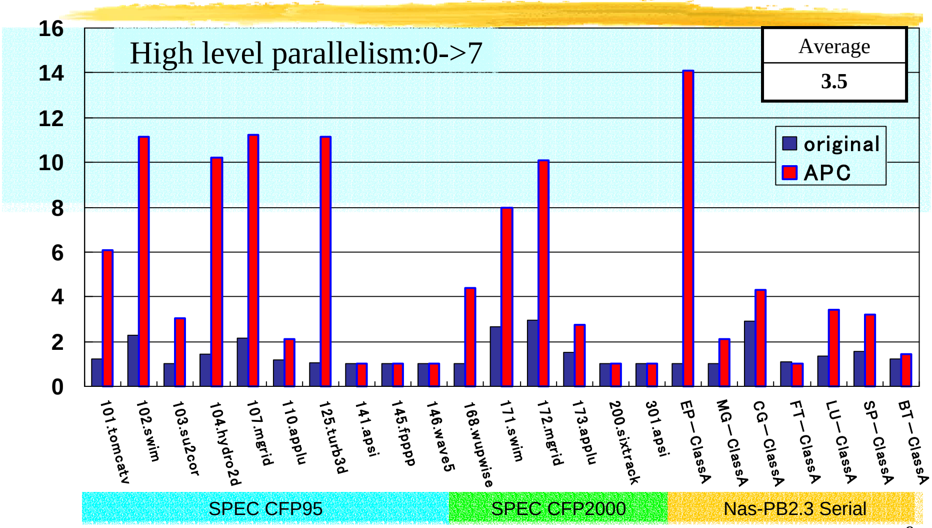

HP_Alpha (8 CPUs)

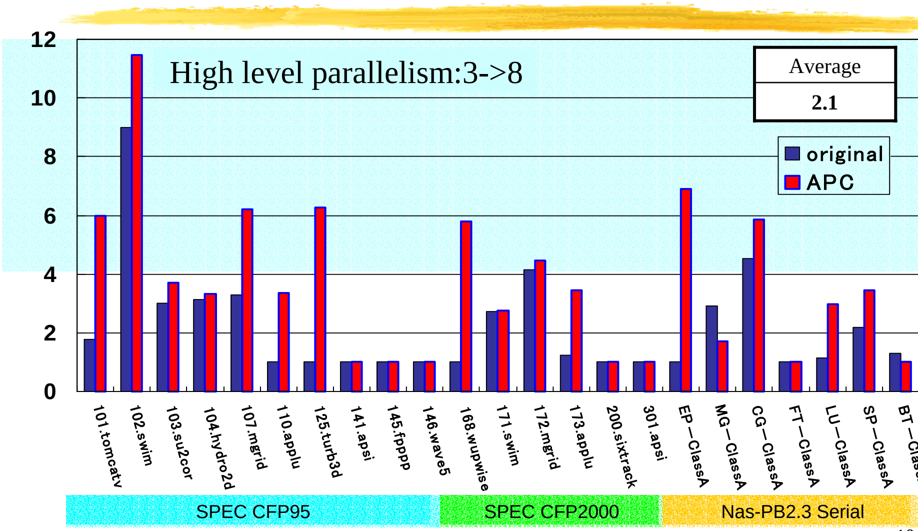

SGI Origin2000 (32 CPUs)

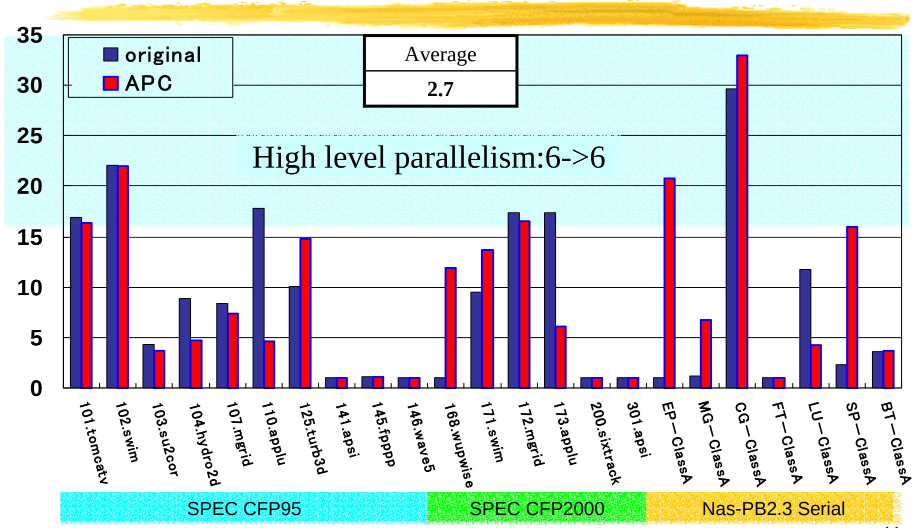

Overall Results

- **IBM RS/6000**
  - #benchmarks highly parallelized:
    APC compiler: **13** (out of 23) original: **4**
  - The APC compiler shows ***2.5*** times faster speedup than the original.
- **IBM pSeries690**
  - #benchmarks highly parallelized:
    APC compiler: **7** (out of 23) original: **0**
  - The APC compiler shows ***3.5*** times faster speedup than the original.
  - Difficult to scale well over 8 processors.
- **HP_Alpha**
  - #benchmarks highly parallelized:
    APC compiler: **8** (out of 23) original: **3**
  - The APC compiler shows ***2.1*** times faster speedup than the original.
  - Difficult to scale well over 4 CPUs.
- **SGI Origin2000**
  - #benchmarks highly parallelized:
    APC compiler: **6** (out of 23) original: **6**
  - The APC compiler shows ***2.7*** times faster speedup than the original.
  - Difficult to scale well over 8 or 16 CPUs.

101.tomcatv (Hierarchical, cache optimize)

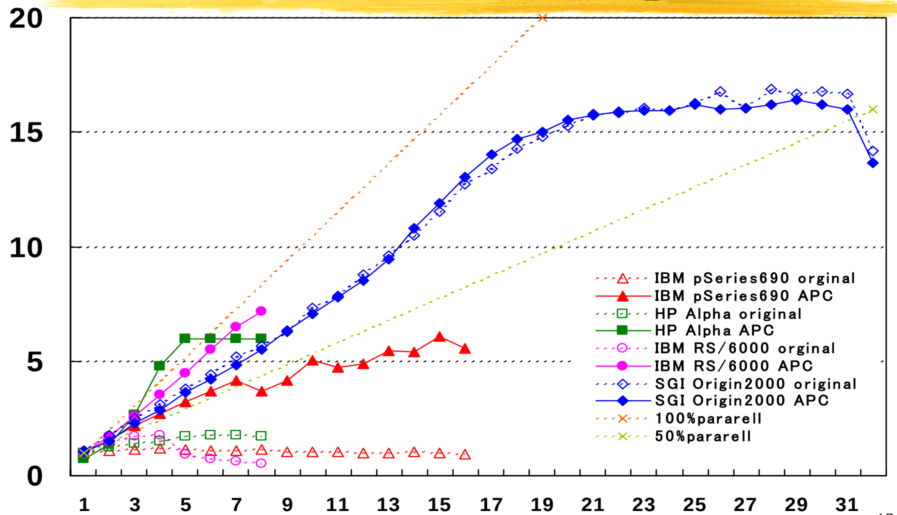

102.swim (Hierarchical, cache optimize)

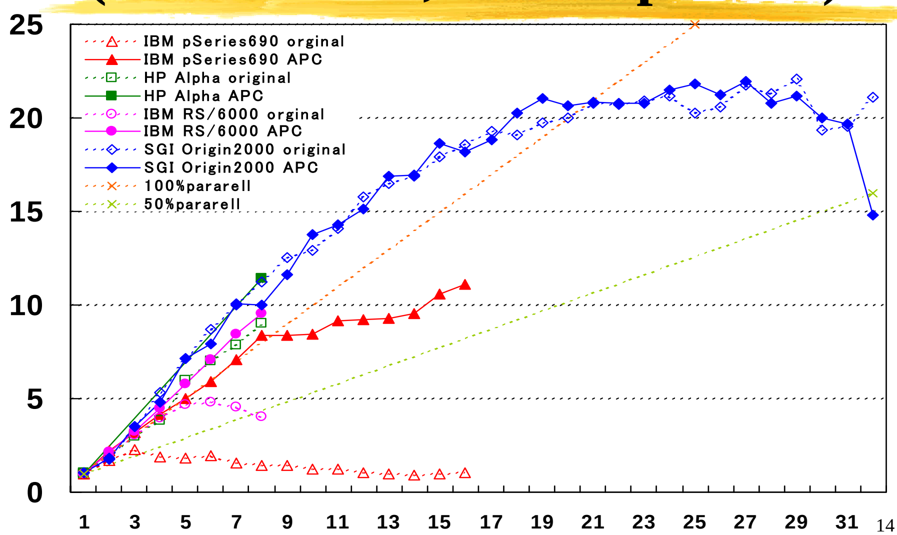

171. Swim (Hierarchical, cache optimize)

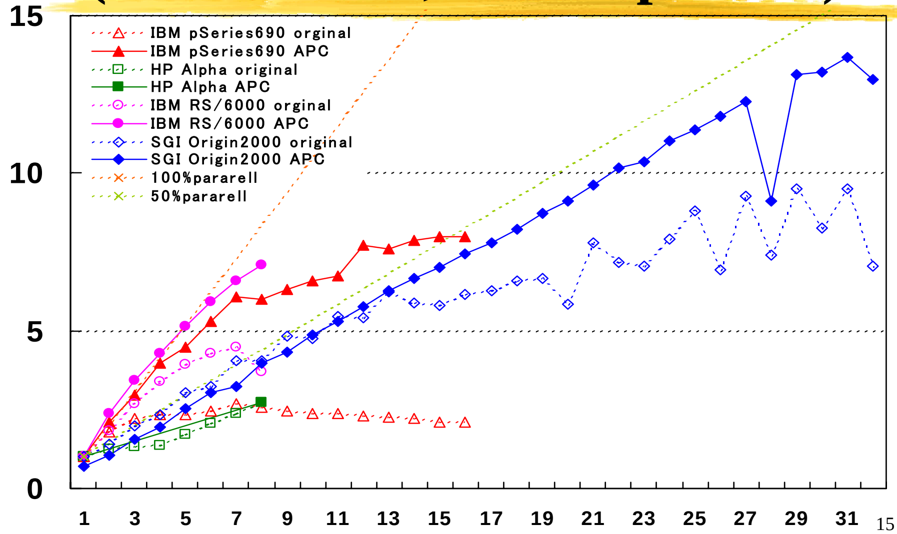

103.su2cor (Hierarchical)

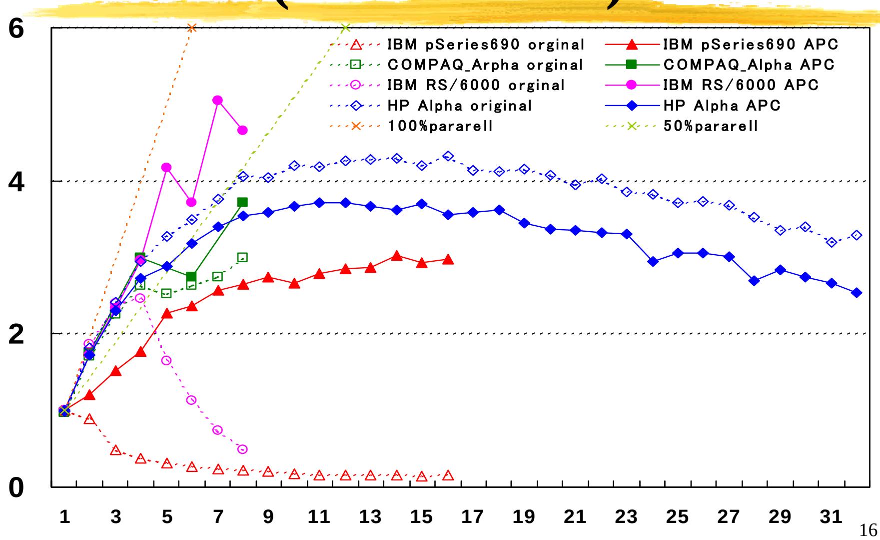

104.hydro2d
(Hierarchical, cache optimize)

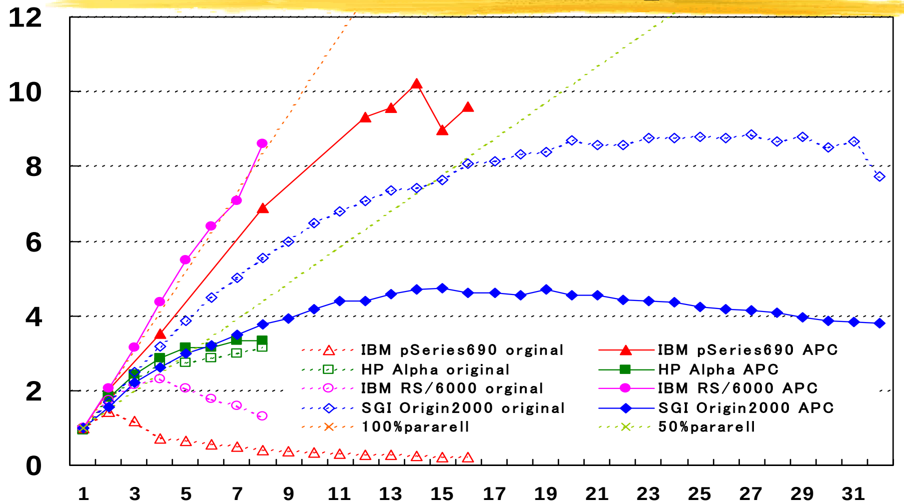

168. wupwise
(Inter-Procedural)

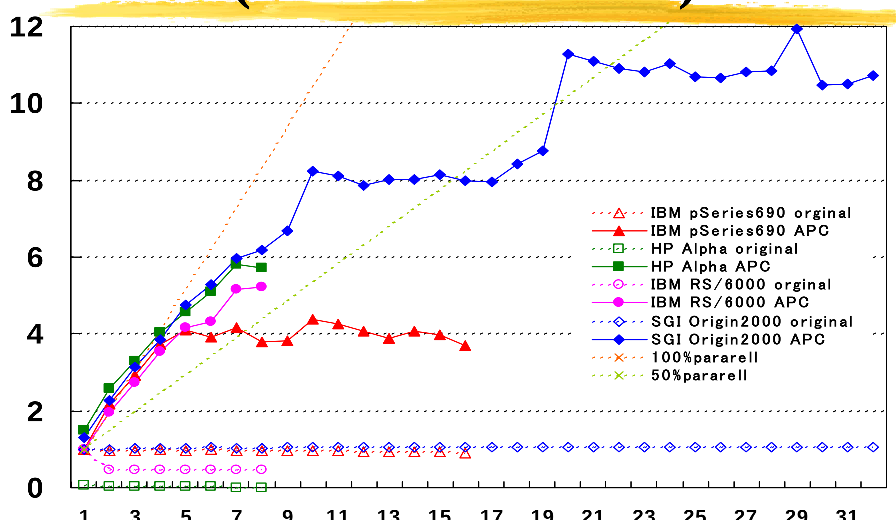

110.applu
(Affine Transformation)

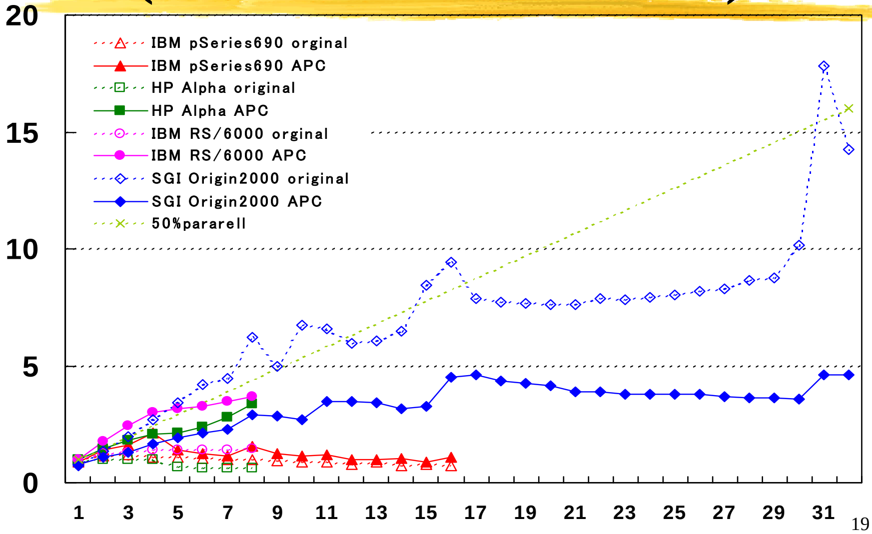

125. Turb3d
(Hierarchical, Inter-procedural)

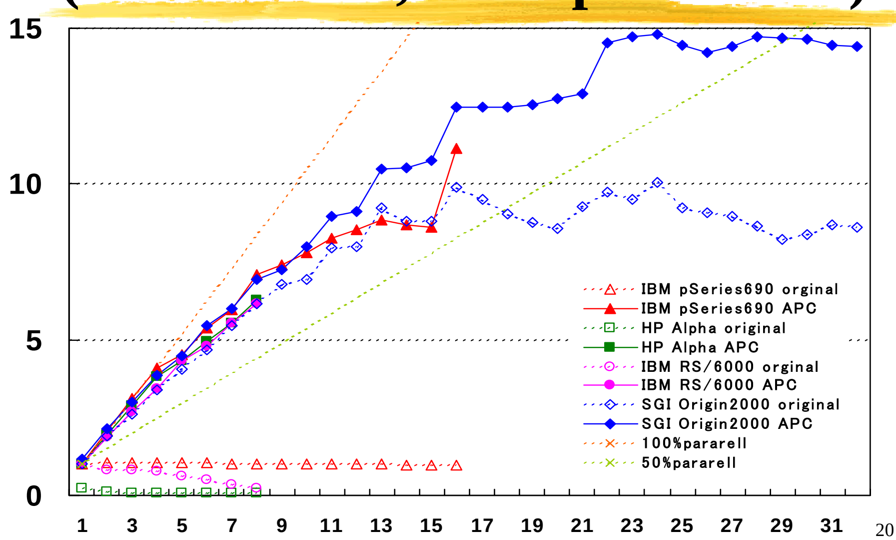

NPB2.3 EP-ClassA (Inter-procedural)

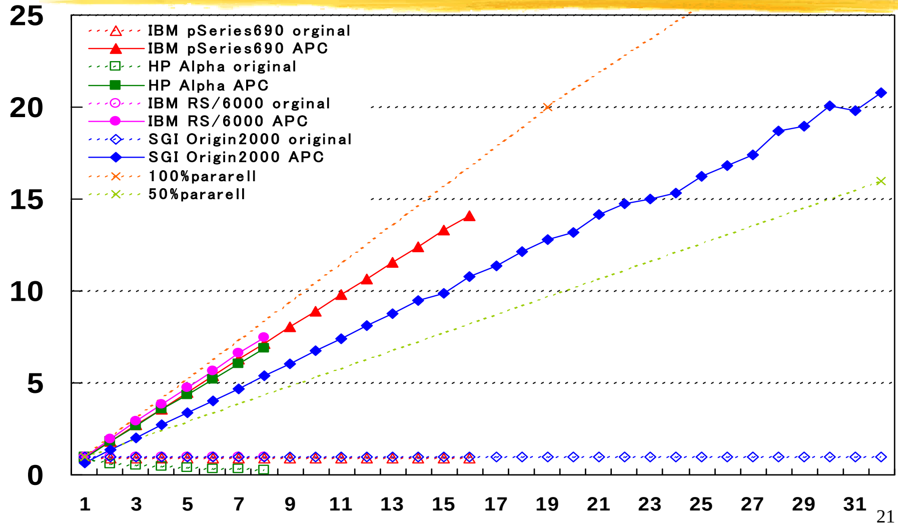

107.mgrid (Hierarchical)

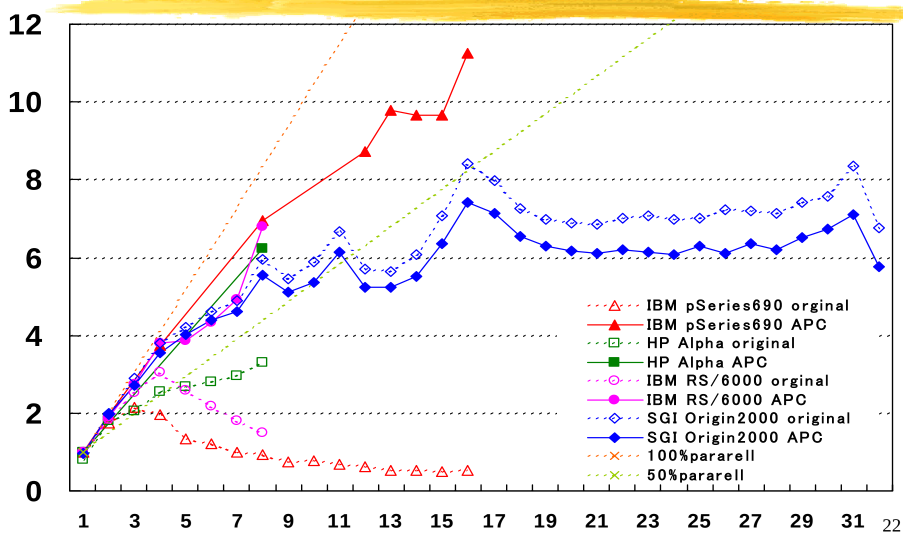

NPB2.3 MG-ClassA (Hierarchical)

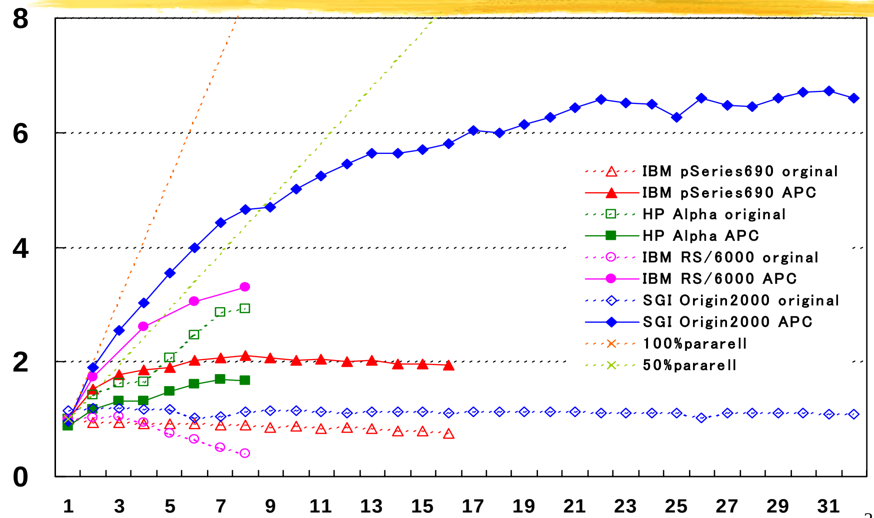

NPB2.3 BT-ClassA (Medium grain)

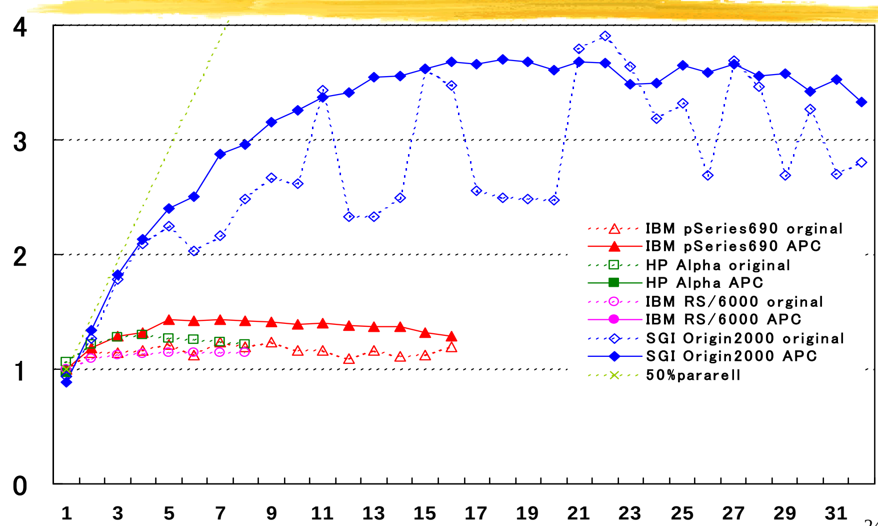

NPB2.3 CG-ClassA (First touch)

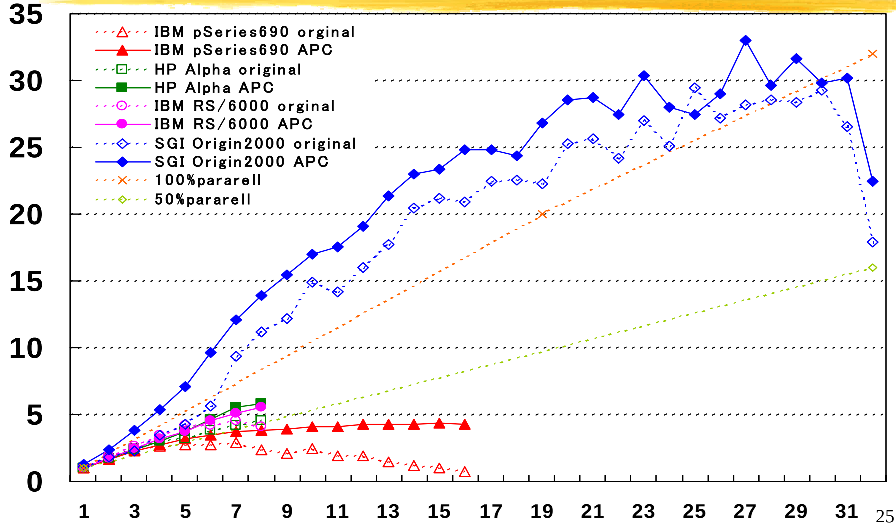

NPB2.3 CG-ClassB (First touch)

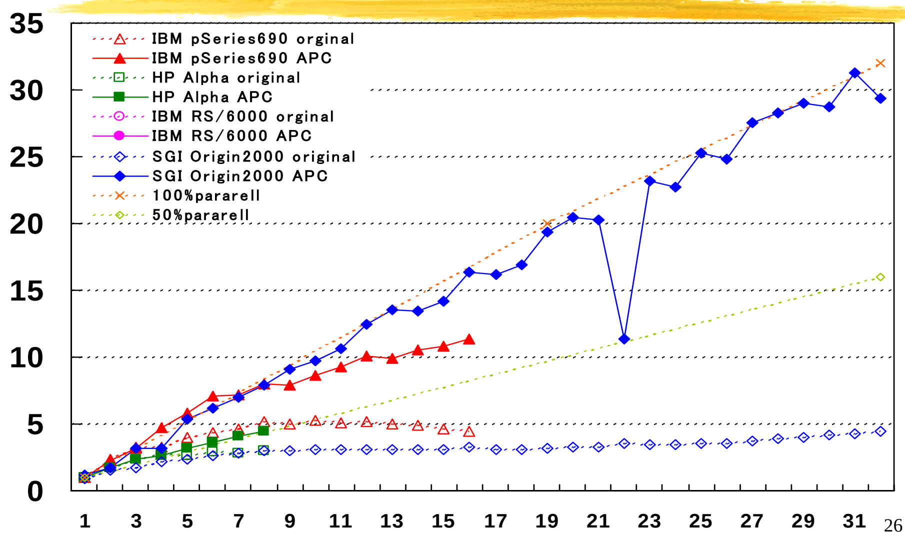

NPB2.3 SP-ClassA (Medium grain)

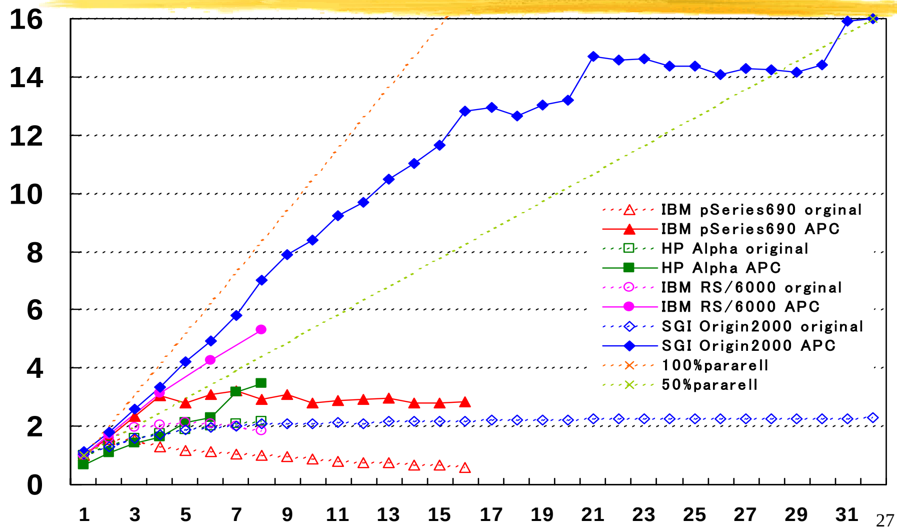

Conclusion

- Performance evaluation of compilers, the APC compilers and the originals, by 23 FORTRAN 77 benchmarks (out of 10 programs in SPEC CFP95, 6 programs in SPEC CFP2000, 7 programs in NAS-PB) on 4 parallel computers.
- Automatic multi-grain parallelization is effective for many benchmarks, even for some of the benchmarks that are difficult to parallelize.
- On every machine, the APC compiler attains more than double performance compared with the original compilers.