

# **Research and Development in Advanced Parallelizing Compiler Technology Project**

**APC Project Leader  
Prof. Hironori Kasahara  
Waseda University**

<http://www.kasahara.elec.waseda.ac.jp>

[kasahara@waseda.jp](mailto:kasahara@waseda.jp)

<http://www.apc.waseda.ac.jp>

## **Needs for Advanced Parallelizing Compiler**

- **Wide use of multiprocessor architectures**
  - From chip multiprocessor to high performance computer such as
    - » IBM Power4 chip multiprocessor
    - » Sun Ultra80 4processor multiprocessor workstation
    - » Sun V880 8 processor entry level server
    - » IBM pSeries690 high end sever (RegattaH)
    - » Earth Simulator world fastest supercomputer
  - Increasing gap between peak and effective performance with increase of the number of processors
  - Increasing speed gap between processor and memory
  - Difficulty in parallel programming and tuning
- **Compilers to improve Effective Performance, Cost Performance and Ease of Use (Software Productivity) are required.**



<http://www.sun.com/desktop/products/ultra80/>

# IBM Power4 Chip Multiprocessor

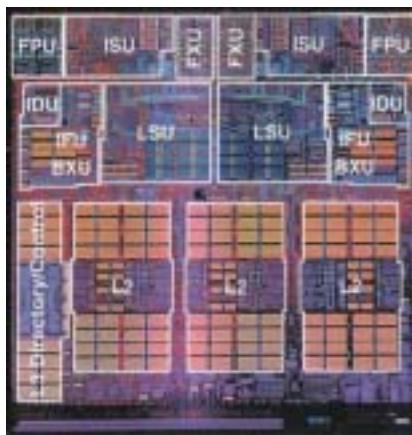
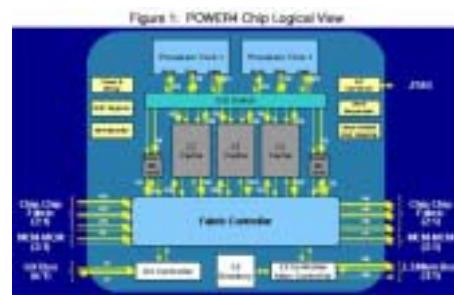


Figure 2  
POWER4 chip photograph showing the principal functional units in the microprocessor core and in the memory subsystem.



<http://www-1.ibm.com/servers/eserver/pseries/hardware/whitepapers/power4.html>

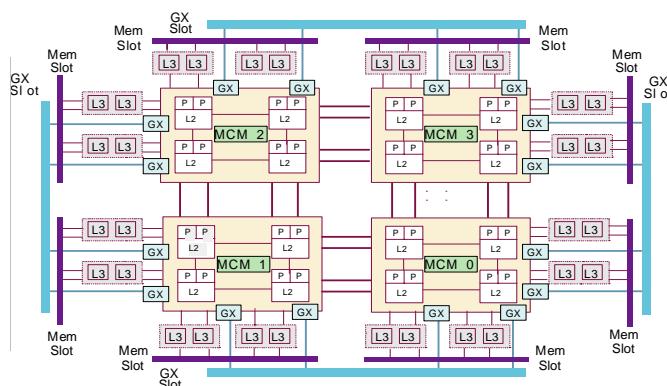
J. M. Tendler, J. S. Dodson, J. S. Fields, Jr., H. Le, and B. Sinharoy,

## "POWER4 system microarchitecture",

IBM Journal of Research and Development, Vol46, No. 1, 2002

# IBM pSeries690 RegattaH

- Up to 16 Power4: 32 way SMP Server
  - L1(D) : 64 KB (32KB/processor, 2 way assoc.), L1(I) : 128 KB (64KB/processor, Direct map)
  - L2 : 1.5 MB (shared cache with 2 procs., 4-8 way assoc.)
  - L3 : 32 MB (external, 8 way assoc.) [x 8 = 256 MB]



Four 8-way MCM Features Assembled into a 32-way pSeries 690

[http://www-1.ibm.com/servers/eserver/pseries/hardware/whitepapers/p690\\_config.html](http://www-1.ibm.com/servers/eserver/pseries/hardware/whitepapers/p690_config.html)

IBM @server pSeries690 Configuring for Performance

## Advanced Parallelizing Compiler Technology Project

2000.9.8 – 2003.3.31

2000: ¥420 Million, 2001: ¥380M, 2002: ¥290M

Performance

1T

Hardware Peak performance

### <Purpose>

Improvement of Effective performance

Cost-performance

Ease of use

1G

Theoretical maximum performance vs. Effective performance of HPC

1980

1990



Slow down

Expand Gap

### Background and Problems

Adoption of parallel processing as a core technology on PC to HPC

Increase of importance of software on IT

Need for improvement of cost-performance and usability

### Contents of Research and Development

R & D of advanced parallelizing compiler

Multigrain, Data localization, Overhead hiding

R & D of Performance evaluation technology for parallelizing compilers

### Goal: Double the effective performance

### Ripple Effect

Development of competitive next generation PC and HPC

Putting the innovative automatic parallelizing compiler technology to practical use

Development and market acquisition of future single-chip multiprocessors

Boosting R&D in the following many fields:

IT, Bio-tech., Device, Earth environment,

Next-generation VLSI design, Financial engineering,

Weather forecast, New clean energy, Space

development, Automobile, Electric Commerce, etc



## IT21 in Japanese Millennium Project

### • Millennium Project:

- Organized by the Prime Minister Keizo Obuchi in the fiscal year of 2000 to promote research and development in the Environment, Aging and Information Processing.

### • IT21:

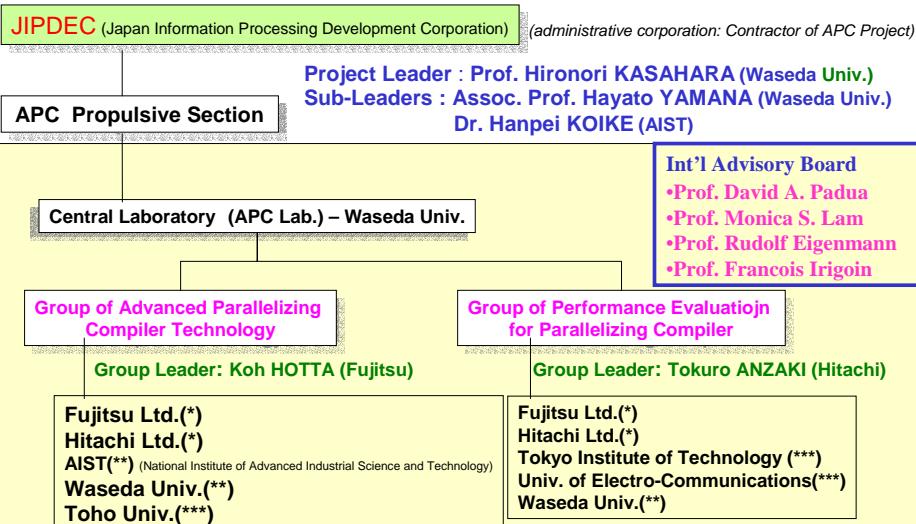
- A project in the Information Technology area in the Millennium Project.
- Supported by METI (Ministry of Economy, Trade and Industry) and MPHPT (Ministry of Public Management, Home Affairs, Posts and Telecommunications).
- 40 projects for hardware and software related with
  - Next generation computing and internet including
    - Bio-informatics, GRID, Internet infrastructure, System on Chip, Parallelizing Compiler and so on.

# Organization of Advanced Parallelizing Compiler Project (1)



Num. of Members : 35	Participating Organization
11	Fujitsu Ltd.
12	Hitachi Ltd.
3	AIST (National Institute of Advanced Industrial Science and Technology)
6 (+ Students)	Waseda Univ.
1	Tokyo Institute of Technology
1	Univ. of Electro-Communications
1	Toho Univ.

# Organization of Advanced Parallelizing Compiler Project (2)



## APC Committees for R&D and Self Evaluation

### ■ Development Promotion Committee:

Highest Decision Making, 3 times for three FYs

### ■ Technology Committee:

Discussing project targets, research plans, etc, 20 times

### ■ Technology Forum :

Discussing detailed research plans, research progress, technology  
46 times in addition to the Networked Centralized R&D activities  
introduced in APC as the first trial.

### ■ International Cooperation Committee with Advisory Board:

Having advices and evaluation for the APC activities from the world  
leading professors to achieve internationally valuable R&D results

2 times (1<sup>st</sup>: Sep. 5-6, 2001, 2<sup>nd</sup>: Mar 19-21, 2003)

### ■ International Advisory Board

Chair: Prof. Padua (Univ. of Illinois at Urbana-Champaign)

Polaris, Cedar, Parafrase

Prof. Irigoin (Ecole des Mines de Paris) PIPS Parallelizer

Prof. Lam (Stanford Univ.) Suif, Suif Explorer,  
NCI(National Compiler Infrastructure)

Prof. Eigenmann (Purdue Univ.) Polaris, SPEC HPC/OMP

## Research Parallelizing Compilers

- Automatic loop parallelization

<Data Dependence Analysis & Restructuring>

◆ Runtime dependence analysis, Symbolic analysis,  
Interprocedure analysis, Unimodular  
transformation, Array privatization, Fusion,  
Unrolling etc.

→ Polaris, SUIF , Promis (Parafrase2) ...

◆ Limitation of loop parallelism

- Sequential loops and parts of program except loops
- Fewer loop iterations than the number of processors

## Multigrain Parallelization

- Improvement of effective performance
- Scalability
  - ◆ Coarse grain parallelism:  
subroutines, loops, basic blocks
  - ◆ Near fine grain parallelism: statements  
in addition to loop parallelism

Multi level parallelism:

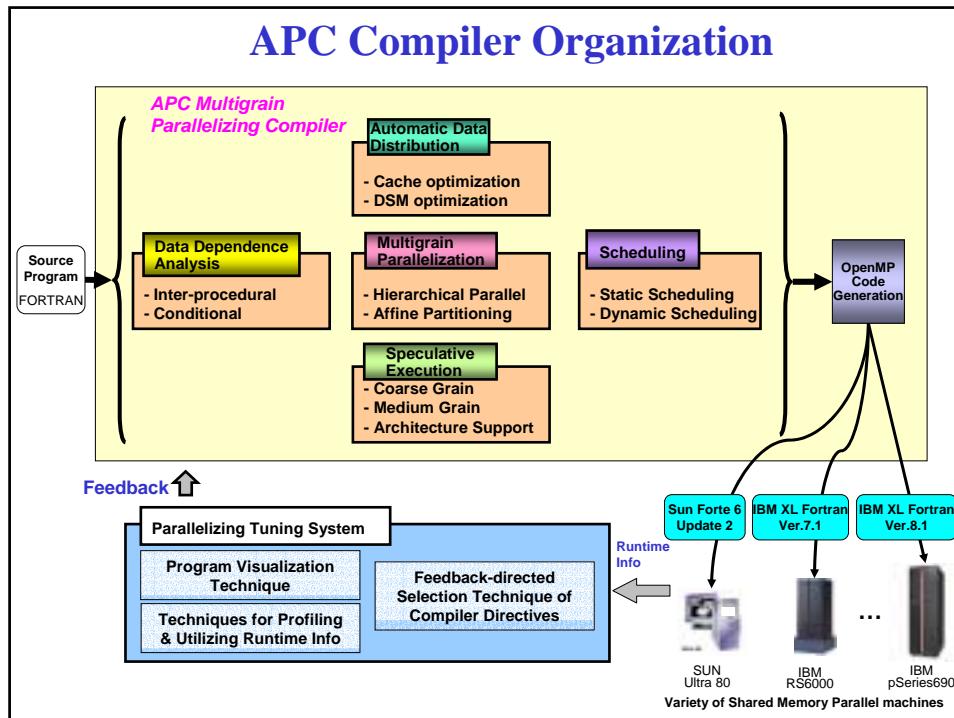
- PROMIS compiler (Parafrase2 + EVE)
- NANOS compiler (NthLIB, extended OpenMP)
- OSCAR Fortran compiler

## Multigrain Parallel Processing in OSCAR parallelizing compiler

- Coarse grain task parallelism
  - Among subroutines, loops & basic blocks
  - Statically or dynamically schedule to processors or PCs (Processor Clusters) at compile or run time
- Loop parallelism
  - Among loop iterations
  - Schedule to processors or PCs
- (Near) Fine grain parallelism
  - Among statements in a basic block
  - Statically schedule to processors in a PC
- Their hierarchical combination

# Research Topics in Advanced Parallelizing Compiler Project

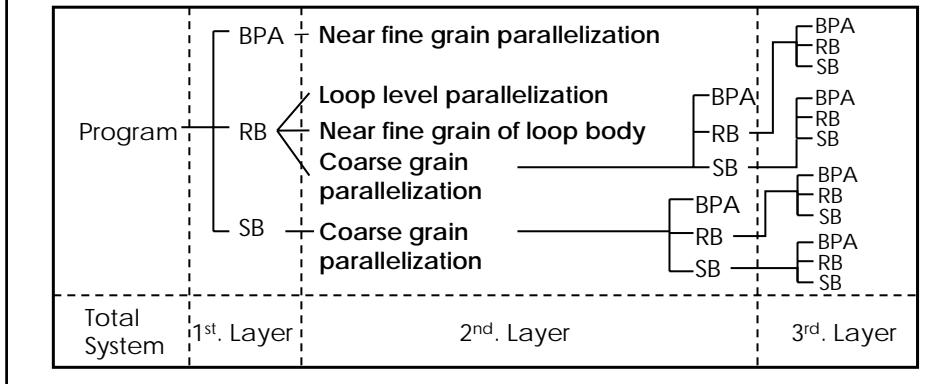
- **Multigrain parallelization technology**
  - Data dependence analysis (Inter-procedural, Runtime)
  - Automatic selection of suitable grain
  - Automatic data distribution (DSM, Cache, Local Mem.)
  - Scheduling (Load balancing, Data transfer overhead minimization)
  - Speculative execution
  - Tuning tools
  - Use of OpenMP+ as an intermediate language
- **Performance evaluation of compiler**
  - Evaluation of individual parallelization technology
  - Evaluation of total compiler performance
    - Selection of benchmark programs which can clearly show performance of compiler



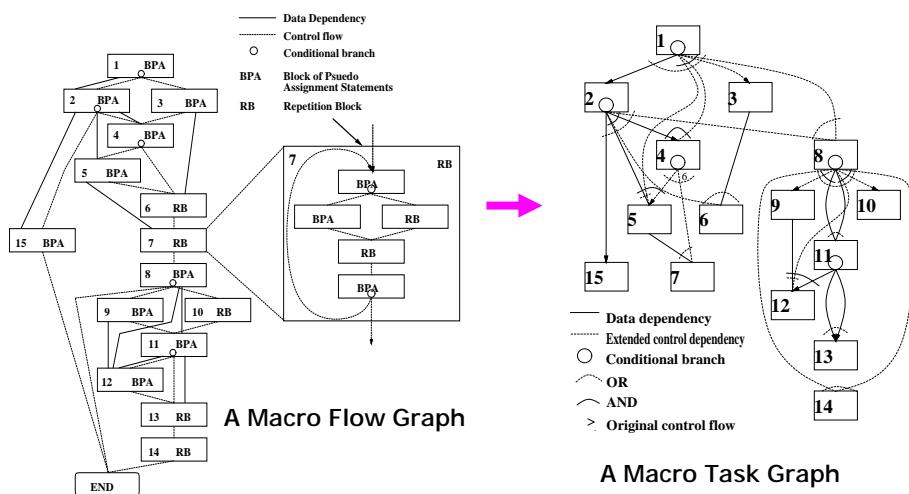
# Generation of Coarse Grain Tasks

## ■ Macro-tasks (MTs)

- ↗ Block of Pseudo Assignments (BPA): Basic Block (BB)
- ↗ Repetition Block (RB) : outermost natural loop
- ↗ Subroutine Block (SB): subroutine

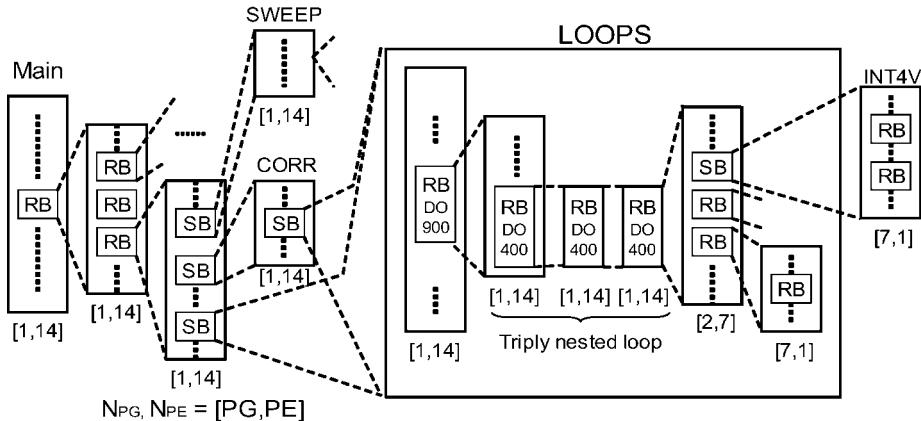


## Earliest Executable Condition Analysis for Coarse Grain Tasks (Macro-tasks)



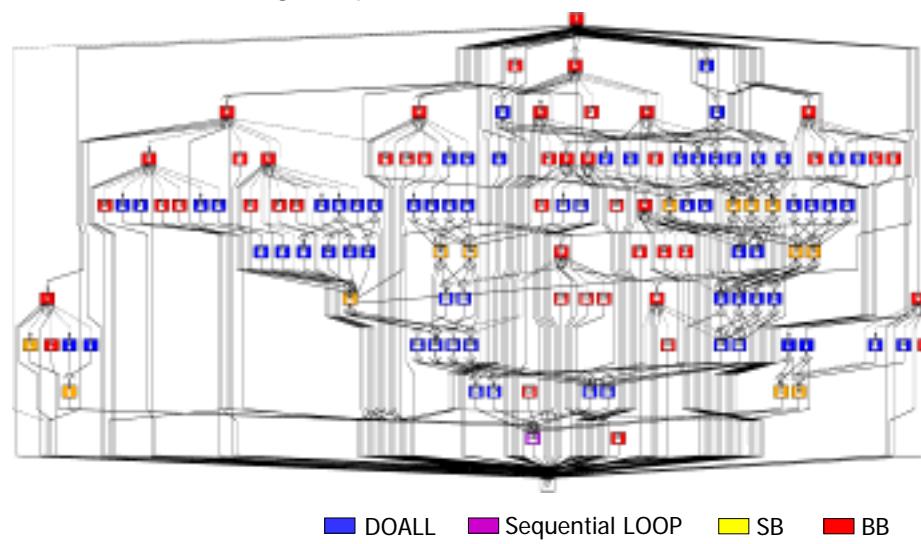
# Automatic processor assignment in 103.su2cor

- Using 14 processors
    - Coarse grain parallelization within DO400 of subroutine LOOPS



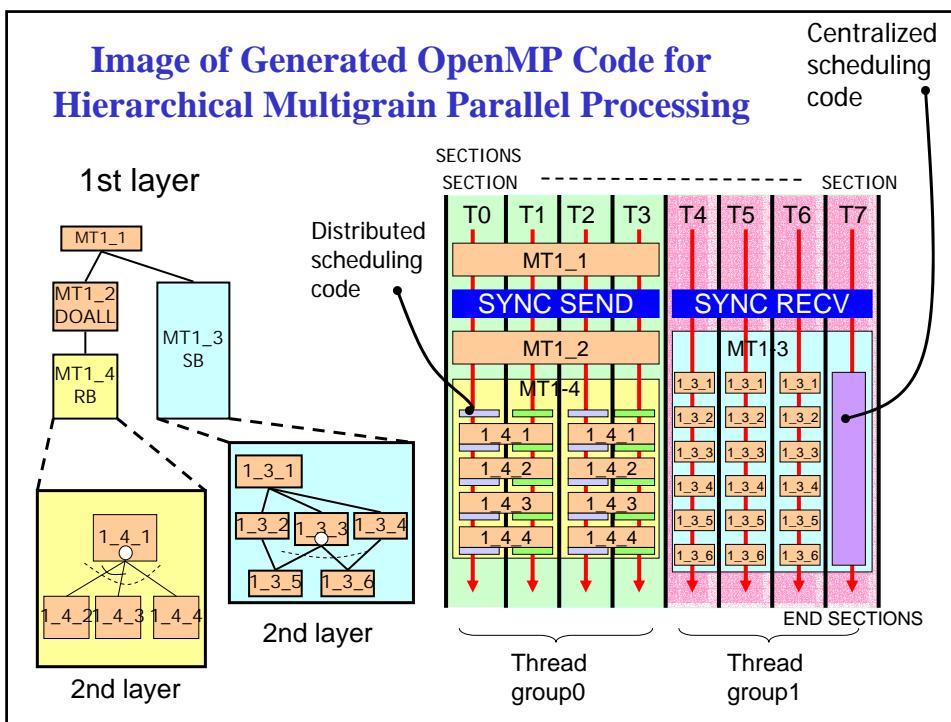
## MTG of Su2cor-LOOPS-DO400

■ Coarse grain parallelism PARA\_ALD = 4.3



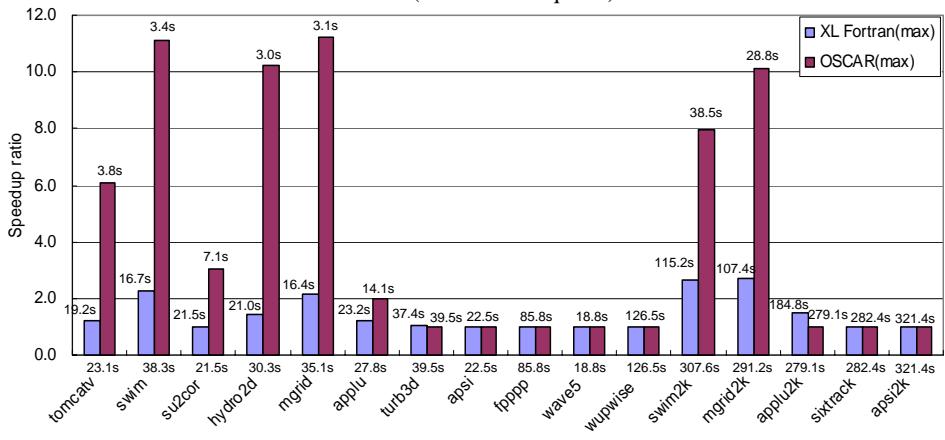
# Code Generation Using OpenMP

- Compiler generates a parallelized program using OpenMP API
  - One time single level thread generation
    - Threads are forked only once at the beginning of a program by OpenMP “PARALLEL SECTIONS” directive
    - Forked threads join only once at the end of program
  - Compiler generates codes for each threads using static or dynamic scheduling schemes
  - Extension of OpenMP for hierarchical processing is not required

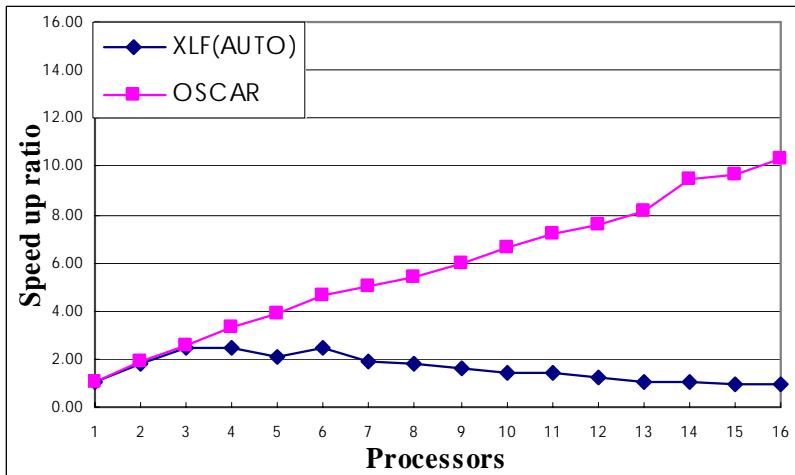


# Performance of Multigrain Parallelization on 16 processor High-end Server IBM pSeries690

- IBM XL Fortran for AIX Version 8.1
  - Sequential execution : -O5 -qarch=pwr4
  - Automatic loop parallelization : -O5 -qsmp=auto -qarch=pwr4
  - OSCAR compiler : -O5 -qsmp=noauto -qarch=pwr4  
(su2cor: -O4 -qstrict)

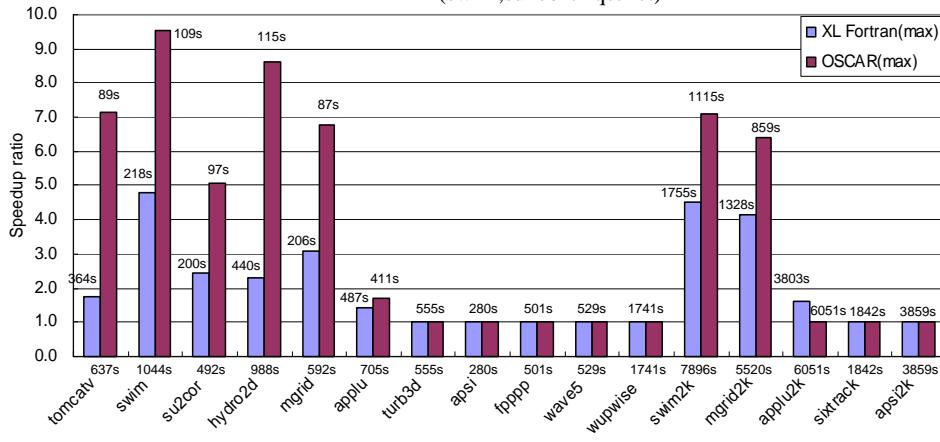


# Performance of Multigrain Parallel Processing for 102.swim on IBM pSeries690



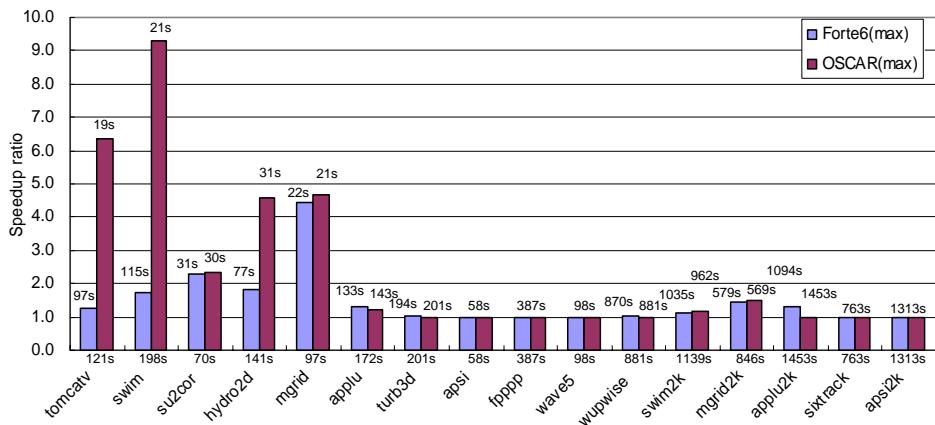
# Performance of Multigrain Parallelization on 8 processor Server IBM RS6000

- IBM XL Fortran for AIX Version 7.1
  - Sequential execution : -O5 -qarch=ppc
  - Automatic loop parallelization : -O5 -qsmp=auto -qarch=ppc
  - OSCAR compiler : -O3 -qsmp=noauto -qhot -qarch=ppc  
(swim,su2cor : -qstrict)



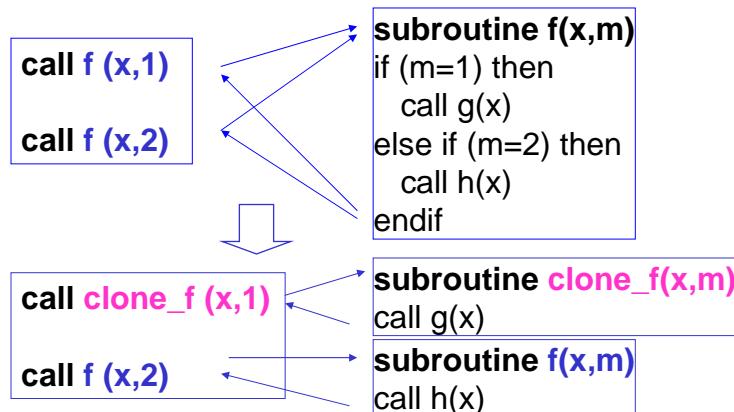
# Performance of Multigrain Parallelization on 4 processor workstation Sun Ultra80

- Sun Forte Developer 6 Update 2
  - Sequential execution : -fast
  - Automatic loop parallelization : -fast -autopar -reduction -stackvar
  - OSCAR compiler : -fast -explicitpar -mp=openmp -stackvar



# Aggressive Constant Propagation

It makes clone procedures,  
propagates constants beyond procedure boundaries,  
and evaluates expressions at compile time.



# Array Region Analysis

It uses a system of linear inequalities  
that enables the representation of complex regions  
such as triangular or stride regions.

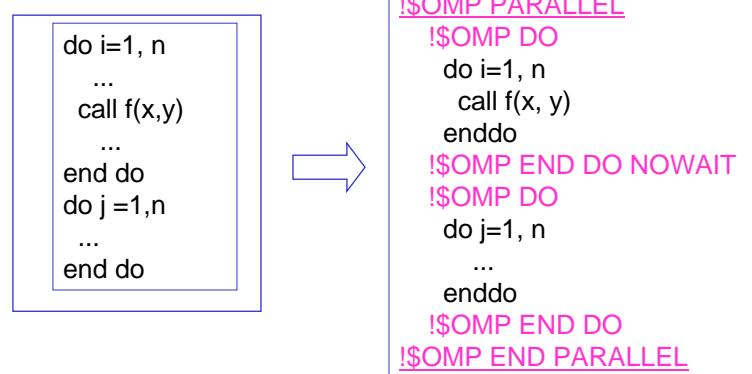
```
do j=1, n  
    do i = j, n  
        x(i, j) = ...
```



{ (i, j) | 1 ≤ i ≤ n, 1 ≤ j ≤ n, j ≥ i }  
region representation for x

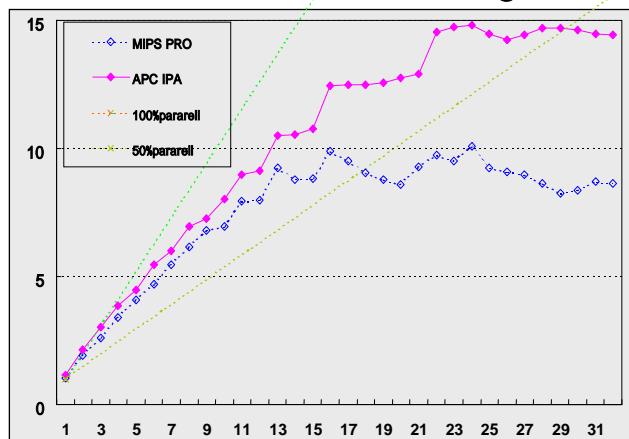
# Interprocedural Parallelization

It can parallelize the loops including procedure calls and enclose contiguous parallel loops in one parallel region.



## Performance of Interprocedural Analysis

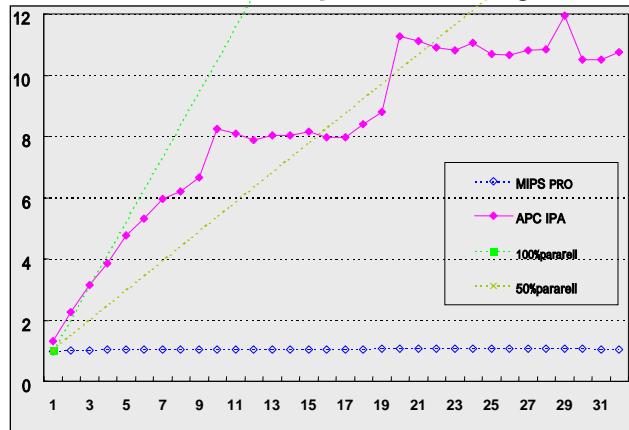
SPEC CFP95 turb3d on Origin 2000



- Interprocedural Analysis in APC compiler
- MIPSPro Fortran V.7.30

## Performance of Interprocedural Analysis

### SPEC CFP 2000 wupwiseon Origin 2000



■ **Interprocedural Analysis in APC compiler**  
■ **MIPSPro Fortran V.7.30**

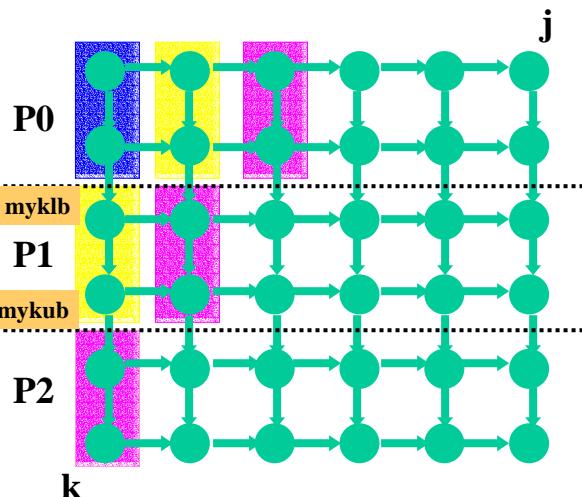
## Affine Partitioning [Lim & Lam97]

- The followings are considered at the same time
  - parallelization
  - improve data locality
  - reduce synchronization overhead
- A lot of transformations can be done automatically
- Extract pipeline parallelism

## Generated pipelined code in OpenMP

Consider parallel execution order

```
do j = 1, ny-2  
    waitPrev()  
    do k = myklb, mykub  
        ...  
    enddo  
    signalToNext()  
enddo
```



## HP Alpha Server (8 processors)

### The Alpha Server GS160 Model 6/73

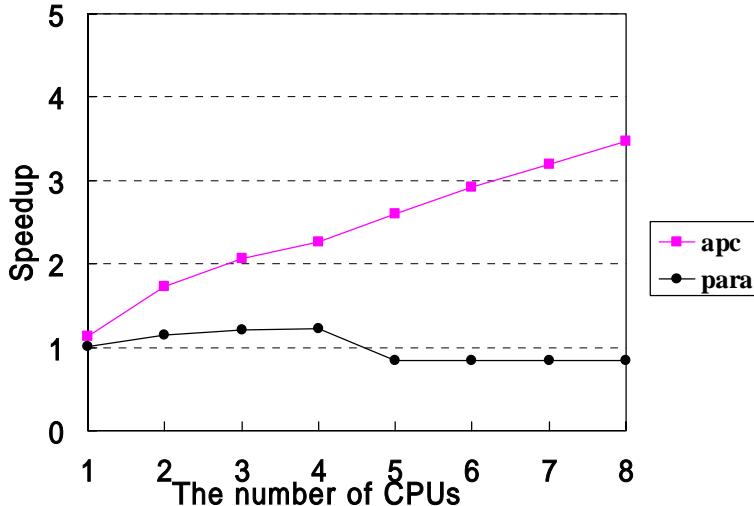
- Alpha 21264 (731MHz) × 8  
(The cc-NUMA machine in which each unit has 4 processors)
- L1-Cache (on-chip)
  - I-Cache 64KB
  - D-Cache 64KB(2-way)
- L2-Cache (direct-map, off-chip) 4MB
- Memory 4GB

### The Alpha Digital Fortran Compiler

- compile options:

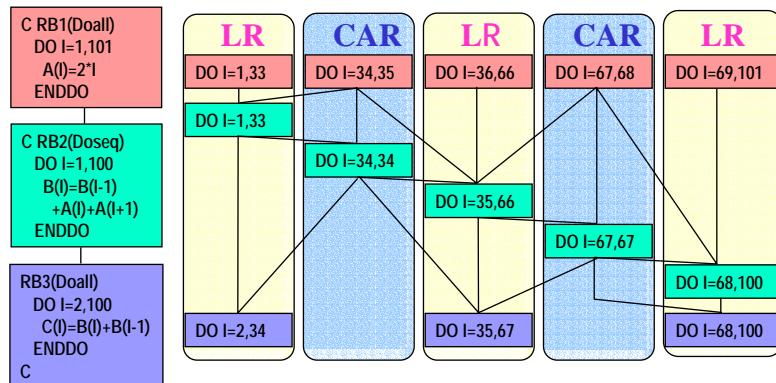
```
parallelized code: -v -arch ev6 -O5 -fkapargs=' -conc -ur=1'  
sequential code: -v -arch ev6 -O5 -fkapargs=' -ur=1'
```

# Speedup of applu on HP Alpha Server (8 Processors) by Affine Partitioning

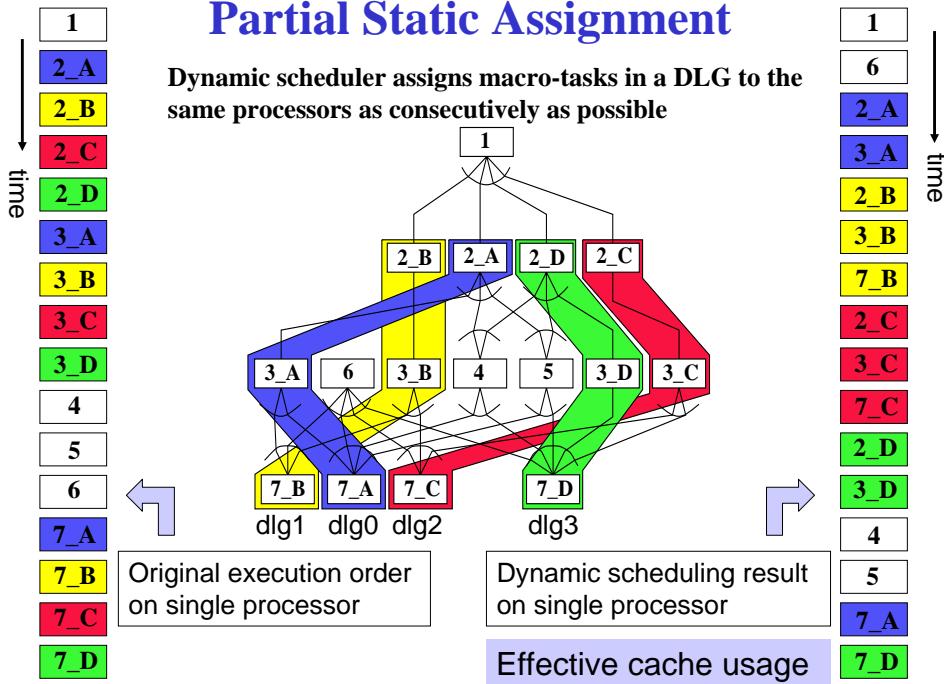


## Data-Localization Loop Aligned Decomposition

- Decompose multiple loop (Doall and Seq) into CARs and LRs considering inter-loop data dependence.
  - Most data in LR can be passed through LM.
  - LR: Localizable Region, CAR: Commonly Accessed Region



# Partial Static Assignment



## An Example of Data Localization for Spec95 Swim

```

DO 200 J=1,N
DO 200 I=1,M
    UNEW(I+1,J) = UOLD(I+1,J)+  

1   TDTSD8*(Z(I+1,J+1)+Z(I+1,J))*(CV(I+1,J+1)+CV(I,J+1)+CV(I,J)  

2   +CV(I+1,J))-TDTSDX*(H(I+1,J)-H(I,J))  

    VNEW(I,J) = VOLD(I,J)-TDTSD8*(Z(I+1,J+1)+Z(I,J+1))  

1   *(CU(I+1,J+1)+CU(I,J+1)+CU(I,J)+CU(I+1,J))  

2   -TDTSDY*(H(I,J)-H(I,J))
    PNEW(I,J) = POLD(I,J)-TDTSDX*(CU(I+1,J)-CU(I,J))  

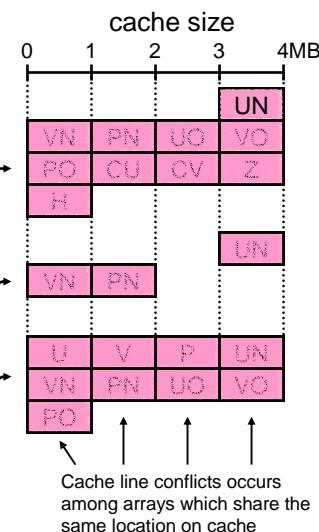
1   -TDTSDY*(CV(I,J+1)-CV(I,J))
200 CONTINUE

DO 210 J=1,N
    UNEW(1,J) = UNEW(M+1,J)
    VNEW(M+1,J+1) = VNEW(1,J+1)
    PNEW(M+1,J) = PNEW(1,J)
210 CONTINUE

DO 300 I=1,M
DO 300 J=1,N
    UOLD(I,J) = U(I,J)+ALPHA*(UNEW(I,J)-2.*U(I,J)+UOLD(I,J))
    VOLD(I,J) = V(I,J)+ALPHA*(VNEW(I,J)-2.*V(I,J)+VOLD(I,J))
    POLD(I,J) = P(I,J)+ALPHA*(PNEW(I,J)-2.*P(I,J)+POLD(I,J))
300 CONTINUE

```

(a) An example of target loop group for data localization



(b) Image of alignment of arrays on cache accessed by target loops

# Data Layout for Removing Line Conflict Misses by Array Dimension Padding

Declaration part of arrays in spec95 swim

## before padding

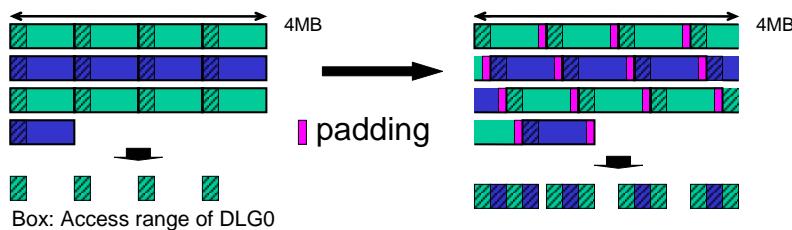
PARAMETER (N1=513, N2=513)

```
COMMON U(N1,N2), V(N1,N2), P(N1,N2),
*   UNEW(N1,N2), VNEW(N1,N2),
1   PNEW(N1,N2), UOLD(N1,N2),
*   VOLD(N1,N2), POLD(N1,N2),
2   CU(N1,N2), CV(N1,N2),
*   Z(N1,N2), H(N1,N2)
```

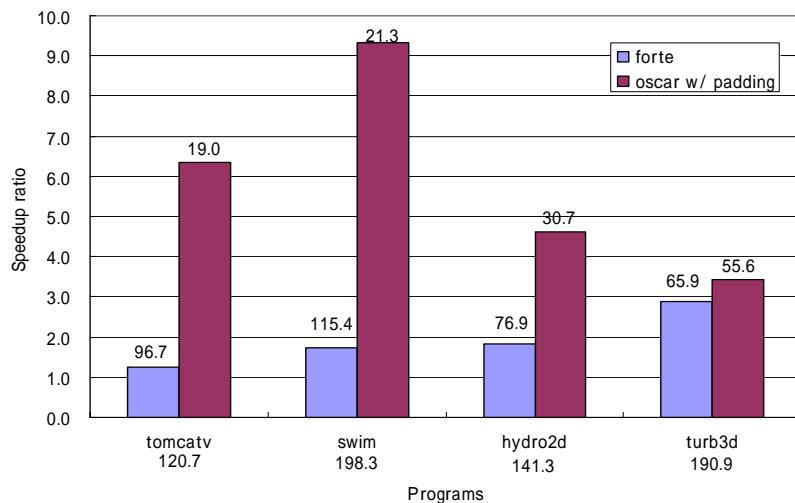
## after padding

PARAMETER (N1=513, N2=544)

```
COMMON U(N1,N2), V(N1,N2), P(N1,N2),
*   UNEW(N1,N2), VNEW(N1,N2),
1   PNEW(N1,N2), UOLD(N1,N2),
*   VOLD(N1,N2), POLD(N1,N2),
2   CU(N1,N2), CV(N1,N2),
*   Z(N1,N2), H(N1,N2)
```



## Performance of Cache Optimization with Padding on Sun Desktop WS Ultra80 (4 processors)



# First Touch Control Method for Distributed Shared Memory Systems

(a) Example program 1

```

1: real A(100)

21: c Initialization loop
22: !$omp parallel do
23: do i=1, 100
24:   A(i)=0
25: enddo
...
31: c Kernel loop
32: do itr=1, 10000
33: !$omp parallel do
34:   do i=41, 100
35:     A(i)= ...
36:   enddo
37: enddo

```

Inserted in the declaration part

(b) Conventional method

```

11: c Data distribution directive
12: c$distribute A(block)

```

(c) First touch control method

```

11: c FTC loop
12: !$omp parallel do
13: do i=41, 100
14:   A(i)=0
15: enddo

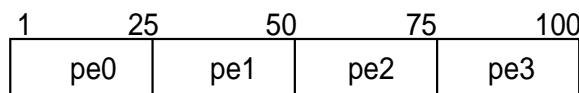
```

Inserted at the head of execution part

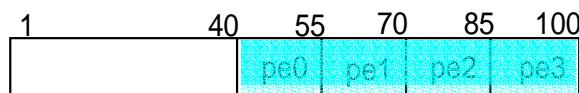
Generate the dummy loop to imitate reference patterns

## Image of Data Distribution by First Touch Control

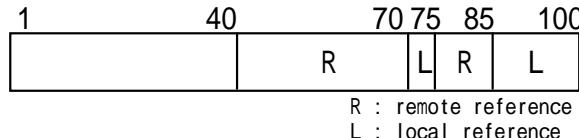
(a) conventional data distribution (both methods)



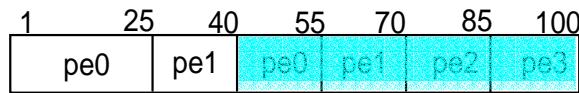
(b) reference region of kernel loop



(c) region of remote reference



(d) data distribution by our method



**100% local reference !!**

# Algorithm and Data of CG

```

!$omp parallel do
do j = 1, n
  do k = rowstr(j), rowstr(j+1)-1
    sum = sum + a(k) * p(colidx(k))
  enddo
  q(j) = sum
enddo

```

1	2	3	4	5	6	7	8
1.1							
2	2.2						
3		3.3		3.6		3.8	
4	4.1		4.4	4.6			
5	5.2		5.5				
6			6.6				
7				7.5	7.7		
8				8.4		8.7	8.8

(a) original sparse matrix

1.1							
2.2							
	3.3		3.6		3.8		
		4.4	4.6				
			5.5				
				6.6			
					7.5	7.7	
					8.4	8.7	8.8

(b) a

1							
2							
	3	6	8				
		4	6				
			5				
				10			
					11		
						13	
							16

(c) colidx

1	(1)						
2	(1)						
	5	(3)					
	8	(3)					
	10	(2)					
	11	(1)					
	13	(2)					
	16	(3)					

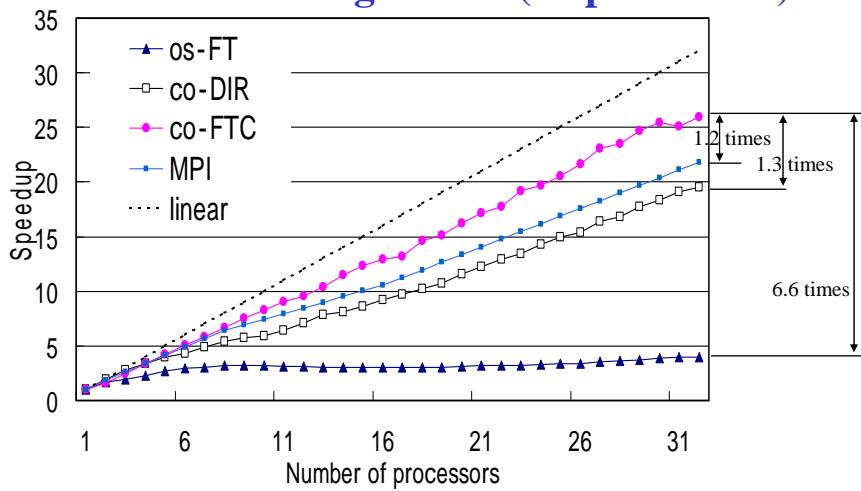
(d) rowstr

(e) on memory

[1.1|2.2|3.3|3.6|3.8|4.1|4.4|4.6|5.2|5.5|6.6|7.5|7.7|8.4|8.7|8.8]

(f) In the case of  
a block distribution

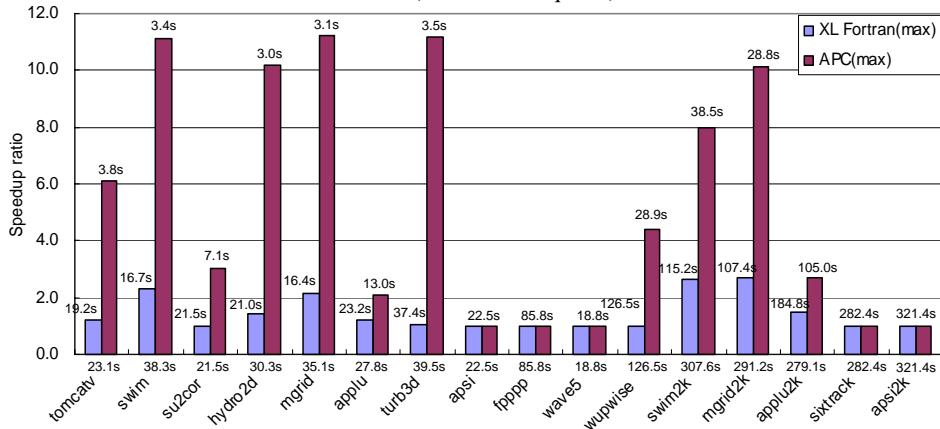
## Speedup for CG (class B) by First Touch Control on Origin 2000 (32 processors)



co-DIR: add data distrib. directive,

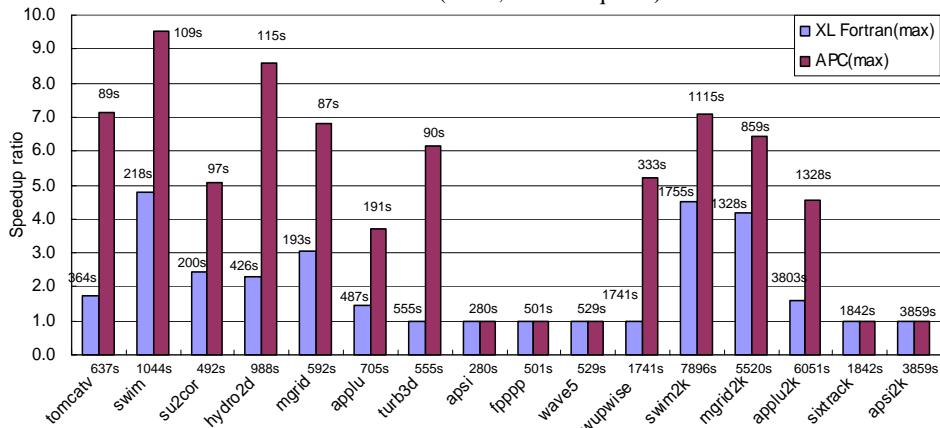
## Performance of APC Compiler on IBM pSeries690 16 Processors High-end Server

- IBM XL Fortran for AIX Version 8.1
  - Sequential execution : -O5 -qarch=pwr4
  - Automatic loop parallelization : -O5 -qsmp=auto -qarch=pwr4
  - OSCAR compiler : -O5 -qsmp=noauto -qarch=pwr4  
(su2cor: -O4 -qstrict)



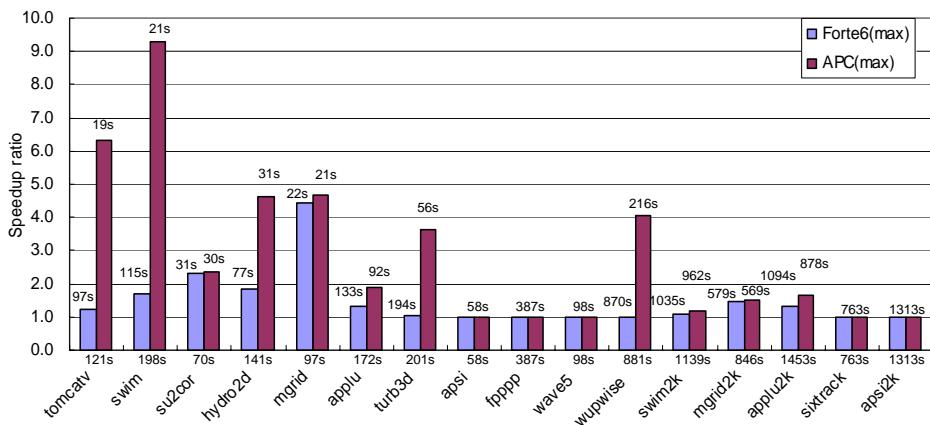
## Performance of APC Compiler on IBM RS6000 604e High-node 8 processor Server

- IBM XL Fortran for AIX Version 7.1
  - Sequential execution : -O5 -qarch=ppc
  - Automatic loop parallelization : -O5 -qsmp=auto -qarch=ppc
  - OSCAR compiler : -O3 -qsmp=noauto -qhot -qarch=ppc  
(swim,su2cor: -qstrict)



## Performance of APC Compiler on Sun Ultra80 4 Processor Workstation

- Sun Forte Developer 6 Update 2
  - Sequential execution : -fast
  - Automatic loop parallelization : -fast -autopar -reduction -stackvar
  - OSCAR compiler : -fast -explicitpar -mp=openmp -stackvar



<Target>

## APC Accomplishments

Double the performance compared with product compilers for SMPs at the September 2000 when APC was started.

<Results & Future>

- **3.5 times speedup in average and 10.7 times at the maximum for 16 FORTRAN programs in SPEC CFP95 and CFP2000 compared with the latest compiler on the latest 16 processor high-end SMP server.**
- **45 reviewed papers, 43 technical reports, 15 short papers, 8 patents, 9 invited talks, 1 invited survey paper, 8 newspaper articles, 13 Web news and 5 PhDs .**
- **Automatic multigrain parallelizing compiler will realize high productivity computing in various R&D field like Environments, Bio-informatics, Automobile, Aerospace, Financial Engineering etc.**
- **Some of the technologies will be incorporate into products.**
- **Multigrain parallelizing compiler will improve the cost performance, software and hardware productivity of chip multiprocessors to be introduced in SoC (System on Chip) like mobile phones, games, PDA, Digital TV and so on.**