

Medium Grain Level Parallelization and A New Array Contraction Technique



Akira Hosoi and Toshihiro Ozawa

**Advanced Parallelizing Compiler Project
Fujitsu LTD.**

1

Overview



- 1. Extracting Loop Level Pipeline Parallelism
and its Evaluation**

- 2. Partial Array Contraction
and Contraction-Oriented Loop Fusion**
Not implemented yet. But the methods are proposed to
improve the performance of applu, BT,
and SP drastically

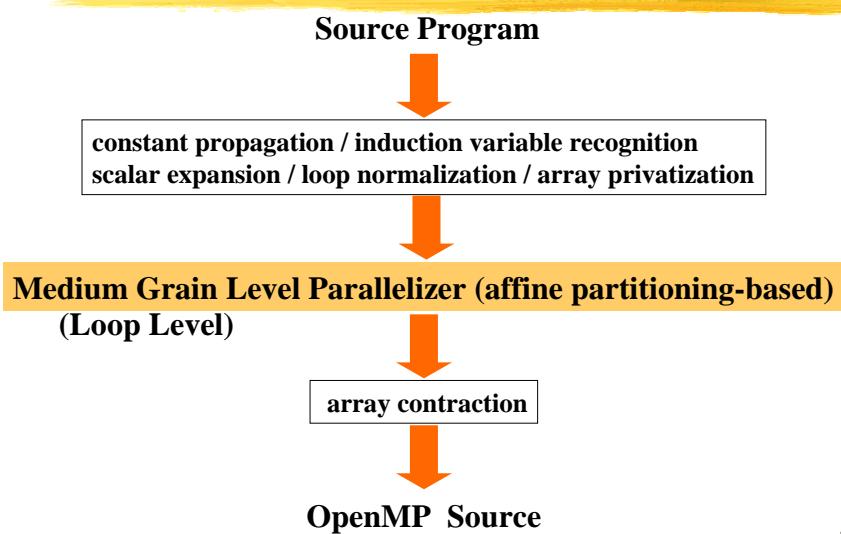
2

Extracting Loop Level Pipeline Parallelism and its Evaluation

1. The structure of the medium grain parallelizer
2. Affine Partitioning
 - 2.1. How to extract pipeline parallelism and its problem
 - 2.2. Refined Algorithm
3. How to implement pipeline parallelism in OpenMP
4. Evaluation

3

Structure of Medium Grain Level Parallelizer



4

Affine Partitioning [Lim & Lam97]

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

5

Pipeline Parallelism extracted by Affine Partitioning

Any imperfectly nested loop nests are transformed as follows:

all the assignment statements are surrounded by
as many fully permutable loops as possible



m loops

- $m-1$ dimensional pipeline parallel execution can be done
- m dimensional tiling can be done

6

Pipeline Parallelism extracted by Affine Partitioning (con't)

Example :

```
do i = 1, N
  do j = 1, i-1
    do k = 1, j-1
      a(i, j) = a(i, j)-a(i, k)*a(j, k)
    enddo
    a(i, j) = a(i, j) / a(j, j)
  enddo
  do k = 1, i-1
    a(i, i) = a(i, i)-a(i, k)*a(i, k)
  enddo
  a(i, i) = sqrt(a(i, i))
enddo
```

```
do i = 1, N
  do j = 1, i
    do k = 1, i
      if (j < i .and. k < j)
        a(i, j) = a(i, j)-a(i, k)*a(j, k)
      if (j < i .and. k == j)
        a(i, j) = a(i, j) / a(j, j)
      if (j == i .and. k < i)
        a(i, i) = a(i, i)-a(i, k)*a(i, k)
      if (j == i .and. k == i)
        a(i, i) = sqrt(a(i, i))
    enddo
  enddo
enddo
```

7

How to extract pipeline parallelism

1. Construct an inequality system $Ax \geq 0$ from array subscripts and loop bounds
2. Solve $Ax \geq 0$ in such a way that $\text{rank } A$ should be as large as possible ($\text{rank } A = \text{the number of fully permutable loops}$)

Problem of extracting pipeline parallelism

As the number of assignment statements in a loop nest increases a little, the solution space becomes very large



It takes a large amount of memory and compile time to solve the inequality system directly

8

Refined Algorithm

1. Assume the number of fully permutable loops in the transformed loop nest

the number of the common surrounding loops $\leq \text{rank } A \leq$ the maximum depth of in the original loop nest

```
do i = 1, N
    do j = 1, i-1
        do k = 1, j-1
            a(i, j) = a(i, j)-a(i, k)*a(j, k)
        enddo
        a(i, j) = a(i, j) / a(j, j)      → 1 <= rank A <= 3
    enddo
    do k = 1, i-1
        a(i, i) = a(i, i)-a(i, k)*a(i, k)
    enddo
    a(i, i) = sqrt(a(i, i))
enddo
```

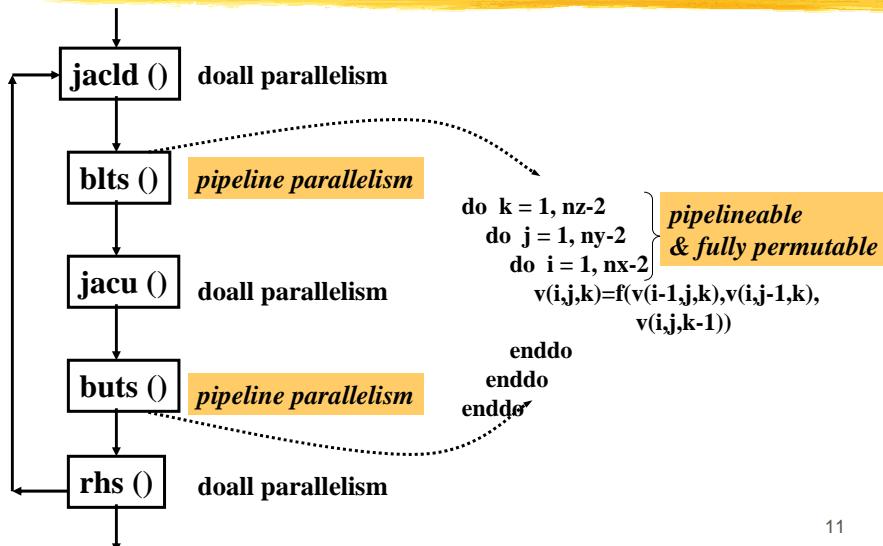
9

Refined Algorithm (con't)

2. Ignore the loop bounds of common surrounding loops to simplify the inequality system

```
do i =  $i_0, i_1$  ignore
    do j = LBj(i), UBj(i)
        A(f(i, j)) = ...
    enddo
    do k = LBk(i), UBk(i)
        ... = A(g(i, k))
    enddo
enddo
```

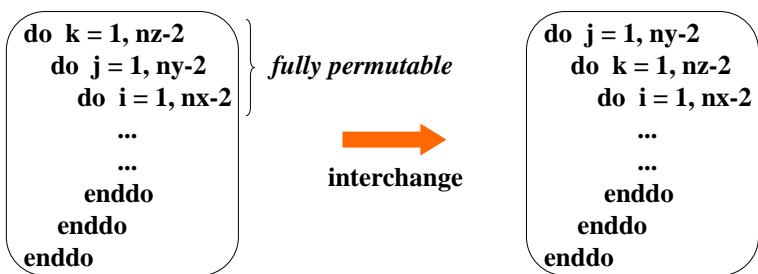
The main loop of SPEC CFP2000 / applu



11

How to generate pipelined code in openMP

1. Interchange



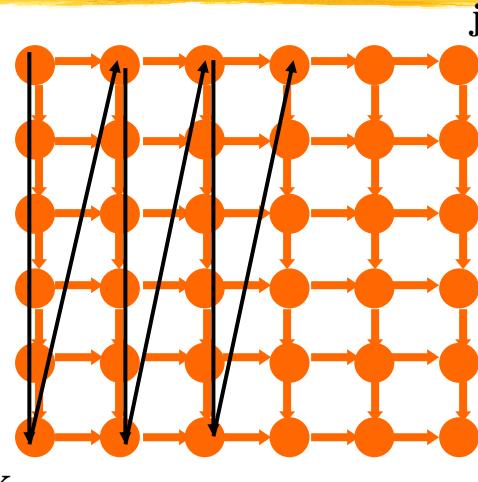
12

How to generate pipelined code in openMP (Cont.)

2. Consider sequential execution order

```
do j = 1, ny-2  
  do k = 1, nz-2  
    ...  
    ...  
  enddo  
enddo
```

- Iteration space
- data dependence
- sequential execution order

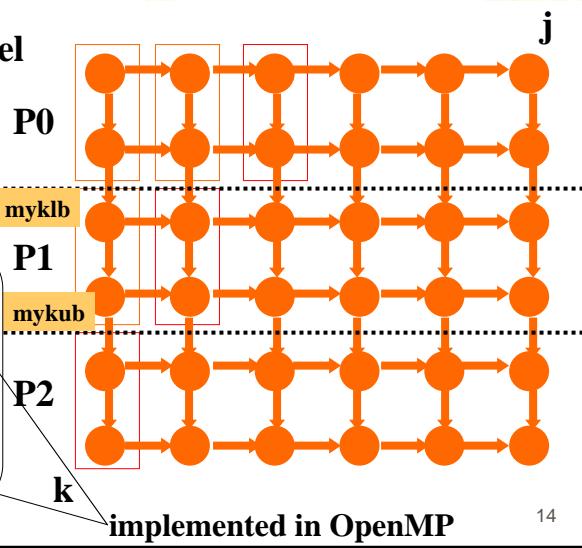


13

How to generate pipelined code in openMP (Cont.)

3. Consider parallel execution order

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



14

The Alpha Server

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

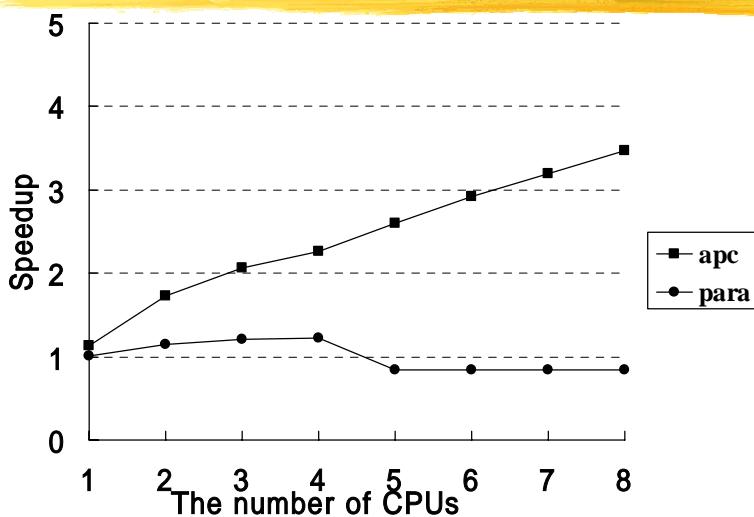
The Alpha Digital Fortran Compiler

- compile options:

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

15

Speedup of applu on the the Alpha Server



16

Conclusion

- Pipeline parallelism are automatically extracted from the complicated imperfectly nested loop
- Pipelined code is implemented in OpenMP
- The performance of SPEC CFP2000 / applu can be 2.5 times faster than that on Alpha Server

17

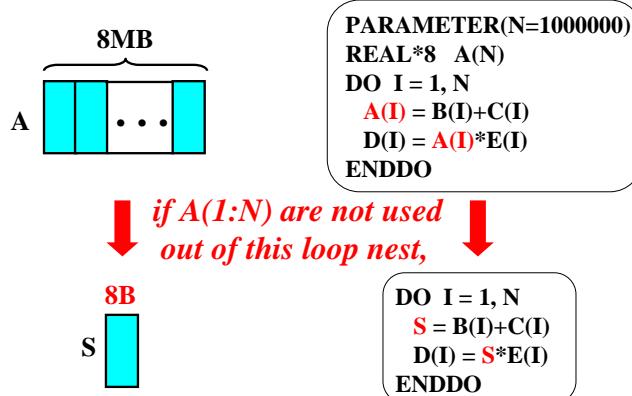
Partial Array Contraction and Contraction-Oriented Loop Fusion

- 1. Introduction**
- 2. Problems of the previous array contraction and its solutions**
- 3. Evaluation**
- 4. Conclusion**

18

Array Contraction

shrinks the size of an array and improves data locality

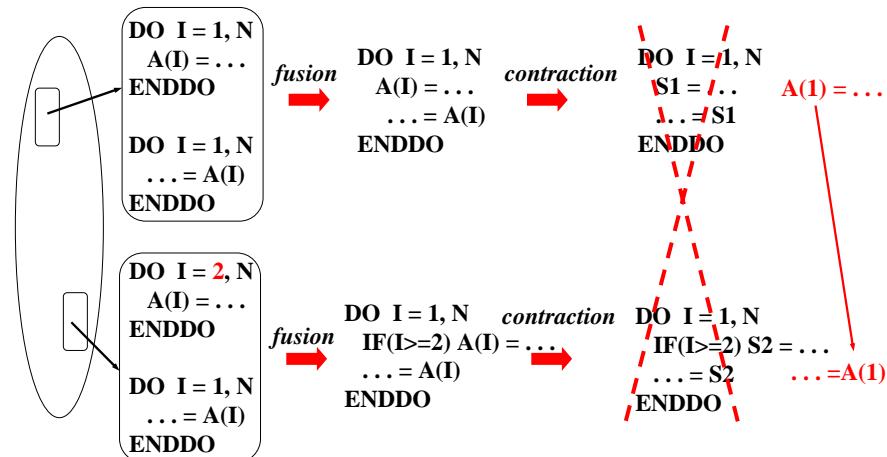


usually combined with *loop fusion*

19

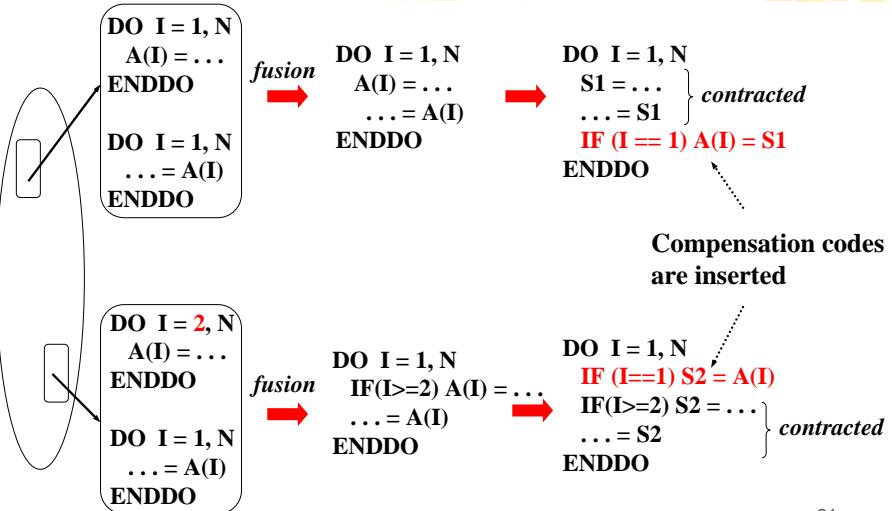
Problems of the previous array contraction(1)

Not so many opportunities of array contraction



20

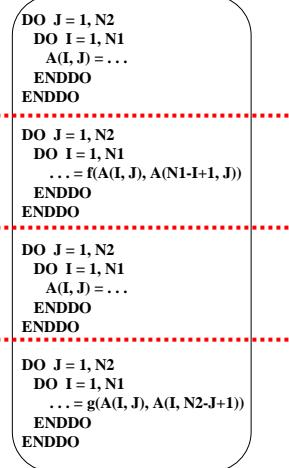
Partial Array Contraction with Save & Restore



21

Problems of the previous array contraction(2)

It is difficult to decide which loop nests should be fused



How we should split a loop nest sequence for better contraction ?

22

Definition: definition / use-dominated

- An array A is *definition-dominated* in a loop nest L :
if no element is used before defined in L .

ex.

```
DO I = 1, N  
A(I) = ...  
ENDDO
```

```
DO I = 1, N  
A(I) = ...  
... = A(I)  
ENDDO
```

- An array A is *use-dominated* in a loop nest L :
if no element is defined before used in L .

ex.

```
DO I = 1, N  
... = A(I)  
ENDDO
```

```
DO I = 1, N  
... = A(I)  
A(I) = ...  
ENDDO
```

23

How to split a loop nest sequence

If A is definition-dominated in L ,
 L becomes the first loop nest
of a new group

```
DO J = 1, N2  
DO I = 1, N1  
A(I, J) = ...  
ENDDO  
ENDDO  
  
DO J = 1, N2  
DO I = 1, N1  
... = f(A(I, J), A(N1-I+1, J))  
ENDDO  
ENDDO  
  
DO J = 1, N2  
DO I = 1, N1  
A(I, J) = ...  
ENDDO  
ENDDO  
  
DO J = 1, N2  
DO I = 1, N1  
... = g(A(I, J), A(I, N2-J+1))  
ENDDO  
ENDDO
```

24

Benchmark Programs

- SPEC CFP2000/applu
- NPB2.3-serial/SP CLASS A
- NPB2.3-serial/BT CLASS A

25

How to apply Our Method to SP

1. split



```
do 100 k = 1, n3-2
do 100 j = 1, n2-2
do 100 i = 1, n1-2
100   lhs(i, j, k, 1:15) = ...
      ...
do 110 k = 1, n3-2
do 110 j = 1, n2-2
do 110 i = 1, n1-2
110   lhs(i, j, k, 1:15) = lhs(i, j, k, 1:15) + ...
      ...
do 120 k = 1, n3-2
do 120 j = 1, n2-2
do 120 i = 0, n1-3
120   lhs(i, j, k, 1:15) = f0(lhs(i, j, k, 1:15),
      lhs(i+1, j, k, 1:15), lhs(i+2, j, k, 1:15))
      ...
do 130 k = 1, n3-1
do 130 j = 1, n2-1
do 130 i = n1-3, 0, -1
130   ... = g0(lhs(i, j, k, 1:15),
      lhs(i+1, j, k, 1:15),
      lhs(i+2, j, k, 1:15))
```

2. renaming

save & restore code are added

```
do k = 1, n3-2
do j = 1, n2-2
  lhs1(0, j, k, 1:15)=lhs(0, j, k, 1:15)
  lhs1(n1-1, j, k, 1:15)=lhs(n1-1, j, k, 1:15)
enddo
enddo

do k = 1, n3-2
do j = 1, n2-2
  lhs(0, j, k, 1:15)=lhs1(0, j, k, 1:15)
  lhs(n1-1, j, k, 1:15)=lhs1(n1-1, j, k, 1:15)
enddo
enddo
```

26

How to apply Our Method to SP (con't)

3. fusion & 4. partial contraction

```

do k = 1, n3-2
do j = 1, n2-2
  lhs1(0, 1:15)=lhs(0, j, k, 1:15)
  lhs1(n1-1, 1:15)=lhs(n1-1, j, k, 1:15)
  do i = 1, n1-2
    lhs1(i, 1:15) = ...
  enddo
  do i = 1, n1-2
    lhs1(i, 1:15)=lhs1(i, 1:15) + ...
  enddo
  do i = 0, n1-3
    lhs1(i, 1:15)=f0(lhs1(i, 1:15),
    lhs1(i+1, 1:15), lhs1(i+2, 1:15))
  enddo
  do i = n1-3, 0, -1
    ... = g0(lhs1(i, 1:15),
    lhs1(i+1, 1:15),
    lhs1(i+2, 1:15))
  enddo
  lhs(0, j, k, 1:15)=lhs1(0, 1:15)
  lhs(n1-1, j, k, 1:15)=lhs1(n1-1, 1:15)
enddo
enddo

```

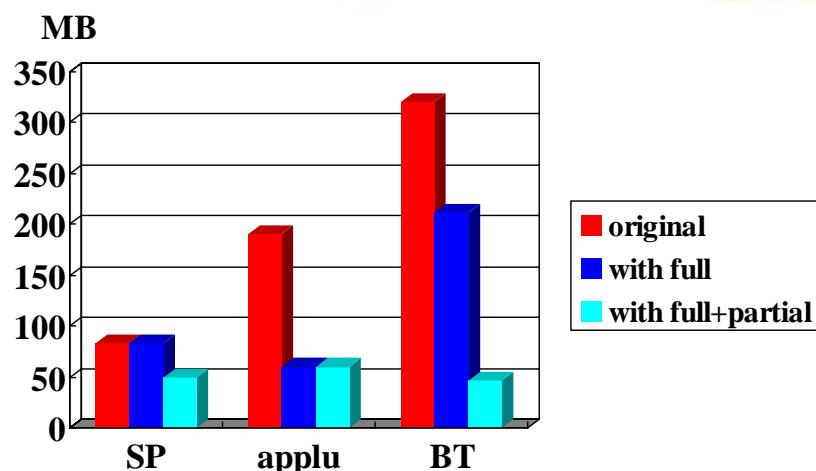
27

lhs(65, 65, 65, 15) / 33MB

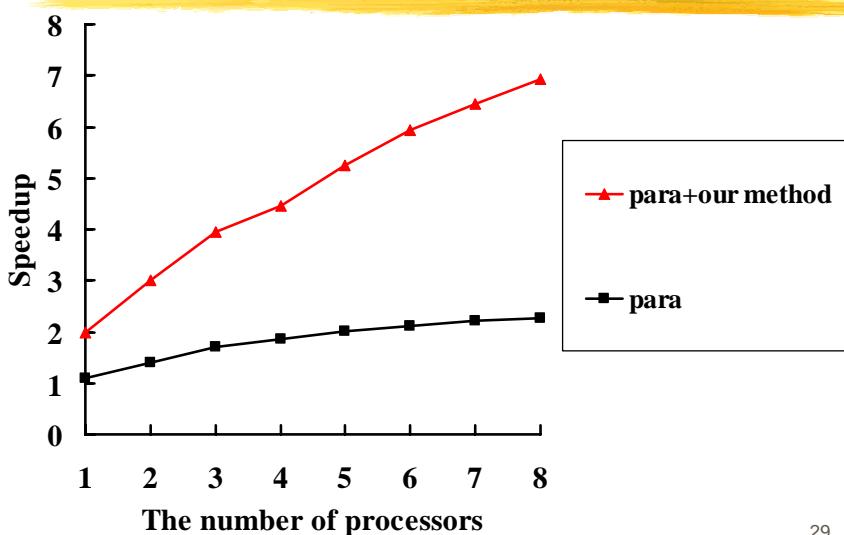
↓
lhs1(65, 15) / 8KB

partially contracted

Array data size without/with contraction

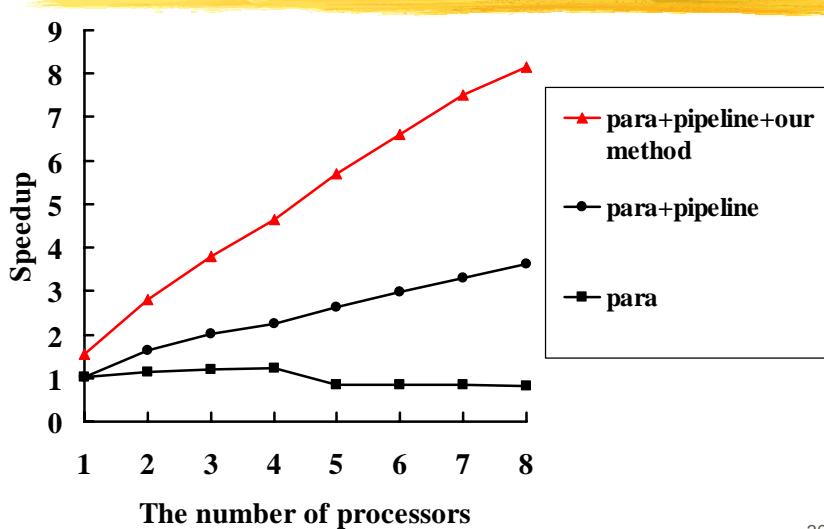


Speedup of SP on the Alpha Server



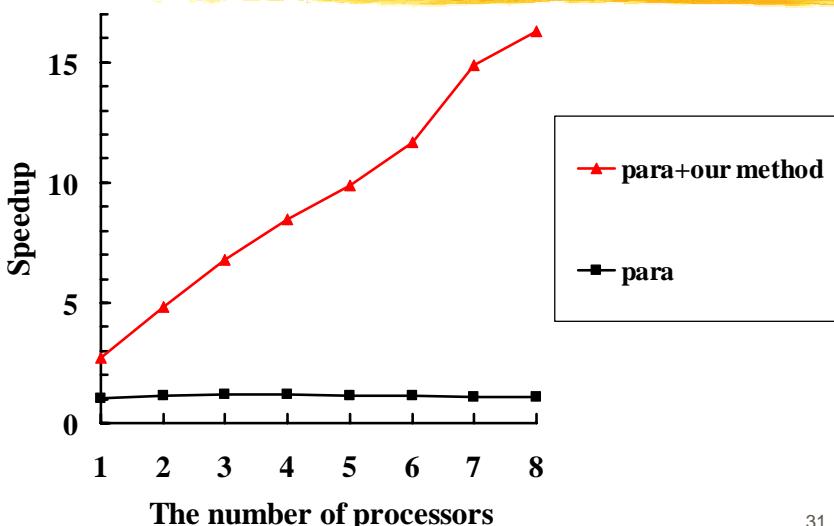
29

Speedup of applu on the Alpha Server



30

Speedup of BT on the Alpha Server



31

Conclusion

- We generalize array contraction and introduce *partial array contraction with save & restore*
→ We have more opportunities of array contraction
- We split a loop nest sequence using *definition / use-dominated*
→ Array contraction can be applied more effectively
- The performance of SP, applu, and BT are drastically improved

32