

SCOPES 2003

Tailoring Software Pipelining For Effective Exploitation Of Zero Overhead Loop Buffer

Gang-Ryung Uh
CS Department
Boise State University

Outline

1. Low-power DSP 16000 and ZOLB
2. Compiler Mission
3. Conventional Approach
4. Alternative approach
5. Intermediate Results
6. Conclusion

Signal Processing Algorithm


$$y_k = \sum_{n=0}^{N-1} x_n w_{k-n}$$
$$y_k = \sum_{n=0}^{N-1} x_n w_{k-n} \quad \text{for } 0 \leq k < N$$
$$y_k = \sum_{n=0}^{N-1} x_n w_{k-n} \cos(2\pi f_0 n) \cos(2\pi f_1 n)$$
$$y_k = \sum_{n=0}^{N-1} x_n w_{k-n} \cos(2\pi f_0 n) \cos(2\pi f_1 n)$$

- I. **Heavy arithmetic** computations
- II. **Can be easily programmed into Tight Small Loops**

DSP (Digital Signal Processor)

- Programmable processor for mathematical operations to manipulate signals with

1. **High performance,**
2. **Minimal power consumption**
3. **Minimal memory footprint**



Finite Impulse Response (FIR)


At the least, compute one **Tap** in a **Single Cycle**

Lucent DSP16000 Architecture Features

1. **Havard Architecture**
2. **Separate AGU from DALU for rich addressing modes**
3. **Zero-wait State High Speed Memory**


Lucent DSP16000 Architecture Features (cont)

4. Compiler (Programmer) Controlled On-Core Instruction Cache – ZOLB (Zero Overhead Loop Buffer) to support high performance with minimal power dissipation



Lucent DSP16000 Instruction Set Design

In order to achieve **performance** & **higher code density**



- Permissible order of operations is very limited
- The register usage is restricted to only a few different registers

Compiler Mission!

Where are the compound/complex instructions?

```
// EDN Benchmarks
fir(const short array1[],
  const short coeff[],
  short output[])
{
  int i,j,sum;
  for(i=0;i < N-ORDER;i++){
    sum=0;
    for(j=0; j < ORDER; j++){
      sum += array1[i+j]*coeff[j];
    }
    output[i]=sum>>15;
  }
}
```

```
A2=0
j=a4
do 50 {
  /* inst 1 */ xh = *(r0 + j)
  /* inst 2 */ yh = *r3++
  /* inst 3 */ r4 = j
  /* inst 4 */ p0 = xh*yh p1 = xl*yl
  /* inst 5 */ a2 = a2+p0
  /* inst 6 */ j = r4+1
}
```

A0 = A0 + P0 P0 = Xh*Yh P1 = Xl * Yl Y = *R0++ X = *PT0++

Experience with Iterative Modulo Scheduling Techniques

EDN Benchmark: FIR Filter

```
fir(const short array1[], const short coeff[], short output[])
{
  int i,j,sum;
  for(i=0;i < N-ORDER;i++){
    sum=0;
    for(j=0; j < ORDER; j++){
      sum += array1[i+j]*coeff[j];
    }
    output[i]=sum>>15;
  }
}
```

```
a2=0
j=a4
do 50 {
  /* inst 1 */ xh = *(r0 + j)
  /* inst 2 */ yh = *r3++
  /* inst 3 */ r4 = j
  /* inst 4 */ p0 = xh*yh p1 = xl*yl
  /* inst 5 */ a2 = a2+p0
  /* inst 6 */ j = r4+1
}
```

Step 1: Resource Initiation Interval

MII = MAX(RecII, ResII)

ResII : Smallest Loop Initiation Interval to meet the system resource requirement

```
do 50 {
  /* inst 1 */ xh = *(r0 + j)
  /* inst 2 */ yh = *r3++
  /* inst 3 */ r4 = j
  /* inst 4 */ p0 = xh*yh p1 = xl*yl
  /* inst 5 */ a2 = a2+p0
  /* inst 6 */ j = r4+1
}
```

ResII: Resource Initiation Interval **2**




Step 2: Recurrence Initiation Interval

MII = MAX(RecII, ResII)


RecII : Smallest Integer Loop Initiation Interval to meet all the deadlines imposed by data dependence circuits.

```
do 50 {
  /* inst 1 */ xh = *(r0 + j)
  /* inst 2 */ yh = *r3++
  /* inst 3 */ r4 = j
  /* inst 4 */ p0 = xh*yh p1 = xl*yl
  /* inst 5 */ a2 = a2+p0
  /* inst 6 */ j = r4+1
}
```



- True Dependence
- Output Dependence
- Anti Dependence

Step 2: Recll (cont)



Adjacency Matrix

Start	Inst1	Inst2	Inst3	Inst4	Inst5	Inst6	End
Start	X	(0)	(0)	(0)	(0)	(0)	(0)
Inst1	X	X	X	X	(0)	X	X
Inst2	X	X	X	X	(0)	X	X
Inst3	X	X	X	X	X	X	(0)
Inst4	X	(0)	(0)	X	X	(0)	X
Inst5	X	X	X	X	(0)	X	X
Inst6	X	(0)	X	(0)	X	X	X

Step 2: Compute MinDIST Matrix

Adjacency Matrix

Start	Inst1	Inst2	Inst3	Inst4	Inst5	Inst6	End
Start	X	(0)	(0)	(0)	(0)	(0)	(0)
Inst1	X	X	X	X	(0)	X	X
Inst2	X	X	X	X	(0)	X	X
Inst3	X	X	X	X	X	(0)	(0)
Inst4	X	(0)	(0)	X	X	(0)	X
Inst5	X	X	X	X	(0)	X	X
Inst6	X	(0)	X	(0)	X	X	(0)

Floyd Algorithm:
 $MinDist[i,j] \leq 0$
 with II (Initiation Interval) 2

Start	Inst1	Inst2	Inst3	Inst4	Inst5	Inst6	End
Start	X	0	0	0	1	2	1
Inst1	X	-1	-1	X	1	2	X
Inst2	X	-1	-1	X	1	2	X
Inst3	X	0	-1	0	1	2	1
Inst4	X	(0)	(0)	X	X	(0)	X
Inst5	X	-4	-4	X	-2	-1	X
Inst6	X	-1	-2	-1	0	1	0

Step 3: Slack Scheduling by computing Estart and Lstart

Floyd Algorithm:
 $MinDist[i,j] \leq 0$
 with II (Initiation Interval) 2

Legal Partial Schedule based on Estart and Lstart

Start	Inst1	Inst2	Inst3	Inst4	Inst5	Inst6	End
Start	X	0	0	1	2	1	2
Inst1	X	0	-1	X	1	2	X
Inst2	X	-1	0	X	1	2	X
Inst3	X	0	-1	0	1	2	1
Inst4	X	-2	-2	X	0	1	X
Inst5	X	-4	-4	X	-2	0	X
Inst6	X	-1	-2	-1	0	1	0
End	X	X	X	X	X	X	X

Operation	Estart	Lstart	Issue Time
Inst-1	0	1	0
Inst-2	0	1	1
Inst-3	0	1	0
Inst-4	1	1	1
Inst-5	0	1	1
Inst-6	1	1	1

Why Modulo Scheduling is not suitable?

Legal Partial Schedule based on SLACK

Operation	Estart	Lstart	Issue Time
Inst-1	0	1	0
Inst-2	0	1	1
Inst-3	0	1	0
Inst-4	1	1	1
Inst-5	0	1	1
Inst-6	1	1	1

```
// inst-1 && inst-3
xh=(r0+j)    r4=j
// inst-2 && inst-4 && inst-5 && inst-6
yh=*r3++ p0=xh*yh p1=x1*y1 a2=a2+p0 j=r4+1
```

No Legal Encoding

Why Modulo Scheduling is not suitable?

Due to limited encoding space, DSP16000 compound instructions that account for $\{Inst-i, Inst-j, Inst-k\}$, but there is NO legal encoding to capture any subset of $\{Inst-i, Inst-j, Inst-k\}$

How to Overcome?

- Software pipelining optimization must be sensitive to **Instruction Selection**
- This requires that the Instruction selection performs the following tasks in a demand driven manner
 - proactively perform **Register Renaming**
 - proactively introduce additional **micro-operations** on the fly

New Compiler Strategy

- 1. **Block Partitioning** to group instructions G_1, G_2, \dots, G_n that instructions can be parallelized in G_k for $k=(i+1)(i+2)\dots n$
- 2. **Block Structure** to group instructions that can be parallelized in a single instruction template
- 3. **Block Partitioning** to group instructions that can be parallelized in a single instruction template

Potentially Pipelineable

Instructions I_i and I_j are **potentially pipelineable** only when the following two conditions can be met.

1. There exists a compound/complex instruction template that can hold (may be more) both effects I_i and I_j in parallel. This implies that I_i and I_j can be potentially combined into a single complex instruction.
2. The Distance from I_i in the instruction Group G_k to I_j can meet the Minimum Distance requirements (**MinDist(i,j)**).

Another FIR Filter from a Customer

Instruction	DF	CF	EF	LF	MF	NF	OF	PF	QF
Str	0X	0	0	0	0	0	0	0	0
Ie1	0X	0	0	0	0	0	0	0	0
Ie2	0X	0	0	0	0	0	0	0	0
Ie3	0X	0	0	0	0	0	0	0	0
Ie4	0X	0	0	0	0	0	0	0	0
Ie5	0X	0	0	0	0	0	0	0	0
Ie6	0X	0	0	0	0	0	0	0	0
Ie7	0X	0	0	0	0	0	0	0	0
Ie8	0X	0	0	0	0	0	0	0	0
Ie9	0X	0	0	0	0	0	0	0	0

Another FIR Filter from a Customer

MinDist = 6
MinDist[i][j] Matrix


Instruction	DF	CF	EF	LF	MF	NF	OF	PF	QF
Str	X	0	1	1	2	3	4	5	6
Ie1	X	-5	1	1	2	3	4	5	6
Ie2	X	-6	-2	0	1	2	3	4	5
Ie3	X	-8	-2	0	1	2	3	4	5
Ie4	X	-9	-3	-1	0	1	3	3	4
Ie5	X	-10	-4	-2	-1	0	1	2	3
Ie6	X	-11	-5	-3	-2	-1	0	1	2
Ie7	X	-12	-6	-4	-3	-2	-1	0	1
Ie8	X	-13	-7	-5	-4	-3	-2	-1	0
Ie9	X	X	X	X	X	X	X	X	X

Step 1: Partition

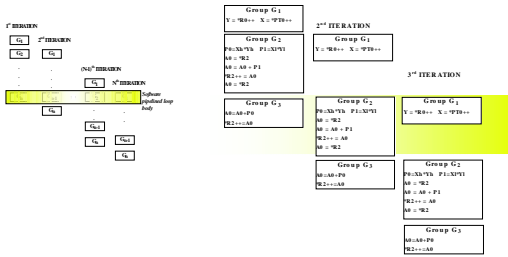
```

/* initialization */
void partition() {
    create a new group G1;
    add an instruction I1 to G1;
    i = 2;
    j = 1;
    FOR each instruction Ii in the loop DO {
        FOR each instruction Ij in Gj {
            IF (Ii and Ij are potentially pipelineable) {
                j = j + 1;
                create a new group Gk;
                Tag Gk and Gj with complex instruction templates;
                break;
            }
        }
        add an instruction Ii to Gj;
        i = i + 1;
    }
}
    
```

Step 2: Project a Maximally Pipelined Loop



Step 2: Project a Perfectly Pipelined Loop

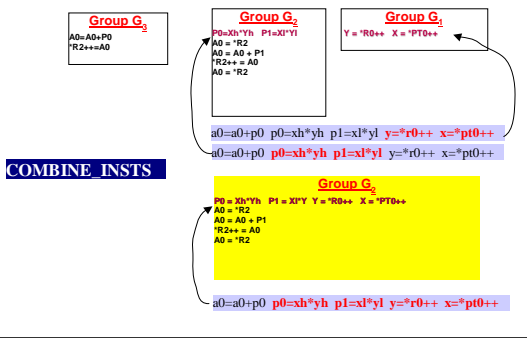


Instruction Selection Algorithm

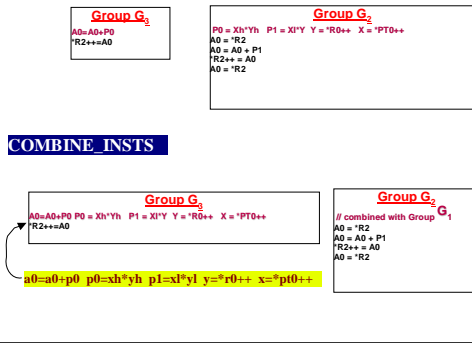
```

MAIN Y
DO
  change = FALSE;
  FOR each instruction group G DO (
    IF (L == NULL)
      CONTINUE;
    F (change == FALSE)
      change = COMBINE_INSTRS(L, G);
    ELSE
      (void) COMBINE_INSTRS(L, G);
      "Advance to the next instruction group"
      L = L + 1;
  ) "end of FOR"
  WHILE (change == TRUE)
    "COMBINE INSTRUCTIONS"
    BOOLEAN Success = FALSE;
    IF (G == NULL)
      "Gi is empty"
      return FALSE;
    F first instruction in Gi
    FOR each instruction Ii in Gi DO (
      IF (there exists FI,FIi instruction template that
        accounts for effects Ii)
        (
          IF (Scheduling Ii at the same time slot for Ii
            does violate a deadline imposed
            by some instruction in Gj)
            "It is not legal schedule"
            continue;
          IF (Ii satisfies the register encoding restrictions)
            replace Ii with Ii
            remove Ii from Gi;
            Success = TRUE;
            break;
          ELSE IF (there exist available register(s) that
            can make Ii satisfy the register
            encoding restrictions)
            perform register renaming;
            tag the Ii with the FI,FIi instruction template;
            remove Ii from Gi;
            Success = TRUE;
            break;
        )
      )
    )
  ) "END FOR"
  F = F + 1;
) "END FOR"
Merge G3 and G2 into one instruction group;
RETURN FALSE;
) "END COMBINE_INSTRS"
  
```

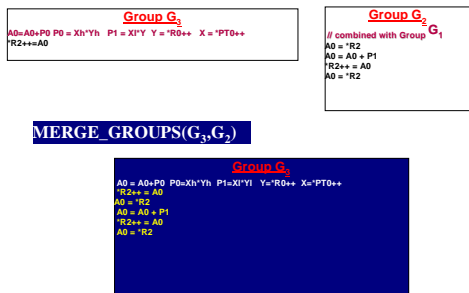
Discover Complex Instruction by Overlapping G₂ and G₁



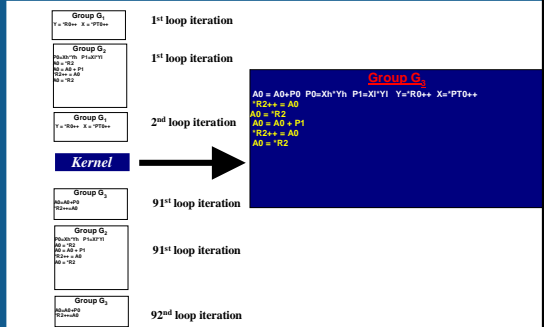
Discover Complex Instruction by Overlapping G₃ and G₂



Merge Instruction Groups G₃ and G₂



Restructured Loop




Results

Buhak Est Dugan	
Dugan	Duofan
at8	Atbo8strings
anskin	Ansksin
arp8	Arp8stringsumfir
fl	FlstringsFFT
fr	Frstringsmedfir
frjstld	Frjstldstringsfir
fr	Frstrings
inns8	Inns8strings
jgnt	Jgntstrings
sak8	Sak8strings
srbs8	Srbs8strings
scpp	Scppstrings


Execution Time

Table 2. Impact on Execution Time



Code Size

Table 3. Impact on Code Size



Conclusion

1. New Compiler Strategy that automatically exploits compound instructions
2. As a result of this work, the Zero Overhead Loop Buffer on Lucent DSP16000 can be further exploited

- SCOPES 2003
- Outline
- Signal Processing Algorithm
- DSP (Digital Signal Processor)
- At the least, compute one Tap in a Single Cycle
- Lucent DSP16000
- Lucent DSP16000
- Example of Using ZOLB on the DSP16000
- Instruction Set Design for Low-power Embedded Processors
- Instruction Set Design for Low-power Embedded Processors