



REVIEW: multiprocessor scheduling, periodic tasks

	τ_1	τ_2	τ_3
T_i	4	8	6
C_i	2	8	3

$$U = \frac{2}{4} + \frac{8}{8} + \frac{3}{6} = 2$$


$$T = \text{gcd}(4, 8, 6) = 2$$

In each slice,

$$T_1 \text{ has } 2 \cdot \frac{2}{4} = 1 \text{ unit}$$

$$T_2 \text{ has } 2 \cdot \frac{8}{8} = 2 \text{ units}$$

$$T_3 \text{ has } 2 \cdot \frac{3}{6} = 1 \text{ unit}$$



REVIEW: Task migration time

Lemma: If $U \leq n$, then **within each time slice** the tasks can meet the migration time requirement without missing deadlines, if the task migration time is one unit.

- Sort tasks according to non-increasing computation time
- If computation block = T
 \Rightarrow allocate processor exclusively.
- If computation time is $< T$:
 - allocate completely on 1 processor if possible \Rightarrow no migration
 - allocate part of computation at the end of proc i , rest at beginning of proc $i+1$
 \Rightarrow gap of at least 1.

BF - ES

- 3 -

REVIEW: Task migration time

Lemma: If $U \leq n$, then **between time slices** the tasks can meet the migration time requirement without missing deadlines, if the task migration time is one unit.

- For each slice, sort tasks according to non-increasing computation blocks
- If computation block = $T \rightarrow$ find processor that executed task at the end of the previous slice
 \rightarrow no migration
- if no such processor exists \Rightarrow assign to same left-over processor at the end (migration time already accounted for in previous slice)


BF - ES

- 4 -

- If computation block $< T$
 - => Find processor j that executed task at the end of the previous slice
 - Assign as much as possible to current processor;
 - if insufficient, use j from beginning. (no migration at beginning, ≥ 1 unit gap within slice)
- If no such processor exists => assign task later (migration time already accounted for in the previous slice).

Example (4 processors)

	Computation block	
τ_1	10	$T=10$
τ_2	9	
τ_3	9	
τ_4	9	
τ_5	3	



Extension: Task migration time

Theorem: Let $T = \text{gcd}(T_1, \dots, T_m)$ and let R be the task migration time. A **sufficient condition** for scheduling the m periodic tasks is that $U \leq n \cdot (T-R+1)/T$.

- Schedule as before, but only use $T-R+1$ units of slice
- If migration time is too short \Rightarrow shift task(s) to the right end of schedule


BF - ES

- 7 -

Example (4 processors)

i	Computation block
τ_1	10
τ_2	9
τ_3	9
τ_4	9
τ_5	3

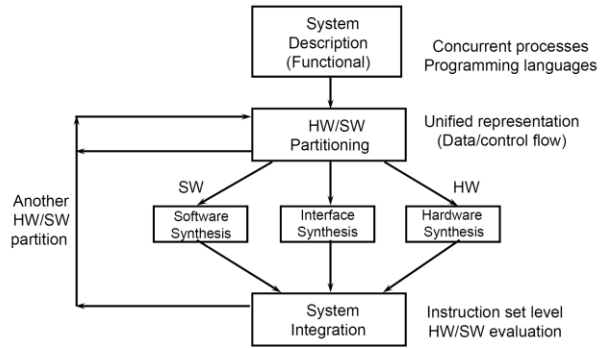
$T=12,$
 $R=3$



BF - ES

- 8 -

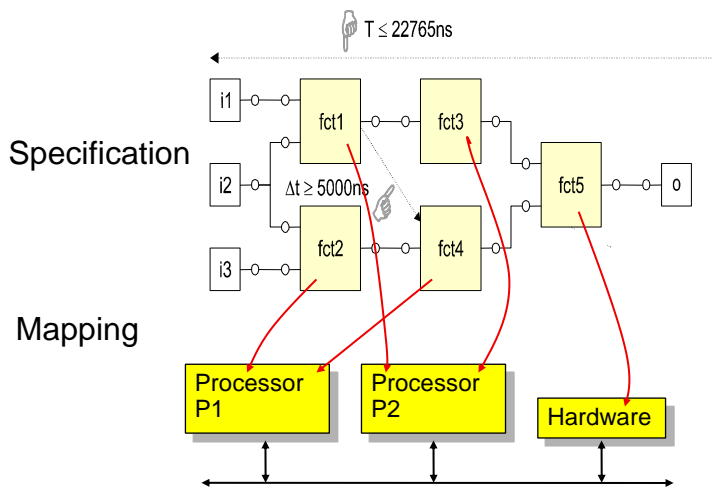
Hardware/Software Codesign



BF - ES

- 9 -

Hardware/software codesign



BF - ES

- 10 -

Objective function

- Cost depends on components selected to implement the application
 - Software Processors: PowerPC, ARM, Pentium,...
 - Hardware: FPGAs, ASIC blocks, ...
 - Communication Infrastructure: buses, networks-on-chip, p2p links, ...
- Multiple metrics, such as cost, power, and performance are weighed against one another
- A function combining multiple metric values into a single value that defines the quality of a partition is called an **Objective Function**, the value returned is called **cost**.

The Partitioning Problem

The **partitioning problem** is to assign n **objects** $O=\{o_1, \dots, o_n\}$ to m **blocks** (also called **partitions**) $P=\{p_1, \dots, p_m\}$ such that

- $p_1 \cup p_2 \dots \cup p_m = O$
- $p_i \cap p_j = \emptyset$ for all $i \neq j$, and
- cost $c(P)$ is minimized.

Partitioning Methods

- **Heuristic methods**
 - Constructive methods
 - Random mapping
 - Hierarchical clustering
 - Iterative methods
 - Kernighan-Lin Algorithm
 - Simulated Annealing
- **Exact methods**
 - Enumeration
 - Integer Linear Programming (ILP)



Constructive Methods

- **Random mapping**
 - Each object randomly assigned to some block
 - Used to find starting partition for iterative methods
- **Hierarchical clustering**
 - Assumes closeness function: determines how desirable it is to group two objects
 - Start with singleton blocks
 - Repeat until termination criterion (e.g., desired number of blocks reached)
 - Compute closeness of blocks (average closeness of object pairs)
 - Find pair of closest blocks
 - Merge blocks
 - Difficulty: find proper closeness function

Example: Hierarchical Clustering

Average closeness;

Termination:
2 blocks



BF - ES

- 15 -

Hw/Sw Partitioning

- Special case: Bi-partitioning $P=\{p_{SW}, p_{HW}\}$
- Software-oriented approach: $P=\{O, \emptyset\}$
 - In software, all functions can be realized
 - Performance might be too low \Rightarrow migrate objects to HW
- Hardware-oriented approach: $P=\{\emptyset, O\}$
 - In hardware, performance is OK
 - Cost might be too high \Rightarrow migrate objects to SW

BF - ES

- 16 -

Greedy Hw/Sw Partitioning

Migration of objects to the other block (HW/SW) until no more improvement

```
repeat
  begin
    P'=P;
    for i=1 to n
      begin
        if (cost(move(P,oi)) < cost(P))
          then P':=move(P,oi);
        end;
      end;
    until (P==P')
```

BF - ES

- 17 -

Iterative Methods: Kernighan-Lin (K-L)

An iterative balanced partitioning (bi-sectioning) heuristic

Given: Two sets A and B, such that $|A|=|B|=n$ and $A \cap B = \emptyset$
cost of edge (a,b) in cut: c_{ab}


While the cost keeps decreasing

- Mark all objects as „unlocked“
- While there are unlocked pairs left
 - Select pair of unlocked objects (a,b) which give the largest decrease or the smallest increase in cut size
 - Mark a and b as „locked“
 - Exchange a and b
 - Record resulting partition and cost
- Continue with the partition with least cost

BF - ES

- 18 -

Example



Step #	object pair	cost reduction	cost
0			5
1	{d, g}	3	2
2	{c, f}	1	1 ← least cost
3	{b, h}	-2	3
4	{a, e}	-2	5


BF - ES

- 19 -

Computing the cost reduction



- External cost of $a \in A$: $E_a = \sum_{v \in B} c_{av}$
- Internal cost of $a \in A$: $I_a = \sum_{v \in A} c_{av}$
- Cost reduction for moving a : $D_a = E_a - I_a$
- Cost reduction for swapping a and b : $g_{ab} = D_a + D_b - 2c_{ab}$
- Update to D -values when a and b are swapped:
 - $D'_x = D_x + 2c_{xa} - 2c_{xb}$ for all $x \in A - \{a\}$
 - $D'_y = D_y + 2c_{yb} - 2c_{ya}$ for all $y \in B - \{b\}$



$$\begin{aligned}
 & c_{xb} - c_{xa} \\
 & c_{xa} - c_{xb} \\
 \hline
 \Delta & 2c_{xa} - 2c_{xb}
 \end{aligned}$$

BF - ES

- 20 -

Weighted Example

C	a	b	c	d	e	f
a	0	1	2	3	2	4
b	1	0	1	4	2	1
c	2	1	0	3	2	1
d	3	4	3	0	4	3
e	2	2	2	4	0	2
f	4	1	1	3	2	0

A={a,b,c}

B={d,e,f}

$$\begin{aligned}
 I_a &= 1+2=3 & E_a &= 3+2+4=9 & D_a &= 9-3=6 \\
 I_b &= 1+1=2 & E_b &= 4+2+1=7 & D_b &= 7-2=5 \\
 I_c &= 2+1=3 & E_c &= 3+2+1=6 & D_c &= 6-3=3 \\
 I_d &= 4+3=7 & E_d &= 3+4+3=10 & D_d &= 10-7=3 \\
 I_e &= 4+2=6 & E_e &= 2+2+2=6 & D_e &= 6-6=0 \\
 I_f &= 3+2=5 & E_f &= 4+1+1=6 & D_f &= 6-5=1
 \end{aligned}$$

BF - ES

- 21 -

Weighted Example

	a	b	c	d	e	f
a	0	1	2	3	2	4
b	1	0	1	4	2	1
c	2	1	0	3	2	1
d	3	4	3	0	4	3
e	2	2	2	4	0	2
f	4	1	1	3	2	0
D	6	5	3	3	0	1

A={a,b,c}

B={d,e,f}

$$\begin{aligned}
 g_{ad} &= D_a + D_d - 2c_{ad} = 6 + 3 - 2 \cdot 3 = 3 \\
 g_{ac} &= 6 + 0 - 2 \cdot 2 = 2 \\
 g_{cf} &= 6 + 1 - 2 \cdot 4 = -1 \\
 g_{bd} &= D_b + D_d - 2c_{bd} = 5 + 3 - 2 \cdot 4 = 0 \\
 g_{be} &= 1 \\
 g_{bf} &= 4 \quad \leftarrow \text{max} \\
 g_{cd} &= 0 \\
 g_{ce} &= -1 \\
 g_{cf} &= 2
 \end{aligned}$$

swap b and f
($g_a = 4$)

BF - ES

- 22 -

Weighted Example

A={a,b,c}

B={d,e,f}

$$A \begin{cases} D'_a = D_a + 2c_{ab} - 2c_{af} \\ \quad \quad = 6 + 2 \cdot 1 - 2 \cdot 4 = 0 \\ D'_c = 3 + 2 \cdot 1 - 2 \cdot 1 = 3 \end{cases}$$

$$B \begin{cases} D'_d = D_d + 2c_{df} - 2c_{db} = 3 + 2 \cdot 3 - 2 \cdot 4 = 1 \\ D'_e = D_e + 2c_{ef} - 2c_{eb} = 0 + 2 \cdot 2 - 2 \cdot 2 = 0 \end{cases}$$

	a	b	c	d	e	f
a	0	1	2	3	2	4
b	1	0	1	4	2	1
c	2	1	0	3	2	1
d	3	4	3	0	4	3
e	2	2	2	4	0	2
f	4	1	1	3	2	0
D	6	5	3	3	0	1

BF - ES

- 23 -

Weighted Example

A={a,b,c}

B={d,e,f}

$$\begin{aligned} g_{ab} &= -5 \\ g_{ac} &= -4 \\ g_{cd} &= -2 \\ g_{ce} &= -1 \leftarrow \text{max} \end{aligned}$$

	a	b	c	d	e	f
a	0	1	2	3	2	4
b	1	0	1	4	2	1
c	2	1	0	3	2	1
d	3	4	3	0	4	3
e	2	2	2	4	0	2
f	4	1	1	3	2	0
D	6	5	3	3	0	1

swap c and e
($\hat{g}_2 = -1$)

$$\begin{aligned} D''_a &= 0 \\ D''_c &= 3 \end{aligned}$$

$$g_{ad} = -3$$

swap a and d
($\hat{g}_3 = -3$)

BF - ES

- 24 -

partial sums: $4, 3, 0$
 \uparrow
 largest partial sum
 \Rightarrow swap b and t .

Iteration 2:
 $\hat{g}_1 = g_{cc} = -1; \hat{g}_2 = g_{ab} = -3, \hat{g}_3 = g_{td} = 4$
 partial sums: $-1, -4, 0$
 largest partial sum: $0 \Rightarrow$ STOP

Kernighan-Lin

Repeat

- Compute D_v für all objects $\leftarrow O(n^2)$
 - Mark all vertices as unlocked
 - For $i=1$ to $n/2$ do $\leftarrow O(n^3)$
 - Compute g_{ab} for all pairs a, b $\leftarrow O(n^2)$
 - Pick unlocked a_i, b_i with largest $g_{ab,i}$
 - Mark a_i, b_i as locked
 - Store gain
 - Update D_v für all objects
 - Find k such that $G_k = \sum_{i=1}^k g_{ab,i}$ is maximal
 - If $G_k > 0$, then move a_1, \dots, a_k from A to B and b_1, \dots, b_k from B to A.
- If repeat loop terminates after r iterations $\Rightarrow O(rn^3)$
 "In practice r is small"



Extensions to K-L

- **Different block sizes**
 - If $|A| < |B|$, add $|B| - |A|$ dummy objects to A.
Dummy objects are connected to each other with infinite weight
 - Apply K-L
 - Remove dummies
- **Objects with size > 1**
 - Replace each object of size s with s objects of size 1
new objects are connected with edges of infinite weight
 - Apply K-L
- **More than 2 blocks**
 - Apply K-L to each pair of blocks

BF - ES

- 27 -

Hypergraphs



BF - ES

- 28 -

Fiduccia-Mattheyses Heuristic (F-M)

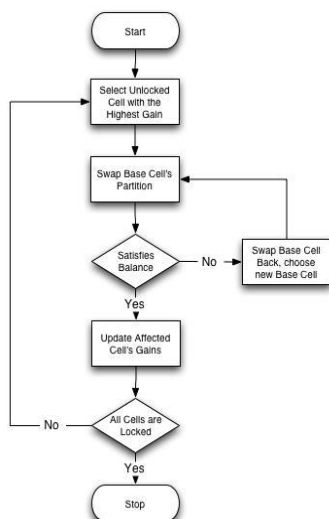
- Objects have size $s(o)$
- Size of block: sum of size of objects
- **Balanced two-way partition:**
Given a fraction r , $0 < r < 1$,
partition a graph into two blocks A and B such that
 $|A| / (|A| + |B|) \approx r$
and cutset is minimized
- **Linear complexity**

Terminology: object=„cell“, hyperedges=„net“

BF - ES

- 29 -

Single pass of the F-M heuristic




- Select the cell with the greatest gains that satisfies balance conditions
- Move the cell and lock it
- Update gains
- Repeat until all cells are locked or will dissatisfy balance conditions

BF - ES

- 30 -

Overall F-M heuristic

- Create an initial partition
- Execute a pass of the F-M heuristic
- Start again using the resulting partition as the initial partition
- Continue until the resulting gain is no longer greater than zero




BF - ES

- 31 -

Calculating Gain

- $g(i) = FS(i) - TE(i)$
- $FS(i)$: The number of nets which contain cell i but no other object in the same partition as i
- $TE(i)$: The number of nets that consist only of i and other cells currently in the same partition as i

object	FS	TE	gain
A	2	0	2
B	1	0	1
C	1	1	0
D	0	1	-1



BF - ES

- 32 -