

Goals of the lecture

Repeated Computation of a Global Function


- Deadlock Detection
- Clock Synchronization
- Distributed Branch and Bound Search
- Distributed Debugging

Desirable Characteristics


- Light Load
 - not more than k messages/time step
- High Concurrency
 - $\log_k N$ time steps
- Symmetry (Equitable Workload)
 - load balancing
 - fairness

Some Possible Approaches


- Centralized



- Ring-based



- Hierarchical




All links are logical connections

Message Flow Table


Static Hierarchy

- Number of nodes (processe) = 7

time step	Messages		
1	1, 3 → 2	5, 7 → 6	
2	1, 3 → 2	5, 7 → 6	2, 6 → 4
3	1, 3 → 2	5, 7 → 6	2, 6 → 4



Overlapping Tress



Message Flow Table

- Revolving Hierarchy

- number of nodes = 7


time step	Messages		idle
1	2 ← 1, 3	6 ← 5, 7	4
2	4 ← 2, 6	5 ← 1, 3	7
3	7 ← 4, 5	1 ← 2, 6	3
4	3 ← 7, 1	2 ← 4, 5	6

- Reorganization of Hierarchy
- Reuse of messages

$$\begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 \\ 5 & 1 & 7 & 2 & 6 & 3 & 4 \end{pmatrix}$$

Requirements for Desired Permutation


- Gather tree constraints
 - interior nodes of $T_i =$ subtree of T_{i+1}
- Fairness constraints
 - No cycle of size less than N .



Interesting but ..

- Does there always exist such a permutation ?
- Is there a systematic method to find it ?
- Is there an efficient implementation for it ?

Method to Generate the Permutation



$next(x) :$

[

$even(x) \rightarrow x' := x/2; (* \text{gather tree constraint} *)$

□

$odd(x) \wedge (x < 2^{n-1}) \rightarrow x' := x + 2^{n-1}; (* \text{fairness constraint} *)$

□

$odd(x) \wedge (x > 2^{n-1}) \rightarrow$

[$x = N \rightarrow x' := (N - 1)/2;$

□

$x \neq N \rightarrow y := x - 2^{n-1} + 2$

$x' := y * 2^{\lceil \log \frac{2^n}{y} - 1 \rceil}$

]

]

Implementation 1

Q: Who should I send message to at time t ?

$$\begin{aligned}msg(x, t) &= next^{-t}(parent(next^t(x))), \text{ if } next^t(x) \text{ is odd} \\ &= nil, \text{ otherwise}\end{aligned}$$


x is in-order label

$next$ is the new position function

$parent$ is the parent function for in-order labeling

parent of $x = x$ with last two bits changed to 10

Implementation 2



$$\begin{aligned}
 msg(x, t) &= new_parent(x + t) - t, && \text{if } (x + t) \text{ is a leaf-node} \\
 &= nil, && \text{otherwise}
 \end{aligned}$$

- Just need to store *new_parent* array

Communication Required

- Communication distance set (CDS)

$$= \{ \mathit{new_parent}(j) - j \mid j \text{ a leaf node} \}$$

- process x will send a message to process y iff $y - x \in CDS$.
 - for $N = 15$

$$CDS = \{1, 5, 8, 10, 13, 14\}.$$

- CDS depends on the *next* function.

Data Gathering and Broadcasting

- a process can send/receive only one message per time step
- require that the same set of messages is used for data gathering and broadcasting.
- Constraints :
 1. fairness constraints
 - equal load
 2. gather tree constraints.
 - $G(t)$ available at $t + \log N$ time step at one node.
 3. broadcast constraints.
 - $G(t)$ available at $t + 2 \log N$ time step at all nodes.

Message Flow Table

time step	Messages			
0	0 → 7	4 → 6	1 → 3	2 → 5
1	7 → 6	3 → 5	0 → 2	1 → 4
2	6 → 5	2 → 4	7 → 1	0 → 3
3	5 → 4	1 → 3	6 → 0	7 → 2
4	4 → 3	0 → 2	5 → 7	6 → 1

- fairness in workload
- four times less messages than static hierarchy

Method to Generate the Permutation

$bcnext(x) ::$

[$b_0 = 1 \rightarrow x' := RS_0(x)$
 (* gather tree *)

□

$(b_0 = 0) \wedge (b_1 = 0) \rightarrow x' := RS_1(x)$
 (* broadcast *)

□

$(b_0 = 0) \wedge (b_1 = 1) \rightarrow x' := LS_1^a ((LS_0^b(x) + 2) \bmod 2^{n-1});$
 (* fairness *)

]

$b_{n-1} \cdots b_0 = x$

$RS_p =$ Right shift with p as m.s.b

$LS_p =$ Left shift with p as l.s.b.

$a =$ number of leading zeros

$b =$ number of leading ones

Algorithm to find Current Minimum in the Network

- Distributed branch and bound
- Distributed simulation
 - process x : $step = 0$
 - *[$dest_msg(x, step) \neq nil \rightarrow send_msg(dest_msg(x, step), mymin)$
 $step := step + 1$
 - $src_msg(x, step) \neq nil \rightarrow recv_msg(src_msg(x, step), hismin)$
recompute $mymin$
 $step := step + 1$

Performance of the Algorithm

- at most k messages handled by a node/time step
- the global function $G(t)$ is available at $t + \lceil \log N \rceil$ time steps.
- a throughput of one global function per times step.
- number of messages required \sim half of that for static hierarchy.
- equal workload distribution

Extensions

- General N
 - use virtual nodes
- General k
 - methods to generate permutations for binary trees generalize to k -ary trees.
- asynchronous messages
 - can be used instead of synchronous messages. Nodes synchronized due to “receives”.

Conclusions

- Useful for algorithms that
 - use hierarchical control
 - run for long time
- main advantages
 - equal workload distribution.
 - reduction in number of messages due to their reuse
- main disadvantages
 - requires that the communication network has more edges than static hierarchy.