

Parallel Nondeterminator

He Yuxiong, Wang Junqing





Objective

Develop Parallel Nondeterminator to check for determinacy races during the parallel execution of a Cilk program.

Two main tasks:

- Develop efficient algorithm for concurrency test
- Evaluate its performance through experiments




Outline

Algorithm: modified NR Labeling (MNR)

- » make it applicable to Cilk
- » correct and systematic proof
- » improve the performance

Simulation:

- » Performance analysis
- » Simulator design and implementation
- » Experimental results



Cilk POEG



Cilk POEG is the partial order execution graph of Cilk Program.

Vertex - sync, spawn operation

Edge - thread

Cilk POEG Unit




basic constructive block



Cilk POEG

Recursive Composition



replace any basic function
F by the series of cilk
POEG units



Cilk POEG

Recursive Composition


replace any basic function
F with the series of cilk
POEG units



MNR Labeling

$T0 \downarrow \langle 1, 50 \rangle$


Basic Label:
 $NR(T0) = \langle 1, 50 \rangle$



MNR Labeling

Syn counter of T:

the number of the sync threads of T in the critical path from initial thread to T




Basic Label:


$NR(T_0) = \langle 1, 50 \rangle$

Sync counter:

$Sc(T_0) = 1$




MNR Labeling



Basic Label:
 $NR(T_0) = \langle 1, 50 \rangle$

Sync counter:
 $Sc(T_0) = 1$

Recursive depth+1:
 $Rd(T_0) = 1$



MNR Labeling



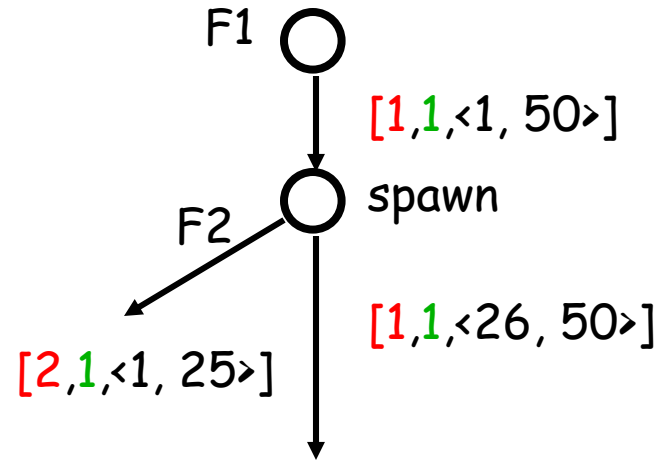
Spawn Operation

Child:

- (1) Divide and inherit the first half of the range in parent thread
- (2) Update recursive depth

Parent:

- (1) Inherit the second half of the range in the parent thread before spawn



MNR Labeling

Sync Operation

Child:

(1) Return sync count



Parent:

(1) Restore the initial range

(2) Update sync count

Concurrency Test



$$T_i || T_j \Leftrightarrow NR(T_i) \cap NR(T_j) = \varepsilon$$



MNR Labeling

Concurrency Test

$$T_i || T_j \stackrel{?}{\Leftrightarrow} NR(T_i) \cap NR(T_j) = \varepsilon$$





MNR Labeling

History OH(T):

Include one-way roots of the thread T

One-way root of T is the most recent sync thread happened before T in its ancestor functions and its own function



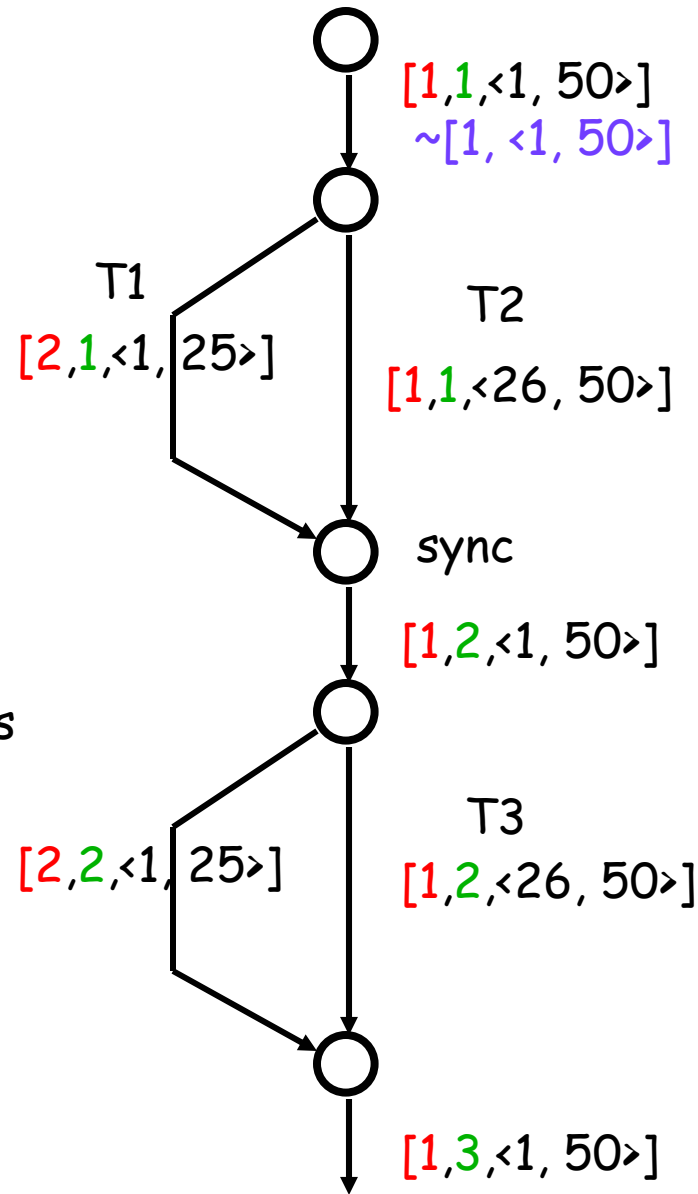


MNR Labeling

History $OH(T)$:

Include one-way roots of the thread T

One-way root of T is the most recent sync thread happened before T in its ancestor functions and its own function



MNR Labeling



Spawn Operation

Child:

- (1) Divide and inherit the first half of the range in parent thread
- (2) Update recursive depth
- (3) Copy OH(parent) to child

Parent:

- (1) Inherit the second half of the range in the parent thread before spawn
- (2) OH does not change



MNR Labeling

Sync Operation

Child:



(1) Return sync count

Parent:

(1) Restore the initial labeling

(2) Update sync count

(3) Update OH



MNR Labeling


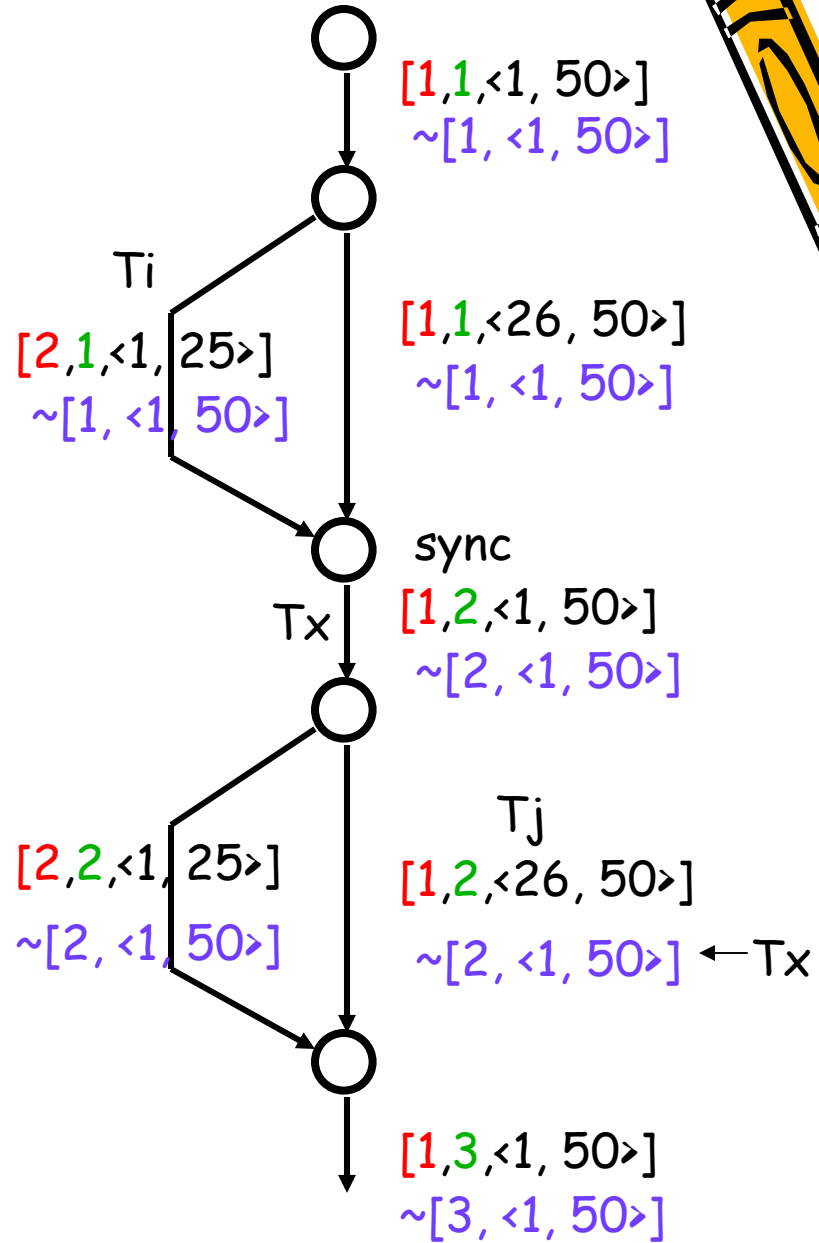
Concurrency Test

if $sc(T_i) = sc(T_j)$
 $T_i || T_j \Leftrightarrow NR(T_i) \cap NR(T_j) = \epsilon$

if $sc(T_i) < sc(T_j)$
 $T_i || T_j \Leftrightarrow NR(T_i) \cap NR(T_x) = \epsilon$

T_x : nearest one-way root of T_j from T_i .

the thread whose sync counter is the smallest one in $OH(T_j)$ such that it is greater than sync counter of T_i



Main Theory for MNR


Theorem 1:

Given two threads T_i and T_j in Cilk POEG

$$T_i || T_j \iff \begin{cases} NR(T_i) \cap NR(T_j) = \varepsilon & \text{if } sc(T_i) = sc(T_j) \\ NR(T_i) \cap NR(T_x) = \varepsilon & \text{if } sc(T_i) < sc(T_j) \\ NR(T_j) \cap NR(T_y) = \varepsilon & \text{if } sc(T_j) < sc(T_i) \end{cases}$$

Proof: (\Rightarrow)

lemma: $T_i || T_j \Rightarrow NR(T_i) \cap NR(T_j) = \varepsilon$



Proof of Main Theorem

Proof: (\Rightarrow Continued)

lemma: $T_i || T_j \Rightarrow NR(T_i) \cap NR(T_j) = \varepsilon$


proof by induction:

base step:


- (1) true in Cilk POEG unit
- (2) true in series of Cilk POEG units

induction step:

suppose lemma is true in Cilk POEG G , prove it is still true in G' which is the Cilk POEG after we do recursive composition on any basic function of G



Main Theory for MNR



Theorem 1:

Given two threads T_i and T_j in Cilk POEG

$$T_i || T_j \iff \begin{cases} NR(T_i) \cap NR(T_j) = \varepsilon & \text{if } sc(T_i) = sc(T_j) \\ NR(T_i) \cap NR(T_x) = \varepsilon & \text{if } sc(T_i) < sc(T_j) \\ NR(T_j) \cap NR(T_y) = \varepsilon & \text{if } sc(T_j) < sc(T_i) \end{cases}$$

Proof: (\Rightarrow)

lemma: $T_i || T_j \Rightarrow NR(T_i) \cap NR(T_j) = \varepsilon$



Proof of Main Theorem

Proof: (\Leftarrow)


Case 1: $sc(T_i) = sc(T_j)$

prove by contradiction

(a) if $T_i \rightarrow T_j$ or $T_j \rightarrow T_i$, match T_i and T_j with the threads in one-way cilk POEG



(b) Prove

$(T_i \rightarrow T_j) \vee (T_j \rightarrow T_i) \Rightarrow NR(T_i) \cap NR(T_j) \neq \varepsilon$
in one-way Cilk POEG



One-way Cilk POEG

Recursive composition
of Cilk POEG units



Proof of Main Theorem

Proof: (\Leftarrow)

Case 1: $sc(T_i) = sc(T_j)$

prove by contradiction

(a) if $T_i \rightarrow T_j$ or $T_j \rightarrow T_i$, match T_i and T_j with the threads in one-way cilk POEG

(b) Prove

$(T_i \rightarrow T_j) \vee (T_j \rightarrow T_i) \Rightarrow NR(T_i) \cap NR(T_j) \neq \varepsilon$
in one-way Cilk POEG



Proof of Main Theorem

Proof: (\Leftarrow Continued)

Case 2: $sc(T_i) < sc(T_j)$

Lemma:

$$T_i \rightarrow T_j \Rightarrow NR(T_i) \cap NR(T_x) \neq \varepsilon$$


proof:

based on $G_s = \text{sub-graph}(LCA(T_i, T_j))$, T_i and T_j are

(a) in the same expanded Cilk POEG unit in G_s

(b) in different expanded Cilk POEG units in G_s

Case 3: similar way with case 2




Variable Monitoring Algorithm

Tasks:

- Report data race
- Update the access history entries

Issues:

- How many history entries per variable?
- How to update the history entries?
- Does it guarantee correct and complete race report?



Variable Monitoring Algorithm

How many history entries per variable?

- two read entries, one write entry

How to update the history entries?

- introduce left-of operator: $T_i < T_j \Leftrightarrow T_i.a < T_j.a$
- prove left-of operator establishes a canonical ordering of relatives

Theorem 2: Given three threads T_i , T_j and T_k , $T_i || T_j || T_k$
 $T_i < T_j < T_k \Rightarrow LCA(T_i, T_k) = LCA(T_i, T_j, T_k)$



Variable Monitoring Algorithm



Does it guarantee correct and complete race report?

Theorem 3-6:

At least one data race will be reported if there are any conflicting, logically concurrent access in a shared variable.

Theorem 7:

Data race on a variable will only be reported if there exists conflicting, logically concurrent access.



Performance

V: # of shared variables

T: maximum parallelism

N: depth of recursion


P: # of processors

- Time Complexity

- label updates

- for each spawn: $O(N)$

- for each sync: $O(1)$




Performance (cont.)

- access history check
 - for each shared variable access: $O(\lg N)$

- Space Complexity

- variable access history
 - for each variable: $O(1)$
- labels
 - for each thread: $O(N)$
 - in total: $O(V + \min(NP, NT))$



Simulation

- Generator

- generate files:

- two files under the same specification, one is no race-checking cilk program, the other is for MNR Nondeterminator.

- requirements:

- random enough
 - general case



Simulation (cont.)

-- specifications:

- # of shared variables
- depth of recursion
- # of functions


For each function:

- # of syncs
- # of spawns per sync
- range of calculation delay


- Simulator Runtime Support

-- labeling

-- readcheck / writecheck operation



Testing Introduction



- Correctness
 - Sample results

Write-Read data race
@ var_9 in line:504
with line:891

Write-Write data race
@ var_29 in line:529
with line:869

Read-Write data race
@ var_29 in line:844
with line:869

....

MNR Nondterminator


Data race @ 0x3ce5c
`var_9' (eg1_s.cilk:408) with
`var_9' (eg1_s.cilk:717)

Data race @ 0x33174
`var_29' (eg1_s.cilk:428) with
`var_29' (eg1_s.cilk:680)

Data race @ 0x33174
`var_29' (eg1_s.cilk:428) with
`var_29' (eg1_s.cilk:700)

....

Serial Nondeterminator in Cilk 5.2.1




Testing Introduction (cont.)



- Result analysis

```
cilk int main(...) {  
    ...  
    spawn function_13(...);  
    ...;  
    spawn function_23(...);  
    ...;  
    sync;  
    ...;  
}
```



```
cilk void function_13(...){  
    ...  
    writecheck(&hist[8],9, flabel, 504);  
    var_9=234;  
    ...  
}  
  
cilk void function_23(...){  
    ...  
    spawn function_13(...);  
    int tmp;  
    readcheck(&hist[8], 9, flabel, 891);  
    tmp=var_9;  
    ...  
}
```



Testing Introduction (cont.)



- Result analysis

```
cilk int main(...) {  
    ...  
    spawn function_13(...);  
    ...;  
    spawn function_23(...);  
    ...;  
    sync;  
    ...;  
}
```

```
cilk void function_13(...){  
    ...  
    writecheck(&hist[8],9, flabel, 504);  
    var_9=234;  
    ...  
}
```

```
cilk void function_23(...){  
    ...  
    spawn function_13(...);  
    int tmp;  
    readcheck(&hist[8], 9, flabel, 891);  
    tmp=var_9;  
    ...  
}
```



Testing Introduction (cont.)



- Result analysis

```
cilk int main(...) {
```

```
...
```

```
spawn function_13(...);
```

```
...;
```

```
spawn function_23(...);
```

```
...;
```

```
sync;
```

```
...;
```

```
}
```

```
cilk void function_13(...) {
```

```
...
```

```
writecheck(&hist[8], 9, flabel, 504);
```

```
var_9=234;
```

```
...
```

```
}
```

```
cilk void function_23(...) {
```

```
...
```

```
spawn function_13(...);
```


```
int tmp;
```

```
readcheck(&hist[8], 9, flabel, 891);
```

```
tmp=var_9;
```

```
...
```

```
}
```



Testing Introduction (cont.)



- Result analysis

```
cilk int main(...) {  
    ...  
    spawn function_13(...);  
    ...;  
    spawn function_23(...);  
    ...;  
    sync;  
    ...;  
}
```

```
cilk void function_13(...) {  
    ...  
    writecheck(&hist[8], 9, flabel, 504);  
    var_9=234;  
    ...  
}  
  
cilk void function_23(...) {  
    ...  
    spawn function_13(...);  
    int tmp;  
    readcheck(&hist[8], 9, flabel, 891);  
    tmp=var_9;  
    ...  
}
```

WR Race

Write-Read data race
@ var_9 in line:504
with line:891



Testing Introduction (cont.)

- Effectiveness


- We try to illustrate the performance of our MNR Nondeterminator from 4 main experiments

- recursion
 - calculation delay
 - # of spawn per sync
 - critical path & parallelism

- all the tests are conducted on ygg


Performance Ratio:

$$\frac{\text{execution time in MNR Nondeterminator}}{\text{corresponding cilk program without race-checking}}$$



Testing (1)

-- change of recursion




- Analysis

- Range from 1 to 1.8, mainly restricted in the interval of 1 to 1.5
- Relative performance drops with the increase of recursion level when the # of proc is large




Testing (2)

-- change of calculation delay




- Analysis

- When $d=1000$, the performance ratio is worst
- With the increment of calculation delay, the performance improves, the ratio is close to 1



Testing (3)

-- change of spawn per sync




- Analysis

- A little fluctuation around 16 proc
- With the increment of spawn number, the lock overhead increases




Testing (4)

-- test on critical path & parallelism




- Analysis

- Critical path ranges from 2 to 3, meanwhile parallelism ranges from 0.5 to 1



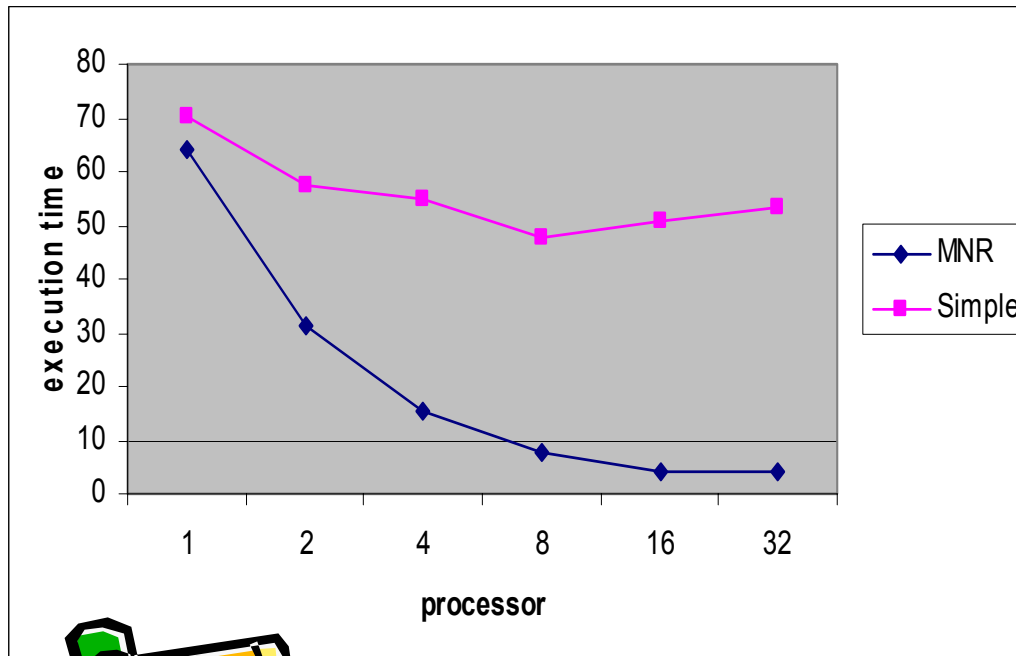
Testing Summary

- Analysis
 - From the above 4 tests, we could find that the execution time of our MNR Nondeterminator ranges generally from 1 to 1.8, compared with the corresponding no race-checking cilk program based on our testcases.
 - Unknown phenomenon that the performance decreases when the # of proc reaches 16





Testing Summary (cont.)

- Bad case




- Analysis

- Testcase specification
- Performance of 32 proc is only 2/3 of that of 1 proc
- Lock/unlock overhead.
- Most accesses will be held on by the locks



Testing Summary (cont.)

- Further testing
 - more testing data
 - benchmark case
 - comparison with serial Nondeterminator



Conclusions

Achievements

--MNR Labeling

time complexity:



$O(1)$ for each sync, $O(N)$ for each spawn

$O(\lg N)$ for worst case of access check

space complexity: $O(V + \min(NT, NP))$

-- Performance

with the change of the parameters, such as recursion, spawn number, calculation delay and processor number, the ratio to the no race-checking cilk program is generally within the range from 1 to 1.8 based on our testcases



Questions ?

