



UCIFF: Unified Clustering, Instruction scheduling and Fast Frequency selection for Heterogeneous Clustered VLIW

Vasileios Porpodas and Marcelo Cintra

University of Edinburgh



Scheduling, Scalability and Energy

- Energy becomes a major design constraint
 - Dynamic Instruction scheduling in hardware consumes a large part of the energy budget
 - Statically scheduled processors are an energy-efficient alternative to dynamically scheduled processors
 - VLIW processors are high-performance statically scheduled processors

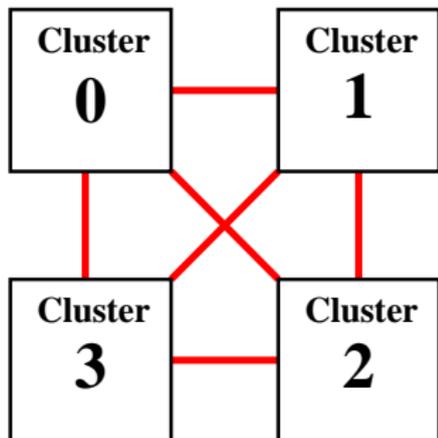


Scheduling, Scalability and Energy

- Energy becomes a major design constraint
 - Dynamic Instruction scheduling in hardware consumes a large part of the energy budget
 - Statically scheduled processors are an energy-efficient alternative to dynamically scheduled processors
 - VLIW processors are high-performance statically scheduled processors
- Resource scalability is necessary for both energy and performance
 - Clustered VLIW processors operate at an attractive energy-performance point
 - No global buses, only point-to-point communication
 - Instruction scheduling done by the compiler

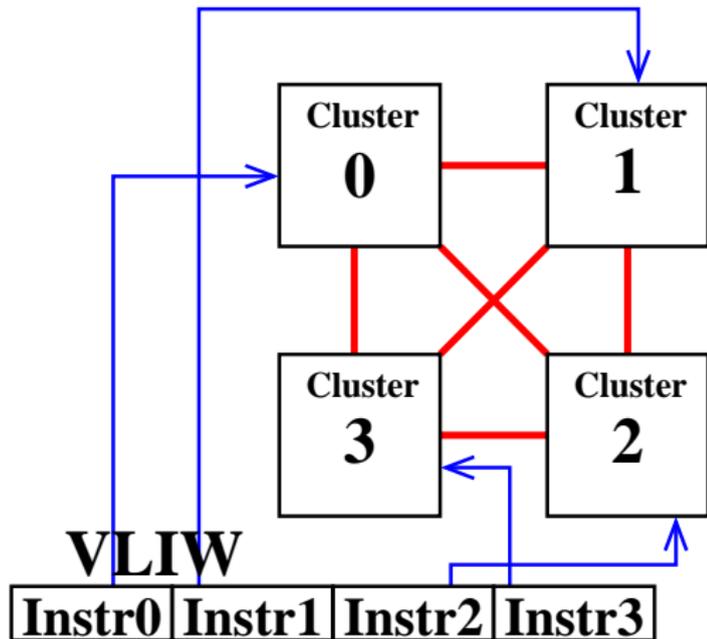
Clustered VLIW

- Statically scheduled
- Scalable
- Energy efficient
- Inter-Cluster delay



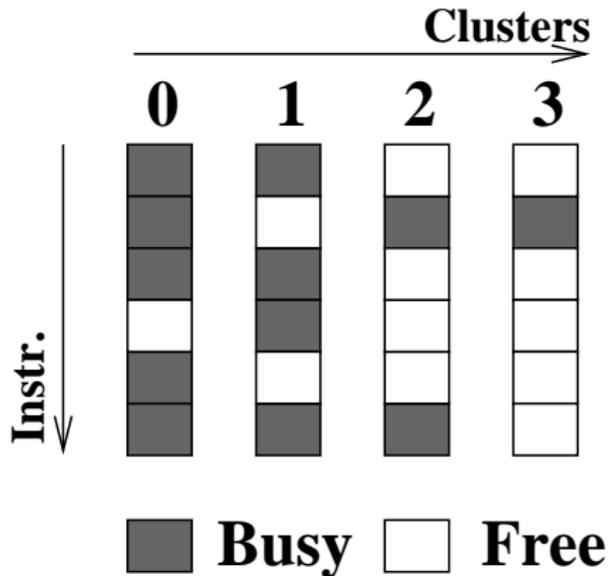
Clustered VLIW

- Statically scheduled
- Scalable
- Energy efficient
- Inter-Cluster delay
- Relies on compiler
- Explicit ILP



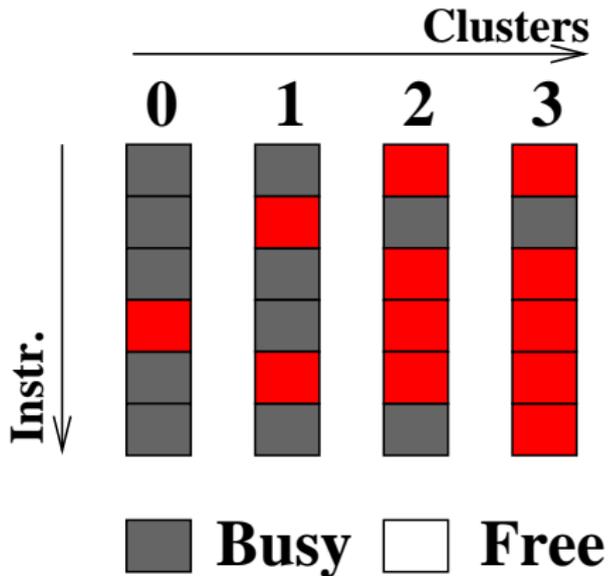
Cluster Utilization

- Few of the clusters are fully utilized



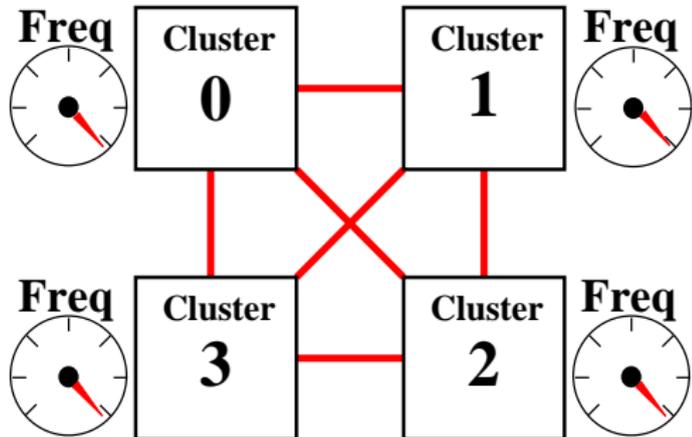
Cluster Utilization

- Few of the clusters are fully utilized
- Slack in schedule
- Opportunity to save energy without performance impact



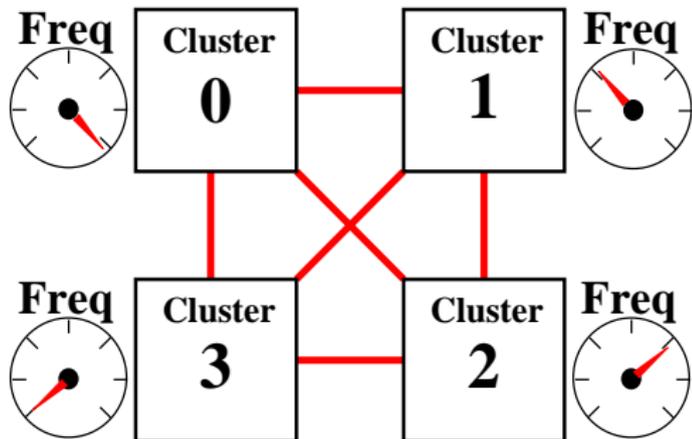
Heterogeneous Clustered VLIW

- Each cluster operates at its own frequency



Heterogeneous Clustered VLIW

- Each cluster operates at its own frequency
- Exploit cluster under-utilization
- Save energy by slowing down under-utilized clusters



The Problem of Energy Efficiency on Clustered VLIW

- Cluster utilization varies
 - Some clusters are fully utilized
 - Others are running idle
- Per-cluster DVFS required for energy efficiency
 - Hardware DVFS not applicable (breaks semantics)
 - Effective compile-time DVFS required

Outline

Introduction

Problem Definition and Existing Solutions

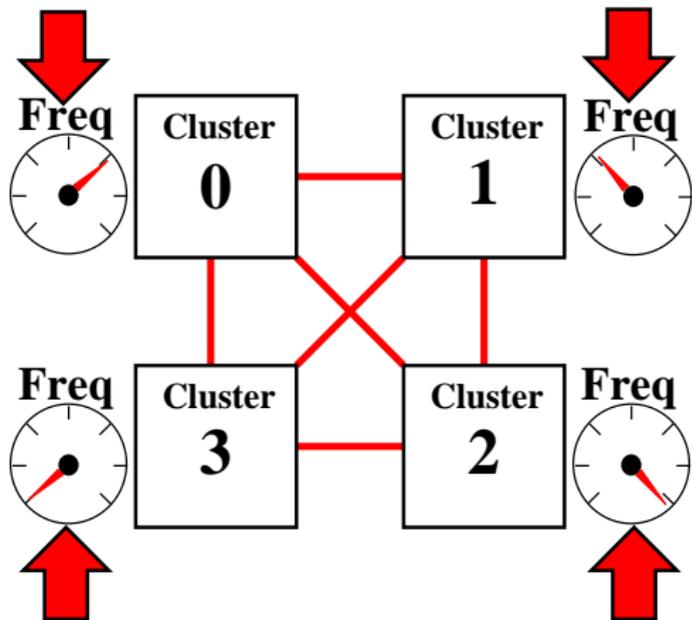
UCIFF

Experimental Setup and Results

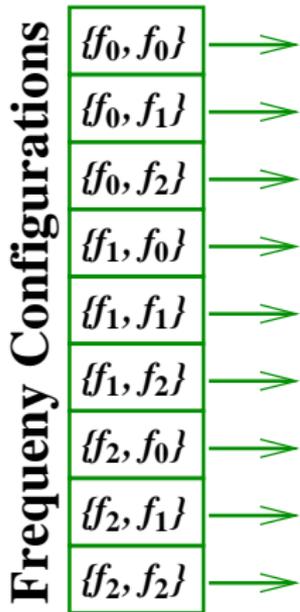
Conclusion

Problem Definition

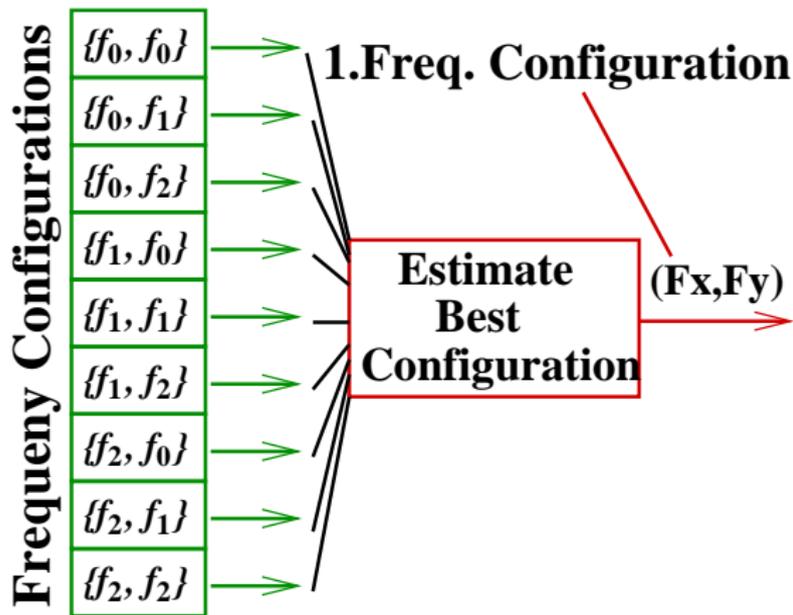
- How to determine “best” frequency for each cluster
- “Best” freq. is the one that leads to best Delay/Energy/ED/ ED^2
- Determine frequencies during Instruction Scheduling



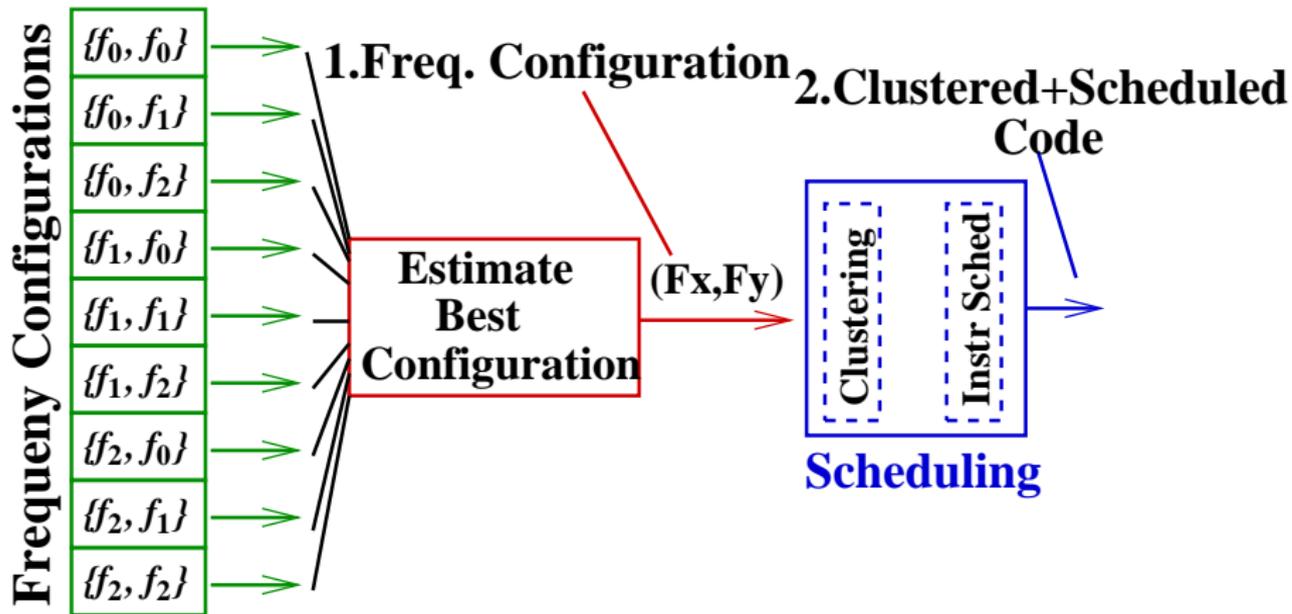
Existing Solution [CGO'07] (Decoupled)



Existing Solution [CGO'07] (Decoupled)

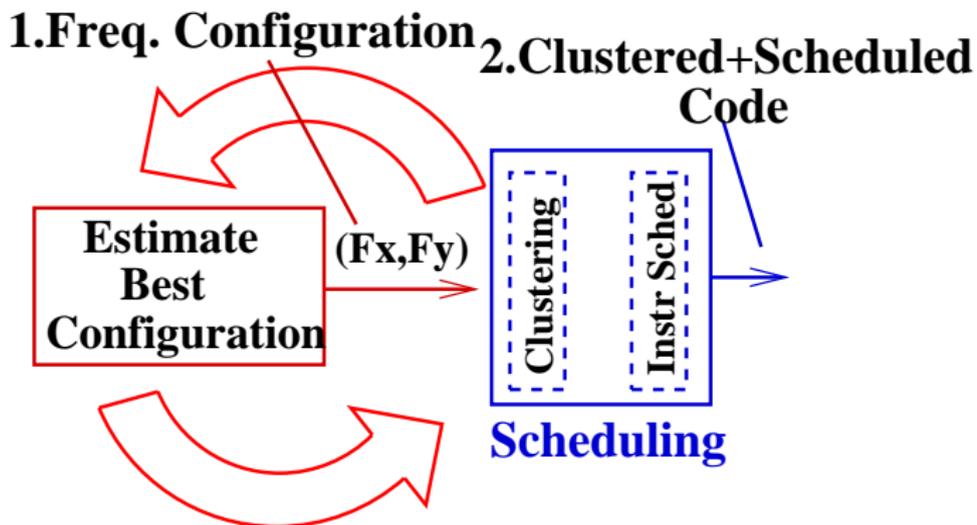


Existing Solution [CGO'07] (Decoupled)



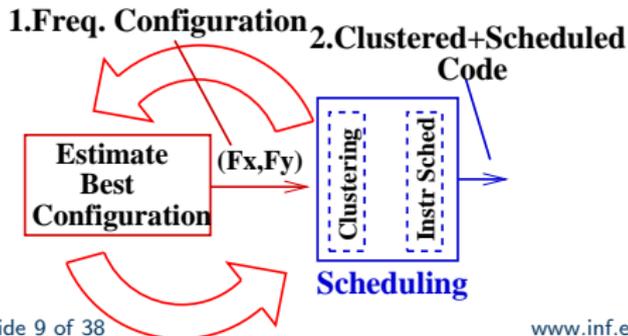
Existing Solution [CGO'07] (Decoupled)

- Phase ordering problem



Existing Solution [CGO'07] (Decoupled)

- Phase ordering problem
- Estimation of Best freq. requires knowledge of Performance and Energy
- Performance and Energy measurement requires schedule
- Scheduling requires that the frequencies are set



Outline

Introduction

Problem Definition and Existing Solutions

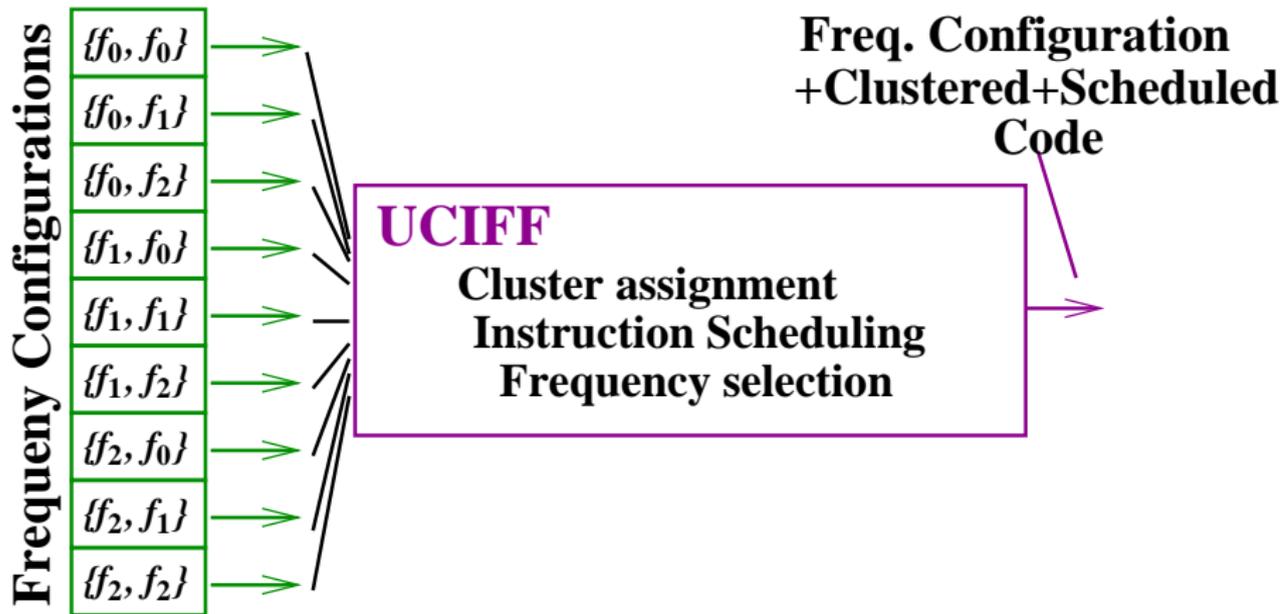
UCIFF

Experimental Setup and Results

Conclusion

UCIFF (unified)

- No cyclic dependency

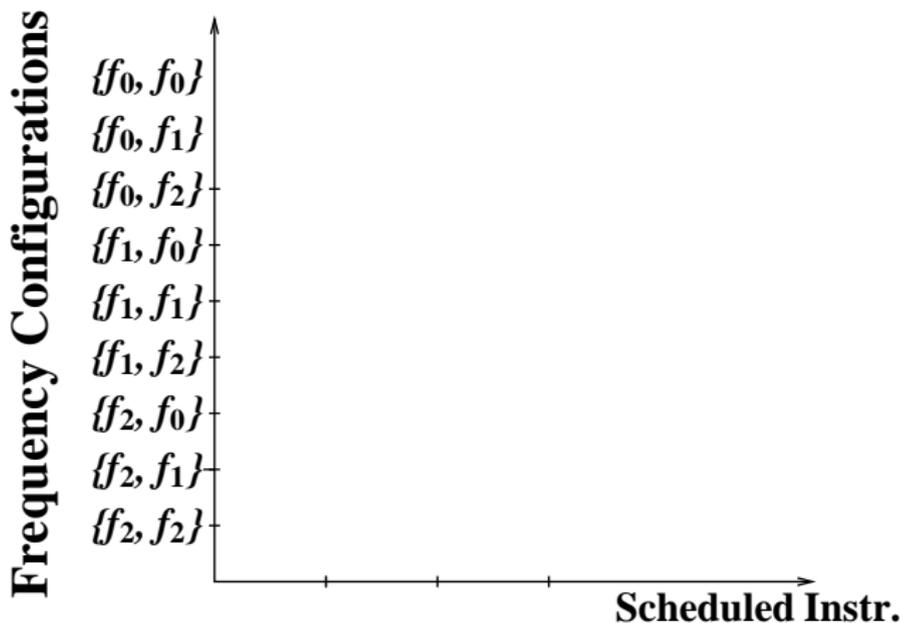




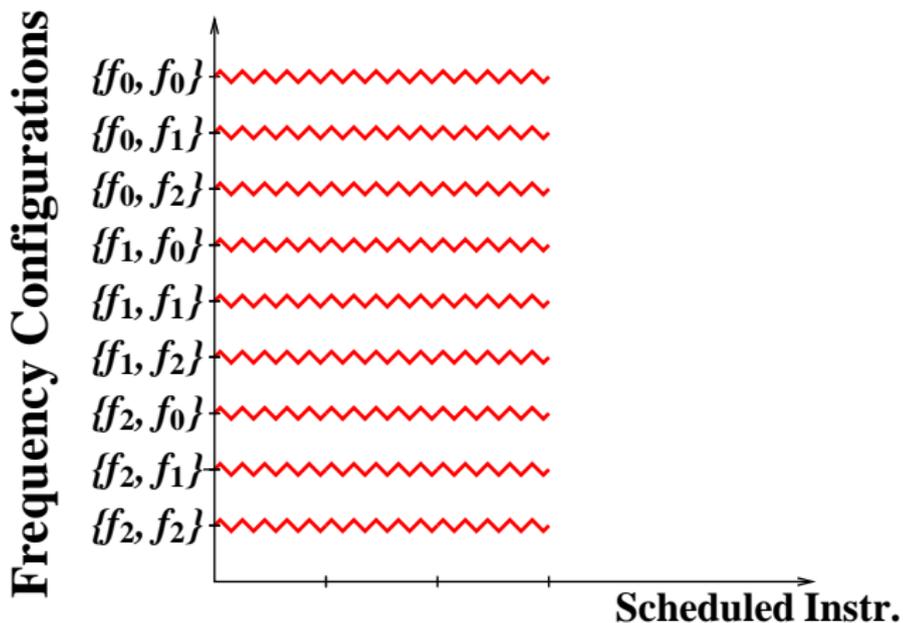
Full-Search Solution

- Brute-force: Try (schedule) all configurations
- After trying all find the best
- Obviously the slowest method

Full-Search Solution

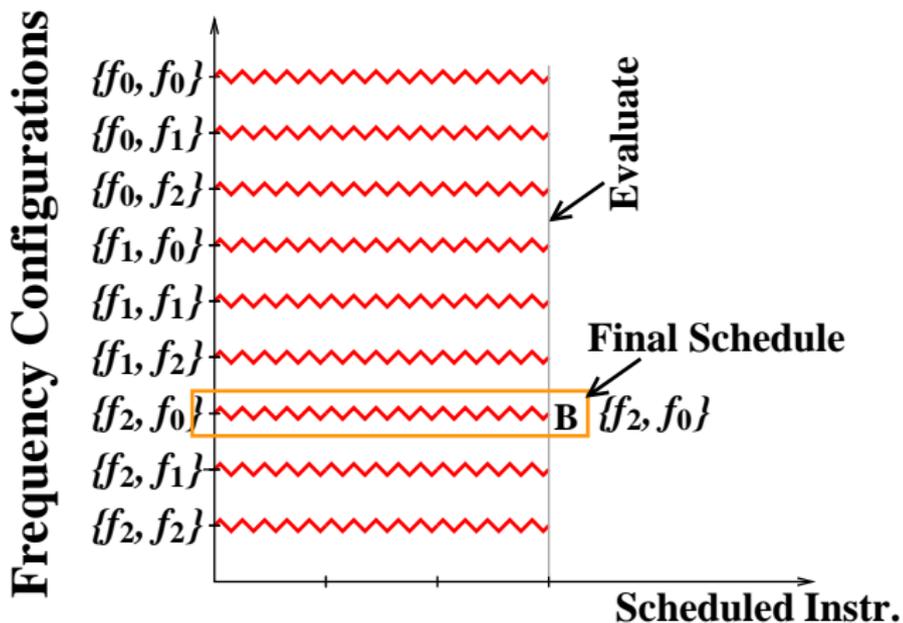


Full-Search Solution



 **Schedule**

Full-Search Solution



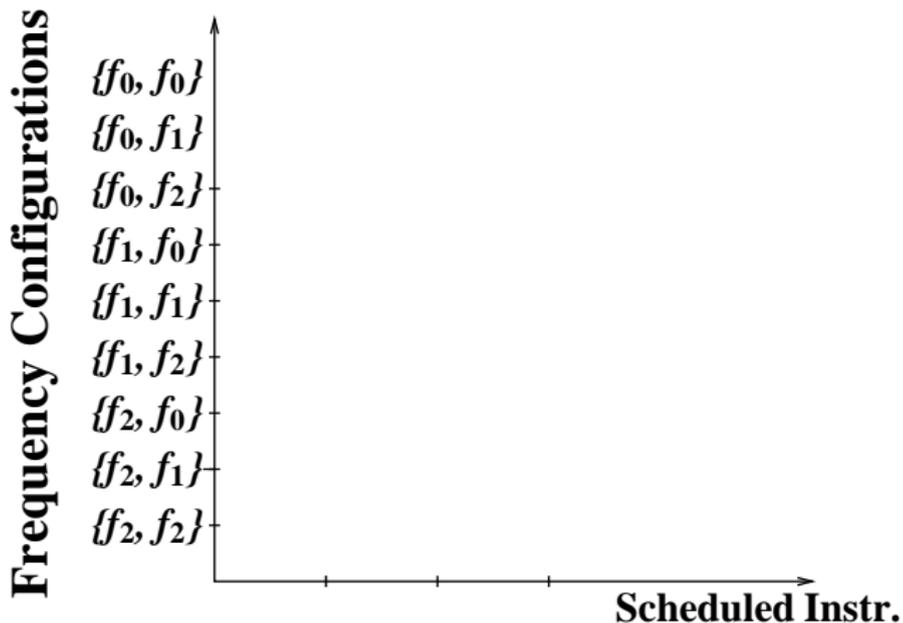
~~~~~ Schedule

**B** Best Freq Configuration

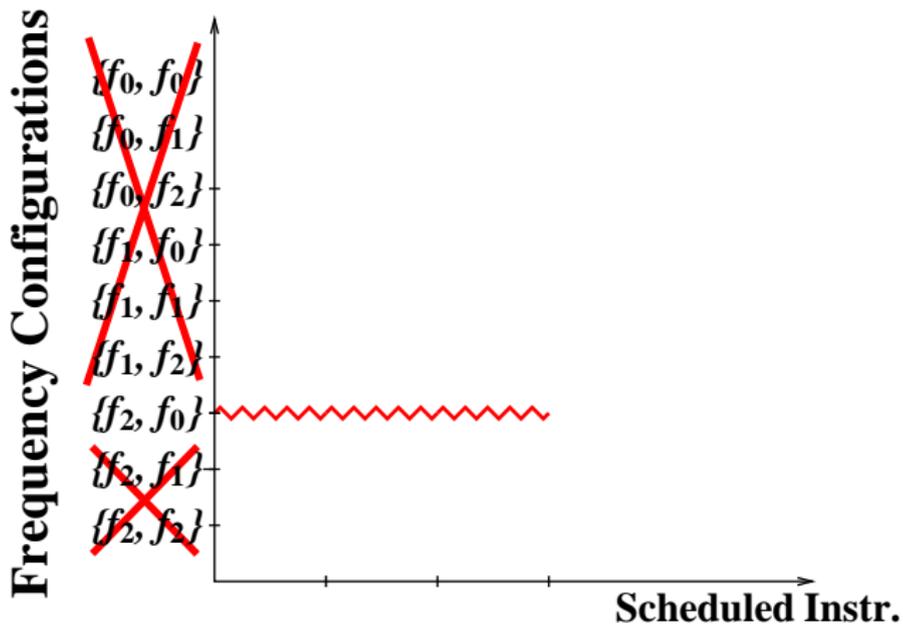
## Theoretical Oracle Solution

- A-priori knowledge of the best frequency configuration
- Schedules only the configuration which will generate the best schedule
- Fastest but Non-implementable

# Theoretical Oracle Solution

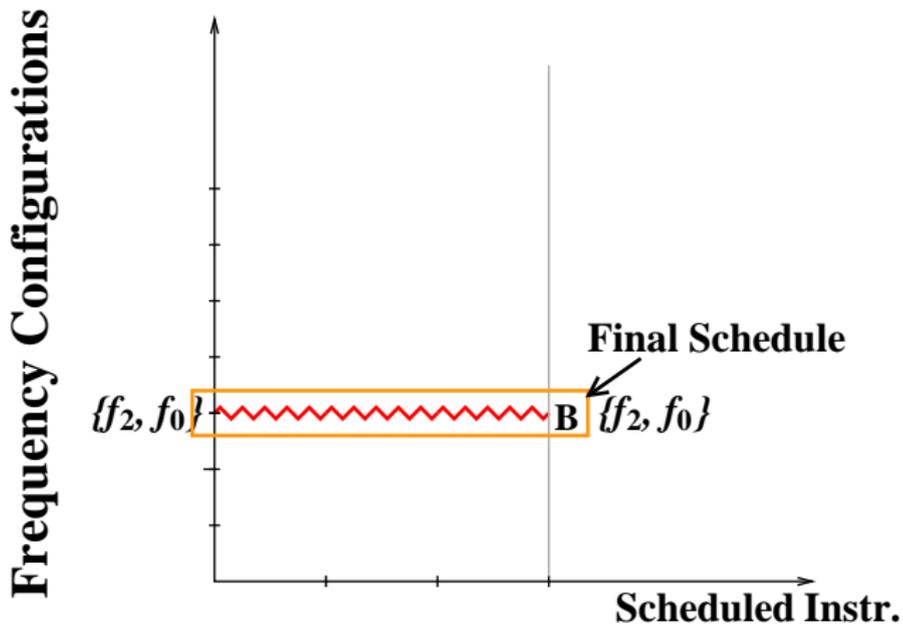


# Theoretical Oracle Solution



~~~~~ Schedule

Theoretical Oracle Solution

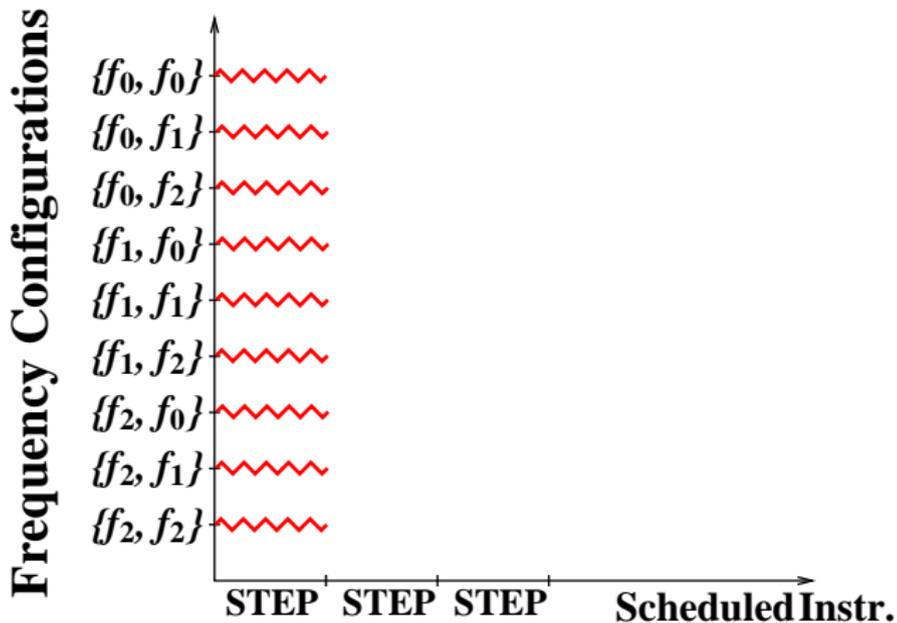


~~~~~ **Schedule**

# UCIFF

- Finds a good solution without resorting to full-search
  - Hill-Climbing over frequency configurations
  - Partial schedules “STEP” scheduling cycles each
  - At the end of each “STEP” it finds the best configuration and schedules it along with its neighbors for the next “STEP”

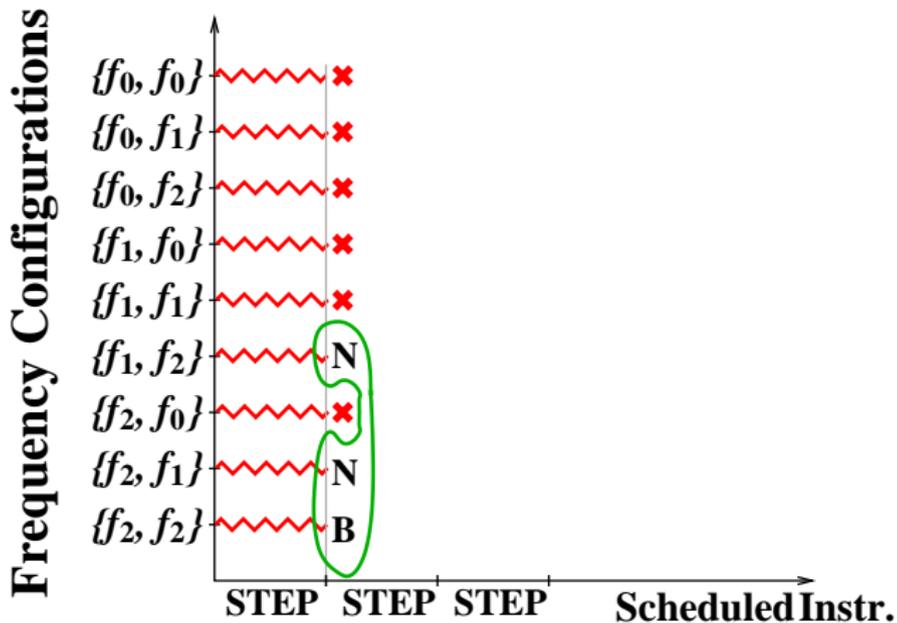
## UCIFF



 **Active Partial Schedule**

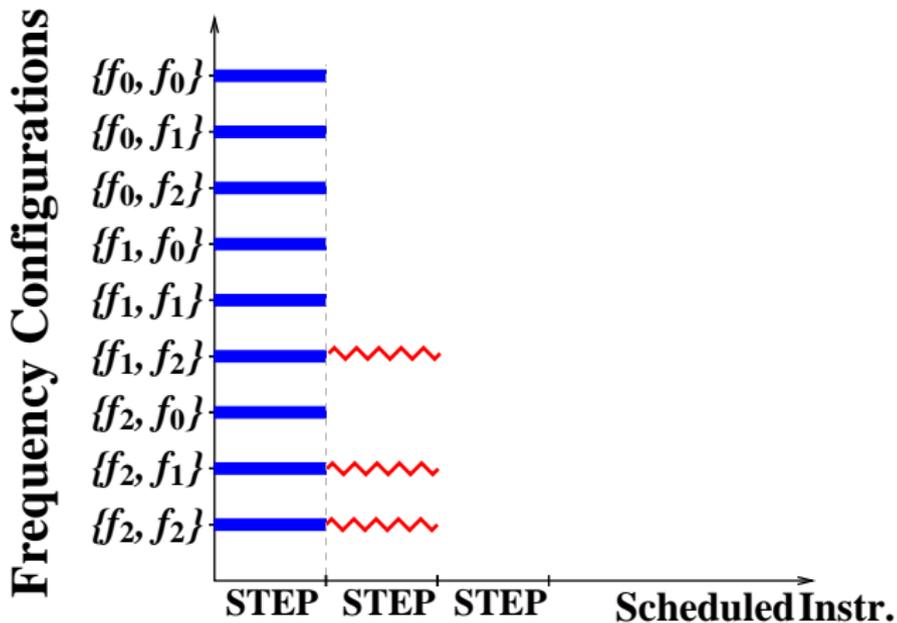


# UCIFF



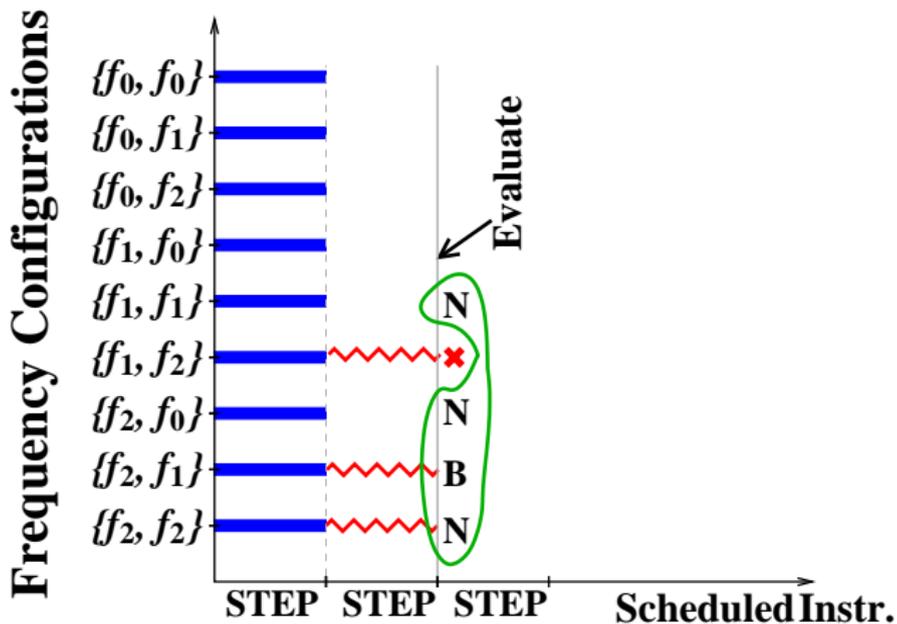
- |  |                          |          |                         |
|--|--------------------------|----------|-------------------------|
|  | Active Partial Schedule  | <b>B</b> | Best Freq Configuration |
|  | Kill Freq. Configuration | <b>N</b> | Neighbor to B           |
|  | Active Set               |          |                         |

# UCIFF



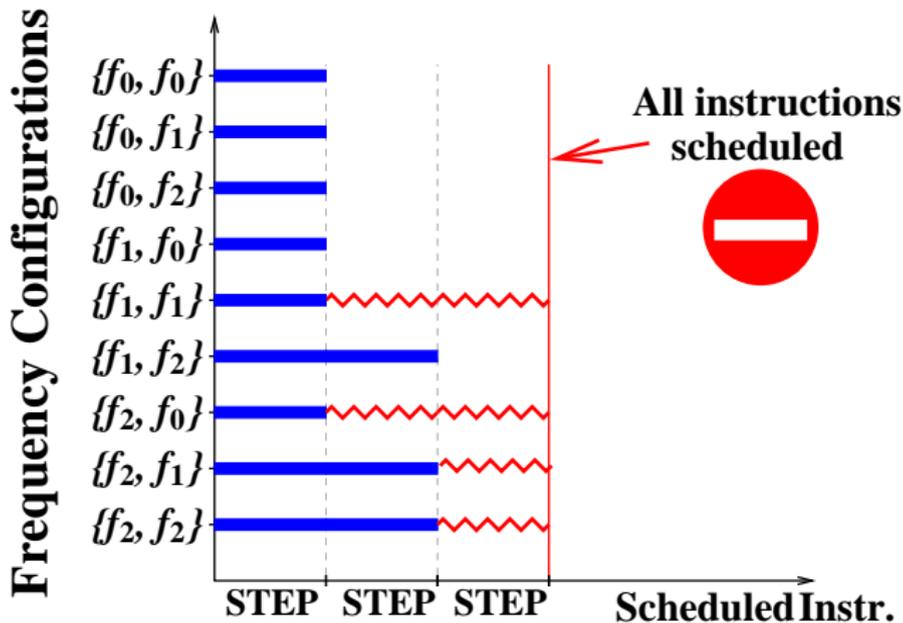
- |  |                           |          |                         |
|--|---------------------------|----------|-------------------------|
|  | Active Partial Schedule   | <b>B</b> | Best Freq Configuration |
|  | Inactive Partial Schedule | <b>N</b> | Neighbor to B           |
|  | Kill Freq. Configuration  |          | Active Set              |

# UCIFF



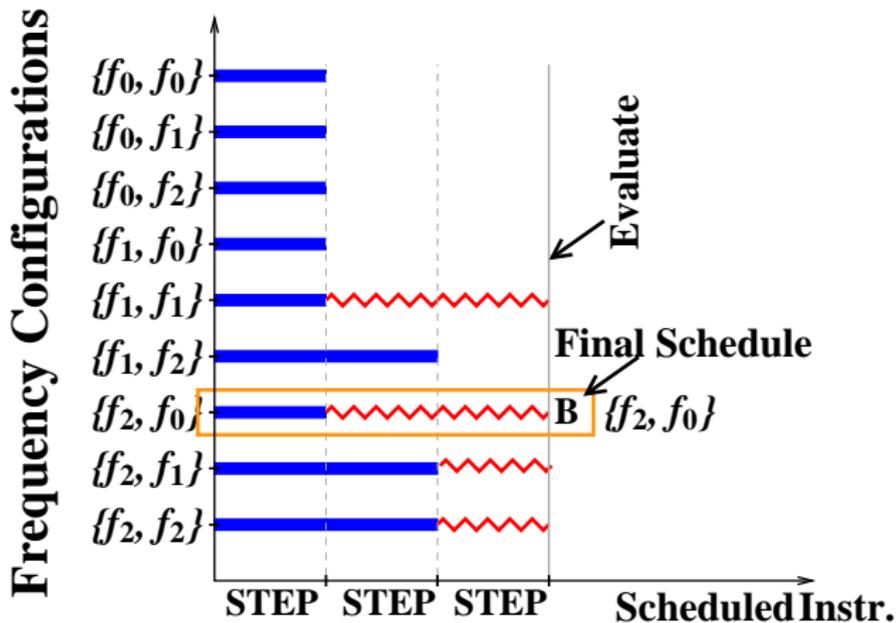
- |                                                                                   |                           |                                                                                   |                         |
|-----------------------------------------------------------------------------------|---------------------------|-----------------------------------------------------------------------------------|-------------------------|
|  | Active Partial Schedule   | <b>B</b>                                                                          | Best Freq Configuration |
|  | Inactive Partial Schedule | <b>N</b>                                                                          | Neighbor to B           |
|  | Kill Freq. Configuration  |  | Active Set              |

# UCIFF



- |  |                           |          |                         |
|--|---------------------------|----------|-------------------------|
|  | Active Partial Schedule   | <b>B</b> | Best Freq Configuration |
|  | Inactive Partial Schedule | <b>N</b> | Neighbor to B           |
|  | Kill Freq. Configuration  |          | Active Set              |

# UCIFF



- |  |                           |          |                         |
|--|---------------------------|----------|-------------------------|
|  | Active Partial Schedule   | <b>B</b> | Best Freq Configuration |
|  | Inactive Partial Schedule | <b>N</b> | Neighbor to B           |
|  | Kill Freq. Configuration  |          | Active Set              |



# UCIFF

- Benefits of UCIFF hill-climbing:
  - More accurate frequency selection than CGO'07 (no phase-ordering problem)
  - Not based on estimation of performance or energy consumption, it measures the actual schedule.
  - Less scheduling time than full-search
  - Accuracy close to full-search

# Outline

Introduction

Problem Definition and Existing Solutions

UCIFF

**Experimental Setup and Results**

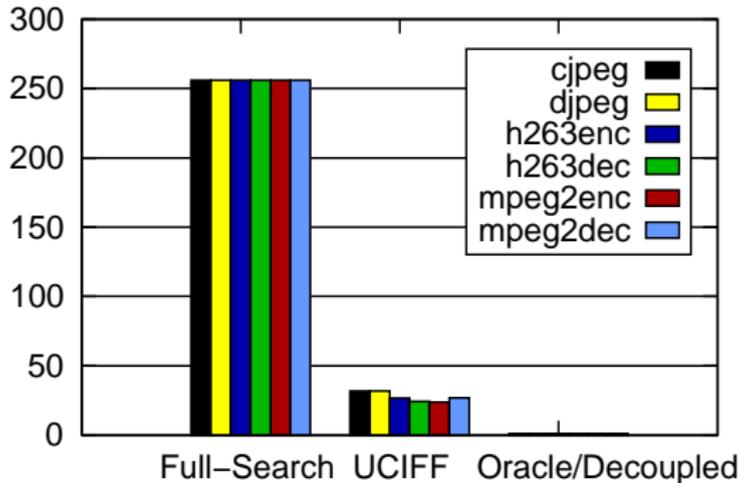
Conclusion

# Experimental Setup

- Compiler
  - GCC-4.5.0
  - Modified Haifa-Scheduler
  - Energy model built into the scheduler
- Architecture
  - IA64-based 4-cluster/4-issue VLIW 1-cycle inter-cluster delay
  - 4 possible frequencies. Fastest:Slowest = 7:4
- Benchmarks
  - MediabenchII Video Benchmark suite
- Compare
  - Decoupled, Full-Search, Oracle, UCIFF

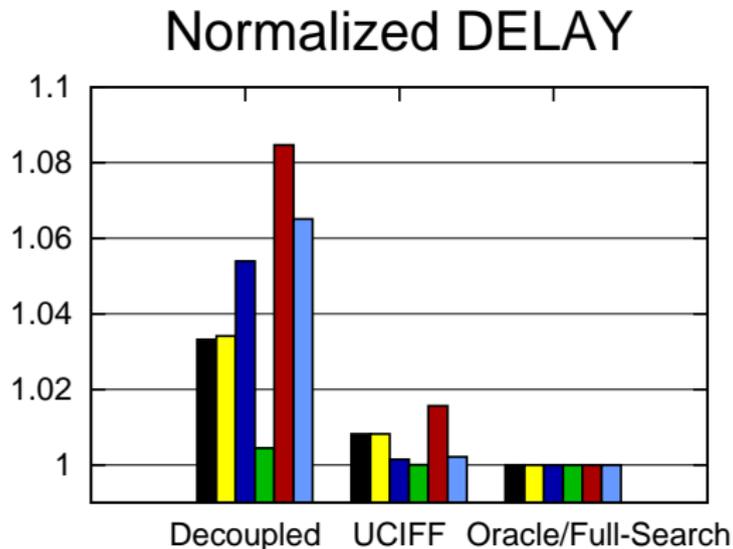
# Time Complexity Comparison

## Normalized Scheduled Instructions



- $5\times$  faster than Full-Search
- $30\times$  slower than theoretical oracle

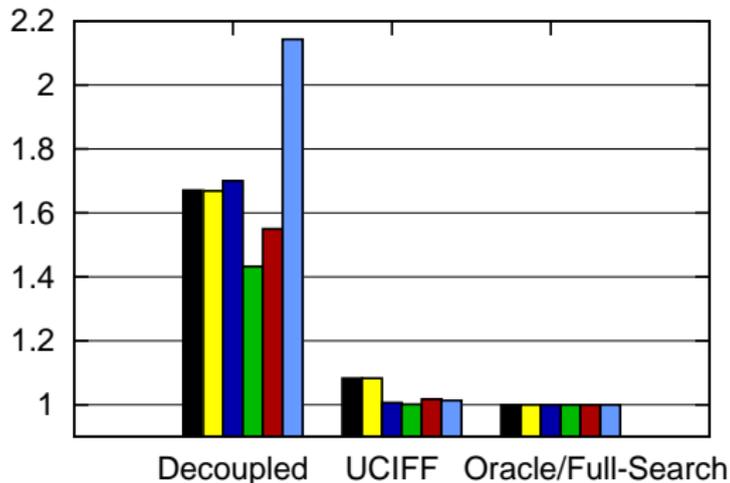
## Delay Comparison



- Decoupled about 5% worse than Oracle
- UCIFF almost identical to Oracle
- Delay is biased towards **high frequencies**

# ED<sup>2</sup> Comparison

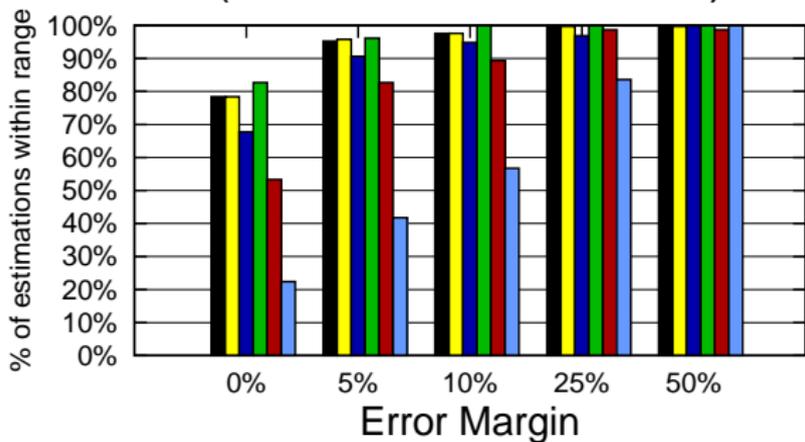
## Normalized ED2



- Decoupled 1.6× worse than oracle
- UCIFF within 5% of Oracle
- *ED<sup>2</sup>* is hard to estimate

# Delay Estimation Accuracy

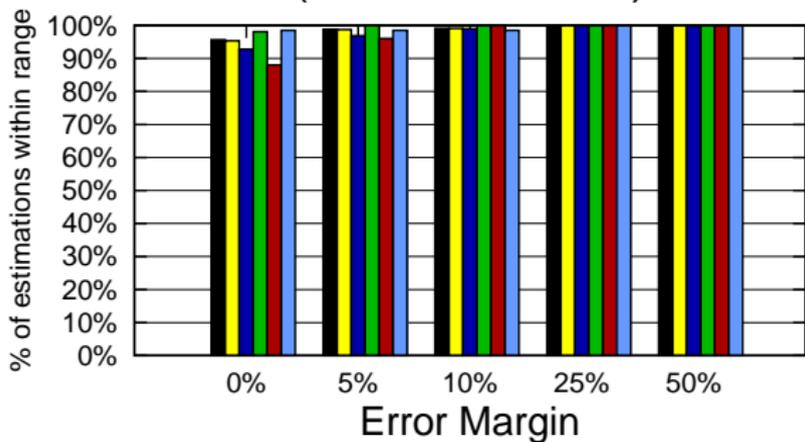
(DECOUPLED, DELAY)



- Decoupled: only 60% of estimations are the same as Oracle

# Delay Estimation Accuracy

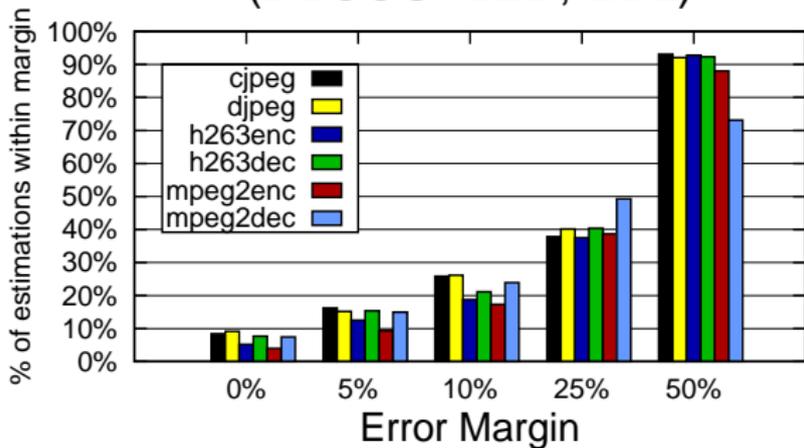
(UCIFF, DELAY)



- Decoupled: only 60% of estimations are the same as Oracle
- UCIFF very close to Oracle (90%)

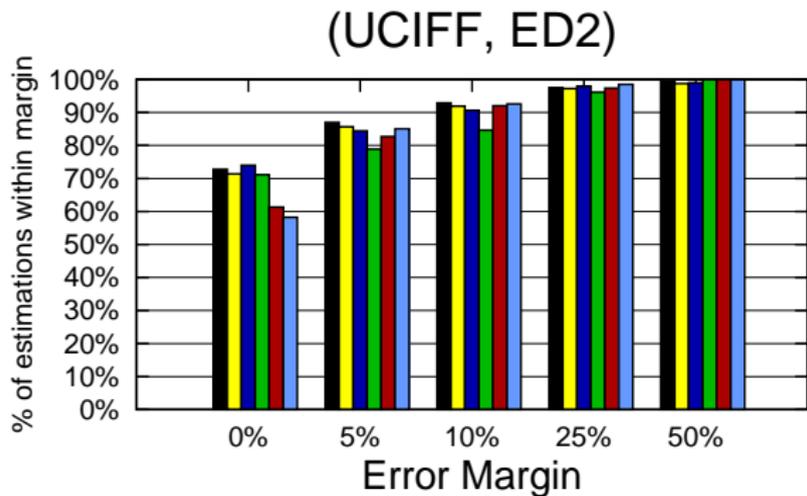
# ED<sup>2</sup> Estimation Accuracy

(DECOUPLED, ED2)



- Decoupled is very inaccurate 40% of estimations within 25% of the Oracle

## $ED^2$ Estimation Accuracy



- Decoupled is very inaccurate 40% of estimations within 25% of the Oracle
- UCIFF is at 95% within 25% of Oracle

## Conclusion

- Proposed UCIFF, a unified scheduling algorithm that
  - Performs cluster assignment
  - Performs instruction scheduling
  - Selects cluster frequenciesin a heterogeneous clustered VLIW
- UCIFF is more accurate and generates better schedules than the current state-of-the-art
- UCIFF is faster than Full-Search while generating code of equivalent quality



# UCIFF: Unified Clustering, Instruction scheduling and Fast Frequency selection for Heterogeneous Clustered VLIW

Vasileios Porpodas and Marcelo Cintra

University of Edinburgh



## Backup slides

- Bibliography
- CGO'07 Estimations
- Scheduling for heterogeneous (various freq.)
- UCIFF Energy Model
- DVFS regions
- UCIFF Neighbors
- UCIFF Algorithm

## Bibliography

- [CGO'07] A. Aleta, J. Codina, A. Gonzalez, and D. Kaeli. Heterogeneous clustered vliw microarchitectures. In CGO, pages 354-366, 2007.

Backup Slides

# CGO'07 Energy & Performance Estimation

- Performance Estimation:

- ① Perform Scheduling on a homogeneous architecture
- ② Cycles = cycles of homogeneous multiplied by the arithmetic mean of the clock periods of the heterogeneous clusters:  $Time = cycles_{hom} \times (\sum_{cl} T_{cl}) / NumOfClusters$

- Energy Estimation (similar to UCIFF except:)

- ① Dynamic energy of cluster is equal to a fraction of that of the homogeneous cluster, proportional to the ratio of the cluster's frequency to the average frequency:

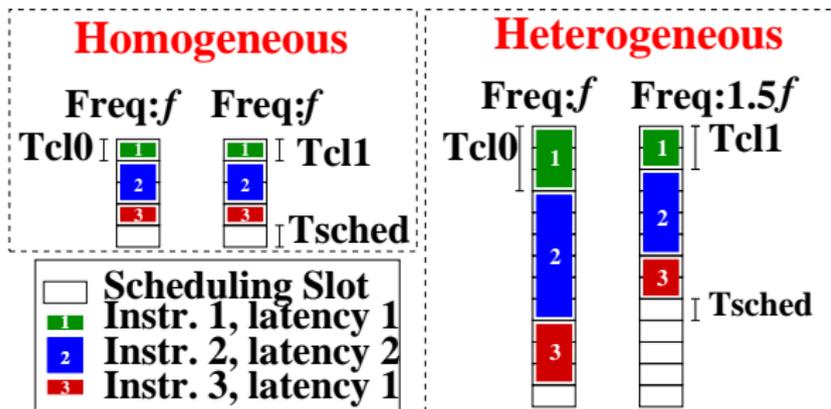
$$E_{dyn,ins}(cl) = E_{dyn,ins\_hom}(cl) \times f_{cl} / [\sum_{cl}(f_{cl}) / NumOfClusters]$$

- ② Energy of interconnect is equal to that of the homogeneous  $E_{dyn,icc} = P_{icc} \times NumICCs_{homogeneous}$

## Scheduling for clusters of various Freq.

- Scheduler's internal frequency is the lowest integer common multiple of all possible frequencies of all clusters ( $T_{sched}$ )
- Instruction latencies are specific to each cluster and are a multiple of the original latencies:

$$T_{cl}/T_{sched} \times \text{Original\_Latency}$$



## DVFS Region

- Scheduling regions are too small for DVFS (H/W limitation)
- Possible solutions:
  - Micro-Architecture: Push DVFS points into a FIFO queue and take the average
  - Software1: Sampling at a rate acceptable by H/W
  - Software2: Get an average single DVFS point for the whole program

# UCIFF Energy Model

- Total Energy:

- $E = \sum_{clusters} [E_{st}(cl) + E_{dyn}(cl)]$

- Static Energy:

- $E_{st}(cl) = P_{st} \times cycles_{cl} \times T_{cl}$

- $P_{st}(cl) = C_{st} \times V_{cl}$

- Dynamic Energy:

- $E_{dyn}(cl) = E_{dyn,ins}(cl) + E_{dyn,icc}$

- $E_{dyn,ins}(cl) = \sum_{ins} [P_{ins}(cl) \times Latency(ins, cl)]$

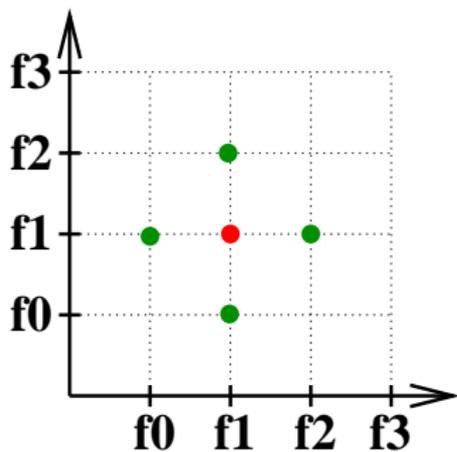
- $P_{ins}(cl) = C_{dyn} \times f_{cl} \times V_{cl}^2$

- $E_{dyn,icc} = P_{icc} \times NumICCs$

- $P_{icc} = C_{dyn} \times f_{fastest} \times V_{fastest}^2$

# UCIFF Neighbors

- The Configuration  $\{f_{na}, f_{nb}, f_{nc}, \dots\}$  is a UCIFF neighbor of  $\{f_a, f_b, f_c, \dots\}$  if  $n_x = x$  for all  $x$  except one (say  $y$ ) such that  $|n_y - y| < NDistance$ .



- **Freq. configuration**
- **Neighbors of •**

---

```
1  /* Unified Cluster assignment Instr. Scheduling and Fast Frequency selection.
2  In1: METRIC_TYPE that the scheduler should optimize for.
3  In2: Schedule STEP instructions before evaluating and getting the best.
4  In3: STEPVAR: Decrement STEP by STEPVAR upon each evaluation.
5  In4: NEIGHBORS: The number of neighbors per cluster.
6  Out: Scheduled Code and Best Frequency Configuration. */
7  uciff (METRIC_TYPE, STEP, STEPVAR, NEIGHBORS)
8  {
9      Schedule for STEP cycles and find the Best Freq Configuration (BFC)
10     do
11         if (BFC not set)  /* If first run */
12             NEIGHBORS_SET = all frequency configurations
13         else
14             NEIGHBORS_SET = neighbors of BFC /*up to NEIGHBORS per cluster*/
15         for FCONF in NEIGHBORS_SET
16             /* Partially schedule the ready instructions of FCONF frequency
17                ↳ configuration for STEP cycles, optimizing METRIC_TYPE */
18             SCORE = cluster_and_schedule (METRIC_TYPE, STEP, FCONF)
19             Store the scheduler's calculated SCORE into SCORECARD [FCONF]
20             Decrement STEP by STEPVAR until 1. /* Variable steps (optional) */
21             BFC = Best Freq Configuration of SCORECARD, clear SCORECARD
22         while there are unscheduled instructions in active set
23         return BFC and scheduled code of BFC
24     }
```

---

```

1  /* In1: METRIC_TYPE: The metric type that the scheduler will optimize for.
2     In2: STEP: Num of instrs to schedule before switching to next freq. conf.
3     In3: FCONF: The architecture's current frequency configuration.
4     Out: Scheduled Code and metric value. */
5  cluster_and_schedule (METRIC_TYPE, STEP, FCONF)
6  {
7     /* Restore ready list for this frequency configuration */
8     READY_LIST = READY_LIST_ARRAY [FCONF]
9     /* Restore current cycle. CYCLE is the scheduler's internal cycle. */
10    CYCLE = LAST_CYCLE [FCONF]
11    Restore the Reservation Table state that corresponds to FCONF
12    while (instructions left to schedule && STEP > 0)
13        update READY_LIST with ready to issue at CYCLE, include deferred
14        sort READY_LIST based on list-scheduling priorities
15        while (READY_LIST not empty)
16            select INSN, the highest priority instruction from the READY_LIST
17            create LIST_OF_CLUSTERS[] that INSN can be scheduled at on CYCLE
18            BEST_CLUSTER=best of LIST_OF_CLUSTERS[] by comparing for each cluster
19                ↳calculate_heuristic(METRIC_TYPE,CLUSTER,FCONF,INSN,IPCL[])
20            /* Try scheduling INSN on the best cluster */
21            if (INSN can be scheduled on BEST_CLUSTER at CYCLE)
22                schedule INSN, occupy LATENCY[FCONF][BEST_CLUSTER][INSN] slots
23                IPCL [CLUSTER] ++ /* count number of instructions per cluster */
24                remove INSN from READY_LIST
25            /* If failed to schedule INSN on best cluster, defer to next cycle */
26            if (INSN unscheduled)
27                remove INSN from READY_LIST and re-insert it at CYCLE + 1
28            /* No instructions left in ready list for CYCLE, then CYCLE ++ */
29            CYCLE ++
30            /* If we have scheduled for STEP cycles, finalize and exit */
31            if (CYCLE $>$ LAST_CYCLE[FCONF] + STEP)
32                Update READY_LIST_ARRAY[], LAST_CYCLE[] and Reservation Table
33                return

```

---

```
1 /* In1: METRIC_TYPE: The metric type that the scheduler will optimize for.
2    In2: CLUSTER: The cluster that INSN will be tested on.
3    In3: FCONF: The architecture's current frequency configuration.
4    In4: INSN: The instruction currently under consideration.
5    In5: IPCL: The Instruction count Per Cluster (for dyn energy).
6    Out: metric value of METRIC_TYPE if INSN scheduled on CLUSTER under FCONF*/
7 calculate_heuristic (METRIC_TYPE, CLUSTER, FCONF, INSN, IPCL[])
8 {
9     START_CYCLE = earliest cycle INSN can be scheduled at on CLUSTER
10    UCIFC_SC = START_CYCLE + LATENCY[FCONF][CLUSTER][INSN]
11    switch (METRIC_TYPE)
12        case ENERGY: return energy (CLUSTER, FCONF, UCIFC_SC, IPCL[])
13        case EDP: return edp (CLUSTER, FCONF, UCIFC_SC, IPCL[])
14        case ED2: return ed2 (CLUSTER, FCONF, UCIFC_SC, IPCL[])
15        case DELAY: return UCIFC_SC
16 }
```

---

Backup Slides