Investigation into Gang scheduling by integration of Cache in Multi core processors

Authors

Rupali¹, Shailja Kumari²

¹Research Scholar, Department of CSA, CDLU Sirsa

²Assistant Professor, Department of CSA, CDLU Sirsa,

Email- SSK.88@rediffmail.com

ABSTRACT

Objective of research is increase efficiency of scheduling dependent task using enhanced multithreading. gang scheduling of parallel implicit-deadline periodic task systems upon identical multiprocessor platforms is considered. In this scheduling problem, parallel tasks use several processors simultaneously. first algorithm is based on linear programming & is first one to be proved optimal for considered gang scheduling problem. Furthermore, it runs in polynomial time for a fixed number m of processors & an efficient implementation is fully detailed. second algorithm is an approximation algorithm based on a fixed-priority rule that is competitive under resource augmentation analysis in order to compute an optimal schedule pattern. Precisely, its speedup factor is bounded by $(2-1/m)$. In computer architecture, multithreading is ability of a central processing unit (CPU) or a single core within a multi-core processor to execute multiple processes or threads concurrently, appropriately supported by operating system. This approach differs from multiprocessing, as with multithreading processes & threads have to share resources of a single or multiple cores: computing units, CPU caches, & translation lookaside buffer (TLB). Multiprocessing systems include multiple complete processing units, multithreading aims to increase utilization of a single core by using thread-level as well as instruction-level parallelism.

Keywords: TLP, Response Time, Latency, throughput, multithreading, Scheduling

1. INTRODUCTION

The multithreading paradigm has become more popular as efforts to further exploit instruction-level parallelism have stalled since late 1990s. This allowed concept of throughput computing to re-emerge from more specialized field of transaction processing; even though it is very difficult to further speed up a single thread or single program, most computer systems are actually multitasking among multiple threads or programs. Thus, techniques that improve throughput of all tasks result within overall performance gains. The simplest type of multithreading occurs when one thread runs until it is blocked by an event that normally would create a long-latency stall. Such a stall might be a cache miss that has to access off-chip memory, that might take hundreds of CPU cycles for data to return. Instead of waiting for stall to resolve, a

threaded processor would switch execution to another thread that was ready to run. Only when data for previous thread had arrived, would previous thread be placed back on list of ready-to-run threads.

Conceptually, it is similar to cooperative multitasking used within real-time operating systems, within which tasks voluntarily give up execution time when they need to wait upon some type of event. This type of multithreading is known as block, cooperative or coarse-grained multithreading. The goal of multithreading hardware support is to allow quick switching between a blocked thread & another thread ready to run. To achieve this goal, hardware cost is to replicate program visible registers, as well as some processor control registers. Switching from one thread to another thread means hardware switches from using one register set to another; to switch efficiently between active threads, each

active thread needs to have its own register set. For example, to quickly switch between two threads, register hardware needs to be instantiated twice. Additional hardware support for multithreading allows thread switching to be done within one CPU cycle, bringing performance improvements. Also, additional hardware allows each thread to behave as if it were executing alone & not sharing any hardware resources with other threads, minimizing amount of software changes needed within application & operating system to support multithreading.

Many families of microcontrollers & embedded processors have multiple register banks to allow quick context switching for interrupts. Such schemes could be considered a type of block multithreading among user program thread & interrupt threads.

Fig 1 Multi thread Execution

The Objective of interleaved multithreading is to remove all data dependency stalls from execution pipeline. Since one thread is relatively independent from other threads, there is less chance of one instruction within one pipelining stage needing an output from an older instruction within pipeline. Conceptually, it is similar to preemptive multitasking used within operating systems; an analogy would be that time slice given to each active thread is one CPU cycle.

For example:

1. Cycle $i + 1$: an instruction from thread B is issued.
2. Cycle $i + 2$: an instruction from thread C is issued.

This type of multithreading was first called barrel processing, within which staves of a barrel represent pipeline stages & their executing threads. Interleaved, preemptive, fine-grained or

time-sliced multithreading are more modern terminology.

In addition to hardware costs discussed within block type of multithreading, interleaved multithreading has an additional cost of each pipeline stage tracking thread ID of instruction it is processing. Also, since there are more threads being executed concurrently within pipeline, shared resources such as caches & TLBs need to be larger to avoid thrashing between different threads.

Simultaneous multithreading

The most advanced type of multithreading applies to superscalar processors. Whereas a normal superscalar processor issues multiple instructions from a single thread every CPU cycle, within simultaneous multithreading (SMT) a superscalar processor could issue instructions from multiple threads every CPU cycle. Recognizing that any single thread has a limited amount of instruction-level parallelism, this type of multithreading tries to exploit parallelism available across multiple threads to decrease waste associated with unused issue slots.

For example:

1. Cycle i : instructions j & $j + 1$ from thread A & instruction k from thread B are simultaneously issued.
2. Cycle $i + 1$: instruction $j + 2$ from thread A , instruction $k + 1$ from thread B , & instruction m from thread C are all simultaneously issued.
3. Cycle $i + 2$: instruction $j + 3$ from thread A & instructions $m + 1$ & $m + 2$ from thread C are all simultaneously issued.

To distinguish other types of multithreading from SMT, term "temporal multithreading" is used to denote when instructions from only one thread could be issued at a time.

In addition to hardware costs discussed for interleaved multithreading, SMT has additional cost of each pipeline stage tracking thread ID of each instruction being processed. Again, shared resources such as caches & TLBs have to be sized

for large number of active threads being processed.

Implementations include DEC (later Compaq) EV8 (not completed), Intel Hyper-Threading, IBM POWER5, Sun Microsystems UltraSPARC T2, MIPS MT, & CRAY XMT.

2. LITERATURE REVIEW

Yeh-Ching Chung wrote on “Applications & Performance Analysis of A Compile-Time Optimization Approach for List Scheduling Algorithms on Distributed Memory Multiprocessors”

They have proposed a compile-time optimization approach, *bottom-up top-down duplication heuristic* (BTDH), for static scheduling of *directed+acyclic graphs* (DAGS) on *distributed memory multiprocessors* (DMMs). In this paper, they discuss applications of BTDH for *list scheduling algorithms* (LSAs). There are two ways to use BTDH for LSAs. BTDH can be used with a LSA to form a new *scheduling algorithm* (LSA/BTDH). It could be used as a *pure optimization algorithm* for a LSA (LSA-BTDH)..

Ishfaq Ahmad¹ & Yu-Kwong Kwok² wrote on “On Parallelizing Multiprocessor Scheduling Problem”

Existing heuristics for scheduling a node & edge weighted directed task graph to multiple processors could produce satisfactory solutions but incur high time complexities that tend to exacerbate within more realistic environments with relaxed assumptions. Consequently, these heuristics do not scale well & cannot handle problems of moderate sizes. The algorithm also exhibits an interesting trade-off between solution quality & speedup while scaling well with problem size.

Maruf Ahmed, Sharif M. H. Chowdhury wrote on List Heuristic Scheduling Algorithms for Distributed Memory Systems with Improved Time Complexity

They present a compile time list heuristic scheduling algorithm called *Low Cost Critical Path algorithm* (LCCP) for distributed memory

systems. LCCP has low scheduling cost for both homogeneous & heterogeneous systems. In some recent papers list heuristic scheduling algorithms keep their scheduling cost low by using a fixed size heap & a FIFO, where heap always keeps fixed number of tasks & excess tasks are inserted within FIFO. When heap has empty spaces, tasks are inserted within it from FIFO. Best known list scheduling algorithm based on this strategy requires two heap restoration operations, one after extraction & another after insertion. Our LCCP algorithm improves on this by using only one such operation for both extraction & insertion, that within theory reduces scheduling cost without compromising scheduling performance. In our experiment they compare LCCP with other well known list scheduling algorithms & it shows that LCCP is fastest among all.

Wayne F. Boyer wrote on “Non-evolutionary algorithm for scheduling dependent tasks within distributed heterogeneous computing environments”

The Problem of obtaining an optimal matching & scheduling of interdependent tasks within distributed heterogeneous computing (DHC) environments is well known to be an NP-hard problem. In a DHC system, task execution time is dependent on machine to which it is assigned & task precedence constraints are represented by a directed acyclic graph. Recent research within evolutionary techniques has shown that genetic algorithms usually obtain more efficient schedules than other known algorithms.

We propose a non-evolutionary random scheduling (RS) algorithm for efficient matching & scheduling of inter-dependent tasks within a DHC system. RS is a succession of randomized task orderings & a heuristic mapping from task order to schedule. Randomized task ordering is effectively a topological sort where outcome may be any possible task order for which task precedence constraints are maintained. A detailed comparison to existing evolutionary techniques (GA & PSGA) shows proposed algorithm is less complex than evolutionary techniques, computes

schedules within less time, requires less memory & fewer tuning parameters. Simulation results show that average schedules produced by RS are approximately as efficient as PSGA schedules for all cases studied & clearly more efficient than PSGA for certain cases.

3 RESEARCH METHODOLOGY

In computing, scheduling is method by which work specified by some means is assigned to resources that complete work. The work may be virtual computation elements such as threads, processes or data flows, that are within turn scheduled onto hardware resources such as processors, network links or expansion cards.

A scheduler is what carries out scheduling activity. Schedulers are often implemented so they keep all computer resources busy (as within load balancing), allow multiple users to share system resources effectively, or to achieve a target quality of service. Scheduling is fundamental to computation itself, & an intrinsic part of execution model of a computer system; concept of scheduling makes it possible to have computer multitasking with a single central processing unit (CPU).

A scheduler may aim at one of several goals, for example, maximizing *throughput* (total amount of work completed per time unit), minimizing *response time* (time from work becoming enabled until first point it begins execution on resources), or minimizing *latency* (the time between work becoming enabled & its subsequent completion), maximizing *fairness* (equal CPU time to each process, or more generally appropriate times according to priority & workload of each process). In practice, these goals often conflict (e.g. throughput versus latency), thus a scheduler would implement a suitable compromise. Preference is given to any one of concerns mentioned above, depending upon user's needs & objectives.

4. CHALLENGES WITHIN RESEARCH

Multiple threads could interfere with each other when sharing hardware resources such as caches or translation lookaside buffers (TLBs). As a result, execution times of a single thread are not improved but could be degraded, even when only one thread is executing, due to lower frequencies or additional pipeline stages that are necessary to accommodate thread-switching hardware.

Overall efficiency varies; Intel claims up to 30% improvement with its HyperThreading technology,^[1] while a synthetic program just performing a loop of non-optimized dependent floating-point operations actually gains a 100% speed improvement when run within parallel. On other hand, hand-tuned assembly language programs using MMX or AltiVec extensions & performing data pre-fetches (as a good video encoder might) do not suffer from cache misses or idle computing resources. Such programs therefore do not benefit from hardware multithreading & could indeed see degraded performance due to contention for shared resources.

From software standpoint, hardware support for multithreading is more visible to software, requiring more changes to both application programs & operating systems than multiprocessing. Hardware techniques used to support multithreading often parallel software techniques used for computer multitasking of computer programs. Thread scheduling is also a major problem within multithreading.

5. PARALLEL COMPUTING

Parallel computing is a type of computation in which many calculations are carried out simultaneously, operating on principle that large problems could often be divided into smaller ones, which are then solved at same time. There are several different forms of parallel computing: bit-level, instruction-level, data, & task parallelism. Parallelism has been employed for many years, mainly in high-performance computing, but

interest in it has grown lately due to physical constraints preventing frequency scaling. As power consumption (and consequently heat generation) by computers has become a concern in recent years, parallel computing has become dominant paradigm in computer architecture, mainly in form of multi-core processors. Parallel computing is closely related to concurrent computing—they are frequently used together, & often conflated, though two are distinct: it is possible to have parallelism without concurrency & concurrency without parallelism.

Fig 2

In parallel computing, a computational task is typically broken down in several, often many, very similar subtasks that could be processed independently & whose results are combined afterwards, upon completion. In contrast, in concurrent computing, various processes often do not address related tasks; when they do, as is typical in distributed computing, separate tasks may have a varied nature & often require some inter-process communication during execution.

6. RESULT AND DISCUSSION

Single processor output

```
disp('problem with 7 tasks and 7 processors');
u=[6 8 3 2 1 7 9];
p=[2 1 2 1 1 6 5];
m=7;
disp(strcat('platform with
',num2str(m),'processors'));
disp('task execution requirements:');
[Slice X]=ghp12(u,p,m);
disp('schedule slices(rows:tasks
1...n;column:schedule slices):');
```

```
disp(X);
disp('Slice is');
disp(Slice);
disp(strcat('schedule length:',num2str((X))));
ganghueristic function
function [Slice X]= ghp(u,p,m)
disp('reuirement');
disp(u);
disp('availability');
disp(p);
disp('no of processor');
disp(m);
c=u;
n=length(u);
disp('length is');
disp(n);
r=c;
disp('remaining processing times');
disp(r);
Slice=[];
X=[];
SchedCompleted=false;
tc=0;
eltn=0;
walltime=zeros(1:4);
k=1;
maxslice=max(r);

%matlabpool open
while(~SchedCompleted)
    used=0;
    sched=zeros(1,n);
    X(k)=maxslice;
    for i=1:n
        tic
        %for j=1:4
            %tc=tc+maxNumCompThreads(2^(j-1));
            if(r(i)>0 && used+p(i)<=m)
                sched(i)=1;
                used=used+p(i);
                X(k)=min(X(k),r(i));
            end
            walltime(i)=toc;
            eltn=eltn+walltime(i);
        %end
```

```

end
Slice(:,end+1)=sched;
r=r-(X(k).*sched);
if(r==0)
    SchedCompleted=true;
else
    k=k+1;
end
end
%matlabpool close
disp('cpu cycle');
disp(tic);
disp('time');
disp(eltm);
end

Code for Multithreaded based processor
disp('problem with 7 tasks and 7 processors');
u=[6 8 3 2 1 7 9];
p=[2 1 2 1 1 6 5];
m=7;
disp(strcat('platform with
',num2str(m),'processors'));
disp('task execution requirements:');
[Slice X]=ghp(u,p,m);
disp('schedule slices(rows:tasks
1...n;column:schedule slices):');
disp(X);
disp('Slice is');
disp(Slice);
disp(strcat('schedule length:',num2str((X))));
modified ganghueristic
function [Slice X]= ghp(u,p,m)
disp('reuirement');
disp(u);
disp('availability');
disp(p);
disp('no of processor');
disp(m);
c=u;
n=length(u);
disp('length is');
disp(n);
r=c;
disp('remaining processing times');

```

```

disp(r);
Slice=[];
X=[];
SchedCompleted=false;
tc=0;
eltm=0;
walltime=zeros(1:4);
k=1;
maxslice=max(r);

matlabpool open
while(~SchedCompleted)
    used=0;
    sched=zeros(1,n);
    X(k)=maxslice;
    for i=1:n
        tic
        for j=1:4
            tc=tc+maxNumCompThreads(2^(j-1));
            if(r(i)>0 && used+p(i)<=m)
                sched(i)=1;
                used=used+p(i);
                X(k)=min(X(k),r(i));
            end
            walltime(i)=toc;
            eltm=eltm+walltime(i);
        end
    end
    Slice(:,end+1)=sched;
    r=r-(X(k).*sched);
    if(r==0)
        SchedCompleted=true;
    else
        k=k+1;
    end
end
matlabpool close
disp('cpu cycle');
disp(tic);
disp('time');
disp(eltm);
end

```

Step 1

```

Command Window
>> example12
problem with 7 tasks and 7 processors
platform with 7 processors
task execution requirements:
requirement
    6    8    3    2    1    7    9

availability
    2    1    2    1    1    6    5

no of processor
    7

length is
    7

remaining processing times
    6    8    3    2    1    7    9

cpu cycle
    205877189538

time
    2.4375e-005

schedule slices(rows:tasks 1...n;column:schedule slices):
    1    1    1    3    2    5    9

Slice is
    1    1    1    1    0    0    0
    1    1    1    1    1    0    0
    1    1    1    0    0    0    0

```

Fig 1 Task Execution requirement

Step 2

```

no of processor
    7

length is
    7

remaining processing times
    6    8    3    2    1    7    9

cpu cycle
    205877189538

time
    2.4375e-005

schedule slices(rows:tasks 1...n;column:schedule slices):
    1    1    1    3    2    5    9

Slice is
    1    1    1    1    0    0    0
    1    1    1    1    1    0    0
    1    1    1    0    0    0    0
    1    1    0    0    0    0    0
    1    0    0    0    0    0    0
    0    0    0    0    1    1    0
    0    0    0    0    0    0    1

schedule length:1 1 1 3 2 5 9

```

Fig 2. Finding schedule length

We are very interested in your feedback regarding these capabilities.
Please send it to parallel_feedback@mathworks.com.

```

Submitted parallel job to the scheduler, waiting for it to start.
Connected to a matlabpool session with 4 labs.
Sending a stop signal to all the labs...
Waiting for parallel job to finish...
Performing parallel job cleanup...
Done.
cpu cycle
    206477613428

time
    0.2661

schedule slices(rows:tasks 1...n;column:schedule slices):
    6    2    1    7    9

Slice is
    1    0    0    0    0
    1    1    0    0    0
    0    1    1    0    0
    0    1    0    0    0
    0    0    1    0    0
    0    0    0    1    0
    0    0    0    0    1

schedule length:6 2 1 7 9

```

7. SCOPE OF RESEARCH

If several threads work on same set of data, they could actually share their cache, leading to better cache usage or synchronization on its values. If a thread gets a lot of cache misses, other threads could continue taking advantage of unused computing resources, that may lead to faster overall execution as these resources would have been idle if only a single thread were executed. Also, if a thread cannot use all computing resources of CPU (because instructions depend on each other's result), running another thread may prevent those resources from becoming idle.

REFERENCE

1. Remzi H. Arpaci-Dusseau; Andrea C. Arpaci-Dusseau (January 4, 2015). "Chapter 7: Scheduling: Introduction, Section 7.6: A New Metric: Response Time". *Operating Systems: Three Easy Pieces (PDF)*. p. 6. Retrieved February 2, 2015.
2. Paul Krzyzanowski (2014-02-19). "Process Scheduling: Who gets to run next?". *cs.rutgers.edu*. Retrieved 2015-01-11.
3. Abraham Silberschatz, Peter Baer Galvin & Greg Gagne (2013). *Operating System Concepts 9*. John Wiley & Sons, Inc. ISBN 978-1-118-06333-0.
4. Here is C-code for FCFS
5. Early Windows at Wayback Machine
6. Sriram Krishnan. "A Tale of Two Schedulers Windows NT & Windows CE".
7. Inside Windows Vista Kernel: Part 1, Microsoft Technet
8. "Vista Kernel Improvements".
9. "Technical Note TN2028 - Threading Architectures".
10. "Mach Scheduling & Thread Interfaces".
11. http://www.ibm.com/developerworks/aix/library/au-aix5_cpu/index.html#N100F6
12. Molnár, Ingo (2007-04-13). "[patch] Modular Scheduler Core & Completely Fair Scheduler [CFS]". *linux-kernel (Mailing list)*.