

Latency-preserving software pipelining of predicated reservation tables for distributed hard real-time applications



Thomas Carle – Dumitru Potop-Butucaru
INRIA Paris-Rocquencourt, FRANCE
Team AOSTE


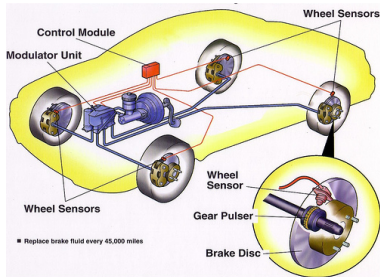
Outline

- Throughput optimization problem
- Previous work (software pipelining)
- System models (pipelined and non-pipelined)
- Pipelining algorithms
- A complex example
- Conclusion and future work

Application areas

Complex embedded control applications:

- Cyclic, periodic execution
- Safety-critical applications
 - Hard Real-Time constraints
 - Focus on functional and temporal correctness
- Distributed implementations



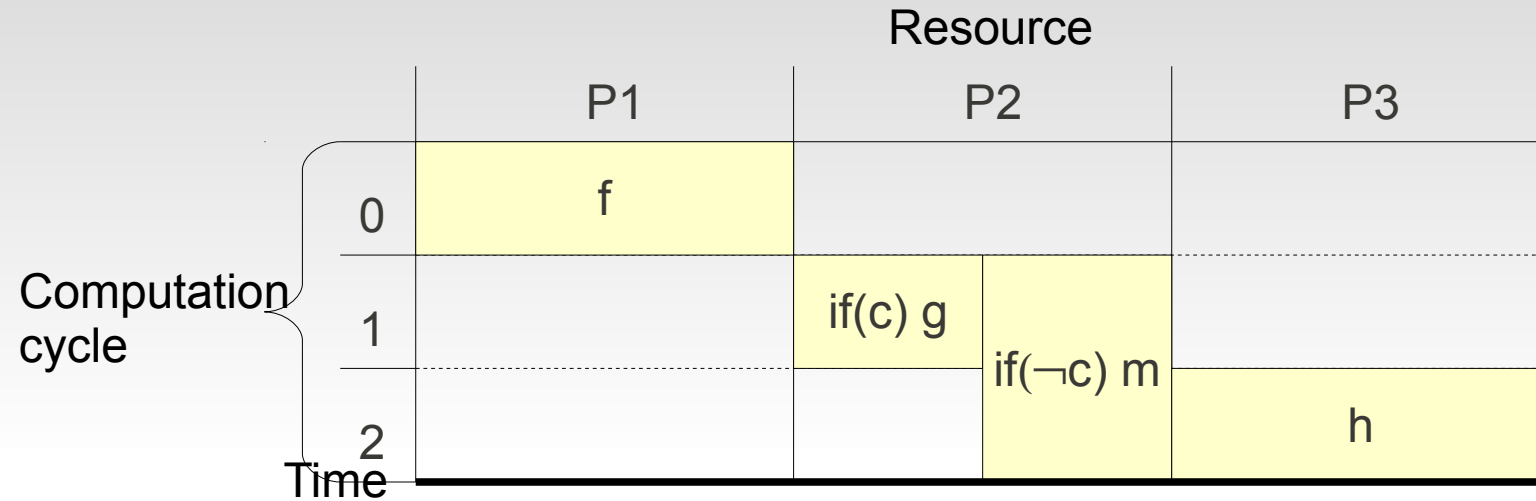
Static scheduling

Our work focuses on static schedules:
scheduling/reservation tables

Validated by industrial standards: **ARINC 653**,
AUTOSAR, FlexRay,...

Defines one cycle of execution, repeated
periodically


Motivating example



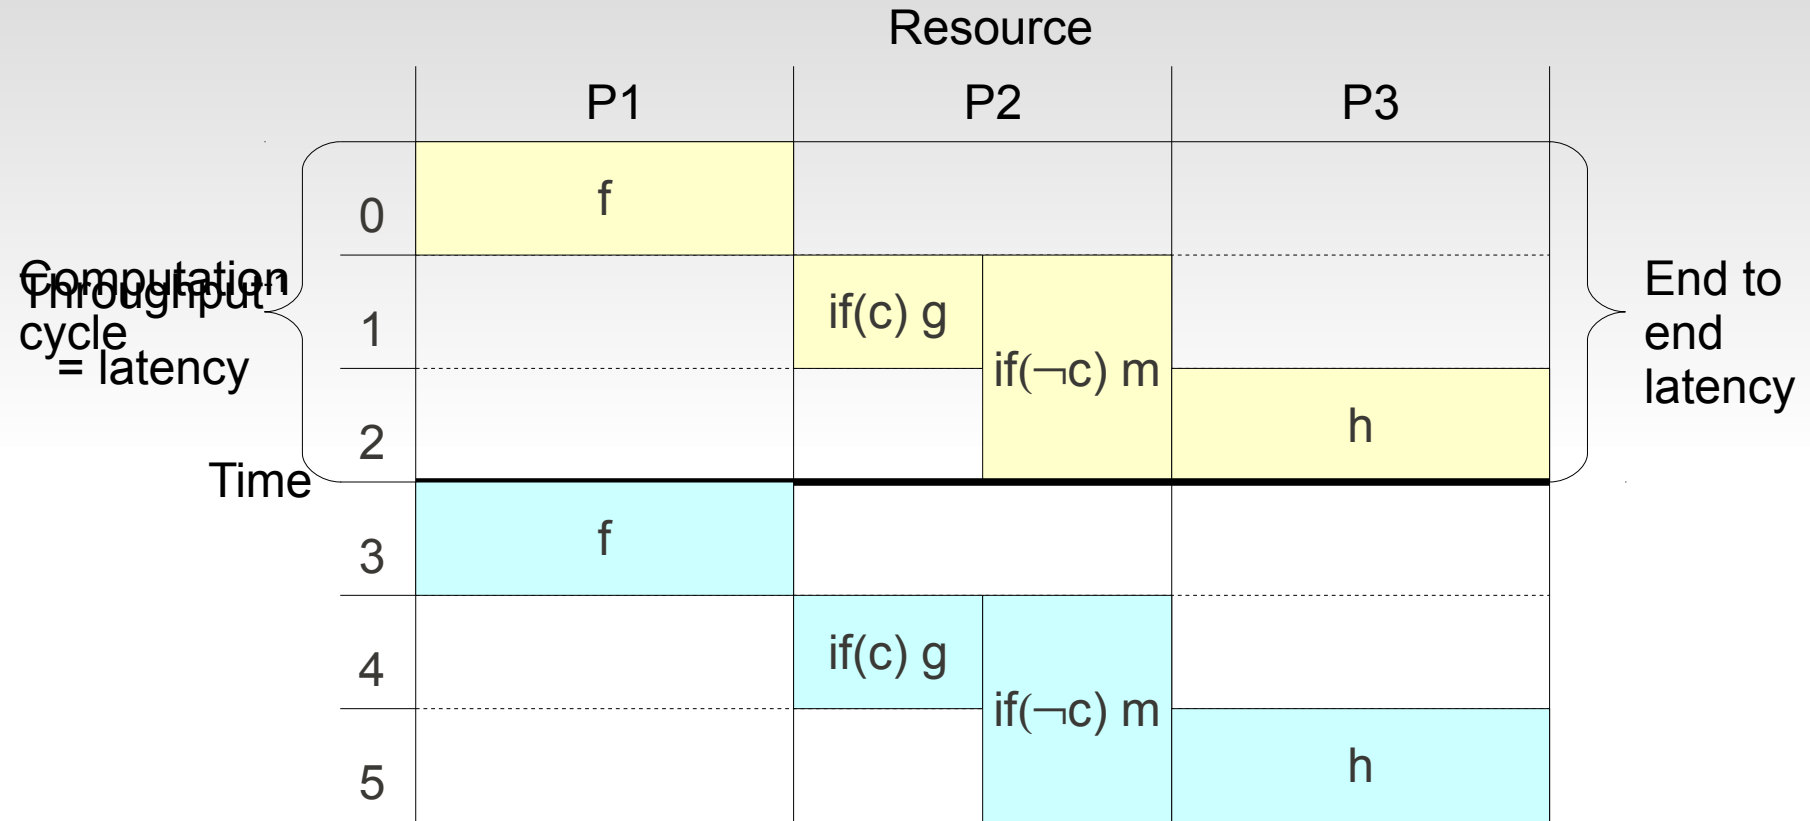
Code:

```

v2:=v2_init;
loop
  (v1,c)=f(v2);
  if c then v2:=g(v1)
  else m(v1);
  h(v2);
end
  
```



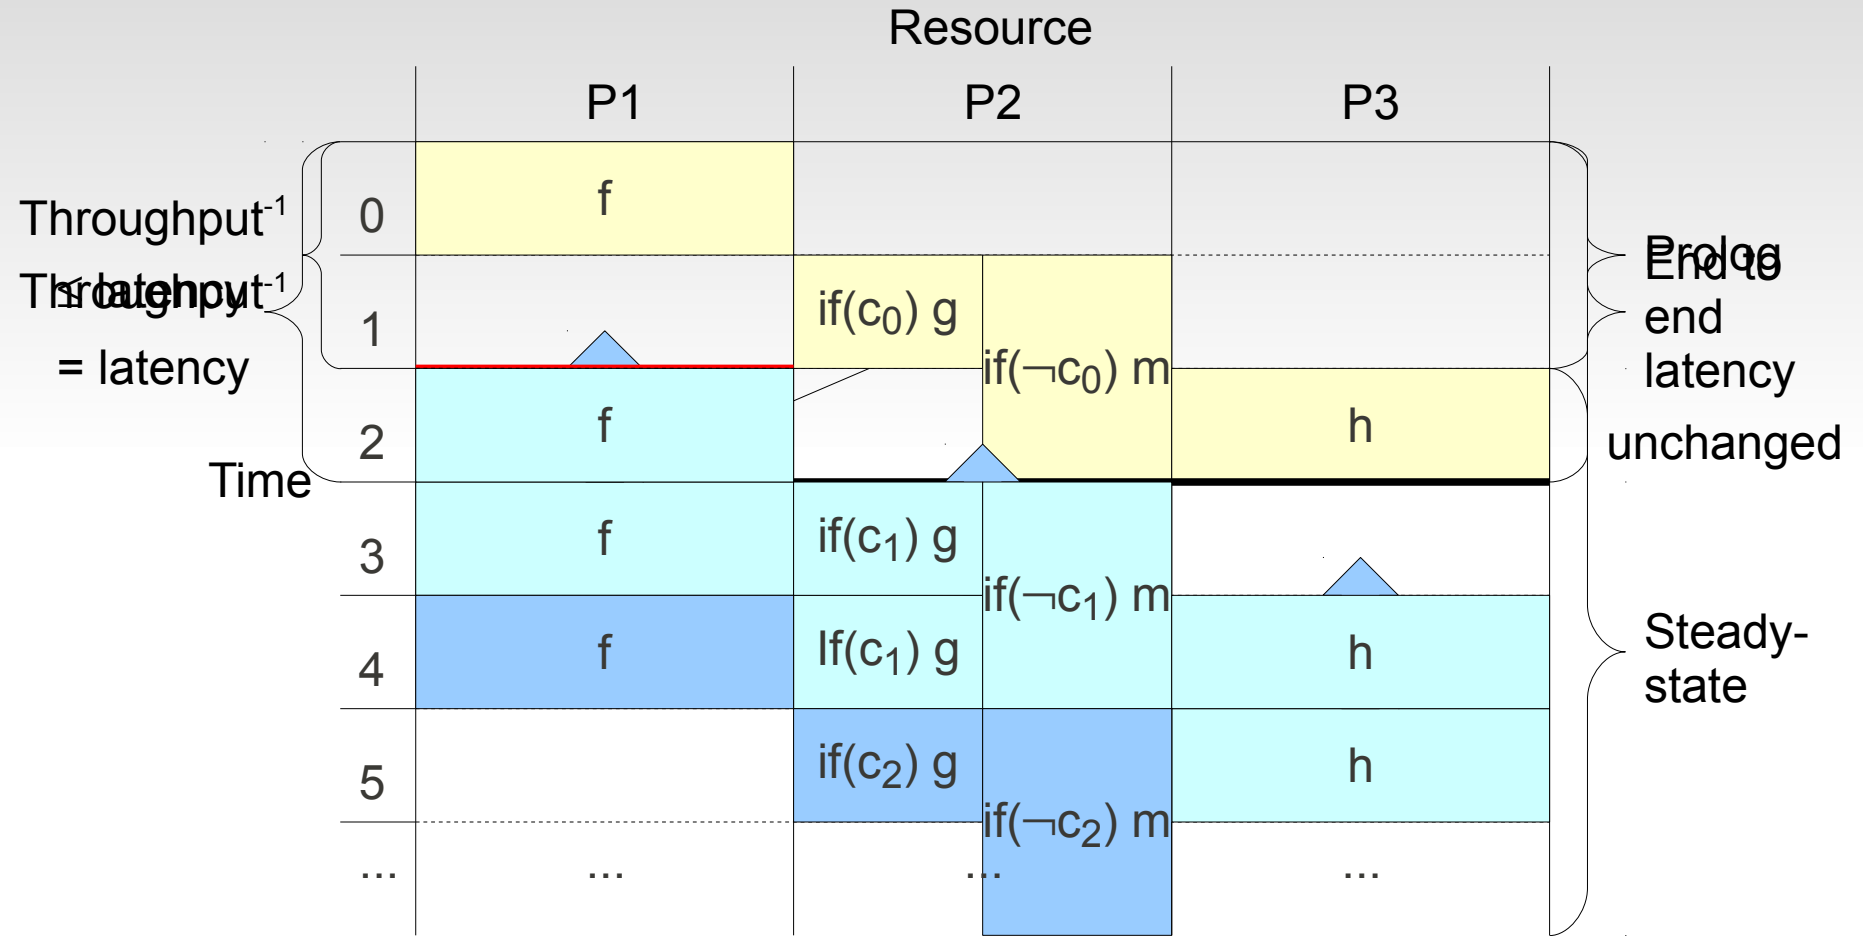
Motivating example



Latency: number of time units between the beginning and the end of the execution of a cycle

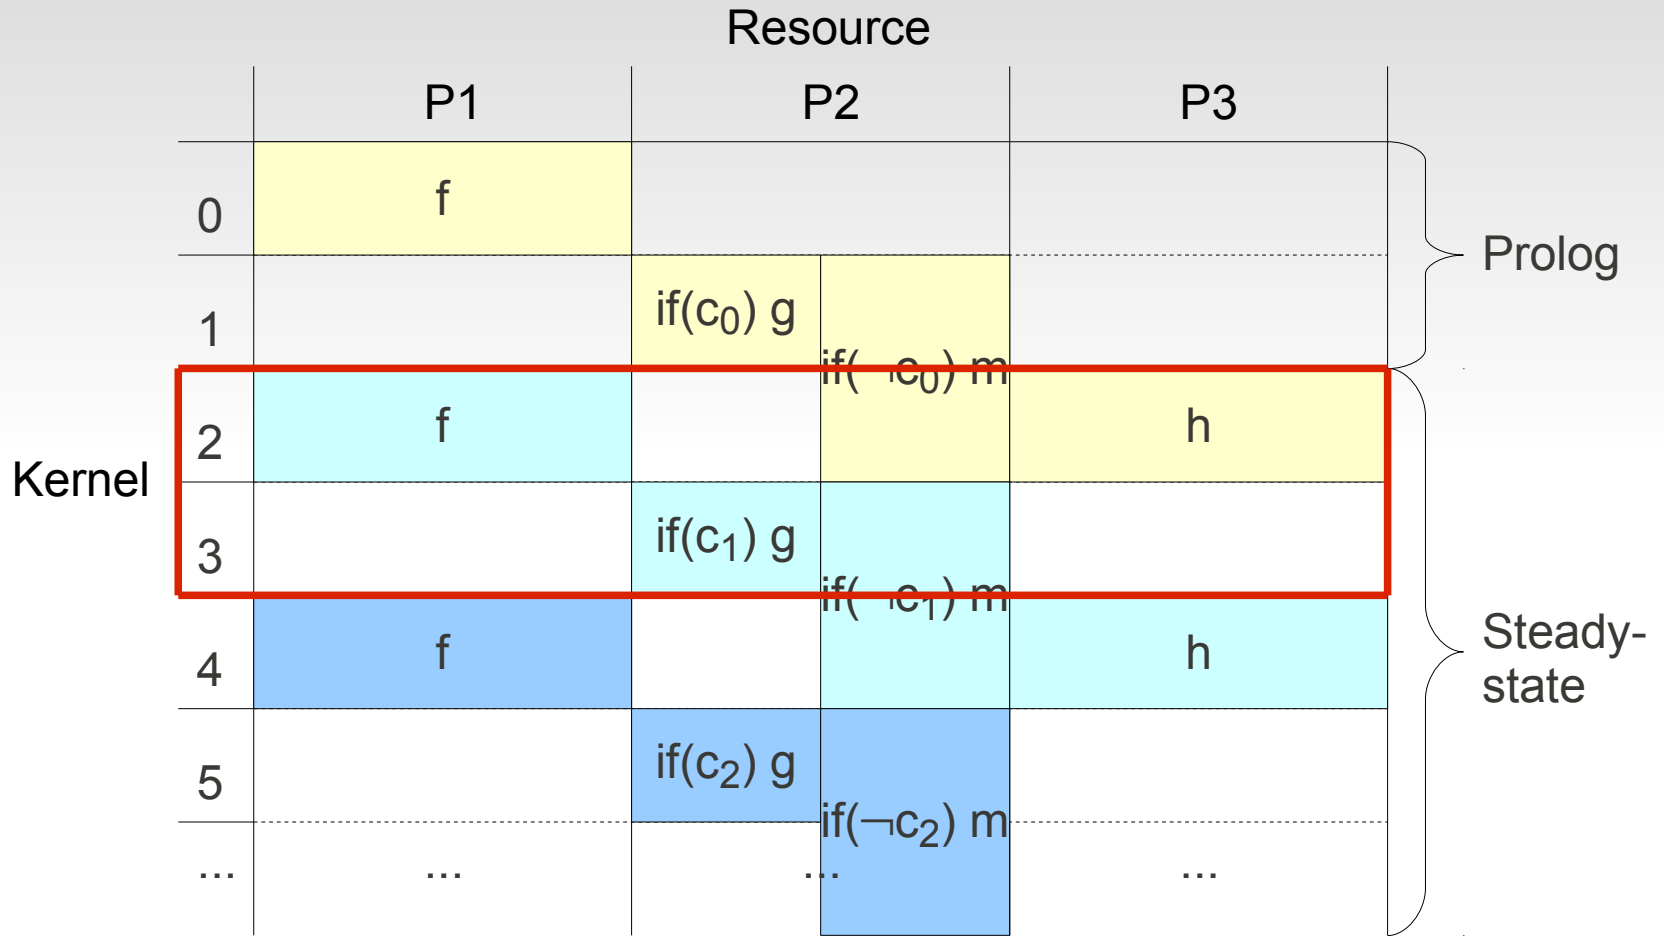
Throughput: number of cycles executed in one time unit

Our objective



Goal : increase throughput while keeping the system's latency, I/O function, and periodic behaviour

Our objective



Prolog and steady-state are instances of the kernel

Previous work(1): Software Pipelining

Scheduling techniques for **parallelizing loop computations**:

- Developed since the 1980's,
- First aimed at massively parallel architectures such as **VLIW** and superscalar machines,
- Now common optimization, present in most **compilers**,
- Similar to hardware pipelining: out-of-order execution,
- Reordering done in the compiler instead of in the processor.

Software Pipelining vs our work

- **Low-level** vs **coarse-grain** code generation technique
- Goal : **optimize average-case throughput by reorganizing operations order to take advantage of parallelism** vs **optimize worst-case throughput without degrading the cycles latency by preserving the intra-cycle scheduling**
- **No periodicity** for applications with data-dependent control vs **preservation of the periodic behaviour of the application**
- **Low degree of control over operators/functional units for conditional execution** vs **exploitation of conditional execution to improve the pipelining process**

Previous work(2): Retiming

- Optimization method in which **registers** in a synchronous circuit are **relocated** in order to **improve the throughput** or memory consumption of an application,
- Very similar to our techniques e.g. **no increase in latency** after applying the retiming techniques, preservation of the **I/O function**,
- Nevertheless: **no support for conditional execution/predication.**

Previous work(3): Real-Time Software Pipelining

- Builds a pipelined schedule for the application,
- Demonstrated on e.g. multimedia streaming applications,
- Again, no optimization for conditional execution

Elements of our approach

Architecture model

Initial non-pipelined
scheduling table

Algorithms

Pipelined
scheduling table

We design low-level implementation models that can be integrated at the end of the development cycle

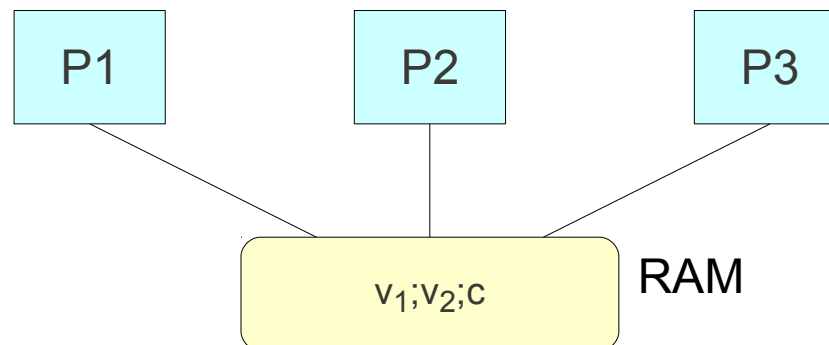
Architecture model

Bipartite undirected graph: $A = \langle P, M, C \rangle$, where:

- P: "processors", i.e. computation and communication resources capable of independent execution (Processors, DMAs, ...),
- M: RAM blocks,
- $(P, M) \in C$ indicates that processor P has direct access to memory block M.

RAM blocks: sets of disjoint untyped memory cells

Example:



Reservation/Scheduling table

S=<p,O,Init>, where :

- p : activation period of execution cycles, equal to the length of the reservation table,
- O : Set of scheduled operations,
- Init : set of initial values of all memory cells (can be *nil* or a constant).

Reservation/Scheduling table

Scheduled operation o :

- $In(o)$: set of **memory cells** whose data is used as **input** by o ,
- $Out(o)$: set of **memory cells written** by o ,
- $Guard(o)$: **execution condition of o** (predicate over memory cells),
- $Res(o)$: **set of "processors"** used during the execution of o ,
- $t(o)$: **start date** of o ,
- $d(o)$: **duration** of o , maximum time budget which can be ensured through WCET analysis.

Reservation/Scheduling table

Well-formed properties:

- **Exclusive resource use:** for O_1, O_2 scheduled on the same resource, if $\text{Guard}(O_1) \wedge \text{Guard}(O_2) \neq \text{false}$, then $t(O_1) \geq t(O_2) + d(O_2)$ or $t(O_2) \geq t(O_1) + d(O_1)$,
- **No data races :** if O_1 writes variable v_1 and O_2 uses (reads or writes) v_1 , then $t(O_1) \geq t(O_2) + d(O_2)$ or $t(O_2) \geq t(O_1) + d(O_1)$, or $\text{Guard}(O_1) \wedge \text{Guard}(O_2) = \text{false}$,
- **Causal correctness.**

Enough to describe non-pipelined schedules.

Pipelined Reservation/Scheduling table

For pipelined schedules, each scheduled operation o also has a **start index** $fst(o)$. It accounts for the prologue phase, where operations progressively start to execute.

If operation o has $fst(o)=n$, it will first be executed in the **pipelined cycle** of index n .

Due to periodicity, the description of the schedule of the kernel with start indexes is enough to describe the whole execution of the system.

Memory elements can be modified to take into account the variable replication process (described later)

Pipelined Reservation/Scheduling table

Resource

	P1	P2	P3	
0	f			Prolog
1		if(C_0) g		
2	f		if($\neg C_0$) m	Steady-state
3		if(C_1) g		
4	f		if($\neg C_1$) m	
5		if(C_2) g	if($\neg C_2$) m	

Pipelined Reservation/Scheduling table

Resource

	P1	P2	P3	
0	f		h fst=1	Pipelined Iteration 0
1		if(C_0) g		
2	f		h fst=1	Pipelined Iteration 1
3		if(C_1) g		
4	f		h	Pipelined Iteration 2
5		if(C_2) g	if($\neg C_2$) m	

Pipelining algorithm

- Constraints:
 - need to respect inter-cycle data dependency
 - no two operations can use a "processor" at the same time
 - no memory cell can be written by an operation and used (written or read) by another at the same time
- Our algorithm
 - Enforces the fulfilment of these constraints
 - Incrementally builds the Data Dependency Graph of the application
 - Takes advantage of guards during pipelining (better than existing work)
 - Specific memory handling

Pipelining algorithm

- Relies on the incremental construction of the **Data Dependency Graph (DDG)** of the application i.e. the set $\{(o_1, o_2, n)\}$ for all o_1 and o_2 such that $\text{In}(o_2) \cap \text{Out}(o_1) \neq \emptyset$, and o_1 **happens n cycles** before o_2 ,
- Uses an **SSA transformation** before performing a symbolic execution of the different iterations in order to construct the DDG.

Pipelining algorithm

	P1	P2	P3	P4
0	$c := \neg c$			
1		if(c) $v_2 := f_1(v_1)$	if($\neg c$)	
2		$w_2 := g_1(w_1)$		
3			if(c)	if($\neg c$)
4			$v_3 := f_2(v_2)$	$w_3 := g_2(w_2)$
5				if(c)
6				$v_1 := f_3(v_3)$
				if($\neg c$)
				$w_1 := g_3(w_3)$

Pipelining algorithm

	P1	P2		P3		P4	
0	$c := \neg c$						
1		if(c)	if($\neg c$)				
2	$c_1 := \neg c$	$v_2 := f_1(v_1)$	$w_2 := g_1(w_1)$				
3		if(c_1)	if($\neg c_1$)	if(c)	if($\neg c$)		
4		$v_2 := f_1(v_1)$	$w_2 := g_1(w_1)$	$v_3 := f_2(v_2)$	$w_3 := g_2(w_2)$		
5				if(c_1)	if($\neg c_1$)	if(c)	if($\neg c$)
6				$v_3 := f_2(v_2)$	$w_3 := g_2(w_2)$	$v_1 := f_3(v_3)$	$w_1 := g_3(w_3)$
7						if(c_1)	if($\neg c_1$)
8						$v_1 := f_3(v_3)$	$w_1 := g_3(w_3)$

Pipelining algorithm

	P1	P2		P3		P4	
0	$c := \neg c$						
1		if(c)	if($\neg c$)				
2	$c_1 := \neg c$	$v_2 := f_1(v_1)$	$w_2 := g_1(w_1)$				
3		if(c_1)	if($\neg c_1$)	if(c)	if($\neg c$)		
4		$v_2 := f_1(v_1)$	$w_2 := g_1(w_1)$	$v_3 := f_2(v_2)$	$w_3 := g_2(w_2)$		
5				if(c_1)	if($\neg c_1$)	if(c)	if($\neg c$)
6	$c_2 := \neg c_1$			$v_3 := f_2(v_2)$	$w_3 := g_2(w_2)$	$v_1 := f_3(v_3)$	$w_1 := g_3(w_3)$
7		if(c_2)	if($\neg c_2$)			if(c_1)	if($\neg c_1$)
8		$v_2 := f_1(v_1)$	$w_2 := g_1(w_1)$			$v_1 := f_3(v_3)$	$w_1 := g_3(w_3)$
9				if(c_2)	if($\neg c_2$)		
10				$v_3 := f_2(v_2)$	$w_3 := g_2(w_2)$		
11						if(c_2)	if($\neg c_2$)
						$v_1 := f_3(v_3)$	$w_1 := g_3(w_3)$

14/12/11

Complete algorithm: the first repetition is fully covered

Pipelining algorithm

	P1	P2		P3		P4	
0	$c := \neg c$						
1		if(c)	if($\neg c$)				
2	$c_1 := \neg c$	$v_2 := f_1(v_1)$	$w_2 := g_1(w_1)$				
3		if(c_1)	if($\neg c_1$)	if(c)	if($\neg c$)		
4		$v_2 := f_1(v_1)$	$w_2 := g_1(w_1)$	$v_3 := f_2(v_2)$	$w_3 := g_2(w_2)$		
5				if(c_1)	if($\neg c_1$)	if(c)	if($\neg c$)
6	$c_2 := \neg c_1$			$v_3 := f_2(v_2)$	$w_3 := g_2(w_2)$	$v_1 := f_3(v_3)$	$w_1 := g_3(w_3)$
7		if(c_2)	if($\neg c_2$)			if(c_1)	if($\neg c_1$)
8		$v_2 := f_1(v_1)$	$w_2 := g_1(w_1)$			$v_1 := f_3(v_3)$	$w_1 := g_3(w_3)$
9				if(c_2)	if($\neg c_2$)		
10				$v_3 := f_2(v_2)$	$w_3 := g_2(w_2)$		
11						if(c_2)	if($\neg c_2$)
						$v_1 := f_3(v_3)$	$w_1 := g_3(w_3)$

14/12/11

Make the repetition periodic : $\text{new_period} = \max_{(o_1, o_2, n) \in DDG} ((t(o_1) + d(o_1) - t(o_2)) / n)$

Pipelining algorithm

	P1	P2	P3	P4	
0	$c := \neg c$				
1		if(c) $v_2 := f_1(v_1)$	if($\neg c$)		
2		$w_2 := g_1(w_1)$			
3	$c_1 := \neg c$		if(c)	if($\neg c$)	
4		if(c_1) $v_2 := f_1(v_1)$	$v_3 := f_2(v_2)$	$w_3 := g_2(w_2)$	
5		$w_2 := g_1(w_1)$		if(c)	if($\neg c$)
6	$c_2 := \neg c_1$		if(c_1) $v_3 := f_2(v_2)$	$v_1 := f_3(v_3)$	$w_1 := g_3(w_3)$
7		if(c_2) $v_2 := f_1(v_1)$	if($\neg c_1$) $w_3 := g_2(w_2)$		
8		$w_2 := g_1(w_1)$		if(c_1) $v_1 := f_3(v_3)$	if($\neg c_1$) $w_1 := g_3(w_3)$
9			if(c_2) $v_3 := f_2(v_2)$		
10			if($\neg c_2$) $w_3 := g_2(w_2)$		
11				if(c_2) $v_1 := f_3(v_3)$	if($\neg c_2$) $w_1 := g_3(w_3)$

14/12/11

new_period := 3

Pipelining algorithm

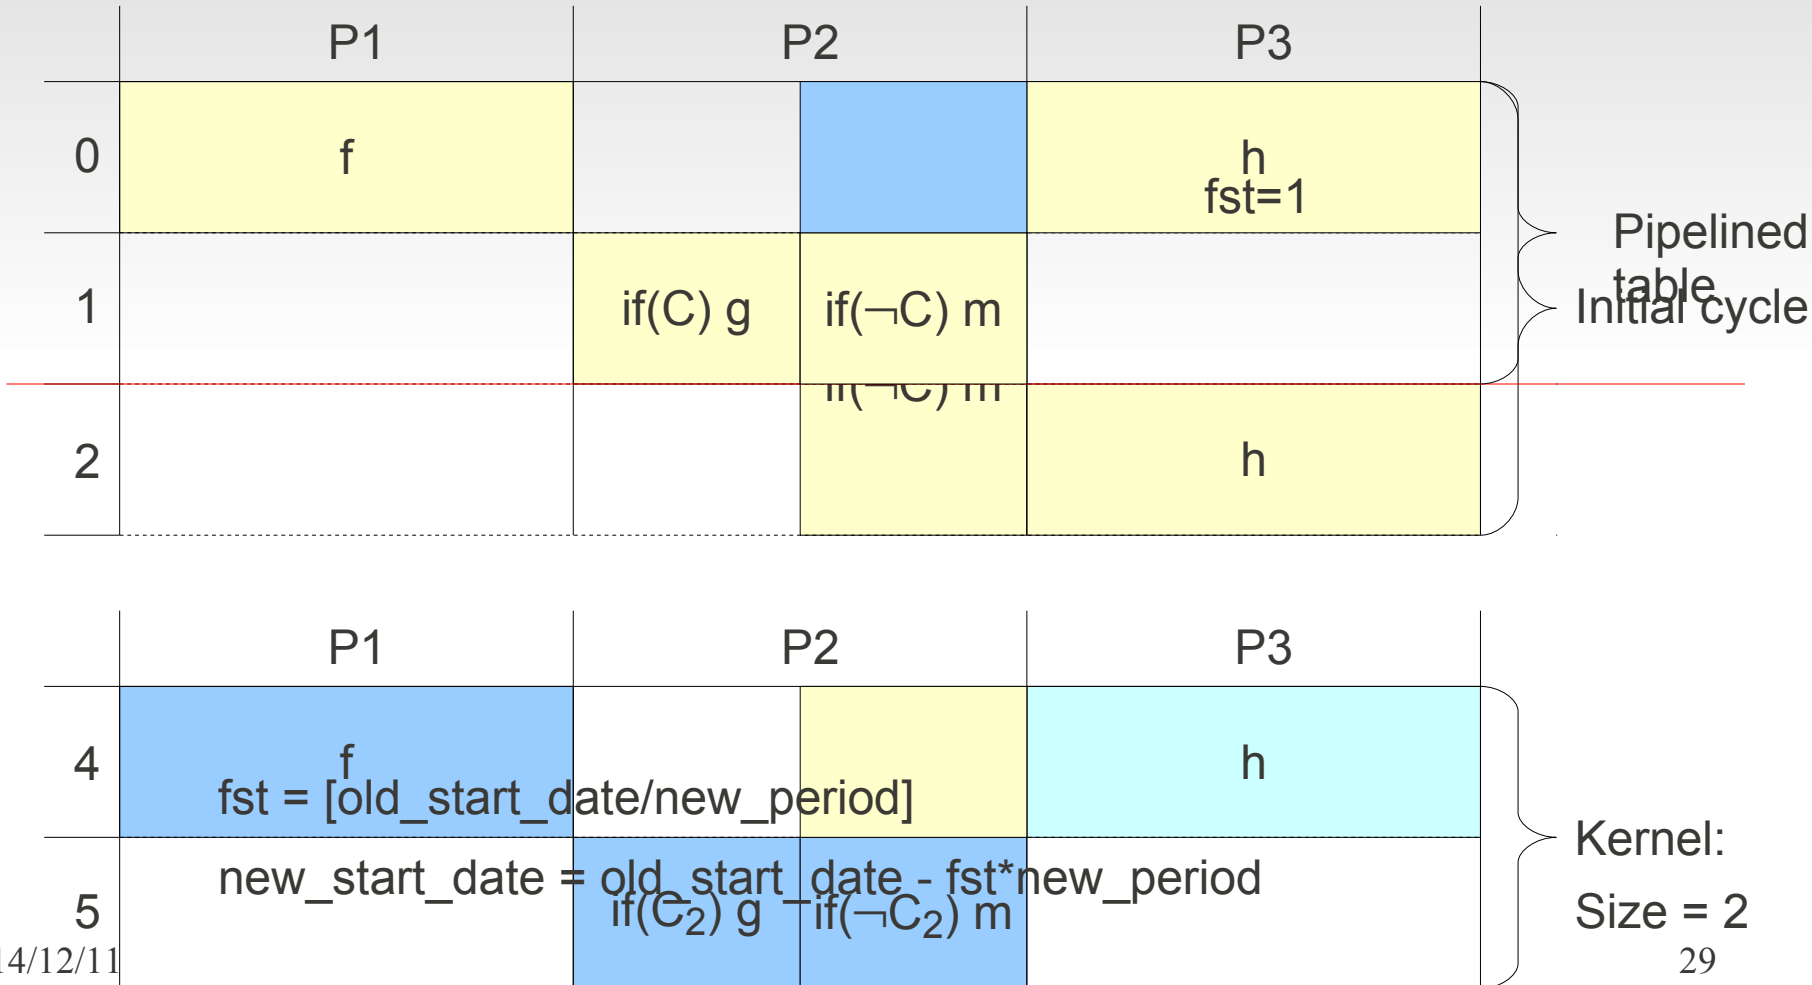
P1	P2		P3		P4	
$C := \neg C$						
	$v_2 := f_1(v_1)$ @C	$w_2 := g_1(w_1)$ @ $\neg C$				
$C := \neg C$			$v_3 := f_2(v_2)$ @C	$w_3 := g_2(w_2)$ @ $\neg C$		
	$v_2 := f_1(v_1)$ @C	$w_2 := g_1(w_1)$ @ $\neg C$			$v_1 := f_3(v_3)$ if(C_{i-2})	$w_1 := g_3(w_3)$ if($\neg C_{i-2}$)
$C_i := \neg C_{i-1}$			if(C_{i-1}) $v_3 := f_2(v_2)$	if($\neg C_{i-1}$) $w_3 := g_2(w_2)$		
	if(C_i) $v_2 := f_1(v_1)$	if($\neg C_i$) $w_2 := g_1(w_1)$			$v_1 := f_3(v_3)$ @C	$w_1 := g_3(w_3)$ @ $\neg C$
			$v_3 := f_2(v_2)$ @C	$w_3 := g_2(w_2)$ @ $\neg C$		
					$v_1 := f_3(v_3)$ @C	$w_1 := g_3(w_3)$ @ $\neg C$

14 Build the kernel


fst=0

fst=1

Pipelined Reservation/Scheduling table construction



Memory aspects



Memory aspects

Problem: if a variable is used at the same time in two or more pipelined iterations, that variable must be **replicated**.

Memory management is achieved the following way:

- The replication factor $\text{rep}(v)$ is computed for each variable v ,
- Each memory cell v of the initial non-pipelined scheduling table is replaced by **$\text{rep}(v)$ memory cells**, allocated on the same memory block as v , in a **cyclic fashion to the successive computation cycles**.

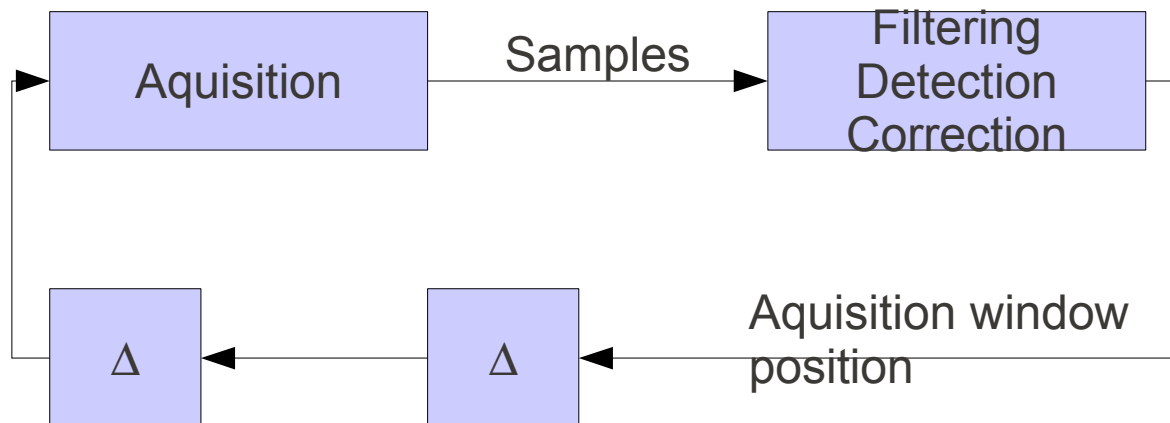
Knock controller example

Industrial case study

Function of the engine control unit for gasoline spark-ignition engines

Objective : choose ignition time for each cylinder at each rotation in order to maximize power output and avoid autoignition as much as possible

High-level description of the control loop function :

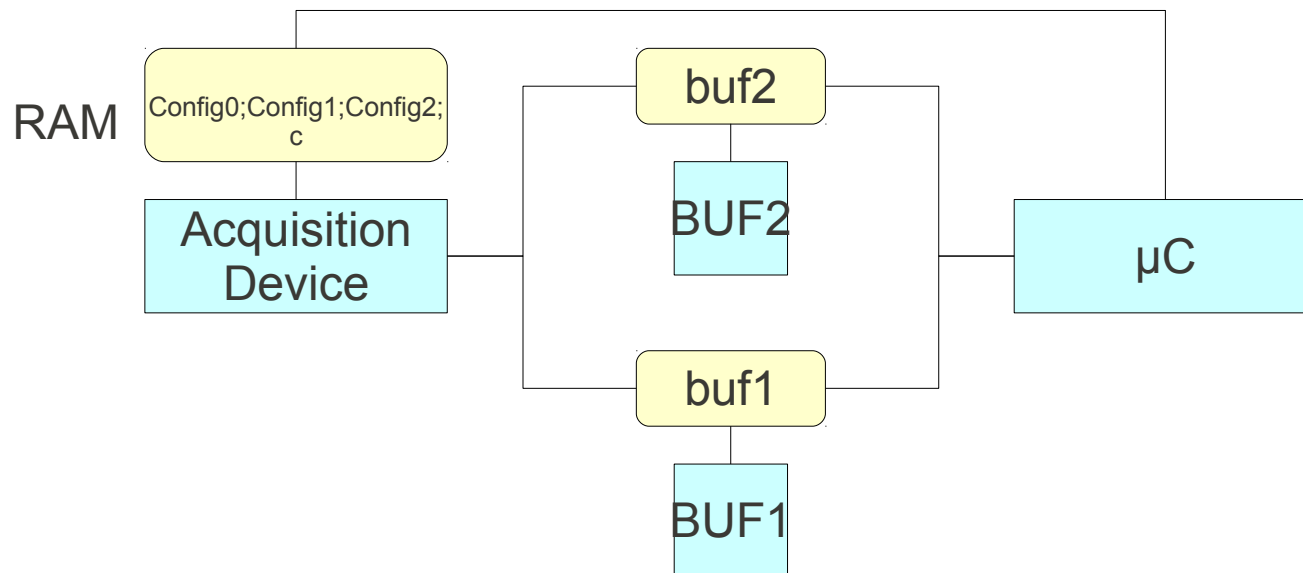


Knock controller example

One acquisition device

One microcontroller for filtering, detection, and correction

Two buffers and their independent controllers



Knock controller example

		Resources						
		AD		BUF 1	BUF 2	μC		
Rotation units	0	book						
	1	if(c) Acq ₁	if(\neg c) Acq ₂	if(c) Acq ₁		if(\neg c) Acq ₂		
	2							
	3							
	4			if(c) FDC ₁		if(\neg c) FDC ₂	if(c) FDC ₁	if(\neg c) FDC ₂
	5							

Resources

		Resources							
		AD		BUF 1	BUF 2	μC			
Rotation units	0	book							
	1	if(c) Acq ₁	if(\neg c) Acq ₂	if(c) Acq ₁			if(\neg c) Acq ₂		
	2								
	3	book							
	4	if(\neg c ₁) Acq ₂	if(c ₁) Acq ₁	if(c) FDC ₁	if(c ₁) Acq ₁	if(\neg c ₁) Acq ₂	if(\neg c) FDC ₂	if(c) FDC ₁	if(\neg c) FDC ₂
	5								
	6	book							
	7	if(c ₂) Acq ₁	if(\neg c ₂) Acq ₂	if(c ₂) Acq ₁	if(c ₁) FDC ₁	if(\neg c ₁) FDC ₂	if(\neg c ₂) Acq ₂	if(c ₁) FDC ₁	if(\neg c ₁) FDC ₂
	8								
			

Knock controller example

		Resources							
		AD		BUF 1	BUF 2		μC		
Rotation units	0	book							
	1	if(c_n) Acq ₁	if($\neg c_n$) Acq ₂	if(c_n) Acq ₁	if(c_{n-1}) FDC ₁ fst=1	if($\neg c_{n-1}$) FDC ₂ fst=1	if($\neg c_n$) Acq ₂	if($\neg c_{n-1}$) FDC ₂ fst=1	if(c_{n-1}) FDC ₁ fst=1
	2								

Experimental results

We demonstrated our algorithms on four examples:

- Embedded control application for the CyCab electric car: 27% reduction in cycle time
- Adaptive equalizer: 6 % reduction
- Knock controller: 50% reduction
- Simple example: 66% reduction. 100% resource usage

Conclusion

- Study of the state of the art in software pipelining
- Definition of formal models
- Definition of algorithms
- Implementation in a prototype
- Evaluation on study cases

Future work

- Enhance memory management: we can forbid memory replication when it is not necessary, but we cannot yet limit it by giving memory sizes,
- Exploit execution guards over partitioned architectures, maybe by using n-synchronous formalism to express and exploit repetition patterns during the pipelining process,
- Integrate pipelining in the initial scheduling process to obtain better trade-offs between latency/response-time, throughput, and resource usage: we would like to do this on a software radio use case we are currently working on.