Precision Timed Computing in the Synchronous Setting

Partha S Roop



Synchron 2010

29 November - 4 December 2010

The following people actively contributed to the PRETzel research. www.ece.auckland.ac.nz/~pretzel

- Sidharta Andalam
- Alain Girault
- Reinhard von Hanxleden
- Claus Traulsen

Layout

Introduction

- Synchronous Approach at Auckland
- Problems and motivations

2 Related Work

- The Berkeley-Columbia Approach
- Reactive Processors

3) The Auckland Approach

- PRET-C based predictable programming
- The Auckland PRET Architecture
- From logical to physical time
- Results

- Model checking: SoC bus protocols, web services composition, synchronous observers, real-time systems.
- Synchronous languages: Reactive and precision timed architectures, compilation, distribution, timing analysis.
- Industrial Informatics: PLC control, industrial buses, industrial PCs, standards and semantics, compilation, applications.



- A year-4 course on embedded system design.
- We teach both the RTOS approach and the synchronous approach.

Layout

Introduction

- Synchronous Approach at Auckland
- Problems and motivations

Related Work

- The Berkeley-Columbia Approach
- Reactive Processors

The Auckland Approach

- PRET-C based predictable programming
- The Auckland PRET Architecture
- From logical to physical time
- Results

Space Exploration





Toyota

- Feb 2010: Global recall of more than 400,000 of the auto-maker's 2010 hybrid models, for problems in their anti-lock braking systems.
- Overall, more than 8 million Toyota cars have been recalled globally due to accelerator problems.



http://edition.cnn.com/2010/BUSINESS/02/09/ japan.prius.recall/index.html

Current design approaches of Time Critical Systems

- Rely on the well known theory of real-time scheduling.
- A set of tasks with timing parameters, which execute on RTOS.
- WCET derived through static analysis.
- Major issue with WCET analysis: "modern processors render WCET virtually unknowable; even simple problems demand heroic efforts" [Edwards and Lee, DAC 2007].

Motivation

Can we rethink the link between computation and timing so that we can:

- Design Precision Timed Architectures (PRET) [Edwards and Lee, DAC 2007], where processors elicit time as a property of computation.
- The goals of PRET would be to simplify static timing analysis.
- To achieve predictability without sacrificing throughput.

Multiple Active Context System (MACS)

B. Cogswell and Z. Segall. MACS: A predictable architecture for real time systems. In Real-Time Systems Symposium. IEEE CS Press, 1991.

- The use of a shared pipelined processor as a possible solution to the problem of predictability and high performance in real time systems.
- Task-level parallelism is used to maintain high processor throughput while individual threads execute at a relatively slow, but very predictable rate.

T1 Image: Constraint of the second seco	Fetch Decode	Execute
T2 T1 Time	T1	
	T2 T1	
T3 T2 T1	T3 T2	T1

Layout

Introduction

- Synchronous Approach at Auckland
- Problems and motivations

Related Work

• The Berkeley-Columbia Approach

Reactive Processors

3) The Auckland Approach

- PRET-C based predictable programming
- The Auckland PRET Architecture
- From logical to physical time
- Results

The PRET Machine Design [Lickly et al., CASES'08]

- PRET machine designed based on the SPARC's ISA.
- Uses multiple active contexts and latency hiding similar to the MACS architecture.
- Memory hierarchy is replaced by statically allocated scratch-pad memories.



< □ > < □ > < □ > < □ > < □ > < □ >

The PRET Programming Model

- Concurrent C programs with shared memory communication.
- Precise timing of threads using the deadi (deadline) instruction.
- Thread-safe programming by time interleaving the shared memory access.





イロト イポト イヨト イヨト

Layout

Introduction

- Synchronous Approach at Auckland
- Problems and motivations

Related Work

- The Berkeley-Columbia Approach
- Reactive Processors

The Auckland Approach

- PRET-C based predictable programming
- The Auckland PRET Architecture
- From logical to physical time
- Results

Reactive Processors

- Alternative platforms for reactive embedded systems.
- Use ISA support for environment interaction instead of interrupts.
- Have been used for direct execution of Esterel
 [www.ece.auckland.ac.nz/~proo003/ReactiveProcessors.php].

Features	Reactive Processors	Conventional Processors
Execution progression	Evolves in discrete instants separated by "tick delimiting instructions"	Evolves continuously
Preemption	Accomplished through event reaction block with implicit priority resolution and context switching in hardware	Accomplished through interrupt mechanism requiring explicit priority resolution, context saving and restoration in software
Concurrency	Synchronous parallel execution and broadcast communication between threads	Asynchronous execution requiring explicit message passing/rendezvous for communication between threads
View of the environment	Changes at discrete instants. Inputs are latched at the beginning and outputs are sustained till the end of a "tick".	Changes continuously Inputs can be read at any time, and outputs can be sustained for any duration.

Reactive Processors

- Alternative platforms for reactive embedded systems.
- Use ISA support for environment interaction instead of interrupts.
- Have been used for direct execution of Esterel
 [www.ece.auckland.ac.nz/~proc003/ReactiveProcessors.php].

Features	Reactive Processors	Conventional Processors
Execution progression	Evolves in discrete instants separated by "tick delimiting instructions"	Evolves continuously
Preemption	Accomplished through event reaction block with implicit priority resolution and context switching in hardware	Accomplished through interrupt mechanism requiring explicit priority resolution, context saving and restoration in software
Concurrency	Synchronous parallel execution and broadcast communication between threads	Asynchronous execution requiring explicit message passing/rendezvous for communication between threads
View of the environment	Changes at discrete instants. Inputs are latched at the beginning and outputs are sustained thrue end IBS are	Changes continuously Inputs can be read at any time, and outputs can be sustained on any UBSER ED

Roop (University of Auckland)



æ

イロト イ団ト イヨト イヨト

Example: The ReMIC Processor



Ξ.

イロト イポト イヨト イヨト

Roop (University of Auckland)

æ

イロト イ団ト イヨト イヨト

• Berkeley-Columbia approach is based on tailored processors that is resource intensive.

- Berkeley-Columbia approach is based on tailored processors that is resource intensive.
- This approach also mixes physical and logical time and is not portable.

- Berkeley-Columbia approach is based on tailored processors that is resource intensive.
- This approach also mixes physical and logical time and is not portable.
- Hard to derive the values of deadlines for thread-interleaved access.

- Berkeley-Columbia approach is based on tailored processors that is resource intensive.
- This approach also mixes physical and logical time and is not portable.
- Hard to derive the values of deadlines for thread-interleaved access.
- Reactive processors, while being predictable, can only execute pure Esterel.

Philosophy



3

イロト イポト イヨト イヨト

• Notion of concurrency: concurrency is logical but execution is sequential very similar to synchronous languages [Benveniste'03].

∃ ► < ∃ ►

- Notion of concurrency: concurrency is logical but execution is sequential very similar to synchronous languages [Benveniste'03].
- Notion of time: time is logical and the mapping of logical to physical time is done using static analysis of code.

- Notion of concurrency: concurrency is logical but execution is sequential very similar to synchronous languages [Benveniste'03].
- Notion of time: time is logical and the mapping of logical to physical time is done using static analysis of code.
- Design approach: Auckland Reactive PRET (ARPRET) architectures are designed by simple customization of soft-core processors.

Layout

Introduction

- Synchronous Approach at Auckland
- Problems and motivations

Related Work

- The Berkeley-Columbia Approach
- Reactive Processors

The Auckland Approach

- PRET-C based predictable programming
- The Auckland PRET Architecture
- From logical to physical time
- Results

Precision Timed C (PRET-C)

Simple set of synchronous extensions to C for:

- light-weight multithreading in C.
- all extensions implemented as C macros.
- provides thread-safe shared memory access.
- supports predictable programming by mapping logical time to physical time through static analysis.

Statement	Meaning		
ReactiveInput I	declares I as a reactive input coming from the		
	environment		
ReactiveOutput O	declares 0 as a reactive output emitted to the		
	environment		
PAR(T1,,Tn)	synchronously executes in parallel the n		
	threads Ti, with higher priority of Ti over		
	Ti+1		
EOT	marks the end of a tick (local or global de-		
	pending on its position)		
[weak] abort P when	immediately kills P when C is true in the pre-		
pre C	vious instant		

Table: PRET-C extensions to C.

æ

▶ ★ 문 ► ★ 문 ►

< 行い



Description

- Robot has to follow the line.
- Data from 3 sensors as an input from the environment.
- Two PWM outputs controlling two motors.



Description

- Three reactive inputs (sensors) from the environment.
- Two reactive outputs (PWM) controlling the motors.
- Two threads (Decoder, Motor Driver) execute concurrently.
- Speed information is shared between threads.

Environment

ReactiveInput int reset; ReactiveInput int sensor1,sensor2,sensor3; ... ReactiveOutput int leftMotor; //PWM ReactiveOutput int rightMotor; //PWM

∃ ► < ∃ ►

Environment

ReactiveInput int reset; ReactiveInput int sensor1,sensor2,sensor3;

• • •

```
ReactiveOutput int leftMotor; //PWM
ReactiveOutput int rightMotor; //PWM
```

main

```
int main(void){
   while(1){
      abort
      PAR(Decoder,MotorDriver);
      if (error1) set LED 1;
      if (error2) set LED 2;
      else set both LED;
      EOT;
   when(reset==1);
} }
```

```
thread Decoder(void) {
  //initialize variables;
  while(1) {
      process sensor data;
      store previous motor instructions;
      /*set motor control
      according to sensor data*/
      if (sensor==some value) {
         /*set goLeft,goRight
           according to the LUT*/
      }
      if(robot is turning){
         stuckCounter++;
      3
      else{reset stuckCounter;}
      if(stuckCounter over threshold){
         set error 2; stop robot;
      if(error > 0){break;}
      else{EOT;}
} }
```

Sensor		Motor		Error	
Left	Middle	Right	Left	Right	One
0	0	0	0	0	0
1	0	0	0	2	0
0	1	0	1	1	0
1	1	0	0	1	0
0	0	1	2	0	0
1	0	1	0	0	1
0	1	1	1	0	0
1	1	1	0	0	0

Table: Look-up table for tracking
Example: Line Following Robot

```
thread Decoder(void) {
  //initialize variables;
  while(1) {
      process sensor data;
      store previous motor instructions;
      /*set motor control
      according to sensor data*/
      if (sensor==some value) {
         /*set goLeft.goRight
           according to the LUT*/
      }
      if(robot is turning){
         stuckCounter++;
      }else{reset stuckCounter;}
      if(stuckCounter over threshold){
         set error 2; stop robot;
      if(error > 0){break;}
      else{EOT:}
} }
```

```
thread MotorDriver(void){
  int PWMCounter = 0;
  while(1){
    if(error > 0){
       stop robot; break;
    }else{
       if(PWMCounter < goLeft){</pre>
         leftMotor = 1; //left high
       }else{
         leftMotor = 0: //left low
       3
       if(PWMCounter < goRight){</pre>
         rightMotor = 1; //right high
       }else{
         rightMotor = 0; //right low
       PWMCounter++;
       if(PWMCounter == maxSpeed){
         PWMCounter = 0;
1 1 1
```

Comparison with Esterel

	_
Esterel	
[
<pre>emit A(0);</pre>	
pause;	
emit A(?B+1)	
11	
<pre>emit B(?A);</pre>	
pause;	
emit B(7)	
]	

PRET-C

PAR(T1, T2); ... void T1() A=0; EOT; A=B+1;

```
void T2()
B=A;
EOT;
B=7;
```

1/0				
Esterel	А	0	8	
	В	0	7	
PRET-C	А	0	1	
	В	0	7	

< ロ > < 同 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ >

3

Layout

Introduction

- Synchronous Approach at Auckland
- Problems and motivations

Related Work

- The Berkeley-Columbia Approach
- Reactive Processors

The Auckland Approach

• PRET-C based predictable programming

• The Auckland PRET Architecture

- From logical to physical time
- Results

ARPRET



Hardware extension (PFU) to the Microblaze (GPP) in order to achieve better throughput while simplifying WCRT analysis. The FastSimplex Link (FSL) provides a predictable communication.

Hardware support

Predictable Functional Unit (PFU)



Thread Table stores:

- priority, local tick.
- alive, suspension.
- spawn count.
- parent ID.

Abort Table stores:

- type of preemption (Weak/Strong)
- nesting of preemptions.
- monitoring signal.

< A

• preemption address.

Layout

Introduction

- Synchronous Approach at Auckland
- Problems and motivations

Related Work

- The Berkeley-Columbia Approach
- Reactive Processors

The Auckland Approach

- PRET-C based predictable programming
- The Auckland PRET Architecture
- From logical to physical time
- Results

Design Flow: PRET-C to WCRT

- PRET-C to Assembly: standard gcc based compilers can be used.
- Assembly to TCCFG : our code analyser.
- TCCFG to Model Checker : our FSM generator (XML).
- CTL temporal logic property checking: bounded integer checking.



Mapping PRET-C to TCCFG





Timed Concurrent Control Flow Graph



Intermediate format captures:

- concurrent control flow.
- preemption using checkaborts.
- execution cost of every node.

Stages

• PRET-C: simple synchronous extension to C (using macros).

Code

```
void main() {
  while(1) {
    abort
    PAR(sampler,display);
    when pre (reset);
    EOT;
}
```

э

→

- **Q** PRET-C: simple synchronous extension to C (using macros).
- O TCCFG : intermediate format.



Overview of the solution

- PRET-C: simple synchronous extension to C (using macros).
- ICCFG : intermediate format.
- TFSM : FSM denoted with execution costs.



- **9** PRET-C: simple synchronous extension to C (using macros).
- ICCFG : intermediate format.
- STFSM : FSM denoted with execution costs.
- Model Checking : calculates the WCRT based on a set of TFSMs and a safety property.



- **9** PRET-C: simple synchronous extension to C (using macros).
- ICCFG : intermediate format.
- STFSM : FSM denoted with execution costs.
- Model Checking : calculates the WCRT based on a set of TFSMs and a safety property.



Problem and motivation

main

```
loop each tick{
        ReadInptuts();
        ReactiveFunction();
        EmitOutputs();
    }
     Tick 0
                 Tick 1
                                       Tick 3
                                                        Tick 3
                            Tick 2
                                                 Emit
                                                 utput
                                                                   utput
                             -
                                                                     Time
                                                         Computation
    Computation
                                        Computation
                 Computation
                           Computation
                             time
                                         time
EOT
            EOT
                        EOT
                               EOT
                                                EOT
```

How to determine the "tight/optimal" worst case tick length of the Reactive Function?

→ < ∃ →</p>

Current approaches

æ

イロト イ団ト イヨト イヨト

Current approaches

 MaxPlus [Boldt et al., SLA++P'07]: CKAG based intermediate format of reactive processor KEP that is analysed. Close to 40% overestimation.

Current approaches

- MaxPlus [Boldt et al., SLA++P'07]: CKAG based intermediate format of reactive processor KEP that is analysed. Close to 40% overestimation.
- WCRT algebra [Mendler et al., DATE'09]: max-plus algebra with support for infeasible path pruning.

Current approaches

- MaxPlus [Boldt et al., SLA++P'07]: CKAG based intermediate format of reactive processor KEP that is analysed. Close to 40% overestimation.
- WCRT algebra [Mendler et al., DATE'09]: max-plus algebra with support for infeasible path pruning.
- Timed KS [Logothetis et al., RTSS'03]: Synchronous program compiled into a timed Kripke structure.

Current approaches

- MaxPlus [Boldt et al., SLA++P'07]: CKAG based intermediate format of reactive processor KEP that is analysed. Close to 40% overestimation.
- WCRT algebra [Mendler et al., DATE'09]: max-plus algebra with support for infeasible path pruning.
- Timed KS [Logothetis et al., RTSS'03]: Synchronous program compiled into a timed Kripke structure.
- ILP formulation [Ju et al., CODES+ISSS'08]: Esterel program mapped to C using CEC and then the C code is analysed.

Current approaches

- MaxPlus [Boldt et al., SLA++P'07]: CKAG based intermediate format of reactive processor KEP that is analysed. Close to 40% overestimation.
- WCRT algebra [Mendler et al., DATE'09]: max-plus algebra with support for infeasible path pruning.
- Timed KS [Logothetis et al., RTSS'03]: Synchronous program compiled into a timed Kripke structure.
- ILP formulation [Ju et al., CODES+ISSS'08]: Esterel program mapped to C using CEC and then the C code is analysed.

Limitation

Most approaches ignore state-dependencies while determining infeasible paths. A model checking based formulation may be used to compute the reachable state-space and this reachability analysis automatically prunes unreachable paths.





















How?

Tight analysis depends on: a) Data/Control dependency.

b) Tick/State alignment across threads.








Tighter

Tighter analysis by taking state and data dependencies into account. We can also keep track of our variables in every tick for further infeasible path pruning.

Tighter

Tighter analysis by taking state and data dependencies into account. We can also keep track of our variables in every tick for further infeasible path pruning.

Model Checking

Synchronous C programs may be represented as a set of concurrent FSMs with transition guards that represent the *execution cost*. This can be exploited by a model checker to determine the reachable state-space and the maximum tick length.











c) By considering data dependencies , tick alignment and by tracking the value of the variables.

Roop (University of Auckland)

PRET



c) By considering data dependencies , tick alignment and by tracking the value of the variables.

Roop (University of Auckland)

PRET



c) By considering data dependencies , tick alignment and by tracking the value of the variables.

Roop (University of Auckland)

PRET



c) By considering data dependencies , tick alignment and by tracking the value of the variables.

Roop (University of Auckland)

PRET



c) By considering data dependencies , tick alignment and by tracking the value of the variables.

Roop (University of Auckland)

PRET



c) By considering data dependencies , tick alignment and by tracking the value of the variables.

Roop (University of Auckland)

PRET

Roop (University of Auckland)

æ

イロト イ団ト イヨト イヨト

• MaxPlus: 28

æ

イロト イポト イヨト イヨト

- MaxPlus: 28
- MaxPlus + Data/Control : 25

æ

-∢∃>

- I A P

- MaxPlus: 28
- MaxPlus + Data/Control : 25
- MaxPlus + Tick alignment: 21

- MaxPlus: 28
- MaxPlus + Data/Control : 25
- MaxPlus + Tick alignment: 21
- MaxPlus + Data/Control + Tick alignment: 21

- MaxPlus: 28
- MaxPlus + Data/Control : 25
- MaxPlus + Tick alignment: 21
- MaxPlus + Data/Control + Tick alignment: 21
- Data/Control + Tick alignment + Track Variable: 20

Model checking formulation

- We generate an input format of the model checker UPPAAL called timed automata (TA). Our timed automaton do not use any clocks and instead have only bounded integers.
- We compose all TAs in parallel and introduce a barrier TA to emulate the synchronous composition.
- We then evaluate a set of queries of the form $AG(gtick \Rightarrow x \le val)$.
- We perform binary search in the interval [*WCRT_{min}*, *WCRT_{max}*] by successively trying different *val* until the tight value is found.

Layout

Introduction

- Synchronous Approach at Auckland
- Problems and motivations

Related Work

- The Berkeley-Columbia Approach
- Reactive Processors

The Auckland Approach

- PRET-C based predictable programming
- The Auckland PRET Architecture
- From logical to physical time
- Results

PRET-C execution:

- Hardware support (ARPRET).
- We have also developed a software model
 - (CEC-like linked-list based scheduler).

▶ ∢ ∃ ▶

PRET-C execution:

- Hardware support (ARPRET).
- We have also developed a software model
 - (CEC-like linked-list based scheduler).

- We have compared with other light-weight C extensions such as:
 - SyncCharts in C.
 - Protothreads.
 - Esterel.

Hardware vs Software

	Hardware			Soft	ware	Gain%		
Example	А	W	U	A	W	А	W	
ABRO	29	58	64	36	94	19.45	38.29	
Channel Protocol	57	88	90	91	122	37.36	27.86	
Reactor Control	64	82	86	98	114	34.69	28.07	
Producer Consumer	42	50	53	43	62	2.32	19.35	
Smokers	224	409	413	328	412	31.70	0.73	
Robot Sonar	73	92	96	130	175	43.85	47.43	
Average						28.23	26.96	

• On average, the throughput on the hardware is about 26% greater than the software implementation.

Quantitative Comparison

	PRET-C		SC		PT		Esterel		AC Gain%		WC Gain%			
Example	A	W	A	W	A	W	А	W	SC	РΤ	ES	SC	PΤ	ES
ABRO	36	94	261	493	53	138	78	109	86	32	62	80	31	13
Channel Protocol	91	122	684	757	139	162	232	313	86	34	75	83	24	61
Reactor Control	98	114	444	520	93	106	112	144	77	-5	42	78	-7	20
Producer Consumer	43	62	355	422	74	86	408	417	87	41	89	85	27	85
Smokers	328	412	589	671	268	520	552	1063	44	-22	59	38	20	61
Robot Sonar	130	175	720	770	194	236	408	417	81	32	82	77	25	58
Average									77	18	68	74	20	50

Quantitative Comparison of the PRET-C software approach with SyncCharts in C (SC), Protothreads and Esterel.

• Memory usage of PRET-C is better than (2.5%)Prothreads and (26%) Esterel, while slightly worse (3.7%) than SC.

Execution Time

Example	WCRT	WCRT	Gain
	(Model Checker)	(Actual Execution)	%
ABRO	89	87	97.75
Channel Protocol	152	149	98.03
Reactor Control	118	114	96.61
Producer-Consumer	92	88	95.65
Smokers	449	430	95.77
Robot Sonar	365	339	97.40
Average			96.87

• On an average, the actual value is approximately 96% of the value obtained from UPPAAL.

æ

▶ ∢ ∃ ▶

NUS approach for timing analysis of Esterel [Suhendra et al., DAC'06]

Contributions

- Timing analysis over Esterel language.
- Information from intermediate representations of the compiler (CEC) are used to analyse data flow.
- Explicit Tick Transition Automaton is extracted to capture tick alignment between synchronous threads.
- Models Caches, Pipelines and other architectural features for tighter analysis.
- Timing analysis over multi-core platforms.
- Formulation is based on Integer linear Programming (ILP).
- Currently can only handle simple data flow analysis (SDFA) for pruning infeasible paths.

(B)



- Timing analysis of Esterel.
- ILP based solution.

• Timing analysis of PRET-C.

< □ > < □ > < □ > < □ > < □ > < □ >

MC based solution.

э

code snippet

```
## analysed by NUS
x=3; if(x){..} //simple conflicting pair
```

```
## not analysed by NUS
x=3; z=3; if(z){..} // not a pair
x=3; z=3; if(x<=z){..} // RHS is not a constant
x=3; EOT; if(x){..} //data across ticks</pre>
```

- NUS approach can only handle conflicting pairs. Were simple set and test pairs are analysed. [NUS'DAC06]
- While testing RHS of the expression has to be a constant.
- Cannot analyse data flow across ticks..
- Always, user guided loop bounds.

< ∃ > <

Comparison with NUS approach. [Andalam et al., DATE'11]

NUS's ILP approach can handle simple data flow analysis. Our MC approach can handle more expressive data flow analysis.

Туре	Infeasible paths	NUS	UoA
(1)	Set and Test	Yes	Yes
(2)	Encoding tick transition	Yes	Yes
(3)	Loops with fixed bounds	Yes	Yes
(4)	Set and Test	No	Yes
	with expressive data-flow analysis		
(5)	Encoding tick transition	No	Yes
	with expressive data-flow analysis		
(6)	Loops bounds requiring	No	Yes
	expressive data-flow analysis		

Results

Example	LOC	Observed	(MAX+)	(NUS)	(UoA)	NUS/
		WCRT	Est	est.	est.	UoA
			WCRT	WCRT	WCRT	
				(1)(2)(3)	(1)to(6)	
Synchronizer	455	238	608	422	268	1.57
ProducerConsumer	567	259	808	523	294	1.78
Smokers	648	437	1309	903	521	1.73
Channel Protocol	727	644	1426	897	685	1.31
Robot Sonar	1081	764	2028	1688	858	1.97
Synthetic1	1569	898	3593	2127	1022	2.08
Synthetic2	1630	786	3617	1752	942	1.86
Average						1.67

- MicroBlaze is the target platform.
- Entire program fits on on-chip memory.

æ

3 × 4 3 ×

< A

NUS vs UoA

Over Estimation



• On an average, the NUS approach overestimates by 89% while UoA approach only overestimates by about 13%.

Context Sensitive Pruning

Tighness



• Increase in context information prunes infeasible paths. Thus reducing the WCRT overestimation.

Roop (University of Auckland)

Context Sensitive Pruning

States Explored



• Increase in Context information also reduces the number of states explored. This reduces the time to analyse the model.

Conclusions

Contributions

- New synchronous language for predictable programming.
- Thread-safe shared memory access via simple semantics.
- Predictable preemption support.
- Hardware accelerator to improve the worst case behaviour.
- Mapping of logical time to physical time through static analysis.
- Side effect: PRET-C excels both in the worst case and average case execution over other light-weight threading libraries.

Future work

- To explore the trade-offs of scratchpads versus caches.
- Support for parallel execution of PRET-C via new semantics.
- Exploring the link between our research and the Berkeley-Columbia research.