

A Low-Latency, Energy-Efficient L1 Cache Based on a Self-Timed Pipeline

Louis-Charles Trudeau¹ Ghyslain Gagnon¹
François Gagnon¹ Claude Thibeault¹ Thomas Awad²
Doug Morrissey²

¹École de technologie supérieure, Montréal, Canada

²Octasic Inc., Montréal, Canada

21st IEEE International Symposium on Asynchronous
Circuits and Systems, 2015

Plan

Introduction

Problematic

Motivations

Scope of Work

Cache Implementation

Architecture and Organization

Operation

Self-Timed Pipeline Design

Design Guidelines

Pipeline Control

Pipeline Operation

Performance Results

Summary

Research Program

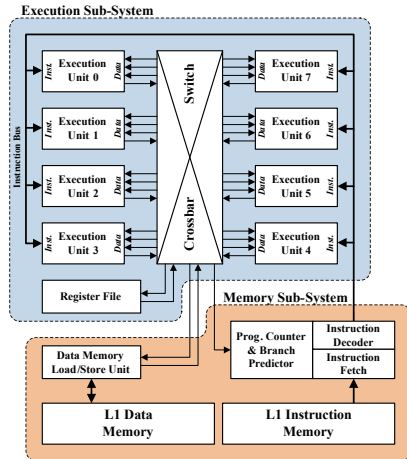
Objective: Adapting Octasic's power-efficient asynchronous architecture in a general purpose processor (ARM v7-A).

Collaborators



Problematic

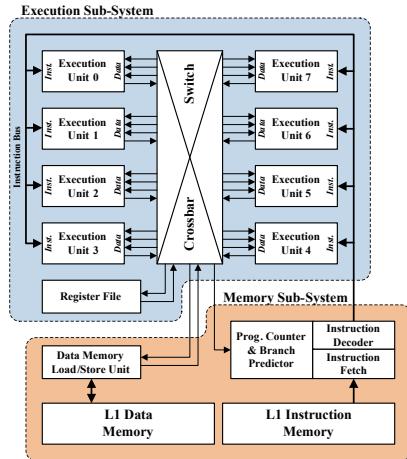
Current architecture separates the **asynchronous CPU** from the **synchronous L1 memory**.



Problematic

Current architecture separates the **asynchronous CPU** from the **synchronous L1 memory**.

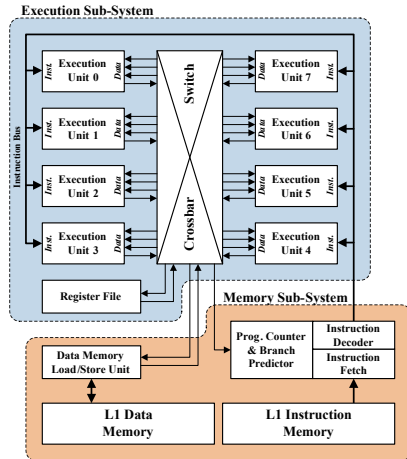
- ▶ **2-cycle synchronization penalty.**



Problematic

Current architecture separates the **asynchronous CPU** from the **synchronous L1 memory**.

- ▶ 2-cycle synchronization penalty.
- ▶ **Energy efficiency is suboptimal.**



Motivations

This work focuses on improving the L1 memory access.

Why Go Asynchronous ?

- ▶ No balanced **clock trees**.
Clocks are point-to-point and skew insensitive.
- ▶ No major **critical path** due to frequency constraints.
Less large/leaky gates.
- ▶ **Less complex** pipeline structure.
Only neighboring stages are connected together.

Scope of Work

Design an asynchronous cache based on a self-timed pipeline.

Objectives

1. Mitigate CPU \leftrightarrow L1 memory access **latency**.
2. Reduce the average memory **access time**.
3. Improve the cache **energy** efficiency.
4. Push the **synchronization barrier** at the L2 memory.

Plan

Introduction

Problematic

Motivations

Scope of Work

Cache Implementation

Architecture and Organization
Operation

Self-Timed Pipeline Design

Design Guidelines

Pipeline Control

Pipeline Operation

Performance Results

Summary

L1 Instruction Cache Design

Dual-fetch, 32kB, 4-way set-associative phased-cache.

Synchronous Cache

- ▶ 5-stage pipeline (hit).
- ▶ Pipeline stall on miss.
- ▶ 2-cycle pipeline reinjection following cache fill.

Asynchronous Cache

- ▶ 4-stage pipeline (hit)
- ▶ Single stage stall on miss.
- ▶ Resource arbitration for concurrent cache fill.

Integration in ARM-like processor

⇒ **Dhrystone** & **Coremark** (armcc compiled).

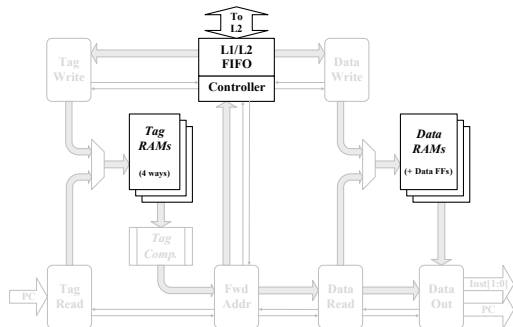
Cache Pipeline

Shared Resources

- ▶ **Tag Memory**
- ▶ **Data Memory**
- ▶ **(L2 Memory)**

Tasks Partitioning

- ▶ 6 pipeline stages.
- ▶ Two-phase handshake protocol.



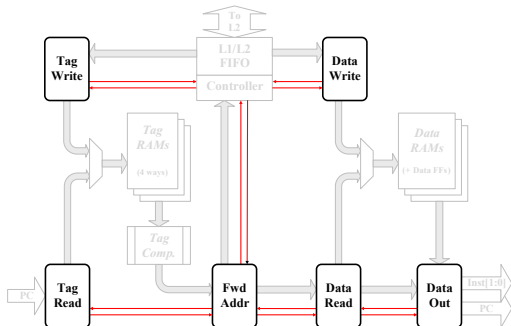
Cache Pipeline

Shared Resources

- ▶ Tag Memory
- ▶ Data Memory
- ▶ (*L2 Memory*)

Tasks Partitioning

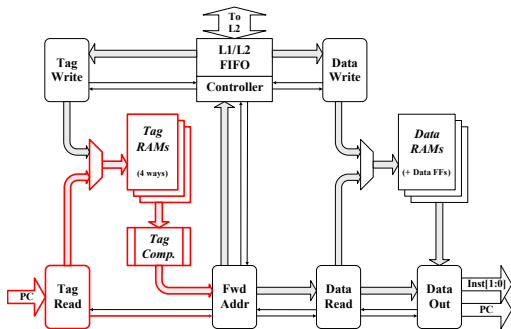
- ▶ **6 pipeline stages.**
- ▶ Two-phase **handshake** protocol.



Cache Operation

Pipeline Stages

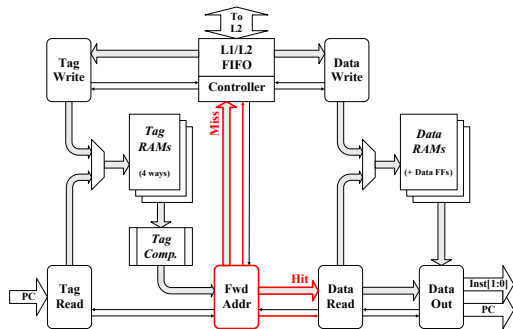
- ▶ Tag Read
- ▶ Forward Address
- ▶ Tag & Data Write
- ▶ Data Read
- ▶ Data Output



Cache Operation

Pipeline Stages


- ▶ Tag Read
- ▶ **Forward Address**
- ▶ Tag & Data Write
- ▶ Data Read
- ▶ Data Output



Cache Operation

Pipeline Stages

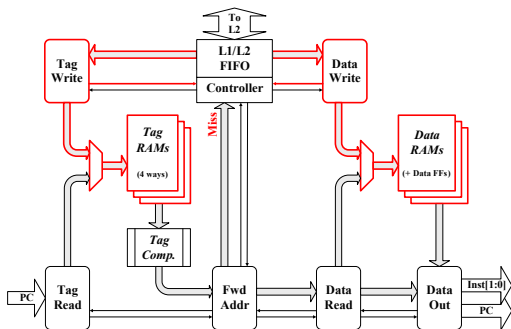
- ▶ Tag Read
- ▶ Forward Address
- ▶ Tag & Data Write
- ▶ Data Read
- ▶ Data Output



Cache Operation

Pipeline Stages

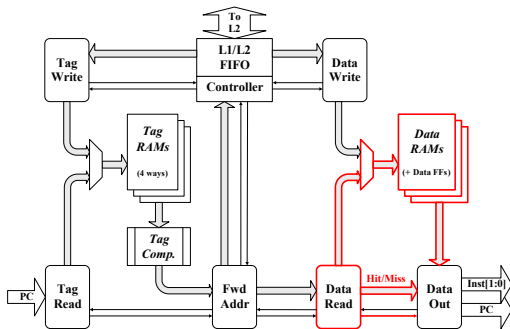
- ▶ Tag Read
- ▶ Forward Address
- ▶ Tag & Data Write
- ▶ Data Read
- ▶ Data Output



Cache Operation

Pipeline Stages

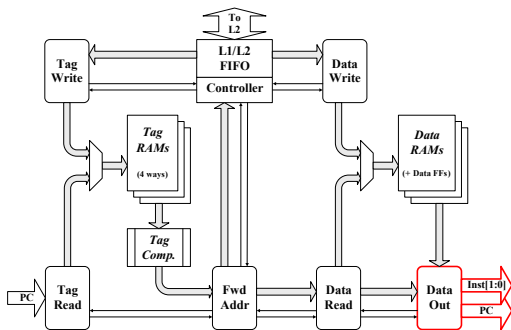
- ▶ Tag Read
- ▶ Forward Address
- ▶ Tag & Data Write
- ▶ Data Read
- ▶ Data Output



Cache Operation

Pipeline Stages

- ▶ Tag Read
- ▶ Forward Address
- ▶ Tag & Data Write
- ▶ Data Read
- ▶ Data Output



Plan

Introduction

Problematic

Motivations

Scope of Work

Cache Implementation

Architecture and Organization

Operation

Self-Timed Pipeline Design

Design Guidelines

Pipeline Control

Pipeline Operation

Performance Results

Summary

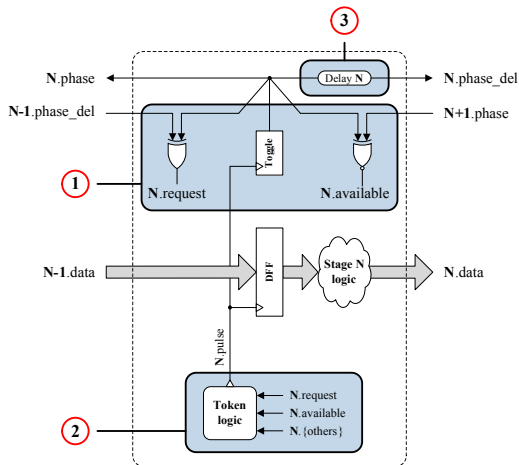
Design Guidelines

Self-timed pipeline design had to follow these **guidelines**:

- ▶ Standard cell libraries.
- ▶ Standard edge-triggered flip-flops.
- ▶ Prioritize High-Voltage Threshold (HVT) cells to limit leakage.


Overview Of A Single Pipeline Stage

1. Click controllers
 - ▶ Two-phase handshake protocol.
2. Token modules
 - ▶ Synchronization;
 - ▶ Pulse generation.
3. Adjustable delays



Click Controllers

- ▶ Based on a two-phase handshake protocol.
- ▶ Stores the stage phase, toggles it upon usage.




Control Signals

- ▶ Request: $N - 1.phase_{del} \oplus N.phase$
- ▶ Available: $\overline{N.phase \oplus N + 1.phase}$

Token Modules

General Idea

To enable a **transaction** with a specific resource, “users” must possess the **resource’s token**.




Token Modules

General Idea

To enable a transaction with a specific resource, “users” must possess the resource’s token.

Operation

1. **Hold** the token



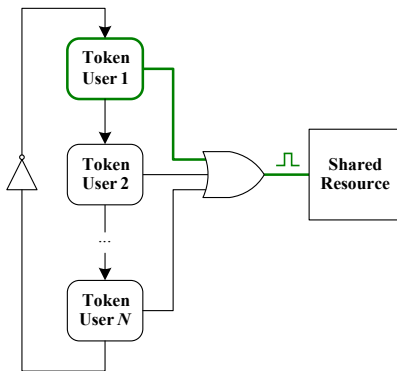
Token Modules

General Idea

To enable a transaction with a specific resource, “users” must possess the resource’s token.

Operation

1. Hold the token
2. **Consume** the token and **access** resource




Token Modules

General Idea

To enable a transaction with a specific resource, “users” must possess the resource’s token.

Operation

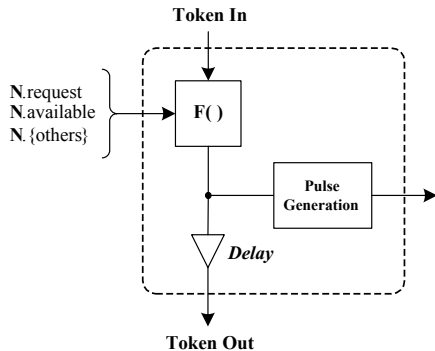
1. Hold the token
2. Consume the token and access resource
3. Pass the token



Token Modules

Internal Structure

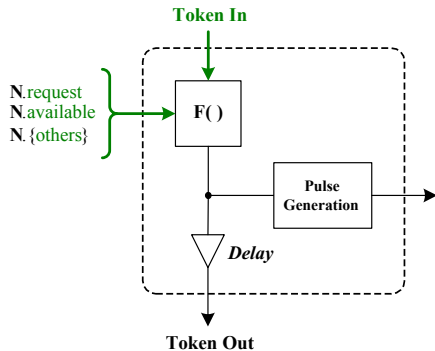
- ▶ Token control function
- ▶ Pulse generation
- ▶ Token ring delay



Token Modules

Operation

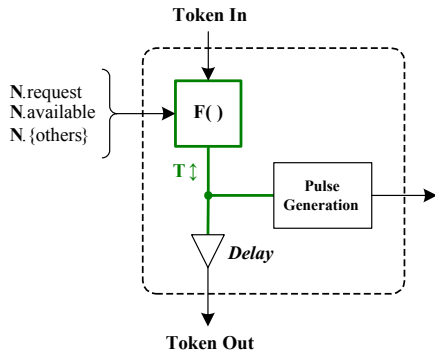
1. Token **conditions** are met.



Token Modules

Operation

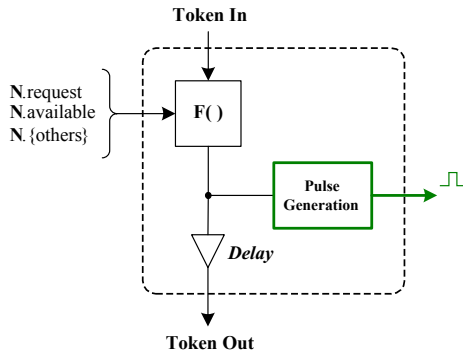
1. Token conditions are met.
2. Token passes through $F()$ thus causing a **transition**.



Token Modules

Operation


1. Token conditions are met.
2. Token passes through $F()$ thus causing a transition.
3. Transition (edge) generates a **pulse** signal.



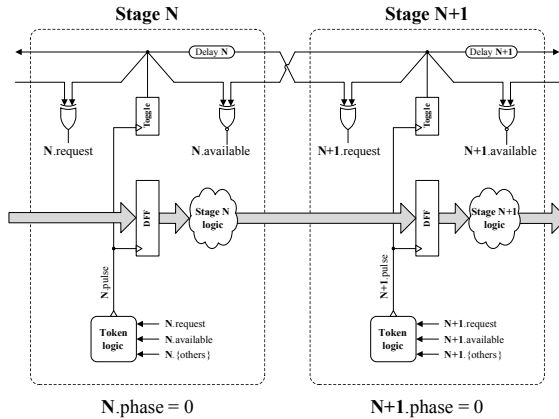
Token Modules

Operation

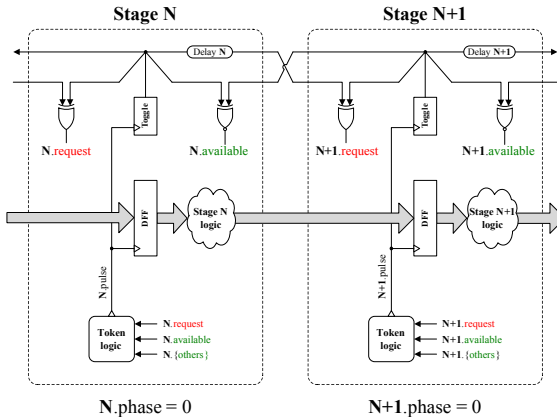
1. Token conditions are met.
2. Token passes through $F()$ thus causing a transition.
3. Transition (edge) generates a pulse signal.
4. Token is **delayed**, then **passed** to next user.



Self-Timed Pipeline Operation

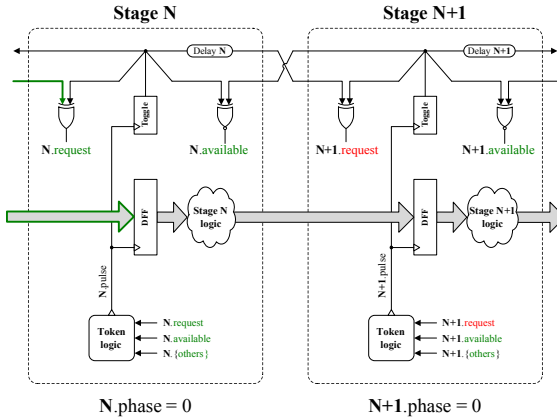


Self-Timed Pipeline Operation



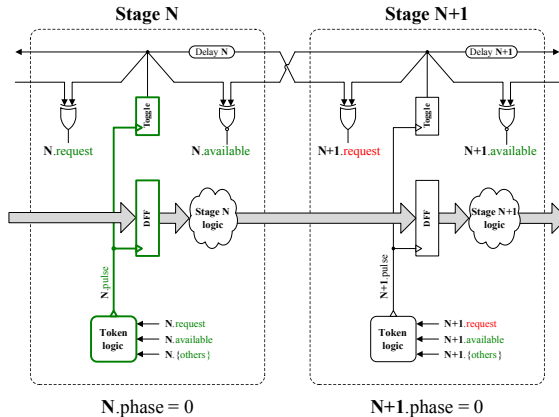
1. Initialization: all stages are **available**.

Self-Timed Pipeline Operation



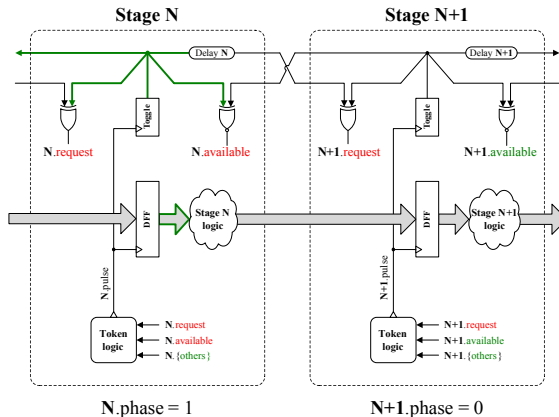
2. Request in: stage N conditions are met.

Self-Timed Pipeline Operation



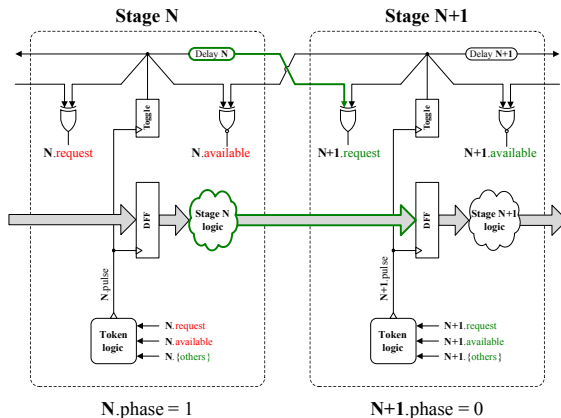
3. Pulse generation: N flops input data & toggles **phase**.

Self-Timed Pipeline Operation



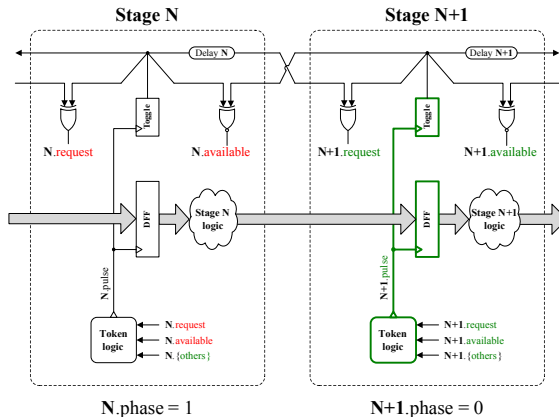
4. Processing: stage N is used, therefore **unavailable**.

Self-Timed Pipeline Operation



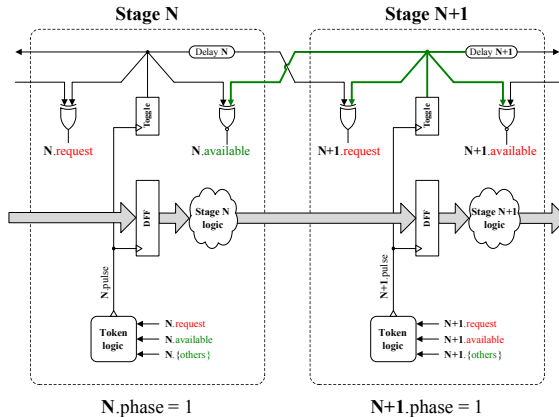
5. Request in: delayed N phase triggers stage N+1 **request**.

Self-Timed Pipeline Operation



6. Pulse generation: N+1 flops input data & toggles **phase**.

Self-Timed Pipeline Operation



7. Stage N is now **available**, N+1 processes data.

Plan

Introduction

Problematic

Motivations

Scope of Work

Cache Implementation

Architecture and Organization

Operation

Self-Timed Pipeline Design

Design Guidelines

Pipeline Control

Pipeline Operation

Performance Results

Summary

Extracting Results

Objective

Use synchronous instruction cache for **baseline results**.


⇒ Keep very similar silicon layout for proper **comparison**.

Performance Metrics


- ▶ Average memory access time
- ▶ Energy consumption
- ▶ Area

Results: Area

Synchronous



Asynchronous



Results: Area

Pipeline	Synchronous		Asynchronous	
	Area (μm^2)	(%)	Area (μm^2)	(%)
<i>Total</i>	11185	100	9900	100
<i>(w/o L2 FIFO)</i>	8455	75,6	7170	72,4
<i>(control)</i>	375	3,35	915	9,24

Table: Pipeline Size Comparison

- ⇒ Pipeline **size** reduced by : **10-15%**.
- ⇒ Pipeline **control** is $\sim 2.5\times$ larger.

Results : Energy

Power Analysis

- ▶ Power estimations based on capacitive switching.
- ▶ Routing estimated from Manhattan distance.
- ▶ L1-L2 interface clock frequency matched.

L1 Cache behavior tests (32kB program)

1. **Miss** : 20k random instructions fetch;
2. **Hit** : 200k, 2M, 20M random instructions fetch.

Results : Energy

	Synchronous	Asynchronous	
Sequence	Energy (<i>nJ</i>)	Energy (<i>nJ</i>)	ΔE (%)
<i>20k inst.</i>	150.6	118.0	21.6
<i>200k inst.</i>	1373.5	1048.3	23.7
<i>2M inst.</i>	13609.1	10357.9	23.9
<i>20M inst.</i>	135917.3	103424.7	24.0

Table: Energy Consumption

⇒ **Energy efficiency** improved by : **> 22%**.

Results : Performance

	Synchronous	Asynchronous	
Sequence	Exec. Time (<i>ms</i>)	Exec. Time (<i>ms</i>)	ΔT (%)
<i>20k inst.</i>	43.1	29.5	31.6
<i>200k inst.</i>	358.1	265.8	25.8
<i>2M inst.</i>	3508.1	2628.3	25.1
<i>20M inst.</i>	34360.0	25769.9	25.0

Table: Average Memory Access Time

⇒ **Access time** reduced by : > 25%.

⇒ **Throughput** at L1-L2 interface : > 40%.

Results : What needs to be addressed

Future Work

- ▶ Further pipeline cache to reach > 1 GHz equivalent.
- ▶ Design L1 data cache based on self-timed pipeline.
- ▶ Integrate asynchronous L1 caches in Octasic's next-generation processors.

Plan

Introduction

Problematic

Motivations

Scope of Work

Cache Implementation

Architecture and Organization

Operation

Self-Timed Pipeline Design

Design Guidelines

Pipeline Control

Pipeline Operation

Performance Results

Summary

Summary

Problematic

⇒ **Asynchronous** CPU accesses **synchronous** L1 memory.

Goals & Results

Design and implement an asynchronous L1 cache:

1. Mitigate CPU \leftrightarrow L1 memory access latency: ✓
2. Improve the cache energy efficiency: > 22%
3. Reduce the average memory access time: > 25%
4. Push the synchronization barrier at the L2 memory: ✓