



Cost-Effective Compiler Directed Memory Prefetching and Bypassing

Daniel Ortega,[†] Eduard Ayguadé[†], Jean-Loup Baer[‡] and Mateo Valero[†]

[†]Departamento de Arquitectura de Computadores,


[‡]Department of Computer Science and Engineering,

Universidad Politécnica de Cataluña – Barcelona

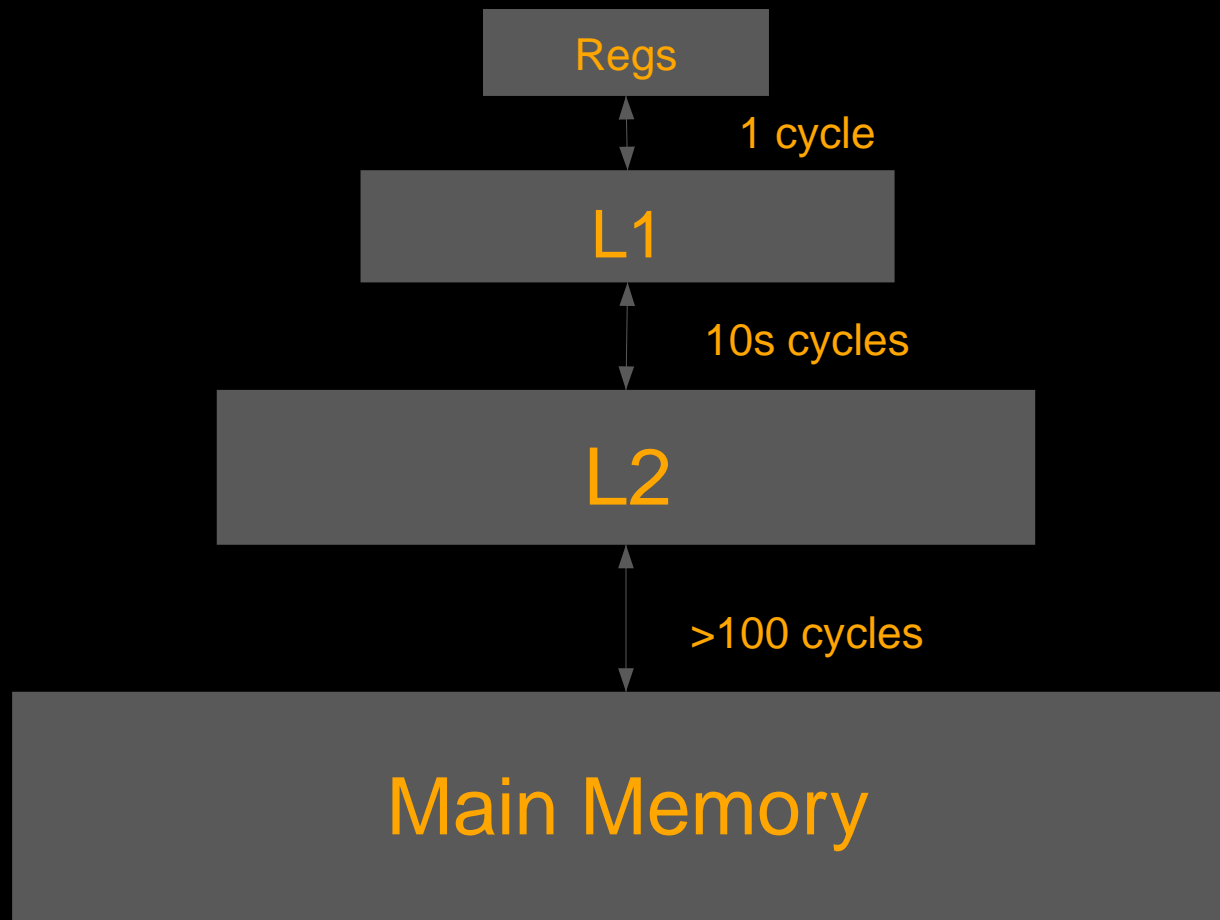
University of Washington – Seattle

{dortega,eduard,mateo}@ac.upc.es

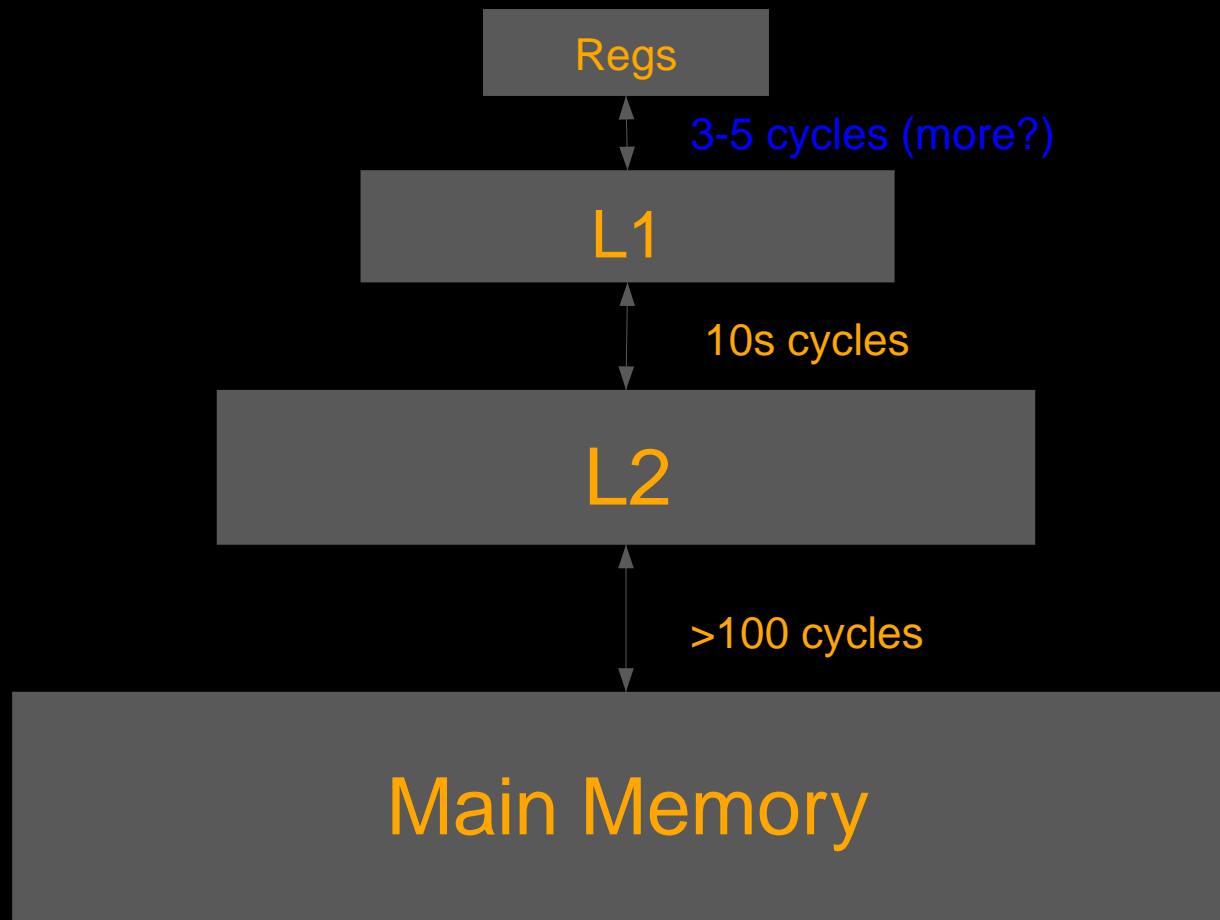
baer@cs.washington.edu



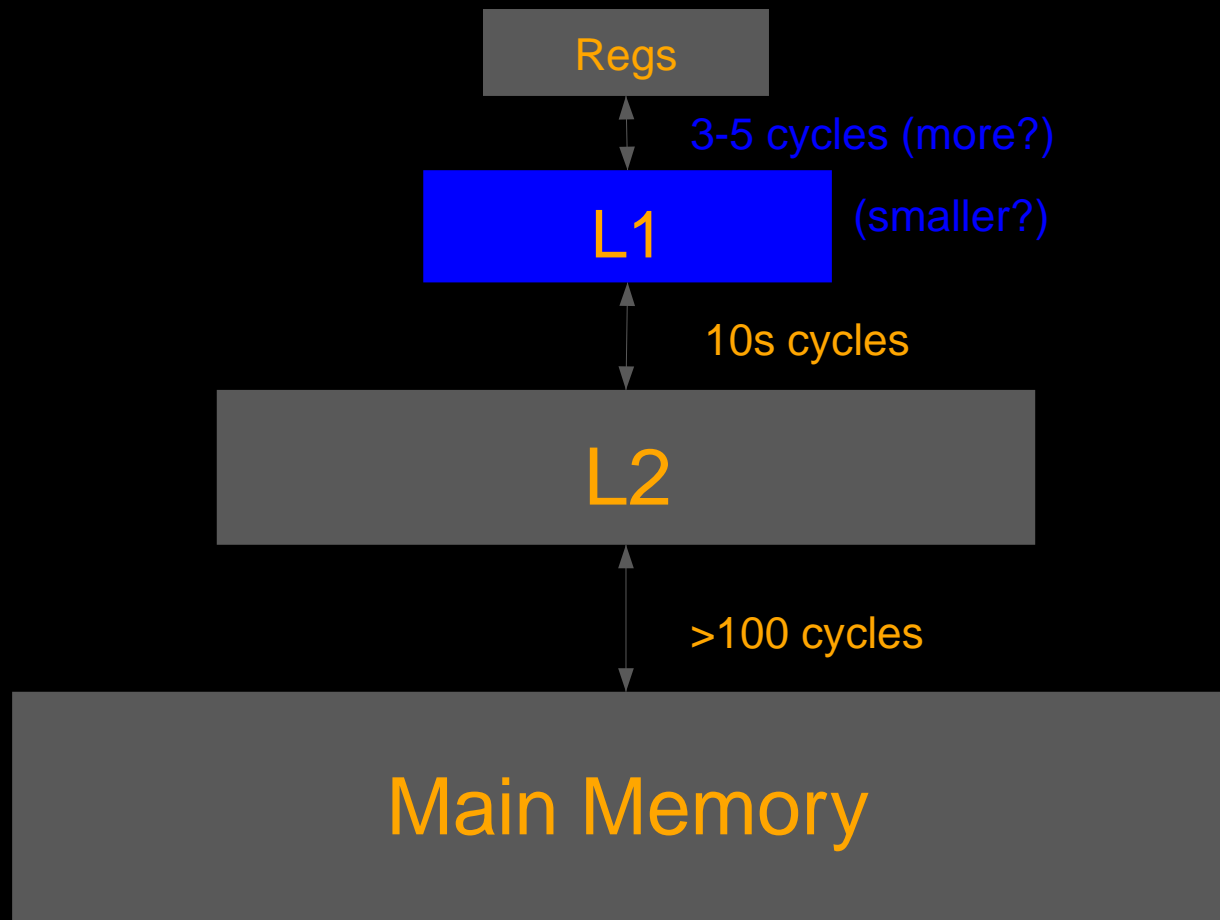
Conventional Cache Hierarchy



Conventional Cache Hierarchy




Conventional Cache Hierarchy




Our approach

- Attack register-L1 gap with memory instruction bypassing
- Use hardware prefetcher for L1-L2 gap
 - ★ Directed by the software
 - ★ thus simple Hardware
 - ★ no recovery needed in case of misprediction




Limit Study




Limit values for a 4-way machine

Index

- Motivation
- Memory Instruction Bypassing
- Compiler Directed Memory Prefetcher
- Comparison with APDP

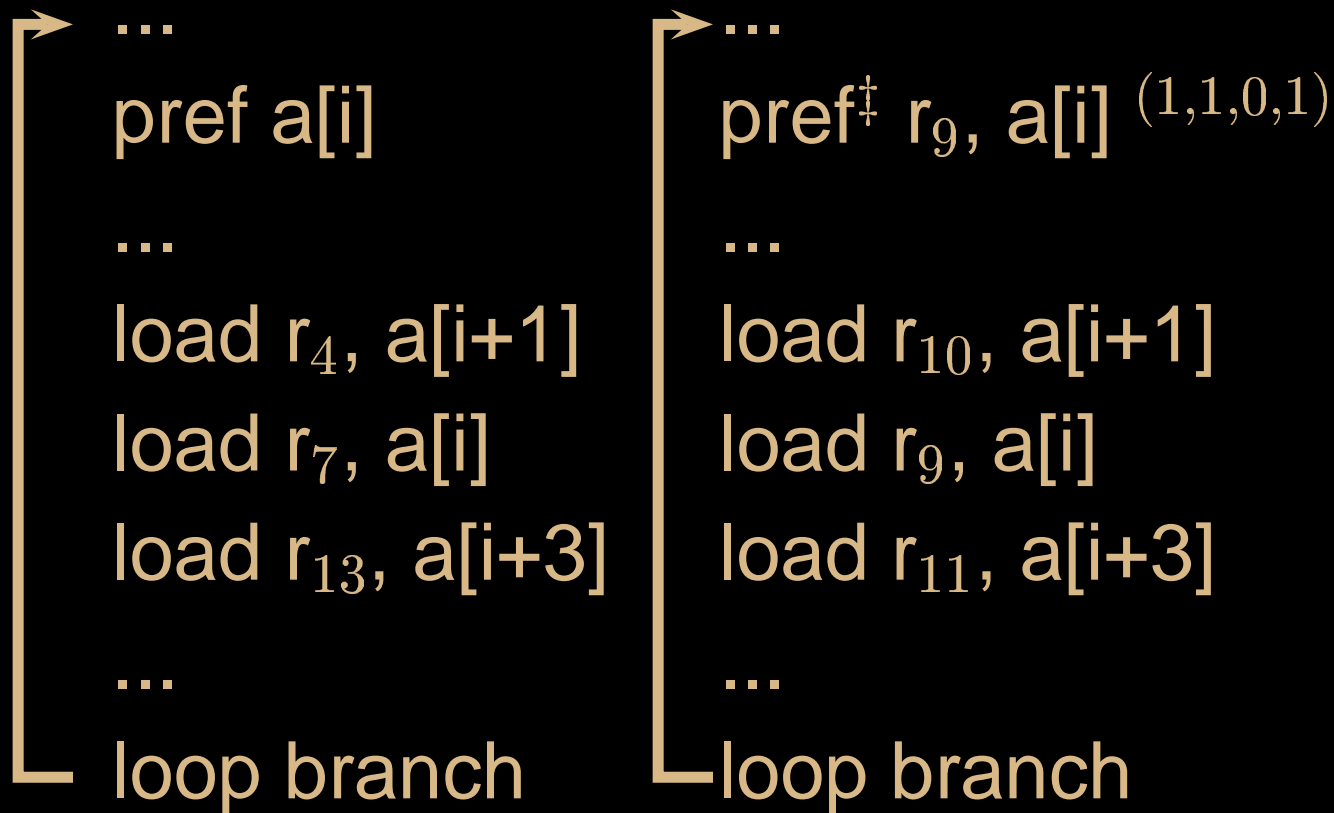


MIB through Renaming I



```
...  
pref a[i]  
...  
load r4, a[i+1]  
load r7, a[i]  
load r13, a[i+3]  
...  
loop branch
```


MIB through Renaming I



MIB through Renaming II




Renaming Table

...	→	...
r ₉	→	f ₁₀
r ₁₀	→	f ₁₄
r ₁₁	→	f ₉
...	→	...

Special Renaming Table

...	→	...
r ₉	→	...
r ₁₀	→	...
r ₁₁	→	...
...	→	...



MIB through Renaming II



...


pref[‡] r₉, a[i] (1,1,0,1)

Renaming Table

...	→	...
r ₉	→	f ₁₀
r ₁₀	→	f ₁₄
r ₁₁	→	f ₉
...	→	...

Special Renaming Table

...	→	...
r ₉	→	f ₂₁
r ₁₀	→	f ₃₂
r ₁₁	→	f ₃₆
...	→	...



MIB through Renaming II



...
pref[‡] r₉, a[i] (1,1,0,1)


...
load r₁₀, a[i+1]

Renaming Table

...	→	...
r ₉	→	f ₁₀
r ₁₀	→	f ₃₂
r ₁₁	→	f ₉
...	→	...

Special Renaming Table

...	→	...
r ₉	→	f ₂₁
r ₁₀	→	...
r ₁₁	→	f ₃₆
...	→	...



MIB through Renaming II



...
pref[‡] r₉, a[i] (1,1,0,1)

...

load r₁₀, a[i+1]


load r₉, a[i]

Renaming Table

...	→	...
r ₉	→	f ₂₁
r ₁₀	→	f ₃₂
r ₁₁	→	f ₉
...	→	...

Special Renaming Table

...	→	...
r ₉	→	...
r ₁₀	→	...
r ₁₁	→	f ₃₆
...	→	...



MIB through Renaming II



...
pref[‡] r₉, a[i] (1,1,0,1)

...

load r₁₀, a[i+1]

load r₉, a[i]


load r₁₁, a[i+3]

Renaming Table

...	→	...
r ₉	→	f ₂₁
r ₁₀	→	f ₃₂
r ₁₁	→	f ₃₆
...	→	...

Special Renaming Table

...	→	...
r ₉	→	...
r ₁₀	→	...
r ₁₁	→	...
...	→	...



MIB through Renaming II



...
pref[‡] r₉, a[i] (1,1,0,1)

...

load r₁₀, a[i+1]

load r₉, a[i]

load r₁₁, a[i+3]

...


loop branch

Renaming Table

...	→	...
r ₉	→	f ₂₁
r ₁₀	→	f ₃₂
r ₁₁	→	f ₃₆
...	→	...


Special Renaming Table

...	→	...
r ₉	→	...
r ₁₀	→	...
r ₁₁	→	...
...	→	...



Index

- Motivation
- Memory Instruction Bypassing
- **Compiler Directed Memory Prefetcher**
- Comparison with APDP




Decoupled Prefetcher

- Compiler inserts prefetching operations
- Instructions bring data closer to the processor
- Memory instruction bypassing takes care of bringing data to the register file from L1
- Compiler also instructs prefetching hardware to prefetch ahead (to L1)
- No.s of prefetches are minimised by compiler control

Prefetching Mechanism



pctag	last @	type
	...	



Prefetching Mechanism

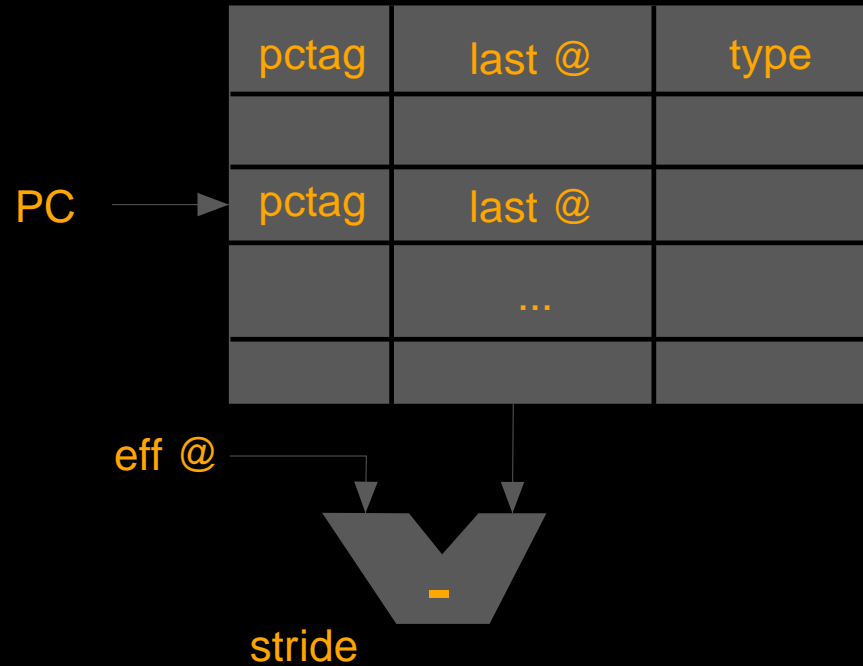
- Decode phase

PC →

pctag	last @	type
pctag	last @	
	...	

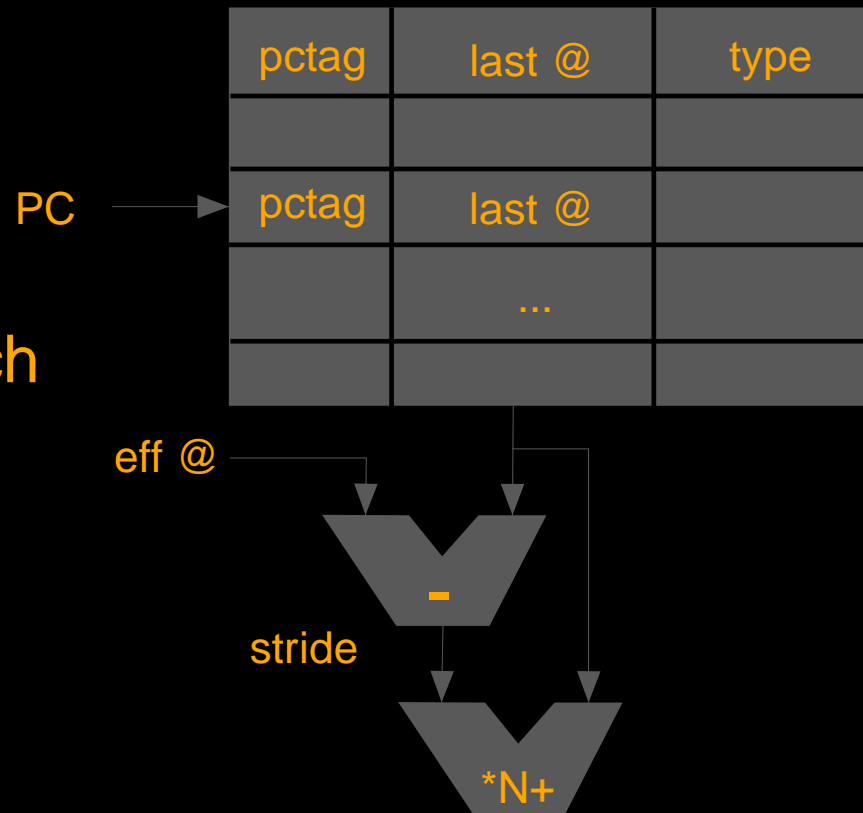
Prefetching Mechanism

- Decode phase
- Address calc.



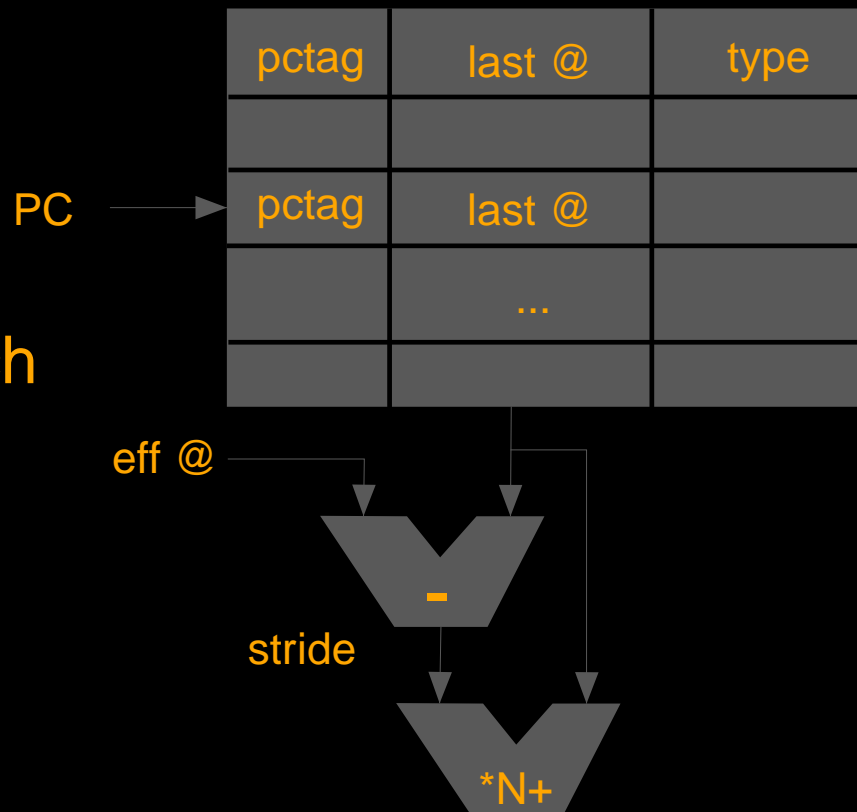
Prefetching Mechanism

- Decode phase
- Address calc.
- Generate new prefetch



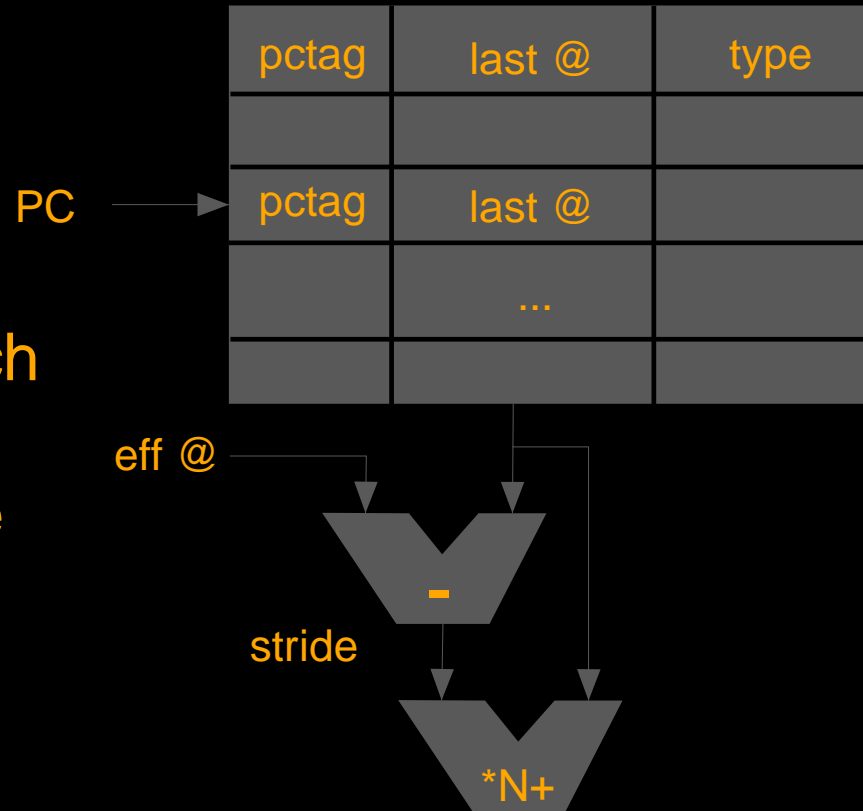
Prefetching Mechanism

- Decode phase
- Address calc.
- Generate new prefetch
 - ★ N depends on type




Prefetching Mechanism

- Decode phase
- Address calc.
- Generate new prefetch
 - ★ N depends on type
 - ★ wait for free port




Performance Results



Effect of number of streams in a 4-way machine


Performance Results II



Effect of lookahead policy in a 4-way machine with 32 entries

Index

- Motivation
- Memory Instruction Bypassing
- Compiler Directed Memory Prefetcher
- Comparison with APDP




Comparison with APDP[‡]



- Address Prediction for Data Prefetching
 - ★ predicts addresses of memory operations
 - ★ makes prediction available as soon as load arrives to decoding
 - ★ needs recovery mechanism in case of misprediction
 - ★ dynamic mechanism


[‡]J. González and A. González, ICS 97



Prefetching Mechanism



address	stride	conf.	val	v
		⋮		



Prefetching Mechanism

- Decode phase

address	stride	conf.	val	v
PC →		⋮		

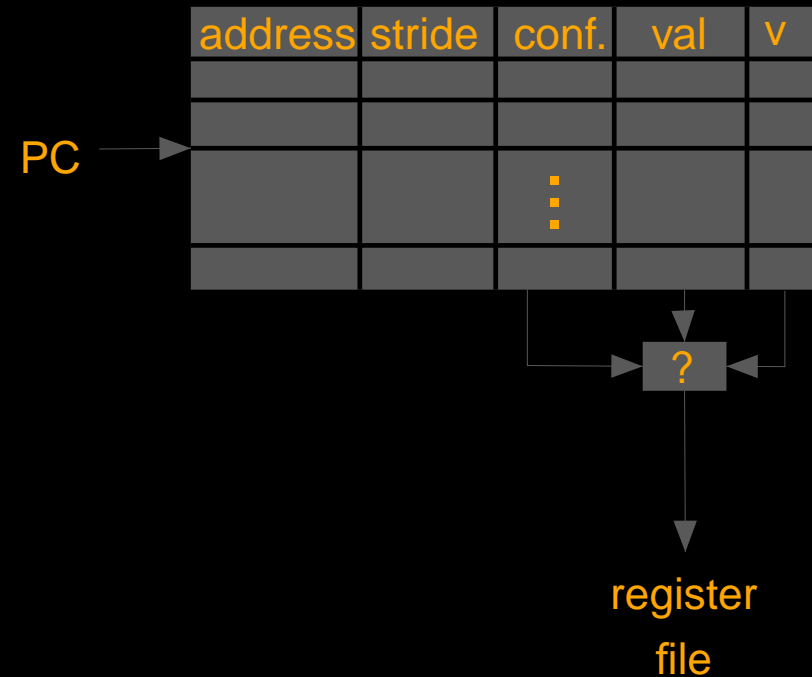
Prefetching Mechanism

- Decode phase
- Is value correct?



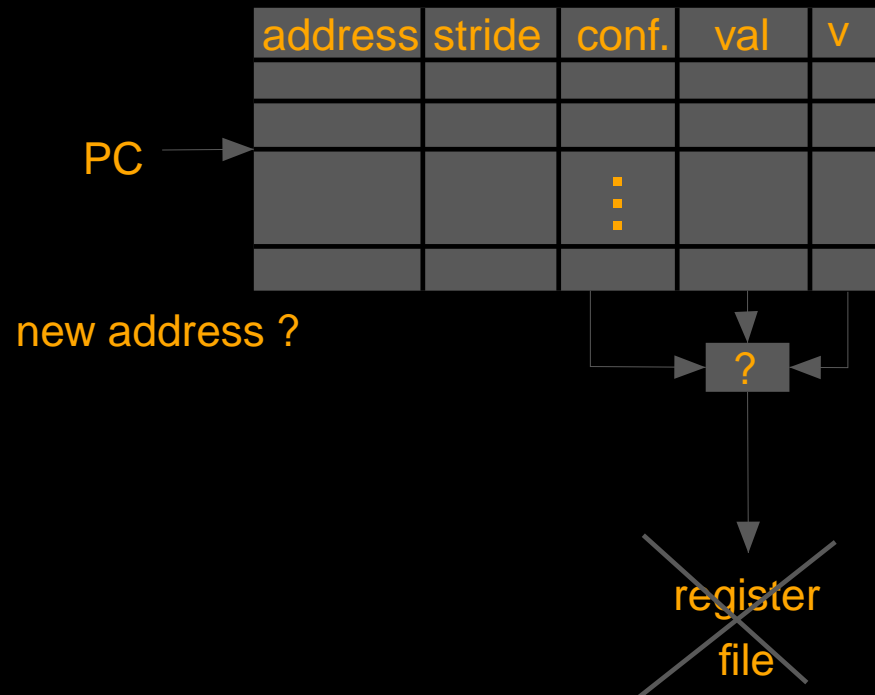
Prefetching Mechanism

- Decode phase
- Is value correct?
- Yes → bypass



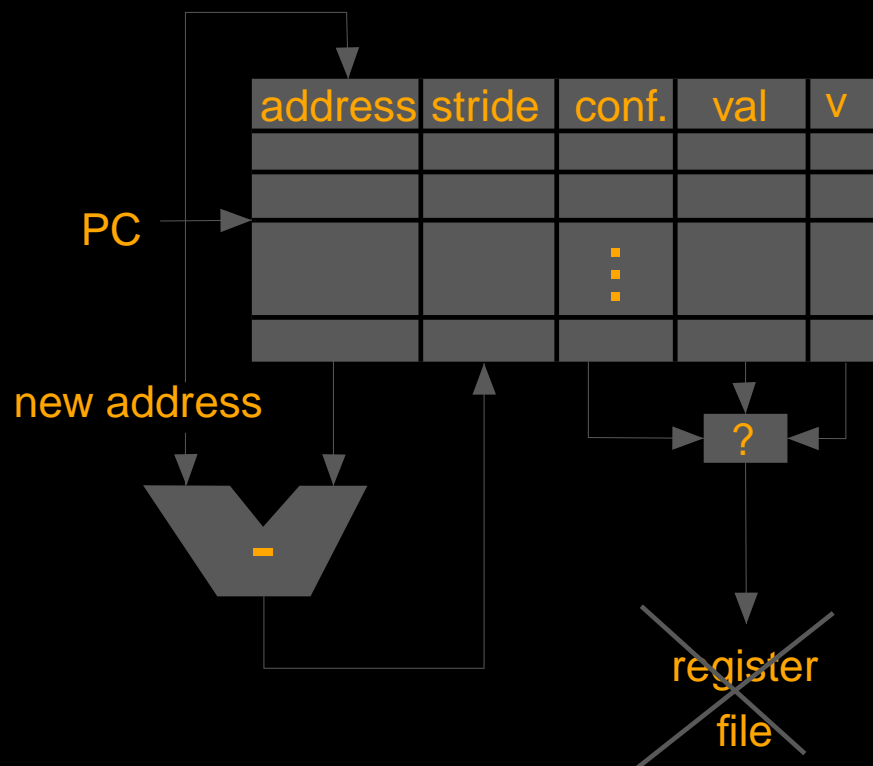
Prefetching Mechanism

- Decode phase
- Is value correct?
- Yes → bypass
- Address calc.



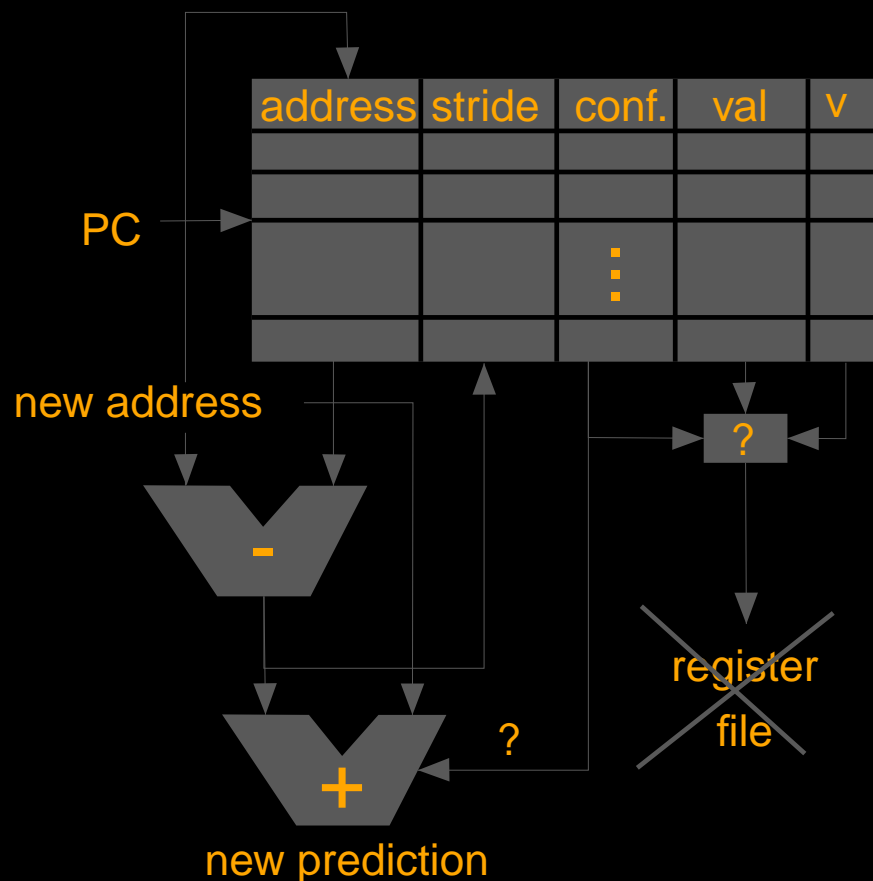
Prefetching Mechanism

- Decode phase
- Is value correct?
- Yes → bypass
- Address calc.
- Update table




Prefetching Mechanism

- Decode phase
- Is value correct?
- Yes → bypass
- Address calc.
- Update table
- Generate prefetch?




Comparison Results



Effect of only prefetching in APDP (4-way and 2 ports)

Comparison Results




4 way machine comparison against APDP

Memory traffic



	applu	apsi	hydro2d	swim	tomcatv
LoadStore	1.00	1.00	1.00	1.00	1.00
APDP sel.	1.14	1.14	1.24	1.21	1.21
APDP sq.	1.14	1.15	1.24	1.21	1.22
Bypassing	0.90	0.74	0.82	0.72	0.89
Proposal	0.98	0.85	0.95	0.83	0.94





Any questions?



Thank you
dortega@ac.upc.es

