

You Can't Always Get What You Want

If you try sometimes you might find you get what you need – *but maybe not*. That's the problem with **caches**. As modern processors have sprouted increasingly complex memory hierarchies, they still have not solved the fundamental problem of ensuring that instructions and data are always ready when they are needed.

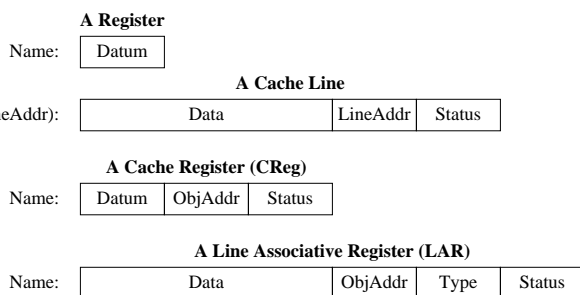
A value kept in a **register** is there when you need it. However, there are lots of memory objects conventional machines do not keep in registers. Instructions are not kept in registers, although there really isn't a good reason why not. There is a widely held belief that registers don't help with "spatial locality" – but that is true only if we assume that a register is just one object wide.

Of course, ambiguously aliased data cannot be kept in registers. For example, if a code refers to `a[i]` and `a[j]` a number of times, there are three possible circumstances:

Alias Analysis	Register Assignment
compiler knows <code>i==j</code>	share one register
compiler knows <code>i!=j</code>	two separate registers
compiler doesn't know	which of the above?

In fact, ambiguously aliased object references are the only reason a processor needs cache: without cache, the store and reload would become slow references to main memory. Even so, it is possible that the cache would be ineffective, because references to other objects might have had addresses hashing to the same cache line slot, thus evicting the desired object from cache. However, everything needed to resolve the ambiguity was in the processor – why not simply modify the register file to automatically update aliased objects in other registers?

CRegs & LARs. Our SC'88 paper, *CRegs: A New Kind of Memory for Referencing Arrays and Pointers*, introduced CACHE-REGISTERS to solve the ambiguous alias flushing problem by *associatively updating* registers whose address fields match. However, data CREGs don't make use of spatial locality – our work on LARs (LINE ASSOCIATIVE REGISTERS) over the past few years combines CREGs with SWAR (SIMD WITHIN A REGISTER) to take full advantage of spatial locality. The basic hardware cell structures are:



An example using Data LARs. In the following C code, suppose `i`, `j` and `k` are ambiguously aliased. Because each LAR records the current object position within its data field, the LAR code reduces memory loads, fewer if any of the lines have

the same base addresses.. The architecture also provides implicit lazy stores. Each LAR is type tagged, and type information need not be encoded within arithmetic instructions. The LAR code as written is scalar; however, if the data are properly aligned, replacing ADDS and ANDS with ADDV and ANDV would perform the SWAR parallel operations on a "line" of data at a time.

C Code	Conventional RISC	LARs
<code>nasty(int*i,</code>	<code>LW \$t1,j(0)</code>	<code>LOADSW d1,d0,d0,j</code>
<code>int*j, int*k)</code>	<code>LW \$t2,0(\$t1)</code>	<code>LOADSW d2,d0,d1,0</code>
<code>{</code>	<code>LW \$t3,k(0)</code>	<code>LOADSW d3,d0,d0,k</code>
<code>*i=*j**k;</code>	<code>LW \$t4,0(\$t3)</code>	<code>LOADSW d4,d0,d3,0</code>
<code>*k=*i&**k;</code>	<code>LW \$t5,i(0)</code>	<code>LOADSW d5,d0,d0,i</code>
<code>}</code>	<code>ADD \$t6,\$t2,\$t4</code>	<code>LOADd d6,d0,d5,0</code>
	<code>SW \$t6,0(\$t5)</code>	<code>ADDS d6,d2,d4</code>
	<code>LW \$t7,0(\$t5)</code>	<code>ANDS d4,d6,d4</code>
	<code>LW \$t8,0(\$t3)</code>	
	<code>AND \$t9,\$t7,\$t8</code>	
	<code>SW \$t9,0(\$t3)</code>	

Instruction LARs. Instruction LARs remove the instruction fetch process from the execution of each instruction, replacing it with separate, explicit, instructions that load of compressed blocks of instructions. If a load requests a block that is already in another instruction LAR, the decoded instruction block is logically copied without any memory activity. Control flow targets are specified using an instruction offset within an instruction LAR, rather than by comparatively lengthy memory addresses. The overall result is a smaller memory footprint, improved utilization of memory bandwidth, and complete freedom from misses during instruction processing.

Status. Krishna Melarkode's 2004 M.S. Thesis, *Line Associative Registers*, first introduced the concepts of LAR as an extension of CREGs. An HDL simulator has been constructed which defines the details of the implementation including data type conversions. An instruction fetch mechanism which includes compression and decompression of instruction blocks has been developed. Work has begun to construct a compiler tool chain for the designed architecture. Watch **Aggregate.Org/LAR** for more information.

This document should be cited as:

```
@techreport{sc09lar,
author={Henry Dietz, Kalyan Ponnala, Nien Yi Lim, Paul
S. Eberhart},
title={You Can't Always Get What You Want},
institution={University of Kentucky},
address={http://aggregate.org/WHITE/sc09lar.pdf},
month={Nov}, year={2009}}
```

