

Título: FILTERING DIRECTORY LOOKUPS IN CMPS

Nombre: Bosque Arbiol, Ana

Universidad: Universidad de Zaragoza

Departamento: Informática e ingeniería de sistemas

Fecha de lectura: 11/11/2011

Mención a doctor europeo: concedido

Programa de doctorado: Ingeniería de sistemas e informática

Dirección:

- > **Director:** VÍCTOR VIÑALS YÚFERA
- > **Director:** JOSÉ MARÍA LLABERIA GRIÑÓ
- > **Director:** PABLO ENRIQUE IBAÑEZ MARÍN

Tribunal:

- > **presidente:** BENJAMÍN SAHELICES FERNÁNDEZ
- > **secretario:** JESÚS ALASTRUEY BENEDÉ
- > **vocal:** SALLY A. MCKEE
- > **vocal:** EMILIO LÓPEZ ZAPATA
- > **vocal:** GIORGOS DIMITRAKOPOULOS

Descriptor:

- > ARQUITECTURA DE ORDENADORES

El fichero de tesis ya ha sido incorporado al sistema

- > <http://zaguan.unizar.es/record/6812>

Localización: UNIVERSIDAD DE ZARAGOZA

Resumen: Nowadays, most computer manufacturers offer chip multiprocessors (CMPs) due to the always increasing chip density. These CMPs have a broad range of characteristics, but all of them support the shared memory programming model. As a result, every CMP implements a coherence protocol to keep local caches coherent.

Coherence protocols consume an important fraction of power to determine which coherence action to perform. Specifically, on CMPs with write-through local caches, a shared cache and a directory-based coherence protocol implemented as a duplicate of local caches tags, we have observed that energy is wasted in the directory due to two main reasons.

Firstly, an important fraction of directory lookups are useless, because the target block is not located in any local

cache. The power consumed by the directory could be reduce by filtering out useless directory lookups.

Secondly, useful directory lookups (there are local copies of the target block) are performed over target blocks that are shared by a small number of processors. The directory power consumption could be reduced by limiting the directory lookups to only the directory entries that have a copy of the block.

Along this thesis we propose two filtering mechanisms. Each of these mechanisms is focused on one of the problems described above: while our first proposal focuses on reducing number of directory lookups performed, our second proposal aims at reducing the associativity of directory lookups. Several implementations of both filtering approaches have been proposed and evaluated, having all of them a very limited hardware complexity. Our results show that the power consumed by the directory can be reduced as much as 30%.