Information Technology

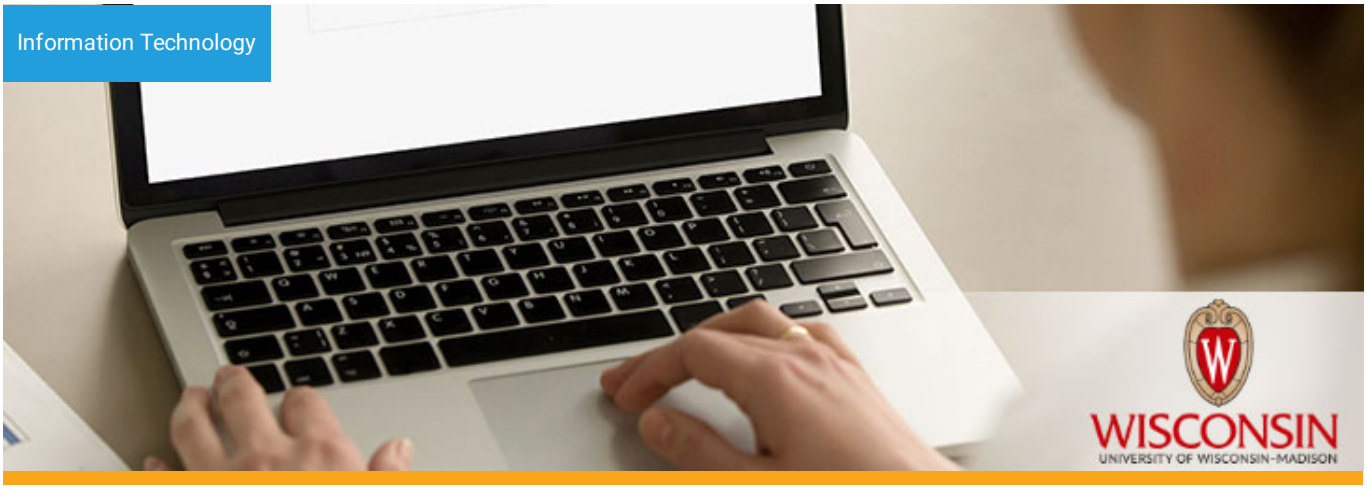

Cache Accessed Using Virtual Addresses; Energy Reduction Significant

View U.S. Patent No. 10,089,240 in PDF format.

WARF: P150177US03

Inventors: Gurindar Sohi, Hongil Yoon

The Wisconsin Alumni Research Foundation (WARF) is seeking commercial partners interested in developing virtual cache architecture that, in experimental results, has shown >90 percent reduction in dynamic energy consumption for virtual to physical address translation via translation lookaside buffers (TLB) compared to current methods.

The improved architecture provides for a memory cache that is accessed not by physical addresses but by virtual addresses directly from running processes. Ambiguities that can result from multiple virtual addresses mapping to the same physical address are handled by dynamically identifying synonym accesses and remapping their virtual addresses to a single key virtual address that is used exclusively for cache access.

Overview

Modern processors and operating systems support virtual memory. Programs running on the system access (read or write) data using virtual addresses that differ from the physical address of the data in memory. The use of virtual addresses greatly improves programmability. Despite its advantages, the use of virtual address space increases the delay in accessing memory by requiring a translation from virtual to physical address space. Normally this translation is done by means of page tables, which may be augmented by a TLB to speed up the process. The TLB is consulted on every memory access when a cache is accessed with physical addresses.

Ideally, one would be able to create virtual caches that, compared to physical caches, have lower access latency and energy consumption by not consulting TLBs on every access. Yet this has not been popular in commercial designs. The crux of the problem is the possibility of synonyms, that is, distinct virtual addresses mapping to the same physical address. Various solutions developed to date have placed large demands on software and offset real energy gains.

The Invention

UW–Madison researchers led by computer architecture expert Gurindar Sohi have recognized that in many important computational tasks, while synonyms are present over the duration of the program, (1) there are very few synonyms mapped to the same physical address, and (2) very few of the accesses to the cache are to data with synonyms during the life of cached data. By exploiting these temporal properties of synonyms, their invention accesses the cache directly with a given virtual address without consulting the TLB. It also efficiently converts the given virtual addresses to a selected key virtual address prior to cache access when they are synonym virtual addresses. It is practical to track synonyms with modest size data structures because of the relative scarcity of synonyms during the life of cached data.

The novel architecture includes a processor, a memory cache, a cache control circuit caching memory data of physical addresses in the memory cache by virtual addresses, and a translation circuit. In addition, the architecture includes a synonym tracking circuit receiving a given virtual address from the processor for access to the cache and: (1) determining if the given virtual address is a synonym, and (2) when the given virtual address is a synonym, accessing the memory cache using the other virtual address as an accessing address for the memory cache.

Applications

- New virtual cache architecture with an added Active Synonym Detection Table (ASDT), Address Remapping Table (ART), Synonym Signature (SS) and associated computing methods

Key Benefits

- Outperforms current methods – dramatically reduced energy consumption
- Accommodates possible aliasing or overlapping
- Permits faster direct access to the cache using virtual memory addresses in the dominant situation where there are no synonyms
- Reduces the delay caused by synonym tracking when there are no synonyms
- Provides a simple and compact method of monitoring synonyms consistent with the expected low number of active synonyms in a short timeframe
- Leverages existing translation circuitry (e.g., the TLB) for the purpose of determining synonyms
- No burdens placed on software to ensure correct operation

Stage of Development

Experimental results based on real world applications (server, database, etc.) show a 90 percent and 86 percent reduction in dynamic energy consumption for instruction-TLB and data-TLB respectively for an example configuration.

Additional Information

For More Information About the Inventors

- Gurindar Sohi

Related Technologies

- WARF reference number P130019US01 describes a cache memory system developed by the researchers to reduce invalidation message traffic.
- For information about parallel processing innovations developed by Prof. Sohi, see WARF reference numbers:
- P100343US01
- P130018US01
- P140373US01

Tech Fields

- Information Technology : Hardware

For current licensing status, please contact Emily Bauer at emily@warf.org or 608-960-9842