

The opinion in support of the decision being entered today was not written for publication in a law journal and is not binding precedent of the Board.

Paper No. 15

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE BOARD OF PATENT APPEALS
AND INTERFERENCES

Ex parte RAVI KUMAR ARIMILLI, JOHN STEVEN DODSON
and JERRY DON LEWIS

Appeal No. 2001-2164
Application No. 09/024,620

ON BRIEF

Before THOMAS, HAIRSTON and KRASS, Administrative Patent Judges.
KRASS, Administrative Patent Judge.

DECISION ON APPEAL

This is a decision on appeal from the final rejection of claims 1-20.

The invention pertains to caches in a data processing system. More particularly, the invention allows multiple caches, i.e., horizontally-oriented caches supporting different processing units, to share a memory value which is not consistent

with system memory. One and only one of these shared-modified cache lines is uniquely identified to enable intervention (i.e., a process whereby a cache can source a requested value in lieu of the system memory sourcing the value) using the uniquely identified cache line. The identification is accomplished by expanding the prior art MESI cache coherency protocol to include a new cache coherency state, the "tagged" state.

This "tagged" state indicates:

1. That the value contained in the "tagged" cache-line is shared, i.e., the same value is also present in the horizontally-oriented cache lines of other processing units.
2. That the value in the "tagged" cache line is in a modified state, i.e., the value is not consistent with system memory, so that if the value is to be supplied to any other processing units, it must come from one of the cache lines sharing the value, and not from system memory.
3. That the specific "tagged" cache in this "tagged" cache coherency state will be responsible for intervening the value, as well as writing the value, to system memory should a flush of the value from that particular cache occur.

The instant claims are directed to a particular implementation of the T-MESI cache coherency protocol in which different priorities are associated with cache coherency responses, including the tagged intervention response.

Appeal No. 2001-2164
Application No. 09/024,620

Representative independent claim 1 is reproduced as follows:

1. A method of improving memory latency associated with a read-type operation in a multi-processor computer system having a plurality of processing units, each processing unit further having an associated cache, the method comprising the steps of:

a requesting processing unit issuing a message to an interconnect of the computer system, indicating that the requesting processing unit desires to read a value from an address of a memory device of the computer system;

each cache snooping the interconnect to detect the message;

each cache transmitting a response to the message, wherein a tagged intervention response is transmitted from a responding cache to indicate (i) that the responding cache contains a modified value corresponding to the address of the memory device that has not been written to the memory device, (ii) that the responding cache shares the modified value with at least one other cache of an adjacent processing unit, and (iii) that the responding cache will exclusively source the value;

associating a priority with each response from each cache;

detecting each response and its associated priority; and

forwarding a response with a highest priority to the requesting processing unit.

The examiner relies on the following references:

Wilson, Jr. et al. (Wilson)	4,755,930	Jul. 5, 1988
Singh et al. (Singh)	5,903,908	May 11, 1999
		(filed Oct. 15, 1996)

Claims 1-20 stand rejected under 35 U.S.C. § 103 as unpatentable over Singh and Wilson.

Appeal No. 2001-2164
Application No. 09/024,620

Reference is made to the briefs and answer for the respective positions of appellants and the examiner.

OPINION

At the outset, we note that, in accordance with appellants' grouping of the claims at page 5 of the principal brief, all claims will stand or fall together. Accordingly, we will focus on instant independent claim 1.

In rejecting claims under 35 U.S.C. 103, it is incumbent upon the examiner to establish a factual basis to support the legal conclusion of obviousness. See In re Fine, 837 F.2d 1071, 1073, 5 USPQ2d 1596, 1598 (Fed. Cir. 1988). In so doing, the examiner is expected to make the factual determinations set forth in Graham v. John Deere Co., 383 U.S. 1, 17, 148 USPQ 459, 467 (1966), and to provide a reason why one having ordinary skill in the pertinent art would have been led to modify the prior art or to combine prior art references to arrive at the claimed invention. Such reason much stem from some teachings, suggestions or implications in the prior art as a whole or knowledge generally available to one having ordinary skill in the

Appeal No. 2001-2164
Application No. 09/024,620

art. Uniroyal, Inc. v. Rudkin-Wiley Corp., 837 F.2d 1044, 1051, 5 USPQ2d 1434, 1438 (Fed. Cir.), cert. denied, 488 U.S. 825 (1988); Ashland Oil, Inc. v. Delta Resins & Refractories, Inc. , 776 F.2d 281, 293, 227 USPQ 657, 664 (Fed. Cir. 1985), cert. denied, 475 U.S. 1017 (1986); ACS Hosp. Sys., Inc. v. Montefiore Hosp., 732 F.2d 1572, 1577, 221 USPQ 929, 933 (Fed. Cir. 1984).

These showings by the examiner are an essential part of complying with the burden of presenting a prima facie case of obviousness. Note In re Oetiker, 977 F.2d 1443, 1445, 24 USPQ2d 1443, 1444 (Fed. Cir. 1992). If that burden is met, the burden then shifts to the applicant to overcome the prima facie case with argument and/or evidence. Obviousness is then determined on the basis of the evidence as a whole and the relative persuasiveness of the arguments. See Id.; In re Hedges, 783 F.2d 1038, 1040, 228 USPQ 685, 687 (Fed. Cir. 1986); In re Piasecki, 745 F.2d 1468, 1472, 223 USPQ 785, 788 (Fed. Cir. 1984); and In re Rinehart, 531 F.2d 1048, 1051, 189 USPQ 143, 146-147 (CCPA 1976). Only those arguments actually made by appellant have been considered in this decision. Arguments which appellant could have made but chose not to make in the brief have not been considered and are deemed to be waived [see 37 CFR 1.192 (a)].

The examiner points to column 10, lines 20 and 25-42, column

Appeal No. 2001-2164
Application No. 09/024,620

3, lines 39-42, column 8, lines 1-2 and Figure 4, items 402 and 403, of Singh, and column 6, lines 51-55 and 57-60, of Wilson in support of the examiner's position that the subject matter of instant claim 1 would have been obvious.

In particular, the examiner states that Singh discloses a method for improving memory latency in a system having a plurality of processing units, each processing unit having an associated cache. The examiner contends that a requesting processing unit issuing a message to an interconnect of the system indicating that the requesting processing unit desires to read a value from an address of the memory device of the system is disclosed by Singh at column 10, lines 20 and 25-42, and in Wilson, at column 6, lines 51-54.

The examiner further contends that each cache snoops the interconnect to detect the message and that this is disclosed at "column 8, lines 1-2" (answer-page 4). While the examiner does not indicate which reference is being cited, it appears to be Singh. The examiner "interprets snooping to be the process of caches checking their contents to see if the requested exists there" (answer-page 4).

The examiner describes the "each cache transmitting a response..." paragraph of claim 1 as being disclosed at column

Appeal No. 2001-2164
Application No. 09/024,620

10, lines 25-42, of Singh, as well as column 6, lines 55 and 57-60, of Wilson. The examiner alleges that the "associating a priority...", "detecting each response..." and "forwarding a response..." limitations of the claim are disclosed at "column 8, lines 33-39" (answer-page 4-though the reference is not identified, we presume the examiner is referring to Singh).

For their part, appellants contend that Singh is "merely pointing out the prior art MESI coherency protocol...Singh does not describe any cache coherency state corresponding to the tagged state..." (principal brief-page 6) because the "tagged state" refers to a cache line which contains a modified value. While the prior art, including Singh, refers to a "Modified (M)" state (e.g., Singh-column 8, line 26), appellants argue that the "primary distinction" between that "modified" state and the "tagged" state of the instant claimed invention "is that the "tagged" state is used when a value is not only modified, but also shared by multiple caches (horizontal caches, i.e., caches associated with different processing units). The 'modified' prior art state can only be used with a single cache; in other words, if a cache is in the prior art 'modified' state, that value is not shared by any other caches in the multi-processor computer system" (principal brief-page 6).

Appeal No. 2001-2164
Application No. 09/024,620

It is our view, that the examiner has not established a prima facie case of obviousness.

Independent claim 1 is very specific as to "a tagged intervention" response transmitted by a cache. This response is transmitted by a responding cache to indicate three things:

1. That the responding cache contains a "modified" value, i.e., a value that has not been written to system memory;
2. That the responding cache "shares" the modified value with at least one other cache of an adjacent processing unit; and
3. That the responding cache will "exclusively" source the value.

Independent claim 11 contains similar recitations. Thus, three specific requirements are present.

We have reviewed the portions of Singh and Wilson cited by the examiner and we do not find therein a disclosure or suggestion of the three claimed requirements. Singh discloses a Shared (S) coherency state from the MESI protocol. As indicated at column 9, lines 4-6, when the data is marked "Shared," this indicates "that although the cache memory has the latest copy of the data, other cache memories in the system also have copies of it." Accordingly, Singh does appear to disclose a cache that

Appeal No. 2001-2164
Application No. 09/024,620

"shares" a value with other caches.

The instant claims require not only that the cache "share" a value with other caches, but that that shared value be a modified value. Accordingly, the tagged intervention response is provided only by a responding cache holding requested data that are both shared and modified.

We find nothing in Singh indicating that the requested data held by a responding cache is "modified," i.e., the value held has not been written to system memory. In fact, every indication in Singh is that the data in the responding caches is consistent with that in the system memory. Column 7, lines 60 et seq., of Singh states that the processor accesses from the instruction fetch unit "are always fetch operations with no intent to modify the data..." Column 8, lines 21-25, of Singh indicates that the processor may update any portion of the cache line it chooses, resulting in a "write command being issued to the data cache memory...with the modified data and the data cache memory...updates its copy of the cache line to reflect the change," but this does not indicate that the data in the cache memory is different, or inconsistent, from the data in the system memory. In fact, it appears to indicate that the data in the

Appeal No. 2001-2164
Application No. 09/024,620

cache memory is updated to be, in fact, consistent with the data in the system memory. Column 10, lines 18-19, of Singh, also states that "the processor uses these micro-operations to maintain the cache consistency" (emphasis added), suggesting, yet, again, that Singh is interested in cache consistency with the processor, or system, memory, rather than the inconsistency required by the instant claims.

Thus, we find no teaching or suggestion in Singh of a tagged state and associated tagged intervention response of the claimed invention.

Turning to Wilson, we find nothing therein that would remedy the deficiency of Singh with regard to the inconsistency of data between the cache and the system memory. While the examiner cites lines 55 and 57-60, of Wilson, we agree with appellants that this portion of Wilson appears only to disclose that a cache having an exclusive copy of data can supply the data to a requesting cache instead of main memory. Appellants allege that it is well known that "if a cache holds data in the Exclusive state, the data are by definition consistent with the main memory and held in only a single one of a plurality of horizontal caches." The examiner offers no convincing rebuttal to this allegation.

Appeal No. 2001-2164
Application No. 09/024,620

The examiner, at page 10 of the answer, contends that while it "may or may not be true" that Wilson does not discuss data being kept inconsistent with cached copies, "appellants have not claimed such a limitation" and even assuming this is claimed, "not sharing the value at the time it sources it is indeed keeping that data inconsistent." The examiner's position is in error. Clearly, the claims require the cached data to be inconsistent with the data in system memory because the claims require a "modified" value and that is what "modified" means, i.e., that the data in the cache has been "modified," or changed, so that it is no longer consistent with the data in the system memory. Wilson clearly does not disclose or suggest the tagged intervention response specifically set forth in the instant claims.

Since neither of the applied references teaches or suggests the tagged intervention response claimed by appellants, the instant claimed subject matter cannot be said to have been obvious over the combination of these references.

We also note that the examiner's rejection under 35 U.S.C. § 103 is deficient for another reason. The examiner indicates that the rejection is based on the combination of Singh and Wilson. But, while the examiner alleges that Singh teaches

Appeal No. 2001-2164
Application No. 09/024,620

certain claimed elements and that Wilson teaches certain claimed elements, with those alleged teachings apparently overlapping, the examiner never indicates exactly what, if anything, is missing from Singh that is taught by Wilson. More importantly, assuming these references are to be combined, since the rejection was based on Singh and Wilson, rather than Singh or Wilson, the examiner never offers any motivation for combining the teachings of these references in any manner nor does the examiner indicate what modification is being made to Singh by Wilson or vice-versa.

Since the examiner has failed to establish a prima facie case of obviousness with regard to the instant claimed subject matter, we will not sustain the rejection of claims 1-20 under 35 U.S.C. § 103.

Appeal No. 2001-2164
Application No. 09/024,620

The examiner's decision is reversed.

REVERSED

JAMES D. THOMAS)	
Administrative Patent Judge)	
)	
)	
)	
)	
KENNETH W. HAIRSTON)	BOARD OF PATENT
Administrative Patent Judge)	APPEALS AND
)	INTERFERENCES
)	
)	
)	
ERROL A. KRASS)	
Administrative Patent Judge)	

EK/RWK

Appeal No. 2001-2164
Application No. 09/024,620

BRACEWELL & PATTERSON, L.L.P.
INTELLECTUAL PROPERTY LAW
P.O. BOX 969
AUSTIN, TEXAS 78767-0969