

The opinion in support of the decision being entered today was not written for publication in a law journal and is not binding precedent of the Board.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. 41

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE BOARD OF PATENT APPEALS
AND INTERFERENCES

Ex parte JEFFREY V. MERKEY

Appeal No. 2001-2553
Application No. 08/512,369

ON BRIEF

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

DECISION ON APPEAL

This is a decision on appeal from the final rejection of claims 50-55, 58-75 and 77.

The invention pertains to the allocation and scheduling of processors in a multiprocessing computer system. More particularly, thread-scheduling creates a strong affinity between

Appeal No. 2001-2553
Application No. 08/512,369

Anderson et al. (Anderson), "The Performance Implications of Thread Management Alternatives for Shared-memory Multiprocessors", IEEE Trans On Computers, Vol. 38, No. 12, Dec. 1989, pp. 1631-1644. (Anderson)

Cheng et al., "Scheduling in Parallel Systems with a Hierarchical Organization of Tasks", School of Comp Science, Carleton Univ, Ottawa, Ontario, Canada, Jul. 1992, pp. 377-386. (Cheng)

Baron et al., "MACH Kernel Interface Manual", Dept. Comp. Science, Carngie-Mellon Univ. August 13, 1990. pp. 1-6. (Baron)

Squillante et al., "Analysis of task Migration in Shared-Memory Multiprocessor Scheduling", IBM Research Div. Feb. 91. pp. 143-155. (Squillante)

Draves et al., "Using Continuations to Implement Thread Management and Communication in Operating Systems", School of Comp. Science, Carngie-Mellon Univ. March, 91. pp. 122-136. (Draves)

Claims 50-55, 58-75 and 77 stand rejected under 35 U.S.C. § 103. As evidence of obviousness, the examiner offers Rudolph and Leung with regard to claims 50 and 51 and Rudolph and Anderson with regard to claims 52 and 75. The examiner offers Rudolph, Anderson and Leung with regard to claim 53 and Rudolph, Anderson and either one of Valencia or Cheng with regard to claims 54 and 55. The examiner cites Valencia and Anderson with regard to claims 58 and 77 and Valencia and Cheng with regard to claim 59. The examiner cites either one of Belo or Baron in view of Rudolph with regard to claims 60 and 64, adding Valencia to

Appeal No. 2001-2553
Application No. 08/512,369

this combination with regard to claims 61, 62, 66 and 68, and further adding Anderson with regard to claim 63. The examiner offers either one of Belo or Baron, with Rudolph and Leung with regard to claim 65 and either one of Belo or Baron, with Rudolph and Cheng with regard to claims 67 and 69-73. The examiner applies Squillante and Draves with regard to claim 74.

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

OPINION

In arguing certain claim limitations, appellant groups the claims into four categories. The first category, including claims 50, 51, 53 and 65, is identified as "sleep request claims" since the claims in this group require identifying as a busiest processor a processor which has received the smallest number of sleep requests during a sampling period. The second category, including claims 52-55, 60-71 and 75, is identified as "minimum count claims" since the claims in this group require identifying as a busiest processor or identifying as a popular processor a processor which has at least a predetermined number of eligible threads. The third category, including claim 74, is identified

Appeal No. 2001-2553
Application No. 08/512,369

as an "idle thread claim" since this claim requires identifying as a busiest processor a processor which has spent a higher proportion of its processing capacity running application threads versus running an idle thread relative to other processors in the system. The fourth category, including claims 58, 59, 71-73 and 77, is identified as "limited frequency claims" since these claims require limiting the relative frequency of global dispatch queue accesses. Some claims, e.g., claim 53, appear in more than one group because they have limitations which encompass more than one group.

We accept appellant's grouping of the claims in this manner and we will treat the claim limitations as argued by appellant.

Turning first to independent claim 50, this claim requires, inter alia, that a busiest processor be identified as a processor which has received the smallest number of sleep requests during a sampling period.

While the examiner employs Rudolph as disclosing most of the claim limitations (see pages 5-6 of the answer), the examiner recognizes that Rudolph does not take into consideration the number of sleep requests, as claimed. Thus, the examiner turns to Leung for the teaching of identifying as a busiest processor a processor which has received the smallest number of sleep

Appeal No. 2001-2553
Application No. 08/512,369

requests during a sampling period, referring to Figure 4 and alleging that "high execution time reduces the number of sleeps that can be accommodated within a specific time period" (answer-page 6). The examiner then concludes that it would have been obvious "to identify the processor with the smallest number of sleep requests since this would indicate the time a processor spends in executing a task" (answer-page 6).

First, we find the examiner's rationale for making the combination insufficient to establish a prima facie case of obviousness because while it may be that identifying the processor with the smallest number of sleep requests indicates the time a processor spends in executing a task, this does not explain why or how the skilled artisan would have been led to apply such an identification of a processor with the smallest number of sleep requests to Rudolph to achieve any particular result.

More importantly, we will not sustain the rejection of claim 50 (or of claim 51, containing a similar limitation) over Rudolph and Leung because we agree with appellant that Leung does not appear to teach what the examiner alleges it to teach. In particular, the examiner points to Figure 4 and to page 411, section 6, first paragraph, for a teaching of "computing the

Appeal No. 2001-2553
Application No. 08/512,369

number of times that a job A needs to enter into sleep" (answer-page 16). The examiner equates this teaching to a teaching of identifying as a busiest processor a processor which has received the smallest number of sleep requests of any of the processors during a sampling period.

Leung refers to the number of times a "job" needs to enter into sleep. Claim 50 refers to a smallest number of sleep requests received by a "processor." While a processor may run a "job," or a plurality of "jobs," a processor is not a job and it makes no sense for the examiner to contend that Leung teaches the claimed "identifying as a busiest processor a processor which has received the smallest number of sleep requests..." when Leung identifies no processor at all.

Accordingly, we will not sustain the rejection of claims 50 and 51 under 35 U.S.C. § 103 over Rudolph and Leung or of claim 53 over Rudolph, Leung, and Anderson, or of claim 65 over Belo or Baron and Rudolph and Leung.

We now turn to claims 52-55, 60-71 and 75, claims 52, 60 and 75 being independent claims, and the limitation of identifying a busiest, or popular, processor as a processor which has at least a predetermined number of eligible threads.

Independent claims 52 and 75 stand rejected under 35 U.S.C.

Appeal No. 2001-2553
Application No. 08/512,369

§ 103 in view of Rudolph and Anderson, while independent claim 60 stands rejected under 35 U.S.C. § 103 over either one of Belo or Baron in view of Rudolph.

It is the examiner's position, with regard to claims 52 and 75, that Rudolph shows associating an unlocked queue of threads, selecting movable threads, and identifying a busiest popular processor, a popular processor being one that has an unlocked local queue containing at least a predetermined number of eligible threads, but does not specifically recite "threads." The examiner turns to Anderson for a teaching of threads as a unit of scheduling, at page 1631, right column, first paragraph, and concludes that it would have been obvious to place the threads in the queues for the reasons set forth in Anderson, at pages 1631-1632.

With regard to claim 60, the examiner contends that Belo or Baron shows processors being assigned to exactly one processor set, a shared memory, a bus connecting the processors with the shared memory and an unlocked local queue of threads associated with each of the processors. The examiner contends that Rudolph shows a global dispatch of threads not presently associated with any of the processors, means for selecting movable threads from the unlocked local queues, etc. and concludes that it would have

Appeal No. 2001-2553
Application No. 08/512,369

been obvious "to determine a popular processor so that efficient load balancing can be performed" (answer-page 12). While it is not clear from this rejection exactly how the alleged combination is to be made, i.e., how the alleged teaching of Rudolph is to be combined, exactly, with the alleged teachings of Belo or Baron, this is not argued by appellant and is not an issue in the case.

Appellant's argument centers on Rudolph and its alleged teaching of a popular processor whose "unlocked local queue contains at least a predetermined number of eligible threads." It is appellant's view that "[t]he load of a workpile is measured as the number of tasks on the workpile," recited at page 240, left column, of Rudolph, fails to teach "at least a predetermined number" of threads in the workpile, even if we consider the workpiles to be unlocked local queues. Appellant further points to Rudolph's recitation of a threshold value, τ , at page 240 and concludes that this compares a thread count in one workpile with the thread count in another workpile, instead of comparing the thread count in one workpile with some predetermined minimum count, so that Rudolph fails to teach "at least a predetermined number" of threads.

Since this is the only argument made by appellant relative to the "minimum count claims," and we do not agree with

Appeal No. 2001-2553
Application No. 08/512,369

appellant's view, we will sustain the rejection of claims 52, 54, 55, 60-64, 66-70 and 75 under 35 U.S.C. § 103. This is because as we read page 240 of Rudolph, the reference recites "A system-wide threshold value, τ , is fixed such that only task workpiles whose length differs by more than τ will perform a balancing operation." Thus, contrary to appellant's contention, Rudolph is not comparing a thread count in one workpile with the thread count, or number of tasks on a workpile, in another workpile, but instead, is comparing a thread count to a predetermined, or threshold, value, τ .

Specifically, with regard to claims 54 and 55, appellant argues that Cheng, applied for a teaching of global and local queues, teaches queues which are neither local nor global.

First, we note that the Cheng and Valencia references are applied in the alternative in rejecting claims 54 and 55. Since appellant presents no argument regarding Valencia, we must take the examiner's position as viable and the rejection of claims 54 and 55 can be sustained for that reason alone.

Moreover, the examiner points out that the "ready queues" of Cheng are local (pointing to Figure 1-task queue organizations) and that "the tasks in a centralized ready queue are local and are in a global queue and Cheng teaches that" (answer-pages 18-

Appeal No. 2001-2553
Application No. 08/512,369

19). While the local and global queues in Cheng may differ from what appellant intends, the plain language of the claims, "local" and "global" queues, appears to be suggested by Cheng.

Accordingly, we will sustain the rejection of claims 52, 54, 55, 60-64, 66-70 and 75 under 35 U.S.C. § 103. Arguments which appellant could have made, but did not, are waived. In re Kroekel, 803 F.2d 705, 231 USPQ 640 (Fed. Cir. 1986).

Turning now to "idle thread" in claim 74, the examiner relies on Squillante and Draves, with Draves being relied on for its teaching of an idle thread. While Draves mentions "an idle thread" in a footnote, at page 133, appellant argues that not only is this passing reference to "idle thread" not enabling, but that Draves' "idle thread" is not the same thing as appellant's "idle thread" (principal brief-page 8). In particular, appellant argues that, at page 20, lines 14-16, and at numerous places on pages 21-32 of the instant specification, appellant defines "idle thread" as a thread associated with a particular queue, and it is a thread that performs searching and/or scheduling functions. Contrary to this, appellant argues, Draves' "idle thread" is "simply a thread that is not running" (principal brief-page 8) and Draves mentions nothing about the idle thread being associated with a given queue, or about its being used for

Appeal No. 2001-2553
Application No. 08/512,369

searching or scheduling functions.

Moreover, appellant argues, the rejection should be reversed because Squillante does not discuss "idle threads," as used in the claims.

We will not sustain the rejection of claim 74 because while Squillante teaches "busy" and "idle" processors, the claim is directed to "application" and "idle" "threads." While Draves does mention, in passing, an "idle thread," it is unclear to us how a mere mention of an "idle thread" would have suggested any modification to Squillante regarding identifying a busiest processor, the busiest processor having spent a higher proportion of its processing capacity running application threads versus running an idle thread relative to a plurality of other processors.

Accordingly, the rejection of claim 74 is not sustained.

Turning, finally, to the rejection of claims 58, 59, 71-73 and 77 in the "limited frequency" group, appellant argues that Valencia teaches local and global run queues and moving processes between queues but does not teach the claimed global-local-global sequence, which requires a step that yields control to a local queue thread at least once between two steps that each yield

Appeal No. 2001-2553
Application No. 08/512,369

control to global queue thread(s).

The examiner's response is to point to page 379, right column, second paragraph, of Cheng. The examiner contends that Tr is the ratio of "the number of tasks moved one level down the tree to the number of processor(?) below the child task queue" (answer-page 21). The examiner concludes, without support, in our view, that "since the combination of the references necessarily includes all of the steps needed for limiting frequency of global dispatch queue accesses, the combination must perform the same function" (answer-page 21).

We find appellant's argument convincing. The examiner has failed to show that Valencia suggests the three control yielding steps for limiting the relative frequency of global dispatch queue accesses.

Accordingly, we will not sustain the rejection of claims 58, 59, 71-73 and 77 under 35 U.S.C. § 103.

Since we have sustained the rejection of claims 52, 54, 55, 60-64, 66-70, and 75 under 35 U.S.C. § 103 and have not sustained the rejection of claims 50, 51, 53, 58, 59, 65, 71-73 and 77 under 35 U.S.C. § 103, the examiner's decision is affirmed-in-part.

Appeal No. 2001-2553
Application No. 08/512,369

No time period for taking any subsequent action in connection with this appeal may be extended under 37 CFR § 1.136(a).

AFFIRMED-IN-PART

JAMES D. THOMAS)	
Administrative Patent Judge)	
)	
)	
)	
)	
ERROL A. KRASS)	BOARD OF PATENT
Administrative Patent Judge)	APPEALS AND
)	INTERFERENCES
)	
)	
)	
ANITA PELLMAN GROSS)	
Administrative Patent Judge)	

EK/RWK

Appeal No. 2001-2553
Application No. 08/512,369

JAMES D. LILES
DINSMORE & SHOHL LLP
255 EAST FIFTH STREET
1900 CHEMED CENTER
CINCINNATI, OH 45202