UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
12/821,784	06/23/2010	Michael J. Branson	ROC920100043US1	4549

46296 7590 10/28/2016
MARTIN & ASSOCIATES, LLC
P.O. BOX 548
CARTHAGE, MO 64836-0548

EXAMINER
DAO, TUAN C.

ART UNIT	PAPER NUMBER
2193	

NOTIFICATION DATE	DELIVERY MODE
10/28/2016	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

derekm@ideaprotect.com

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

Ex parte MICHAEL J. BRANSON
and JOHN M. SANTOSUOSSO

Appeal 2015-000104
Application 12/821,784
Technology Center 2100

Before ALLEN R. MacDONALD, JOHN F. HORVATH, and
KEVIN C. TROCK, *Administrative Patent Judges*.

HORVATH, *Administrative Patent Judge*.

DECISION ON APPEAL

STATEMENT OF THE CASE

Appellants seek review, under 35 U.S.C. § 134(a), of the Examiner's rejection of claims 1, 2, 4–8, 16, 17, and 19–22. We have jurisdiction under 35 U.S.C. § 6(b).

We AFFIRM.

SUMMARY OF THE INVENTION

The invention is directed to dynamically allocating a job on a multi-nodal, parallel computer system. Spec. ¶ 2.

Claim 1, reproduced below, is illustrative of the claimed subject matter:

> 1. An apparatus comprising:
>
> a plurality of nodes of a multi-nodal computer system, wherein the plurality of nodes are connected by a plurality of networks, where each of the plurality of nodes has at least one central processing unit (CPU) coupled to a memory;
>
> an application having a plurality of jobs, each with at least one processing unit executing on the plurality of nodes;
>
> a collector collecting metrics of the system, nodes, application, jobs and processing units in order to determine how to allocate the jobs on the system;
>
> a job optimizer dynamically changing the allocation of processing units on the plurality of nodes based on the collected metrics, wherein the job optimizer dynamically changes the allocation of the processing units by splitting a job running on a first node into multiple jobs allocated to different nodes of the plurality of nodes, wherein an interprocess communication between processing units of the job is within a node before splitting the job, and after splitting the job the interprocess communication is then handled by the multi-nodal computer system over a communication link between the different nodes.

REFERENCES

Fellenstein	US 2006/0152756 A1	July 13, 2006
Dillenberger	US 2008/0216087 A1	Sept. 4, 2008
Kaminsky	US 2009/0048998 A1	Feb. 19, 2009
Abu-Ghazaleh	US 2009/0083390 A1	Mar. 26, 2009
Barsness	US 2009/0313636 A1	Dec. 17, 2009
Votta	US 2010/0011254 A1	Jan. 14, 2010

REJECTIONS

Claims 1, 16, and 21 stand rejected under 35 U.S.C. § 103(a) as unpatentable over Abu-Ghazaleh and Fellenstein. Final Act. 2.

Claims 2 and 17 stand rejected under 35 U.S.C. § 103(a) as unpatentable over Abu-Ghazaleh, Fellenstein, and Barsness. Final Act. 9.

Claim 4 stands rejected under 35 U.S.C. § 103(a) as unpatentable over Abu-Ghazaleh, Fellenstein, and Dillenberger. Final Act. 10.

Claims 5, 6, 8, 19, 20, and 22 stand rejected under 35 U.S.C. § 103(a) as unpatentable over Abu-Ghazaleh, Fellenstein, and Votta. Final Act. 11.

Claim 7 stands rejected under 35 U.S.C. § 103(a) as unpatentable over Abu-Ghazaleh, Fellenstein, Votta, and Kaminsky. Final Act. 15.

ANALYSIS

Claims 1, 2, 4, 8, 16, 17, 21, and 22

Issue 1: Whether Abu-Ghazaleh teaches or suggests a job optimizer that dynamically changes the allocation of processing units by splitting a job running on a first node into multiple jobs allocated to different nodes as required by claim 1.

The Examiner finds Abu-Ghazaleh (hereafter "Abu") teaches or suggests this limitation. Final Act. (citing, *inter alia*, Abu ¶¶ 17, 125); Ans. 4–9 (citing, *inter alia*, Abu ¶¶ 17, 39, 94, 125). Appellants argue that although Abu teaches partitioning a job into sub-tasks and assigning different sub-tasks to different nodes, "Abu does not teach or suggest splitting a job running on a first node into multiple jobs allocated to different nodes." App. Br. 7. That is, Appellants argue that the allocation of tasks

and subtasks in Abu is *predetermined and fixed*, and "[t]here is nothing about splitting a job as claimed." *Id.* We are not persuaded by Appellants' arguments.

Abu discloses a self-organizing, distributed computing grid that can "tolerate a significant portion of its participating resources to be *dynamically added or removed, even those that are being used by active computations*." Abu ¶ 17 (emphasis added). Abu divides jobs into tasks, and allocates tasks to the nodes of the grid in order to "accomplish [the] tasks which are partitioned to [the] various nodes, and also *subject to maintenance in the event that the underlying resources change in availability* or use." *Id.* (emphasis added). Abu teaches that "a job is a set of tasks . . . carried out on a set of resources, and a task is an atomic operation to be performed on a resource." *Id.* ¶ 94. Abu further teaches that nodes "may be placed within the [grid] hierarchy or removed from the hierarchy *while the distributed task is in progress*, and the new node *allocated a portion of the distributed task*." *Id.* ¶ 125.

In view of these disclosures, we are not persuaded by Appellants' argument that Abu's teaching is limited to allocating tasks in a predetermined and fixed manner, and that Abu fails to teach or suggest dynamically splitting a job or task that had been allocated to a first node into multiple jobs or tasks that are allocated to different nodes as required by claim 1. "The test for obviousness is not . . . that the claimed invention must be expressly suggested in any one or all of the references. Rather, the test is what the combined teachings of the references would have suggested to those of ordinary skill in the art." *In re Keller*, 642 F.2d 413, 425 (CCPA 1981).

> *Issue 2: Whether Abu teaches or suggests that an interprocess communication between processing units occurs within a node before splitting a job and over a communication link between different nodes after splitting the job as required by claim 1.*

The Examiner finds Abu teaches or suggests this limitation. Final Act. 5 (citing Abu ¶¶ 50, 61, 62, 67, 71, 72, and 122–125); Ans. 11–12. Appellants argue the portions of Abu cited by the Examiner "describe the super-node and how the communication between nodes is handled by the self organizing grid system," and that "[n]othing in these cited sections describe [that] interprocess communication between processing units of the job is within a node before splitting the job," and is handled "over a communication link between the different nodes" after splitting the job. App. Br. 11–12. We are not persuaded by Appellants' arguments.

As noted *supra*, "[t]he test for obviousness is not . . . that the claimed invention must be expressly suggested in any one or all of the references. Rather, the test is what the combined teachings of the references would have suggested to those of ordinary skill in the art." *In re Keller*, 642 F.2d at 425. Appellants' "interprocess communication" limitation is little more than a tautological statement that the communications between two processing units occurs between the locations where the two processing units are executing. That is, when the two processing units are executing within a node, they *communicate* between locations *within the node*. When the job is split, and the two processing units are executing on different nodes, they *communicate over a link between* the locations of *the different nodes*. Abu teaches or suggests this tautological truth. For example, Abu teaches that when two tasks are executing on nodes in the same neighborhood (e.g.,

nodes 2 and 6 in layer 0 of Figure 1), they communicate over links within the neighborhood. However, when the tasks are executing on nodes in different neighborhoods (e.g., nodes 4 and 6 in layer 0 of Figure 1), they communicate through their respective supernodes (e.g., nodes 2 and 15 in layer 1 of Figure 1) over links between the neighborhoods. *See* Abu ¶ 62, Fig. 1. Thus, we find the sections of Abu cited by the Examiner teach or suggest the "interprocess communication" limitations to a person of ordinary skill in the art.

Accordingly, for the reasons discussed above, we sustain the Examiner's rejection of claim 1, as well as of claims 2, 4, 8, 16, 17, 21, and 22, not separately argued. *See* App. Br. 7–13, 15.

Claims 5–7, 19, and 20

Issue 3: Whether the combination of Abu and Votta teaches or suggests collecting metrics associated with the processing units, application, and nodes, where the metrics are selected from CPU utilization, memory utilization, throughput, latency, and heap size.

Claims 5–7, 19, and 20 require the collected metrics to be associated with the processing units (claims 5/19), the application (claims 6/20), and the nodes on which the processing units are executing (claim 7). The claims also require the collected metrics to be selected from CPU utilization (all claims), memory utilization (all claims), throughput (claims 5/19 and 6/20), latency (claims 5/19 and 6/20), and heap size (claim 7).

The Examiner finds the combination of Abu and Votta teaches or suggests these limitations. Final Act. 11–13, 15. In particular, the Examiner

finds Abu teaches collecting resource utilization metrics, and Votta teaches resource utilization metrics can include CPU utilization, memory utilization, throughput, and latency. *Id.* (citing Abu ¶¶ 30, 39, 93; Votta ¶ 53; Fig. 3).

Appellants argue the Examiner erred in rejecting claims 5–7, 19, and 20 because Votta's metrics are general CPU utilization metrics rather than metrics for "CPU utilization **by the processing unit**" as recited in claims 5/19, metrics for "CPU utilization **by the application**" as recited in claims 6/20, and metrics for "CPU utilization **for the node**" as recited in claim 7. App. Br. 14, 16. We are not persuaded by Appellants' arguments.

Abu teaches allocating system resources to optimize the performance of the grid as a whole, and individual applications running on the grid. Abu ¶ 30. Abu further teaches that effective resource allocation requires considering different metrics and constraints, such as resource utilization for different jobs, and for different storage, computing, and network resources. *Id.* ¶ 93. Abu further teaches that peak performance can be achieved by minimizing job execution time and communication delays, and maximizing resource utilization. *Id.* ¶ 94. Votta teaches enhancing throughput in grid computing systems based on monitored software parameters including, but not limited to, the load on one or more processors, CPU utilization, memory utilization, transaction latency, system throughput, and other monitored performance metrics gathered from the operating system. Votta ¶¶ 3, 53.

As noted *supra*, "[t]he test for obviousness is not . . . that the claimed invention must be expressly suggested in any one or all of the references. Rather, the test is what the combined teachings of the references would have suggested to those of ordinary skill in the art." *In re Keller*, 642 F.2d at 425. Appellants' claims 5–7, 19, and 20 recite collecting various performance

metrics at different levels of granularity, namely, from individual tasks or processing units, from the application as a whole (i.e., from the sum of the processing units), and from the sum of the nodes on which the processing units, and therefore application, are executing. Votta similarly teaches or suggests gathering performance metrics at different levels of granularity based on parameters obtained from the operating system (e.g., the load on one or more processors, individual transaction latencies, and system throughput). Accordingly, we are not persuaded by Appellants' argument that the combination of Abu and Votta fails to teach or suggest gathering various performance metrics (e.g., CPU utilization, latency, throughput) for the processing units (claims 5/19), for the application (claims 6/20), or for the nodes (claim 7).

Accordingly, for the reasons discussed, we sustain the Examiner's rejections of claims 5–7, 19, and 20.

DECISION

The rejection of claims 1, 16, and 21 under 35 U.S.C. § 103(a) as unpatentable over Abu-Ghazaleh and Fellenstein is affirmed.

The rejection of claims 2 and 17 under 35 U.S.C. § 103(a) as unpatentable over Abu-Ghazaleh, Fellenstein, and Barsness is affirmed.

The rejection of claim 4 under 35 U.S.C. § 103(a) as unpatentable over Abu-Ghazaleh, Fellenstein, and Dillenberger is affirmed.

The rejection of claims 5, 6, 8, 19, 20, and 22 under 35 U.S.C. § 103(a) as unpatentable over Abu-Ghazaleh, Fellenstein, and Votta is affirmed.

The rejection of claim 7 under 35 U.S.C. § 103(a) as unpatentable over Abu-Ghazaleh, Fellenstein, Votta, and Kaminsky is affirmed.

No time period for taking any subsequent action in connection with this appeal may be extended under 37 C.F.R. § 1.136(a)(1)(iv).

AFFIRMED