



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

Table with 5 columns: APPLICATION NO., FILING DATE, FIRST NAMED INVENTOR, ATTORNEY DOCKET NO., CONFIRMATION NO. Includes application details for JU WANG and examiner information for LEE, PHILIP C.

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):

- IPDocket@shb.com
docket.shb@clarivate.com
usdocket@microsoft.com

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

Ex parte JU WANG, BRADLEY GENE CALDER, and
ARILD E. SKJOLSVOLD

Appeal 2019-000962
Application 13/366,039
Technology Center 2400

Before MICHAEL J. STRAUSS, DANIEL N. FISHMAN, and
NABEEL U. KHAN, *Administrative Patent Judges*.

STRAUSS, *Administrative Patent Judge*.

DECISION ON APPEAL

STATEMENT OF THE CASE¹

Pursuant to 35 U.S.C. § 134(a), Appellant² appeals from the Examiner's decision to reject claims 1–10, 12–19, 21, and 22. Appeal Br. 3. Claims 11 and 20 are canceled. *Id.* We have jurisdiction under 35 U.S.C. § 6(b).

We AFFIRM-IN-PART.

¹ We refer to the Specification, filed February 3, 2012 (“Spec.”); Final Office Action, mailed September 1, 2017 (“Final Act.”); Appeal Brief, filed April 2, 2018 (“Appeal Br.”); Examiner’s Answer, mailed September 26, 2018 (“Ans.”); and Reply Brief, filed November 15, 2018 (“Reply Br.”).

² We use the word Appellant to refer to “applicant” as defined in 37 C.F.R. § 1.42. Appellant identifies the real party in interest as MICROSOFT TECHNOLOGY LICENSING, LLC. Appeal Br. 3.

CLAIMED SUBJECT MATTER

The claims are directed to managing partitions in a scalable environment. Spec., Title. Claims 1 and 9, reproduced below with bracketed reference labels added and disputed limitations emphasized in *italics*, are illustrative of the claimed subject matter:

1. A method for performing computations in a distributed computing environment, comprising:
 - [(i)] receiving one or more application-defined partitioning system interfaces;
 - [(ii)] creating a plurality of master role instances including the one or more application-defined partitioning system interfaces, the master role instances corresponding to a master storage object;
 - [(iii)] assigning a lease for the master storage object, each master role instance competing for the lease, the master role instance that is assigned the lease being the dictator master role instance;
 - [(iv)] assigning, by the dictator master role instance, a group of partitions to a plurality of partition servers;
 - [(v)] performing one or more computations corresponding to an application using the plurality of partition servers;
 - [(vi)] receiving, by the dictator master role instance, a message from a first partition server comprising partitions the first partition server reports to be currently serving; and
 - [(vii)] *breaking, by the dictator master role instance, a lease of the first partition server on a corresponding first storage object, the breaking of the lease of the first partition server being at the first partition server responsive to detecting, by the dictator master role instance, that the first partition server is currently serving at least one partition the partition server should not currently be serving as indicated by a conflict between one or more partitions assigned to the first partition server in a partition table and the partitions in the received message.*

9. A method for performing computations in a distributed computing environment, comprising:

- [(i)] executing a computation comprising at least two namespaces and at least two master role instances, each master role instance corresponding to a different namespace, each master role instance being the dictator for a corresponding namespace and holding a dictator lease on a master storage object for the corresponding namespace, *the master role instances instantiating a common master module and common fixed-interfaces, while each being of a different master role and having a different application-defined interface;*
- [(ii)] assigning a machine that provides failover service for the master role instances and comprises a backup used for each master role instance of the master role instances in response to a failure event being detected for the master role instance, the backup including the common master module and the common fixed-interfaces;
- [(iii)] in response to detecting a failover event for a master role instance of the master role instances:
- [(iv)] incorporating the different application-defined interface of the master role instance, needed by the backup on the assigned machine to take over for the master role instance, into the backup; and
- [(v)] assigning the backup comprising the common fixed interfaces and the incorporated different application-defined interface as the dictator for the namespace corresponding to the failover event.

REFERENCES

The prior art relied upon by the Examiner is:

Name	Reference	Date
Powers	US 20040203378 A1	Oct. 14, 2004
Nakahara	US 2005/0193227 A1	Sept. 1, 2005
Berg	US 2006/0245433 A1	Nov. 2, 2006
Wookey	US 2007/0171921 A1	July 26, 2007
Van Biljon (“Biljon”)	WO 2011/159842 A2	Dec. 22, 2011
Deshmukh	US 2012/0042030 A1	Feb. 16, 2012
McAlister	US 8,392,482 B1	Mar. 5, 2013

REJECTIONS

Claims 1, 3, 4, and 14–19 stand rejected under pre–AIA 35 U.S.C. § 103(a) as being unpatentable over McAlister, Biljon, and Berg. Final Act. 10–19.

Claim 2 stands rejected under pre–AIA 35 U.S.C. § 103(a) as being unpatentable over McAlister, Biljon, Berg, and Deshmukh. Final Act. 22–23.

Claims 5–8 stand rejected under pre–AIA 35 U.S.C. § 103(a) as being unpatentable over McAlister, Biljon, Berg, Deshmukh, and Powers. Final Act. 24–29.

Claims 9, 10, 12, 13, and 21 stand rejected under pre–AIA 35 U.S.C. § 103(a) as being unpatentable over Biljon and Nakahara. Final Act. 4–9.

Claim 22 stands rejected under pre–AIA 35 U.S.C. § 103(a) as being unpatentable over McAlister, Biljon, Berg, and Wookey. Final Act. 29.

OPINION

Claims 1–8, 14–19, and 22

The Examiner finds the combination of McAlister, Biljon, and Berg teaches or suggests the limitations of independent claim 1. Final Act. 10–14. In connection with the disputed lease breaking limitation (vii), the Examiner finds the limitation is taught or suggested by Berg’s disclosure of a first server receiving from a second server a routing table update message indicating that an assigned partition has been moved to another server. Final Act. 13–14. According to the Examiner, the update message from the second server instructing the first server to update its dynamic routing table “discloses a dictator master instance to instruct a server to break a lease on an incorrect partition.” *Id.*

Appellant contends Berg fails to teach the disputed lease breaking limitation because, *inter alia*,

the Office does not clearly articulate, nor does Appellant believe Berg teaches “. . . the first partition server is currently serving at least one partition the partition server should not currently be serving ***as indicated by a conflict between one or more partitions assigned to the first partition server in a partition table and the partitions in the received message.***”

Appeal Br. 10.

The Examiner responds, finding Berg’s receipt by server 1 (equated to a dictator master role instance) of messages from other servers (i.e., partition servers) reporting their currently served partitions teaches or suggests the update message of limitation (vi). Ans. 4. The Examiner further finds Berg discloses a partition server 2 that, due to an outdated routing table entry, will (i) identify and forward a misrouted message to the correct partition server 4 and (ii) send a routing table update message to server 1, i.e., Appellant’s

dictator master role instance. *Id.* The Examiner finds, when routing server 1 updates the routing table, it breaks the lease of partition server 2. *Id.*

According to the Examiner, partition server 2 is considered to be still serving moved partition 2 because it takes action in response to the misrouted message concerning partition 2 including, *inter alia*, sending a routing table update to server 1. *Id.* The Examiner finds “the update of the dynamic routing table cause[d] the actual breaking of the lease of [partition] server 2 to partition 2 [which] took place at server 2.” *Id.* at 5. The Examiner concludes Berg discloses the disputed lease breaking limitation (vii). *Id.*

Appellant replies, contending,

the rejection does not address claimed features related to a partition that a server “***should not currently be serving***, as indicated by a conflict between one or more partitions assigned to the first partition server in a partition table and the partitions in the received message.” For example, the present application describes that what partition a partition server should be serving may be dictated by a partition table/leases managed by a dictator master role instance. *See, e.g.*, paragraph [0058] of the present application as published. In contrast, in Berg, a conflict with the dynamic routing table merely means the partition has moved, not that any server “***should not currently be serving***” the partition, as recited in the claim.

Reply Br. 2–3. Appellant further contends,

the portions of Berg that are relied upon [for teaching the message from a first partition server comprising partitions the first partition server reports to be currently serving] only state “the various different servers are queried to populate the dynamic routing table with dynamic routing information (step 620).” Berg does not expressly or impliedly disclose any message associated with a query, nor does it expressly or impliedly disclose what a message may comprise.

Reply Br. 3.

We agree with Appellant that the rejection of claim 1 is improper. Lease breaking limitation (vii) requires detecting that the first partition server is currently serving at least one partition the partition server should not currently be serving *as indicated* by a conflict between one or more partitions assigned to the first partition server in a partition table and the partitions indicated in a message received from the first partition server. The Examiner fails to provide sufficient evidence or reasoning to persuade us that Berg's server 1, identified as the claimed dictator master role instance of claim 1, detects conflicts as indicated by a conflict between (i) partitions assigned to server 2 and (ii) messages received from server 2 reporting its current serving partitions. Rather than server 1 identifying a conflict between assigned (i.e., table information) and partition server supplied partition information, Berg discloses partition server 2 identifies routing errors and supplies instructions to server 1 for updating the routing table. Accordingly, Berg fails to teach or suggest detecting that a partition server should not currently be serving a partition *as indicated* by a conflict between one or more partitions assigned to the first partition server in a partition table and the partitions in the received message as required by the disputed lease breaking limitation (vii). Because we agree with at least one of the arguments advanced by Appellant, we need not reach the merits of Appellant's other arguments in connection with claim 1.

For the reasons discussed above, we do not sustain the rejection of independent claim 1 under U.S.C. § 103(a) over McAlister, Biljon, and Berg. For the same reasons, we do not sustain the rejection of independent claim 14 which includes similar limitations or the rejections of dependent claims 2–8, 15–19, and 22 which stand with their respective base claims.

Claims 9, 10, 12, 13, and 21

The Examiner finds the combination of Biljon and Nakahara teaches or suggests the limitations of claim 9, with Nakahara applied for disclosing the disputed portion of limitation (i) reciting “the master role instances instantiating a common master module and common fixed-interfaces, while each being of a different master role and having a different application-defined interface” (hereinafter “the disputed limitation”). Final Act. 4–8. In particular, the Examiner finds the operating systems of Nakahara’s servers teach the disputed master module and Nakahara’s network interfaces teach the claimed fixed-interfaces. *Id.* at 7. The Examiner further finds Nakahara’s backup servers that take over applications from failed servers teach the recited common master module and common fixed-interfaces. Finally, the Examiner finds Nakahara’s servers run multiple different application programs thereby teaching each master role instance being of a different master role and having a different application-defined interface as recited by claim 1. *Id.* at 6.

Appellant contends the Examiner’s interpretation of the disputed limitation is improper in light of the Specification. Appeal. Br. 12. Appellant argues the broadest reasonable interpretation of application-defined interface and common fixed-interfaces is constrained by descriptions provided in the Specification:

An “application-defined interface” refers to a computation, operation, or other function defined by a client for performance by a role instance. . . . An application-defined interface is in contrast to a “fixed interface” for a partitioning system. A fixed interface refers to an application interface that is provided as part

of the partitioning system. A client cannot modify the action of a fixed interface.

Spec. ¶ 21.

Appellant further contends “the combination of the references does not teach or suggest *‘the master role instances ... each being of a different master role ... ’*” Appeal Br. 14. Appellant argues the portion of Nakahara cited by the Examiner only describes a server structure, not the disputed limitation. *Id.* (citing Final Act. 6 and Nakahara ¶ 33).

The Examiner responds:

[U]nder the broadest reasonable interpretation, the “application-defined interface” is interpreted as an application program because one of ordinary skill in the art would [have] readily recognized that an application program is a computation, operation, or other function defined by a developer (i.e. client) and the application program is to be executed (i.e. performed) by a processor (i.e. a role instance). And, the “fixed interface”, under the broadest reasonable interpretation, would be a physical interface, such as a fixed network interface in a system.

Ans. 6.

Addressing Appellant’s contention “the references do not teach ‘the master role instances . . . each being of a different master role,’” (*id.* at 7), the Examiner finds:

Nakahara discloses each of the servers to store different applications (which would cause the servers to perform different tasks/roles) [*para. 33 and 35*]. Nakahara further discloses one server 203g is assumed to be a failure management server, while another server 203a is assumed to be failure management target server that are subject to failure management [*para. 36*]. i.e. the servers are assumed different roles in the failover system.

Id. (brackets in original).

Appellant's arguments are unpersuasive. In connection with the Examiner's interpretation of the terms application-defined interface and common fixed-interfaces, Appellant provides insufficient evidence or argument explaining why the Examiner's interpretation is overly broad. For example, Appellant provides no evidence that the Examiner's interpretation of the terms is inconsistent with the ordinary and customary meaning given those terms as would be understood by one of ordinary skill in the art in the context of the entire disclosure. Nor does Appellant explain why the Examiner's interpretation is inconsistent with the Specification. In the absence of persuasive evidence or argument to the contrary, we conclude the Examiner's interpretation of the disputed terms is reasonable.

We are also unpersuaded the Examiner erred in finding Nakahara teaches or suggests each master role instance being of a different master role. Appellant fails to persuade us of error or deficiency in the Examiner's reasoning for finding Nakahara's servers, by storing different applications, perform different tasks or roles thereby teaching or suggesting the disputed limitation.

For the reasons discussed, Appellant's contentions are unpersuasive of reversible Examiner error. Accordingly we sustain the rejection of independent claim 9 under 35 U.S.C. § 103(a) over Biljon and Nakahara together with the rejection of dependent claims 10, 12, 13, and 21 which are not argued separately with particularity.

CONCLUSION

The Examiner's rejections of claims 1–8, 14–19, and 22 are reversed. The Examiner's rejection of claims 9, 10, 12, 13, and 21 is affirmed.

DECISION SUMMARY

Claims Rejected	35 U.S.C. §	Reference(s)/Basis	Affirmed	Reversed
1, 3, 4, 14–19	103(a)	McAlister, Biljon, Berg		1, 3, 4, 14–19
2	103(a)	McAlister, Biljon, Berg, Deshmukh		2
5–8	103(a)	McAlister, Biljon, Berg, Deshmukh, Powers		5–8
9, 10, 12, 13, 21	103(a)	Biljon, Nakahara	9, 10, 12, 13, 21	
22	103(a)	McAlister, Biljon, Berg, Wookey		22
Overall Outcome:			9, 10, 12, 13, 21	1–8, 14–19, 22

TIME PERIOD FOR RESPONSE

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

AFFIRMED-IN-PART