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.
13/899,258	05/21/2013	Doddaballapur N. Jayasimha	026517-0068P	8903

34284 7590 07/12/2017

Rutan & Tucker, LLP
611 ANTON BLVD
SUITE 1400
COSTA MESA, CA 92626

EXAMINER	
WANG, EDWARD	

ART UNIT	PAPER NUMBER
2139	

NOTIFICATION DATE	DELIVERY MODE
07/12/2017	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):

patents@rutan.com

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

Ex parte DODDABALLAPUR N. JAYASIMHA
and DREW E. WINGARD

Appeal 2017-004585
Application 13/899,258
Technology Center 2100

Before JOHN A. JEFFERY, DENISE M. POTHIER, and LINZY T. MCCARTNEY, *Administrative Patent Judges*.

POTHIER, *Administrative Patent Judge*.

DECISION ON APPEAL

STATEMENT OF THE CASE

Appellants appeal under 35 U.S.C. § 134(a) from the Examiner's Final Rejection of claims 1–20. *See* App. Br. 1, 5.[1] We have jurisdiction under 35 U.S.C. § 6(b). We affirm.

---

[1] Throughout this Opinion, we refer to (1) the Final Office Action (Final Act.) mailed October 7, 2015, (2) the Appeal Brief (App. Br.) filed March 7, 2016, (3) the Examiner's Answer (Ans.) mailed November 25, 2016, and (4) the Reply Brief (Reply Br.) filed January 25, 2017.

The Invention

Appellants' invention concerns a technique for "[m]aintaining cache coherence in a System-on-a-Chip with both multiple cache coherent master IP cores (CCMs) and non-cache coherent master IP cores (NCMs)." Spec., Abstract.

Claim 1 is reproduced below with the disputed limitations emphasized:

> 1. An apparatus, comprising:
>
> *a plug-in cache coherence manager, coherence logic in one or more agents, and an interconnect for a System on a Chip are configured to provide a scalable cache coherence scheme for the System on a Chip that scales to an amount of cache coherent master intellectual property cores in the System on a Chip*, where the plug-in cache coherence manager and coherence logic maintain consistency of memory data stored in one or more local memory caches including a first local memory cache for a first cache coherent master intellectual property core and a second local memory cache for a second cache-coherent master intellectual property core, *where two or more master intellectual property cores including the first and second intellectual property cores are configured to send read or write communication transactions over the interconnect to an IP target memory core, as well as a third intellectual property core in the System on a Chip that is a non-cache-coherent master intellectual property core, which is also configured send read or write communication transactions over the interconnect to the IP target memory core.*

The Examiner relies on the following as evidence of unpatentability:

Kinter	US 2009/0083493 A1	Mar. 26, 2009

Jari Nurmi & Drew Wingard, *Interconnect-Centric Design for Advanced SOC and NOC - Socket-Based Design Using Decoupled Interconnects* (Chapter 14) 1–30 (2002) (hereafter "Wingard").

The Rejections

Claims 1–3, 5–15, and 17–20 are rejected under 35 U.S.C. § 102(a)(1) as anticipated by Kinter. Final Act. 2–8.

Claims 4 and 16 are rejected under 35 U.S.C. § 103(a) as unpatentable over Kinter and Wingard. Final Act. 8–10.

THE ANTICIPATION REJECTION OVER KINTER

Regarding independent claim 1, the Examiner finds Kinter teaches all the elements of claim 1. Final Act. 2. Appellants argue Kinter fails to disclose the recited (1) coherence logic, (2) interconnect for a System on a Chip, and (3) plug-in cache coherence manager providing "scalable coherence scheme for the System on a Chip that scales to an amount of cache coherent master intellectual property cores in the System on a Chip." App. Br. 19–27. Appellants also assert Kinter fails to disclose two or more master intellectual property (IP) cores and a non-cache-coherent IP target memory core are configured to send read or write communication transactions over the same interconnect to an IP target memory core. App. Br. 27–28. Appellants further contends Kinter does not disclose a cache coherence manager and coherence logic "that scales to an amount of cache coherent master intellectual property cores in the System on a Chip" as recited in claim 1. App. Br. 29–31.

ISSUES

Under § 102, have Appellants shown the Examiner erred in rejecting claim 1 by finding Kinter discloses:

(I) "a plug-in cache coherence manager, coherence logic . . . , and an interconnect for a System on a Chip are configured to provide a scalable cache coherence scheme for the System on a Chip that scales to an amount of cache coherent master intellectual property cores in the System on a Chip" and

(II)

> two or more master intellectual property cores including the first and second intellectual property cores are configured to send read or write communication transactions over the interconnect to an IP target memory core, as well as a third intellectual property core in the System on a Chip that is a non-cache-coherent master intellectual property core, which is also configured send read or write communication transactions over the interconnect to the IP target memory core[?]

ANALYSIS

Claims 1, 3, 5, 7–9, and 12[2]

I. "[A] plug-in cache coherence manager, coherence logic . . . , and an interconnect for a System on a Chip . . . configured to provide a scalable cache coherence scheme for the System on a Chip that scales to an amount of cache coherent master intellectual property cores in the System on a Chip"

A. *A Plug-In Cache Coherence Manager*

Appellants generally assert "the Office Action fails to disclose a *plug-in* cache coherence manager." App. Br. 19. Appellants discuss "[t]he law

[2] Claims 3, 5, 7–9, and 12 are not separately argued. *See* App. Br. 19–31. We select independent claim 1 as representative. 37 C.F.R. § 41.37(c)(1)(iv).

instructs the claim limitation plug-in must be given some meaning when being construed." App. Br. 24.

The Examiner maps cache coherence manager 200 in Kinter to the recited "plug-in cache coherence manager." Final Act. 2 (citing Kinter ¶¶ 36, 47, Fig. 2A). In construing the term "plug-in" within the phrase "plug-in cache coherence manager" of claim 1, the Examiner determines Kinter's cache coherence manager 200 corresponds to "a plug-in cache coherence manager" as recited, because such a manager is swappable within the multi-core environments shown in the Figures 3, 8, and 9 embodiments. *See* Ans. 4. We determine the Examiner's determination is reasonable.

The Specification does not define the term "plug-in cache coherence manager," but discusses that "[t]he plug-in cache coheren[ce] manager maintains the consistency of instances of instructional operands stored in the memory IP target core and each local cache of the memory." Spec. ¶ 22. Although this discussion in the Specification informs our construction of the disputed "plug-in cache coherence manager," we decline to import such restrictions into claim 1, which fails to recite the plug-in cache coherence manager perform this specific function.

Turning to the figures, Figures 1 and 5 of the disclosure are described as illustrations of a snoop-based cache coherence manager but not a "plug-in" cache coherence manager. Spec. ¶¶ 8, 20, Figs. 1, 5. On the other hand, "CM," shown in Figure 4 as "CM B," is described as an illustration of a plug-in cache coherence manager. Spec. ¶¶ 19, 45, Fig. 4. However, this Figure and its accompanying discussion do not define the phrase, "plug-in," such that the term imparts any particular meaning or structure into the recitation "plug-in cache coherence manager" of claim 1.

A relevant technical dictionary defines "plug in" as "[a]n accessory program that enhances a main application." *Plug-in*, *ComputerUser High Tech Dictionary, available at* http://www.computeruser.com/dictionary (last visited July 23, 2017). We therefore construe the phrase "plug-in cache coherent manager" to include a cache coherent manager that has an accessory program for enhancing the cache coherent manager. Figures 3, 8, and 9 in Kinter show cache coherence manager 200 having various units that add functionality to manager 200 (e.g., enhances the manager with accessory programming), including maintaining consistency of data stored on local cache as recited through reads/writes exiting and entering units 205, 210, 215. *See* Kinter, Figs. 3, 8, and 9 (items 205, 210, 215, 220), *cited in* Ans. 4; *see also* Kinter ¶¶ 34 (discussing cache coherency information used to synchronize data with other caches), 38–40, 58–59. Appellants do not rebut the Examiner's explanation or distinguish sufficiently Kinter's cache coherence manager 200 from the recited "plug-in cache coherence manager." *See generally* App. Br., Reply Br.

Based on the evidence of record, the Examiner did not err in determining Kinter discloses the disputed "plug-in cache coherence manager" in claim 1.

B. *An Interconnect for a System on a Chip*

Appellants also assert that "the Office Action fails to inherently disclose the components of Figure 2A in Kinter to the . . . 'interconnect for a System on a Chip' recited in Claim 1." App. Br. 19. More specifically, Appellants assert that the rejection does not specify clearly "which one of the components of Figure 2A is the [recited] 'interconnect for a System on a

Chip'" and "fails to meet the evidence requirement to inherently disclose an interconnect with the limitations as recited in Claim 1." App. Br. 27.

We disagree. The Examiner maps this recited "interconnect" to elements 225, 245, 240, and 265, which are located in Figures 3, 8, and 9—not Figure 2A as asserted. Final Act. 2 (citing Kinter ¶¶ 38 (describing Figure 3), 65, 69). The Examiner elaborates in the Examiner's Answer that the "interconnect for the System on a Chip" comprises "the chip's bus and ports which interconnect[] all the components of the system to the System on a Chip." Ans. 4–5. For example, Kinter shows ports 225, 245, 240, and 265 in Figure 3 and bus 30 in Figures 2A and 5. *See* Kinter, Figs. 2A, 3, and 5. We determine these components reasonably correspond to "an interconnect" as recited.

Regarding the disputed "System on a Chip" component of the recited "interconnect for a System on a Chip," Appellants assert that Kinter's Figures do not show a System on a Chip (SoC) with a communication interconnect between the master cores and target memory core, but rather show only a cache coherence manager. App. Br. 27. Notably, claim 1 does not recite an interconnect on, or that is a part of, a SoC but rather an interconnect "for a System on a Chip." As such, claim 1 does not positively recite a SoC.

In any event, Kinter teaches "the above described embodiments of the processors of the present invention," including processor 100 shown in Figures 2A, 4, and 5, "may be represented as . . . computer-usable programs and data files . . . . Such programs and data files may be . . . used to integrate embodiments of the invention with other components to form . . . system on a chip (SoC)." Kinter ¶ 65. As such, although Kinter does not

show a SoC with a reference numeral in Figure 2A (*see* App. Br. 19), Kinter teaches that processor 100 shown in Figures 2A, 4, and 5 can be represented as programs and data files that are integrated with other components to form a SoC. *See id.* Kinter further shows in Figures 2A processor 100 may consist of cores $105_{1\text{-}4}$ and cache coherent manager 200 and in Figure 3 that ports 225, 245, 240, and 265 are between cores and cache coherence manager 200. *See id.*, Fig. 2A and 3. Kinter, thus, teaches an embodiment where processor 100, which consists of cores $105_{1\text{-}4}$, cache coherent manager 200, and ports 225, 245, 240, and 265, is integrated to form a SoC.

Thus, the mapped interconnect reasonably corresponds to "an interconnect for a System on a Chip" recited in claim 1.

C. *Plug-in cache coherence manager and coherence logic are configured to provide a scalable cache coherence scheme for the SoC that scales to an amount of cache coherent master IP cores in the SoC*

Next, Appellants argue that the rejection fails to disclose the plug-in cache coherence manager and coherence logic "are configured to provide a scalable cache coherence scheme for the System on a Chip that scales to an amount of cache coherent master intellectual property cores in the System on a Chip." App. Br. 29. Specifically, Appellants dispute that Kinter's paragraphs 36 and 39 do not discuss the scaling concept. App. Br. 29–30.

The Examiner responds in the Answer, explaining Kinter discloses its coherence scheme scales to 2, 3, 4, or even more cache coherent cores. Ans. 5 (citing Kinter ¶ 32). Kinter teaches processor 100, which includes cache coherence manager 200 as shown in Figures 2A and 3 in communication with other system components, including logic, has four cores but "may include more or fewer than four cores." Kinter ¶ 32. We, thus, agree with the Examiner that Kinter demonstrates "a plug-in cache coherence manager"

and "coherence logic" "are configured to provide a scalable cache coherence scheme . . . that scales to an amount of cache coherent master intellectual property cores" as argued and recited in claim 1.

Lastly, Appellants state the Examiner "is mistaken" that claim 1 does not recite a scheme that scales "the amount of coherent cache to an amount of cache the master IP cores in the System on a Chip have." App. Br. 30 (emphasis omitted) (citing Final Act. 10). We are not persuaded. Notably, the Examiner was pointing out claim 1 does not specifically recite which cache coherent master IP cores are scaled; instead, claim 1 recites an unspecified "amount of cache coherent master intellectual property cores." Additionally, any argument related to features not recited, such as "a System on a Chip having multiple master IP cores such that scales [sic] the amount of coherent cache to an amount of cache the master IP cores in the System on a Chip have" (App. Br. 30), are unavailing. *See* Ans. 5.

II.

"two or more master intellectual property cores including the first and second intellectual property cores are configured to send read or write communication transactions over the interconnect to an IP target memory core" and "a third intellectual property core in the System on a Chip that is a non-cache-coherent master intellectual property core, which is also configured send read or write communication transactions over the interconnect to the IP target memory core"

In the Response to Arguments section of the Final Office Action, the Examiner states that Kinter describes including the invention in an IP core. Final Act. 10 (citing Kinter ¶ 69). Appellants conclude that (1) this statement made by the Examiner logically means the entire device in Kinter as presented in the rejection is a single core, and (2) there is no structure/module left in Kinter to map to "the additional three IP cores of the

two [cache coherent] intellectual property cores and the non-cache coherent] [sic] intellectual property cores sending communication over the interconnect to the memory IP core." App. Br. 24 (emphasis and underlying omitted) (bracketing other than [sic] in the original). Essentially, Appellants "assert that if all of the components of the Kinter['s] invention make up the [single] IP core of assumedly the plug-in cache coherence manager and the associated coherence logic, then the other components cannot make up the additional three IP cores." App. Br. 23–24 (second set of brackets in original); *see also* App. Br. 25.

We are not persuaded. As the Examiner notes, the statements on page 10 of the Final Office Action respond to Appellants' argument that Kinter fails to disclose a SoC or an interconnect for a SoC. Ans. 3–4 (stating the "response has no relation to the mapping of the IP cores to the claim"); *compare* Final Act. 10 *with* Final Act. 2. To be clear, the processing core 110_1, 110_2, and 460 in Kinter are respectively mapped to the recited "first cache coherent master intellectual property core," the "second cache-coherent master intellectual property core," and the "third intellectual property core" of claim 1 in the rejection. Final Act. 2 (citing Kinter ¶¶ 34, 54, Figs. 2A, 5 and referring to Processing Core in Fig. 2A and Processor in Fig. 5). Given the rejection's mapping, the entire device in Kinter's Figure 2A or 5 does not correspond to only one of the claimed cores, such that there is no structure left to correspond to the other recited cores in claim 1 as argued. *See* App. Br. 22–24.

Furthermore, based on the rejection's mapping, we disagree with Appellants that the Examiner ignores words in the claim as asserted. Reply Br. 6. Notably, claim 1 positively recites an apparatus with three elements:

(1) a plug-in cache coherence manager, (2) coherence logic in one or more agents, and (2) an interconnect for a SOC. App. Br. 38 (Claims App'x). The remaining portions of claim 1 recite functional limitations or the functions these three elements are capable of performing (e.g., "configured to . . . .") For example, claim 1 recites the functional limitation of

> the plug-in cache coherence manager and coherence logic maintain consistency of memory data stored in one or more local memory caches including a first local memory cache *for* a first cache coherent master intellectual property core and a second local memory cache for a second cache-coherent master intellectual property core.

Id. (emphasis added). The further limitations in claim 1 concerning the IP cores merely claim how these cores are configured, but these disputed IP cores are not positively recited. As such, claim 1 does not positively recite "the same interconnect connecting to a memory IP core (e.g., 'system memory') also connects to the coherent master IP cores as well as non-coherent master IP cores" as asserted. Reply Br. 6.

Appellants further contend that "none of the Figures of Kinter" show "a communication interconnect between the master cores and the target memory core but rather show merely a cache coherency manager." App. Br. 27; *see also* Reply Br. 7–14. Appellants argue Kinter's Abstract clarifies that traffic flows within a coherence manager only among processing cores that have a coherency relationship, but fails to disclose two or more master intellectual property cores are configured to send read or write communication transactions over the same interconnect to the IP target memory core as a non-cache-coherent IP target memory core that is also configured to send read or write communication transactions to the IP target memory core. App. Br. 27–28. For example, Appellants "assert that it is

entirely possible and probable that Kinter's non-cache-coherent processor 460 uses a completely separate communication bus interconnect . . . than the Processing Cores 110_1, and 110_2." App. Br. 29.

We are not persuaded. Courts have "repeatedly emphasized that an indefinite article 'a' or 'an' in patent parlance carries the meaning of 'one or more' in open-ended claims containing the transitional phrase 'comprising.'" *KCJ Corp. v. Kinetic Concepts, Inc.*, 223 F.3d 1351, 1356 (Fed. Cir. 2000); *see also In re Varma*, 816 F.3d 1352. 1362–63 (Fed. Cir. 2016). On the other hand, "when the claim language and specification indicate that 'a' means one and only one, it is appropriate to construe it as such even in the content of an open-ended 'comprising' claim." *Harari v. Lee*, 656 F.3d 1331, 1341 (Fed. Cir. 2011).

Here, the Specification shows or indicates at least two components (e.g., coherent control fabric and dataflow fabric) may be considered to be the recited "interconnect,"[3] as well as other intervening elements (e.g., TA, RD/WR Ch) between the cache coherent master IP cores and the IP target memory core. Spec., Fig. 1. As such, the recited "interconnect" does mean only one or a single interconnect. Rather, the interconnect can comprise multiple components and does not require a direct connection or

[3] Appellants map the recited "interconnect" in Figure 1 to the box with rounded edges. Reply Br. 8. However, we note that the disclosure does not describe this box as an interconnect. Rather, the disclosure states an "interconnect is composed of 1) data flow bus fabric separate from 2) its coherence command and signaling fabric that couples to a flexible implementation of a cache coherence manager." Spec. ¶ 26; *see also id.* ¶ 83 and claim 2. Although this is not a definition of the term, this informs our construction that an interconnect is not all the components identified in the annotated Figure 1 in the Reply Brief.

communication between the IP target memory core and the first cache coherent master IP core, the second cache coherent master IP core, and the third non-cache-coherent master IP core, all three of which are all configured to send read or write communications over the interconnect.

Nor is claim 1 restricted such that the components of the recited "interconnect" cannot have intervening elements. That is, as broadly as recited and as explained above, "an interconnect" is not limited to a single interconnect for receiving read or write communications to the IP target memory core. *See* Ans. 5 (stating the claim "does not preclude the interpretation where the separate communication bus interconnect is part of the interconnect"). Thus, much of Appellants' argument related to cache coherence manager 200 physically and logically separating the cache coherent domains using different buses is not availing. *See* Reply Br. 9–13.

Also, the Examiner states "the interconnect . . . essentially comprises the chip's bus and ports[,] which interconnects all the components of the system to the System on Chip." Ans. 4–5; *see also* Final Act. 2. Thus, even presuming, without agreeing, processor 460 in Figure 5 (e.g., a processing core) sends read or write communications to the IP target memory core with a separate bus interconnect than the IP cores $110_{1,2}$, this does not sufficiently demonstrate that ports 225, 245, 240, 265, and 255 and related buses as shown in Figure 5, for example, cannot collectively correspond to the recited "interconnect" of claim 1.

Lastly, Appellants argue Kinter does not disclose "coherence logic in the agents . . . that is configured to facilitate both the cache coherent & non-cache-coherent master intellectual property cores . . . ." App. Br. 28 (emphasis omitted). We are not persuaded. Claim 1 does not recite the

coherence logic *facilitates* any component or function of claim 1, let alone both cache coherent master IP cores and/or non-cache-coherent master IP cores as argued. Rather, claim 1 recites the coherence logic is a component "configured to provide a scalable cache coherence scheme" and is a component to "maintain consistency of memory data stored in one or more local memory caches." For reasons previously discussed, Kinter teaches this feature of claim 1.

Based on the record before us, we find no error in the Examiner's rejection of (1) representative claim 1 and (2) independent claim 12[4] and dependent claims 3, 5, and 7–9, which are not separately argued.

Claims 13–15 and 17–20

Although independent claim 13 varies in scope from independent claim 1 (App. Br. 38, 43 (Claims App'x)), Appellants refer to the reasons presented for claim 1 in asserting that independent claim 13 is patentable. App. Br. 32. To the extent claim 13 includes the disputed limitations, we are not persuaded for the previously stated reasons.

Accordingly, we find no error in the Examiner's rejection of independent claim 13 and claims 14, 15, and 17–20, which depend directly or indirectly from claim 13 and are not separately argued.

[4] Notably, independent claim 12 differs in scope from independent claim 1. Claim 12 recites a "non-transitory computer readable medium containing instructions . . . configured to cause [a] machine to generate a *software representation* of the apparatus of claim 1." App. Br. 42 (Claims App'x) (emphasis added).

Claims 2, 10, and 11

Claim 2 recites

> the interconnect is composed of 1) a data flow bus fabric separate from 2) its coherence command and signaling fabric . . . , where the coherence command and signaling fabric is configured to convey signaling and commands to maintain the system cache coherence scheme and where the data flow bus fabric is configured to carry non-coherent traffic and all traffic transfers between the three or more master intellectual property cores and the IP target memory core . . . .

App. Br. 38–39 (Claims App'x). According to Appellants, the Examiner maps ports 225, 240 to the coherence command and signaling fabric, but that port 225 receives both coherent and non-coherent read/write requests. App. Br. 32. In Appellants' view, claim 2 requires separate and different ports to receive coherent and non-coherent requests and contrasts with Kinter's teaching. *See id.* at 32–33.

Claim 2 recites "the coherence command and signaling fabric is configured to convey signaling and commands to maintain the system cache coherence scheme." There is no limitation related to the coherence command and signaling fabric configured to convey only coherent requests. Rather, the fabric is configured to convey unspecified signals and commands to maintain the system cache coherence scheme "that scales to an amount of cache coherent master intellectual property cores" as required by claim 1. *See* Ans. 6 (noting a similar construction for claim 2).

Turning to the rejection, the Examiner maps the recited "data flow bus fabric" to elements 245, 265 in Kinter and the recited "coherence command and signaling fabric" to elements 225 and 240 in Kinter. *See* Final Act. 2–3 (citing Kinter ¶¶ 38–39). Kinter teaches port 225 receives read and write requests from cores 105 and port 240 receives intervention messages issued

to the cores that are to remain coherent with the requesting core. Kinter ¶¶ 38–39. As broadly as recited, these requests include "signaling and commands" such that these ports are reasonably a "coherence command and signaling fabric . . . configured to convey signaling and commands to maintain the system cache coherence scheme" as recited in claim 1. Granted, Kinter further teaches requests are received by request unit 205 from cores 105_i and Kinter shows non-coherent requests exiting request unit 205, suggesting port 225 also receives non-coherent requests. *See* Kinter ¶ 39, Fig. 3. Yet, claim 2 does limit such non-coherent traffic exclusively to the "data flow bus fabric" and does not exclude such traffic from being conveyed using the coherence command and signaling fabric.

Additionally, Kinter teaches port 245 outputs data stored in the cache memory of a core that is requested by another core and port 265 forwards requested data to the requesting core. Kinter ¶ 39, Fig. 3. Kinter further shows coherent read/write requests exiting intervention unit 210 in response to coherent intervention response request from unit 205. *Id.* Thus, although Kinter shows coherent—as opposed to non-coherent traffic—exiting intervention unit 210, Kinter does not exclude non-coherent traffic from being carried on port 245. Moreover, claim 2 recites that the "data flow bus fabric is configured to carry non-coherent traffic" or is capable of carrying non-coherent traffic. As noted above, port 245 is capable of carrying various types of traffic, including the "non-coherent traffic" as well as "all data traffic transfers" recited in claim 2.

Based on the evidence of record, we find no error in the Examiner's rejection of claim 2 and dependent claims 10 and 11, which are not separately argued.

Claim 6

Claim 6 depends from claim 3 and further recites a "snoop filter-based cache coherence manager performs . . . periodic snooping to check on a state on cache coherent data in each local cache." App. Br. 40 (Claims App'x). Appellants argue Kinter does not teach the recited "periodic snooping" feature of claim 6. App. Br. 34.

Notably, the rejection cites to different passages in Kinter (Final Act. 4 (citing Kinter ¶¶ 56, 58)) than Appellants discuss fails to disclose "periodic snooping" (App. Br. 34 (discussing Kinter ¶ 39)). Kinter teaches a core or unit issues requests, which causes an intervention message to be sent to all or some of the cores to establish the multitude of coherence domains, to ensure coherency is maintained, and to reduce the number of cache lookup. Kinter ¶¶ 56–57. Kinter further discusses "a hit resulting from a search in L1 cache tag array" will cause an intervention message to be sent. *Id.* ¶ 58. These teachings in Kinter disclose that intervention messages (e.g., snooping messages) are created or occur to establish the coherence domains from time to time. *See id.* ¶¶ 56, 58.

In the Response to Argument section of the Examiner's Answer, the Examiner explains that "periodic" can mean both "occurring at regular intervals of time" and "happening repeatedly." Ans. 7 (citing to *Merriam-Webster Dictionary*). The Examiner states Kinter teaches the snooping (e.g., the intervention message) is being repeatedly sent. Appellants do not dispute this understanding of "periodic" or that Kinter teaches sending intervention messages repeatedly. *See generally* Reply Br.

Additionally, an alternative understanding of "periodic" is "[a]ppearing or occurring at intervals." *Periodic*, Oxford Dictionaries,

*available a*t https://en.oxforddictionaries.com/definition/us/periodic (def. 1) (last visited July 23, 2017). As such, the record establishes a broad, but reasonable understanding of "periodic snooping" in claim 6 to include snooping at (1) intervals, regular or not, or (2) repeatedly. Given the above discussion of Kinter, we agree with the Examiner that Kinter teaches the intervention messages (e.g., snooping messages) appear or occur at intervals (e.g., periodically), such as from time to time. Kinter ¶¶ 56–58.

Given the evidence in the record, we find no error in the Examiner determining Kinter discloses performing "periodic snooping" feature of claim 6.

In conclusion, Appellants have not persuaded us of error in the rejection of claims 1–3, 5–15, and 17–20.

OBVIOUSNESS REJECTION OVER KINTER AND WINGARD

Claims 4 and 16 are rejected under 35 U.S.C. § 103 over Kinter and Wingard. Final Act. 8–10. Appellants repeat that Kinter does not disclose the limitation of claim 1, from which claim 4 depends indirectly. App. Br. 35. For the above-explained reasons, we are not persuaded. Thus, we need not consider whether Wingard fails to remedy any purported deficiency of Kinter. *See id.*

For the foregoing reasons, we sustain the rejection of claims 4 and 16.

DECISION

We affirm the Examiner's rejections of (1) claims 1–3, 5–15, and 17–20 under § 102 and (2) claims 4 and 16 under § 103.

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