

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

ADVANCED MICRO DEVICES, INC.,
Petitioner,

v.

REALTEK SEMICONDUCTOR CORPORATION,
Patent Owner.

Case No. IPR2023-00920
U.S. Patent No. 8,006,218

**DECLARATION OF VIVEK SUBRAMANIAN IN SUPPORT OF
PETITIONER'S REPLY TO PATENT OWNER'S RESPONSE**

I, Vivek Subramanian, hereby declare as follows:

I. INTRODUCTION

1. I have prepared this declaration in connection with Advanced Micro Devices, Inc.’s Petition for *Inter Partes* Review of U.S. Patent No. 8,006,218 (“the ’218 Patent”).

2. I previously prepared declarations in support of Advanced Micro Devices, Inc.’s petition in this case. EX1003, EX1034. The documents I reviewed, my compensation and my qualifications are set forth in my first declaration. EX1003 ¶¶3-13.

3. I have been asked to provide my opinions on a few limited points, and have done so below.

II. “SEPARATELY”


4. Claim 12 of the ’218 Patent recites, in part “a global power mesh, coupling to the first local power mesh and the second local power mesh, to provide the power needed by the first power domain and the second domain *separately*.” EX1001 at 7:63-66. I have been informed that Realtek construes the term “separately” to mean “through separate electrical paths,” and I have been asked to apply that meaning of “separately” for the purposes of my analysis in this declaration.

5. With this understanding in mind, I have been asked to opine on whether a POSITA would understand there to be any further restrictions on the provision of power “separately,” such as restricting “separately” to separate electrical paths within the “global power mesh.” I see no support for any such restriction.

6. The language of the claim 12 limitation at issue simply states “a global power mesh, coupling to the first local power mesh and the second local power mesh, to provide the power needed by the first power domain and the second domain *separately*,” EX1001 at 7:63-66, without any further restrictions on how providing the power needed “separately” must occur. I also do not see any further restrictions on the “separately” term elsewhere in claim 12 or anywhere in the ’218 Patent. *Id.* at 7:56-8:11; *see also generally id.* Absent any such restriction, one of skill would understand several different ways for providing the power needed “separately” through separate electrical paths, including separate electrical paths within a global power mesh, separate electrical paths going from the global power mesh, and separate partial local power meshes between power domains and a global power mesh. This is further supported by the ’218 Patent’s specification which, discloses separate electrical paths within a global power mesh, separate electrical paths going from the global power mesh, and separate partial local power meshes between the power domains and the global power mesh. EX1001 at Figs. 3-6, 4:49-52 (“As shown in FIG. 5, the global power mesh 400 is placed above the local power mesh

300 and the partial local power meshes 310, 320 of the local power mesh are placed *separately* above each power domain.” (emphasis added)), 5:8-20 (explaining “power supply wire 401 that couples to the external power V_{DD1} uses via (not shown in the figure) to couple to the power wire 311 located below so as to conduct the external power V_{DD1} to the chip circuitry” and “power supply wire 403 that couples to the external power V_{DD2} uses via to couple to the power wire 321 located below so as to conduct the external power V_{DD2} to the chip circuitry”).

7. I previously opined that Jiang discloses this limitation. EX1003 ¶¶78-85. Adopting Realtek’s construction of “separately” as using “separate electrical paths” does not change that opinion. It is my opinion that Jiang shows “separate electrical paths.” Jiang shows a global power mesh with separate couplings to separate partial local power meshes as illustrated below.



EX1004, Fig. 1(b) (annotated). A POSITA would understand that these separate couplings to separate local power meshes are “separate electrical paths.” Indeed, Jiang discloses separate sleep transistors to couple the separate partial local power meshes (and their underlying power domains) to the global power mesh. EX1004 at 1 (“In the power-gated chip, we have local power meshes for all macros, and sleep transistors are used to connect the local mesh to the global mesh.”), Fig. 1(b). When the sleep transistors are on, power flows from Jiang’s global power mesh, through each of the sleep transistors and its respective partial local power mesh to the power domain underneath. *Id.* at 1-2, Fig. 1(b).

III. “SEPARATELY” IS AT LEAST OBVIOUS OVER JIANG

8. I previously testified that it would have been obvious to supply multiple independent voltage in the power meshes of Jiang and that a POSITA would have been motivated to do so to save power consumption on the chip, among other reasons. EX1003 ¶¶91-92 (citing EX1014, EX1012, EX1010, EX1005). I have been asked to clarify whether the power savings mentioned in my previous declaration would be in addition to the “47% leakage power saving” mentioned in Jiang. EX1004 at 1. A POSITA would have known that the potential power savings of using multiple independent voltages in the power meshes of Jiang would be separate and in addition to the “47% leakage power saving” mentioned by Jiang.

9. Jiang’s power savings due to avoiding leakage is achieved “by shutting off the idle blocks.” EX1004 at 1. When power is shut off to an idle block it will not experience the same quantity of leakage current as if the idle block were not performing function but still exposed to full Vdd. That is how Jiang achieves its power savings. However, using multiple voltages permits areas of the chip that need to operate to use less power *while operating*. As I explained in my previous declaration, powering portions of a chip at a lower voltage that do not require the same voltage as other portions of the chip because they are not timing-critical blocks or are not input-output circuits, for example, will save power. EX1014, 3 (“*In active state*, supply voltage can be controlled to reduce power, albeit in the limited ranges

allowed in aggressively scaled technologies. ... Similarly to the multiple-threshold scheme, timing-critical transistors can be powered at a high voltage, while most transistors are connected to the low voltage supply.” (emphasis added)); EX1012, 16-17 (“Since dynamic power is proportional to V_{DD}^2 , lowering V_{DD} on selected blocks helps reduce power significantly. Unfortunately, lowering the voltage also increases the delay of the gates in the design.”).

10. I have also been asked to clarify whether supplying multiple independent voltages would require Jiang’s domains to receive different voltage levels at all times. A POSITA would understand that supplying multiple independent voltages would allow, but does not require, Jiang’s power domains to receive different voltage levels. When appropriate, these multiple independent voltages could be set to the same level, thereby supplying Jiang’s power domains with the same voltage level.

IV. JIANG – CLAIM 15

11. I previously testified that it would have been obvious to a POSITA to route a ground connection through Jiang’s global power mesh. EX1003 ¶109. I have been informed that Realtek argues that Jiang’s Figure 2, below shows grounding in the substrate.




Figure 2. Power gating modeling

EX1004, Fig. 2. I have been asked to opine whether grounding in the substrate would dissuade one of skill from routing ground through Jiang’s global mesh. It is my opinion that substrate grounding would **not** dissuade one of skill from routing ground through Jiang’s global mesh.

12. One of skill would understand that simply showing ground in the substrate does not address or explain how the ground connection got to the substrate. Integrated circuits are typically wrapped in a “package.” The “package” electrically insulates the integrated circuit and requires that signals and power (including ground) be distributed into the integrated circuit through bumps or pins affixed to the package. The connections to those bumps or pins are routed to the lower portions of the integrated circuit through the interconnect layers across the top of the

integrated circuit (also known as the “interconnect stack”). For a ground connection to get to the substrate, it must be routed through the interconnect stack.

13. As I explained in my earlier declaration, the simplest and most common way for a ground connection to be routed to the substrate would be through the global mesh of the interconnect stack, as shown in the Li and Chiu references. *See* EX1003 ¶109; EX1005, 493-94, Fig. 3; EX1006, 2:50-61, Fig. 7; *see also* EX1018, 9:17-23. I see nothing in Jiang teaching how a ground connection is routed to the substrate in Jiang’s Figure 2, or any discussion of preferred or disfavored methods of routing ground to the substrate. *See generally* EX1004. Accordingly, substrate grounding would not dissuade one of skill from routing ground through the global mesh, as was well known in the art. Routing ground through the global mesh of the interconnect stack is typically the most efficient way to distribute ground on a chip because it can be routed on metal layers that are already being used to route power, requiring no extra allocation of layers.

V. A POSITA WOULD FIND IT OBVIOUS TO USE AN APR TOOL WITH JIANG

14. I previously testified that it would have been obvious for a POSITA to form and couple Jiang’s power meshes as shown in Jiang’s Figure 1(b) with an automatic placement and routing (APR) tool. EX1003 ¶¶126-129. I have been asked to opine whether the simplicity of the design of Jiang’s power routing structure would affect my opinion. It would not. A POSITA would understand that APR

tools would be useful in designing power routing structures of any complexity. In fact, the '218 patent itself describes a “simple algorithm” for an APR tool where:

the local power mesh 300 is evenly distributed on top of the entire chip circuitry. Then, the partial local power meshes in the local power mesh 300 corresponding to the power domains 110, 120 are defined (such as partitioning the local power mesh 300 based on the positions of the power domains 110, 120 so that the plurality of partitioned power meshes are not electrically coupled to each other) to form the above mentioned partial local power meshes 310, 320. On the other hand, the approach to form the global power mesh 400 is much easier for the APR tool by just evenly distributing the global power mesh 400 on metal layers that are positioned higher.

EX1001, 5:60-6:4. Jiang already discloses precisely this arrangement as I testified in my prior declaration, EX1003 ¶¶105-07 (partitioned local power mesh based on positions of the power domains), 140 (evenly distributed global power mesh on higher layers), and thus would have been a candidate for an APR tool as explained by the '218 Patent itself.

15. I have also been asked to opine on whether Jiang's teachings apply to actual chip designs. A POSITA would understand that Jiang teaches applying its power mesh design to actual chips. The study in Jiang performed real circuit analysis with regards to the benefits and costs of a power-gating design and determined ultimate benefits “from the power-gating technique in future technology nodes.” EX1004, 1. Indeed, Jiang explained that power-gating had a “promising future” and

that trends suggested “that in the future the low-power designs will favor the power-gating techniques.” *Id.* at 8.


16. Moreover, Jiang expressly included an experiment applying its power meshes to an actual chip design using “26 circuits (from MCNC and ISCAS benchmarks respectively) as the basic logic blocks.” EX1004, 6-7. In this experiment, Jiang “gated half of the circuits (blocks whose geometric centers were on the left half of the chip) and compared it with non-power-gated chip in terms of power and area” and then studied the effects of power-gating on this actual circuit design. *Id.* Such an experiment would further confirm a POSITA’s understanding that Jiang’s teachings were meant to be applied to actual chip designs.

VI. LI – CLAIM 16

17. I have been asked to opine whether all ground references within Li’s chip are connected to one another. It is my opinion that Li’s grounds are connected. Li discloses that “BCM1280 has several power supplies that all share the same chip ground reference.” EX1005 at 495. By teaching that its power supplies in the chip share “the same chip ground reference,” a POSITA would understand that all grounds within Li (when not turned off by, for example, a power switch) would be connected to one another through the same chip ground reference. This is further confirmed by Li’s disclosure that the “BCM1280 is designed with two power supply inputs and *one ground input* at the package level.” *Id.* at 493 (emphasis added).

Ultimately all ground supplies in Li's BCM1280 design must connect to the "one ground input" and thus to each other.

18. Li's Figure 4 showing two grounds also confirms my opinion. The first ground "Substrate GND" is the chip ground (GND) supplied through the global mesh in Figure 3. EX1005 at 494 ("BCM1280 has a total of four core grounds connected to the P-wells inside the four triple wells, respectively, and *one chip ground connected to the P-substrate.*" (emphasis added)). The second ground "core_gnd," which is representative of the four core grounds, is connected to the Substrate GND via a power switch when the switches are turned on as shown in Figure 4 below. A POSITA would understand that all grounds in Li's BCM1280 are interconnected when the power switches are turned on and that all ultimately connect to the GND supply in the global power mesh in Figure 3.



EX1005, Fig. 4.

VII. CONCLUSION

19. In signing this Declaration, I recognize that the Declaration will be filed as evidence in a contested case before the Patent Trial and Appeal Board of the

United States Patent and Trademark Office. I also recognize that I may be called upon to testify in deposition or before the Patent Trial and Appeal Board. If called upon to testify, I would testify truthfully and competently to all matters stated in this Declaration.

20. I reserve the right to supplement my opinions in the future to respond to any arguments that Patent Owner raises and to take into account new information as it becomes available to me.

21. I hereby declare that all statements made herein of my own knowledge are true, and that all statements made on information and belief are believed to be true and that such statements are made with the knowledge that willful false statements and the like are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Dated: May 24, 2024



Vivek Subramanian