

NOTE: This disposition is nonprecedential.

United States Court of Appeals for the Federal Circuit

APPLE INC.,
Appellant

v.

ZENTIAN LTD.,
Appellee

2024-2207

Appeal from the United States Patent and Trademark
Office, Patent Trial and Appeal Board in Nos. IPR2023-
00036, IPR2023-01195.

Decided: July 21, 2026

BRIAN ROBERT MATSUI, Morrison & Foerster LLP,
Washington, DC, argued for appellant. Also represented
by SETH W. LLOYD; ALEXANDRA M. AVVOCATO, New York,
NY; BITA RAHEBI, REBECCA WEIRES SETRAKIAN, Los Ange-
les, CA.

KATHERINE E. RHOADES, Bartlit Beck LLP, Chicago, IL,
argued for appellee. Also represented by JESSICA R.
BERNHARDT, NEVIN M. GEWERTZ; KAYVAN B. NOROOZI, No-
roozi PC, Los Angeles, CA.

Before DYK, MAYER, and PROST, *Circuit Judges*.

DYK, *Circuit Judge*.

Apple Inc. (“Apple”) appeals a decision of the Patent Trial and Appeal Board (the “Board”) concluding that claims 1, 2, 4, 6–9, 11, 24, and 29 of U.S. Patent No. 10,839,789 (the ’789 patent) were not shown to be unpatentable. We vacate the Board’s decision and remand.

BACKGROUND

I

Zentian Ltd. (“Zentian”) owns the ’789 patent, which is directed toward systems, circuits, and coprocessors used in speech recognition. Claims 1 and 29 are representative for purposes of this appeal. Claim 1 recites:

An acoustic coprocessor for processing data associated with an audio signal, comprising:

a first interface for receiving at least one feature vector, wherein the feature vector is determined from the audio signal;

an acoustic model memory for storing an acoustic model defining a plurality of acoustic states;

a calculating apparatus for calculating distances indicating a similarity between the at least one feature vector and respective acoustic states of the acoustic model read from the acoustic model memory; and

a second interface for sending at least one distance calculated by the calculating apparatus;

wherein the calculating apparatus and the acoustic model memory are fabricated on a single integrated circuit.

'789 patent, claim 1 (emphases added). Claim 29 claims a “speech recognition system” wherein an “acoustic coprocessor” and “at least one [search] processor are fabricated on a single integrated circuit,” and wherein the acoustic coprocessor contains “one or more result memories,” “a status register,” and “a direct memory access (DMA) controller capable of copying [data] from the one or more result memories to a buffer memory.” *See id.* claim 29 (depending from claims 18, 25, and 26).

The speech recognition circuit analyzes speech sounds by breaking down an incoming digital audio stream into short “frames,” or segments of the audio stream. *Id.* col. 12 ll. 55–57. The audio in each frame is then converted into a “feature vector,” which represents the various characteristics of the audio. *Id.* col. 12 ll. 58–65. A prototypical feature vector is a series of coordinates that “represents a point in an N-dimensional space,” where each dimension corresponds to an audio characteristic, such as a particular audio frequency band. *Id.* col. 12 ll. 58–65, col. 13 ll. 22–26. Each of these feature vectors is then compared against an “acoustic model,” which is a library of many acoustic “states” representing known and expected sounds in speech. *See id.* col. 4 ll. 12–17. In a preferred embodiment, each state contains a probability distribution mapped onto the same N-dimensional space as the feature vectors. *Id.* col. 13 ll. 26–40. This allows a “distance” calculation to be performed between the feature vector and each state; the shorter the distance, the more likely the audio represented by the feature vector is the known speech sound represented by the state. *See id.* col. 14 ll. 1–10. Once likely sounds have been identified, a “search stage” identifies likely words reflecting what was said in the digital audio stream. *Id.* col. 6 ll. 6–9.

Three disclosed memories perform different functions. First, the “acoustic model memory” holds the acoustic model to which feature vectors are compared. *E.g., id.* col. 25 ll. 42–56. Second, the “results memories . . . stor[e]

[the] calculated distances” between the feature vectors and the states of the acoustic model. *Id.* col. 25 ll. 58–61. Third, the “buffer memory” is additional memory space for holding the calculated distances if the results memories require additional storage space. *Id.* col. 6 ll. 9–15. Claim 29 also requires a “status register” that “indicate[s] the completion of calculation of [a particular] set of distances.” *Id.* (depending from claim 25).

Relevant to this appeal, both claims 1 and 29 require “a single integrated circuit.” *Id.* Claim 1 requires that the integrated circuit contain both the claimed acoustic model memory and the claimed “calculating apparatus” (i.e., the processor that calculates distances between feature vectors and the states of the acoustic model). Claim 29 requires that the integrated circuit contain both the claimed “acoustic coprocessor” (also referring to a processor that calculates distances) and the claimed “at least one [search] processor” (a separate processor that analyzes the calculated distances to identify likely words). *Id.* (depending from claims 10 and 28). The acoustic coprocessor of claim 29 contains “one or more result memories” and a “direct memory access (DMA) controller” that connects to the buffer memory. *Id.* (depending from claims 18 and 26). Claim 29 also requires an acoustic model memory, but it is not required to be part of the integrated circuit.

II

Apple petitioned for *inter partes* review of claims 1, 2, 4, 6–14, 16–18, 20–29, 35, 37–39, and 42–45 of the ’789 patent. In the petition, Apple argued that claim 1 was obvious over prior art reference U.S. Patent No. 5,819,222 (“Smyth”) or the combination of Smyth and U.S. Patent No. 6,832,194 (“Mozer”) and that claim 29 was obvious over Smyth, Mozer, and several other references.

A

Smyth teaches a similar speech recognition circuit wherein a digital audio stream is broken into frames, each frame is converted into a feature vector, and each feature vector is compared with a set of acoustics states, contained in the “state memory,” to determine their similarity for the purpose of eventually identifying likely words. Smyth, col. 5 ll. 19–67. Smyth’s “state memory” is a component of its “classifier,” which also contains a “classifying processor” that calculates the distances between feature vectors and states in the state memory. *Id.* col. 5 ll. 55–57. As Smyth describes, the classifying processor could be “a single suitably programmed digital signal processor (DSP) device (such as the Motorola DSP 56000 . . .) or similar device.” *See id.* col. 5 ll. 49–53, col. 6 ll. 5–9. Smyth also discloses a “sequencer,” which takes the calculated distances and uses them to match the audio frames to likely words. *Id.* col. 6 ll. 10–12, 35–42.

Mozer teaches another similar speech recognition circuit wherein the feature vector being analyzed is stored in a memory on the integrated circuit along with a “template vector,” which functions similarly to a single state in an acoustic model. Mozer, col. 7 ll. 53–58, col. 9 l. 22–col. 10 l. 20.

B

Smyth alone and the Smyth and Mozer combination disclose most of the claim limitations. The parties’ dispute concerns the “acoustic model memory” and “integrated circuit” limitations of claims 1 and 29. As to claim 1, the petition argued that Smyth’s state memory and disclosure of the Motorola DSP56000 disclosed the claimed acoustic model memory on an integrated circuit with a calculating apparatus. Alternatively, the petition argued that the combination of Smyth and Mozer disclosed the acoustic model memory on an integrated circuit with a calculating apparatus.

In response, Zentian raised two separate arguments. First, as to Smyth alone, Zentian argued that the Motorola DSP56000 could not hold Smyth's large-vocabulary model, as Apple's expert testified that the "typical size" of such a model would be "around 4,700 bytes," J.A. 3804,¹ whereas the undisputed testimony was that the Motorola DSP56000 only contained around 1,500 bytes of available memory. In reply, Apple argued that the claims did not require a large-vocabulary model and that an "abbreviated model," such as one that is "configured to recognise certain phrases or words for example a string of digits," would be much smaller than a complete model. J.A. 368 (quoting Smyth, col. 6 ll. 61–66). Apple argued that the Smyth/Mozer combination, which disclosed a larger 4,000-byte memory, also disclosed this claim limitation. As pointed out in the reply, Zentian's expert testified in deposition that a model for something "like a digits recognizer or something . . . like, I think, Smith [sic] was doing" would be an "abbreviated model" and could be "very small," J.A. 3380.

Second, as to the Smyth/Mozer combination, Zentian argued that the claims required fabrication of the integrated circuit and that a person of ordinary skill in the art was not capable of fabricating a new circuit. On reply, Apple argued that the patent does not require the skilled artisan to perform the step of fabricating the circuit and that the skilled artisan could simply obtain "known already fabricated integrated circuits" with the appropriate specifications. J.A. 387–88.

Implicitly construing the claimed acoustic model memory to require holding a large-vocabulary model, the Board found that Apple had failed to show that Smyth

¹ In contrast, Zentian's expert testified that an acoustic model could "require more than 720 kBytes [720,000 bytes] of memory." J.A. 3846.

alone disclosed the claimed acoustic model memory on an integrated circuit because the Motorola DSP56000 only had 1,500 bytes of available memory, which was insufficient to hold Smyth's state memory, even accepting Apple's expert's assertion that a typical large-vocabulary model would be around 4,700 bytes. The Board then concluded that Apple's argument that an abbreviated model would have been smaller was untimely as not raised "in the Petition or in the Reply" and in any event, there was insufficient evidence to show that an abbreviated model could fit on the 1,500 bytes of the DSP56000.

The Board then determined that the petition's proposed combination of Smyth and Mozer was also insufficient. First, the Board found that the proposed combination required the skilled artisan to personally fabricate the integrated circuit and that there was insufficient evidence that the skilled artisan would be able to do that. The Board concluded that "even if [Apple] had shown there was a reason to combine Smyth and Mozer," a proposed combination that required fabricating a new circuit combining the integrated circuit disclosed in Mozer with a larger memory would have been "beyond the level of ordinary skill." J.A. 38–39. Second, even assuming that the claims do not require fabricating a new circuit, the Board found that there was insufficient evidence that the skilled artisan would have been able to use a prefabricated circuit because Mozer's disclosure only included a memory of 4,000 bytes, which would not have been large enough to hold a 4,700-byte large-vocabulary acoustic model. The Board therefore found claim 1, as well as several dependent and similar claims,² not unpatentable for the same reasons.

² These are claims 2, 4, 6–9, 11, and 24.

C

As to claim 29, the petition argued that the claim was disclosed over a combination of several prior art references, including Smyth and Mozer. The petition mapped Smyth's classifier onto the claimed acoustic coprocessor and Smyth's sequencer onto the claimed at least one search processor. Zentian did not defend the patentability of claim 29 in either its response or sur-reply.

Nevertheless, the Board found that Apple had not shown that claim 29 was unpatentable, finding that the integrated circuits in the prior art were not capable of holding Smyth's classifier because Smyth's classifier contained the acoustic model memory, which the Board had determined was too large to fit on either the Motorola DSP56000 or Mozer's 4,000-byte memory.

Apple appeals the Board's conclusion that claims 1, 2, 4, 6–9, 11, 24, and 29 of the '789 patent were not shown to be unpatentable. We have jurisdiction under 28 U.S.C. § 1295(a)(4)(A).

DISCUSSION

On appeal, Apple argues that the Board erred in its analysis of claim 1 and that each of the remaining claims is unpatentable as obvious. Zentian agrees that the Board's decision as to claim 29 must be set aside if the decision as to claim 1 is set aside and does not advance any alternative arguments to support the Board's decision as to that claim.³ We therefore solely address Apple's arguments as to claim 1, which we find sufficient to resolve the issues on appeal.

³ To the extent the Board found that the petition's proposed combination did not disclose the claimed status register, results memories, or buffer memory of claim 29, those findings were not supported by substantial evidence.

“We review the Board’s factual findings for substantial evidence and review its legal conclusions de novo.” *Uber Techs., Inc. v. X One, Inc.*, 957 F.3d 1334, 1337 (Fed. Cir. 2020) (internal quotation marks omitted). As obviousness is a question of law, we review the Board’s ultimate obviousness determination de novo. *Fed. Express Corp. v. Qualcomm Inc.*, 174 F.4th 910, 913–14 (Fed. Cir. 2026).

I

As the Board found, the limitation that allegedly distinguishes claim 1 from the prior art is that it stores the acoustic model on chip, that is, in an integrated circuit with the processor. The prior art disclosed similar speech recognition circuits that had on-chip memory but nonetheless utilized off-chip memory to store the acoustic model. The theory of the petition was that it would have been obvious to use on-chip memory to store the acoustic model. The Board made no finding that a standard memory of the appropriate size would be unable to serve as the acoustic model memory. Neither did the Board find that storing data on the on-chip memory rather than off-chip memory was anything other than routine. Instead, the Board found that amongst the commercially available circuits presented by Apple, none had an on-chip memory large enough to accommodate a large-vocabulary acoustic model.

Apple argues that the Board erroneously required the use of a large-vocabulary model when a smaller model would have fit on an integrated chip at the time of the invention. Zentian concedes that claim 1, which only requires an “acoustic model,” does not require a large-vocabulary model and would be satisfied by an abbreviated model such as a digits recognizer model that can recognize spoken numbers but not other words. The claim language only requires that the acoustic model “defin[e] a plurality of acoustic states,” which would include an abbreviated model that contains fewer states than what would be needed to model an entire language. *See* ’789 patent,

claim 1. In requiring that Apple show that a large-vocabulary model—rather than a smaller abbreviated model—would have fit on the integrated circuit in the prior art, the Board committed legal error. An abbreviated model such as a digits recognizer model satisfies the claim limitation. The question remains whether the on-chip memory disclosed in Smyth and Mozer would be sufficient to store an abbreviated acoustic model.

Apple argues that the abbreviated model of a digits recognizer would fit on the Motorola DSP56000 chip disclosed in Smyth. Contrary to the Board, Apple raised this response to Zentian’s argument in its reply. J.A. 368 (“[T]he Motorola DSP56000 included a memory size sufficient to store an ‘abbreviated model’ as would be required for Smyth’s digit recognizer” (emphasis omitted)). The Board erred in concluding that the argument was not made in Apple’s reply.

However, the Board also concluded that Apple failed to prove that an abbreviated model would fit onto the Motorola DSP56000’s 1,500-byte memory. We need not determine whether this conclusion was supported by substantial evidence because Apple made alternative arguments concerning Mozer, which discloses a larger on-chip memory than the Motorola DSP56000.

Apple argues that the Board erred with respect to its theory based on a combination of Smyth and Mozer. We agree. As to the Smyth/Mozer combination, the Board again erroneously required a large-vocabulary model. While the Board addressed the question of whether Mozer’s disclosed 4,000-byte memory could hold a large-vocabulary model, it did not address whether it could hold an abbreviated digits recognizer model. As a result, while the Board found that an abbreviated digits recognizer model would not fit on the Motorola DSP56000’s 1,500-byte memory, the Board made no findings as to whether it would fit on Mozer’s memory. Mozer’s 4,000-byte memory is much

larger than the Motorola DSP56000 and is nearly large enough to store the 4,700-byte large-vocabulary model described by Apple's expert. Given that Zentian's expert testified that an abbreviated model could be "very small" compared to a large-vocabulary model, J.A. 3380, there is evidence in the record that Mozer's 4,000-byte memory was large enough to store an abbreviated model. On remand, the Board must address this factual question in the first instance and consider whether the combination of Smyth and Mozer would result in the claimed abbreviated acoustic model on an integrated circuit.

In this connection, Apple presented evidence that integrated circuits combining processors and memory were well known and Zentian does not argue otherwise. Apple also presented undisputed evidence that there were known benefits to using on-chip memory rather than off-chip memory. *See* J.A. 3395 (Zentian's expert acknowledging that "it was known that . . . one way of making a system more efficient[] was to reduce the off-chip memory accesses" by placing memory on chip). The patent indicates that the claimed acoustic model memory is generically described and does not require unique features. *See, e.g.*, '789 patent, col. 34 ll. 52–66 (noting that while the acoustic model is "typically held in a non-volatile memory" it can also be stored in "on[-]chip" RAM memory); *id.* col. 35 ll. 35–39 ("There is no requirement for a memory that is large enough to hold the entire decompressed acoustic model . . .").

II

Zentian argues that the Board rejected the Smyth/Mozer combination for another reason that Zentian argues sustains the Board's decision. The Board held that Apple was required to show that the ordinarily skilled artisan could personally fabricate a new circuit implementing the combination of a memory and processor. The relevant claim language requires that "the calculating apparatus

and the acoustic model memory are fabricated on a single integrated circuit.” ’789 patent, claim 1. While the petition argued that a skilled artisan would have found it obvious to apply “Mozer’s teachings to fabricate a calculating apparatus and acoustic model memory on a single integrated circuit,” the petition did not argue that a skilled artisan would be capable of fabricating the circuit without assistance. J.A. 171. Nothing in the patent indicates that fabrication is required to be performed as a step in the claims. The Board’s own description of the level of skill in the art assumes fabrication is not required because, as the Board found, the ordinarily skilled artisan was not shown to know how to fabricate a circuit with a generic memory on it. See *Liquid Dynamics Corp. v. Vaughan Co.*, 449 F.3d 1209, 1224 (Fed. Cir. 2006) (patent must enable the ordinarily skilled artisan to use the claimed invention). There is therefore nothing to suggest that the skilled artisan must personally fabricate the circuit rather than obtaining a prefabricated circuit. The Board’s reading of the claims was erroneous and incorporated an unclaimed requirement into its analysis.

CONCLUSION

We vacate the Board’s decision as to claims 1, 2, 4, 6–9, 11, 24, and 29 and remand for further proceedings consistent with this opinion.⁴

VACATED AND REMANDED

COSTS

Each party to bear its own costs.

⁴ We note that Apple raises on appeal other alleged errors with respect to the Board’s claim 29 determination that are not addressed in our opinion and would have to be addressed on remand if the Board were to conclude that claim 1 was not unpatentable.