IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent of:	Erez Galil	
U.S. Patent No.:	8,447,132	Attorney Docket No.: 39521-0044IP1
Issue Date:	May 21, 2013	
Appl. Serial No.:	12/961,400	
Filing Date:	December 6, 2010	
Title:	DYNAMIC RANGE CORRECTION BASED ON IMAGE CONTENT	

SUPPLEMENTAL DECLARATION OF ALAN C. BOVIK, PH.D., P.E

1. My name is Alan C. Bovik. I provided my credentials and background in my earlier declaration (APPLE-1003, ¶¶1-4) as to U.S. Patent No. 8,447,132 ("the '132 patent," APPLE-1001). That information has not materially changed since my earlier declaration, dated March 18, 2019 ("original declaration," APPLE-1003), and I incorporate that information into this declaration.

2. I have been retained on behalf of Apple Inc. to offer technical opinions relating to Petitioner's Reply to Patent Owner's Response filed in the *Inter Partes* Review concerning the '132 patent. For the purposes of this declaration, I have reviewed the following:

- U.S. Patent No. 8,447,132 to Galil ("the '132 patent," APPLE-1001);
- U.S. Publication No. 2002/0181801 to Needham et al. ("Needham," APPLE-1004);

- U.S. Publication No. 2008/0007634 to Nonaka et al. ("Nonaka," APPLE-1005);
- U.S. Patent 6,891,977 to Gallagher ("Gallagher," APPLE-1006);
- U.S. Publication No. 2008/0291287 to Dvir ("Dvir," APPLE-1007);
- Qualcomm Incorporated's Patent Owner Response (POR);
- Declaration of John Villasenor (EX-2001);
- Transcript of Deposition of Alan Bovik, Ph.D. (EX-2004);
- Transcript of Deposition of Dr. John Villasenor (APPLE-1021); and
- U.S. Provisional Application No. 60/285,063 ("the '132 provisional," APPLE-1025)

3. I have considered these materials through the lens of a person of ordinary skill in the art ("POSITA") related the '132 patent at the time of the earliest purported priority date of the '132 patent, which I understand to be December 9, 2009 (the "critical date"), as discussed in my original declaration (APPLE-1003, ¶¶6-7). The level of ordinary skill that I have applied here is the same level of ordinary skill in the art that I applied in my original declaration.

4. My opinions, as explained below, are based on my education, experience, and background. Unless otherwise stated, my testimony below refers to the knowledge of a POSITA in the field as of the critical date.

5. I have no financial interest in either party or in the outcome of this proceeding. I am being compensated for my work as an expert on an hourly basis, for all tasks involved. My compensation is not dependent in any manner on the outcome of these proceedings or on the content of my opinions.

6. My understanding of the legal standards for prior art, anticipation, and obviousness are set out in my original declaration (APPLE-1003, ¶¶22-44).

I. Claim Construction

7. I understand that Qualcomm proposes that the term "gain," used alone and in the phrase "gain map," should mean "an attribute that relates to brightening or darkening a region of an image." POR, 8. In my opinion, Qualcomm's construction is not "the broadest reasonable interpretation" because the construction improperly imports limitations from the specification. A POSITA would not understand the '132 patent's use of "gain" as limited to only "brightening or darkening a region of an image."

8. I agree with Dr. Villasenor that the '132 patent identifies different types of correction. APPLE-1021, 52:9-53:5, 115:10-116:6. For example, the '132 patent discloses:

> Note that the content-based technique introduced here is not necessarily limited to determining a gain adjustment or to DRC. For example, other types of corrections or adjustments can be made based on image content as described herein, such as

> contrast enhancement, color manipulation, sharpness adjustment, noise reduction, skin smoothing, etc.

APPLE-1001, 3:59-64.

9. I also agree with Dr. Villasenor's statement that it is possible to perform the different corrections identified in the '132 patent (e.g., smoothing, contrast enhancement, and color manipulation) using the structure of FIG. 2B, which includes a gain lookup table, or by determining a digital gain from a gain lookup table that outputs a gain map. APPLE-1021, 122:7-127:10. The '132 patent provides such an example:

> Multiplying each color component of the input image by one or more pixel dependent gain or attenuation factors, using a nonlinear mapping function that may either lighten shadow regions or darken highlights or change the chromaticity of pixels in either of these regions, generates the output image.

APPLE-1001, 3:49-53 (emphasis added).

> Although the following discussion is given in terms of a single gain factor for each pixel, more generally this can be one or more pixel dependent gain or attenuation factors or other factors, such as changing the chromaticity of pixels in either shadow regions or highlights, that can enhance the perception of image details at the extreme ends of a system's dynamic range.

APPLE-1001, 4:4-25 (emphasis added).

10. Based on my knowledge and experience in the field and my review of the '132 patent, Qualcomm's proposed construction of "gain" is not the broadest reasonable interpretation as understood by a POSITA after reading the entire patent, and the terms "gain" and "gain map" should be given their ordinary meanings under the broadest reasonable interpretation.

II. Needham discloses that the correction is matched to the predetermined type as required by claim 1.

11. I understand that Qualcomm argues that "Needham fails to disclose that a correction is determined that is ***matched*** to the 'predetermined type' of physical object detected on the image as required by Claim 1." POR, 14 (emphasis in original). I also understand that Dr. Villasenor states that "'matched' in this context means a particular type of correction is predetermined or assigned to a particular image feature." EX-2001, ¶51. In my opinion, Qualcomm and Dr. Villasenor are taking a very narrow view of the term "matched" because the '132 patent provides no requirement that the determined correction must be predetermined or assigned to the predetermined type.

12. Based on my review of the '132 patent, I agree with Dr. Villasenor that the term "matched" and the phrase "a particular type of correction is predetermined or assigned to a particular image feature" appear nowhere in the '132 patent's specification. APPLE-1021, 45:20-47:18. In spite of this,

Qualcomm states that the '132 patent requires that the "correction is *matched* to the particular image feature type in a predetermined manner." POR, 14 (emphasis in original). Qualcomm also states that:

> The '132 Patent describes examples of the different correction amounts as including different gain curves, such as shown in figure 7, that map input pixel intensities to output pixel intensities. *Id*. at 3:12-27, 5:52-60. The correction applied in the first and second amounts is "matched to the predetermined type" of physical object determined in the first portion of the image at step [1a]. For example, the '132 Patent describes that a lookup table provides gain values for use with a particular detected object type, and additional lookup tables can provide gain values for other object types. *Id*. at 5:65-67, 6:22-25.

POR, 10 (emphasis added). However, the portions of the '132 patent cited by Qualcomm do not disclose that a particular type of correction is predetermined for the particular image feature type, but rather only that, for whatever type of correction, the correction amount or the gain values are predetermined for different object types and other image content.

13. I understand that Qualcomm states that:

> Needham discloses no matching of these different corrections to the different types of features. This lack of matching is demonstrated in two ways in Needham. First, Needham discloses multiple objects of the same type receiving different

> corrections. *See id.* at [0032] ("The operational parameter(s) may also specify that different occurrences of the same image feature type (e.g., different human faces) in an image receive a differing correction operation.").

POR, 12. In my opinion, Qualcomm mischaracterizes Needham and ignores disclosure surrounding the quoted sentence, which states:

> As another example, to enhance the contrast 740 of an image feature, an intensity dynamic range 760 may be specified as an operational parameter so that the contrast in the detected image feature will be scaled to the specified dynamic range. Similarly, the specified dynamic range 760 may be defined as a function of statistical properties of one or more image features detected from the input image 110. The operational parameter(s) may also specify that different occurrences of a same image feature type (e.g., different human faces) in an image receive a differing correction operation. For example, the intensity dynamic range used for correcting a particular human face may be determined according to the size of the face detected. The larger the face (e.g., closer to the camera), the larger the dynamic range that may be used (so that the face can be seen more clearly).

APPLE-1004, [0032]. This example of Needham teaches multiple objects of the same predetermined type (e.g., human faces) receiving the same type of correction (dynamic range) in different amounts specified by operational parameters (intensity of the dynamic range correction). Thus, Needham discloses that

dynamic range correction is matched to human faces since multiple human faces receive dynamic range correction.

14. I understand that Qualcomm further states that:

> Additionally, corrections would be applied regardless of whether any particular object type or even whether any object was detected as confirmed by Petitioner's expert Dr. Bovik. Ex. 2004 at 112:3-16.

POR, 12.

> In fact, Petitioner's expert admitted that Needham does not disclose this element when Dr. Bovik stated that the corrections would be applied regardless of whether any object (e.g., a building) was detected:
>
> Q. In paragraph 22, Needham doesn't say that it's applying an intensity dynamic range correction to the face and the building because there's a face in the image; does it?
>
> A. I don't think that Needham is saying that the example that begins with "for instance" in paragraph 22 requires that the input image have faces in it. They may or may not.
>
> Q. And in paragraph 22, Needham doesn't say that it's applying an intensity dynamic range correction to the face and the building because there's a building in the image; does it?
>
> A. I think the algorithm would be applied whether or not there's a building in the image.
>
> Ex. 2004 at 112:3-16. Dr. Bovik further recited that corrections

> would be applied regardless of whether a face was detected: "Q. Would that algorithm also be applied whether or not there's a face in the image? A. I already said that. Yeah." Ex. 2004 at 112:18-20. If Needham's correction happens even in the absence of an object, then that correction is not matched to the object.

POR, 15-16. As an initial matter, Qualcomm mischaracterizes my testimony as I did not admit that Needham does not disclose this element of claim 1. Indeed, Needham's disclosure is similar to the '132 patent and its provisional in that the corrections would be applied regardless of whether any particular object type or even whether any object was detected.

15. For example, the '132 patent discloses the following:

> The content map 119 in one embodiment is a two-dimensional data structure indicating the locations of <u>any objects</u> of the predetermined type or types in the input image data. The content map 119 can be, for example, a matrix of binary values, each corresponding to a different pixel in the input image, where a binary one (1) value indicates that the corresponding pixel is part of an object of a predetermined type and <u>a binary zero (0) value indicates that the corresponding pixel is not part of an object of a predetermined type</u>.

APPLE-1001, 4:64-5:5 (emphasis added).

> FIG. 6 illustrates a general process that may be performed by the DRC unit 117 [in] accordance with one embodiment of the

> technique introduced here. Here the example of face detection is used to facilitate description, although other types of objects could be detected in addition or instead. At step 601, the process inputs a pixel intensity value(s) (see FIG. 2) from the input image. At step 602 the process determines whether the pixel is part of a group of adjacent pixels that collectively represent a face. If the outcome of this determination is negative, then at step 603 the process selects a normal gain value G(s). If the outcome is affirmative, then the process branches to step 605, where an alternate gain value is selected. The process then applies the selected gain value, G(s), (either directly or indirectly) to the input pixel intensity value L(s) at step 604.

APPLE-1001, 6:36-41 (emphasis added).

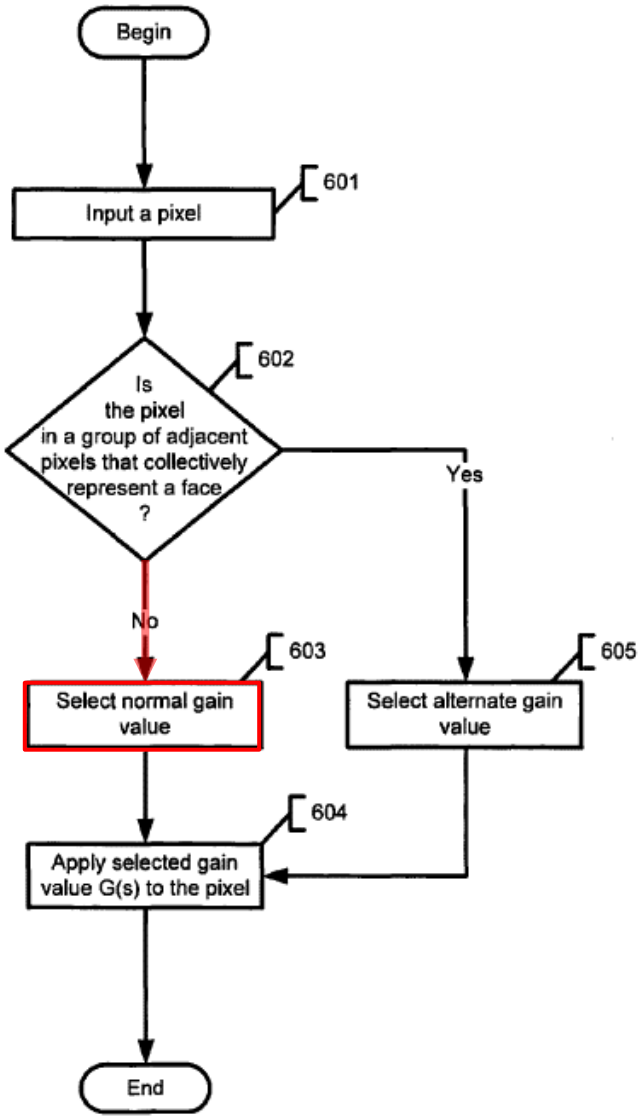

Fig. 6

FIG. 6 (annotated). Nowhere does the '132 patent describe applying no correction or not correcting the image if no predetermined object is found in the image. Like Needham, therefore, in the '132 patent, correction is applied regardless, just in an amount specified by the "normal gain value[s]."

16. The '132 provisional discloses the following:

> For normal image data (non-face), Z-Light applies a digital gain which is calculated accor[d]ing to a user defined gain LUT. The LUT defines the input/output relation (digital gain), fu[r]thermore Z-Light will make sure local contrast is maintained when applying the digital gain. The gain LUT, is designed by the camera manufacturer to best fit normal content (without specific priority).
>
> For specific image content (i.e. face), Z-Light will use a gain LUT which is optimized to brighten the face to a desired user level. The gain LUT will be calculated according to the face image data which can be obtained from the image data (input to Z-Light) and the face location (ou[tp]ut of Face Detection module).

APPLE-1025, 8. Nowhere does the '132 provisional describe applying no correction or not correcting the image if no predetermined object is found in the image. Like Needham, therefore, in the '132 provisional, correction is applied regardless, just in an amount specified for the "normal image data."

17. In my original declaration, I cited multiple portions of Needham which support Needham's disclosure of the claimed matching. APPLE-2003, ¶¶45-52, 59 (citing APPLE-1004, [0001]-[0003], [0015], [0017]-[0022], [0024], [0026]-[0033], [0035], FIGS. 2 and 5-8). For example, paragraph [0020] discloses:

> One or more types of correction operations may also be defined for each image feature. For example, both contrast and brightness may be corrected for a specific image feature such as a human face. Further, when one or more correction operations are applied to individual image features, different types of correction operations may be applied to different image features. For example, maximizing the intensity dynamic range (i.e., enhance the contrast) may be performed on a human face image feature and a different correction operation that increases the brightness may be applied to the image feature that represents the sky.

APPLE-1004, [0020] (emphasis added). Here, Needham clearly discloses a type of correction being "defined" for, i.e., "matched" to, an image feature, e.g., a human face image feature. Needham's example describes different types of corrections being applied to different types of image features—the operation of maximizing dynamic range is applied to human faces and the operation of increasing brightness is applied to the sky—which indicates that the different corrections are matched to the different image features. In addition, paragraph [0032] discloses "different occurrences of a same image feature type (e.g., different human faces) in an image receive a differing correction operation" such as different amounts of "dynamic range" correction, which also establishes that

dynamic range correction is matched to human faces as I discussed above in ¶13.

More generally, paragraph [0026] discloses:

> In operation, the automated feature detection unit 250 detects, from the input image 110, the types of image features that are specified by the feature types 220. Such feature detection may be performed using any technique or algorithm for detecting features in an image as should be known to those skilled in the art. For each detected image feature, the automated feature detection unit 250 may construct a corresponding feature description 260. A feature description may characterize the visual properties of the corresponding detected image feature(s). Such characterization may include the location, the size, or the statistical properties (e.g., average intensity or minimum and maximum intensity values) of the detected image feature. Based on the feature description 260, the feature-based correction unit 270 performs one or more specified correction operations according to the correction parameter(s) 240.

APPLE-1004, [0026] (emphasis added.) This is illustrated in FIG. 2:

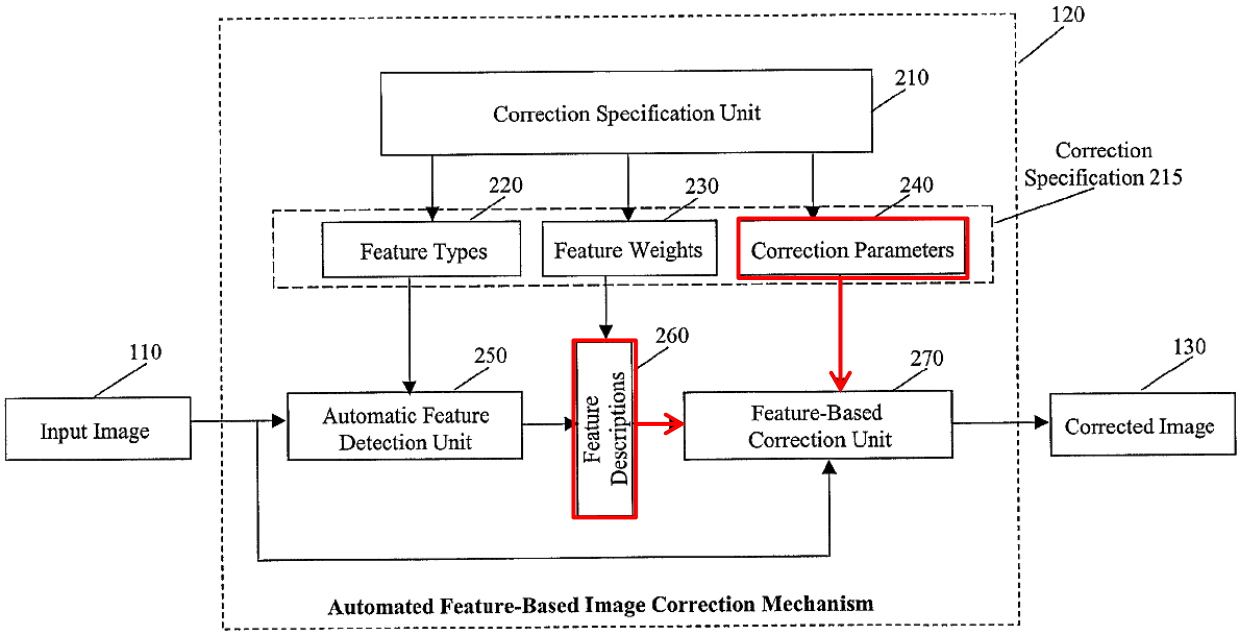

Fig. 2

APPLE-1004, FIG. 2 (annotated). Paragraph [0030] also discloses:

> FIG. 6 illustrates an example construct of a feature description 260 of an embodiment of the present invention. A feature description is used to characterize one or more detected instances of an image feature. In FIG. 6, a feature description 260 includes a feature type 610, a location descriptor 620, a shape descriptor 630, and statistical properties 640 of the image feature. An image feature corresponds to an area in the input image 110. Such an area may occupy a region of an arbitrary shape at a certain location in the image. For example, a detected human face is located somewhere in an image and the location of that face may be described using a center point of the face. In addition, a human face normally has an elongated shape, which may be characterized by a curve along the boundary of the human face,

> representing the precise shape of the detected face. Furthermore, statistics about the visual properties of the face may be computed within the boundary of the detected face. Such descriptions may be used to determine where and how to apply one or more specified correction operations.

APPLE-1004, [0030] (emphasis added). Paragraph [0031] further discloses that the correction operation defines the type of correction (e.g., brightness or contrast enhancement) to be performed:

> FIG. 7 shows an example construct of the correction parameters 240 according to an embodiment of the present invention. Each correction operation is defined by both an operation mode 705 and an operation definition 710. The operation mode 705 may specify whether the correction operation is applied to the entire image or to individual image features. The operation definition 710 defines the correction to be performed. For example, a correction operation may be defined as increasing the brightness 730 (of either the entire image or an image feature or both, depending on the operation mode 705). A correction operation may also be defined as enhancing the visual contrast 740 (of either the entire image or an image feature or both, depending on the operation mode 705).

APPLE-1004, [0031] (emphasis added). Needham [0026] and [0030]-[0031], in conjunction with Needham [0020] and [0032], specify that a feature type, e.g., human faces, is matched to a specified correction operation, e.g., dynamic range

correction, and the location descriptor, the shape descriptor, and the statistical properties are used to determine where and how to apply the specified correction operation. In addition, paragraph [0035] discloses "[f]or each detected image feature, the corresponding correction parameters are retrieved" and "[t]he correction operation(s) is then performed 870 on each image feature." Paragraph [0019] discloses that "[c]orrection operations may also be defined on individual image features" such as "detected human faces." Paragraph [0020] discloses "[o]ne or more types of correction operations may also be defined for each image feature" and "different types of correction operations may be applied to different image features." Paragraph [0029] discloses "[t]he correction is made in accordance with the specified correction parameters and the detected image features." Based on my knowledge and experience in this field and my review of Needham, Needham's disclosures of a correction operation and/or a correction parameter being "defined," "specified," "based on," "in accordance with," and "retrieved" for, as well as "corresponding" to, an image feature and an image feature type, is sufficient to satisfy the claimed "match[ing]."

III. Needham discloses applying the matched correction to two portions of the image in different amounts as required by claim 1.

18. As I discussed above in ¶13 and ¶17, Needham, particularly the disclosure in paragraphs [0020] and [0032], establishes that dynamic range

correction is "defined" for, i.e., "matched" to, human faces. Needham [0022] further discloses that "the intensity dynamic range used for correcting human faces in an image ***may be larger (corresponding to higher contrast) than that*** used for correcting buildings in an image." I understand that Qualcomm argues that:

> Needham (e.g., paragraph [0022]) fails to, for example, explicitly define that the same operation is performed in two different amounts on two different portions. *See id.* Paragraph 22 describes that a higher contrast is applied to change the intensity dynamic range of the face portion, but fails to specify that a contrast correction is necessarily applied to the building portion of the image. APPLE-1004 at [0022]. Instead, Needham's paragraph 22 is ambiguous about what correction is applied to the building to adjust the intensity dynamic range of the building. *See id.*; Villasenor Decl. ¶45.

POR, 21. I also understand that Dr. Villasenor states that:

> Needham's disclosure that "the intensity dynamic range used for correcting human faces in an image may be *larger* (corresponding to higher contrast) than that used for correcting buildings in an image" is ambiguous in light of the Needham disclosure in the very same sentence describing "maximizing intensity dynamic range."

EX-2001, ¶45. Based on my knowledge and experience in this field and my review of Needham, I do not agree that Needham [0022] is ambiguous.

19. I agree with Dr. Villasenor that "'[m]aximizing' the intensity dynamic range means using the full range of intensity values." EX-2001, ¶45. I also agree with Dr. Villasenor's discussion of ways of adjusting dynamic range, which include extending one end of the dynamic range (e.g., increasing the brightest value):

> Q. Are there different ways to adjust dynamic range of an image or of part of an image, like a face?
>
> A. Yes, there are.
>
> Q. Can you give me some examples?
>
> A. So one example would be to perform an adjustment such that after the adjustment, the region of interest used the full range of possible values from using the convention agreed on before from 0 to 255, but there are other ways as well. You could adjust it by -- by, for example, leaving the -- the darkest value in the image, not changing that, but only making – you know, doing -- doing an increase in brightness in proportion to how much brighter a pixel was relative to the darker pixel.
>
> So, in other words, the darkest pixels would stay unchanged and the brightest pixels would get adjusted the most or you could do the opposite where you would keep the brightest pixels unchanged and then darken the darkest pixels, and -- but the ones that are half way in between you darken less than the ones that are already as dark.
>
> You could also shift -- just shift things. So if -- you know, if

> -- again, using 0 to 255 as the full possible range, if an area had pixel values from a hundred to 160, you could just add 20, so it goes from 120 to 180; so there's any number of ways that you could adjust the dynamic range.

APPLE-1021, 200:22-202:2 (emphasis added).

20. Based on my knowledge and experience in this field, Needham [0022] discloses to a POSITA applying a correction of "maximizing intensity dynamic range" using a "larger" intensity dynamic range for human faces "(corresponding to higher contrast)" ***and*** applying the same correction of "maximizing intensity dynamic range" using a smaller intensity dynamic range (corresponding to lower contrast) for buildings in the same image. The "larger" intensity dynamic range used for human faces would correspond to a larger increase in the brightest value of the dynamic range, resulting in higher contrast for faces. The smaller intensity dynamic range used for buildings in the same image would correspond to a smaller increase in the brightest value of the dynamic range, resulting in smaller contrast for buildings. The correction would then maximize, i.e., use the full range of intensity values, within the larger intensity dynamic range determined for human faces when correcting human faces and maximize within the smaller intensity dynamic range determined for buildings when correcting buildings in the same image, i.e., apply two different amounts of the same type of correction, anticipating the claims of the '132 patent. Thus, Needham [0022] provides a clear and concrete

example of "applying a first amount of the correction to the first portion of the digital image data" and "applying a second amount of the correction to a second portion of the digital image data." Indeed, this is similar to the '132 patent's disclosure regarding the histograms 802A, 802B, and 802C in FIGS. 8A-8C. APPLE-1001, 9:54-10:28 (shift in "peak of histogram" "reflects the more optimal brightness level of face area"), FIGS. 8A-8C.

IV. Needham discloses determining different amounts of gain for different portions as required by claim 7.

21. As I have articulated above in ¶13 and ¶17, Needham, particularly the disclosure in paragraphs [0020] and [0032], establishes that dynamic range correction is "defined" for, i.e., "matched" to, human faces. As I have articulated above in ¶¶18-20, Needham [0022] further discloses applying a larger intensity dynamic range for human faces when correcting human faces, and applying a smaller intensity dynamic range for buildings when correcting buildings in the same image. As Dr. Villasenor stated, and I agree, dynamic range correction involves changing the intensity of one or more pixel values:

> THE WITNESS:
>
> A person of ordinary skill in the art would know that there's – that Needham does describe enhancing the contrast in that that would -- and that would involve changing, generally speaking, in changing the intensity of -- of the pixel value or more than one.
>
> BY MR. AMON:

Q. Would a person of skill in the art know that Needham describes in enhancing the brightness and that that would involve changing -- generally changing the intensity of the pixel value or more than one pixel value?

A. Well, I mean, almost -- almost -- most operations, almost any operation involves changing intensities. It's -- it's -- you know, increasing the brightness would be a different, you know, correction. It's described as a different correction operation than enhancing the contrast, but in both cases you would change pixel values.

Q. And that would include changing pixel intensity for both brightness and contrast, fair?

A. Well, for brightness it certainly would, but -- but for contrast, you know, it would depend on how you did it; but -- but, like I said, almost any -- any -- any manipulation of images could involve some change to the pixel values.

APPLE-1021, 136:9-137:14 (emphasis added). I also agree with Dr. Villasenor that ways of adjusting dynamic range include adjusting the brightest value (an attribute related to brightening) and/or the darkest value (an attribute related to darkening) of the image portion to be corrected:

Q. Are there different ways to adjust dynamic range of an image or of part of an image, like a face?

A. Yes, there are.

Q. Can you give me some examples?

A. So one example would be to perform an adjustment such

> that after the adjustment, the region of interest used the full range of possible values from using the convention agreed on before from 0 to 255, but there are other ways as well. You could adjust it by -- by, for example, leaving the -- the darkest value in the image, not changing that, but only making – you know, doing -- doing an increase in brightness in proportion to how much brighter a pixel was relative to the darker pixel.
>
> So, in other words, the darkest pixels would stay unchanged and the brightest pixels would get adjusted the most or you could do the opposite where you would keep the brightest pixels unchanged and then darken the darkest pixels, and -- but the ones that are half way in between you darken less than the ones that are already as dark.
>
> You could also shift -- just shift things. So if -- you know, if -- again, using 0 to 255 as the full possible range, if an area had pixel values from a hundred to 160, you could just add 20, so it goes from 120 to 180; so there's any number of ways that you could adjust the dynamic range.

APPLE-1021, 200:22-202:2 (emphasis added). Based on my knowledge and experience in this field and my review of Needham, a POSITA would have understood that Needham determines different amounts of gain for dynamic range correction of the face portion and the building portion of the image, even under Qualcomm's proposed construction of "gain."

V. Nonaka discloses that the correction is matched to the predetermined type as required claim 1.

22. I understand that Qualcomm argues that:

> Nonaka describes that the correction (applied by "optimization processing section 5b") is still applied when a face is not detected, such as "when the face detecting section 11 cannot detect the face . . . to detect the face, the optimization processing section 5b and the exposure control section 12a perform processing and control so as to brighten the dark part." APPLE-1005 at [0042]. Apple's expert Dr. Bovik confirms that the optimization processing section's correction is being performed regardless of whether a face is detected. Ex. 2004 at 152:22-153:1. Dr. Bovik also confirms that paragraph [0042] described here and paragraph [0058] cited in the Petition are just "describing different scenarios where the processing happens differently for good reasons." Ex. 2004 at 156:3-5. Thus, even though the Petition's cited portions of Nonaka (e.g., paragraphs [0029], [0033], [0050], [0058], [0064], Fig. 1) may disclose applying optimizations to a face, Nonaka cannot be interpreted such that Nonaka's optimization processing is matched to the object type "face." If the brightening is performed regardless of whether the face is detected, then that brightening cannot be matched to the detection of a face object type.

POR, 26-27. However, Qualcomm misinterprets Nonaka and mischaracterizes my testimony.

23. Nonaka describes three different optimization processing options: "an image subjected to the optimization processing for the face display in the step S17, an image subjected to the intense optimization processing in the step S20 and an image subjected to the emphasis processing for the face detection in the step S15." APPLE-1005, [0070]. FIG. 8 shows these three options:

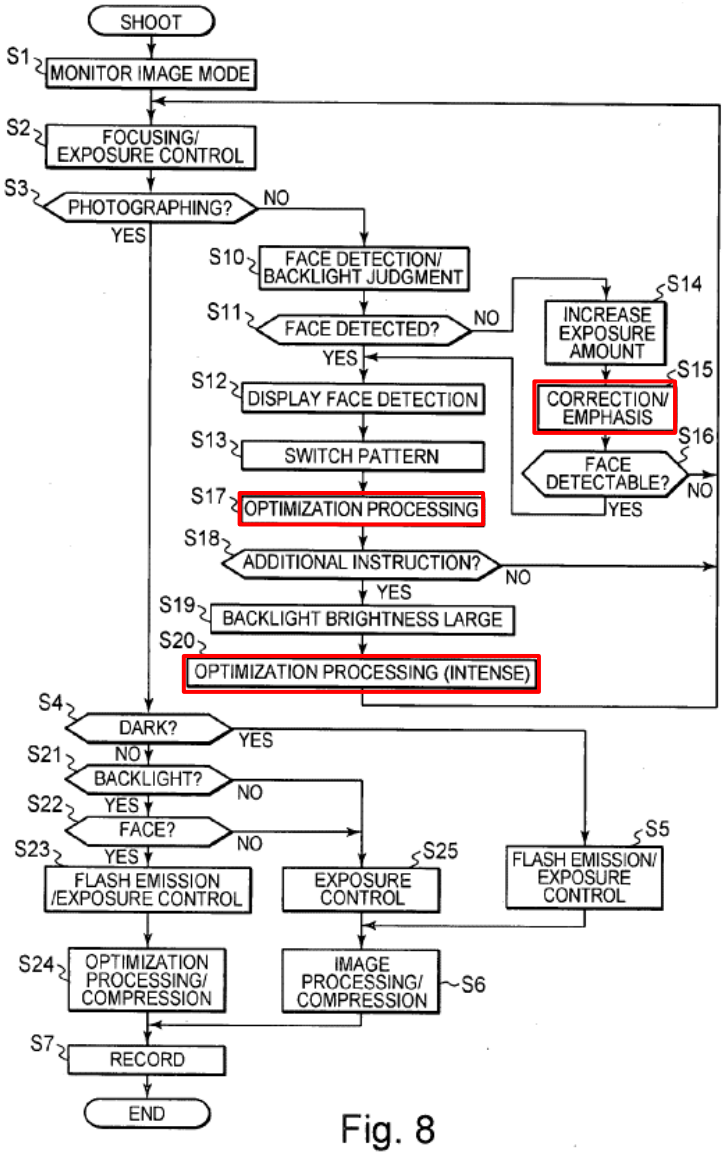

Fig. 8

APPLE-1005, FIG. 8 (annotated).

24. Nonaka discloses not performing the "optimization processing" of step S17 when no face is detected:

> FIG. 6 is a diagram showing a scene of only backlight landscape without any person. In such a scene, the optimization processing section 5b does not perform the optimization processing. The optimization processing section 5b performs the optimization processing of the scene in which the person is detected at the backlight part as shown in FIG. 3, and performs the exposure control.

APPLE-1005, [0058]. FIG. 8 shows that when "NO" is the answer to "FACE DETECTED?" at step S11 and "NO" is the answer to "FACE DETECTABLE?" at step S16, Nonaka does not perform the optimization processing of step S17:

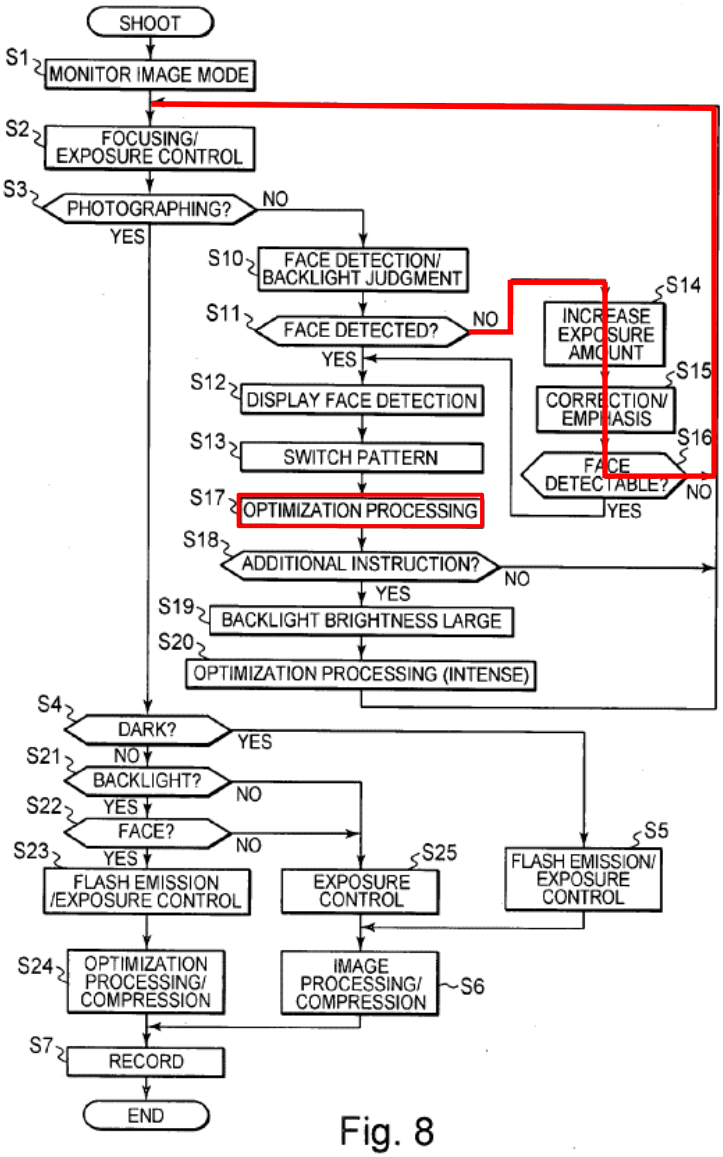

Fig. 8

APPLE-1005, FIG. 8 (annotated), *see also* [0067] ("when the face cannot be detected in the step S11 (step S11 NO)"), [0069] ("when the face cannot be detected (step S16 NO), the processing returns to the step S2"), FIG. 8.

25. On the other hand, Nonaka describes performing the "emphasis processing" of step S15 when "face detecting section 11 cannot detect the face" in step S11:

> In addition, in a case where the main subject is positioned at a dark part of the screen of the backlight scene or the like, the image is blackened, and a fine difference between the brightness and the darkness cannot be seen. In this state, the face detecting section 11 cannot detect the face. In such a backlight scene, to detect the face, the optimization processing section 5b and the exposure control section 12a perform processing and control so as to brighten the dark part as described later in detail.

APPLE-1005, [0042] (emphasis added).

> The face detection is performed by the face detecting section 11, and the backlight judgment is performed by the scene judgment section 12b, respectively. Moreover, the face detecting section 11 can judge whether or not the face detection can be performed (step S11).

APPLE-1005, [0063].

> On the other hand, when the face cannot be detected in the step S11 (step S11 NO), the processing is performed so as to realize the face detection. For example, in the image in which the face is too dark to be detected as shown in FIG. 2B, the exposure and the processing of the image are switched to obtain the image in which background is whitened to disappear, the dark part is

> regarded as important and the face can be detected as shown in FIG. 2C.
>
> First, the exposure control section 12a lengthens an exposure time for the face detection to increase the exposure amount (step S14). Furthermore, the optimization processing section 5b performs emphasizing correction processing (amplifying correction, contrast emphasis or the like) for the face detection (step S15).

APPLE-1005, [0067]-[0068]. This is shown in FIG. 8:

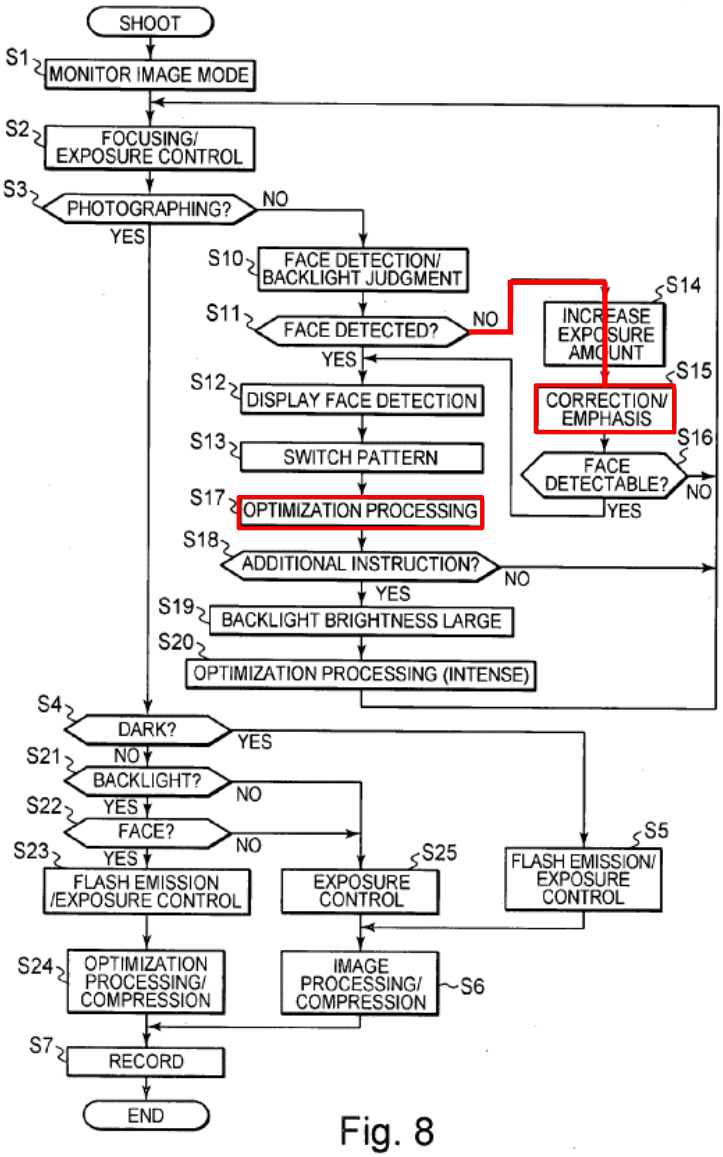

Fig. 8

APPLE-1005, FIG. 8 (annotated).

26. The "emphasis processing" of step S15 ***is a different correction*** than the "optimization processing" of step S17. APPLE-1005, [0064] ("optimization processing of correcting the brightness of the part and emphasizing the contrast (step S17)"), [0068] ("emphasizing correction processing (amplifying correction,

contrast emphasis or the like) for the face detection (step S15)"). Based on my review of Nonaka, Nonaka's system does not perform the "optimization processing" of step S17 when no face is detected. APPLE-1005, [0058], [0067], [0069], FIG. 8. Because the "optimization processing" of step S17 is performed when a face is detected and not performed when no face is detected, the "optimization processing" of step S17 is matched to faces (a predetermined type of object).

27. Regarding my testimony in EX-2004 at 152:22-153:1 that Qualcomm relies on, I was merely confirming what is stated in Nonaka [0042], ***not*** that the "optimization processing" of step S17 is performed regardless of whether a face is detected:

> Q. Paragraph 42 says if a face is not found in the dark area, then the optimization processing section 5b performs the optimization processing section on the dark area; correct?
>
> A. It says that, yeah.

EX-2004, 152:2-153:1. Additionally, Qualcomm's reliance on my testimony in EX-2004 at 156:3-5 regarding Nonaka's [0042] and [0058], reproduced below, is misplaced.

> Q. So in Nonaka, paragraph 48 and -- sorry -- paragraph 42 and paragraph 58 can't both be true if they're on the same path; can they?

A. Paragraph 58 is talking about a situation where there isn't any person, meaning no face, because there isn't a person. Paragraph 42 is talking about there is a face potentially, but it's dark, so it can't detect the face. It's a different scenario than whether, you know, the other one there's just no person; here, it's saying it's dark. Okay?

And certainly, you know, it goes on to say that, "The optimization processing section and the exposure control section perform processing control so as to brighten the dark part as described later in detail."

So it's perfectly acceptable to have a variation where you detect that something is so dark that you can't decide whether there's a face, then you brighten it, then you try face detection again. I don't see any contradiction at all.

Q. Do you agree both paragraphs 42 and 58 are addressing the possibility in which the face detection unit can't initially determine whether there is a face in the dark part of the image?

A. They can't determine, but for different reasons. The conditions of the image are different. So just because they can't detect the face doesn't mean that the camera can't, you know, determine that, "Well, it's really dark there. Maybe that's why I didn't detect a face."

So it doesn't mean the processing has to be the same or that those two scenarios are the same, and that's why there's different paragraphs.

Q. So given the scenario 42, when the dark section is so dark

> it can't detect a face, if you apply paragraph 58 to that same scenario, wouldn't 42 and 58 reach different conclusions: one of which would brighten the dark area and one of which would not brighten the dark area?
>
> A. (Reviewing document.) I don't think that paragraph 58 and 42 are generally -- you know, you can't really comingle them, because paragraph 58 doesn't talk in any way about there being a dark part. Okay? It's not even a same scenario.
>
> I mean, not such a dark part you couldn't find the face. Let's be careful. It does say "produced that the dark part is emphasized." Okay?
>
> So I don't think that those two paragraphs are meant to be commingled or combined or one superimposed on the other one artificially in any such way. I think they're describing different scenarios where the processing happens differently for good reasons.

EX-2004, 154:4-156:5 (emphasis added). This portion of my testimony actually supports that Nonaka's "emphasis processing" of step S15 and "optimization processing" of step S17 are two different corrections, not the same correction as Qualcomm argues.

VI. A POSITA would have found it obvious to combine Needham with Nonaka.

28. I understand that Qualcomm argues that:

> Needham's algorithm already maximizes the dynamic range.

> *See* APPLE-1004 at [0019] ("a correction operate that maximizes the intensity dynamic range may be defined as the correction operation"); Villasenor Decl. ¶77. With the dynamic range of the image already maximized, a POSITA would not have been motivated to modify Needham to further improve dynamic range. Villasenor Decl. ¶77. With no anticipated benefit to dynamic range achieved by modifying Needham with Nonaka, there would be no motivation to implement additional processing that may have a negative impact on the system design as described further below.

POR, 28-29, *see also* 36-37. As I have stated in my original declaration, a POSITA would have been led by Nonaka to modify Needham to apply amplifying correction and contrast emphasis (i.e., dynamic range correction), enabling both the person and landscape portions of an image to be more easily distinguished and perceived, and a POSITA would have found it desirable to implement Nonaka's teachings in Needham's camera to process image data at a high speed so that the image signal from an image pickup device is substantially displayed in real time and process image data to improve the visibility of the subject during the display of the image for monitoring while the user is shooting the image. APPLE-1003, ¶¶62-63. Qualcomm does not address these improvements and benefits. The desirability of such improvements and benefits, as taught in Nonaka, would have motivated a POSITA to combine Nonaka's and Needham's teachings.

29. I understand that Qualcomm also argues that Needham's and Nonaka's architectures are incompatible. Specifically, Qualcomm argues that:

> A POSITA would not seek to modify Needham's single-path image processing based on the disclosure of a two-path processing system in Nonaka because a POSITA would recognize the solutions of Needham and Nonaka to be two alternate solutions, not complimentary solutions, to the same problem of backlight control. Villasenor Decl. ¶¶68, 72-76.

POR, 29-30.

> A POSITA would recognize more specifically that Needham is directed to post-processing techniques for processing images after capture by a camera. Villasenor Decl. ¶¶71-72. In contrast, Nonaka is directed to an architecture aimed at enabling differential processing for two copies of an image signal at the time of image capture. Villasenor Decl. ¶¶66-67. These two techniques require different implementations in a digital image processing system, and thus a POSITA would not have combined these references. Villasenor Decl. ¶¶68-73.

POR, 31.

> A POSITA would not combine Nonaka with Needham because Needham is directed to post-processing of images whereas Nonaka is directed to a system for combined processing of a monitored-image and a photographed-image. See APPLE-1005 at [0033] ("the optimization processing section 5b subjects

> the image to correction processing such as the amplification of the luminance for each region or the contrast emphasis in order to improve the visibility of the subject during the display of the image for monitoring"). These two architectures are incompatible.

POR, 32-33. In making these arguments, Qualcomm fails to focus on the teachings of Needham and Nonaka. As Dr. Villasenor confirms and I agree, both Needham and Nonaka operate on digital image data:

> Q. Dr. Villasenor, as it relates to the Needham reference, are corrections under the method described in Needham applied to digital image data?
>
> …
>
> A. It's being applied to digital image data is the way I would put it, yeah.

APPLE-1021, 133:14-135:2.

> THE WITNESS:
>
> Nonaka is operating on -- on digital data, digital image data, yes.

APPLE-1021, 155:2-20. In addition, Nonaka teaches not performing optimization processing when no face is detected in digital image data, using gain as a correction parameter for dynamic range correction of digital image data, and applying different amounts of gain to different portions of digital image data in a single image. APPLE-1003, ¶¶56, 57, 60-62. A POSITA would have been

prompted to combine these teachings of Nonaka with Needham for the reasons stated in my original declaration and repeated above. APPLE-1003, ¶¶62-63.

30. Qualcomm also argues that "[a] POSITA would not combine Needham and Nonaka because a POSITA would not understand how to combine Needham with Nonaka in a manner to provide a reasonable expectation of success." POR, 32-35. As one reason, Qualcomm states that:

> Nonaka's image optimization processing is intended to brighten dark regions that contain faces, but Nonaka's system cannot detect that a face exists without first brightening the image. APPLE-1005 at [0042]. In contrast, Needham's correction operations are only applied when objects are detected. APPLE-1004 at [0022]. Nonaka's gains are dependent on the scene already brightened enough to identify the face. That pre-brightening does not exist in Needham, and thus Nonaka's gains may not be useable values in Needham's system.

POR, 33. As discussed above in ¶25, Nonaka [0042], which Qualcomm argues describes "pre-brightening," in actuality describes performing the "emphasis processing" of step S15 because the "face detecting section 11 cannot detect the face" in step S11. APPLE-1005, [0042], [0063], [0067]-[0068], FIG. 8. Contrary to Qualcomm's arguments, Nonaka does not require "pre-brightening" (i.e., "emphasis processing" of step S15) prior to "full correction" (i.e., "optimization processing" of step S17) when a face is detected in the pre-processed image data

("FACE DETECTED?" is "YES" in step S11), as shown in FIG. 8 and described in the accompanying description:

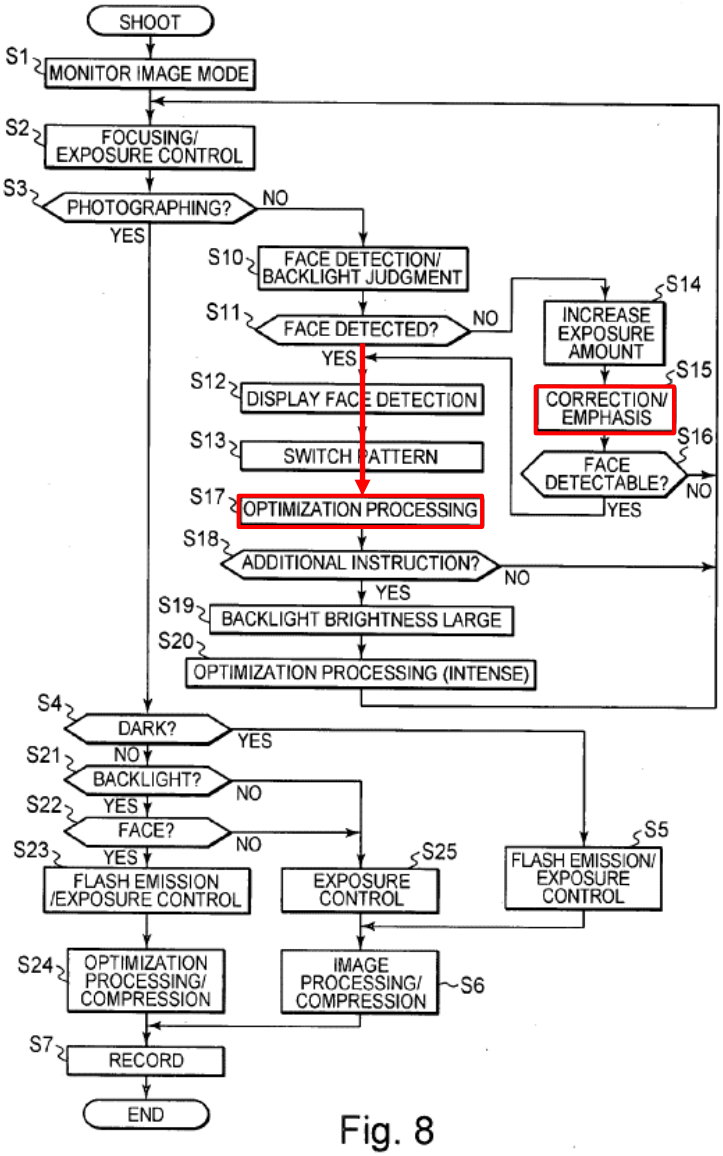

Fig. 8

APPLE-1005, FIG. 8 (annotated).

Moreover, when the face is detected (step S11 YES), as shown in FIG. 4A, the optimization processing section 5b subjects the face part and the person to the optimization

> processing of correcting the brightness of the part and emphasizing the contrast (step S17). The image subjected to the optimization processing including the correction of the brightness of the face part and the emphasizing of the brightness is displayed in the display section 8. During the monitor display prior to this release, the image is constantly switched. Therefore, any disturbance due to any degree of noises is not noticed. It is important whether or not facial expression can be seen. To solve the problem, the optimization processing is performed so as to improve the visibility, so that the person of the subject can well be seen in the display section 8.

APPLE-1005, [0064]. As shown above, Nonaka's gains are applicable to the processing path where no "emphasis processing" (step S17) is performed when the face is detected at step S11, and Nonaka does not require the gains to be dependent on "pre-brightening." Thus, Nonaka's gains are useable values in Needham's system, with the weights controlling the gains as taught by Needham and Nonaka. APPLE-1003, ¶¶61-62 ("Needham's teachings of applying weights to control correction parameters" and Nonaka's teachings of "using gain as a correction parameter").

31. Qualcomm mistakenly looks to my testimony for additional support. POR, 34-35. Specifically, Qualcomm relies on the following portions of my testimony:

Q. Is it your opinion that the image for monitoring is the same as the image that is being photographed?

A. In the sections I'm referencing, these are from the first embodiment, keep in mind that when you're attempting to capture a photograph with an image, you've got to monitor. What's displayed on the monitor is rarely, if ever, the exact same as the image that is -- would be captured because it goes through a series of processing steps prior to display.

So, generally, what's shown in the monitor is a representation of what you might be capturing, okay, on any camera. That's a way of guiding the photographer as they're capturing a picture. What's on that monitor won't be exactly the same as what will be captured generally ever in most cameras.

EX-2004, 137:17-138:7.

Q. So paragraph 42, when it says that the face can't be detected, then the optimization processing section 5b and the exposure control section 12a perform processing and control so as to brighten the dark part, that's referring to the image for monitoring; correct?

A. So in paragraph 42, it says, "The face detection section extracts the characteristic point from information during focusing and the image for monitoring described above" -- so that's information -- "based on the image data" -- meaning the image for monitoring -- "output from the image processing section 5 to detect the face."

EX-2004, 141:25-142:11.

> Q. Nonaka's paragraph 42, just to close the loop, is it discussing processing of the image to be monitored or the photographed image?
>
> A. (Reviewing document.) I'm reading the paragraph now. My apologies. Well, this signal in paragraph 42 is being discussed. It refers to the image for monitoring. Okay? "The image data," parentheses, "the image for monitoring, output from the image processing section 5." Okay?

EX-2004, 150:24-151:7. These portions of my testimony discuss the differences between the monitored image and the photographed image and that Nonaka [0042] refers to the monitored image. These portions of my testimony do not support Qualcomm's arguments that "pre-brightening" "is applied to the monitored image prior to the full correction that is applied to the photographed image." POR, 34.

32. Qualcomm further relies on the following testimony:

> Q. Okay. Hypothetically, if the blackened image of paragraph 42 were run through the steps of paragraph 58, Nonaka wouldn't be able to detect the face; would it?
>
> A. I mean, we started out with "hypothetically," which means I'm stepping outside the embodiment. In any such hypothetical thing, it's a different embodiment. I don't know what's going to happen. Okay. I don't know, for example, what the threshold for darkness is. Okay?

EX-2004, 180:7-15. My testimony provides no support for Qualcomm's argument that "[w]ithout such preliminary brightening, Nonaka's system may not detect faces and would not operate as expected in an image that included a face in a dark region." POR, 34-35.

VII. Dvir discloses determining a gain amount "to maintain local contrast of the first portion of the digital image data" as required by claim 8.

33. I understand that Qualcomm argues that:

> Needham does not disclose determining gains. Dvir likewise does not disclose determining gains, let alone determining gains that maintain a local contrast as required by claim 8.

POR, 40-41. As discussed above in ¶21, Needham discloses determining gains. Dvir also discloses determining gains. For instance, Dvir discloses using "a cascade of nonlinear edge preserving filters, and pixel point operations, to calculate pixel gain … that can be employed to make image detail more visible." APPLE-1007, [0006]; APPLE-1003, ¶64. Dvir [0012], [0014], [0016], [0018], [0019], and [0025] also describe determining a gain amount. APPLE-1003, ¶¶64-68.

34. Qualcomm and Dr. Villasenor both argue that "to 'maintain' local contrast requires specific consideration of the contrast before the processing as well as affirmative steps to ensure a similar output contrast—something that is not discussed in Dvir." POR, 41; EX-2001, ¶96. I agree with Dr. Villasenor that the '132 patent at 3:28-32 and 9:28-53 discuss maintaining local contrast and relate to

considering local contrast before the processing steps. APPLE-1021, 93:13-98:16. There appears to be no other portions of the '132 patent specification that discuss "local contrast," much less "maintaining local contrast." As I pointed out in my original declaration (APPLE-1003, ¶64) and worth reemphasizing here, the '132 patent at 3:28-32 and 9:28-53 are substantially the same as Dvir [0012] and [0033]-[0036], respectively. *Compare* APPLE-1001, 3:28-32, 9:28-53 *with* APPLE-1007, [0012], [0033]-[0036]. Thus, contrary to Qualcomm's and Dr. Villasenor's contentions, Dvir [0012] and [0033]-[0036] discuss maintaining local contrast and relate to considering local contrast before the processing steps.

35. These portions of the '132 patent and Dvir describe a "general embodiment of FIGS. 3 and 4," which is an embodiment of block 205 shown in FIGS. 2A-D of the '132 patent and in FIG. 2 of Dvir. APPLE-1001, 5:31-51, 9:23-24, FIGS. 2-4; APPLE-1007, [0017]-[0018], [0033], FIGS. 2-4. As I noted in my original declaration, the '132 patent's FIG. 2B shows a DRC unit 117 that includes the same blocks 205 and 207 and in exactly the same configuration as the DRC unit 117 shown in FIG. 2 of Dvir.

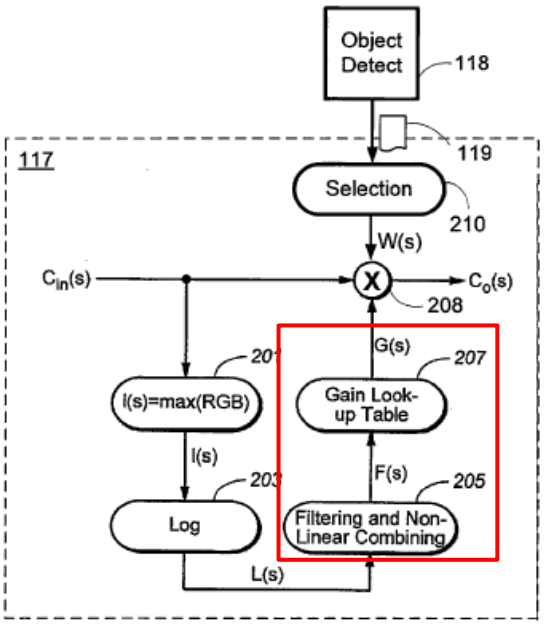

FIG. 2B

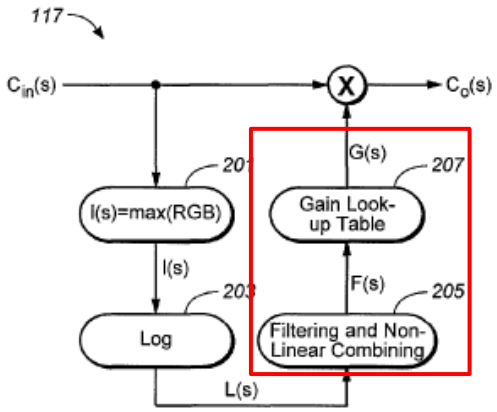

FIG. 2

'132 patent, FIG. 2B **Dvir, FIG. 2**

Compare APPLE-1001, FIG. 2B with APPLE-1007, FIG. 2; APPLE-1003, ¶¶16, 64, 68. I agree with Dr. Villasenor that, in the '132 patent's FIG. 2B, box 205 provides information to the gain look-up table 207, which the gain look-up table 207 uses to produce a gain that is applied on pixels in the digital image to change the pixel values, which can result in a change in pixel values in a first portion of the digital image. APPLE-1021, 112:23-115:8; APPLE-1001, 6:45-8:8. Because Dvir's FIG. 2 is exactly the same as the '132 patent's FIG. 2B with respect to blocks 205 and 207, Dvir discloses determining a gain amount "to maintain local contrast of the first portion of the digital image data" in the same manner as the '132 patent. *Compare* APPLE-1001, 6:45-8:15, FIG. 2B *with* APPLE-1007, [0019]-[0025], FIG. 2.

VIII. A POSITA would have found it obvious to combine Needham (alone or in view of Nonaka) with Dvir.

36. I understand that Qualcomm argues the purported "reengineering necessary," "add[ed] complexity," "trade-offs," and "degradation to the image" "would lead a POSITA away from combining Dvir's edge-preserving filters with Needham." POR, 42-43. However, neither Qualcomm nor Dr. Villasenor point to any teachings of the references to establish clear discouragement of the combination of Dvir with Needham. Additionally, Dvir does not lack definiteness or certainty about the result of using a cascade of nonlinear edge preserving filters for dynamic range correction and gives more than mere general guidance. *See generally* APPLE-1007; APPLE-1003, ¶¶64-68.

37. As I stated in my original declaration and repeated here, a POSITA would indeed have recognized that Dvir's teachings of dynamic range correction using a cascade of nonlinear edge preserving filters could have been combined with Needham's dynamic range correction to determine gain amounts, with Needham's weights controlling, in one way or another (e.g., before, during, or after Dvir's filtering), the gain amounts to apply to different portions of the image based on object type in an image. APPLE-1003, ¶¶52, 59, 64-72. In addition, Dr. Villasenor testified and I agree that it could be possible to apply the '132 patent's FIG. 2B structure, which includes Dvir's filters as block 205, to correct for contrast

enhancement (which Needham discloses corresponds to dynamic range correction). APPLE-1021, 122:7-125:1; *compare* APPLE-1001, 6:45-8:15, FIG. 2B *with* APPLE-1007, [0019]-[0025], FIG. 2; APPLE-1004, [0022]. .

IX. Gallagher discloses "applying the determined amount of gain to apply to the first portion of the digital image data comprises obtaining the amount of gain to apply to the first portion of the digital image data from the gain map" as required by claim 13.

38. I understand that Qualcomm argues that:

> In contrast to the gain map of claim 13, Gallagher discloses a "gain map" "used to determine the level of sharpening to apply on a pixel-by-pixel basis." APPLE-1006 at 6:26-31; Villasenor Decl. ¶¶108-110. Sharpening is an operation that enhances edge visibility, which is a separate aspect from the brightening or darkening of a region containing the edge. Villasenor Decl. ¶109. Because Gallagher's gain map is not used for controlling the brightening or darkening of a region of an image, Gallagher does not remedy Needham's deficiency in disclosing or suggesting the use of a gain map.

POR, 46. Under any interpretation of "gain map," Qualcomm's arguments are irrelevant and does not focus sufficiently on the teachings of Needham and Gallagher. Dr. Villasenor stated and I agree that both Needham and Gallagher teach adjusting pixel intensity values of digital image data:

> Q. …in the context of sharpness adjustment, is -- and what you've just described about identifying edges and manipulating

edges, is that pixel -- is that an adjustment of pixel intensity values?

A. Well, generally if you were going to enhance an edge, for example, you would have to change the pixel intensity value in order to achieve that goal. I mean, it depends on the context, but that would be normally what you would do.

APPLE-1021, 129:10-130:17 (emphasis added).

THE WITNESS:

A person of ordinary skill in the art would know that there's – that Needham does describe enhancing the contrast in that that would -- and that would involve changing, generally speaking, in changing the intensity of -- of the pixel value or more than one.

BY MR. AMON:

Q. Would a person of skill in the art know that Needham describes in enhancing the brightness and that that would involve changing -- generally changing the intensity of the pixel value or more than one pixel value?

A. Well, I mean, almost -- almost -- most operations, almost any operation involves changing intensities. It's -- it's -- you know, increasing the brightness would be a different, you know, correction. It's described as a different correction operation than enhancing the contrast, but in both cases you would change pixel values.

Q. And that would include changing pixel intensity for both brightness and contrast, fair?

A. Well, for brightness it certainly would, but -- but for

> contrast, you know, it would depend on how you did it; but -- but, like I said, almost any -- any -- any manipulation of images could involve some change to the pixel values.

APPLE-1021, 135:3-137:14 (emphasis added).

39. Gallagher is relied on for its general teaching of generating a gain map that indicates an amount of correction to be applied to image pixels and using the gain map to correct the image. As discussed in my original declaration and repeated here, a POSITA would have been led by Gallagher to modify Needham and/or the Needham-Nonaka combination to generate a gain map for dynamic range correction, and using the gain map to correct the image, as suggested by Gallagher. APPLE-1003, ¶¶80-83.

X. A POSITA would have found it obvious to combine Needham (alone or in view of Nonaka) with Gallagher.

40. I understand that Qualcomm argues that:

> However, Petitioner has not taken into consideration design incentives that would lead a POSITA away from combining Gallagher with Needham and/or Needham-Nonaka. Villasenor Decl. ¶¶101-103. For example, the goals of Gallagher are sharpening an image rather than feature-based enhancement of dynamic range contrast as in Needham. Villasenor Decl. ¶¶103, 108. There is no incentive to adapt a technique used for sharpening an entire image for use in enhancing dynamic range contrast of portions of an image. Villasenor Decl. ¶¶103-104.

> When considered as a whole, these design considerations weigh against combining Needham with Gallagher.

POR, 47. Here again, neither Qualcomm nor Dr. Villasenor point to any teachings of the references to establish clear discouragement of the combination of Gallagher with Needham. In addition, Gallagher does not lack definiteness or certainty about the result of using a gain map to correct an image and gives more than mere general guidance. *See generally* APPLE-1006; APPLE-1003, ¶¶75-79. As I stated in my original declaration and repeated here, a POSITA would have recognized that Gallagher's teachings of generating a gain map that indicates an amount of correction to be applied to image pixels and using the gain map to correct the image could have been combined with Needham's dynamic range correction to generate a gain map for dynamic range correction and use the gain map to correct the image. APPLE-1003, ¶¶80-81.

XI. CONCLUSION

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

Dated: 8/8/2019 By: 

Alan C. Bovik, Ph.D., P.E.