



US 20020120195A1

(19) **United States**(12) **Patent Application Publication****Hossack et al.**(10) **Pub. No.: US 2002/0120195 A1**(43) **Pub. Date: Aug. 29, 2002**(54) **MEDICAL DIAGNOSTIC ULTRASOUND
IMAGING METHODS FOR EXTENDED
FIELD OF VIEW****Related U.S. Application Data**

(60) Division of application No. 09/536,215, filed on Mar. 27, 2000, now Pat. No. 6,364,835, which is a continuation-in-part of application No. 09/196,986, filed on Nov. 20, 1998, and which is a continuation-in-part of application No. 09/384,707, filed on Aug. 26, 1999, now Pat. No. 6,352,508.

(75) Inventors: **John A. Hossack**, Charlottesville, VA (US); **Linyong Pang**, Stanford, CA (US); **Thilaka S. Sumanaweera**, San Jose, CA (US); **John W. Allison**, Los Altos, CA (US); **Kutay F. Ustuner**, Mountain View, CA (US); **Charles E. Bradley**, Burlingame, CA (US)

Publication Classification

(51) **Int. Cl.⁷** **A61B 8/00**
(52) **U.S. Cl.** **600/443**

Correspondence Address:

Attn: Elsa Keller
SIEMENS CORPORATION
Intellectual Property Department
186 Wood Avenue South
Iselin, NJ 08830 (US)

(73) Assignee: **Acuson Corporation**(21) Appl. No.: **10/081,662**(22) Filed: **Feb. 21, 2002**(57) **ABSTRACT**

A medical diagnostic ultrasound imaging system aligns substantially co-planar two-dimensional images to form an extended field of view using improved compounding methods. Compounding with a finite impulse response is used for more versatile compositing. The compounding is adaptive, such as through adapting the image regions, weighting, or type of compounding as a function of correlation, location within the image, estimated motion or combinations thereof. A user warning is provided as a function of the correlation between images.




Fig. 1




Fig. 2





Fig. 3




Fig. 4








Fig. 5








Fig. 9b





Fig. 11




Fig. 12

Fig. 13




FIG. 14





FIG. 15



MEDICAL DIAGNOSTIC ULTRASOUND IMAGING METHODS FOR EXTENDED FIELD OF VIEW

CROSS REFERENCE TO RELATED APPLICATION

[0001] This application is a continuation-in-part of copending U.S. patent application Ser. Nos. 09/196,986, filed Nov. 20, 1998, and 09/384,707, filed Aug. 26, 1999, which are hereby incorporated by reference in their entirety.

BACKGROUND

[0002] The present invention relates to medical diagnostic ultrasonic imaging, and in particular to improved methods used in connection with the combination of two or more partially overlapping images into an extended field of view image.

[0003] Hossack et al. in U.S. Pat. No. 6,014,473, filed Aug. 22, 1997, assigned to the assignee of the present invention and hereby incorporated by reference in its entirety, describe systems for acquiring, aligning and registering multiple medical diagnostic ultrasound images. Such alignment is used both to determine the motion between two selected images as well as to provide the information needed to composite an extended image from two or more selected ultrasound images. That is, two coplanar tracking images can be aligned and in this way the relative motion of the transducer between the times of the two tracking images can be obtained. Similarly, two or more substantially coplanar diagnostic images can be aligned and then composited to form an extended field of view.

[0004] In one embodiment, Hossack et al. disclose adaptively determining the number of image data frames that are collected between consecutive tracking frames (column 22, lines 18-50). The number of image data frames collected between tracking frames varies in response to the estimate of motion between each image data frame.

[0005] The collected frames are compounded. The compounding method described in U.S. Pat. No. 6,014,473 interpolates data from boundary portions of different image data frames with variable weights (column 29, lines 5-22). The weights applied to the image data frames vary linearly as a function of distance where the weights for one frame are one minus the weights for the other frame. Using these weights, a previous frame is compounded with a subsequent frame. The composition of the previous frame is not discussed, but Hossack et al. note that compounding can be used for accumulating image data.

[0006] Likewise, Weng et al. in U.S. Pat. No. 5,575,286 disclose a "ramp compounding" which gives weight ramps for both a new image frame and an existing compound image in the overlapping area (column 8, lines 21-30). Weng et al. do not discuss the composition of the existing compound image. Weng et al. also disclose alternatives to the ramp compounding, such as using only new pixels for non-overlapping regions or recursively averaging the new image frame with the existing compound image.

[0007] The present invention is directed in part to an improved compounding method that provides versatility and is quick to execute.

SUMMARY

[0008] By way of introduction, the preferred embodiments described below provide methods and systems for combin-

ing multiple frames of data for an extended field of view. Using finite-impulse response combinations, adaptive compounding or combinations thereof, the amount of speckle is reduced, and/or the signal-to-noise ratio is increased for the extended field of view images.

[0009] The present invention is defined by the following claims. This paragraph has been provided merely by way of introduction, and is not intended to define the scope of the following claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] FIG. 1 is a flow chart of a method that incorporates a presently preferred embodiment of this invention.

[0011] FIG. 2 is a schematic diagram of a parent image and a test block suitable for use in the method of FIG. 1.

[0012] FIG. 3 is a schematic diagram of an extended field of view generated by the method of FIG. 1.

[0013] FIGS. 4 and 5 are diagrammatic views of two different test blocks suitable for use in the method of FIG. 1.

[0014] FIGS. 6-8 are diagrammatic views of test blocks rotated by 0, 1 and 2 pixels, respectively.

[0015] FIGS. 9a and 9b combine to form a flow chart of another preferred embodiment of the method of this invention.

[0016] FIG. 10 is a flow chart of another preferred embodiment of the method of this invention.

[0017] FIGS. 11 and 12 are diagrammatic views showing the spatial relationship of B-mode and color Doppler regions in a parent image.

[0018] FIG. 13 is a flow chart of another preferred embodiment of the method of this invention.

[0019] FIG. 14 is a diagrammatic view showing one embodiment of a compounding mask.

[0020] FIG. 15 is a diagrammatic view showing another embodiment of a compounding mask.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0021] Various methods for estimating motion between two or more component images for forming an extended field of view image may be used. Some of these methods are discussed below. After an introduction to these motion estimation methods, compounding methods are discussed.

[0022] Turning now to the drawings, FIG. 1 shows a block diagram of one implementation of a method for estimating motion between two images. In a first step 10 at least two medical diagnostic ultrasound images are acquired, using any suitable technology. For example, the Sequoia or Aspen systems of Acuson Corporation can be used to acquire these two images in any suitable manner. As discussed in greater detail below, these images may be in any desired imaging mode, including B-mode, Color Doppler mode, and fundamental mode or harmonic mode (including harmonic imaging of contrast agent and harmonic imaging of tissue without

added contrast agent). On occasion, these originally acquired images will be referred to in this specification as parent images.

[0023] The two images are preferably substantially coplanar and partially overlapping, and can be obtained as separate frames acquired as a transducer probe is translated and potentially rotated in the XZ plane. The widest variety of transducer probes can be used, including conventional external probes as well as intracavity probes, such as transducer probes designed for insertion into the esophagus or other body cavities, intravascular and other catheter-based transducer probes, and intraoperative transducer probes (transducer probes designed for use during surgical operations in direct contact with internal organs).

[0024] In step 12, test blocks are selected from the images acquired in step 10. FIG. 2 shows one example, in which a test block 22 is shown as a central strip taken from a parent image 20. As described below, the test block 22 can be shaped as desired, and various techniques can be used to modify the test block as compared to the parent image. For example, the test block can be reduced in complexity by various filtering, decimation, and other techniques as described below. FIG. 2 shows the conventional X (azimuthal) and Z (range) dimensions as they relate to the parent image and the test block 22.

[0025] Returning to FIG. 1, the next step 14 is to find the translation in the XZ plane that matches the test blocks from the two images. As explained below, the match can be measured in many ways, including the minimum sum of absolute differences (MSAD) technique and various correlation techniques that utilize multiplications. Note that in step 14 the test blocks are merely translated with respect to one another and are not rotated. As described below, the location, range, and orientation of the translations that are used in step 14 can be selected adaptively to minimize the search time.

[0026] In step 16, the test blocks are aligned in translation using the matching translation of the preceding step and then rotated about a central axis 24 (FIG. 2) to find the angle of rotation that best matches the test blocks. As before, the angles of rotation that are searched can be selected adaptively based on previous searches to speed the search.

[0027] In the method of FIG. 1, the matching translation and rotation angle are found independently. This offers an advantage over a method that detects multiple translations from among several blocks and infers the rotation subsequently. In practice, motion may be irregular in either, or both, translation and rotation in a completely independent fashion. Typically, translation errors are of lesser consequence since they result in less serious cumulative geometric errors. Angular errors can have a more serious impact since the errors tend to accumulate and an angular error results in increasing translational errors in subsequently acquired frames. Therefore, it is preferable to independently control the search range for valid angular motions between frames.

[0028] Finally, in step 18 the matching translation and rotation from the previous steps are used as registration information to composite at least parts of the original parent images. FIG. 3 shows one example in which two parent images 20, 20' are composited with one another. The parent images 20, 20' overlap in an overlapping region 20'', and the registration information (ΔX , ΔZ , ΔO) is used to align the parent images 20, 20'.

[0029] The method of FIG. 1 preferably takes advantage of the nature of real motion of a transducer over the tissue. Typically, translation along the X (azimuth) axis is far greater than translation along the Z (range) axis; the transducer slides along the surface and is not pressed into the tissue.

[0030] Rotation effects are generally relatively modest compared to translation effects. This will certainly be true for the most common applications—such as scanning the abdominal surface or scanning along necks or legs.

[0031] The test block shape is preferably optimized to facilitate rotation determination. The block is made narrow in the azimuthal direction (e.g. 32 pixels along the X axis) and long in the range direction (e.g. 320 pixels—or almost the entire available image depth along the Z axis). The reason for this is that as the block is rotated by small amounts, the effect of rotation can be approximated by translating rows of pixels in the top and bottom regions of the test block. If the pixel block were square, one would have to translate the pixels in both X and Z, since the motion in the Z direction would be non-negligible.

[0032] FIGS. 4 and 5 illustrate this effect. In FIG. 4 the test block 22 is rotated by an angle θ . To first order, the motion of each row of pixels along the X-axis is simply translation along the X-axis, and none of the pixels of the test block 22 moves substantially in the Z direction. In contrast, when a wider test block 22' is used as shown in FIG. 5, a comparable rotation through the angle θ produces substantial motion along the Z-axis for pixels near the corners of the test block 22'. For this reason, the ratio of the range extent to the azimuthal extent of the test block is preferably greater than 5, more preferably greater than 9, and in the embodiments described below substantially equal to 10.

[0033] The search of step 14 is typically primarily in the azimuthal direction, but the search can also contain a small search (a few pixels) in the range direction. The search is preferably conducted firstly in the azimuthal direction only, followed by a fine scale, two-dimensional search in both X and Z.

[0034] The search of step 16 is most efficiently performed by generating a pixel offset table for each angular search, since high accuracy and high speed are both important. Also, small rotations are far more likely than large rotations. Therefore the rotations are typically in terms of single pixel motion increments at the extremes (top and bottom) of the test block.

[0035] Preferably, a pixel offset table is generated once when the program is started. A predefined angular search range is defined, e.g. -5 to $+5$ degrees in 0.2 degree steps. For every angular step one increments down the pixel block in the range direction and calculates the associated azimuthal (X) direction offset for the particular angle and pixel location in the range direction. The pixel offset is defined by the following equation:

$$\text{Round}((\text{pixel_index} * \tan(\theta))),$$

[0036] where Round is a function returning the nearest whole number, pixel_index is the index of the pixel in the range direction with respect to the center of the test block (i.e. the axis of rotation), and θ is the particular angle. In

this example, pixels are assumed to be square. If the pixels are not square, the above equation is modified to correct for the pixel aspect ratio, as for example as follows:

$$\text{Round}\left(\left(\frac{\text{pixelwidth}}{\text{pixelheight}}\right)\text{pixel_index}\tan(\theta)\right).$$

[0037] Consider as an example, an image which extends over 301 pixels in the range direction. The center pixel in the range direction is 151. Consider a pixel 10 pixels down in the range direction and a selected angle equal to +0.5 degrees. The pixel offset entry for this pixel depth and selected angle is:

$$\text{Round}((151-10)\tan(0.5^\circ))=\text{Round}(1.23)=1$$

[0038] This process is repeated for all pixel depths and rotation angles to fill the pixel offset table, and the process is performed only once per program run.

[0039] FIGS. 6, 7 and 8 illustrate three such pixel offset tables in simplified form. FIG. 6 shows a ten-pixel offset table appropriate for zero rotation. As shown at 26, the rotated test block is identical to the original test block. FIG. 7 shows an offset table for a rotation of one pixel in a clockwise direction about the axis 24. The motion ranges from ± 1 pixel to -1 pixel over the pixels 1 through 10, and the rotated test block is shown at 26'. FIG. 8 shows an offset table appropriate for rotation of 2 pixels in a clockwise direction, and the motion plotted in the table ranges from ± 2 . The rotated test block 26" shows the manner in which pixels are translated along the X direction, without translation along the Z direction. Similar tables can be created for anti-clockwise rotations.

[0040] For both steps 14 and 16 of FIG. 1, it is beneficial to calculate the motion to sub-step-size accuracy for both translation and rotation. This can be done by fitting a quadratic to the values around the MSAD minimum. Having found the equation for the fitting quadratic, the true minimum to sub-step-size accuracy is determined from the zero in the quadratic equation after differentiating. In this specification step size corresponds to the pixel search resolution for translation and the angular step resolution for rotation.

[0041] Once the matching translation and rotation are determined, the motions of all points along the central axis of the sub-block are known. If one considers the rotation as an indication of slope and the translation as a constant, we have the essential components for defining motion in terms of $y = mx + c$ where c is the translation amount, m is the slope or $\tan$ (angle of rotation), and in this case x is the range component and y is the translation amount at that range along the X or azimuthal axis. In this way, the motions of all points on the plane may be found in order to move one image into registration with another.

[0042] The extended field of view processing described above can be performed immediately after image acquisition and displayed on the ultrasound display of the acquisition system. Alternately, such extended field of view processing can be performed after the fact on a remote computer using stored image data.

[0043] Further details regarding preferred and alternative implementations of the method described above and related

methods are set out below, including methods for estimating motion and methods for compounding two or more component images to form an extended field of view of image.

[0044] 1. An iterative search can be used as set out below and as shown in FIG. 9.

[0045] Process

[0046] (a) In step 30, search first for a match in terms of translation (either X only or X and Z).

[0047] (b) In step 32, search for match in terms of rotation, taking account of the translation found in step (a).

[0048] (c) Taking account of the match from steps 30 and 32, determine in step 34 if a translation (X and/or Z) can be found that provides a better match.

[0049] (d) Taking account of the match for translation and rotation obtained from steps 30, 32 and 34, search in step 36 for a rotation that provides a better match.

[0050] (e) Repeat steps 34 and 36 using updated estimates of translation and rotation.

[0051] It steps 30, 32, 34 and 36, it is preferred to find a best match in each case. Alternately, in each of these steps a high speed search can be used that provides a better match, though the search is too fast or approximate to find the best match. As used herein, the term "match" is intended broadly to encompass best matches as well as improved matches.

[0052] The number of times steps 34 and 36 are repeated can be determined in various ways. In the illustrated embodiment, M (the number of times steps 34 and 36 have been executed) is compared with a constant, equal to 2 in this embodiment. The constant can be greater if desired. Alternately, steps 34 and 36 can be repeated until a measure of the quality of the match (such as the ratio of the minimum SAD to the mean SAD) reaches some predetermined acceptable level, or alternately until a larger number of repetitions has been completed.

[0053] By way of example, the test block width may be set to $16 \times N$, where N is a small integer. For $N=2$, this results in a width in the azimuthal direction of 32. The first translation matching step 30 of the iterative method of FIG. 9 can be in the X direction only, e.g. over a range of ± 32 pixels. The search region may be varied adaptively based on previous estimates of motion in the X direction, and the range may be made less than ± 32 pixels. The first rotation matching step 32 of the iterative method of FIG. 3 can be over a range of ± 3 degrees. The second translation matching step 34 can be over a range of ± 2 pixels in the X direction and ± 1 pixels in the Z direction.

[0054] The second rotation matching step 36 can be over a range of ± 0.5 degrees.

[0055] 2. The test block need not be tall and narrow, but as described above, such a shape for the test block provides advantages.

[0056] Consider a rectangular test block 30 pixels wide and 300 pixels high. Assume individual pixels are square. For this test block, for rotations up to 2.42 degrees there is no requirement for any pixel to shift in the Z direction by more than 0.5 pixels. Therefore, for rotations up to 2.42

degrees Z motion calculations can be ignored. Being able to ignore Z motion calculations has two advantages:

[0057] (a) Z motion calculations are not required; this saves computation time.

[0058] (b) Pixel rows in the reference frame (old frame) and the test position frame (new frame) both occupy consecutive memory locations (assuming pixels are ordered along rows in memory, that is, as a 2D array is read, pixel address first increments in X, followed by increments in Y after each row has been read). This ordering makes the programming simpler and faster. In particular it greatly facilitates the use of pointer-based programming and also makes the manipulations of blocks of pixels in parallel fashion far more tractable. For example, using Intel Pentium™ microprocessors with MMX™, this approach allows for manipulation of eight 8-bit binary numbers in parallel. The fact that pixels are grouped continuously in rows in both the old and the new frames facilitates easier implementation.

[0059] In contrast, if the pixel block is 300 pixels wide by 300 pixels high, then rotations of more than 0.19 degrees result in Z motions of more than 0.5 pixels in at least some pixels. Experience indicates that rotational motions between successive frames are frequently greater than 0.19 degrees but rarely greater than 2.42 degrees. In cases where Z motion is greater than 0.5 pixels, then it is still quite possible to reassign pixels from the Z shifted locations to the new frame pixel block so that an accurate rotation motion calculation can be performed.

[0060] 3. Color Doppler ultrasound image data (Velocity, Energy or combinations) may form part or all of the images used with this method. If both B-mode data and Color data are available for display, then the B-mode data for the entire pixel block can be used to determine image registration information, and the motion calculated for a particular B-mode region can also be applied to the related (superimposed) Color region. Color data, such as velocity or energy data, may also be compounded as discussed below to form the extended field of view image.

[0061] 4. Fundamental or harmonic data may form part or all of the images used with this method.

[0062] 5. Image data can be pre-processed prior to motion estimation and post-processed after composition (image registration). Examples include the following:

[0063] (a) Gray scale remapping, i.e. the as-acquired 0-255 range of gray scale values is remapped using a non straight line function.

[0064] (b) As an example of the above, contrast enhancement (e.g. histogram equalization) is used. As another example, 24 bit RGB data is decimated to 8 bit gray data.

[0065] (c) Pixel contrast resolution is decimated to speed up motion estimation, e.g. 0-255 pixel levels are mapped to 0-15, or even to 0-1 (i.e. only two binary levels) prior to motion estimation.

[0066] (d) Image brightness is altered. The resultant image is passed through a 0-255 mapping function

that increases the mean pixel level (but clips the upper end at 255 to prevent wrapping beyond 255 and back to low integer numbers (associated with near black)).

[0067] (e) Image data is filtered in some way (e.g. low-pass filtered).

[0068] (f) The output image is speckle reduced (e.g. low-pass filtered).

[0069] 6. The preferred motion merit function is typically MSAD (Minimum Sum of Absolute Differences); however, other matching techniques, such as finding the matching translation/rotation using cross-correlation (sum of multiplies) techniques may also be used.

[0070] 7. Sub-pixel estimation can be applied to translation, and sub-angular-step-size estimation can be applied to rotation calculations.

[0071] As an example, the levels of the sum of absolute differences (SAD) for the minimum and its neighbors are compared to determine the likely position of the minimum to sub-pixel resolution. A simple way to do this is to fit a quadratic equation ($y=ax^2+bx+c$) to three points (the pixel y_2 with the minimum SAD plus the next-neighboring pixels y_1 and y_3 on each side). The x values for these three data points are x_1 , x_2 and x_3 , and these x values are typically separated by unity (the pixel spacing for translation searches and the angular step size for rotation searches). Data for three points (x_1, y_1), (x_2, y_2), (x_3, y_3) are solved simultaneously for the equation $y=ax^2+bx+c$. In this way a , b and c (c not strictly needed) are found. The derived quadratic equation is differentiated and set equal to zero to solve for:

$$2ax+b=0.$$

[0072] In this way, the value of x to sub-step-size resolution is found.

[0073] 8. The pixel data used in the test block may be pre-scan-conversion acoustic line data or post-scan-conversion video data. The pixel data for the test block can be envelope detected, raw RF, baseband quadrature (I,Q), ideally processed to give coherency between adjacent lines as described in Wright U.S. Pat. No. 5,623,928, assigned to the assignee of the present invention. Likewise, the pixel data used for compositing the images may comprise one or more of these variations.

[0074] 9. The images used to create the test blocks may be acquired at the rate of one receive line per transmit line or at the rate of multiple receive lines per transmit line as described in Wright U.S. Pat. No. 5,685,308, assigned to the assignee of the present invention.

[0075] 10. As described in U.S. Pat. No. 6,014,473, poor estimates of motion can be identified. A poor estimate of motion may be determined based on a comparison of the level of the MSAD compared with the mean SAD. Also, a poor estimate of motion is often associated with a dissimilarity between the present estimate and previous estimates. As described in U.S.

[0076] Pat. No. 6,014,473, these various available confidence factors can be combined (e.g. using fuzzy logic) or used independently (e.g. if a given MSAD is unacceptable, a previous frame-to-frame motion estimate can be used).

[0077] As one example, the MSAD can be compared with a threshold to determine if the corresponding estimate of motion is reliable. One suitable threshold is equal to the number of pixels in the test block multiplied by a constant, such as 5 for example. In the event the MSAD for a particular estimate of motion is less than the threshold and is therefore indicative of a reliable estimate of motion, then the calculated estimates of motion (i.e. ΔX , ΔZ , and $\Delta\theta$) are used and three variables are updated as follows:

$$\begin{aligned}\text{history_}\Delta X &= (\text{history_}\Delta X + 3 * \Delta X) / 4, \\ \text{history_}\Delta Z &= (\text{history_}\Delta Z + 3 * \Delta Z) / 4, \\ \text{history_}\Delta\theta &= (\text{history_}\Delta\theta + 3 * \Delta\theta) / 4.\end{aligned}$$

[0078] The weights used in the weighted averages of the foregoing equations are merely examples of possible choices. Other alternatives can be used. In the event the MSAD associated with a particular estimate of motion is greater than the threshold, and therefore indicative of an unreliable estimate of motion, then the history variables are not updated and the estimated motion is replaced with values that vary predominantly with the history variables discussed above according to the following equations:

$$\begin{aligned}\Delta X &= (3 * \text{history_}\Delta X + \Delta X) / 4, \\ \Delta Z &= (3 * \text{history_}\Delta Z + \Delta Z) / 4, \\ \Delta\theta &= (3 * \text{history_}\Delta\theta + \Delta\theta) / 4.\end{aligned}$$

[0079] Once again, the weights used in the weighted averages are only examples.

[0080] 11. Frames being compared need not be consecutive. Instead of every frame being used, every Nth frame can be used (N=2,3 etc.).

[0081] 12. Frames being registered need not be consecutive. Instead of every frame being used, every Mth frame can be used (M=2,3 etc.).

[0082] 13. If every Nth frame is being used for motion estimation (N=2,3 etc.) then every Mth frame (M not necessarily equal N) can be used for registration or rendering.

[0083] 14. The rendered image regions that are composited to form the extended field of view need not necessarily comprise the entirety of the respective image. Preferably only a subset of the image is composited, such as an image block extending over the full Z dimension and over a width in the X direction that is slightly wider than the search area. Preferably, the central region of the image is composited, since it generally has the best resolution.

[0084] 15. The compositing method preferably includes the use of texture mapping, as found in the OpenGL Applications Programming Interface. See OpenGL Programming Guide (M. Woo et al.) published by Addison Wesley (1997). OpenGL is a trademark of Silicon Graphics.

[0085] 16. A visible or acoustic warning is preferably given to the user in event of a poor motion estimate or a series of poor motion estimates, as described in U.S. Pat. No. 6,014,473. Such a warning can indicate to a user that a re-scan is required.

[0086] 17. A user-visible icon (e.g. a variable length arrow) is preferably used to guide the user to make a scan at an optimal speed, as described in U.S. Pat. No. 6,014,473.

[0087] 18. The registration method can be used in combination with ECG and/or respiratory triggering/gating.

(Triggering is when acquisition is triggered by ECG or other events. Gating is when potential triggering events are analyzed to see if they are valid, i.e. they fit some predefined criteria.) For example, the selected images for registration alignment may be only those that are acquired in response to a selected trigger.

[0088] When forming the extended field of view of a pulsing object (e.g., an artery such as the carotid artery), it may be preferable to acquire only images during a particular portion of the cardiac cycle—such as end systole or some time interval after observed end systole. It is well known in ultrasound imaging to acquire an ECG signal using 3 ECG electrodes connected to the chest of the individual. The ECG signal exhibits some easily recognized peaks, such as the R wave. The R wave occurrence can be observed using an electronic trigger set to fire after each ECG voltage pulse surpassing some preset threshold is detected. This detected electronic trigger pulse can have a user-selectable delay applied to it and the resulting delayed electronic trigger used to trigger the start of an ultrasound image frame acquisition. Hence, ultrasound image frames are acquired once per heart cycle and at presettable delays from the detected R wave. It is also possible to use a gating technique wherein successive detected trigger pulses are compared with previously detected pulses and simple computer logic used to determine whether the ECG signal is gated or not. Typically, an ECG signal is gated only if the detected R wave interval falls within certain pre-programmed valid bounds. For example, suppose the R waves are occurring at 1000 mS intervals. A range of valid R waves can be set so that only R wave intervals in the range 900-1100 mS are gated. If one R wave signal is at 800 mS from the previous R wave, then it is associated with an irregular heart beat and ignored (since it may cause an image artifact due to irregular motion). This technique is discussed in McCann "Multidimensional Ultrasonic Imaging for Cardiology" Proc. IEEE Vol. 76, No. 9, September 1998, pp. 1063-1073.

[0089] 19. During image composition, pixel values are preferably interpolated or averaged based on the pixel levels in the overlapping regions of the two registered frames, as described in U.S. Pat. No. 6,014,473.

[0090] 20. The images are preferably ultrasonic medical diagnostic images acquired with 1D arrays employing a fixed focus in the elevation direction, one or more fixed foci in azimuth (transmit) and dynamic receive focus in azimuth. Ideally, a 1.5D array is employed that is focused in elevation, typically using a fixed transmit focus and a dynamically updated receive focus. Acuson Plano-Concave transducers as described in the related Hanafy U.S. Pat. No. 5,651,365, assigned to the assignee of this invention, can also be used.

[0091] Any suitable technology can be used for the transducer array, including piezoelectric ceramics (e.g. PZT), piezoelectric polymers (e.g. PVDF), relaxor ferroelectric ceramics (e.g. PMN-PT), and electrostatic transducers. Electrostatic transducers are described in "Micromachined Capacitive Transducer Arrays for Medical Ultrasound Imaging," X. C. Jin et al, Ultrasonic Symposium (Sendai, Japan, Oct. 5-8, 1988). The present inventors have recognized that such electrostatic transducers can be used in all conventional medical diagnostic ultrasonic imaging modes, including B-mode, Color Doppler mode, pulse wave mode, continuous wave mode, fundamental imaging mode, and harmonic imaging mode, with or without added contrast agent.

[0092] 21. Registration and composition processing can be performed on-line (on the ultrasound system processor) or can be done off-line (at an arbitrary time after acquisition). If performed off-line, data compression such as JPEG compression can be used to speed data transfer.

[0093] 22. The positioning of the search region for the test block can be varied adaptively: if the previous optimal value was found at 5 pixels to the right then the next search is preferably centered at 5 pixels to the right. Further details can be found in U.S. patent application Ser. No. 08/916,585.

[0094] For example, the current search can be centered on the offset estimated from the previous search. This is reasonable since the user typically uses a smooth motion in which a rapid change of velocity between successive frames is not expected. Suppose the first search is over the range ± 20 pixels. This is a 41 pixel search. However, if the optimum match (MSAD) is found at +10 pixels, then the subsequent search may be set to search from +5 pixels to +15 pixels (an 11 pixel search). In this way the search area is minimized and the overall speed improved. As well as changing the bias or center of the search, it is also preferred to vary the size of the search. If the successive searches are approximately uniform (e.g., 8, 7, 8, 9, 7 pixels all to the right), then it may be preferable to search the range +6 to +10. Alternately, if the successive searches are semi-random (e.g., 5, 7, 9, 3 pixels, all to the right), a better search range would be +1 to +15. Further, the size of the search can be made to vary according to the quality of the previous searches. If the ratio of MSAD (minimum SAD) to mean SAD is close to one, this means that the estimates are of poor quality, and that a high degree of uncertainty exists. In this case, larger search regions are preferred. Similar techniques can be applied to the rotational search.

[0095] 23. The size of the test block can be varied adaptively. A small test block that yields good results is preferred because processing time is low. Otherwise, a bigger test block should be used to improve quality (e.g. the ratio of MSAD to mean SAD, or the similarity with respect to previous estimates). The decision to increase or decrease the test block size can be made on the basis of any quality measure, including those mentioned above.

[0096] 24. Image resolution in the test block can be decimated (e.g. use only every Nth pixel in X and every Mth pixel in Z). The new pixel dimensions are then taken into account when determining motion in real units (mm). Preferably the image is low-pass filtered prior to decimation. Non-decimated images can be used for compositing even if decimated images are used for motion estimation.

[0097] 25. The size of the image motion test block can be adaptively altered based on lack of signal at depth. If the lower (deeper) part of the test block is noisy (randomly varying levels—not like acoustic speckle which follows a pattern determined in part by the acoustic system point spread function) or if the deeper range part of the image is black, there is lack of good signal at greater depths, and the size of the test block can be reduced to ignore this region. For any given transducer and frequency there is a range (which can be experimentally measured) beyond which the displayed signal is unreliable because the acoustic signal is approximately the same size as or less than the electronic noise threshold of the system. This range can be pre-computed for all transducers and operating conditions or

measured on test objects and then stored in a look up table. Whenever the system defines a test block to be used for motion estimation, it selects only pixels lying above (shallower) than the noise threshold region. An alternative method involves measuring acquired acoustic line data and detecting randomly varying signals. This can be performed using cross-correlation of successive signals, similar to the method used for Doppler detection of blood flow. If the lines are fired often enough, even if there is image motion the cross-correlation of signals will reveal a genuine detectable signal in the near field and no correlation in regions dominated by electronic noise.

[0098] 26. If using harmonic data, the subject either may or may not contain non-linear scattering agent. Further the composited image may use a combination of fundamental and harmonic in the near field and fundamental only in the far field.

[0099] 27. The rate of decorrelation can be used as an approximate estimate of transducer speed with respect to the imaged region and can be provided as a guide to the user. The rate of decorrelation can be derived using Doppler processors (correlators). Effectively, if a Doppler power signal is detected then there is motion. The size of this decorrelation signal is a function of speckle decorrelation.

[0100] Doppler cross-correlators can be used to correlate successive lines fired along one acoustic line direction. If the signal decorrelates by more than certain level (i.e., maximum cross-correlation level is below a threshold, e.g., 0.7), then this is an indication that the transducer has been moved too fast and the system displays a warning to the user. The advantage of this approach is that it allows the already-existing processing power of the ultrasound machine to be used to estimate whether the image acquisition is likely to be effective before the image acquisition is complete and the potentially time-consuming image transfer process has begun.

[0101] The rate of motion can be estimated using the Color Doppler processors already present in most ultrasound machines. This allows an indication to the user of poor speed control (i.e. transducer movement that is fast or out of plane) before the motion estimation process. In some cases, it is easier to do the motion estimation after complete image sequence collection and hence there is value in having an approximate estimate or warning of possibly non-optimal acquisition.

[0102] Effectively, the Doppler processors are used to measure pulse to pulse decorrelation. In one embodiment, beamformed 1,0 signals are applied to the B-mode processor and to the Color Doppler processor. Generally, different acquisitions are used for the B-mode and Color Doppler signals (Doppler signals are typically more narrowband). Once the baseband, beamformed signals are filtered, Doppler processing is performed on the I and Q signals to derive Doppler-related frequency estimates resulting from motion. Doppler processing is well known in the art. See for example "Real time two dimensional blood flow imaging using an autocorrelation technique," Kasai et al. *Tran. Sonics and Ultrasonics*, Volume SU-32, pages 458-464 (1985). See also the detailed discussion in Maslak, U.S. Pat. No. 5,555,534, assigned to the assignee of this invention, the disclosure of which is incorporated herein by reference.

[0103] The Color Doppler processor typically includes low-pass filters, clutter filters and an autocorrelator. The

clutter filters are typically formed as delay line cancellers, and they are used as low frequency rejection filters to eliminate large echo signals from stationary and slow moving objects which have low or zero Doppler frequency shift. The autocorrelator autocorrelates signals from the clutter filters and produces output signals for a Doppler Energy calculator, a Doppler Variance calculator and a Doppler Velocity calculator.

[0104] In the current application, the Doppler Energy signal is of greatest interest. Generally, a significant Doppler Energy signal may be associated with too rapid motion of the transducer relative to the region being imaged and used to generate a warning to the user. Similarly, a sudden motion will cause a color flash to appear. Color flash is well known in diagnostic ultrasound. Essentially, it is this flash indication which is being used here to highlight a sub-optimal motion. Since the application does not specifically require a determination of Doppler Energy over a 2D region (except when forming an extended view image of Doppler Energy images), a very narrow examination of the image may be made to determine if there is excessive Doppler Energy present. Typically one scans numerous acoustic lines to map out a 2D region. However, for the current application it is sufficient to fire Doppler associated acoustic lines in a single direction interleaved with the B-mode beams used to acquire the B-mode parent image. Ideally, these lines are oriented at a non-zero angle with respect to the normal to the transducer face.

[0105] This improves the Doppler angle for azimuthal motion. (Doppler processors detect the component of motion parallel to the beam and hence it is preferable to orient the beam so that it has at least some component in the azimuthal direction.) By experimentation, one determines the correct setting for the clutter filter to remove near DC components, and the Doppler Energy threshold above which the machine produces an indication that transducer velocity may be excessive. Typically, to minimize noise one can integrate the Doppler Energy signals derived over time as the Doppler processor determines the Doppler Energy level for increasing range associated with a Doppler line firing. (Conventionally, the Doppler Energy levels are compared to a threshold as they are acquired and if they exceed a threshold a colored pixel is drawn corresponding to the associated position in range and azimuth on the B-mode Image.) When Doppler Energy is used for determining sub-optimal motion, one can tailor the clutter filter setting and the energy threshold above which one assumes that non-optimal motion has occurred.

[0106] Experiments with different transducers and different frequencies can be performed to determine the relationship between speckle decorrelation (between two line firings at a known time interval along the same line axis) and speed of motion in the elevation direction. Experiments can also be performed to determine the optimal transducer speed (or maximum and maximum workable speeds) and these can be compared with the decorrelation values for different speeds. These values are preferably stored in a look up table in the system. During operation, the system takes account of transducer type and frequency and estimates speed of transducer motion from the measured line-to-line decorrelation value. If this speed is too high or too low, an indication is graphically displayed on the system screen. Alternatively, the estimated speed is displayed.

[0107] The output of the evaluation of the Doppler Energy determination can be presented to the user in a number of ways.

[0108] (a) A binary indicator can be displayed advising the user when sub-optimal transducer motion has been detected.

[0109] (b) A numerical output indicating relative Doppler Energy level can be displayed—either the maximum value for a particular frame to frame case or the mean Doppler Energy level over the entire scan.

[0110] (c) A Color Doppler Energy representation can be displayed on the screen in the conventional manner. In this case, a line-like Color Energy region appears on the image, and the user simply observes whether it flashes.

[0111] In addition to using the Doppler power or energy signal to detect poor motion estimates by means of the flash signal, it is also possible to use the Doppler velocity estimators to estimate transducer motion. Preferably, a Doppler velocity acquisition line is oriented at a non-zero angle with respect to a line oriented perpendicular to the transducer face (e.g. 45 degrees). When a velocity of for example 10 mm/s is detected, one can infer the scanning velocity in the pure azimuthal direction since we can assume that the actual transducer motion is parallel to the azimuthal direction. The velocity estimated for the transducer motion after correction for the Doppler angle (45 degrees) is $10 / \cos(45 \text{ degrees}) = 14.1 \text{ mm/s}$. This value can be output to the display numerically or graphically. Alternatively, the detected azimuthal velocity is compared to a threshold velocity (above which image alignment processing becomes problematic), and a warning is presented to the user by visual or audible means.

[0112] Similarly, the PW or CW Spectral Doppler processor can be used to estimate the velocity of the transducer probe relative to the tissue being imaged (the “sweep velocity”). The velocity is derived from the frequency component occurring with highest signal level from among the various frequency components examined.

[0113] In all cases, the Doppler parameters such as filter settings, pulse shapes and threshold levels are preferably optimized by modeling or experimental techniques as is well known in the art.

[0114] FIG. 10 provides an example of a method for using Doppler signals to estimate the rate of transducer motion in an image alignment method. As shown in FIG. 10, the first steps 50, 52 are to acquire multiple ultrasound images and to align the ultrasound images. Preferably, this alignment is performed as discussed above. However, for this aspect of the invention it should be understood that any suitable alignment method can be used, including those described in co-pending U.S. patent application Ser. No. 08/916,585 (U.S. Pat. No. 6,014,473) and in Weng U.S. Pat. No. 5,575,286, the disclosures of which are herein incorporated by reference.

[0115] The rate of transducer motion is estimated in step 54 from Doppler signals associated with the images. This motion estimate can be made as described above, and the motion estimate is then used in step 56 to signal an excessive rate of transducer motion to the user. As pointed out above,

this can be done by a visual or an audible alarm, or by indicating the actual estimate of transducer motion to the user visually.

[0116] In addition or alternative to providing a user alarm as a function of the estimate of motion, poor correlation indicates the desire to repeat the scan of the target as discussed above. Various measurements of correlation may be used, such as the decorrelation discussed above, a cross-correlation or the ratio of the minimum to the average sum of absolute differences as discussed below. For poorly correlated component images, an audible or visible alarm is provided to the user. The alarm is provided regardless of whether the associated estimate of motion is satisfactory or within operational constraints or settings.

[0117] 28. Pixel values can be interpolated between acquired values to yield a higher accuracy motion estimate (effectively giving sub-original pixel resolution). This interpolation process is well known and is often performed by the scan-converter. The interpolation may be linear (either in X or Y or both) or it may be curve-based. For example, a cubic spline can be fit to available data. Linear interpolation is usually adequate and is often faster. In practice sub-pixel estimation using the quadratic fit technique described above is often more efficient.

[0118] 29. When multiple, partially overlapping frames are combined after registration as described above, the combination can be performed in many ways. For example, pixel $I(i, j)$ of the combined image can be formed as a simple average:

$$I(i, j) = (I_{(i, j)}^{Frame1} + I_{(i, j)}^{Frame2} + I_{(i, j)}^{Frame3}) / 3.$$

[0119] Alternately, non-uniform weights can be used, e.g. [0.25, 0.5, 0.25]. Also, recursive compounding techniques (i.e. infinite impulse response (IIR) techniques) can be used:

$$I'_{(i, j)} = \alpha(I_{(i, j)}^{FrameN}) + (1 - \alpha)I_{(i, j)},$$

[0120] where I_{FrameN} is a newly acquired frame of data, $I(i, j)$ is the currently existing pixel data, and $I'(i, j)$ is the modified pixel data that takes into account the existing pixel data, the newly acquired data, and α . In general, α is less than one and may be equal to 0.3, for example. (If $I(i, j)$ has not been created, then $\alpha x = 1$). With this approach, each compounded frame overwrites or modifies the previously-compounded pixel values from earlier frames based on the new pixel values. The compounded image can then be optimized using histogram equalization to improve contrast or using the method described in Ustuner U.S. Pat. No. 5,479,926, assigned to the assignee of the present invention. Alternatively, self normalizing recursive weights can be used. $I'_{(i, j)} = [I_{(i, j)}^{FrameN} + nI_{(i, j)}] / (n+1)$. In this example, pixel (i, j) comprises a summation of pixels from N component images, divided by N . To obtain the above equation, α is set to $1/(n+1)$.

[0121] As shown above in the simple average example, the images are combined in one embodiment as a function

of finite impulse response (FIR) filtering. As used herein, FIR and IIR filtering comprise any device for implementing FIR or IIR compounding techniques (i.e. implementing FIR or IIR filtering). For FIR filtering, a device provides weighted averaging of a finite number of component images (e.g. frames of data). For example, the combined extended field of view image is free of recursive compounding for the overlapping region. No one of the compounded images comprises a previously compounded image, but the images may have been previously used in other compounding.

[0122] In the three images example above, the data from each image for a pixel is multiplied by equal (e.g. 1) or unequal weights, summed and then divided by the number of images (i.e. averaged or weighted summation). This pixel of the extended field of view image (i.e. compounded image) is calculated once by averaging all the pixels from different frames coincident at this pixel location (i.e. not recursive). FIR compounding reduces persistence of undesired signals, such as noise, artifacts, or flash, associated with recursive compounding.

[0123] In one embodiment, the weights for FIR or IIR compounding are different for different portions of the component image being compounded. For example, the weights applied to data for a component image are different as a function of location within the image. Side portions may be emphasized over a center portion of an image by applying greater weights for pixels representing the side portions and lesser weights for pixels representing the center portion. For Vector® or sector images, the side portion decorrelate faster than center portions, so are preferably emphasized. Various weight curves as a function of location within the image may be used, such as experimentally determined curves depending on the application. Alternatively, center portions or other portions of the images may be emphasized. In one embodiment for FIR compounding, the overlap of the images is such that the weights total a value of one for each pixel location within the extended field of view image even with weights that vary for each component image as a function of location.

[0124] Alternatively, a portion of one or more of the component images is clipped or eliminated (e.g. a zero weight is applied to a portion of the component image). Such variations as discussed above for estimating motion using different image portions may also be used for compounding the images. For example, only side portions, such as sides of images in the direction of motion, are compounded. Where a same weight is applied to each image, the contribution of some images is emphasized where pixels from one of the images are not included in the average. A different degree of compounding is provided. In one embodiment, the subset of an image used for compounding comprises a block of pixel data representing a thin azimuth portion and a substantially full range portion (e.g. 32 pixels in azimuth and 350 pixels in range from a 480 by 640 pixel image).

[0125] FIG. 14 shows one embodiment of a mask for clipping component images. The mask applied for selecting or clipping the data is shaped like a 5 milk carton. The milk carton shape is defined by a (1) top width, (2) bottom width, (3) angle of the sides and (4) height of the milk carton. Other shapes may be used, such as square, rectangular, symmetric, asymmetric, circular, oval or polygonal shapes. The milk carton shape of FIG. 14 is oriented so that the height

corresponds to the range dimension. Other orientations may be used. The height of the milk carton mask comprises a total depth of the component image. Heights less than the total depth may be used.

[0126] The mask of **FIG. 14** is applied to clip data for all or a subset of the component images. In one embodiment, a different mask is applied to a subset of the component images. For example, the mask shown in **FIG. 14** is applied to the component images compounding to form a center portion of the extended field of view image. For component images associated with edge or side portions of the extended field of view image, one of the side portions of one or more component images is selected to comprise the side portion of the extended field of view image. For the left side of the extended field of view image, the left portion of the component images is maintained by the mask shown in **FIG. 15**. For the right side of the extended field of view image, the right portion of the associated component images is maintained by the a mask, such as the mirror image of the mask of **FIG. 15**. The mask of **FIG. 15** is shaped for use with Vector® scan formats, but other scan formats and associated masks may be used.

[0127] Since the component images represent different transverse locations of the target, fewer pixels are available to compound for the edge portions of the extended field of view image. To avoid having the extended field of view image vary in appearance (e.g. generation of compounding related bands), the mask varies as a function of relative position of the component images. For the first component image (i.e. the component image with the leftmost transverse location), the compounding mask comprises the mask shown in **FIG. 15**.

[0128] The mask applied to adjacent component frames varies, such as by making the line on the left side of the mask more vertical. The angle of the line is more vertical for component images further away from the side portion of the extended field of view image. The mask of **FIG. 15** is varied as a function of the position of the component image until the mask of **FIG. 14** is applied. This smooth transition between masks avoids processing artifacts.

[0129] The mask variation process is repeated for component images associated with the other side of the extended field of view image.

[0130] In one embodiment, one or more component images are eliminated to change the effective weighting applied to other images. One or more pixels of the extended field of view image is compounded from a plurality of component images. For example, 10 component images per pixel are compounding for most or all of the extended field of view image. To change the effective weighting, one or more of the component images is eliminated. Each remaining image contributes more heavily to the extended field of view image.

[0131] The values of the weights applied are adjusted to sum to one, for FIR compounding or sum to a value less than one.

[0132] As an example of eliminating a component image, component images are decimated in response to the estimated motion of successive images. If the next component image is associated with a small amount of transducer movement, this component image is eliminated.

[0133] In one example, two variables, CumulativeCx and CumulativeTx, are set to zero. The motion between two component images is estimated, such as estimating azimuthal and range translations and an angle of rotation from the middle of a tracking block. An estimate of the azimuthal translation is also estimated for the top-middle portion of the tracking block. CumulativeCx is set equal to CumulativeCx plus the x translation from the middle (Cx) of the tracking block, and CumulativeTx is set equal to CumulativeTx plus the x translation from the top-middle (Tx) of the tracking block. The maximum of the absolute value of CumulativeCx and CumulativeTx is determined. If this maximum is above a threshold, the component image is included in the extended field of view image and CumulativeCx and CumulativeTx are reset to 0.0. If the maximum is below the threshold, the associated component image is not included in the extended field of view image. This process is repeated for each possible component image. Translation in range may additionally or alternatively be examined to determine whether to eliminate a possible component image.

[0134] As discussed above, OpenGL software and/or hardware is used in one embodiment to provide low cost and efficient weighted summation compounding (FIR filter). A single image buffer configured to include an opacity channel and color channels accumulates alpha blended values. Where a personal computer, motherboard, off-line processor or extra processor within a system performs calculations for the extended field of view image, OpenGL hardware and software are cheaply implemented. For example, see U.S. Pat. No. _____ (U.S. Ser. No. 09/196,207, filed Nov. 20, 1998), the disclosure of which is incorporated herein by reference, which discloses integrating a Perspective™ or other computer within an ultrasound system housing. Substantially real-time or quick compounding is provided with commercially available OpenGL accelerator cards. Alpha blending avoids the need for prior input regarding the number of component images to be compounded. FIR compounding is provided without substantial memory required to separately store each of the component frames prior to compounding.

[0135] For alpha blending with OpenGL, the color of a source pixel, represented as an R_S, G_S, B_S value, is combined with the color of a destination pixel. The source pixel comprises a pixel from a component image, and the destination pixel represents a pixel from another component image or a pixel that is the partial sum of a plurality of other component images. The destination pixel comprises the R_D, G_D, B_D values stored in the image buffer. In addition to RGB values, the source and destination are further represented by alpha or opacity values, A_S and A_D .

[0136] To compound the source pixel with the destination pixel, source and destination factors are specified. The factors comprises a pair of RGBA quadruplets (S_R, S_G, S_B, S_A and D_R, D_G, D_B, D_A) to be multiplied by respective components of the source and destination pixels. This blending is represented by:

$$(R_S S_R + R_D D_R, G_S S_G + G_D D_G, B_S S_B + B_D D_B, A_S S_A + A_D D_A).$$

[0137] Each component of this combined quadruplet is clamped to [0,1].

[0138] As a default, OpenGL combines the source and destination pixels by adding the component parts (see pages

214-219 of Open GL Programming Guide, Second Edition, Woo et al., ISBN 0-201-46138-2). The quadruplet factors are set to one as the default. Alpha blending provided by OpenGL processing performs equal or unequal weighting of the component images. The weighting is set by selecting values for S_R S_G S_B S_A and D_R D_G D_B D_A .

[0139] For equal weighting, the components of each of these quadruplets is be equal, such as set to one or set to a specific alpha value. For example, S_R S_G S_B $S_A=1, 1, 1, 1$ and D_R D_G D_B $D_A=1, 1, 1, 1$. As another example, S_R S_G S_B $S_A=a_r, a_g, a_b, a_a$ and D_R D_G D_B $D_A=1, 1, 1, 1$.

[0140] For compounding, the initial framebuffer or destination values are set to 0. The source pixel of each component image is iteratively blended with the destination values. For example, let S_R S_G S_B $S_A=1, 1, 1, 1$ and D_R D_G D_B $D_A=1, 1, 1, 1$ where N component images are compounded. The opacity of each source pixel is set to 1 in one example, but other values may be used. A table representing the R (red) and A (alpha) channels of the destination is shown below:

TABLE 1

R_d (Red Channel in the I framebuffer)	A_d (Alpha Channel in the framebuffer)
0 0	0
1 $R_s^{(1)} + 0 \cdot 0$	$1^2 + 0$
2 $R_s^{(2)} + (R_s^{(1)})$	$1^2 + 1^2$
3 $R_s^{(3)} + (R_s^{(2)} + R_s^{(1)})$	$1^2 + (1^2 + 1^2)$
4 $R_s^{(4)} + (R_s^{(3)} + (R_s^{(2)} + R_s^{(1)}))$	$1^2 + (1^2 + (1^2 + 1^2))$
N $\sum_{i=1}^n R_s^{(i)}$	$n1^2$

[0141] This blending for the R channel is represented mathematically as:

$$R_c = \frac{1}{n} \sum_{i=1}^n R_s^{(i)}$$

[0142] Substituting for destination quantities, the blending function simplifies to:

$$R_c = \left(\frac{1}{n} \right) \left(\sum_{i=1}^n R_s^{(i)} \right) = \frac{R_d}{A_d}$$

[0143] To obtain the equal weighted compound value for the extended field of view image, the RGB values of the destination are divided by the destination alpha value.

[0144] Where the source factor quadruplet S_R S_G S_B S_A are set to a_r, a_g, a_b, a_a where a is not equal to one or to scale the result, the RGB values as calculated above are multiplied by the respective "a" value or scale value (e.g. $R_x = aR_d/A_d$). To avoid clamping where a is not equal to one, the value of "a" is small, such as 0.05 for compounding 20 source pixels (i.e. 20 component images overlapping at that pixel).

[0145] For unequal weighting, the components of each of these quadruplets are set to be unequal. Where the S_R S_G S_B

S_A and D_R D_G D_B D_A values are arbitrary or unequal, the destination pixel is represented mathematically as:

$$R_d = \frac{\sum_{i=1}^n R_s^{(i)} a_i}{\sum_{(s)} a_i}$$

[0146] Where α_1 is the weighting at that pixel from image i. Before alpha blending as discussed above, each pixel is converted as a function of the pixel's alpha value α_1 . For example, R_s is converted to $\sqrt{\alpha_1} R_s$. In one example, S_R, S_G, S_B, S_1 is set as $\sqrt{\alpha_i}, \sqrt{\alpha_i}, \sqrt{\alpha_i}, \sqrt{\alpha_i}$, and D_R, D_G, D_B, D_A is set 1, 1, 1, 1.

[0147] The table below represents alpha blending the red and alpha channels using the converted pixel information:

TABLE 2

R_d (Red Channel in the I framebuffer)	A_d (Alpha Channel in the framebuffer)
0 0	0
1 $\sum_s^n \alpha_i$	$\sum_s^n \alpha_i$
2 $\sum_s^n \alpha_i$	$\alpha_2 + \alpha_1$
3 $\sum_s^n \alpha_i$	$\alpha_3 + (\alpha_2 + \alpha_1)$
4 $\sum_s^n \alpha_i$	$\alpha_4 + (\alpha_3 + (\alpha_2 + \alpha_1))$
N $\sum_s^n \alpha_i$	$\sum_s^n \alpha_i$

[0148] Substituting for calculated quantities, the blending function simplifies to:

$$R_c = \frac{R_d}{A_d}$$

[0149] To obtain the unequal weighted compound value for the extended field of view image, the source pixel values are converted as a function of respective weights, and the RGB values of the destination are divided by the destination alpha value.

[0150] In one embodiment, the pixel data is scaled down to avoid saturation of the frame or destination buffer. For example, each R, G or B value is within the range of 0 to 255, and the frame buffer is also 8 bits per channel. If the R, G or B values are large and/or the number of component images is large, the frame buffer may saturate.

[0151] To avoid saturation, only the higher order N bits (downshifted by M-N bits where M is the depth of each

color channel (for example, $M=8$) of the RGB channels are used in the component images during compounding. For example, if only the highest 4 bits are used, the scaled down component pixel data range from 0 to 15. Other numbers of bits, for example, from 2 to 7, may also be used. Each overlapping pixels of these component images are then accumulated and multiplied by the quantity, q , given by:

$$q=2^{m-1}/2^{n-1}$$

[0152] where α is the residual value in the alpha channel at the pixel. Although the information content is reduced in the component images due to scaling down, dithering during compounding component images restores some or all of the lost information in the final compounded image.

[0153] In addition to color saturation, the alpha channel may also saturate if the number of component images is sufficiently large. This situation is unlikely in most practical cases.

[0154] In another embodiment, the accumulation buffers defined in OpenGL are used for compounding. Since these buffers are typically wider (for example 16 bits) than the color buffers (for example 8 bits), saturation artifacts are less likely in the final compounded image.

[0155] To further avoid saturation of any channel, the number of component images used for compounding is controlled as discussed above. Since OpenGL alpha blending operates without a previous knowledge of the number of frames to be compounded, decimation of component images is implemented without adverse effects on compounding. Masking or clipping portions of component images as discussed above also reduces the amount of information at a pixel to avoid saturation.

[0156] In one embodiment for compounding with OpenGL alpha blending, a plurality of variables are controlled to avoid saturation and to generate an clinically acceptable extended field of view image. For example, a number of bits to downshift, a top width of a mask, a bottom width of the mask, an angle of the mask, a number of component images used for compounding the ends of the final compounded image, and a threshold amount of motion used for automatically skipping or eliminating a possible component image are selected. In one embodiment, the threshold used for eliminating possible component images is selected as a fraction of the top width of the mask, such as 0.5 of the width. Other variable may be used, such as any of the variable discussed herein. The optimum variables are determined as a function of the clinical application and system. Multiple settings for any give application and system may be provided for user selection. The set of parameters is independent of the probe geometry and type, but may account for these factors. Other techniques and Application Programming Interfaces that support alpha blending may also be used.

[0157] Characteristics of the component images may be optimized for compounding through variation. One component image has different characteristics than another component image. For example, two component images are characterized by different acquisition frequencies. Acquisition frequencies include transmit frequencies, receive frequencies and combinations thereof. Compounding component images associated with different acquisition frequencies reduces the effects of speckle. Other character-

istics for variation include aperture size (transmit or receive), filter settings (transmit or receive), focusing parameters or other variables controlled by the acquisition ultrasound system or a post-acquisition system.

[0158] The variation is repetitive in one embodiment. Every other component image is similar, so that two different types of component images are provided. Other cycles of repetition or number of types of component images may be used, such as every third or fourth component image is similar in characteristics. A randomized process may be used.

[0159] In one embodiment, motion is estimated between images with similar characteristics. The speckle patterns of such images is likely more correlated for more accurate motion estimation. Motion is estimated for other component images by interpolation, curve fitting, or separately. Alternatively, motion is estimated using images not compounded to form the extended field of view image.

[0160] The same or different component images may be processed differently for compounding than for estimating motion. For example, two component images are selected. Motion is estimated between the two selected component images without any, without further or with different filtering. The two component images are filtered, such as low pass spatial filtering, for compounding. Other types of filtering, such as high pass or band pass filtering, or smoothing processes, such as histogram equalization or other gray scale mapping modifications (e.g. changing the shape of the gray scale curve), may be used. The filtering is applied on line data (one dimension) or the image data (two dimensions). Other smoothing or filtering processes may be used, such as disclosed in U.S. Pat. No. 5,479,926, the disclosure of which is herein incorporated by reference. In one embodiment, the filtering or other processing highlights specular reflectors or other image attributes. In other embodiments, the component images are filtered for estimating motion, but not filtered or filtered differently for compounding, as discussed above.

[0161] As also discussed above and in U.S. Pat. _____ No. (Ser. No. 09/384,707), the disclosure of which is incorporated herein by reference, one or more of the component images are transformed or scaled to compensate for scanning rate error. Scanning rate error compensation spatially corrects the component images and the amount of detected frame to frame motion to account for transducer motion while image data is acquired for a single scan. This correction avoids inaccurate motion estimation and resulting compounding of improperly aligned data. By compensating for the scanning rate as a function of the frame rate, the extended field of view image more accurately represents the scanned target.

[0162] The FIR or IIR compounding may be adaptive. Adaptive compounding comprises changing a characteristic of the compounding in response to one or more features. For example, (1) the width or other dimension of a component image used for compounding (e.g. the mask), (2) the weight or weights, (3) the type of compounding (e.g. IIR, FIR or no compounding), or (4) other compounding characteristic is changed in response to (a) an estimation or estimations of motion between two or more component images, (b) a correlation or correlations between two or more component

images, (c) image pixel data, (d) a location of the pixel within the component image, (e) other calculated features or (f) combinations thereof.

[0163] The mask applied for compounding is adaptive in one embodiment. As discussed above for compounding, a portion of a component image is compounded. The portion of the image is selected as discussed above for compounding or for selecting a tracking block.

[0164] The width, height, size, shape or combinations thereof of the portion or mask is determined in response to an estimate of motion. The image texture width (e.g. for azimuthal translation) is set so that a set number of component images are compounded for a given pixel. For example, if the estimated motion is 10 pixels to the right azimuthally and 10 component images are to be compounded, a 100 pixel wide portion of the component images is compounded. For any given azimuthal pixel location, 10 component images are compounded. The 10 right most lines of pixel data from the first frame overlap with the 10 left most lines of pixel data from the tenth frame. Where the estimated motion varies, the mask is varied to account for the changing motion. The actual number of component images compounded may vary from the set number, such as 10. Additionally or alternatively, the number of images for compounding is set by eliminating component images.

[0165] As another alternative or additionally, the weights or opacity used for IIR or FIR compounding may adapt in response to estimated motion or correlation between two or more component images. The varied weights effectively reduce the influence of some component images, resulting in an extended field of view image that is similar to an image formed by compounding a set number of component images. The same weight is used for each pixel of the component image or varies for different pixels of the component image. The image texture width or mask is fixed or also varies as discussed above.

[0166] For IIR compounding, the recursive weight or opacity is adapted as a function of the estimated motion. A table or other listing of possible weights as a function of the estimated motion is determined based on the system and application. Typically, a high motion estimate is associated with a higher weight, α , applied to the most recent component image and with a lower weight (i.e. $1-\alpha$) applied to the previously composited image. Where the estimate of motion varies between component images, the weights are varied correspondingly. In other embodiments, FIR compounding weights adapt to the estimates of motion.

[0167] The IIR weights or FIR weights may adapt as a function of the location of the pixel being compounded within the component image or the extended field of view image. For example and as discussed above, lesser weights are applied to pixels associated with azimuthally center portions of an image.

[0168] The weights may also adapt as a function of a combination of the location of the pixel being compounded and a correlation or estimate of motion. Local estimates of motion or correlations are calculated. The weighting applied for that local area varies as a function of the local estimate of motion or correlation.

[0169] Varying weights as a function of the correlation between images avoids blurring the extended field of view

image. Correlation indicates the quality of a match between aligned images. The correlation may be determined as a function of unaligned images.

[0170] If the component images poorly correlate, less compounding or more unequal weights are applied. In the extreme, the pixel data is overwritten rather than compounded. If the component images have a high correlation, more equal weighting is applied. For example, if the normalized cross-correlation coefficient is 0.5, the opacity (i.e. IIR or FIR weight) is set to 0.6, and if the normalized cross-correlation coefficient is 0.9, the opacity is set to 0.1. Other values may be used, such as values representing a curve of opacity to cross-correlation coefficient determined empirically.

[0171] Compounding as a function of correlation is independent of compounding as a function of estimated motion. The correlation value used comprises a cross-correlation value, an inverse of a decorrelation value, a minimum sum of absolute differences or other value indicating a similarity between component images. For example, the minimum sum of absolute differences is compared to the average sum of absolute differences. This comparison is less computationally intensive to calculate than a cross-correlation, especially where the minimum sum of absolute differences is calculated to estimate motion. A curve or table of weights as a function of a ratio or other function relating the minimum to average sum of absolute differences is determined empirically. For example, a ratio of 0.5 is associated with an IIR opacity or alpha value for OpenGL FIR compounding of 0.4, and a ratio of 0.8 is associated with an IIR opacity of 0.8. Other values may be used. In one embodiment, correlation, such as speckle decorrelation, is used to estimate out-of-plane movement (elevational movement) and also used to adapt the degree of compounding.

[0172] As an alternative to adapting the weight or weights as a function of correlation, the number of component images compounded is adapted as a function of correlation. Varying the number of component images compounded effectively varies the weights. For a high rate of decorrelation, fewer images are compounded together. Alternatively, the mask size, shape or application is varied as a function of the correlation.

[0173] As discussed above, characteristics of the compounding may adapt as a function of a change in the correlation. The magnitude of the change determines the weights, image mask or type of compounding.

[0174] In one embodiment, the type of compounding is adapted to one or more of the features discussed herein. For example, one of IIR, FIR or no compounding is used for one or more pixels in response to correlation.

[0175] Other processes may be adaptive. For example, the selection of the component images is adaptive. The frequency selection of possible component images within a sequence adapts as a function of estimated motion. For example, every other or every third possible component image is selected where the amount of motion is low. Possible component images are more frequently selected where the amount of motion is high. As another example, the frequency of the selection of possible component images decreases for high correlation between the component images. Highly correlated component images, such as a

minimum to average sum of absolute differences ratio less than 0.1 or a cross-correlation coefficient higher than 0.9, provide less speckle reduction through compounding, so images less likely highly correlated (e.g. every other, third or fourth component image) are selected. The selected images are then compounded or adaptively compounded as discussed above. In one embodiment, the compounded extended field of view image is processed after compounding but prior to generation of a responsive display. For example, histogram equalization, high pass filtering, low pass filtering, adaptive filtering (e.g. filtering as disclosed in U.S. Pat. No. 5,478,926), gray scale curve remapping or other signal processing is performed on the compounded extended field of view image.

[0176] FIG. 13 provides a flow chart of an imaging method that utilizes the image acquisition, alignment, and composition techniques described above to create an extended longitudinal section of a tubular organ such as a blood vessel. In step 70 a catheter-mounted probe is introduced into a tubular organ such as a vessel, typically by inserting the probe into a lumen of the vessel. For example, the probe can include an ultrasonic transducer array mounted on a catheter for insertion into a human vein or artery. In step 72 ultrasound images are acquired as the probe is moved along the vessel, either into or out of the body of the subject. In step 72 the images are acquired in respective image planes that are oriented generally parallel to the path along which the probe is moved. In step 74 the images acquired in step 72 are aligned, using any of the alignment techniques now known or later developed by those skilled in the art. In step 76 the aligned images are composited, and in step 78 the extended field of view image as composited in step 76 is displayed. The composited image of step 76 is in this example an extended longitudinal sectional view of the vessel.

[0177] The techniques described above allow for viewing of long sections of vessels and surrounding anatomy and also allow for viewing landmark tissue to allow the user to orient a region of interest, which may be a small area in many situations, in a context of a larger image including such landmark tissue. Examples of suitable catheter-mounted transducer probes are described in U.S. patent applications Ser. Nos. 08/791,598 and 08/802,621, both assigned to the assignee of the present invention and incorporated by reference herein.

[0178] 31. When an extended view image is being acquired using a low frame rate and a relatively high manual transducer sweep rate across the tissue surface, a dimensional error is encountered. This is because the ultrasound system typically scans the array from one side to the other side of the transducer array. If the transducer array is 50 mm wide, the frame rate is 10 frames per second, and the transducer is scanned at a rate of 5 mm/s, then in the time it takes to scan one frame (0.1 s) the transducer has moved 0.5 mm. When the transducer array is scanned in one direction this error corresponds to a 0.5 mm image extension and when scanned in the other direction it corresponds to a 0.5 mm image contraction in the azimuthal motion. (This is a 1% dimensional error.) This error scales with scan rate and inversely with frame rate. During extended field imaging, we can determine the image motion and hence it is possible to correct for this effect.

[0179] Consider the case where the system scans from transducer element #1 (at one end of the array) to transducer element #N (at the other end of the array). If the motion estimator detects a motion between two frames of 1 mm in the direction of element 1, then the real image region which has been scanned has contracted by 1 mm. Hence we must scale the width of the region by a scale factor K, where

$$K=1-(1\text{ mm/Image width(mm)}),$$

$$K=1-1/50=0.98\text{ for a }50\text{ mm wide image.}$$

[0180] Similarly, for motion in the other direction, the scale factor K is set equal to $1+1/50$ to extend the width of the imaged region. The motion estimate (translation) is preferably corrected by the same factor, i.e. the estimated translation is multiplied by 0.98 in the above example for transducer motion in the direction of element #1. See the discussion of image correction in U.S. patent application Ser. No. 08/916,358, assigned to the assignee of the present invention and hereby incorporated by reference herein.

[0181] The motion of the probe also affects the determination of rotation, and estimates of rotation should also be modified. Typically, the angular estimate is derived from the angle associated with the relative motion of pixels in the azimuthal direction at the top and bottom of the search block. This angle must now be corrected to take account of the fact that the X dimensions of the pixels are different from that originally assumed. For the small angles of interest here, this can be effected approximately by multiplying the angles produced before trying to determine the transducer probe velocity error by the factor determined above for the correction to the image azimuthal dimension and motion estimate, i.e. the estimated rotation is multiplied by 0.98 in the example above for transducer motion toward transducer element #1.

[0182] The simplest way to implement this correction is to modify the texture mapping width when doing the image compositing using OpenGL texture mapping.

[0183] 32. When compositing an image including Color Doppler data (velocity or energy), it is preferable to acquire separate B-mode and Color regions as is conventionally done inside the system. Typically, an ultrasound machine acquires a complete frame of B-mode and a partial frame of Color. Depending on the level of Color Energy and/or Velocity, particular pixels are overwritten or mixed with color values. Color regions do not in themselves contain the speckle data used in motion estimation of the type described above. Therefore, it is preferable to use the raw B-mode images for motion estimation prior to the Color superimposing step.

[0184] On occasion, only the color superimposed images are available for estimating transducer motion. There are a number of methods that can be used, as follows:

[0185] (a) The Color pan boxes can be restricted in size to a small part of the parent image. The color pixel data, which contains no speckle motion data, does not corrupt the motion estimate to any significant extent.

[0186] (b) A color superimposing step can be used which mixes Doppler Color and B-mode speckle. Thereafter, the B-mode speckle is processed by referring to the Color lookup table and inferring the

underlying B-mode value. If necessary, the B-mode values are rescaled if they were scaled during the color mapping process.

[0187] (c) The Doppler Color regions can be restricted to specified regions of the parent image. Typically only a portion of the parent image is Doppler processed since Doppler processing is associated with slow data acquisition and associated slow frame rates. Preferably, the Color pan box is limited to only a portion of the available image, and preferably the Color pan box is prohibited from the center of the test block which is being used for motion estimation. Although this may sound unduly burdensome, it is not necessary that the majority of the image be in color since only small portions of the image are used in composition to form an extended field of view. If these small regions are color regions, then an extended color image will still result. This assumes that the composited regions are different from the region used for motion estimation. In this process, motion is estimated for the B-mode region. One then takes account of the difference between the center of the motion estimation block and the center of the color composited region during image composition. The simplest way to do this is to define the composited region to include the center of the motion estimation region but to define that the portion of the image composition region outside that containing color as completely transparent. OpenGL™ texture mapping allows for controlled transparency and hence this technique is easily implemented.

[0188] (d) The Color pan box region alternately can be allowed to occur anywhere in the frame as long as it does not occupy the entire frame. Prior to motion estimation, the program uses knowledge of the location of the color regions (which may be passed to it from the image acquisition controller of the ultrasound machine), and adaptively alters the position of the motion estimation block to position it on one or other side of the Color region. In this way, the region for the motion estimation process is spaced from the Doppler Color region of the parent image. Moving the Color box during acquisition can be prohibited, and to simplify implementation, this adaptive step is performed only at the beginning of the process.

[0189] FIG. 11 shows one preferred relative arrangement of regions within a parent image. In this case, the parent image 60 includes two B-mode regions 62, 64 and a Color Doppler region 66. The B-mode regions 62, 64 are well suited for motion estimation using any of the techniques described above, and the relative positions of the B-mode regions 62, 64 and the color region 66 are well known. As shown in FIG. 12, when two parent images 60, 60' of the type shown in FIG. 11 are composited, the B-mode regions 62, 62' can be used for motion estimation and alignment of the two parent images 60, 60', while the color regions 66, 66' can be used for extending the field of view.

[0190] 33. Typically, the test block is smaller than the entire parent image, but it is also possible to use the entire image as the test block. In this case, it is preferred to account for the fact that there is incomplete matching data for any

non-zero motion (i.e. if the frames move from left to right, then one frame will lack data for comparison on the left hand side and the other will lack data on the right hand side). In the MSAD calculation (or cross correlation), it is important that valid data be present. Therefore, when using a full frame test block, one preferably suppresses SAD calculations for invalid regions. However, in order to calculate the minimum SAD, one preferably uses some form of normalization to take account of the fact that some comparisons use more pixels (since they are less overlapping). One approach is to replace the search for minimum SAD with a search for minimum normalized SAD, where the normalized SAD operation normalizes the SAD by dividing the SAD by the number of pixel-to-pixel comparisons being used.

[0191] As an example, consider a full image block of 100 by 100 pixels, where the search is ± 3 pixels (left and right).

[0192] When calculating the SAD with a shift of three pixels to left, the three pixels on the left are invalid (not defined) on one frame, and three pixels on the right are invalid (not defined) on other frame. Therefore, the number of comparisons is $(100-3-3)*100=9400$. In this case, the normalized SAD is equal to the SAD divided by 9400.

[0193] When calculating the SAD with a zero pixel shift between the test blocks, no pixels on the left are invalid (not defined) on one frame and no pixels on the right are invalid (not defined) on the other frame. Therefore, the number of comparisons is $(100-0-0)*100=10,000$. In this case, the normalized SAD is equal to SAD divided by 10,000. Of course, a full-frame test block may be sub-sampled as described above, e.g. decimated.

[0194] 34. If image acquisition is not ECG gated (as is most often the case), then it is preferable to take steps to minimize perceived pulsatility in the resulting color Doppler extended field of view image. (Pulsatility will appear as a color region with rapidly changing width as a function of azimuthal position and the respective component image). One approach is to persist the Color Doppler data by separating it from the underlying B-mode and then persisting it during the compositing step. If the Color Doppler data is not available independently from the B-mode data from the ultrasound image processor, then the Color data can be separated from the combined Color+Gray scale (B-mode) data by detecting regions in the images in which R, G and B are not balanced (RGB balanced=gray). These regions (identified as Color) are then persisted such that the maximum Color value detected for any overlapping region during composition overwrites any previous Color value or the Gray value if no Color value already exists at that point. Alternatively, the 10 Color data is separated into a separate Color image zone comprising Color values and a high opacity (possibly opacity=1), and non Color values (everywhere else) with opacity=0. During OpenGL texture mapping, the Color regions will overwrite such that the peak flow condition for any particular region will dominate and perceived pulsatility is suppressed. When using this technique, it is preferable that the Color map used is continuous and exhibits a monotonic type variation for increasing Color Doppler Energy or Color Doppler Unsigned Velocity. (If signed velocity is used then dramatically changing colors corresponding to changes in flow direction may result in an undesirable image with blotches of different color.) Other embodiments for separating Color data are disclosed in U.S.

Pat. No. _____ (U.S. Ser. No. 09/370,060, filed Aug. 6, 1999), the disclosure of which is incorporated herein by reference.

CONCLUSION

[0195] The improvements and modifications described in sections 1-34 above can be used with any suitable algorithm for registering and/or compounding two or more substantially coplanar ultrasound images, such as taught in U.S. Pat. Nos. 6,014,473, 5,575,286, 5,566,674, _____ (U.S. Ser. No. 09/384,707) and _____ (U.S. Ser. No. 09/196,987), the disclosures of which are herein incorporated by reference. The embodiments described above may be implemented on an ultrasound imaging system or on an offline processor. Various preset selections of appropriately or experimentally determined groups of compounding parameters, such as the types of compounding and/or variables discussed herein, are provided to the user. For example, different mask shapes, mask sizes and degrees of compounding are provided to the user for each of a plurality of selectable clinical applications or studies.

[0196] Once an extended field of view image has been formed as described above, it can be manipulated in well known ways, as for example by zooming or rotating using OpenGL. All of the techniques described herein may be used with tracking image frames of the type discussed in U.S. patent application Ser. No. 08/916,585. The foregoing detailed description has described only a few of the many forms that this invention can take. For this reason, this detailed description is intended by way of illustration and not by way of limitation. It is the following claims, including all equivalents, that are intended to define the scope of this invention.

1. A medical diagnostic ultrasound method for forming an extended field of view of a target, the method comprising the acts of:

- (a) selecting first and second medical ultrasonic images, the first and second images partially overlapping;
- (b) compounding the first image with the second image in an overlapping region in response to a finite impulse response function; and
- (c) generating an extended field of view image responsive to (b);

wherein at least a portion of the extended field of view image corresponding to the overlapping region is free of recursive compounding.

2. A medical diagnostic ultrasound method for forming an extended field of view of a target, the method comprising the acts of:

- (a) selecting at least three medical ultrasonic images, the first and second images partially overlapping; and
- (b) compounding the at least three images in an overlapping region in response to a finite impulse response function.

3. The method of claims 1 or 2 wherein (b) comprises alpha blending.

4. The method of claim 3 wherein (b) comprises using OpenGL commands.

5. The method of claims 1 or 2 wherein (b) comprises averaging.

6. The method of claim 5 wherein (b) comprises averaging with non-uniform weights.

7. The method of claims 1 or 2 wherein (b) comprises applying weights to the images, the weights larger for portions of the images spaced from the center than for center portions.

8. The method of claims 1 or 2 wherein (b) comprises compounding side portions of the images without compounding center portions.

9. The method of claims 1 or 2 wherein (a) comprises selecting images responsive to different frequencies.

10. The method of claims 2 further comprising:

- (c) generating an extended field of view image responsive to (b).

11. The method of claims 1 or 10 further comprising:

- (d) signal processing the compounded information prior to (c);

wherein the extended field of view image is responsive to the signal processing.

12. The method of claims 1 or 2 wherein (b) is adaptive as a function of motion.

13. The method of claims 1 or 2 wherein (b) is adaptive as a function of correlation.

14. The method of claims 1 or 2 further comprising compensating for scanning rate error.

15. A medical diagnostic ultrasound method for forming an extended field of view of a target, the method comprising the acts of:

- (a) selecting first and second medical ultrasonic images, the first and second images partially overlapping; and
- (b) alpha blending the first image with the second image in an overlapping region.

16. The method of claim 15 wherein (b) comprises using OpenGL commands.

17. The method of claim 15 further comprising:

- (c) performing (b) with an OpenGL accelerator.

18. The method of claim 15 wherein (b) comprises blending image data as a function of an alpha value.

19. The method of claim 18 further comprising:

- (c) blending alpha values for each frame of image data; and
- (d) dividing the blended image data of (b) with the blended alpha values of (c).

20. The method of claim 19 further comprising:

- (e) multiplying the output of (d) by a value.

21. The method of claim 15 further comprising:

- (c) using a sub-set of bits comprising image data for (b), the sub-set comprising higher bits.

22. The method of claim 15 wherein (a) comprises selecting images as a function of an estimate of motion.

23. The method of claim 15 further comprising:

- (c) masking the first and second images;

wherein (b) is responsive to the masked first and second images.

24. The method of claim 23 further comprising:

- (d) applying a different mask to the first image than to the second image.

25. The method of claim 24 wherein (d) comprises varying the mask as a function of a image position within an extended field of view image.

26. A medical diagnostic ultrasound method for forming an extended field of view of a target, the method comprising the acts of:

- (a) selecting first and second medical ultrasonic images, the first and second images partially overlapping;
- (b) compounding the first image with the second image in an overlapping region; and
- (c) emphasizing side portions of the first and second images in (b).

27. The method of claim 26 wherein (c) comprises applying a greater weight for the side portions than for center portions.

28. The method of claim 26 wherein (b) and (c) comprise compounding only side portions of the first and second images.

29. A medical diagnostic ultrasound method for forming an extended field of view of a target, the method comprising the acts of:

- (a) selecting first and second medical ultrasonic images, the first and second images partially overlapping and characterized by different receive frequencies;
- (b) compounding the first image with the second image in an overlapping region; and
- (c) generating an extended field of view image responsive to (b).

30. A medical diagnostic ultrasound method for forming an extended field of view of a target, the method comprising the acts of:

- (a) selecting first and second medical ultrasonic images, the first and second images partially overlapping;
- (b) filtering the first and second images;
- (c) generating an extended field of view image from the filtered first and second images; and
- (d) estimating motion from the unfiltered first and second images.

31. A medical diagnostic ultrasound method for forming an extended field of view of a target, the method comprising the acts of:

- (a) selecting first and second medical ultrasonic images, the first and second images partially overlapping;
- (b) adaptively compounding the first image with the second image in an overlapping region; and
- (c) generating an extended field of view image responsive to (b).

32. The method of claim 31 wherein (b) comprises adaptively changing an image width as a function of estimated motion.

33. The method of claim 32 wherein (b) comprises selecting the image width as a function of a number of images to be compounded.

34. The method of claim 31 wherein (b) comprises adaptively changing a weight as a function of estimated motion.

35. The method of claim 34 wherein (b) comprises compounding the first image with the second image recursively; and

wherein the weight comprises a recursive weight.

36. The method of claim 34 further comprising:

- (d) changing the weight as a function of location within the first image.

37. The method of claim 31 wherein (b) comprises adaptively changing an opacity as a function of estimated motion.

38. The method of claim 31 wherein (b) comprises adaptively changing a weight as a function of a correlation between the first and second images.

39. The method of claim 38 wherein (b) comprises applying higher weights for lower correlations.

40. The method of claim 38 wherein (b) comprises adaptively changing the weight as a function of a minimum sum of absolute differences and an average sum of absolute differences.

41. The method of claim 31 wherein (b) comprises adaptively changing a weight as a function of location in the first image.

42. The method of claim 41 wherein (b) comprises applying greater weight to side portions of the first image.

43. The method of claim 31 wherein (b) comprises changing a type of compounding.

44. The method of claim 43 wherein (b) comprises at least two types of compounding selected from the group consisting of: no compounding, finite impulse response compounding and recursive compounding.

45. The method of claim 31 wherein (b) comprises adaptively compounding as a function of a correlation between the first and second images.

46. The method of claim 45 further comprising:

- (d) determining the correlation as a function of lower frequency components of the first and second images.

47. The method of claim 45 further comprising:

- (d) determining the correlation for a plurality of locations in the overlapping region; and

wherein (b) comprises changing the compounding as a function of the correlation at each of the plurality of locations.

48. The method of claim 45 wherein (b) comprises adaptively compounding as a function of a change in the correlation.

49. The method of claim 31 wherein (b) comprises adaptively compounding as a function of estimated motion.

50. The method of claim 49 wherein (b) further comprises adaptively compounding as a function of correlation.

51. The method of claim 31 wherein (a) comprises adaptively selecting the first and second images from a plurality of images as a function of estimated motion.

52. The method of claim 51 wherein (a) comprises selecting the first image separated by a greater number of images from the second image for smaller motion estimates.

53. The method of claim 31 wherein (b) comprises compounding with a finite impulse response.

54. The method of claim 53 wherein (b) comprises compounding in response to OpenGL commands.

55. The method of claim 53 wherein (b) comprises compounding with alpha blending.

- 56.** The method of claim 31 further comprising:
- (d) scaling the first image as a function of a scan rate and a frame rate.
- 57.** The method of claim 31 further comprising:
- (d) estimating motion between the first and second images; and
- (e) applying different processing to the first and second images for (d) than for (b).
- 58.** The method of claim 57 wherein (e) comprises filtering the first and second images for (b) and maintaining the first and second images substantially free of filtering for (d).
- 59.** The method of claim 31 further comprising:
- (d) determining a correlation between the first and second images; and
- (d) presenting to a user an alarm indicative of poor correlation.
- 60.** The method of claim 31 wherein (a) comprises selecting the first image responsive to a different frequency than the second image.
- 61.** The method of claim 60 further comprising:
- (d) estimating motion between the third and fourth images, the third and fourth images responsive to a substantially same frequency; and
- (e) determining motion between the first and second images as a function of the estimated motion of (d).
- 62.** The method of claim 31 further comprising:
- (d) signal processing the compounded information prior to (c);
- wherein the extended field of view image is responsive to the signal processing.
- 63.** A medical diagnostic ultrasound method for extended field of view data collection, the method comprising the acts of:
- (a) selecting first and second medical ultrasonic images, the first and second images partially overlapping;
- (b) determining a correlation between the first and second images; and
- (c) presenting to a user an alarm indicative of poor correlation in response to the correlation.
- 64.** The method of claim 63 wherein (c) comprises generating an audible sound.
- 65.** The method of claim 63 wherein (c) comprises requesting a repeat scan.
- 66.** The method of claim 63 wherein (b) comprises calculating a minimum sum of absolute differences and an average sum of absolute differences between the first and second images.
- 67.** The method of claim 31 further comprising:
- (d) estimating motion between the first and second images; and
- (e) scaling the estimated motion as a function of a scan rate and a frame rate.
- 68.** The method of claim 1 or 2 wherein the first and second images comprise coherently adjacent line data.
- 69.** A medical diagnostic ultrasound method for forming an extended field of view of a target, the method comprising the acts of:
- (a) selecting first and second medical ultrasonic images, the first and second images partially overlapping and comprising coherently adjacent line data;
- (b) compounding the first image with the second image in an overlapping region; and
- (c) generating an extended field of view image responsive to (b).
- 70.** The method of claim 1, 2, 15 and 31 further comprising providing a plurality of user selectable applications wherein (b) is responsive to user selection of one of the selectable applications.

* * * * *

专利名称(译)	用于扩展视野的医学诊断超声成像方法		
公开(公告)号	US20020120195A1	公开(公告)日	2002-08-29
申请号	US10/081662	申请日	2002-02-21
[标]申请(专利权)人(译)	阿库森公司		
申请(专利权)人(译)	ACUSON CORPORATION		
当前申请(专利权)人(译)	ACUSON CORPORATION		
[标]发明人	HOSSACK JOHN A PANG LINYONG SUMANAWEEERA THILAKA S ALLISON JOHN W USTUNER KUTAY F BRADLEY CHARLES E		
发明人	HOSSACK, JOHN A. PANG, LINYONG SUMANAWEEERA, THILAKA S. ALLISON, JOHN W. USTUNER, KUTAY F. BRADLEY, CHARLES E.		
IPC分类号	A61B8/06 A61B8/12 G01S7/52 G01S15/89 G06T3/40 G06T7/00 A61B8/00		
CPC分类号	A61B8/06 A61B8/12 A61B8/13 G01S7/52034 G01S7/5205 G01S7/5206 G01S7/52065 G01S7/52071 G01S7/52074 G01S7/52085 G01S15/8927 G01S15/8979 G06T3/0075 G06T7/0026 G06T2207/10132 G06T2207/30004 G06T2207/30101 G06T7/32		
其他公开文献	US6641536		
外部链接	Espacenet USPTO		

摘要(译)

医学诊断超声成像系统使用改进的复合方法对准基本上共面的二维图像以形成扩展的视野。具有有限脉冲响应的复合用于更通用的合成。复合是自适应的，例如通过调整图像区域，加权或复合类型作为相关性，图像内的位置，估计的运动或其组合的函数。根据图像之间的相关性提供用户警告。

