



US 20100145210A1

(19) **United States**

(12) **Patent Application Publication**
Graff et al.

(10) **Pub. No.: US 2010/0145210 A1**

(43) **Pub. Date: Jun. 10, 2010**

(54) **MULTI PARAMETRIC CLASSIFICATION OF
CARDIOVASCULAR SOUNDS**

(76) Inventors: **Claus Graff, Klarup (DK); Samuel
Schmidt, Aalborg (DK); Johannes
Struijk, Terndrup (DK)**

Correspondence Address:
**JACOBSON HOLMAN PLLC
400 SEVENTH STREET N.W., SUITE 600
WASHINGTON, DC 20004 (US)**

(21) Appl. No.: **12/308,755**

(22) PCT Filed: **Jun. 26, 2006**


(86) PCT No.: **PCT/DK2006/000374**

§ 371 (c)(1),
(2), (4) Date: **Oct. 5, 2009**

Publication Classification

(51) **Int. Cl.**
A61B 5/02 (2006.01)
(52) **U.S. Cl.** **600/528**
(57) **ABSTRACT**

The present application relates to a method for classifying a cardiovascular sound recorded from a living subject. The method comprises the step of extracting at least two signal parameters (309) from said cardiovascular sound, said at least two signal parameters characterizes at least two different properties of at least a part of said cardiovascular sound. The method further comprises the step of classifying said cardiovascular sound using said at least two signal parameters in a multivariate classification method (310). Furthermore, the application relates to a system, stethoscope and server for classifying a cardiovascular sound recorded from a living subject, where the above-described method has been implemented.






Fig. 1




Fig. 2




Fig. 3




Fig. 4a




Fig. 4b

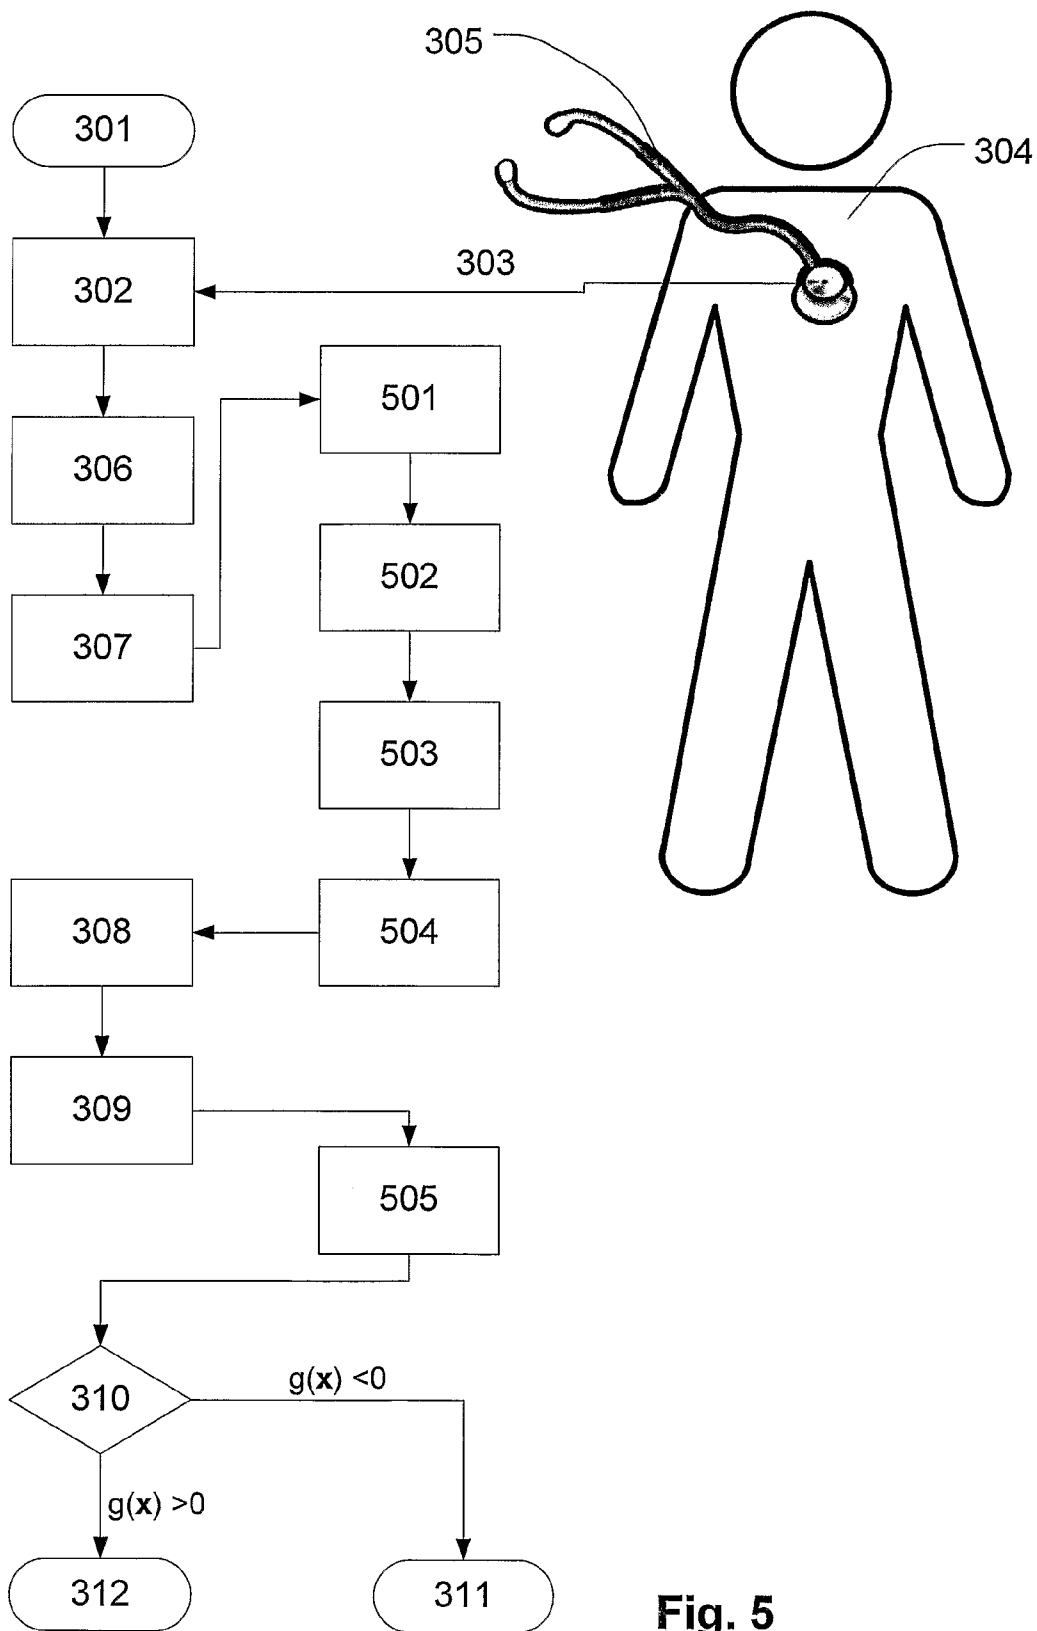


Fig. 5

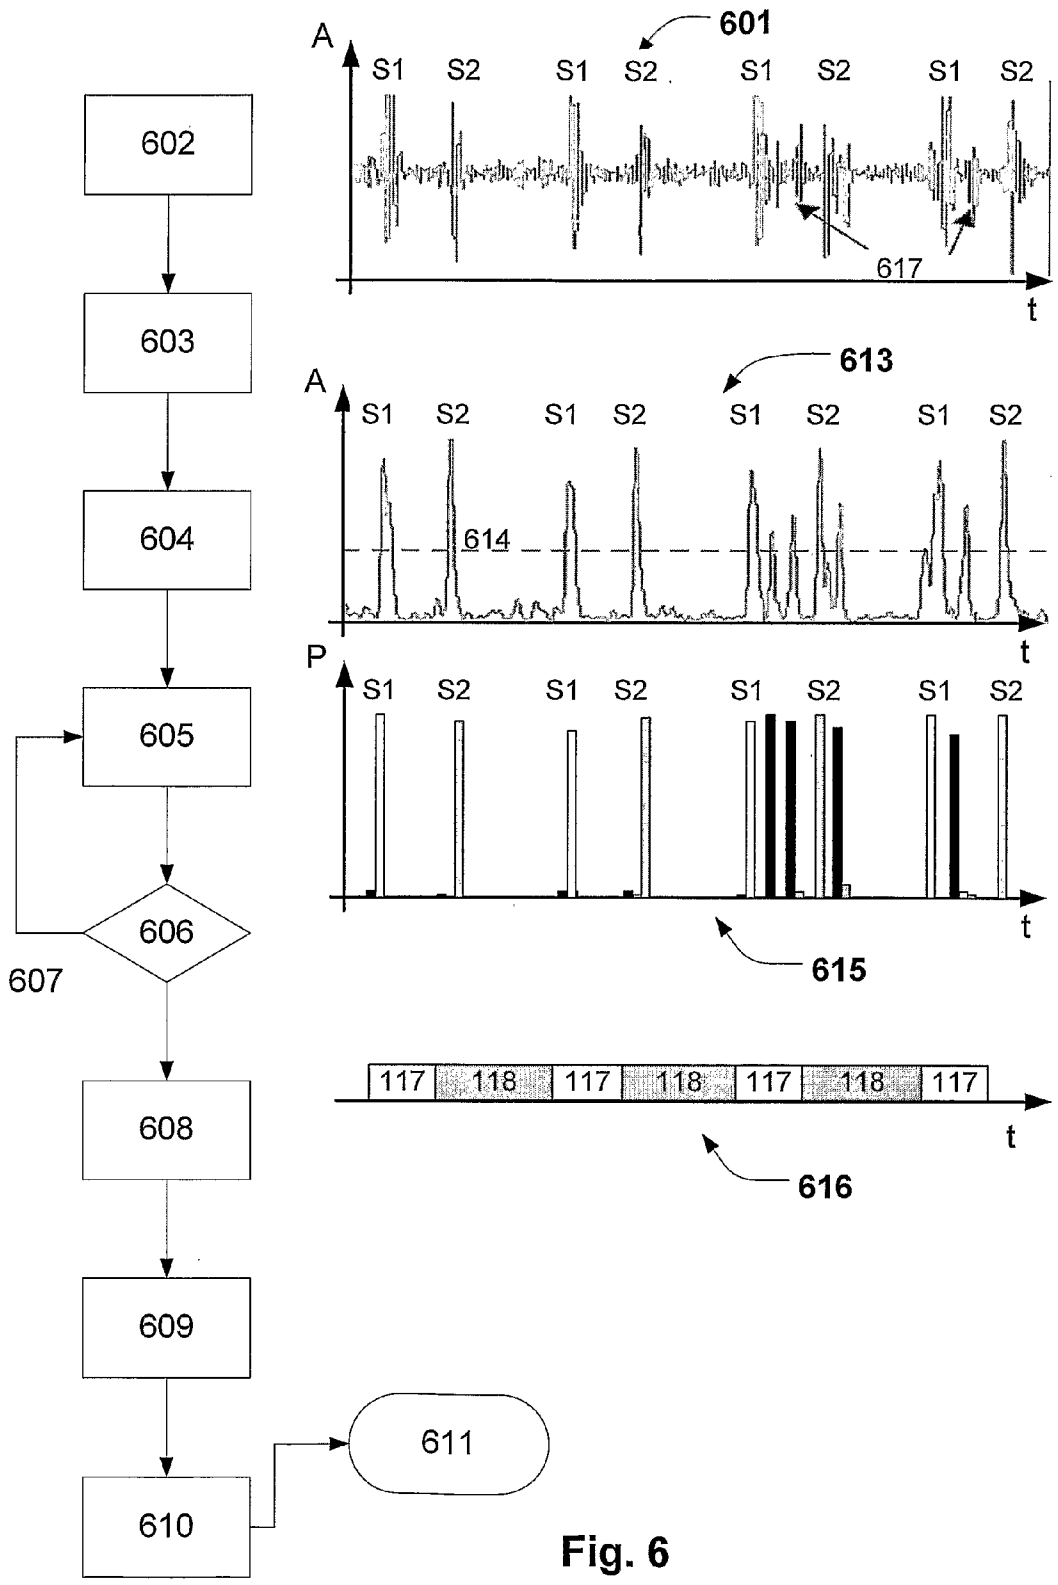


Fig. 6




Fig. 7a




Fig. 7b




Fig. 7c



Fig. 7d




Fig. 8

MULTI PARAMETRIC CLASSIFICATION OF CARDIOVASCULAR SOUNDS

FIELD OF THE INVENTION

[0001] The present invention relates to methods and systems for classification of heart sounds recorded from a living subject into classes describing whether or not murmurs due to coronary artery stenosis is present in the heart sound.

BACKGROUND OF THE INVENTION

[0002] Coronary artery disease is the single most common cause of death from cardiovascular disease in the western world. The heart muscle receives its blood supply through the coronary arteries, and atherosclerosis is the most common pathophysiological process occurring in the coronary arteries giving rise to coronary artery disease (CAD). Atherosclerosis is a process that builds up plaques within the artery, and the blood flow can therefore be reduced or even blocked by the plaque. The constantly working heart requires a continuous and efficient blood supply in order to work properly. Defects in the blood supply may be very severe and even fatal. Increasing degrees of luminal diameter reduction or stenosis of the coronary artery will first limit reserve flow, then reduce flow at rest and may finally totally occlude the vessel.

[0003] There is a need for measuring/detecting coronary artery stenosis for clinicians and other medical professionals to diagnose CAD. Once a diagnose has been made a cure/treatment could be started.

[0004] Today several non-invasive techniques for measuring/detecting the severity of a stenosis or its presence inside a coronary artery exist. This can be done by magnetic resonance imaging (MRI), in vivo intravascular ultrasound (IVUS) or optical coherence tomography (OCT). However, the above-mentioned techniques are all rather complicated and expensive to use and therefore only patients with specific symptoms are offered such examinations. The consequence is that most patients have a critical stenosis when examined.

[0005] Clinicians and other medical professionals have long relied on auscultatory sounds to aid in the detection and diagnosis of physiological conditions. For instance, a clinician may utilize a stethoscope to monitor and record heart sounds in order to detect heart valve diseases. Furthermore, the recorded heart sounds could be digitized, saved and stored as data files for later analysis. Devices have been developed that apply algorithms to electronically recorded auscultatory sounds. One example is an automated blood-pressure monitoring device. Other examples include analysis systems that attempt to automatically detect physiological conditions based on the analysis of auscultatory sounds. For instance, artificial neural networks have been discussed as one possible mechanism for analyzing auscultatory sounds and providing an automated diagnosis or suggested diagnosis. Using these conventional techniques, it is difficult to provide an automated device for diagnosis of coronary stenosis using auscultatory sounds. Moreover, it is often difficult to implement the conventional techniques in a manner that may be applied in real-time or pseudo real-time to aid the clinician.

OBJECT AND SUMMARY OF THE INVENTION

[0006] The object of the present invention is to solve the above-mentioned problems.

[0007] This is achieved by a method for classifying a cardiovascular sound recorded from a living subject. The method

comprises the step of extracting at least two signal parameters from said cardiovascular sound, said at least two signal parameters characterize at least two different properties of at least a part of said cardiovascular sound. The method further comprises the step of classifying said cardiovascular sound using said at least two signal parameters in a multivariate classification method.

[0008] Hereby a simple method for classifying cardiovascular sounds is achieved and the method is furthermore very robust since different properties of the cardiovascular sound is taken into account and used in a multivariate classification method. The cardiovascular sound related to turbulence consists of at least two components: a broad band component caused by turbulent blood flow colliding with the arterial wall and a narrow banded component related to the resonance frequency of the artery wall, therefore different variables describing different properties are needed in order to perform a robust classification. The different properties describe different characteristics of the cardiovascular sound and would therefore be uncorrelated and therefore provide different information of the cardiovascular sound. Different properties could for instance be the time duration of the diastolic segment of cardiovascular sound, the time duration of the systolic cardiovascular sound, the most dominant frequency component of the sound, the bandwidth of different frequency components, the energy in two frequency bands, the mobility of part of the signal, the complexity of the signal, the power ratio between different parts of the signal, e.g. two different segments or two different frequency bands, morphological characteristics such as correlation ratios between different segments or amplitude change over time. The method could easily be implemented in any kind of data processor unit and therefore be e.g. integrated in a software program which clinicians and doctors could use in order to classify the cardiovascular sound. Furthermore, the method could be integrated in a digital stethoscope and the stethoscope could therefore be used in order to classify a patient's cardiovascular sound. Since doctors and other clinicians are familiar with a stethoscope, they could easily be taught to use the stethoscope to classify the cardiovascular sound. The result is that the classification could assist the doctor or other clinicians to diagnose whether or not the patient suffers from CAD.

[0009] In another embodiment of the method, at least one of said at least two signal parameters is a frequency parameter describing a property in the frequency domain of at least a part of said cardiovascular sound. Hereby the frequency components of the cardiovascular sound could be used as a parameter in the multivariable classification method. Frequency parameters are very good parameters for classifying whether or not murmurs due to stenosis are present in a cardiovascular sound because the stenosis would change the frequency components of the cardiovascular sound.

[0010] In another embodiment of the method, at least one of said at least two signal parameters describes a property in the time domain of at least a part of said cardiovascular sound. Hereby time properties of the cardiovascular sound could be used as a parameter in the multivariable classification method. Time properties like the mobility or number of turning points are good indicators, whether or not murmurs due to stenosis are present in cardiovascular sound. Furthermore, by using both time and frequency parameters a very robust classification of the cardiovascular is achieved since time and frequency properties are often uncorrelated.

[0011] In another embodiment of the method, at least one of said frequency is parameters is a frequency level parameter describing a frequency level property of at least a part of said cardiovascular sound. Hereby it is achieved that a frequency level property of the cardiovascular sound is used in the multivariable classification method. The murmurs would typically change the frequency level of the cardiovascular sound, and by using parameters describing the frequency level of the sound a robust classification of the cardiovascular sound could be achieved.

[0012] In another embodiment of the method, at least one of said at least two signal parameters is a frequency bandwidth parameter describing a frequency bandwidth property of at least a part of said cardiovascular sound. Hereby the bandwidth of, for instance, dominating frequency components could be used in the multivariable classification method. The advantage of using a frequency bandwidth property of the cardiovascular sound is that murmurs often has a limited frequency bandwidth, and the frequency bandwidth parameter would therefore be a good indicator of whether or not murmurs due to stenosis are present in the cardiovascular sound.

[0013] In another embodiment of the method, at least one of said frequency level properties characterizes the most powerful frequency component of at least a part of said cardiovascular sound. This parameter is a very useful parameter as the murmurs due to stenosis typically have a dominating frequency component between 200-800 Hz. And if the most powerful frequency component is inside this interval, it would be a good indication of the presence of murmurs due to stenosis.

[0014] In another embodiment of the method, at least one of said frequency bandwidth properties characterizes the bandwidth of the most powerful frequency component of at least a part of said cardiovascular sound. Hereby the bandwidth of the most powerful frequency component could be used in the multivariable classification method. This bandwidth would most likely depend on whether or not murmurs due to stenosis are present in the cardiovascular sound.

[0015] In another embodiment of the method, at least one of said time parameters is a property characterizing the mobility of at least a part of said cardiovascular sound. The mobility is a good indicator of whether or not murmurs due to stenosis are present in the cardiovascular sound. The mobility describes the variance of the sound, and since murmurs would cause larger variance in the sound the mobility would be a good indicator.

[0016] In another embodiment of the method, the method further comprises the step of dividing said cardiovascular sound into at least one sub-segment and at least one of said signal parameters is extracted from said at least one sub-segment. Hereby it is achieved that the cardiovascular sound could be divided into sub-segments, e.g. into a systolic part and a diastolic part. Thereby relevant sub-segments could be used to extract the above-described different parameters.

[0017] In another embodiment of the method, the method further comprises the step of modelling at least a part of said cardiovascular sound and at least one of said signal parameters is extracted from said model. Hereby time models and frequency models of the cardiovascular sound or sub-segments of the sound could e.g. be used to extract the above-described parameters. The advantage of using models is that the models could enhance the signal properties, e.g. by using an envelope function or an autoregressive model. Further-

more, models would simplify and optimize the calculation process when the method is implemented in a data processor.

[0018] In another embodiment of the method, the multivariate classification method is a discriminant function. Hereby a simple and fast implementation of the classification method is achieved. Furthermore, any number of parameters could be used in the discriminant function, and the different parameters could also be weighted differently depending on the parameters' significance. The discriminant function could also be trained using cardiovascular test sounds is recorded from patients suffering from stenosis and healthy patients. Thereby the weights of the different parameters could be optimized to experimental data.

[0019] The invention further relates to a system for classifying a cardiovascular sound recorded from a living subject, said system comprises processing means for extracting at least two signal parameters from said cardiovascular sound, said at least two signal parameters characterizes at least two different properties of at least a part of said cardiovascular sound; processing means for classifying said cardiovascular sound using said at least two signal parameters using a multivariate classification method. Hereby a system for classifying a cardiovascular sound can be constructed and hereby the same advantages as described above are achieved.

[0020] In a further embodiment of the system, said processing means for extracting at least two signal parameters from said cardiovascular sound is adapted to extract at least one frequency parameter describing a property in the frequency domain of at least a part of said cardiovascular sound. Hereby the same advantages as described above are achieved.

[0021] In a further embodiment of the system, said processing means for extracting at least two signal parameters from said cardiovascular sound is adapted to extract at least one time parameter describing a property in the time domain of at least a part of said cardiovascular sound. Hereby the same advantages as described above are achieved.

[0022] In a further embodiment of the system, said processing means adapted to extract at least one of said frequency parameters are further adapted to extract at least one frequency level parameter describing a frequency level property of at least a part of said cardiovascular sound. Hereby the same advantages as described above are achieved.

[0023] In a further embodiment of the system, said processing means adapted to extract at least one frequency parameter is further adapted to extract at least one frequency bandwidth parameter describing a frequency bandwidth property of at least a part of said cardiovascular sound. Hereby the same advantages as described above are achieved.

[0024] In a further embodiment of the system, said processing means adapted to extract at least one frequency level property is further adapted to extract the most powerful frequency component of at least a part of said cardiovascular sound. Hereby the same advantages as described above are achieved.

[0025] In a further embodiment of the system, said processing means adapted to extract at least one of said frequency bandwidth properties are further adapted to extract the bandwidth of the most powerful frequency component of at least a part of said cardiovascular sound. Hereby the same advantages as described above are achieved.

[0026] In a further embodiment of the system, said processing means for extracting at least one time parameter are

further adapted to extract the mobility of at least a part of said cardiovascular sound. Hereby the same advantages as described above are achieved.

[0027] In a further embodiment of the system, said system further comprises processing means for dividing said cardiovascular sound into at least one sub-segment and at least one of said signal parameters is extracted from said at least one sub-segment. Hereby the same advantages as described above are achieved.

[0028] In a further embodiment of the system, said system further comprises processing means for modelling at least a part of said cardiovascular sound and in that said processing means for extracting at least two signal parameters from said cardiovascular sound are further adapted to extract at least one of said parameters from said model. Hereby the same advantages as described above are achieved.

[0029] In a further embodiment of the system, said multi-variate classification method used by said processing means for classification of said cardiovascular sound is a discriminant function. Hereby the same advantages as described above are achieved.

[0030] The invention further relates to a computer-readable medium having stored therein instructions for causing a processing unit to execute a method as described above. Hereby the same advantages as described above are achieved.

[0031] The invention further relates to a stethoscope comprising recording means adapted to record a cardiovascular sound from a living subject, storing means adapted to store said recorded cardiovascular sound, a computer-readable medium and a processing unit, said computer-readable medium having stored therein instructions for causing said processing unit to execute a method according to claims 1-12 and thereby classify said recorded cardiovascular sound. Hereby the method according to the present invention can be implemented in a stethoscope and the above-described advantages are achieved.

[0032] The invention further relates to a server device connected to a communication network comprising receiving means adapted to receive a cardiovascular sound recorded from a living subject through said communication network, storing means adapted to store said received cardiovascular sound, a computer-readable medium and a processing unit, said computer-readable medium having stored therein instructions for causing said processing unit to execute a method as described above and thereby classify said received cardiovascular sound. Hereby the method according to the present invention can be implemented in a server connected to a communication network. The server could then perform the above-described method and the above-described advantages are achieved.

[0033] In another embodiment of the server, said receiving means are further adapted to receive said cardiovascular sound from a client connected to said communication network. Hereby a clinician/doctor could send a cardiovascular sound to the server using a client device such as a laptop. The server could thereafter classify the received cardiovascular sound. The above-described advantages are hereby achieved.

[0034] In another embodiment of the server, the server device further comprises means for sending said classification of said cardiovascular sound to at least one client unit connected to said communication network. Hereby the result of the classification can be sent back to a client, and the

clinician/doctor can therefore receive the result of the classification. The above-described advantages are hereby achieved.

BRIEF DESCRIPTION OF THE DRAWINGS

- [0035]** FIG. 1 illustrates a graph of a typical heart sound,
[0036] FIG. 2 illustrates a fluid dynamic model of an arterial stenosis,
[0037] FIG. 3 illustrates an overview in form of a flow diagram of the method according to the present invention,
[0038] FIG. 4 illustrates an embodiment of the system according to the present invention,
[0039] FIG. 5 illustrates another embodiment of the method according to the present invention,
[0040] FIG. 6 illustrates a flow diagram of the segmentation method,
[0041] FIG. 7 illustrates for a heart sound the relationship between the envelope autocorrelation of a cardiac cycle and the cardiac cycle,
[0042] FIG. 8 illustrates the implementation of a Bayesian network used to calculate the probability of a sound being an S1, S2 and noise sound.

DESCRIPTION OF EMBODIMENTS

[0043] FIG. 1 illustrates a graph of a typical heart sound recorded by a stethoscope and shows the amplitude (A) of the sound pressure at the y-axis and time (t) at the x-axis. The heart sounds reflect events in the cardiac cycle: the deceleration of blood, turbulence of the blood flow and the closing of valves. The closing of the valves is typically represented by two different heart sounds, the first (S1) and the second (S2) heart sound. The first and second heart sounds are illustrated in the figure, and S1 marks the beginning of systole which is the part of the cardiac cycle in which the heart muscle contracts, forcing the blood into the main blood vessels, and the end of the diastole which is the part of the heart cycle during which the heart muscle relaxes and expands. During diastole, blood fills the heart chambers. The duration of systolic segments is nearly constant around 300 ms for healthy subjects. Given a pulse of 60 beats per minute the duration of a cardiac cycle will be one second on average, and the duration of the diastole will be 700 ms. However, the diastolic durations are not constant, but will vary depending on the subject's pulse. In addition, smaller variations of the diastolic duration are introduced due to neural regulation and the effects of respiration.

[0044] FIG. 2 illustrates a fluid dynamic model of an arterial stenosis and shows an artery (201) with a stenotic lesion (202). The arrows (203) indicate the blood flow through the artery. Vortices (204) will occur when high velocity blood exits a stenotic lesion (202). These vortices collide with the arterial wall (205) and are transformed into pressure vibrations that cause the arteries to vibrate at their resonance frequencies. The result is that soundwaves in the form of murmurs (206) with a frequency corresponding to the arterial wall's resonance frequencies are created and emitted from the arterial wall. Resonance frequencies in the arterial segment are increased if a stenosis is present and their frequencies depend on the diameter of the stenotic segment compared to the diameter of the artery. As the severity of a stenosis increases, so does the resonance frequency. The resonance frequency of a partial occluded stenotic artery is most likely between 200 Hz to 1100 Hz. The intensity of the vortice

fluctuations depends on the blood flow so that murmurs from the left coronary arteries are most intense during diastole, when the blood flow through these arteries is highest. Murmurs from the right coronary arteries are most intense during diastole if there is a stenosis in branches of the right coronary artery supplying the right-sided cavities, whereas the murmur more likely will be systolic from those branches of the right coronary artery giving arterial blood to the left ventricle. The intensity of murmurs not only depends on the blood flow, but also on the frequency content of a murmur. High murmur frequencies are more suppressed by the chest wall compared to low frequencies. The murmurs caused by the arterial vibrations would affect the graph of a heart sound recorded by e.g. a stethoscope.

[0045] FIG. 3 illustrates an overview in form of a flow diagram of the method according to the present invention. The method could for instance be implemented as a software program running on a computer or on a microcontroller implemented in a stethoscope. In short, the method starts with an initialization (301), receiving a test signal (302), dividing the test signal into relevant segments (306), filtering the relevant segments (307); calculating/developing a model of the signal (308) in relevant segments; extracting different parameters from the signal and the model (309), performing an analysis of the signal (310) using the extracted parameters and classifying the relevant segments into two groups: one indicating that the signal contains murmurs due to stenosis (311), and one indicating that the signal does not contain murmurs due to stenosis (312).

[0046] After the method has been initialized (301) the method receives the test signal (302) as a data file (303). The test signal would be the heart sound from a person (304) recorded and digitalized into a data file, e.g. by a digital stethoscope (305). The test signal would be similar to the heart sound illustrated in FIG. 1, however, the duration of the test signal would typically be 5-15 times longer than the signal shown in FIG. 1. Once the test signal has been received (302), segmentation (306) is performed in order to detect and divide the test signal into segments. The segmentation process would typically detect the heart sounds S1 and S2 and thereafter divide the test signal into systolic and diastolic parts. Hereafter the test signal is filtered (307), and the filtration process includes an autoregressive filter that reduces white noise in the signal and a band pass filter that only lets frequencies between 450-1100 Hz pass. The test signal would thereafter contain the frequencies caused by the vibrations of the arterial wall when stenosis is present in the artery. The autoregressive filter could be implemented as a Kalman filter that is a powerful estimator of past, present and future states and it can do so even when the precise nature of the modelled system is unknown. This is a desirable feature in the present application when reducing the effects of noise since the exact composition of a murmur is unknown. A first order Kalman filter can reduce the effects of white noise and smooth the noisy heart sound recordings for further processing. The band pass filter could be implemented as a wavelet filter. In another embodiment the Kalman filter is omitted in order to simplify the implementation of the method in e.g. a microprocessor and further to reduce the number of calculations performed by the microprocessor.

[0047] When the signal has been filtered (307), relevant segments are selected for further analysis. In one embodiment a part of the diastolic segment is selected for further analysis as the murmur due to stenosis is most likely to be audible in the diastolic segment.

[0048] A mathematical model of the signal in the selected segment is hereafter calculated/developed (308) using the sampled heart sound in the data file. The model is used to extract parameters that characterize the sound in the segment and could be used to categorize whether or not the murmurs due to stenosis exist in the sound segment. In the present embodiment an autoregressive all-pole parametric estimation (AR-model) is used to model the signal. In the AR-model the sampled sound signal, y , from the data file is modelled as a linear combination of M past values of the signal and the present input, u , driving the sound generating process. The model can be described by the following equation:

$$y(n) = - \sum_{p=1}^M a_p y(n-p) + u(n) \quad [3.1]$$

[0049] where M represents the model order, A_p the AR coefficients and n the sample number. The AR coefficients are determined through an autocorrelation and by minimizing the error associated with the model.

[0050] The AR model in this embodiment is used to extract frequency parameters describing the heart sound. A second order model $M=2$ is preferred because it makes a better separation between the frequency parameters extracted from a heart sound with murmurs present and the frequency parameters extracted from a heart sound with murmurs present.

[0051] Thereafter different parameters are extracted (309) from the sampled signal and the AR model using signal processing techniques. Some parameters could be extracted from the selected segments. Each parameter characterize the heart sound in the selected segments and could therefore be used to categorize the heart sound, e.g. whether or not murmurs due to stenosis are present in the heart sound. The parameters can in this embodiment be the number of turnings points per signal length, TP; the mobility of the signal, MB; pole magnitude, PM; normalized AR-peak frequency, NF; and AR spectral ratio, SR.

[0052] The number of turning points TP is extracted from the sampled signal in the time domain, and it is found by calculating the number of turns the signal performs in the time domain per unit time. This could be done by determining the amount of local maxima in a time period. Thus:

$$TP = \frac{\text{number of turns}}{\text{signal length}} \quad [3.2]$$

[0053] The mobility MB is extracted from the sampled signal in the time domain and found by calculating the variance, σ_y , of the signal in the time domain and the variance of the signal's first derivative, $\sigma_{y'}$. The mobility is hereafter found by:

$$MB = \frac{\sqrt{\sigma_{y'}^2}}{\sqrt{\sigma_y^2}} = \frac{\sigma_{y'}}{\sigma_y} \tag{3.3}$$

[0054] The pole magnitude PM is found by transforming the AR-model into the z-domain and calculating the magnitude of the poles in z-domain described by the AR-spectrum.

[0055] The normalized AR peak frequency NF is based on the assumption that murmurs due to stenosis are more likely to be found in the diastolic segment than in the systolic segment. The NF is found by calculating the angle of the poles in the AR-spectrum in the z-plane and transforming this into a frequency of both a diastolic segment and a systolic segment. If the absolute difference between the two is less than 25 Hz, which is typical in cases where no murmurs due to stenosis are present, then 25 Hz is subtracted from the diastolic peak frequency. If the average diastolic frequency is more than 50 Hz greater than the average systolic peak frequency, which is typical when murmurs due to stenosis are present, then 25 Hz is added to the average peak diastolic frequency.

[0056] The AR spectral ratio SR is found by calculating the ratio of the energy in the frequency rang 200-500 Hz to the energy in the frequency range 500-1000 Hz of a diastolic segment.

[0057] The extracted parameters are thereafter used in a multiparametric discriminant function in order to classify whether or not the sound segment contains murmurs due to stenosis (310). In this embodiment a linear discriminant function is used to classify the sound segments. The linear discriminant function combines weighted features into a discriminant score $g(x)$ and could be described by:

$$g(x) = w_1x_1 + w_2x_2 + w_3x_3 + \dots + w_kx_k + w_0 = w^T x + w_0 \tag{3.4}$$

where x is the feature vector consisting of the extracted parameters, k represents the number of features, i represents the classes and w is a weight vector that holds the discriminant coefficients. In the case where only two classes must be separated, a single discriminant function is used. A two class classifier is called a dichotomizer. A dichotomizer normally classifies the feature vectors with the decisions boarder $g(x) = 0$ (due to the constant w_0). If the discriminant score $g(x)$ is greater than zero the segment is assigned to class 1, otherwise it is assigned to class 2. Since g is a linear function $g(x) = 0$ it defines a hyperplane decision surface, dividing the multi dimensional space into two half sub spaces. The discriminant score $g(x)$ is the algebraic distance to the hyper-plane. The discriminant function needs to be trained in order to find the weights values, w , and make a safe and robust classification of the sound segments. The discriminant training procedure needs to be performed before using the system, and the purpose of the procedure is to find the optimal weights values of w so that the hyper plane separates the feature vectors optimally. The training procedure is in one embodiment carried out by using 18 test sounds recorded from 18 test persons where nine test persons have coronary stenosis and nine test persons do not have coronary stenosis. The discriminant training procedure is performed by using the statistical software program SPSS v.12.0 for windows (SPSS inc., Chicago Ill., USA). The above-mentioned parameters are extracted from the 18 training sounds and used as statistical inputs to the software program. The resulting discriminant could be:

$$g(x) = 164.709MB - 0.061NF - 78.027PM + 27.188SR + 91.878TP + 33.712 \tag{3.5}$$

where MB is the mobility of the signal, NF the AR-peak frequency, PM the pole magnitude, SR the AR spectral ratio and TP the number of turning point.

[0058] If the result of the discriminant function is larger than zero ($g(x) > 0$) then the sound segment does not contain murmurs due to stenosis (312). On the other hand, if the discriminant function is smaller than zero ($g(x) < 0$) then the sound segment contains murmurs due to stenosis (311).

[0059] The discriminant function could by a person skilled in the art easily be adjusted to include additional or fewer parameters in order to develop a proper discriminant function that can be used to classify the heart sound. Further parameters could for instance be:

[0060] The Complexity, CP, of the sampled signal in the time domain. This parameter is based on the ratio of the mobility of the first derivative of the signal to the mobility of the signal itself where y'' is the second derivative of the filtered heart sound signal. The complexity measure is relatively sensitive to noisy signals since it is based on the second derivative.

$$CP = \frac{MB_{y'}}{MB_y} = \frac{\sigma_{y''} / \sigma_{y'}}{\sigma_{y'} / \sigma_y} \tag{3.6}$$

[0061] Further, the AR-peak frequency (PF) could be extracted and used in the discriminant function. The AR-peak frequency could be found by calculating the angle of the AR poles in the z-plane.

[0062] The parameters used in the discriminant function could be extracted from different segments of the heart sound, e.g. a number of different diastolic segments where a number of parameters is extracted from each diastolic segment. Thereafter an average value of each parameter could be calculated and used as input in the discriminant function.

[0063] FIG. 4a illustrates an embodiment of the system according to the present invention where a server (401) is programmed to execute the method described in FIG. 3. Furthermore, the server is connected to a network (402), e.g. the Internet and adapted to on request to receive and analyze heart sound. Clinicians or other medical professionals would record the heart sound from a patient by a digital stethoscope (305) and thereafter transmit the digitalized heart sound to a personal computer (403). The clinician can hereafter send a request to the server in order to have the heart sound analyzed. Once the server has analyzed the heart sound the result is automatically sent back to the clinician. FIG. 4b illustrates a flow diagram of the process and the communication between the personal computer (403) and the server. The left hand side represents the client side (410) and the right hand side represents the server side (411). First the client sends a heart sound in digital form to the server (412). Thereafter the server performs the method illustrated in FIG. 3 and sends (413) the result of the analysis back to the client where it is displayed (414) to the clinician. The clinician could hereafter evaluate the result in order to choose the right treatment of the patient.

[0064] The system according to the present invention could also be implemented as an all in one digital stethoscope. The stethoscope would therefore automatically perform the analysis described in FIG. 3 when a heart sound has been recorded. This means that the method described in FIG. 3

needs to be implemented in stethoscopes' processing means, and the result of the analysis could e.g. be displayed on a small LCD integrated in the stethoscope. An advantage of this embodiment is that most clinicians are familiar with a digital stethoscope and could therefore easily learn to use the stethoscope to diagnose whether or not the patient has a coronary stenosis.

[0065] FIG. 5 illustrates another embodiment of the method described in FIG. 3. When the signal has been filtered (307), relevant segments are selected for further analysis. In one embodiment a part of the diastolic segment is selected for further analysis as the murmur due to stenosis is most likely to be audible in the diastolic segment. In this embodiment the diastolic segment comprising respiration sounds is discarded (501). This is done by calculating the energy level of the diastolic segment in the frequency band 200-440 Hz and comparing this energy level with the median energy level of the entire diastolic segment. The diastolic segment would be discarded if the energy level of the 200-440 Hz frequency band is a factor 1.1 larger than the energy level in the entire diastolic segment.

[0066] The remaining diastolic segments are hereafter divided into sub-segments (502) with a duration of 37.5 ms or 300 samples. This is done because the blood flow in the coronary artery is not constant during a diastole, and the murmurs due to stenosis would therefore not be constant.

[0067] The variance of the signal in all sub-segments is then calculated and the sub-segments with a variance larger than 1.3 of the median variance of all sub-segments are then discarded (503). Hereby sub-segments comprising high noise spikes are removed.

[0068] Thereafter (504) none stationary sub-segments are removed. This is done by dividing the sub-segment into sub-sub-segments with a duration of 3.75 ms or 30 samples and then calculate the variance of each sub-sub-segments. Thereby an outline of the variance throughout the sub-segment is constructed. The variance of the outline is then calculated and the sub-segment is removed if the variance of the outline is larger than 1.

[0069] At this point a number of sub-segments have been discarded in order to remove noisy and none stationary sub-segments. This would typically result in 30-50 sub-segments from a cardiovascular recording of approximately 10 seconds.

[0070] The remaining sub-segments are thereafter used in step (308) and (309) as described in FIG. 3 in order to extract parameters describing different properties of the cardiovascular signal. Thereafter the median of each parameter is calculated using the values of the parameter from each sub-segment (505). The median of each parameter is thereafter used in the multiparametric discriminant function as described in FIG. 3. In this embodiment the following parameters are used: the mobility, the power-ratio and the pole-amplitude of a 3 pole in an AR model of order 6.

[0071] FIG. 6 illustrates a flow diagram of the segmentation method (306) according to the present invention used to automatic divide a heart sound (601) into sub-segments. The heart sound (601) has been recorded by a stethoscope and the signal has been digitized in order to digitally process the signal. The graph shows the amplitude (A) of the sound intensity as a function of time (t). The heart sounds reflect events in the cardiac cycle; the deceleration of blood, turbulence of the blood flow and the closing of valves. The closing of the valves is typically represented by two different heart sounds, the first

(S1) and the second (S2) heart sound. The first and second heart sounds are illustrated in the figure, and (S1) marks the beginning of systole, which is the part of cardiac cycle in which the heart muscle contracts, forcing the blood into the main blood vessels, and the end of the diastole which is the part of the heart cycle during which the heart muscle relaxes and expands. During diastole, blood fills the heart chambers.

[0072] The purpose of the segmentation method is to classify the recorded heart sound into systolic, diastolic and noise segments. The illustrated method includes steps of noise reduction (602) followed by envelope creation (603). The noise reduction could be implemented as a high-pass filter followed by removal of high amplitude friction noise spikes due to external noise like movement of the stethoscope during recording and thereafter a low pass filter. The purpose of the envelope creation is to enhance the trend of the signal. The envelope is in this embodiment created by calculating the Shannon energy of the signal:

$$se(n)=x(n)^2 \cdot \log x(n)^2$$

where x is the signal and se is the Shannon energy. The high amplitude components in the signal are weighted higher than low amplitude components when calculating the Shannon energy. The envelope (613) of the heart sound (601) calculated by using the Shannon energy is shown in figure (613), and it can be seen that the heart sounds S1 and S2 are enhanced.

[0073] In order to classify the detected sounds into systolic segments, diastolic segments and noise components based on interval durations on either side of the heart sounds S1 and S2, it is necessary to know how long the intervals between S1's and S2's are. Therefore, the durations of the heart cycles (systolic and diastolic intervals) are extracted from an autocorrelation of the envelope (604). This process is described in detail in FIG. 7.

[0074] Candidates S1's and S2's are then detected (605) using the time intervals extracted above and a threshold (614) on the envelope (613). To reduce the number of detected noise spikes, a minimum requirement is applied to the candidate segments, which effectively removes some of the erroneously detected noise spikes. In some recordings there is a big difference between the intensity of S1 and S2 sounds. This causes a problem since some of the low intensity sounds may be missed by the threshold. As a result the segmentation method performs a test for missing S1 and S2 sounds (606). If it can be determined that some segments are missing, the threshold procedure is rerun (607) using lower local thresholds.

[0075] Once the signal has been divided into segments as described above interval parameters and frequency parameters for each segment are then extracted (608). The parameters aid in the classification of the sounds into systolic segments and diastolic segments.

[0076] The interval parameters are four Boolean parameters extracted for each sound by comparing the time duration to the previous sound and to the next sound with the time intervals extracted using the autocorrelation. The parameters are:

[0077] AfterDia: Is true if the sound is succeeded by a second sound after a period corresponding to the duration of a diastole,

[0078] AfterSys: Is true if the sound is succeeded by a second sound after a period corresponding to the duration of a systole,

[0079] BeforeDia: Is true if the sound follows a second sound after a period corresponding to the duration of a diastole,

[0080] BeforeSys: Is true if the sound follows a second sound after a period corresponding to the duration of a systole.

[0081] The frequency parameter divides the sounds into low frequency and high frequency sounds by calculating the median frequency of the sound. This is useful information as the first heart sound is expected to be a low frequency sound and the second heart sound is expected to be a high frequency sound.

[0082] The parameters are parsed into a Bayesian network where the probability of a segment being a S1, S2 and noise sound is computed (609). The figure illustrates a bar chart (615) of the probability calculated for each sound in the heart signal (601). Each sound would typically have one dominating probability indicating the type (S1, S2 or noise) of the sound. Thereby all sounds are classified into S1, S2 and noise sounds. However, the probability of the three types would in some cases be more or less equal and in such cases it is not possible to classify the sound into a S1, S2 or noise sound using the Bayesian network.

[0083] The probabilities are used in the last step (610) to divide and verify the heart signal into systole and diastole segments. This is done by using the position of the identified S1 and S2 sounds to mark the beginning of a systolic and diastolic sound segment respectively

[0084] The final result of the method (611) is the beginnings and ends of all identified systoles and diastoles. Therefore a “train” (616) of alternating systoles (617) and diastoles (618) can be created. Once the systoles and diastoles have been identified they can be used in further data handling, e.g. to extract further parameters from these segments and thereafter use the parameters to classify the medical condition of the recorded heart sound.

[0085] FIG. 7 illustrates the relationship between the envelope autocorrelation and the cardiac cycle, and how the intervals between heart sounds S1 and S2 can be found from the autocorrelation.

[0086] FIG. 7a illustrates the envelope autocorrelation with the normalized autocorrelation at the y-axis (NA) and the displacement (m) of the shifted envelope at the x-axis.

[0087] FIG. 7b illustrates the displacement (m1) when the shifted envelope (701) is displaced by the duration of the systole corresponding to the unshifted envelope (702). The y-axis shows the amplitude (A) of the envelope and the x-axis the time (t). The S1’s in the displaced envelope are multiplied by the S2’s in the unshifted envelopes resulting in the first peak (703) seen in the autocorrelation.

[0088] FIG. 7c illustrates the displacement (m2) when the shifted envelope (701) is displaced by the duration of the diastole corresponding to the unshifted envelope (702). The displaced S2’s are multiplied by the S1’s in the unshifted envelope resulting in the second peak (704) seen in the autocorrelation.

[0089] FIG. 6b illustrates the displacement (m3) when the shifted envelope (701) is displaced by the duration of the cardiac cycle corresponding to the unshifted envelope (702). The S1’s in the displaced envelope are multiplied by the S1’s in the unshifted envelope, and the S2’s in the displaced envelope are multiplied by the S2’s in the unshifted envelope. When this occurs the dominating peak (705) in the autocorrelation is produced.

[0090] The interval between the heart sounds could therefore be found by measuring the distance between the peaks in the autocorrelation as described above.

[0091] FIG. 8 illustrates the implementation of the Bayesian network used to calculate the probability of a sound of being an S1, S2 and noise sound in step (809). The basic concept in the Bayesian network is the conditional probability and the posterior probability. The conditional probability describes the probability of the event a given the event b.

$$P(a|b)=x_c \tag{8.1}$$

[0092] If the above equation describes the initial conditional probability, the posterior probability would be:

$$P(b|a)=x_p \tag{8.2}$$

[0093] According to Bayes’ rule the relation between the posterior probability and the conditional probability is:

$$P(b|a) = \frac{P(a|b)P(b)}{P(a)} \tag{8.3}$$

where P(a) is the prior probability for the event a, and P(b) is the prior probability for the event b. Equation [8.3] only describes the relation between one parent and one child, but since the event a can be the combination of several events {a₁, a₂, . . . , a_n} the equation can be expanded to:

$$P(b|a_1, a_2, . . . , a_n) = \frac{P(a_1, a_2, . . . , a_n|b)P(b)}{P(a_1, a_2, . . . , a_n)} \tag{8.4}$$

[0094] Since the goal is to find the probability for the different states of b when a₁ and a₂ are known, P(a₁, a₂, . . . , a_n) is just a normalizing constant k and [7.4] can be simplified to:

$$P(b|a_1, a_2, . . . , a_n) = k \cdot P(a_1, a_2, . . . , a_n|b)P(b) \tag{8.5}$$

[0095] If child events (a₁, a₂ . . . a_n) are conditionally independent, equation [8.5] can be generalized to:

$$P(b|a_1, a_2, . . . , a_n) = k \cdot P(b) \prod_{i=1}^N P(i|b) \tag{8.6}$$

where N is the number of known events a. Equation [8.6] is useful in determining the probability of the event b if the states of all a events are known and if all a events are conditionally independent. A Bayesian network based on equation [8.6] is called a naive Bayesian network because it requires conditional independency of the children.

[0096] The task for the Bayesian network is to evaluate the type of each detected sound above the detection threshold. For each of these sounds, the posterior probability of being an S1 sound, an S2 sound or a noise component is calculated and the Bayesian network is constructed using one parent and five children. The parent is a sound above the envelope threshold (801), and the children are the five parameters described above: Frequency (802), AfterSys (803), AfterDia (804), BeforeSys (805) and BeforeDia (806). When determining the posterior probability for the type of a particular sound, the prior probability for the different states of a sound type P(S) and the conditional probabilities must be known, i.e. the conditional probabilities that “AfterSys” is in a given state

when S is a given type, P(AfterSys|S). This posterior probability requires definition of P(S), P(AfterSys|S), P(AfterDial|S), P(BeforeSys|S), P(BeforeDial|S) and P(Frequency|S) before the equation [8.6] can be used to calculate the posterior probability of a sound being a particular type of sound.

[0097] The prior probability that a sound is an S1, S2 or a noise component changes between recordings. In the optimal recording, where no noise components are detected, the prior probability for noise is zero, P(S=*Noise*)=0. If this is the case and an equal number of S1's and S2's are detected, the prior probability that the detected sound is an S1 is 50%, and similar for S2. Therefore, P(S=*S1*)=P(S=*S2*)=0.5 if P(S=*noise*)=0. However, this optimal condition cannot be assumed for real signals, and noise sounds would be detected. This will increase the prior probability that a given sound is noise.

[0098] The exact probability of a detected sound being noise, P(S=*noise*) can be defined if the number of detected noise sounds, N_{noise} and the total number of detected sounds, N_{sounds} are known. For instance, if it is known that four noise sounds are detected, $N_{noise}=4$, and the total number of detected sounds is 20, the probability that the sound being examined is a noise sound is P(S=*Noise*)=4/20. However, in most signals N_{noise} is unknown and an estimate of N_{noise} is therefore necessary, and this estimate can be based on already available information since the duration of a heart cycle is known from the envelope autocorrelation (804). The expected number of cardiac cycles in one recording can therefore be calculated by dividing the length of the recording with the length of the cardiac cycles. The number of S1's and S2's in a recording is therefore twice the number of cardiac cycles in a recording. The prior probability of the sound type would therefore be:

$$P(S_{=noise}) = \frac{N_{noise}}{N_{sound}} \quad [8.7]$$

and the prior probability that the detected sound is an S1 or S2:

$$P(S_{=s1}) = P(S_{=s2}) = \frac{1 - P(S_{=noise})}{2} \quad [8.8]$$

[0099] The conditional probability that an S1 is followed by an S2 sound after an interval corresponding to the duration of a systole, P(AfterSys|S=*S1*), depends on several factors. The S1 sounds will normally be followed by S2 sounds after an interval of duration equal to the systole. Deviations from this can also occur, e.g. when S1 is the last sound in the recording, or if S2 is missing because it is not detected by the threshold. It may also occur that a weak (below threshold) S2 is detected because noise occurs in the tolerance window associated with those sounds. The probability that "AfterSys" is false if the sound is an S1 sound may thus be calculated as

$$\frac{P(\text{AfterSys}_{=false}|S_{=S1})}{\text{NoiseInWin}} = P(\text{EndSound} \cup \text{Singlesound}), \quad [8.9]$$

where "EndSound" is an event describing that the sound is the last sound in the recording. "SingleSound" describes that S1 is not followed by S2 as the next S2 sound is not detected due to sub-threshold amplitude. "NoiseInWin" describes noise occurrence in the window, where the S2 sound was expected.

The conditional probability that "AfterSys" is true given that the examined sound is an S1 sound is given by:

$$P(\text{AfterSys}_{=true}|S_{=S1}) = 1 - P(\text{AfterSys}_{=false}|S_{=S1}) \quad [8.10]$$

[0100] If the examined sound is an S2 sound it is not likely that any sound occurs after an interval corresponding to the systolic duration since the next S1 sound will occur after the duration of the diastole. An exception is if a noise sound occurs in the window P(NoiseInWin) or if the systole and diastole durations are equal. If the duration of the diastole is equal to the duration of the systole, the S1 sound which follows the S2 sound after the duration of a diastole occurs in both the systole tolerance window and in the diastole tolerance window. This will happen if the heart rate of the subject is high. The probability that a sound occurs in both tolerance windows (overlap) is equal to the degree of the overlap between the systole and diastole tolerance window. This probability is termed P(Overlap). Therefore, the conditional probability that a sound occurs in the window after systole duration if the examined sound is an S2 sound is:

$$P(\text{AfterSys}_{=true}|S_{=S2}) = P(\text{Overlap} \cup \text{NoiseInWin}) \quad [8.11]$$

[0101] The conditional probability that a sound does not occur after a systole duration, if the examined sound is an S2, is the opposite of the conditional probability that it does occur:

$$P(\text{AfterSys}_{=false}|S_{=S2}) = 1 - P(\text{AfterSys}_{=true}|S_{=S2}) \quad [8.12]$$

[0102] The conditional probability that a detected noise sound is followed by another sound after the systole duration is based on the probability that a sound of any kind is present in a segment with the length of the used tolerance window. This can be estimated from the ratio of the tolerance window length multiplied by the number of detected sounds minus one to recording length.

$$P(\text{SoundInWin}|S_{=S2}) = 1 - P(\text{AfterSys}_{=true}|S_{=S2}) \quad [8.12]$$

[0103] The conditional probability that a detected noise sound is followed by another sound after the systole duration, P(AfterSys|S=*noise*), is based on the probability that a sound of any kind is present in a segment with the length of the used tolerance window. This can be estimated from the ratio of the tolerance window length multiplied by the number of detected sounds minus one to recording length. The conditional probability that a noise sound is followed by another sound after a systole duration is therefore:

$$P(\text{AfterSys}_{=true} | S_{=noise}) = P(\text{SoundInWin}) \quad [8.13]$$

$$= \frac{(N_{sound} - 1) \cdot 2 \cdot \text{Sys}_{tot}}{\text{RecLength}}$$

[0104] where N_{sound} is the number of sounds within the recording, Sys_{tot} is the duration of a systole and RecLength is the length of the recording. The conditional probability that a noise is not followed by another sound after the systole interval is the opposite:

$$P(\text{AfterSys}_{=false}|S_{=noise}) = 1 - P(\text{SoundInWin}) \quad [8.14]$$

[0105] The conditional probabilities for P(AfterDial|S), P(BeforeSys|S) and P(BeforeDial|S) are based on the same assumptions used to define P(AfterSys|S). These conditional probabilities can be found in the tables below:

	False	True
<u>P(AfterSys/S)</u>		
S1	$P((\text{EndSound} \cup \text{SingleSound}), \text{NoiseInWin})$	$1 - P((\text{EndSound} \cup \text{SingleSound}), \text{NoiseInWin})$
S2	$1 - P(\text{Overlap} \cup \text{NoiseInWin})$	$P(\text{Overlap} \cup \text{NoiseInWin})$
Noise	$1 - P(\text{SoundInWin})$	$P(\text{SoundInWin})$
<u>P(AfterDial/S)</u>		
S1	$1 - P(\text{Overlap} \cup \text{NoiseInWin})$	$P(\text{Overlap} \cup \text{NoiseInWin})$
S2	$P((\text{EndSound} \cup \text{SingleSound}), \text{NoiseInWin})$	$1 - P((\text{EndSound} \cup \text{SingleSound}), \text{NoiseInWin})$
Noise	$1 - P(\text{SoundInWin})$	$P(\text{SoundInWin})$
<u>P(AfterSys/S)</u>		
S1	$1 - P(\text{Overlap} \cup \text{NoiseInWin})$	$P(\text{Overlap} \cup \text{NoiseInWin})$
S2	$P((\text{EndSound} \cup \text{SingleSound}), \text{NoiseInWin})$	$1 - P((\text{EndSound} \cup \text{SingleSound}), \text{NoiseInWin})$
Noise	$1 - P(\text{SoundInWin})$	$P(\text{SoundInWin})$
<u>P(AfterDial/S)</u>		
S1	$P((\text{EndSound} \cup \text{SingleSound}), \text{NoiseInWin})$	$1 - P((\text{EndSound} \cup \text{SingleSound}), \text{NoiseInWin})$
S2	$1 - P(\text{Overlap} \cup \text{NoiseInWin})$	$P(\text{Overlap} \cup \text{NoiseInWin})$
Noise	$1 - P(\text{SoundInWin})$	$P(\text{SoundInWin})$

[0106] It has previously been found that the frequency parameter classified 86% of the S1 sounds as low frequent and 80% of the S2 sounds as high frequent. 85% of all noise sounds were classified as high frequent. This information was used as the conditional probabilities between the frequency parameter P(Frequency|S):

P(Frequency S)	Low	High
S1	0.86	0.14
S2	0.20	0.80
Noise	0.15	0.85

[0107] When all conditional probabilities are found, equation [8.6] is used by the Bayesian network to calculate the posterior probabilities for all detected sounds. This way, three probabilities are calculated for each sound that reflect how likely it is that the current sound is a given type.

[0108] It should be noted that the above-mentioned embodiments rather illustrate than limit the invention, and that those skilled in the art will be able to suggest many alternative embodiments without departing from the scope of the appended claims.

1. A method for classifying a cardiovascular sound recorded from a living subject, said method comprises the steps of:

extracting at least two signal parameters from said cardiovascular sound, said at least two signal parameters characterizes at least two different properties of at least a part of said cardiovascular sound,

classifying said cardiovascular sound using said at least two signal parameters in a multivariate classification method.

2. A method according to claim 1 characterized in that at least one of said at least two signal parameters is a frequency parameter describing a property in the frequency domain of at least a part of said cardiovascular sound.

3. A method according to claim 1 characterized in that at least one of said at least two signal parameters describing a property in the time domain of at least a part of said cardiovascular sound.

4. A method according to claim 2 characterized in that at least one of said frequency parameters is a frequency level parameter describing a frequency level property of at least a part of said cardiovascular sound.

5. A method according to claim 2 characterized in that at least one of said at least two signal parameters is a frequency bandwidth parameter describing a frequency bandwidth property of at least a part of said cardiovascular sound.

6. A method according to claim 4 characterized in that at least one of said frequency level properties characterizes the most powerful frequency component of at least a part of said cardiovascular sound.

7. A method according to claim 5 characterized in that at least one of said frequency bandwidth properties characterizes the bandwidth of the most powerful frequency component of at least a part of said cardiovascular sound.

8. A method according to claim 3 characterized in that at least one of said time parameters is a property characterizing the mobility of at least a part of said cardiovascular sound.

9. A method according to claim 1 characterized in that said method further comprises the step of dividing said cardiovascular sound into at least one sub-segment and at least one of said signal parameters is extracted from said at least one sub-segment.

10. A method according to claim 1 characterized in that said method further comprises the step of modelling at least a part of said cardiovascular sound and at least one of said signal parameters is extracted from said model.

11. A method according to claim 1 characterized in that said multivariate classification method is a discriminant function.

12. A system for classifying a cardiovascular sound recorded from a living subject, said system comprises:

processing means for extracting at least two signal parameters from said cardiovascular sound, said at least two signal parameters characterizes at least two different properties of at least a part of said cardiovascular sound, processing means for classifying said cardiovascular sound using said at least two signal parameters using a multivariate classification method.

13. A system according to claim 12 characterized in that said processing means for extracting at least two signal parameters from said cardiovascular sound is adapted to extract at least one frequency parameter describing a property in the frequency domain of at least a part of said cardiovascular sound.

14. A system according to claim 12 characterized in that said processing means for extracting at least two signal parameters from said cardiovascular sound is adapted to extract at least one time parameter describing a property in the time domain of at least a part of said cardiovascular sound.

15. A system according to claim 13 characterized in that said processing means adapted to extract at least one of said frequency parameters are further adapted to extract at least one frequency level parameter describing a frequency level property of at least a part of said cardiovascular sound.

16. A system according to claim 13 characterized in that said processing means adapted to extract at least one frequency parameter are further adapted to extract at least one frequency bandwidth parameter describing a frequency bandwidth property of at least a part of said cardiovascular sound.

17. A system according to claim 13 characterized in that said processing means adapted to extract at least one frequency level property are further adapted to extract the most powerful frequency component of at least a part of said cardiovascular sound.

18. A system according to claim 13 characterized in that said processing means adapted to extract at least one of said frequency bandwidth properties are further adapted to extract the bandwidth of the most powerful frequency component of at least a part of said cardiovascular sound.

19. A system according to claim 14 characterized in that said processing means for extracting at least one time parameters are further adapted to extract the mobility of at least a part of said cardiovascular sound.

20. A system according to claim 12 characterized in that said system further comprises processing means for dividing said cardiovascular sound into at least one sub-segment and at least one of said signal parameters is extracted from said at least one sub-segment.

21. A system according to claim 12 characterized in that said system further comprises processing means for modeling at least a part of said cardiovascular sound and in that said processing means for extracting at least two signal parameters from said cardiovascular sound are further adapted to extract at least one of said parameters from said model.

22. A system according to claim 12 characterized in that said multivariate classification method used by said processing means for classification of said cardiovascular sound is a discriminant function.

23. A computer-readable medium having stored therein instructions for causing a processing unit to execute a method according to claim 1.

24. A stethoscope comprising:

recording means adapted to record a cardiovascular sound from a living subject,

storing means adapted to store said recorded cardiovascular sound,

a computer-readable medium and a processing unit, said computer-readable medium having stored therein instructions for causing said processing unit to execute a method according to claim 1 and thereby classify said recorded cardiovascular sound.

25. A server device connected to a communication network comprising:

receiving means adapted to receive a cardiovascular sound recorded from a living subject through said communication network,

storing means adapted to store said received cardiovascular sound,

a computer-readable medium and a processing unit, said computer-readable medium having stored therein instructions for causing said processing unit to execute a method according to claim 1 and thereby classify said received cardiovascular sound.

26. A server device according to claim 25 characterized in that said receiving means are further adapted to receive said cardiovascular sound from a client connected to said communication network.

27. A server device according to claim 25 characterized in that said server device further comprises means for sending said classification of said cardiovascular sound to at least one client unit connected to said communication network.

* * * * *