

DYNAMIC CIRCULAR BUFFERING: A TECHNIQUE FOR EQUILIBRIUM GATED BLOOD POOL IMAGING

JUAN JOSÉ VAQUERO,*† HELLMER RAHMS,* MICHAEL V. GREEN‡ and FRANCISCO DEL POZO*

* Grupo de Bioingeniería y Telemedicina, E.T.S.I. Telecomunicación, Universidad Politécnica de Madrid, Ciudad Universitaria s/n, 28040 Madrid, Spain; and ‡ Chief, Computer Systems and Imaging Physics, Department of Nuclear Medicine, Clinical Center, National Institutes of Health, Bethesda, MD 20892, U.S.A.

(Received 1 June 1995; revised form 2 October 1995)

Abstract—We have devised a software technique called “dynamic circular buffering” (DCB) with which we create a gated blood pool image sequence of the heart in real time using the best features of LIST and FRAME mode methods of acquisition/processing. The routine is based on the concept of independent “agents” acting on the timing and position data continuously written into the DCB. This approach allows efficient asynchronous operation on PC-type machines and enhanced capability on systems capable of true multiprocessing and multithreading.

Gated blood pool imaging
function

Real time

Computer processing

Ventricular

INTRODUCTION

Gated blood pool imaging of the heart [1] is a commonly used nuclear medicine diagnostic imaging procedure intended to visualize and quantify cardiac function during rest, exercise or other interventions. The information needed to create a gated blood pool image sequence consists of the X , Y coordinates of events occurring within the field of view of a scintillation camera placed to view the heart, a logic signal that indicates the moment of occurrence of each R-wave in the patient's ECG signal (the beginning of each heartbeat), and a train of clock signals that measure the passage of time.

Additional hardware needed for these studies is described in [2], and consists of a QRS detector, a set of buffers for the temporary storage of data, and additional processors for managing the events in real time.

Two methods have been proposed for acquiring and manipulating these data to create an image sequence that spans an average cardiac cycle. The most common method, frame mode acquisition, has the clinical virtue of yielding a complete, ready-for-inspection, image sequence at the end of data collection. This method, however, exhibits data fall-off at the end of the reconstructed image sequence [3], and in its pure implementation does not allow “bad beat” rejection; some systems add a buffer in front of the frame mode sorting in order to solve this problem. List mode acquisition, the second method, does allow image data from cardiac cycles of abnormal length to be rejected and also allows accurate reconstruction of events occurring late in the average cardiac cycle. Nevertheless, this method, as usually implemented, does not yield a ready-to-view image sequence at the end of data collection but requires post-processing to achieve this result. Moreover, both methods are usually carried out using specialised computer systems embedded in costly commercial scintillation camera imaging systems. Lear and Pratt [4] present an alternative technique that, like ours, needs a minimum of hardware added to the computer system, but our implementation exploits the asynchronous mode advantages as will be shown. Nevertheless, the technical requirements for the implementation are minimised if off-the-shelf components are used.

† Author to whom correspondence should be addressed.

Given these considerations, we have devised a processing method called “dynamic circular buffering” (DCB) that combines the best features of both of these acquisition/processing methods to yield a gated blood pool image sequence in real-time from which bad beats have been rejected and in which late diastolic events are accurately portrayed. In addition, this approach runs efficiently in asynchronous mode on PC-type machines and is designed to exhibit enhanced performance in a true multitasking/multithreading environment. The method is based on the concept of independent “agents” that work to construct the image sequence from position and timing data continuously written into the DCB. In the following report, we describe this agent-based method, detail the performance of the method in several different computing environments and show results obtained with this method in clinical practice.

METHODS

Data acquisition conditions

The DCB technique was implemented assuming a set of nominal acquisition and processing conditions. First, we assumed that the blood pool was labelled with 20–25 mCi of Tc-99m administered in one of several forms. The scintillation camera was assumed to be equipped with an all purpose, low energy (LEAP), or a low energy high resolution (LEHR) collimator, and the system pulse height analyser set to bracket the Tc-99m photopeak with a $\pm 15\%$ window. The nominal image sequence created by the DCB technique was set to contain 32 equal duration images, each 64×64 pixels in size. This sequence, moreover, was to be created from image data occurring only during cardiac cycles with lengths within $\pm 15\%$ of the mean cycle length (the mean being determined by averaging the lengths of beats during the first 10 seconds of data collection).

Data collection is carried out continuously until about 10 million total events have been acquired, somewhere between 10 and 20 min for the average patient. In addition, a software “zoom” feature was included that allows image data occurring outside the immediate vicinity of the heart to be excluded from the image sequence, thereby improving the speed of subsequent processing. In aggregate, these collection parameters and (default) conditions insure that, under normal circumstances, an image sequence will be acquired that exhibits a high signal-to-noise ratio and in which cardiac structures and dynamics are accurately portrayed.

Dynamic circular buffering

An analysis of the logical processes required to create a gated blood pool image sequence indicates that these processes can be visualised as a series of logically independent “agents” acting on the raw position and timing data written into a “dynamic circular buffer” (Fig. 1). A program constructed from such agents should be able to take advantage of any parallel processing capability available on the chosen computer platform. The effectiveness of these agent-directed processes is limited potentially by the time required for each agent to carry out the logical task associated with that agent at count rates typical of blood pool imaging, i.e. $8\text{--}16 \text{ K s}^{-1}$.kcps (kilocounts per second).

The DCB is a variable length linear segment of 16-bit wide memory. Several types of agents act to introduce or extract data from this DCB. Three agents introduce data into the DCB.

Acquisition agent. Each time a valid scintillation event is detected, the X , Y position of the event is encoded into a computer word (since the maximum image size in this system is 256×256 pixels, one byte per coordinate is enough) and written into the DCB at the location indicated by a pointer. This agent determines whether the event occurs within the zoomed field-of-view and, at the same time formats the event so it can be addressed on the frame-buffer and on the video buffer simultaneously without additional processing. This agent also creates a reference image which allows the technician to monitor the acquisition procedure in real time.

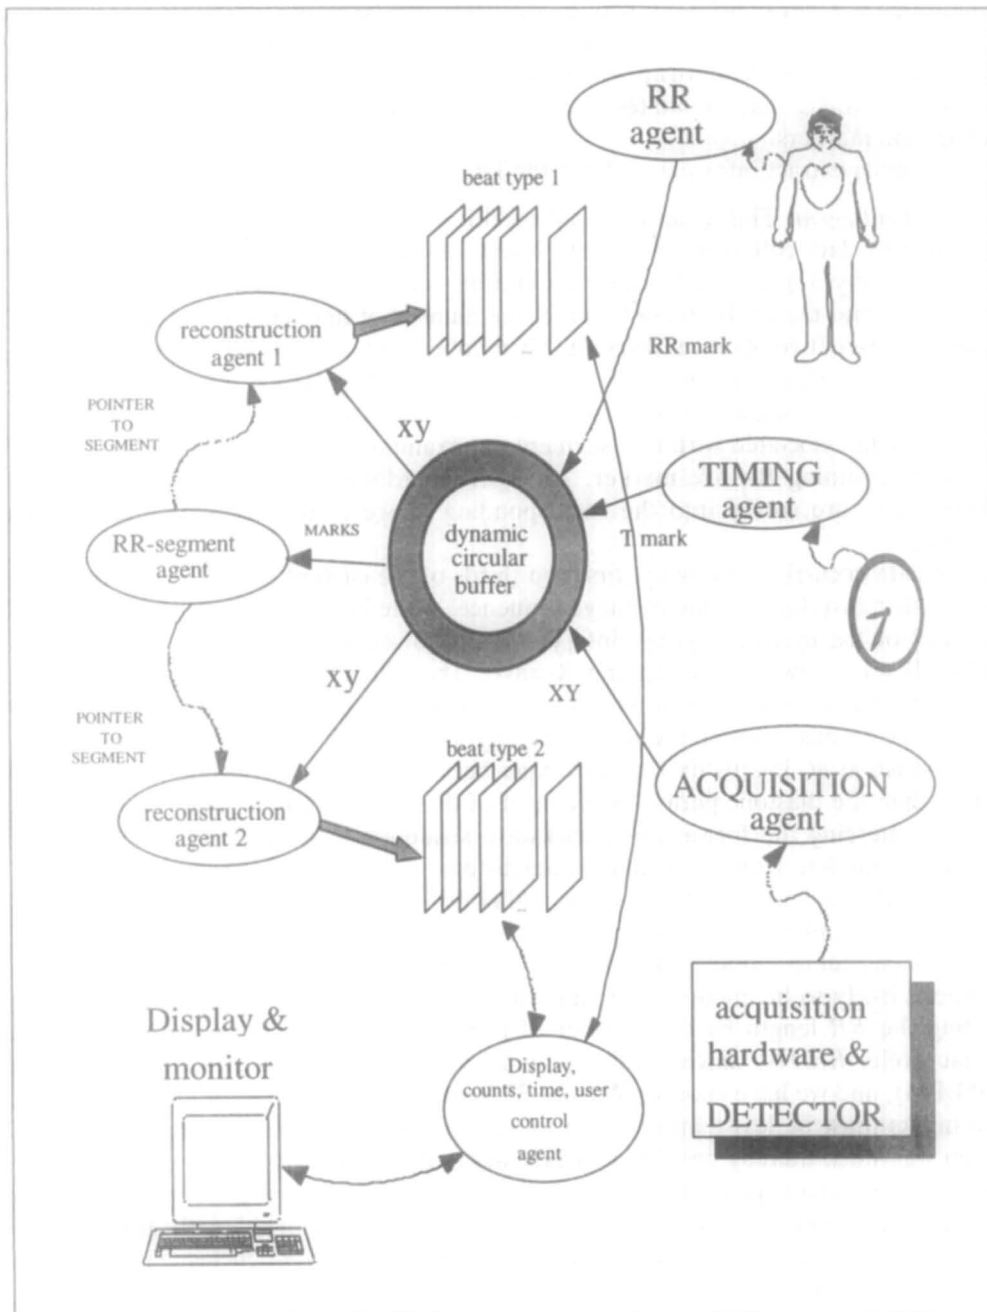


Fig. 1. Architecture of the dynamic circular buffer, showing the agents that introduce data (RR, timing and acquisition) and the ones that extract them (reconstruction and RR-segment). The figure describes the relation between the different hardware elements and the software, and among the different software modules.

Timing agent. The timing agent introduces equally-spaced, clock-generated timing marks into the DCB. The spacing between timing marks can range from a minimum of 4 ms to a maximum of 20 ms, the exact value being determined by dividing the mean cycle length by 32 and rounding to the nearest integer value. This value is exported to the reconstruction agent (see below). The effective reduction of time marks achieved by selecting a longer spacing speeds the reconstruction process since fewer timing marks need to be considered.

R-R agent. Each time an R-wave is detected, this agent introduces a marker into the DCB. This marker (and the timing markers) are fixed 16-bit words chosen such that their

magnitudes can never be X , Y coordinates; this can be easily accomplished since the FOV of the camera is circular and the area outside the circle near the vertices will never contain events, since the hardware zoom is not enabled. The software zoom provides an additional masking that eliminates any event coming from outside the FOV, taking into account the markers.

Two agents extract information from the DCB.

R-R-segment agent. This agent is the most critical to the reconstruction process. At the beginning of data collection, an empty image sequence is established that contains N images (usually 32) of equal temporal duration (nominally 14 ms). The R-R segment agent looks into the DCB and determines the number of time marks that have occurred between consecutive R-R markers. If this number indicates a cycle length outside the preselected ($\pm 15\%$) beat length window, all data occurring during that beat are ignored and the next R-R segment is examined. If within the preselected beat length window, the X , Y data associated with that segment are examined starting with the first element after the beginning R-wave marker. All X , Y coordinate values falling between two timing marks are sorted into the corresponding image in the accumulating N image sequence.

X , Y data occurring during the first two thirds of the current cycle are sorted directly into the first two thirds of the N image sequence, while image data occurring during the last third of the cycle are sorted into the N image sequence but backwards from the trailing R-wave towards the leading R-wave. The effect of this sorting method is to merge a forward framed and backward framed image sequence at the two thirds point of the average cardiac cycle, thereby eliminating data drop-off during late diastole.

Bacharach *et al.* [5] justify this technique because the systolic period is much more regular than the diastolic period. Since most of the cycle length variation occurs during diastasis, merging the forward and backward sequences near, or in, this interval will produce a complete sequence with minimal distortion due to cycle length variations. This technique also causes some data in long cardiac cycles to be lost while some data in short cardiac cycles contribute to more than one image in the sequence.

It is also useful to estimate the shortest clock period necessary to sample cardiac cycle variations likely to be encountered clinically. The frame duration, I , is determined by dividing the RR length by the number of frames (FN) required to span the average cardiac cycle, RR/FN . Since cardiac periods as short as 0.25 s can occur in practice (240 BPM), and we have chosen FN to be 32 frames, it follows that the spacing between time marks must be less than or equal to 7.82 ms. As shown earlier, the limit for our system was 4 ms, thereby meeting this extreme requirement.

The R-R segment agent allows several different average cardiac cycles to be produced simultaneously simply by varying the condition for acceptable beat length. It is possible, therefore, to examine the cardiac function at several cardiac frequencies in patients with cardiac arrhythmias.

This agent updates a set of pointers that indicate the start and end of a cycle validated for a given class. The system allows the reconstruction of several cycle types: normal cycles, slow and rapid. The limitation with respect to the number of simultaneous sequences that can be reconstructed depends on the CPU that is being used: in our case an Intel 486 DX2, 66 MHz CPU was used, permitting two cycle types for counting rates that do not exceed 30 kcps (mean).

Reconstruction agent. This agent transfers events from the LIST buffer to the respective frame buffers. For this purpose, this agent uses the pointers generated by the R-R agent, sequentially reading each coordinate event and incrementing the proper pixel in the image indicated by the timing agent. This simplicity allows the existence of several reconstruction agents with minimal cost in terms of CPU time.

Finally, a third type of low-priority agent exists that interacts with the DCB, the control agent. This agent performs record keeping (counts accumulated so far, etc.), display functions (shows the user an updated image of the field of view during

acquisition, etc.) and execution of orderly or emergency terminations of data collection by handling status conditions. Figure 2 and the following pseudo code illustrates the way these status conditions are handled:

```

Acquisition status:  if COUNT
                      if ECG
                        R-R-Agent
                      else
                        Acquisition Agent
                      else
Reconstruction status:if RECONSTRUCTION < ACQUISITION
                      if R-R or TMark
                        R-R-Segment Agent
                      else
                        Reconstruction Agent
                      end if

```

This pseudo code can easily be implemented in sequential functions (as has actually been done), but it is possible to port it to a true multiprocessing environment. In such a case the piping has to be done using interprocess communications (IPC, semaphores, etc.) and shared memory segments instead of global variables accessed in a single program or process. The optimal environment would be a true multiprocessing operating system in which a shared memory segment is established, and where all the agents run as separate processes. Additionally, it would be possible to program the acquisition module as a *daemon*, which can be controlled by a specific superserver, so no CPU time is consumed while the status conditions do not allow the DCB process to be activated. This approach provides optimal usage of the available processing bandwidth, as the status conditions thread the tasks in their preferred running priority, while still keeping all of the actual process table small.

For the visualization of the images we used the system SVGA, with 1 megabyte memory, and a video mode with a resolution of 800×600 pixels and 64 gray levels (VGA standard). Although this pixel depth is sufficient for the cardiac application, other nuclear medicine applications, such as bone scans, need 256 gray levels. Thus, we modified the standard SVGA adding a 24 bit (8×3 bits) RAM-DAC, in place of the conventional 18 bits component. This modification to the system has its own driver that handles 256 gray levels conveniently, with programmable gamma correction, and is an inexpensive mean of obtaining flicker free movies of the working heart.

RESULTS

Figure 3 depicts performance of the DCB approach on several different platforms. We measured the quality of the software architecture as the percentage of lost counts, since this is the real throughput of the whole system; the time-outs of the gammacamera are independent. The measurement conditions were:

Maximum clock cycles per count:	254 (80486; without counting context switches)
Peak count rate:	60 kcps
Measured CPUs:	80386/25 MHz w/80387, 80386/33 MHz w/80387, 80486/33 MHz, and 80486/66 MHz.

The 386 CPUs behave poorly due to the “context switch cutoff” discussed in the next section, and this is the only cause for saturation at these count rates. For 486 CPUs with a 33 MHz clock losses at 50 kcps are 17%, the limit of acceptability in this application. A 66 MHz clock reduces these losses to 2%, an acceptable value. At 60 kcps the slowest processor is saturated and all the others have losses (except the SPARC 20, but those data are just estimations deduced from the simulations).

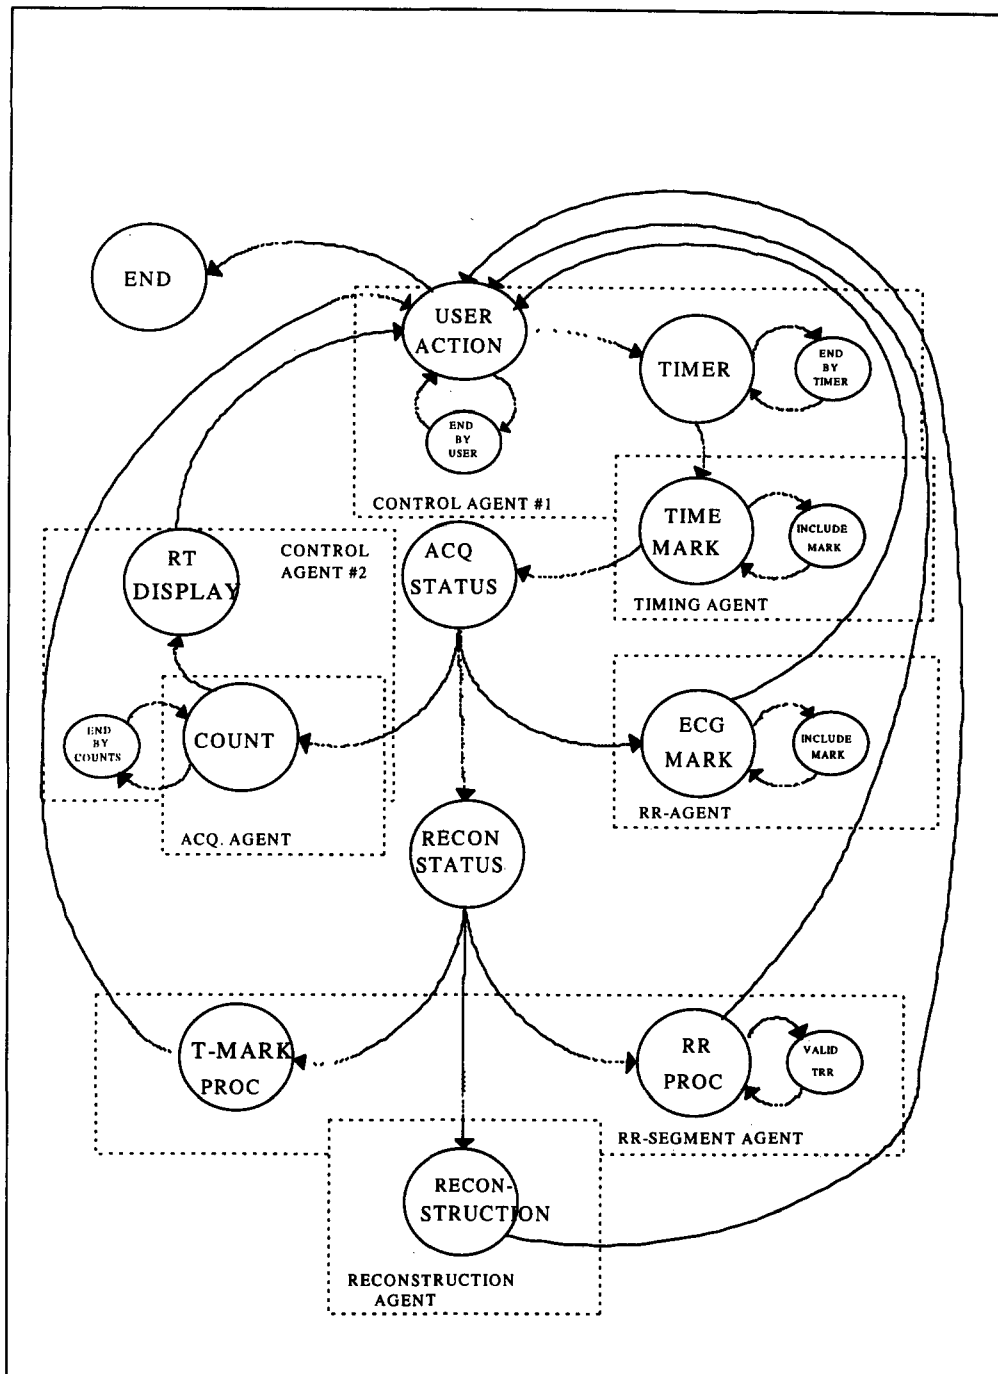


Fig. 2. State diagram. In a serial implementation the different possible states of the system should be carefully managed for a good performance. Boxes indicate which agent manages each state according to the described protocol.

Figure 4 shows an example of an acquired series: 32 images form the sequence starting from end diastole, going to systole, and ending in end diastole again. The graphic besides the images is a plot of the activity within an area of interest traced around the left ventricle for each image versus time during the cardiac cycle. The valley corresponds to end systole, while the peaks occurs at end diastole. The background activity is plotted as an indication of the correctness of the reconstruction: if there were a favoured frame, background activity would not be constant during the cycle.

In Fig. 5 cardiac anatomy is clearly evident at end diastole: both ventricles are filled




Fig. 3. CPU performance comparison. Arrows indicates the first cutoff frequency due to the CPU performance limitation, while diamonds show the limitation imposed by the context switching for interruptions handling.

with labelled blood, so the cavities can be differentiated. The system displays monitor mode images in real time similar to this one, and flicker-free cine with 256 gray levels or colour in full screen presentation. This quality is good enough for on-screen image evaluation. Pseudo-colour display is also provided, and LUT manipulation (window height and width) is available whenever images are displayed.

DISCUSSION

The program was initially developed for an Intel 80386/80486 with DOS operating system, but the final target is intended to be a multiprocessing system like a SPARC 20 mainboard running Solaris 2.X. An important difference between the 80386 and 80486 architectures, as far as our system is concerned, is the existence of 32 bit Interrupt




Fig. 4. On the left are displayed the 32 frames reconstructed during the acquisition without any postprocessing. Acquisition starts at end diastole and runs for a cardiac cycle. On the right time-activity curves for the left ventricle and the background are shown. Background activity remains almost constant indicating that there are no "faint" image frames due to the reconstruction.




Fig. 5. A zoom of the left anterior oblique (LAO) image of the heart at the end of the diastole (ED) displays the anatomy: left ventricle (LV) and right ventricle (RV) can be seen clearly.

Look-Up tables for the 80486. This allows excellent performance of the control agent and the timing agent, since the interrupts are handled directly in the 32 bit mode, while in 80386 systems, the interrupts are intercepted in real mode (16 bits) and then passed over to the 32 bit interrupt handler. This operation implies two context commutations per time-mark, which at a rate of 10 ms per mark, is time consuming and decreases the overall performance of the system. Other context commutations appear at the BIOS interrupt handling (User Interaction, count display). A context commutation can consume more than 1000 clock cycles (depending on the architecture); consequently, the 80486 architecture has been chosen, as the price increment is not noticeable in the full price of the complete system. Additionally, a 66 MHz clock main board has been used, since it assures a minimum loss of counts during acquisition. The SPARC 20 architecture under a multiprocessing O.S. optimises resources and task switching, so significant performance improvements should be attainable.

CONCLUSION

We have designed a technique called “dynamic circular buffering” (DCB), and implemented it fully in software. This technique is efficient in constructing gated blood pool image sequences of the heart in real time and combines the advantages of both frame and list mode acquisition methods.

When implemented on low cost serial PC-type computer hardware, it performs these studies fast enough to avoid count losses up to a count rate of 60 kcps. On parallel type machines, the program should be even more efficient.

The technique has been evaluated in a clinical environment, with positive results, and is currently integrated in a commercial system.

Acknowledgements—This work has been supported by the Comunidad Autónoma de Madrid, Ayuda de Investigación C070/90, and Investigación Técnica Industrial, S.A. (ITISA), Madrid.

REFERENCES

1. C. Schiepers and J. J. Almasi, Equilibrium gated blood pool imaging at rest and during exercise. In *Effective Use of Computers in Nuclear Medicine*, M. J. Gelfand and S. R. Thomas, Eds, p. 136. McGraw-Hill, New York (1988).
2. S. L. Bacharach, M. V. Green, J. S. Borer, M. A. Douglas *et al.*, A real-time system for multi-image gated cardiac studies. *J. Nucl. Med.* **18**, 79–84 (1977).
3. S. L. Bacharach, M. V. Green, D. Vitale, G. White *et al.*, Optimum Fourier filtering for cardiac data; a minimum-error method, *J. Nucl. Med.* **24**, 1176–1184 (1983).
4. J. L. Lear and J. P. Pratt, Real-time list-mode processing of gated cardiac blood pool examinations with forward-backward framing, *Eur. J. Nucl. Med.* **19**, 177–180 (1992).
5. S. L. Bacharach, M. V. Green and J. S. Borer, Instrumentation and data processing in cardiovascular Nuclear medicine: evaluation of ventricular function, *Seminars in Nuclear Medicine*, **IX**, 257 (1979).
6. L. L. Johnson, R. A. Rodney, R. A. Vaccarino, P. Egbe *et al.*, Left ventricular perfusion and performance from a single radiopharmaceutical and one camera, *J. Nucl. Med.* **33**, 7 1411–1416 (1992).
7. R. H. Jones, P. McEwan, G. E. Newman, S. Port *et al.*, Accuracy of diagnosis of coronary artery disease by radionuclide measurement of left ventricular function during rest and exercise, *Circulation* **64**, 586–600 (1981).
8. J. Candell and D. Ortega, *Cardiología Nuclear*. Ediciones Doyma, Madrid (1992).

About the Author—JUAN JOSÉ VAQUERO received his Ph.D. degree in Telecommunication Engineering from the Technical University of Madrid (UPM). After working for Siemens as Head of the Cardiology Group, he joined the Grupo de Bioingeniería y Telemedicina of the UPM in 1988, as a research fellow. His major interest is now magnetic resonance imaging processing and its instrumentation.

About the Author—HELLMER RAHMS has been a research fellow at the Grupo de Bioingeniería y Telemedicina of the Technical University of Madrid (UPM), since 1990. His main working areas are medical image processing and telemedicine, where he has been involved in research and development in several EU projects. His present interests focus on broadband communications for medical applications and services.

About the Author—MICHAEL V. GREEN is head of the Imaging Physics Section, Nuclear Medicine Department, Clinical Center, National Institutes of Health, Bethesda, Maryland. During his tenure at NIH, Mr Green has worked in nuclear cardiology, positron emission tomography and is currently involved in the development of advanced tracer imaging systems for use in animals and man.

About the Author—FRANCISCO DEL POZO is Full Professor at the Technical University of Madrid (UPM), and Head of the Grupo de Bioingeniería y Telemedicina. After many years of bioengineering research, mainly on chronobiology, diabetes, health care telematics and medical imaging, he is presently leading several EU projects focused on Telemedicine and Medical Imaging. He is a member of several European and International Societies.