

Extending the dynamic range of phase contrast magnetic resonance velocity imaging using advanced higher-dimensional phase unwrapping algorithms

M.F Salfity, J.M Huntley, M.J Graves, O Marklund, R Cusack and D.A Beauregard

J. R. Soc. Interface 2006 **3**, 415-427
doi: 10.1098/rsif.2005.0096

References

[This article cites 15 articles](#)

<http://rsif.royalsocietypublishing.org/content/3/8/415.full.html#ref-list-1>

Email alerting service

Receive free email alerts when new articles cite this article - sign up in the box at the top right-hand corner of the article or click [here](#)

To subscribe to *J. R. Soc. Interface* go to: <http://rsif.royalsocietypublishing.org/subscriptions>

Extending the dynamic range of phase contrast magnetic resonance velocity imaging using advanced higher-dimensional phase unwrapping algorithms

M. F. Salfity^{1,*}, J. M. Huntley¹, M. J. Graves², O. Marklund³, R. Cusack⁴
and D. A. Beauregard⁴

¹ Wolfson School of Mechanical and Manufacturing Engineering, Loughborough University, Loughborough LE11 3TU, UK

² Department of Radiology, Addenbrooke's Hospital, Hills Road, Cambridge CB2 2QQ, UK

³ Luleå University of Technology, 97187 Luleå, Sweden

⁴ MRC Cognition and Brain Sciences Unit, 15 Chaucer Road, Cambridge CB2 2EF, UK

Phase contrast magnetic resonance velocity imaging is a powerful technique for quantitative *in vivo* blood flow measurement. Current practice normally involves restricting the sensitivity of the technique so as to avoid the problem of the measured phase being ‘wrapped’ onto the range $-\pi$ to $+\pi$. However, as a result, dynamic range and signal-to-noise ratio are sacrificed. Alternatively, the true phase values can be estimated by a phase unwrapping process which consists of adding integral multiples of 2π to the measured wrapped phase values. In the presence of noise and data undersampling, the phase unwrapping problem becomes non-trivial. In this paper, we investigate the performance of three different phase unwrapping algorithms when applied to three-dimensional (two spatial axes and one time axis) phase contrast datasets. A simple one-dimensional temporal unwrapping algorithm, a more complex and robust three-dimensional unwrapping algorithm and a novel velocity encoding unwrapping algorithm which involves unwrapping along a fourth dimension (the ‘velocity encoding’ direction) are discussed, and results from the three are presented and compared. It is shown that compared to the traditional approach, both dynamic range and signal-to-noise ratio can be increased by a factor of up to five times, which demonstrates considerable promise for a possible eventual clinical implementation. The results are also of direct relevance to users of any other technique delivering time-varying two-dimensional phase images, such as dynamic speckle interferometry and synthetic aperture radar.

Keywords: phase unwrapping; phase contrast velocity imaging; blood flow measurement; image processing; magnetic resonance imaging; speckle interferometry

1. INTRODUCTION

Many measurement techniques across a variety of engineering, scientific and medical disciplines deliver quantitative information in the form of phase images. For example, in the field of optical interferometry, phase extraction through the recording of phase-shifted intensity images has become the standard method for automatically analysing the shapes of two interfering wavefronts. In experimental mechanics, similar analysis methods are used routinely to process fringe patterns from techniques such as moiré photography and photoelasticity. Synthetic aperture radar maps terrain and deformation of the Earth's surface through phase images on length-scales some 10^4 – 10^6 times

greater than the first two examples. One of the areas of fastest growth is probably in phase contrast medical imaging techniques, notably magnetic resonance imaging (MRI), X-ray interferometry and optical coherence tomography. In all these fields, advances in hardware and reduction in cost of acquisition mean that three-dimensional datasets (typically three spatial axes for tomographic measurements, two spatial axes and a time axis for dynamic problems) is now commonplace and may soon become the default choice over the two-dimensional datasets that were hitherto the norm. The parallel development of numerical algorithms able to handle such three-dimensional phase volumes—and exploration of potential benefits from algorithms working on datasets with still higher dimensionality—is important if one is to maximize the potential benefits to society of these techniques.

*Author for correspondence (m.f.salfity@lboro.ac.uk).

The most difficult of these algorithms, and the one that is central to all the techniques mentioned above, is that of phase unwrapping. This is the process of adding integral multiples of 2π to the measured ‘wrapped’ phase (i.e. phase values that lie in the range $-\pi$ to $+\pi$) at all the pixels or voxels in the image. In the case of MRI, for example, the measured phase is wrapped because we can only determine the final phase of the spins, and not the number of rotations taken to get there. While the phase is proportional to the physical quantity of interest for small changes, for larger changes, it is wrapped onto the range $-\pi$ to $+\pi$.

Phase unwrapping is a trivial process on data that is free of noise and has been sampled in accordance with the Shannon sampling theorem; however, practical datasets almost always contain regions that are noisy or sub-sampled, leading to ‘singular’ points (in two dimensions) or lines (in three dimensions) in the phase field. In two dimensions, the singular points are identified by counting the number of 2π phase jumps around all possible 2×2 pixel squares in the image. A positive or $+1$ singularity is present if the result of the integration is $+1$, and a negative or -1 singularity is present if the result is -1 . The only other possible resulting value is 0, in which case no singularity is present. The term singularity (as opposed to the terms ‘pole’ and ‘residue’ which are sometimes used in the literature) will be used throughout this paper for consistency with Huntley (2001). The effect of such singular points on the unwrapping process is illustrated in the wrapped two-dimensional phase map of figure 1. The number of 2π phase jumps required to unwrap point Q, given the phase at P, is path-dependent: path A crosses three phase jumps, whereas path B crosses two. The presence of the two phase singularities, points 1 and 2, caused by local undersampling of the phase map, can, therefore, result in the global propagation of large (2π) phase errors.

1.1. Existing one-dimensional, two-dimensional and three-dimensional phase unwrapping algorithms

While from a rigorous mathematical viewpoint it is impossible to recover the unwrapped phase unambiguously in the presence of singularities, additional *ad hoc* assumptions can often be made to come to a solution that is acceptably close to the true unwrapped phase in many situations. A large variety of different two-dimensional unwrapping algorithms has been developed to reduce the influence of the singular points (Ghiglia & Pritt 1998). They can be broadly classified as path-following methods, which guide the unwrapping path in order to prevent error propagation, and minimum-norm methods, which seek to minimize a cost function that measures the difference between the gradients of the original wrapped phase and the estimated unwrapped phase. The minimum-norm methods have some drawbacks such as requiring the use of iterative solution methods and in certain cases introducing systematic errors to the reconstructed phase field (Ghiglia & Pritt 1998).

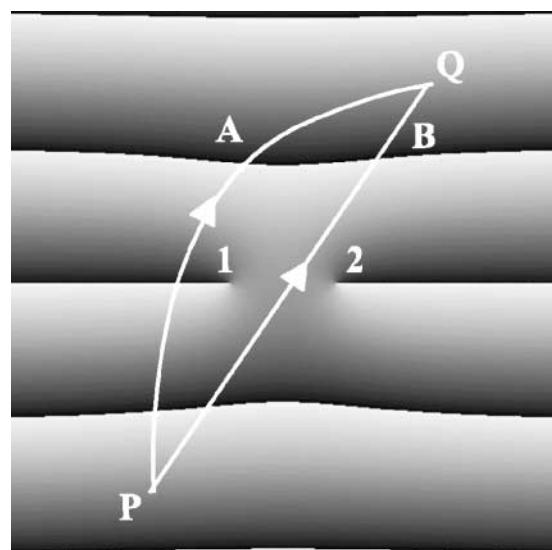


Figure 1. Wrapped phase map (black and white representing, respectively, $-\pi$ and $+\pi$) containing two singular points 1 and 2 results in path-dependent unwrapped phase at point Q. (Reprinted, by permission, from Huntley & Saldner 1993).

Among the path-following methods, one of the most direct is to place branch cuts within the phase map between pairs of singular points of opposing signs to make the unwrapped results independent of the unwrapping path (Huntley 1989). In the example shown in figure 1, point 1 would be joined to point 2 by the cut and path B would not be allowed. All paths not crossing the branch cut then agree on the integral multiple of 2π to be added to the phase at Q.

There are relatively few published papers on the full three-dimensional phase unwrapping problem. The simplest approach is to unwrap each slice independently using one of the two-dimensional methods. However, this requires correction of the resulting phase offset among slices. Furthermore, taking into account the three-dimensional nature of the data can be an advantage to improve the unwrapping process, as in the extension to three dimensions of the two-dimensional branch cut method (Huntley 2001).

In three dimensions, the singularities form closed phase singularity loops (PSLs). A path-independent unwrapping process requires branch surfaces limited by the closed loops to be established. The advantage of using the full three-dimensional volume as opposed to unwrapping each two-dimensional slice independently is the elimination of the ambiguity that arises when pairing singularities of opposing signs in a two-dimensional phase map. As pointed out by Huntley (2001), the presence of knot points, i.e. cubes of $2 \times 2 \times 2$ voxels into which two PSLs enter and out of which two PSLs leave, still represents an ambiguity in that it is not known whether the singularities form a single loop or two separate but linked loops. Subsequent developments in the algorithm (Marklund *et al.* 2005) now use a recursive tree structure to find the set of loops with minimum length. The concept of loop structures within three-dimensional phase volumes was also discovered independently by Cusack & Papadakis (2002) and by

Chavez *et al.* (2002), though it was not developed into a full path-independent algorithm in either case.

An N -dimensional phase unwrapping method was presented by Jenkinson (2003). This method divides the volume into regions that contain no phase wraps. A cost function that measures the difference in phase values at the interface of adjacent regions is minimized. Those regions for which the cost function is minimum are merged after adding the optimum offset found, and the process continues until there is only one region left. Although intuitively appealing, one significant problem with this approach is that if the measured wrapped phase is undersampled there will be 2π phase jumps missing which will never be recovered by Jenkinson's algorithm.

Temporal unwrapping, originally proposed for speckle interferogram analysis (Huntley & Saldner 1993, 1997a) and subsequently introduced independently by Xiang (1995) in MRI applications, is another method that can be used to unwrap CINE phase contrast magnetic resonance imaging (PC-MRI) phase volumes. The phase of each pixel as a function of time is unwrapped independently of the rest of the pixels in the image. The advantage of this one-dimensional approach is that, as opposed to spatial unwrapping, boundaries and regions with poor signal-to-noise ratio do not affect the unwrapping result in regions where the data is good. However, it relies on adequate sampling along the time axis. Another potential problem is cardiac motion, which can cause unwrapping errors due to misregistration of the images from frame to frame. Yang *et al.* (1996) introduced a motion-registered temporal unwrapping method followed by spatial phase unwrapping to overcome this problem.

1.2. Velocity mapping by phase contrast MRI

The phase imaging technique used throughout this paper (PC-MRI) is based on detecting changes in the phase of the transverse magnetization of the blood as it moves in the presence of a magnetic field gradient. PC-MRI uses bipolar gradient pulses to produce a linear relationship between the velocity of blood v and the phase shift of the magnetic resonance signal ϕ :

$$v(\phi) = \frac{\phi \cdot V_{\text{enc}}}{\pi}, \quad (1.1)$$

where V_{enc} is a constant called the velocity encoding parameter. V_{enc} is related to the amplitude and duration of the bipolar gradient pulse, and is equal to the velocity that results in a phase shift of π radians (McRobbie *et al.* 2003). Combining PC-MRI with cardiac triggering, using CINE PC-MRI, results in a temporal series of two-dimensional phase contrast MR images of a single slice at multiple time points (phases) throughout the cardiac cycle to form a three-dimensional (two spatial axes and one time axis) phase volume.

When V_{enc} is smaller than the peak blood velocity, the phase is wrapped onto the range $[-\pi, \pi]$, and it is necessary to unwrap it in order to obtain the true velocity values. The standard procedure at present is to avoid the phase wrapping problem by choosing V_{enc}

values equal to or larger than the maximum velocity of the blood. However, this may compromise the accurate measurement of the lower velocities close to the walls of the arteries, an important requirement for measuring wall shear stresses, which are proportional to the radial derivative of the velocity at the wall.

The purpose of this paper is to present the results of a comparison between the performances of the path-independent three-dimensional phase unwrapping algorithm and the temporal unwrapping method on a set of phase contrast CINE MRI images obtained with different encoding velocities. A modification to the first algorithm is introduced in order to process particular PSLs that may appear in spatial-temporal volumes. Exploitation of data with still higher dimensionality is considered through a variation on the temporal approach, in which velocity encoding is used as the unwrapping direction along, in effect, a fourth dimension. The performance of this new algorithm is evaluated. The improvements in signal-to-noise ratio which result from the lower V_{enc} values that such unwrapping algorithms permit are quantified, and typical velocity and wall shear strain rate profiles are presented. Finally, a general proof of the continuity of the PSLs is presented for the first time. Although the application considered here is specific to PC-MRI, the results are also of direct relevance to users of any technique delivering time-varying two-dimensional phase images.

2. METHODS

2.1. Phase contrast CINE MRI

Six phase contrast CINE MRI acquisitions through the ascending aorta and descending aorta of a single subject were obtained using a 1.5 T magnetic resonance scanner (Echospeed, GE Healthcare Technologies, Waukesha, WI), with velocity encoding values of 25, 50, 75, 100, 125 and 150 cm s^{-1} . Velocity encoding values of 150 and 125 cm s^{-1} were chosen to avoid phase wrapping, since the peak blood velocity was $V_p = 120 \text{ cm s}^{-1}$. As V_{enc} decreased, the phase was increasingly wrapped (figure 2). Details of the PC-MRI acquisitions and preprocessing can be found in appendix A. Although maximum-likelihood filters can be highly effective at reducing the number of noise-induced singular points in techniques such as speckle interferometry where there are strong inter-pixel variations in the signal magnitude (Huntley 1997), this is not the case in PC-MRI. The resulting loss of spatial resolution can have the counter-productive effect of increasing the size of the larger loops. No smoothing filter was, therefore, used in the data analysis described here.

2.2. Three-dimensional phase unwrapping algorithm

As explained in §1, the three-dimensional phase unwrapping algorithm is an extension to three dimensions of a two-dimensional version in which branch cut lines are set between pairs of phase singularity points of

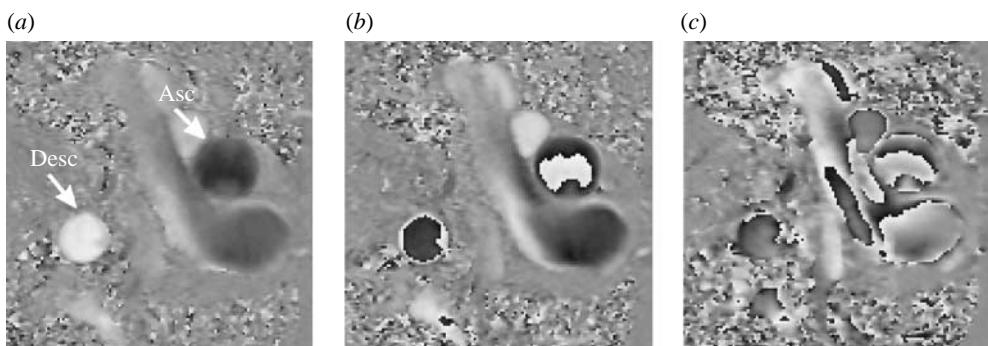


Figure 2. Wrapped phase maps obtained for the same point in the cardiac cycle, with velocity encoding values of (a) 125, (b) 75 and (c) 25 cm s^{-1} . The ascending (Asc) and descending (Desc) aorta are identified in (a).

opposing signs. The example shown in figure 1 was somewhat over-simplistic in that it contained only a single pair of singularities. When more than one pair is present in the image, a major problem arises: which positive singularity should be joined to each of the negative singularities? In two dimensions, it is not possible to answer this question unambiguously, and it is necessary to use statistical arguments. Buckland *et al.* (1995), for example, showed that under certain circumstances the maximum-likelihood pairing is that which minimizes the sum of the squares of the cut lengths, and an efficient algorithm was presented to achieve this.

In three dimensions, additional information is provided by the fact that singular points can now be calculated around 2×2 squares oriented normal to each of three different perpendicular directions. Huntley (2001) showed that under sufficiently low singularity densities, the singular points identified in this way must link up to form continuous PSLs. The presence of $+1/-1$ singularity pairs or ‘dipoles’ in two dimensions may then be interpreted as the intersection of a loop in three dimensions with the two-dimensional surface (see figure 3). The two-dimensional problem of how to pair the singular points does not arise in three dimensions, since the points are unambiguously linked by the loop. The placement of branch cut surfaces across all loops in the phase volume results in a path-independent solution. In appendix B, we relax the low-singularity-density requirement, and prove for the first time that all singular points must link up to form continuous loops, no matter how high the singularity density. This is an important result when dealing with the lowest V_{enc} phase volumes. Despite the additional information compared to the two-dimensional case, ambiguities still arise due to the problem of knot points (elemental $2 \times 2 \times 2$ voxel cubes into which two loops enter and out of which two loops leave). The ambiguity arises when sorting the singularities into loops, because different loop configurations are possible. In these cases, a recursive tree structure is used to find the set of loops with shortest lengths (Marklund *et al.* 2005).

Loops that end on the phase volume boundary, called partial loops, are closed by connecting their ends with artificial singularities added along the phase volume boundary. PSLs that extend along one of the axes but whose top and bottom parts are not

included in the phase volume, as shown in figure 4a, do not normally arise in phase volumes with three spatial axes, but can occur in datasets with two spatial and one temporal dimension. In these cases it is not correct to connect the ends of each vertical loop segment by adding artificial singularities along the boundary (figure 4b). Instead, the two portions should be joined together. To achieve this, the vertical partial loops that end on the first and last slices are identified. The Hungarian algorithm (Buckland *et al.* 1995) is used to pair loop segments that run in opposite directions, using a cost function equal to the sum of the distances between the ends of the two loop segments on the first and last slices. This type of loop occurred in regions of high noise and was not common; fewer than 10 cases were found in the data presented here. After all loops were identified and closed, the branch surface for each was set by shrinking the loop towards its geometric centre. Once all the flags indicating the branch surface locations had been set, the unwrapping was carried out with a flood fill algorithm. This algorithm is robust in the sense that the presence of branch surfaces on the PSLs makes the results independent of the unwrapping path. If a phase volume contains PSLs and the corresponding branch surfaces are not placed, the unwrapped phase volume will have errors and it will not be possible to extract meaningful information from it. Robustness of the algorithm against MRI artefacts, motion, or low signal-to-noise ratio, were not evaluated, as these effects can generally be corrected for during the acquisition process. It is worth noting that this algorithm is fully automatic and does not require user interaction nor optimization of any parameters.

The data was processed on a Sun Blade 100 workstation, with a 500 MHz UltraSPARC IIe processor and 1 GB RAM. The time taken for the three-dimensional phase unwrapping algorithm to unwrap each phase volume of $256 \times 256 \times 20$ voxels, masked by applying a threshold on the corresponding magnitude volume, varied between 15 and 60 s, being higher for the lowest velocity encoding value where more singularities were present. The routine that sorts the singularities into loops is currently the most time-consuming step, but this could be improved by using more efficient sorting methods.

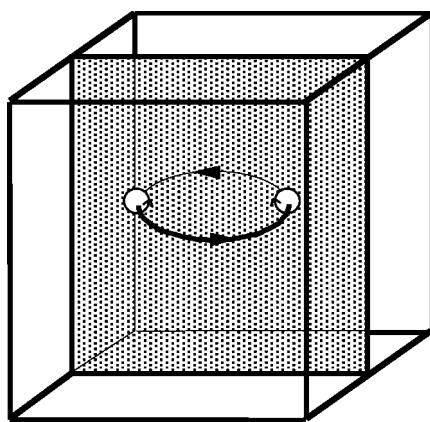


Figure 3. Dipole pair of singular points occurs on a two-dimensional surface intersected by a phase singularity loop. (Reprinted with permission from Huntley (2001)).

2.3. Temporal unwrapping algorithm

Unwrapping the three-dimensional phase data along the time axis was carried out by the following steps. The input to the algorithm was a set of two-dimensional wrapped phase maps, $\phi_w(x, y, t)$, where (x, y) refers to the pixel location and t represents non-dimensional time ($t=0, 1, \dots, s$). Subscripts 'w' and 'u' will be used to denote wrapped and unwrapped phases, respectively.

The phase change at a given pixel between times i and j is given by

$$\Delta\phi(i, j) = \phi(i) - \phi(j), \quad (2.1)$$

where the pixel coordinates have been dropped for reasons of clarity. This phase change is wrapped onto the range $(-\pi, \pi)$ by the unwrapping operator U :

$$\Delta\phi_w(i, j) = U\{\Delta\phi(i, j), 0\}. \quad (2.2)$$

$U\{\phi_1, \phi_2\}$ subtracts an integral multiple of 2π from ϕ_1 such that $\phi_1 - \phi_2$ lies in the range $-\pi$ to $+\pi$:

$$U\{\phi_1, \phi_2\} = \phi_1 - 2\pi \text{NINT}\left[\frac{\phi_1 - \phi_2}{2\pi}\right], \quad (2.3)$$

where $\text{NINT}[\dots]$ denotes rounding to the nearest integer. The term unwrapping operator is used because, although in this special case ($\phi_2=0$) it results in wrapped phases, in general it is used to carry out the unwrapping operations required in §3.

The total unwrapped phase change at the time of the t th phase map measurement can then be calculated by simply summing the wrapped phase differences:

$$\phi_u(t) = \sum_{t'=1}^t \Delta\phi_w(t', t'-1), \quad (2.4)$$

where $\phi(0)$ is defined to be equal to 0.

2.4. 'Temporal' unwrapping algorithm along velocity encoding axis

The temporal phase unwrapping method can also be applied with velocity encoding instead of time as the unwrapping direction. This approach makes use of four dimensions: x , y , t and a parameter $g = V_{\text{enc}}^{\max} / V_{\text{enc}}$, where V_{enc}^{\max} is the maximum V_{enc} value used. V_{enc}^{\max} is

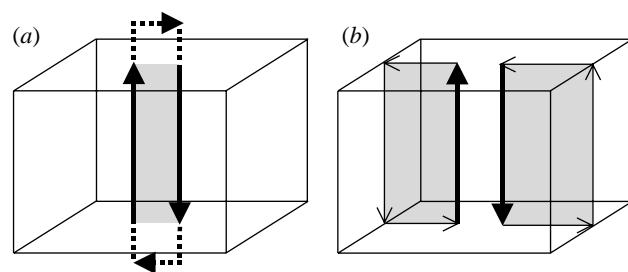


Figure 4. (a) A single loop that appears as two distinct segments within the phase volume. Each segment ends on the first and last slices of the volume. The shaded area represents the correct branch surface. (b) Each segment is closed by joining its ends with artificial singularities added along the phase volume boundary thereby setting the wrong branch surface.

chosen to be equal to or slightly greater than the maximum blood velocity so that no phase wraps occur in the $g=1$ phase volume. The method makes use of the fact that the unwrapped phase at any given (x, y, t) voxel is by equation (1.1) inversely proportional to V_{enc} . The phase at that voxel is unwrapped independently of the rest of the voxels. The method, therefore, offers a significant benefit over the other two techniques in that it does not depend on an adequate sampling rate along any of the spatial or temporal axes.

In order to ensure sufficient sampling along the g dimension one should choose the V_{enc} values such that g forms the linear sequence $g=1, 2, 3, \dots, s$ (Huntley & Saldner 1997a). Such an approach involves an increase in acquisition time over the previous unwrapping methods by a factor s which is generally undesirable. However, the fact that the expected value of the unwrapped phase scales with g allows one to under-sample along the g -axis. Two different strategies—originally proposed for the analysis of data from an optical shape measurement system (Huntley & Saldner 1997a)—are compared in this paper. In the first two values, $g=1$ and $g=s$ are used. In the second, a growing exponential sequence $g=1, 2, 4, 8$ (i.e. $s=8$) is used which requires four velocity encoding values.

In the first case, the volume with lower V_{enc} ($g=s$) is unwrapped with respect to the higher V_{enc} volume ($g=1$), scaled by the ratio s :

$$\phi_u(s) = U\{\phi_w(s), s\phi_w(1)\}, \quad (2.5)$$

where the spatial and temporal indices have been dropped for reasons of clarity. The 150 cm s^{-1} volume was used to unwrap the lowest V_{enc} volume of 25 cm s^{-1} . This approach was also used in Lee *et al.* (1995), Herment *et al.* (2000) and Xiang (2001). Xiang (2001) showed that only three acquisitions were needed by taking the reference volume only once.

For the second approach, a growing exponential sequence of V_{enc}^{-1} values of 0.005, 0.01, 0.02 and 0.04 s cm^{-1} (i.e. velocity encoding values equal to 200, 100, 50 and 25 cm s^{-1} , corresponding to $g=1, 2, 4, 8$) was used to unwrap each voxel. The factor of 2 increase in sensitivity between each acquisition is not mandatory, but was found by Huntley & Saldner (1997a) to be close to optimal. A smaller ratio, for example, reduces the risk of an unwrapping failure

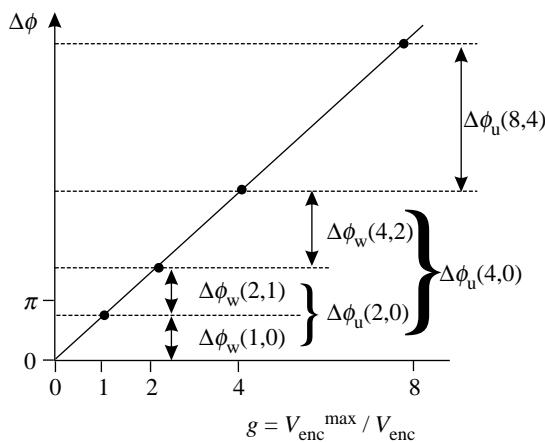


Figure 5. Forward exponential implementation of the ‘temporal’ phase unwrapping algorithm along the velocity-encoding axis in which a given (x, y, t) voxel is unwrapped.

between any two steps, but the larger number of steps required can then result in an overall higher risk of an error in the final unwrapped phase. Since the volume with $V_{\text{enc}} = 200 \text{ cm s}^{-1}$ had not been acquired, it was simulated by re-scaling the 150 cm s^{-1} phase values by a factor $150/200$. With these four points, the unwrapping proceeds as follows. The first wrapped phase value in the exponential sequence is used to unwrap the phase difference between the first and second values. In turn, this unwrapped phase difference is used to unwrap the difference between the second and third points in the sequence, which is finally used to unwrap the difference between the third and fourth points.

Expressed mathematically, the wrapped increment $\Delta\phi_w(2,1)$ is unwrapped using the phase value $\phi_w(1)$, as follows:

$$\Delta\phi_u(2,1) = U\{\Delta\phi_w(2,1), \phi_w(1)\}, \quad (2.6)$$

where $\phi_w(1)$ is not wrapped and, therefore, equal to $\phi_u(1)$.

$\Delta\phi_u(2,0)$ is calculated as

$$\Delta\phi_u(2,0) = \Delta\phi_u(2,1) + \phi_w(1). \quad (2.7)$$

$\Delta\phi_u(2,0)$ can now be used to unwrap $\Delta\phi_w(4,2)$:

$$\Delta\phi_u(4,2) = U\{\Delta\phi_w(4,2), \Delta\phi_u(2,0)\}, \quad (2.8)$$

from which $\Delta\phi_u(4,0)$ is calculated as:

$$\Delta\phi_u(4,0) = \Delta\phi_u(4,2) + \Delta\phi_u(2,0). \quad (2.9)$$

This process is repeated using the phase values $\phi(g)$ ($g = 1, 2, 4, 8, \dots, s$). Equations (2.8) and (2.9) can be rewritten for the general case:

$$\Delta\phi_u(2g,g) = U\{\Delta\phi_w(2g,g), \Delta\phi_u(g,0)\}, \quad (2.10)$$

$$\Delta\phi_u(2g,0) = \Delta\phi_u(2g,g) + \Delta\phi_u(g,0). \quad (2.11)$$

The method is illustrated schematically in figure 5.

As pointed out by Huntley & Saldner (1997b), the phase error can be further reduced by least-squares fitting a line to the unwrapped phase values of the exponential sequence and re-estimating the value of the last point in the sequence from the best-fit gradient. The fitting is done along the g dimension and, therefore, involves no reduction in spatial resolution.

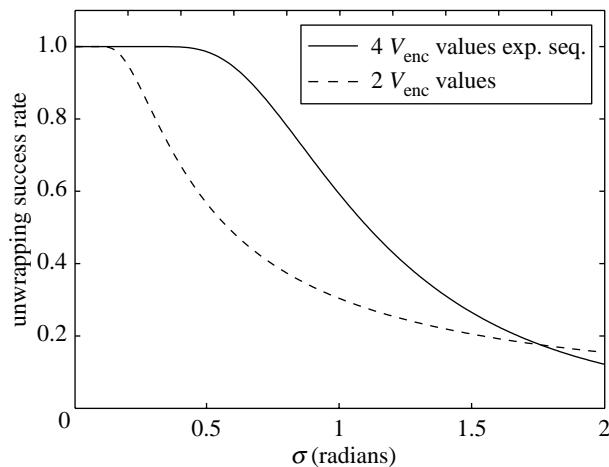


Figure 6. Unwrapping success rates for the two velocity-encoding methods using two V_{enc} values (25 and 200 cm s^{-1}) (dotted line) and an exponential sequence of four V_{enc} values (from 200 to 25 cm s^{-1}) (full line), as a function of phase noise for the case $v=0$.

Unwrapping success rates for the two methods are derived in appendix C. Comparison of the success rates shows that using an exponential sequence of four values results in improved unwrapping success rate in the presence of noise when compared to using only two velocity encoding values. This is illustrated in figure 6. The reason for this result is that in the latter method when the lower sensitivity data is scaled to perform the unwrapping, the noise is also amplified, whereas the exponential method has no noise amplification.

3. RESULTS

The four wrapped phase volumes ($V_{\text{enc}} = 100, 75, 50$ and 25 cm s^{-1}) were unwrapped using the methods described in §2.

3.1. Three-dimensional spatial unwrapping

The three-dimensional wrapped phase volumes with V_{enc} values of $100, 75$ and 50 cm s^{-1} were successfully unwrapped by the three-dimensional noise-immune phase unwrapping algorithm (figure 7). Application to the $V_{\text{enc}} = 25 \text{ cm s}^{-1}$ dataset resulted in some unwrapping errors (figure 7), which are related to undersampling of the data and to the difficulty of setting up proper branch surfaces in complex loops (see §4). The PSLs that appear in the same sub-volume for three velocity encoding values are shown in figure 8. It is clear that the number and length of the PSLs increases as the value of V_{enc} decreases. Also noticeable is the presence of horizontal PSLs in the lower V_{enc} volumes.

3.2. Temporal unwrapping

The volumes with $V_{\text{enc}} = 100$ and 75 cm s^{-1} were successfully unwrapped with the temporal unwrapping method, and the results were similar to those obtained with the three-dimensional spatial unwrapping algorithm (figure 7). However, the volumes with V_{enc} values of 50 and 25 cm s^{-1} could not be correctly unwrapped

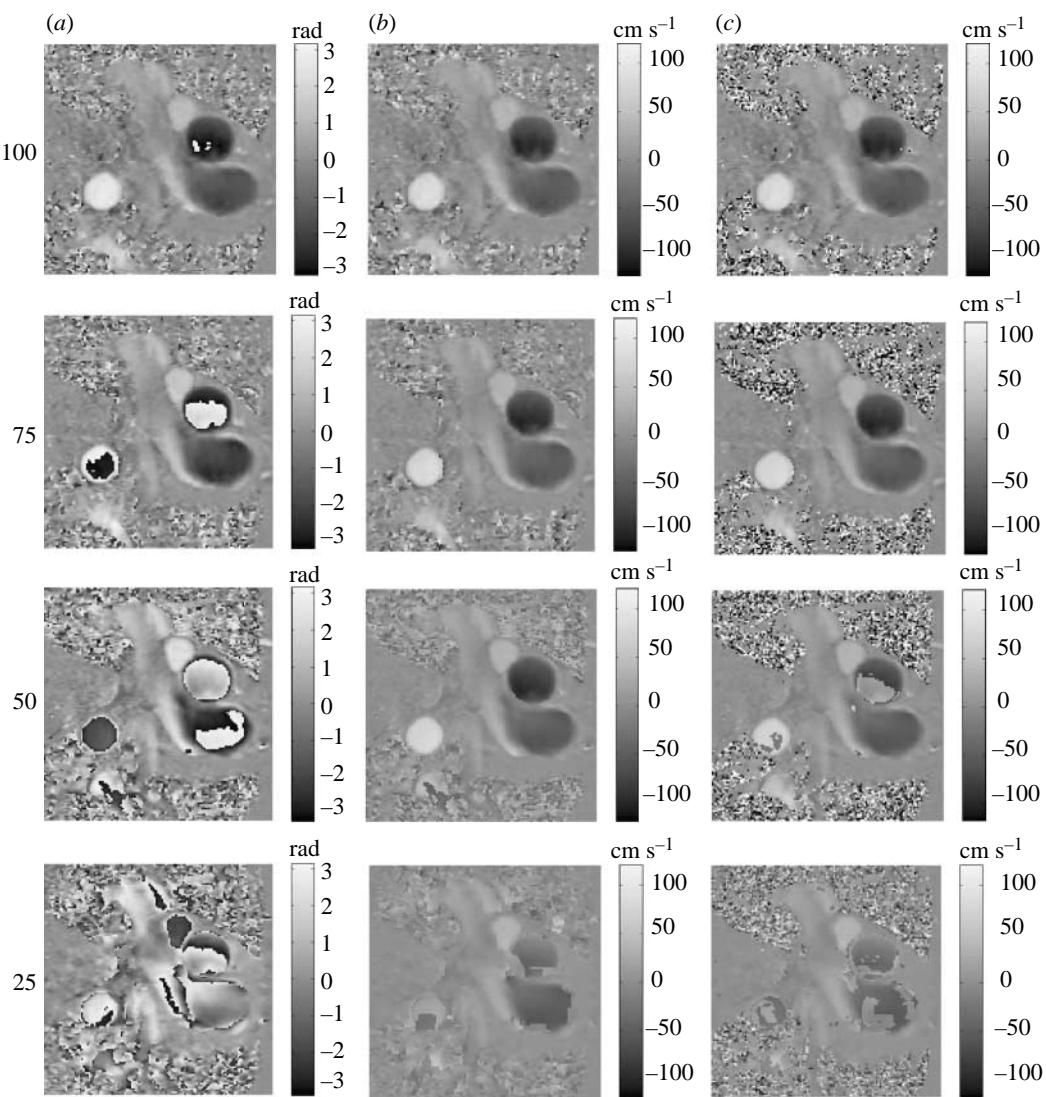


Figure 7. Frames corresponding to the highest blood velocity in the volumes obtained with V_{enc} values of 100, 75, 50 and 25 cm s^{-1} . (a) Wrapped phase. (b) Unwrapped phase with the three-dimensional noise immune phase unwrapping algorithm. (c) Unwrapped phase with the temporal phase unwrapping algorithm.

with this method due to temporal undersampling (figure 7), as will be explained in more detail in §4.

3.3. Velocity encoding unwrapping

The volume with V_{enc} value of 25 cm s^{-1} , which could not be unwrapped with the temporal or three-dimensional spatial unwrapping methods, was successfully unwrapped with the velocity encoding unwrapping approach using two V_{enc} values (figure 9a) and an exponential sequence of four V_{enc} values without (figure 9b) and with least-squares fitting along the velocity encoding direction (figure 9c).

4. DISCUSSION

The advantage of phase contrast images obtained with low velocity encoding is the improved velocity-to-noise ratio. This can be appreciated in the velocity surface inside the ascending aorta obtained with $V_{\text{enc}} = 150 \text{ cm s}^{-1}$ (figure 10a) and $V_{\text{enc}} = 25 \text{ cm s}^{-1}$ (figure 10b). Small velocity changes across the vessel, which are masked by the noise in figure 10a, become

apparent in figure 10b. This improvement is quantified in figure 11 by calculating the root mean square (r.m.s.) velocity fluctuations across the ascending aorta at each V_{enc} value at two points in the cardiac cycle, one where the velocity field is approximately zero and another where the mean velocity was -87 cm s^{-1} . A factor of approximately six times improvement (5.9 for the almost zero velocity field, 6.7 for the high velocity field) is obtained at the lowest V_{enc} compared to the highest of 150 cm s^{-1} , and a factor of approximately five times improvement (4.6 for the almost zero velocity field, 4.7 for the high velocity field) is obtained at the lowest V_{enc} compared to the 125 cm s^{-1} volume which is the lowest V_{enc} volume without aliasing. These values are in accordance with the theoretical linear relationship between velocity-to-noise ratio and V_{enc} derived by Lee *et al.* (1995). The fact that the r.m.s. values decreases with decreasing V_{enc} can be attributed to the noise rather than the signal dominating the r.m.s. calculations.

The wall shear rate measurement involves differentiating the velocity field, a numerical procedure which

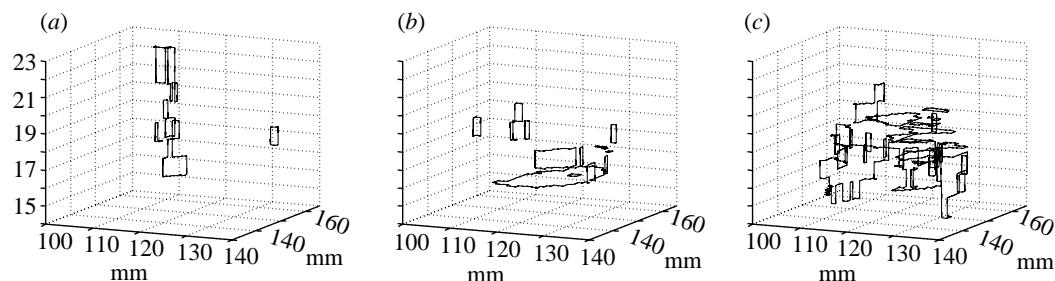


Figure 8. Phase singularity loops in a sub-volume corresponding to the ascending aorta, when the velocity encoding is (a) 75, (b) 50 and (c) 25 cm s^{-1} .

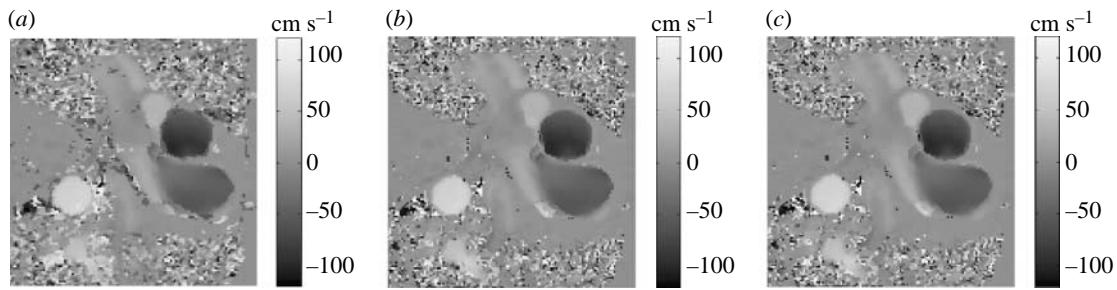


Figure 9. Frame corresponding to the highest blood velocity in the volume obtained with $V_{\text{enc}} = 25 \text{ cm s}^{-1}$. (a) Unwrapped phase with the velocity encoding method using only two velocity encoding values. (b) Unwrapped phase with the velocity encoding unwrapping method with an exponential sequence of four velocity encoding values. (c) Unwrapped phase with the velocity encoding unwrapping method with an exponential sequence of four velocity encoding values plus least-squares fitting.

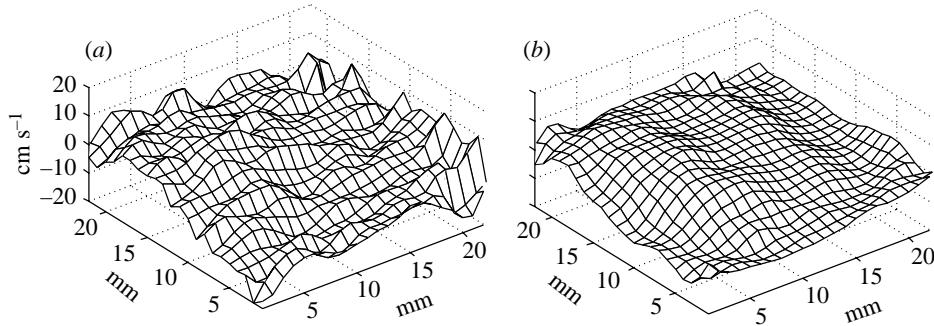


Figure 10. Velocity distribution inside the ascending aorta during diastole obtained with encoding velocities equal to (a) 150 and (b) 25 cm s^{-1} . The velocity distributions in (a) and (b) were obtained from the wrapped phase volumes.

amplifies noise and which would, therefore, be expected to result in better quality data through the use of low values of V_{enc} . This is indeed the case, as illustrated in figure 12a,b. These show the magnitude of the gradient vector in the radial direction with respect to the centre of the vessel. The radial derivative was calculated from the local gradients in the x - and y -directions obtained using a kernel of size equal to 3×3 pixels. The higher radial derivative values define the vessel's wall more sharply and with a smoother edge when V_{enc} of 50 cm s^{-1} (figure 12b) is used compared to the results obtained with V_{enc} of 150 cm s^{-1} (figure 12a). There are of course several other important issues that may affect the calculation of wall shear rate, such as partial volume effects, determination of vessel wall position, voxel size, intravoxel dephasing and blood turbulence. Although a discussion of these phenomena lies beyond the scope of the current paper, it is worth pointing out that the resulting errors will become increasingly important in relative terms as V_{enc} is reduced.

By contrast with the wall shear stress calculations, measurement of flow rate involves integrating the velocity field across the vessel, a procedure that suppresses the noise. Flow versus time graphs obtained from the studies with different encoding velocities are, therefore, seen to be comparable (figure 13).

The three unwrapping methods used gave comparable results for those sequences unwrapped successfully, with r.m.s. values of the difference being below 1 cm s^{-1} in the regions of the ascending and descending aorta. The temporal and spatial unwrapping methods can fail to successfully unwrap a volume when there is undersampling of the data, i.e. when the velocity encoding parameter is below a threshold value. Temporal undersampling is determined by the peak blood acceleration a_p . The magnitude of the velocity change between successive frames must be less than V_{enc} to avoid undersampling, so the Nyquist condition reads $|a_p| \cdot \Delta t \leq V_{\text{enc}}$, where Δt is the inter-slice sample time. Similarly, spatial undersampling depends

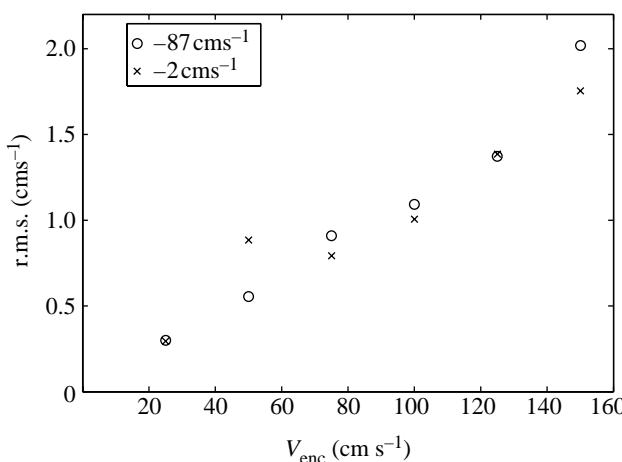


Figure 11. R.m.s. velocity fluctuations across the ascending aorta for different V_{enc} values at two points in the cardiac cycle: one where the mean velocity is approximately zero (-2 cm s^{-1}); and another where the mean velocity is -87 cm s^{-1} . The r.m.s. values for the $V_{\text{enc}}=25 \text{ cm s}^{-1}$ volume were obtained from the unwrapped phase using the velocity encoding unwrapping with a forward exponential sequence method.

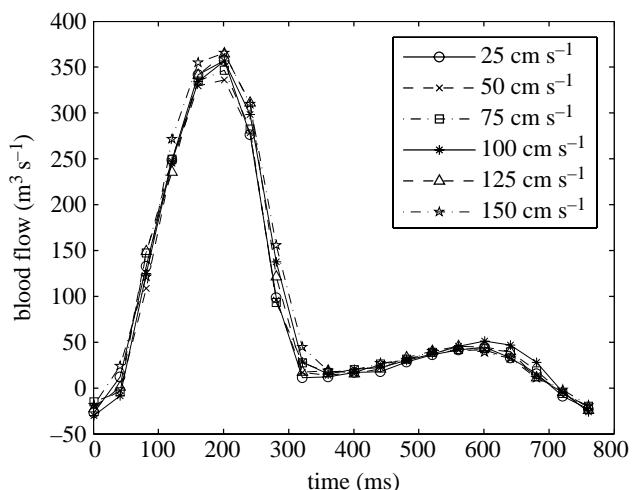


Figure 13. Flow versus time graphs of the descending aorta obtained from images with different encoding velocities. The velocity maps in images with $V_{\text{enc}}=100, 75$ and 50 cm s^{-1} were obtained from the volumes unwrapped with the three-dimensional spatial unwrapping algorithm, while the velocity map for $V_{\text{enc}}=25 \text{ cm s}^{-1}$ was obtained by using the four-dimensional unwrapping approach.

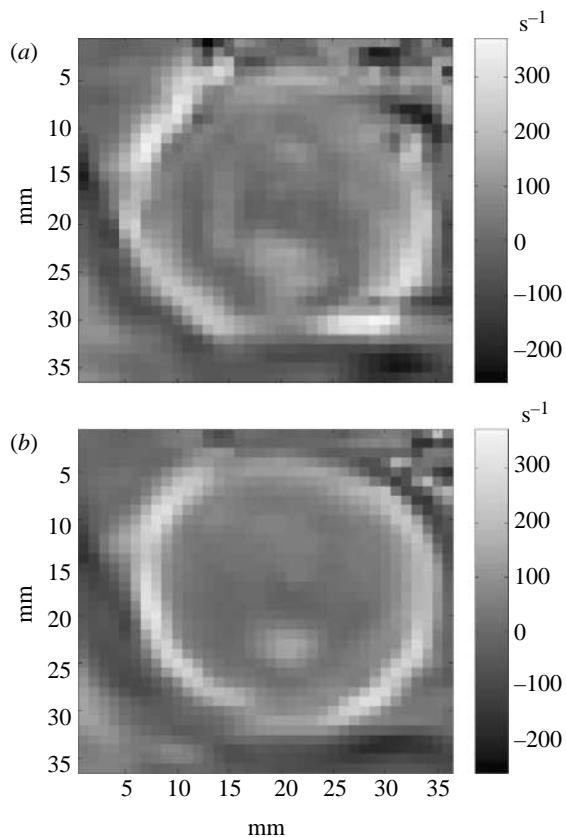


Figure 12. Wall shear strain rate inside the ascending aorta calculated from images with encoding velocities equal to (a) $V_{\text{enc}}=150 \text{ cm s}^{-1}$ and (b) $V_{\text{enc}}=50 \text{ cm s}^{-1}$ and peak velocity of 104 cm s^{-1} (phase unwrapped with three-dimensional spatial unwrapping algorithm).

on the peak velocity gradients, expected at the walls of the aortas, and the Nyquist condition implies that $|\partial v/\partial x|_p \cdot \Delta x \leq V_{\text{enc}}$ and $|\partial v/\partial y|_p \cdot \Delta y \leq V_{\text{enc}}$, where Δx and Δy are the inter-pixel distances in the x - and

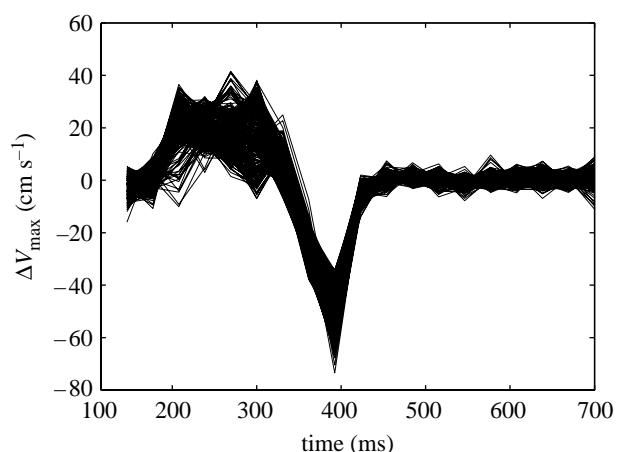


Figure 14. Velocity increments between consecutive frames for pixels in ascending aorta region ($V_{\text{enc}}=150 \text{ cm s}^{-1}$). The maximum absolute velocity increment is 71 cm s^{-1} .

y -directions, respectively, and subscript 'p' refers to the peak value.

The magnitude of the peak velocity change between successive frames was estimated from the $V_{\text{enc}}=150 \text{ cm s}^{-1}$ phase volume, which has no phase wraps (see figure 14), as 71 cm s^{-1} (equivalent to a peak acceleration of 17.8 m s^{-2}). This value is consistent with the results obtained for the temporal unwrapping method, which unwraps successfully for V_{enc} values down to 75 cm s^{-1} , but fails for velocity encoding values of 50 and 25 cm s^{-1} . When there is temporal undersampling, the PSLs tend to have more singularities oriented in directions perpendicular to the time axis (figure 8b,c). Figure 15a shows one large and several small horizontal loops due to temporal undersampling. Temporal unwrapping fails here because the unwrapping path crosses regions of more than 2π phase difference when traversing the horizontally oriented loops, causing unwrapping errors (figure 15b).

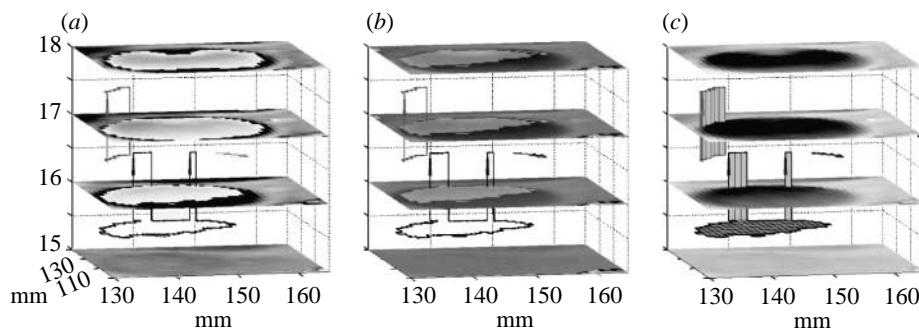


Figure 15. (a) Wrapped phase ($V_{\text{enc}}=50 \text{ cm s}^{-1}$) for four consecutive frames in the region of the ascending aorta, and the phase singularity loops that arise due to temporal and spatial undersampling. (b) Temporal unwrapping failure due to temporal undersampling (horizontal loops). (c) Three-dimensional spatial unwrapping is successful after placing branch surfaces in the phase singularity loops.

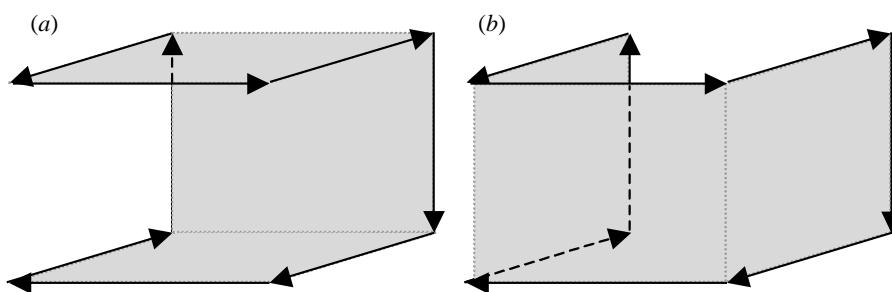


Figure 16. A C-shaped loop is ambiguous because it can support two different surfaces, (a) and (b), which block different unwrapping paths. If the incorrect surface is used, localized unwrapping errors result.

The three-dimensional spatial unwrapping method successfully unwraps this volume because branch surfaces are placed to prevent the unwrapping path from traversing the loop (figure 15c). Even if there is no temporal undersampling, when spatial undersampling is severe, the loops that are formed, though mainly oriented vertically, can have some horizontal components, which again will cause temporal unwrapping errors. It should be pointed out that since the phase variation with time is cyclic, simple error detection on the temporally unwrapped data could be carried out by checking that the phase offset between first and last frames lies in the range $[-\pi, \pi]$.

In the three-dimensional spatial unwrapping method, spatial and/or temporal undersampling will not cause any unwrapping errors as long as the branch surfaces are properly set, as in the case of figure 15c. However, ambiguous situations may arise, where more than one surface can be placed in the loop, blocking different unwrapping paths (figure 16). In this example, both surfaces have the same area, and from the loop geometry alone there is no reason to prefer one branch cut surface to the other. If the wrong surface is chosen, localized unwrapping errors result. The present algorithm unwraps successfully phase volumes with some degree of undersampling, as in $V_{\text{enc}}=50 \text{ cm s}^{-1}$, but not with severe spatial undersampling as in $V_{\text{enc}}=25 \text{ cm s}^{-1}$, where the loops become larger and more convoluted. Figure 17a shows one of the complex loops that appear in the sub-volume of the aorta in the $V_{\text{enc}}=25 \text{ cm s}^{-1}$ study. The unwrapped phase in this region presents errors because the wrong surface was set up for

this particular loop (figure 17b) and for a second similarly complex loop in the same area.

The velocity encoding unwrapping method successfully unwrapped the volume with $V_{\text{enc}}=25 \text{ cm s}^{-1}$, which could not be unwrapped by the other two methods. The use of four velocity encoding values has the disadvantage of requiring three extra acquisitions, but produces results with improved velocity-to-noise ratio (14% improvement) when compared to using only two velocity encoding values which requires only one extra acquisition (figure 18). The improvement is greater (26%) when the least-squares fit is performed on the exponential sequence. Another possibility to make the method faster is to acquire only one reference volume, as proposed by Lee *et al.* (1995). This would increase acquisition time by a factor of 2.5 if four velocity encoding values are used, or by 1.5 for the case of two values. Despite such potential improvements, the fact remains that the method involves some sacrifice of temporal resolution.

5. CONCLUSIONS

CINE phase contrast MRI obtained with velocity encoding values smaller than the maximum blood velocity results in wrapped phase volumes. Even though the velocity-to-noise ratio increases as the velocity encoding decreases, low V_{enc} values are generally not used in order to avoid phase wrapping. This limits the dynamic range of the images and results in a lower velocity-to-noise ratio. This study compared the performance of four different unwrapping strategies on such phase volumes over a range of V_{enc} values.

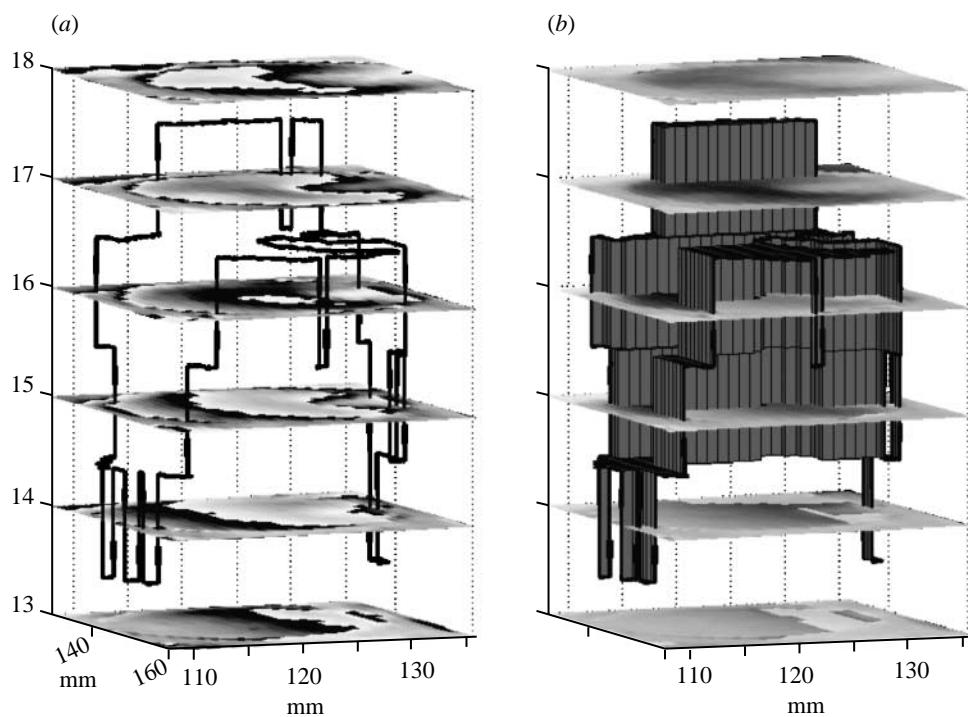


Figure 17. (a) Wrapped phase of six consecutive frames of the volume with $V_{\text{enc}} = 25 \text{ cm s}^{-1}$ in the region of the ascending aorta, and a complex ambiguous loop due to undersampling. (b) Unwrapped phase with the three-dimensional noise immune phase unwrapping algorithm. Unwrapping errors occur near the loop due to the incorrect selection of the two possible branch surfaces.

The simplest method based on unwrapping along the time axis worked well, provided that no temporal undersampling occurred. The theoretical criterion based on the peak acceleration, $|a_p| \cdot \Delta t \leq V_{\text{enc}}$, was found to be consistent with the observed onset of failure of this technique. A path-independent three-dimensional algorithm based on identification of PSLs, and modified to take account of loops traversing the entire time axis of the volume, gives significantly better performance in the presence of moderate amounts of both spatial and temporal undersampling. A factor of $2.5 \times$ improvement in sensitivity was achieved compared with the no-wrapping acquisitions. This algorithm is, therefore, recommended as the best option when only a single phase volume can be acquired.

Attempts to increase the sensitivity parameter s (defined as the ratio between maximum and minimum velocity encoding values) caused rapid increases in the size and density of the PSLs, resulting in localized unwrapping errors. The process of setting up branch surfaces can become ambiguous as the loop size increases. Although further research may improve the performance of the algorithm on severely undersampled data, it seems likely that the only reliable way to increase s further is to increase the sampling rate along both spatial axes by a corresponding amount. This will result in a temporal resolution that scales with s^2 . To prevent temporal undersampling, however, one requires a temporal resolution scaling as $1/s$, and it is the disparity between these two scaling laws that sets a practical limit on the maximum value of s that can be achieved.

Two velocity encoding unwrapping approaches were also studied: one using two velocity encoding values; and a second using an exponential sequence of four

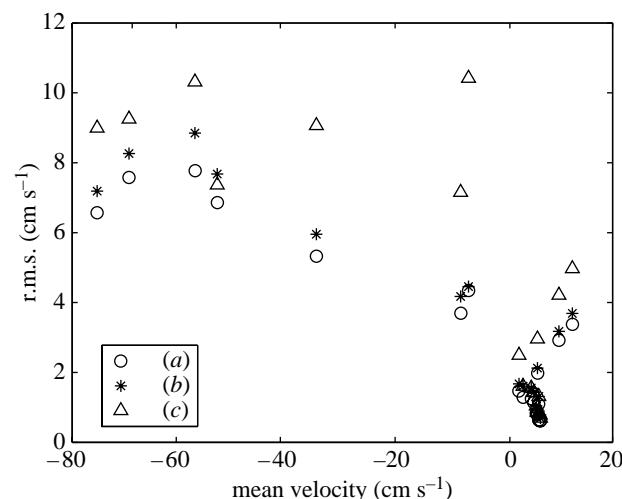


Figure 18. R.m.s. velocity fluctuations inside the ascending aorta at each temporal frame for the $V_{\text{enc}} = 25 \text{ cm s}^{-1}$ volume unwrapped along the velocity encoding dimension using an exponential sequence of four V_{enc} values with (a) and without (b) least-squares fitting along the g -axis and (c) two V_{enc} values.

velocity encoding values. The latter has the advantage of increased unwrapping success rate compared to the former, and produced results with increased velocity-to-noise ratio, though it has the disadvantage of longer acquisition time. However, the temporal resolution with this approach scales as $[\log_2(s) + 1]$ —a much weaker function of s than for the three-dimensional algorithm—and since neither temporal nor spatial undersampling directly affects the unwrapping reliability, significantly higher s values may ultimately be achievable. A factor of five times improvement in velocity-to-noise ratio was

demonstrated, and no global failure of the algorithm was observed on any of the datasets studied. The ultimate limit with this approach will come when the phase gradient within a voxel becomes too high for accurate phase measurements to be made.

Finally, while further tests need to be carried out to evaluate the performance of the algorithms on a larger number of datasets, we believe the results achieved so far demonstrate considerable promise for a possible eventual clinical implementation.

We are pleased to acknowledge financial support from the Engineering and Physical Sciences Research Council (EPSRC) through research grant GR/R88489/01. J.M.H. is also grateful to the Royal Society and Wolfson Foundation for a Royal Society–Wolfson Research Merit Award.

APPENDIX A

This appendix presents detailed technical information on how the PC-MRI data were obtained and preprocessed.

A retrospectively gated, segmented CINE phase contrast sequence using two views-per-segment (two k -space lines were acquired per repetition time (TR) for each velocity encoding) was used. Other imaging parameters were: field-of-view 280×280 mm; slice thickness 5.0 mm; imaging matrix 256×128 with 3.0 NEX; receiver bandwidth 31.3 kHz, flip angle 30°. The TE/TR (TE: echo time) ranged from 3.4 to 4.2/6.8 to 7.6 ms, depending upon the velocity encoding. The study was conducted in compliance with the regulations of the local research ethics committee.

Twenty temporal phases were reconstructed with a temporal resolution of approximately 40 ms. The volumes were registered both spatially and temporally as follows. The flow versus time curves of the different volumes were compared to the flow versus time curve, corresponding to the highest sensitivity volume. It was necessary to shift the frames of the different volumes by 1–4 frames to correct for temporal misalignment. Each magnitude volume was then spatially registered to the highest sensitivity volume, using SPM's (Statistical Parametric Mapping, <http://www.fil.ion.ucl.ac.uk/spm/>) coregistration tool, which uses a rigid model, with mutual information as a cost function, and a third degree spline as re-slice interpolation function. The phase volumes were coregistered using the affine transformation matrices obtained from the coregistered magnitude volumes.

APPENDIX B

In this appendix, we demonstrate that PSLs cannot terminate within the phase volume. It is more general than the proof by Huntley (2001), in that it does not rely on the assumption that singular points in two-dimensional phase maps occur in closely spaced singularity pairs or dipoles. It is, therefore, equally applicable to situations in which global undersampling occurs, which in the case of two-dimensional phase maps can result in isolated singularities or ‘monopoles’.

Consider an elemental cube of $2 \times 2 \times 2$ voxels (see figure 19). The presence of a singularity s_i on a given

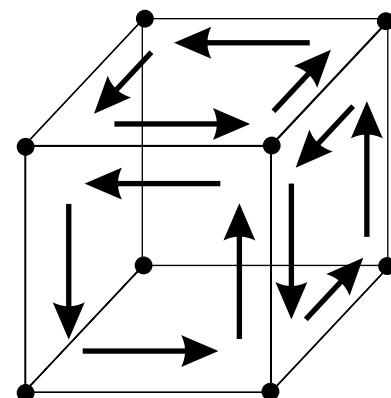


Figure 19. Unwrapping paths for calculation of the singularity distribution on three adjacent faces of a cube element.

face i of the cube is detected by summing the phase discontinuities that are crossed as one follows a closed path along the four edges of the cube face. If we label the four voxels on this face 1, ..., 4, ordered in the direction given by the right-hand screw rule with the thumb pointing along the outward normal, and with corresponding phase values ϕ_1, \dots, ϕ_4 , then

$$s_i = \Delta_{21} + \Delta_{32} + \Delta_{43} + \Delta_{14}, \quad (B1)$$

where

$$\Delta_{kl} = \text{NINT} \left(\frac{\phi_k - \phi_l}{2\pi} \right), \quad (B2)$$

s_i can take the value $+1, -1$, or 0 (Bone 1991). Note that the sign of s_i is reversed if it is calculated with the thumb pointing along the inward normal to the cube.

Now consider the sum of the singularities over the six faces of the cube, $S = \sum_{i=1}^6 s_i$. For every contribution Δ_{kl} to S from face i , there is a contribution Δ_{lk} to S from the face that shares the common edge between voxels k and l . In figure 14, this is illustrated for three faces by the opposing directions of the unwrapping arrows on their adjacent edges. However, by definition (equation (B 2)), $\Delta_{kl} = -\Delta_{lk}$ and, therefore, S must be identically equal to zero.

The fact that S is zero means that a singularity s_i on a given face i of the cube must be balanced by a singularity of opposite sign $s_j = -s_i$ on one of the other five faces of the cube, denoted j . Face j is shared with a neighbouring cube of $2 \times 2 \times 2$ voxels, and as calculated on this neighbour's face (for which the outward normal is in the opposite direction to face j on the first cube), the singularity, therefore, takes the value s_i from face i of the first cube.

The argument can then be repeated as above for the second cube, and then a third and so on, the process terminating only when we reach either face i of the first cube again, or the boundary of the phase volume. The list of cubes built up in this way constitutes the path of the PSL. Each cube in the list contains a face with a -1 singularity and a second face with a $+1$ singularity, both measured in the outward normal direction. The -1 singularity corresponds to a $+1$ singularity measured in the inward normal direction and can be thought of as the face through which the loop ‘enters’ the cube, whereas the $+1$ singularity occurs on the face

through which the loop ‘leaves’ the cube. The arrows shown on the figures in this paper indicate the loop direction as defined in this way.

APPENDIX C

This appendix derives expressions for unwrapping success rates of the two one-dimensional phase unwrapping strategies presented in §2.4.

C.1. Strategy one: two values $g=1$ and $g=s$ are used

An expression for the overall unwrapping success rate R or probability of correct unwrapping of this method was derived by Huntley & Saldner (1997a), by considering that the unwrapping will be successful if the random variable

$$X = s\phi_u(1) - \phi_u(s) \quad (\text{C } 1)$$

falls in the range $(-\pi, \pi]$. The expected values of $s\phi_u(1)$ and $\phi_u(s)$ are identical, and if $\phi_u(1)$ and $\phi_u(s)$ are Gaussian random variables with standard deviation σ , then X is also a Gaussian random variable with zero mean and standard deviation equal to $(s^2 + 1)^{1/2}\sigma$. Therefore, the overall success rate R is

$$R = \text{erf}\left(\frac{\pi}{\sigma\sqrt{2(s^2 + 1)}}\right), \quad (\text{C } 2)$$

where erf is the error function

$$\text{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-u^2} du. \quad (\text{C } 3)$$

C.2. Strategy two: a growing exponential sequence $g=1,2,4,8$ (i.e. $s=8$) is used

The overall unwrapping success rate R (Huntley & Saldner 1997a) can be derived by considering that the unwrapping will only be successful if the random variable

$$Y = \Delta\phi_u(2g, g) - \Delta\phi_u(g, 0) = \phi_u(2g) - 2\phi_u(g) \quad (\text{C } 4)$$

falls in the range $(-\pi, \pi]$ for each of the unwrapping steps after the first. Y is a Gaussian random variable with mean equal to zero and standard deviation equal to $\sqrt{5}\sigma$, where σ is the standard deviation of the phase. The overall success rate R is

$$R = \frac{1}{2} \left[\text{erf}\left(\frac{\pi(1 - v/V_{\text{enc}}^{\max})}{\sqrt{2}\sigma}\right) + \text{erf}\left(\frac{\pi(1 + v/V_{\text{enc}}^{\max})}{\sqrt{2}\sigma}\right) \right] \times \left[\text{erf}\left(\frac{\pi}{\sqrt{10}\sigma}\right) \right]^{\log_2 s}, \quad (\text{C } 5)$$

where the first term on the right-hand side is a measure of the probability of failure of the first unwrapping step (Huntley & Saldner 1997a). From equation (C 5), it is clear that the two requirements to obtain high success rates are that the blood speed $|v| < V_{\text{enc}}^{\max}$, i.e. that the

phase volume with highest velocity encoding, is free from wrapping, and that the noise level is low.

REFERENCES

Bone, D. J. 1991 Fourier fringe analysis: the two-dimensional phase unwrapping problem. *Appl. Opt.* **30**, 3627–3632.

Buckland, J. R., Huntley, J. M. & Turner, S. R. E. 1995 Unwrapping noisy phase maps by use of a minimum-cost-matching algorithm. *Appl. Opt.* **34**, 5100–5108.

Chavez, S., Xiang, Q. S. & An, L. 2002 Understanding phase maps in MRI: a new cutline phase unwrapping method. *IEEE Trans. Med. Imaging* **21**, 966–977. (doi:10.1109/TMI.2002.803106)

Cusack, R. & Papadakis, N. 2002 New robust 3D phase unwrapping algorithms: application to magnetic field mapping and undistorting echoplanar images. *NeuroImage* **16**, 754–764. (doi:10.1006/nimg.2002.1092)

Ghiglia, D. C. & Pritt, M. D. 1998 *Two-dimensional phase unwrapping: theory, algorithms, and software*, 1st edn, pp. 100–278. New York: Wiley.

Herment, A., Mousseaux, E., Jolivet, O., DeCesare, A., Frouin, F., Todd-Pokropek, A. & Bittoun, J. 2000 Improved estimation of velocity and flow rate using regularized three-point phase-contrast velocimetry. *Magn. Reson. Med.* **44**, 122–128. (doi:10.1002/1522-2594(200007)44:1<122::AID-MRM18>3.0.CO;2-C)

Huntley, J. M. 1989 Noise-immune phase unwrapping algorithm. *Appl. Opt.* **28**, 3268–3270.

Huntley, J. M. 1997 Random phase measurement errors in digital speckle pattern interferometry. *Opt. Lasers Eng.* **26**, 121–150. (doi:10.1016/0143-8166(95)00109-3)

Huntley, J. M. 2001 Three-dimensional noise-immune phase unwrapping algorithm. *Appl. Opt.* **40**, 3901–3908.

Huntley, J. M. & Saldner, H. O. 1993 Temporal phase unwrapping algorithm for automated interferogram analysis. *Appl. Opt.* **32**, 3047–3052.

Huntley, J. M. & Saldner, H. O. 1997a Shape measurement by temporal phase unwrapping: comparison of unwrapping algorithms. *Meas. Sci. Technol.* **8**, 986–992. (doi:10.1088/0957-0233/8/9/005)

Huntley, J. M. & Saldner, H. O. 1997b Error reduction methods for shape measurement by temporal phase unwrapping. *J. Opt. Soc. Am. A* **14**, 3188–3196.

Jenkinson, M. 2003 Fast, automated, N -dimensional phase-unwrapping algorithm. *Magn. Reson. Med.* **49**, 193–197. (doi:10.1002/mrm.10354)

Lee, A. T., Pike, G. B. & Pelc, N. J. 1995 Three-point phase contrast velocity measurements with increased velocity-to-noise ratio. *Magn. Reson. Med.* **33**, 122–126.

Marklund, O., Huntley, J. M. & Cusack, R. 2005 Robust unwrapping algorithm for 3-D phase volumes of arbitrary shape containing knotted phase singularity loops. Research report, Luleå University of Technology, Sweden, Nr. 2005-16, ISSN 1402-1528, ISRN LTU-FR-05/16-SE.

McRobbie, D. W., Moore, E. A., Graves, M. J. & Prince, M. R. 2003 *MRI from picture to proton.*, pp. 264–269. Cambridge, UK: Cambridge University Press.

Xiang, Q. S. 1995 Temporal phase unwrapping for cine velocity imaging. *J. Magn. Reson. Imaging* **5**, 529–534.

Xiang, Q. S. 2001 Anti-aliasing acquisition for velocity encoded cine. *Proc. 9th ISMRM Annual Meeting, Glasgow*, p. 1978.

Yang, G. Z., Burger, P., Kilner, P. J., Karwatowski, S. P. & Firmin, D. N. 1996 Dynamic range extension of cine velocity measurements using motion-registered spatiotemporal phase unwrapping. *J. Magn. Reson. Imaging* **6**, 495–502.