

X-ray all-sky monitoring and transient detection using a coded sphere telescope

A.J. Bird and M.R. Merrifield

Department of Physics, University of Southampton, Highfield, Southampton, S017 1BJ, UK

Received August 22; accepted October 20, 1995

Abstract. — We present a design for a hard X-ray all-sky monitor which consists of a spherical structure whose surface is partially covered by circular units which serve the dual purpose of detector and mask elements, forming a coded aperture telescope with full 4π steradian coverage. We demonstrate that a simple practical implementation of this geometry results in an instrument which is capable of measuring the fluxes in the 10–20 keV bandpass of the hundred brightest sources in the sky at least every two days. We also show how this instrument can be used to image burst and transient sources; obtain their positions with arcminute accuracy; and follow their light curves throughout their lifetimes.

Key words: instrumentation: detectors — techniques: image processing — telescopes — X-rays: general — X-rays: bursts

1. Introduction

Astronomy at hard X-ray energies (10–250 keV) faces the basic difficulty that reflective and refractive optics cannot be used at such short wavelengths. However, coded mask telescopes provide a method for imaging the sky in this waveband: the images are reconstructed from the geometric shadows cast by opaque mask elements when detecting elements are illuminated by X-ray sources. This technique is now well established in practice (e.g. Skinner 1984), and it has opened up a new window on the universe.

Observations of cosmic hard X-ray sources provide us with direct information on some of the universe's most energetic phenomena. The short timescale variability seen in this emission underlines the compactness, and hence the intensity, of these sources. Long-term monitoring in this energy band would provide information on source variability over a wide range of timescales, from which we would be able to infer much about the structure of the emitting region. Unfortunately, most coded mask telescopes to-date have had a relatively small field of view, and so continuous monitoring of many sources over long periods has not been possible. Even wide-field instruments, such as those on the SAX satellite (Scarsi 1993), cover only a small fraction of the sky at any given time. Much new science would be possible if we were able to monitor sources all over the sky simultaneously.

Even more extreme examples of source variability are provided by transients and bursters. These sources can

rise from undetectable flux levels to become the brightest sources in the sky with durations from milliseconds in the case of bursters up to years for the more long-lived transients. The BATSE experiment on CGRO has revolutionized the study of these sources by providing data for many hundreds of bursts (Fishman et al. 1994). However, the geometric techniques used for deriving the location of the burst from this instrument are fundamentally limited by systematic effects to an accuracy of ~ 4 degrees, which is rather too large for quiescent counterparts (if they exist) to be identified in other wavebands. The origin of gamma ray bursts remains a mystery.

In this paper, we present a proposal for a new type of hard X-ray instrument, the coded sphere telescope. This design combines coded mask techniques with the all-sky geometric source location techniques of BATSE, and so it can be used to monitor known sources anywhere in the sky, and to detect and localize transients to within a small error box. Section 2 describes the basic geometry of the telescope, and defines the design parameters for a practical implementation of the concept. In Sect. 3, we show one method by which this design can be used to monitor many sources simultaneously, and present the analysis of simulated observations to show the sensitivity of the method. Section 4 describes the use of the telescope for detecting transients and bursts, and presents simulations which show the accuracy with which these sources can be located.

2. The coded sphere telescope

2.1. Basic geometry

The geometry of the proposed telescope design is illustrated in Fig. 1. It consists of a transparent spherical structure on the surface of which are mounted circular pixel elements. These elements contain disk detectors which face inward toward the centre of the sphere, with their outside faces shielded against incident photons. Thus, the elements on one side of the telescope act as masks to the detectors on the other side of the sphere. The detection plane is very incomplete, but, as we shall see, it is still quite practical to form images with such an instrument.

Planar coded mask systems generally have their mask elements arranged in a uniformly redundant array (URA). This arrangement has the desirable property that artifact-free images can be reconstructed by simple correlation techniques (Gunson & Polychronopoulos 1976; Fenimore & Cannon 1978). If the detector plane is not completely filled, many of the advantages of a URA mask disappear, and correlation techniques produce images with significant artifacts (Byard & Ramsden 1994). For the proposed spherical geometry with a significantly incomplete detection “plane”, it is not immediately apparent whether analogues of the URA can be constructed. However, for an all-sky monitor, there are more powerful techniques for deriving the fluxes of sources than simple correlation (see Sect. 3), and so the absence of a URA mask pattern is not terribly restrictive. We therefore place the mask/detector pixels at random locations on the surface of the sphere by analogy with the random pin-hole camera (Dicke 1968).

If the sphere has a radius R_t , and the N_d detector elements have a radius R_d , then the filling factor of the sphere occupied by detectors is

$$g_d = \frac{N_d R_d^2}{4 R_t^2}, \quad (1)$$

and the effective sensitive area in any given direction is

$$A_d = \pi R_t^2 g_d (1 - g_d). \quad (2)$$

The throughput is thus maximized at 25% when $g_d = 0.5$.

The angular resolution of the telescope is given by the angle subtended by a mask element as seen from a detector element, and is simply

$$\delta\theta \sim R_d/R_t \text{ radians} \quad (3)$$

(full width at half maximum). Notice that the telescope geometry means that this resolution is independent of detector position: toward the edge of the instrument in any given direction, the detector and mask elements get closer together, but the projected size of the mask elements also decreases in such a way that the angular resolution remains the same (see Fig. 1b).

2.2. Implementation

As can be seen from Fig. 1, the basic structure of the coded sphere is very simple. A set of identical mask/detector pixels must be mounted at random locations on the surface of a spherical supporting structure. This structure could either be a solid shell of material which is transparent in the energy range of interest, or it could form a skeletal framework like a geodesic sphere.

The choice of materials for the mask/detector pixels depends on the operational energy range for the proposed telescope. For a moderately sized instrument, we must use small mask/detector elements to achieve reasonable spatial resolution (see Eq. (3)), and so we are constrained to use detectors which can be implemented as small pixels. Detectors which might be used in a coded sphere telescope include:

- i) *Silicon pin photodiodes*, which can be used as rugged low-cost X-ray detectors. Standard devices with depletion depths of up to 500 μm offer useful detection efficiency up to ~ 30 keV. A low energy threshold of ~ 5 keV is attainable with standard electronics at room temperature; cooling the detector and associated electronics to ~ -30 C reduces the detection threshold to ~ 1 keV. At these low energies, a simple passive lead or tungsten shield at the rear of each detector would be adequate to provide both background shielding and the necessary mask function. Care would have to be taken to provide a supporting structure which minimizes the amount of extra attenuating material along the path of incident photons
- ii) *Cadmium telluride* detectors would allow the telescope to operate at higher energies. This high-density room-temperature semiconductor would provide good detection efficiency over the range ~ 15 –200 keV with a detector thickness of ~ 2 mm. In this case, a more sophisticated lead-tin-copper graded shield would be needed to avoid fluorescence photons from the shield falling within the operational energy band of the detectors.
- iii) *Scintillator-photodiode combinations* would extend the sensitive range of the telescope to even higher energies. However, it would not be possible to keep such detectors thin and maintain useful detection efficiency. Further, passive shielding is no longer practical at these higher energies and active veto systems would be required. As a result, the size and mass of detectors operating at higher energies would undoubtedly increase rapidly.

2.3. Case study: The BASIS telescope

In order to establish the viability of this basic design for an astronomical telescope, we have to demonstrate that it can be implemented in a form which meets useful observational goals.

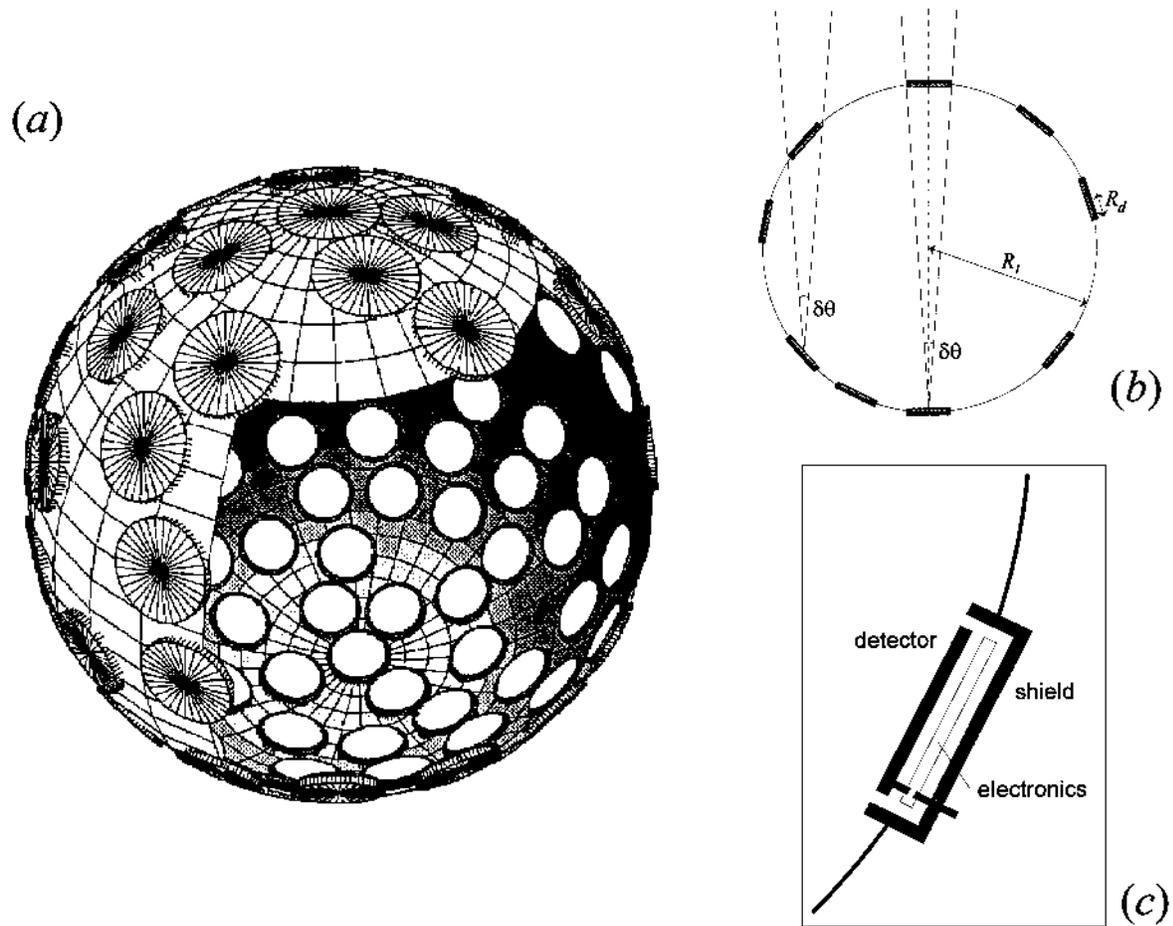


Fig. 1. Geometry of a coded sphere telescope. a) cutaway view illustrating inward-facing detector elements. b) cross-sectional view. c) detailed view of a detector/mask element

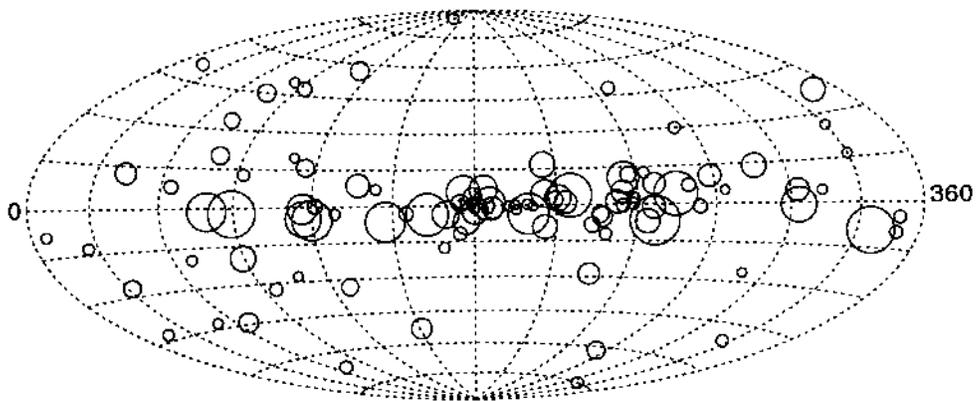


Fig. 2. Aitoff projection on the sky of a sample source distribution used in testing the coded sphere telescope. The size of each circle is proportional to the log of the source flux

As a simple model for a Burst and All-Sky Imaging Sphere (BASIS) telescope, we consider a sphere of radius $R_t = 30$ cm with detector/mask elements of radius $R_d = 1$ cm. Such a geometry gives a spatial resolution of 2 degrees, which is entirely adequate for monitoring the brightest few hundred sources in the sky. For detection elements in the model, we employ silicon photodiodes covered by a passive shield on their outer sides as described above. To achieve the nominal 50% coding of a sphere of this size requires 1800 detector/mask pixels distributed across the surface of the sphere, which produces a telescope with an effective area of $A_d \sim 700$ cm² in all directions. A preliminary mass estimate for this instrument yields a value of somewhat less than 100 kg; this mass is consistent with use as either a monitor instrument on a large mission or as a small mission stand-alone telescope.

3. All-sky monitoring

3.1. Flux estimation by maximum likelihood analysis

In an all-sky monitoring instrument, a large part of the observational programme will consist of continuous observations of a set of pre-determined sources. The positions of these sources will be known either from previous X-ray telescopes with high angular resolution, or by identification in other wavebands. We therefore do not wish to reduce the effective signal-to-noise ratio of our observations by leaving the positions of the sources as free parameters when analyzing observations from this instrument, and so we need an analysis technique which makes full use of the a priori positional information.

When observing a finite number of sources at known locations with a coded mask telescope, Skinner & Nottingham (1993) have shown how source fluxes can be estimated from the observed photon event data by maximum likelihood analysis. In the case of an all-sky monitoring instrument, the event rates are likely to be too high to be telemetered to the ground in their raw form, and will be integrated on-board over some finite observation time, t_{int} (perhaps in a number of energy bands). The observed quantities are then the integrated number of counts, F_i , in each of the N_d detectors, and the RMS uncertainty on each such quantity, σ_i , is just the Poisson noise on this number.

For such a set of observations, we can simplify the maximum likelihood analysis to a linear least-squares problem: the best estimates for the N_s source fluxes, f_j , $j = 1 \dots N_s$, and count rates due to the isotropic diffuse background, B , are obtained by minimizing

$$C^2 = \sum_{i=1}^{N_d} \frac{\left[F_i - \left(\sum_{j=1}^{N_s} f_j t_{\text{int}} R_{ij} + B t_{\text{int}} \right) \right]^2}{\sigma_i^2}, \quad (4)$$

with respect to these parameters. In this equation, R_{ij} is the count rate in detector pixel i due to a unit flux source

at location j (allowing for the shadowing due to mask elements). The unknown parameters occur only linearly inside the square brackets, and so the usual machinery for linear least squares analysis (e.g. Press et al. 1986) can be applied to the problem. If our model for the source distribution is correct, then C^2 should follow a χ^2 distribution with $N_d - (N_s + 1)$ degrees of freedom, so we can use this statistic to see if we are missing any significant sources of X-rays. We can also readily include other linear unknown terms which may occur such as detector noise, diffuse components associated with the galactic plane, etc.

A standard least-squares analysis also enables us to calculate the uncertainties in the derived fluxes, and the covariances between the different estimates. If, for example, two sources are barely resolved by the telescope, then their individual fluxes may be highly uncertain whereas the sum of their fluxes will still be well constrained; by calculating the full covariance matrix, we can explicitly calculate such coupled errors.

3.2. Simulated source distribution

Since each detector element in the coded sphere records flux from a wide range of directions, the telescope's ability to recover the brightnesses of astronomical sources depends on their distribution around the sky. We have therefore simulated source distributions which approximate the true distribution of hard X-ray sources in the sky. Of the N_s sources in this simulated distribution, 50% are assumed to come from galactic sources with a distribution on the sky,

$$\mu_{\text{gal}}(l, b) = \mu_0 e^{-b^2/[2(5 \text{ deg})^2]} e^{-l/(60 \text{ deg})}, \quad (5)$$

where $\{l, b\}$ are the usual galactic coordinate system. The remainder of the sources are assumed to be of extragalactic origin, and are randomly distributed about the sky. The source fluxes are simulated from a $\log(N)/\log(S)$ relation with a slope of -0.5 for the galactic sources (appropriate for a disk distribution), and -1.5 for the extragalactic sources (as expected for a uniform distribution of sources). For such power-law flux distributions, a lower limit cut-off, F_{min} must also be specified. Figure 2 shows a simulated distribution with $N_s = 100$ and $F_{\text{min}} = 10$ mCrab in the 10–20 keV bandpass. It appears very similar to the true distribution of the 100 brightest sources in the sky in this energy range (McHardy et al. 1981).

3.3. Results of simulations

We have used the above recipe for generating source distributions to test the viability of the simple BASIS telescope described in Sect. 2.3. The simulated observations were made in the bandpass from 10–20 keV, and the isotropic cosmic background flux in this bandpass was taken to be 1.3 mCrabs per square degree (Dean 1991), which corresponds to $B \sim 4$ cts s⁻¹ for this telescope geometry. Using the source distribution shown in Fig. 2, we simulated

a 48 hour integration with BASIS. The maximum likelihood fit to this observation returned a minimum value of $C^2 = 1750$, which is entirely consistent with its expected χ^2 distribution with 1699 degrees of freedom. Figure 3 shows the corresponding estimated fluxes and errors.

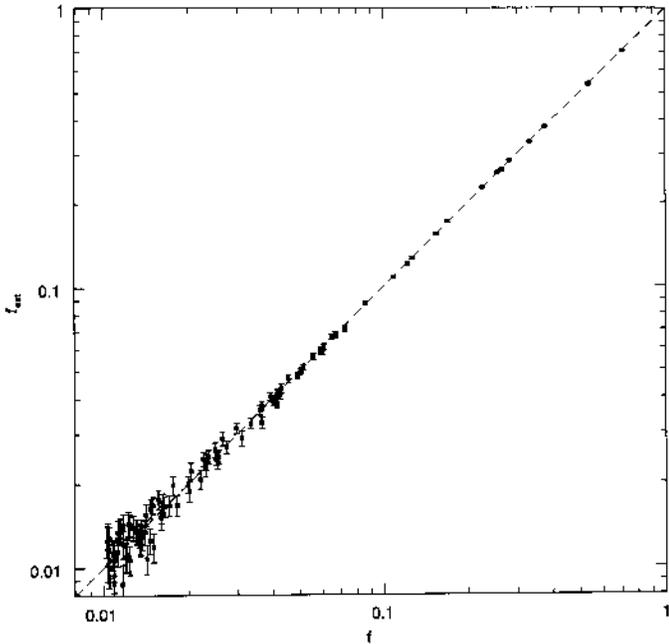


Fig. 3. Flux estimates versus true source fluxes for a 48 hour integration with the BASIS telescope of the source distribution in Fig. 2

Clearly, we obtain more than enough signal to monitor all of the 100 brightest sources in the sky every 48 hours. For the brighter sources, the integration time can be significantly reduced, and so variability in these sources can be measured on much shorter timescales by the same analysis. The properties of still fainter sources could readily be explored by adding them to the source list and analyzing data obtained over longer periods.

We have also attempted to make the simulation more realistic by adding more faint sources to the observed data without including them in the list of known sources for the flux estimation process. As expected, such experiments produce poorer χ^2 statistics, but they do not significantly degrade the best-fit flux estimates — most of the additional flux is ascribed by the maximum likelihood process to a higher effective background, B .

4. Burst and transient detection

The other major use for a coded sphere telescope is in the detection and localization of X-ray bursts and other transient phenomena. In an all-sky monitor, the simplest way to detect that a burst has occurred is by measuring the overall change in count rate in the instrument. If the

burst has a duration τ_b and a mean flux level f_b , then we can use this method to detect the increased count rate at a signal-to-noise ratio of

$$S/N = \frac{f_b A_d \tau_b}{\sqrt{B N_d \tau_b}}. \quad (6)$$

For a burst with a flux of ~ 0.5 photons $\text{cm}^{-2} \text{s}^{-1}$ in the 10–20 keV bandpass [of which there are maybe 50 per year (Lei et al 1995)], we would hence obtain a signal-to-noise ratio of $\sim 3(\tau_b/\text{second})^{1/2}$ from the BASIS telescope. Thus, we can use this method to detect the presence of an X-ray burst at this flux level as long as its duration exceeds a few seconds.

This method is clearly a very inefficient way of analyzing the available data, since it makes no use at all of the directional information in the detected photons. The crudest directional information in the coded sphere comes from noting that all the burst counts will only be detected in one side of the sphere. We can use this phenomenon to obtain an estimate of the direction to the burst: if detector i points in the direction given by the unit vector $\hat{\mathbf{d}}_i$, and it registers a total of n_i counts during the observed burst, then a simple estimate for the direction to the burst is given by

$$\hat{\mathbf{b}} = \frac{\sum_{i=1}^{N_d} n_i \hat{\mathbf{d}}_i}{\left| \sum_{i=1}^{N_d} n_i \hat{\mathbf{d}}_i \right|}. \quad (7)$$

This method of localization is analogous to that adopted for data from the BATSE detectors on CGRO.

Even this approach is not making full use of the data that we have available to us, since it does not include the constraints on the possible location of the burst imposed by the telescope mask coding. This extra information can most readily be utilized via the photon back projection method described by Fenimore (1987). In this method, a map of the sky is made by projecting each detected photon back toward the sky in all possible directions. If the path back to the sky does not encounter a mask element, then a count is added to the sky map at that location. This method will produce a large peak in the resulting sky map at the location of the source. It is computationally very expensive to map out large regions of the sky by this method, but we can use the estimated location returned by Eq. (7) to narrow down the area of sky that we have to search to a manageable region.

Figure 4 shows the image of an X-ray burst obtained by back projecting a simulated observation made using the BASIS telescope. This source was detected in the instrument count rate (Eq. (6)) at a signal-to-noise ratio of ~ 15 and Eq. (7) returned an estimated burst location within 4 degrees of the true source. Back projection then produced an image at the spatial resolution of the telescope (~ 2 degrees), allowing us to centroid the burst location to within a few arcminutes.

Longer-lived X-ray transient sources would initially be detected in a similar manner. Once their positions had

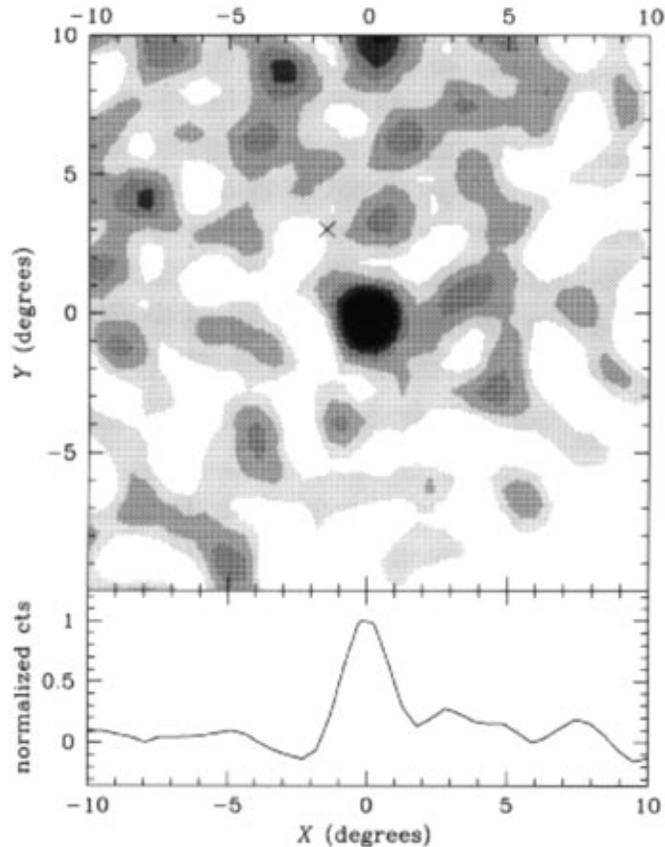


Fig. 4. Back projection map derived from simulated observations with the BASIS telescope of an X-ray burst with a flux of $f_b = 0.5 \text{ photons cm}^{-2} \text{ s}^{-1}$ over a duration of $\tau_b = 10$ seconds. The true location of the source is at the centre of the map, the cross shows the estimate of the burst location derived using Eq. (7), and the greyscale show the image derived by back projection. The lower panel shows a horizontal cut through the centre of the image

been localized by back projection imaging, they would be added to the list of sources to be monitored by the maximum likelihood analysis of Sect. 3, enabling their light curves to be measured. Once they have faded from view, they could be pruned from the list to keep the monitoring analysis manageable.

5. Discussion

In this paper, we have demonstrated the application of a “coded sphere” telescope for whole-sky imaging in the X-ray domain. Maximum likelihood analysis allows this telescope to monitor simultaneously the fluxes from sources all over the sky, and the combination of BATSE-like geometric location and back-projection imaging means that the telescope can find the positions of bursts and transient sources with unprecedented accuracy. The back-projection technique also allows us to use the telescope to map out any extended structures, and search for faint previously undetected sources.

The main drawback of the back-projection imaging technique is that it is a computationally intensive process. Although increasingly powerful computers mean that this limitation is not critical, more efficient imaging techniques are still desirable. It would therefore be very useful to establish the spherical analogue of the cross-correlation technique used in planar coded mask systems, since this method can make use of the computationally highly efficient Fast Fourier Transform. It would also be interesting to explore different mask patterns in the spherical geometry to see if it is possible to recover the improved sensitivity that the uniformly redundant array offers in planar systems.

A further improvement to the design introduced here would come from using pixellated detectors within each mask element. Such detectors would increase the mask-detector sampling ratio and hence improve the point source location accuracy of the imaging system. Pixellated semiconductor detector arrays have already been demonstrated (e.g. Perez & Olmos 1994), and so are very promising for this application.

Finally, we note that although the coded sphere has been designed to meet the need for an all-sky monitoring instrument in X-ray astronomy, the basic concept is equally applicable in other wavebands, and to terrestrial applications where omni-directional imaging is required. In order to explore the properties of coded spheres in more detail and assess their viability for a variety of applications, we now intend to build a number of small prototypes operating at both optical and X-ray wavelengths. The results from these experiments will be reported in future publications.

Acknowledgements. Dr. Merrifield is supported by a PPARC Advanced Fellowship (B/94/AF/1840). The authors would like to thank Mr. A. Gomm for help with preparation of the figures in this paper.

References

- Byard K., Ramsden D., 1994, Nucl. Instrum. Meth. A342, 600
- Dicke R.H., 1968, ApJ 153, L101
- Dean A.J., Lei F., Knight P.J., 1991, Space Sci. Rev. 57, 109
- Fenimore E.E., 1987, Appl. Opt. 26, 2760
- Fenimore E.E., Cannon T.M., 1978, Appl. Opt. 17, 337
- Fishman G.J., et al., 1994, ApJS 92, 229
- Gunson J., Polychronopoulos B., 1976, MNRAS 177, 485
- Lei F., Palmer M.J., Jupp I.D., Ramsden D., 1995, Exp. Astron. (in press)
- McHardy I.M., Lawrence A., Pye J.P., Pounds K., 1981, MNRAS 197, 893
- Perez J.M., Olmos P., 1994, Nucl. Instrum. Meth. A348, 536
- Press W.H., Flannery B.P., Teukolsky S.A., Vetterling W.T., 1986, Numerical Recipes. Cambridge University Press, Cambridge
- Scarsi L., 1993, A&AS 97, 371
- Skinner G.K., 1984, Nucl. Instrum. Meth. 221, 33
- Skinner G.K., Nottingham M.R., 1993, Nucl. Instrum. Meth. A333, 540