The POINT-AGAPE survey – II. An unrestricted search for microlensing events towards M31

V. Belokurov,[1] J. An,[1] N. W. Evans,[1]* P. Hewett,[1] P. Baillon,[2] S. Calchi Novati,[3]
B. J. Carr,[4] M. Crézé,[5,6] Y. Giraud-Héraud,[6] A. Gould,[7] Ph. Jetzer,[3] J. Kaplan,[6]
E. Kerins,[8] S. Paulin-Henriksson,[6] S. J. Smartt,[1,9] C. S. Stalin,[6] Y. Tsapras[4]
and M. J. Weston[4] (The POINT-AGAPE† collaboration)

[1]*Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA*
[2]*European Organization for Nuclear Research CERN, CH-1211 Genève 23, Switzerland*
[3]*Institut für Theoretische Physik, Universität Zürich, Winterthurerstrasse 190, CH-8057 Zürich, Switzerland*
[4]*Astronomy Unit, School of Mathematical Sciences, Queen Mary, University of London, Mile End Road, London E1 4NS*
[5]*Université Bretagne-Sud, Campus de Tohannic, BP 573, F-56017 Vannes Cedex, France*
[6]*Laboratoire de Physique Corpusculaire et Cosmologie, Collège de France, 11 Place Marcelin Berthelot, F-75231 Paris, France*
[7]*Department of Astronomy, Ohio State University, 140 West 18th Avenue, Columbus, OH 43210, USA*
[8]*Astrophysics Research Institute, Liverpool John Moores University, 12 Quays House, Egerton Wharf, Birkenhead CH41 1LD*
[9]*Department of Pure and Applied Physics, The Queen's University of Belfast, Belfast BT7 1NN*

ABSTRACT

An automated search is carried out for microlensing events using a catalogue of 44 554 variable superpixel light curves derived from our 3-yr monitoring programme of M31. Each step of our candidate selection is objective and reproducible by a computer. Our search is unrestricted, in the sense that it has no explicit time-scale cut. So, it must overcome the awkward problem of distinguishing long time-scale microlensing events from long-period stellar variables.

The basis of the selection algorithm is the fitting of the superpixel light curves to two different theoretical models, using variable star and blended microlensing templates. Only if microlensing is preferred is an event retained as a possible candidate. Further cuts are made with regard to: (i) sampling, (ii) goodness of fit of the peak to a Paczyński curve, (iii) consistency of the microlensing hypothesis with the absence of a resolved source, (iv) achromaticity, (v) position in the colour–magnitude diagram and (vi) signal-to-noise ratio. Our results are reported in terms of first-level candidates, which are the most trustworthy, and second-level candidates, which are possible microlensing events but have a lower signal-to-noise ratio and are more questionable. The pipeline leaves just three first-level candidates, all of which have very short full-width at half-maximum time-scales ($t_{1/2} < 5$ d) and three second-level candidates, which have time-scales of $t_{1/2} = 31$, 36 and 51 d. We also show 16 third-level light curves, as an illustration of the events that just fail the threshold for designation as microlensing candidates. They are almost certainly mainly variable stars.

Two of the three first-level candidates correspond to known events (PA 00-S3 and 00-S4) already reported by the POINT-AGAPE project. The remaining first-level candidate is new. This algorithm does not find short time-scale events that are contaminated with flux from nearby variable stars (such as PA 99-N1).

Key words: gravitational lensing – stars: variables: other – galaxies: individual: M31 – dark matter.

1 INTRODUCTION

Microlensing experiments towards M31 have now been in progress for a number of years and are beginning to report results. Ongoing surveys include the POINT-AGAPE (e.g. Aurière et al. 2001; Calchi

Novati et al. 2003; Paulin-Henriksson et al. 2003), the WeCAPP (Riffeser et al. 2001, 2003) and the MEGA (de Jong et al. 2004) collaborations.

The data set used by the POINT-AGAPE collaboration consists of three seasons (1999–2001) of imaging taken with the Wide Field Camera on the Isaac Newton Telescope (INT). In each season, data were taken in two passbands (*r*, together with one of *g* and *i*) for 1 h per night for roughly 60 nights during the 6 months that M31 is visible at low airmass. The field of view of the Wide Field Camera is 33×33 arcmin^2. There are two fields, located north and south of the centre of M31, as shown in fig. 1 of An et al. (2004b). The field locations are motivated by the suggestion of Crotts (1992) that the event rate to sources in the near and far discs is different. The lines of sight to the far disc as compared with the near disc are longer and pass through more of the dark halo. So, if the halo dark matter is in compact form, this should manifest itself in a greater microlensing rate towards sources in the far disc. We have already found a number of interesting individual microlensing events towards M31 (e.g. Aurière et al. 2001; Paulin-Henriksson et al. 2002, 2003; An et al. 2004a), carried out a survey for classical novae (Darnley et al. 2004) and reported the locations, periods and brightness of ∼35 000 variable stars (An et al. 2004b). The variable stars also show a substantial asymmetry between the near and far disc caused by the effects of extinction and dust. So, An et al. (2004b) proposed a new idea – namely, that it is better to look for an east–west asymmetry in microlensing events rather than a near–far asymmetry. The variable star distributions are east–west symmetric to a good approximation.

In the context of the POINT-AGAPE experiment, there are three relevant goals for a survey of microlensing events. First, one might want to identify any genuine microlensing events (without regard to the selection criteria). This is a worthy goal, as the microlensing events are of interest in their own right (e.g. Aurière et al. 2001; An et al. 2004a). Secondly, one might want to identify candidate microlensing events such that the detection probability is spatially homogeneous. This would enable any asymmetry present to be empirically constrained. Thirdly, one might want to identify candidate microlensing events such that the probability that a specified microlensing event is detected can ultimately be computed. Achieving the third aim involves substantially more effort than the first aim. It may also carry the penalty that some genuine microlensing events are omitted from the final list of candidates.

Although the simplest simulations of unresolved M31 microlensing experiments predict that the typical full-width at half-maximum time-scale ($t_{1/2}$) is quite short, it is desirable not to restrict ourselves to selecting short time-scale candidates a priori. However, the optimum choice of candidate selection algorithm may be different, according to whether the sought-after events have long or short time-scales. A serious difficulty in identifying long time-scale events is contamination of the sample from long-period variable stars (e.g. Miras), which – given the sampling – may not repeat during the 3-yr baseline of the experiment. For short time-scale events, confusion with stellar variables is less problematic. Short-period variables (e.g. Cepheids) show multiple bumps in their light curves over the 3-yr baseline. For this reason, it makes sense to devise different algorithms for candidate selection for long and short time-scale events.

Paulin-Henriksson et al. (2003) have reported four candidates in a survey for high signal-to-noise (S/N) ratio, short time-scale microlensing events in the first 2 yr of INT data. Specifically, they restricted attention to events with a full-width at half-maximum time-scale $t_{1/2} < 25$ d and with a flux variation Δf in the *R* band that exceeds that of an $R = 21$ mag star [i.e. $R(\Delta f) < 21$]. More

recently, Paulin-Henriksson et al. (2004) and Calchi Novati et al. (in preparation) have extended their search to the full 3-yr data set. The advantage of the time-scale and magnitude cuts is that the microlensing nature of bright, short events is usually unambiguous. The disadvantage is that the cuts restrict attention to ∼15 per cent of all the variable light curves in the POINT-AGAPE data base.

The purpose of this paper is to complement the short time-scale searches of Paulin-Henriksson et al. (2003, 2004) and Calchi Novati et al. (in preparation) by devising a candidate selection algorithm that is not limited to short time-scale events. We remove all restrictions on full-width at half-maximum time-scale and search for fainter flux variations in the entire catalogue [$19 < R(\Delta f) < 24$]. We pay particular attention to the difficulty of distinguishing long time-scale microlensing events from long-period Miras. The selection algorithms of Paulin-Henriksson et al. (2003, 2004) and this paper are entirely independent, although both the imaging data and the constructed light curves are the same. In both Paulin-Henriksson et al. (2003, 2004) and this paper, the aim is to carry out a fully automated survey for which the efficiency can ultimately be calculated.

This paper begins with a description of the catalogue of superpixel light curves constructed from our monitoring programme of M31 during 1999–2001. Then, the selection procedure for microlensing candidates is described in Section 3. The M31 source stars are generally not resolved and so the selection algorithm is different from those employed in classical microlensing experiments such as MACHO (Alcock et al. 1997) and EROS (Aubourg et al. 1995). The survivors of the seven cuts of the selection procedure are a set of three first-level candidates and three second-level candidates. The first-level candidates are the most trustworthy and, notably, are all short. Their light curves, locations and microlensing parameters are presented in Section 4. The three second-level candidates form a looser selection of possible, if questionable, events. Also extracted are a set of 16 third-level light curves that lie just below our threshold for designation as microlensing candidates. They are discussed in Section 5.

2 THE CONSTRUCTION OF THE CATALOGUE

The *r*-band light curves are available for all three seasons. The *g*-band monitoring was largely discontinued after the first season (1999) and replaced with *i*-band monitoring during 2000 and 2001. There are some, albeit few, *i*-band data for 1999. Full details of the data acquisition and analysis can be found in An et al. (2004b). As in that paper and in Paulin-Henriksson et al. (2003, 2004), the superpixel method is used to process the data and construct the light curves. The basic idea is to smooth the data with a median filter to estimate the background, to apply an empirical correction for seeing variations and hence to build light curves for 7×7 superpixels. Henceforth, we deal exclusively with the set of derived superpixel light curves rather than the images themselves. We select the variable superpixel light curves by requiring that they have at least three consecutive data points at least 3σ above the baseline in the *r* band. This gives us a raw catalogue of 97 280 *r*-band selected variable light curves.

The typical seeing (full-width at half-maximum) for the INT data exceeds 3 pixel and the range of seeing present in even the processed data is a full factor of 2. The superpixel method involves a block-averaging scheme with a scalelength of 7 pixel (2.3 arcsec). One of the drawbacks of the method is that it does not work when a pixel is within approximately the smoothing length of either bad pixels or bright resolved stars. Therefore, the superpixel light curves within ∼ 6 pixel of resolved objects are full of 'variations' due to the

Table 1. The fractions of each field and CCD removed by the masking of the resolved stars on the reference frame. The field and CCD placings are shown in fig. 1 of An et al. (2004b).

Field	CCD	Fraction	Field	CCD	Fraction
Northern	1	0.0987	Southern	1	0.0611
Northern	2	0.0766	Southern	2	0.0772
Northern	3	0.0438	Southern	3	0.1582
Northern	4	0.0702	Southern	4	0.0598

changing seeing or variability of the nearby resolved object. Thus, the statistical properties of the light curves from within this portion of the CCD are very different from those that are well separated from resolved objects.

An et al. (2004b) removed such spurious light curves by utilizing a spectral analysis based on Lomb's periodogram. Here we adopt a different procedure. This is partly because we need to calculate the actual area of M31 covered by our survey for the purposes of estimating the efficiency, and partly because short-$t_{1/2}$ microlensing candidates fail to pass the spectral analysis cut of An et al. (2004b). We construct a mask of the known CCD defects, together with regions around all resolved stars detected in the reference frame. We note that masking the bad pixels is intrinsically different from masking the areas around the resolved stars, as the former is fixed with respect to the CCD, while the latter is fixed with respect to the sky. Whilst the stability of the pointing of the INT is good, the pointing none the less does vary slightly from frame to frame. The masking is therefore applied in two stages.

First, the areas within 6 pixel of any star brighter than $R = 20.5$, together with annuli between 1.5 and 3 pixel around stars fainter than $R = 20.5$ are incorporated into the mask. We chose these numbers after examining fig. 7 of An et al. (2004b), which shows a surfeit of variable light curves at these pixel distances from resolved stars. Table 1 shows the fraction of area lost by masking for each field and CCD combination. After masking, 44 554 variable superpixel light curves remain. Over half the original light curves have been removed, though they occur over <10 per cent of the total surveyed area. This illustrates the fact that the statistical properties of the removed light curves are very different from those that remain. This is why it is beneficial to have excised them from the catalogue, even though – as we will see – some bona fide microlensing events have been lost through masking. In principle, this difficulty can be significantly reduced by techniques such as difference image analysis, albeit at some computational cost.

Secondly, the registration of each frame is known, so the CCD coordinates corresponding to any superpixel at that epoch can be calculated. If these lie on a known CCD defect, the data point is eliminated from the light curve. The final catalogue contains 44 554 cleaned variable light curves, in which all data points lying on bad pixels have been removed. This is the catalogue through which we search for microlensing events.

3 THE MICROLENSING SELECTION CRITERIA

3.1 Overall strategy

The selection criteria used to identify microlensing light curves in experiments such as MACHO require just one significant bump in the light curve, generally by insisting on goodness of fit to the standard Paczyński curve for a point source microlensed by a point lens (e.g. Alcock et al. 1997). In the POINT-AGAPE experiment, pixel light curves may receive contributions from more than one

variable source located in the same superpixel, and so the variety of light curves is very rich. The development of selection criteria for identifying microlensing events from a data set of pixel light curves involves different challenges from those presented by conventional microlensing experiments in which the sources are resolved.

In this paper, the focus is on carrying out a search that has no explicit restrictions on time-scale. This necessitates the devising of algorithms to distinguish long time-scale microlensing events from contaminating long-period variable stars. The goodness of fit to standard microlensing light curves alone cannot resolve this problem. The selection algorithm begins by comparing light curves to two different theoretical models – namely, a simple variable star template and a standard Paczyński microlensing curve. Only if the microlensing fit is better than the alternative is a light curve retained as a possible candidate. Although we do not pursue it here, this line of reasoning is capable of further elaboration; for example, a very general algorithm for identifying microlensing events in pixellated data might begin by comparing goodness of fit to a number of competing models (e.g. variable stars, novae, supernovae, microlensing, microlensing with a bumpy baseline caused by variable contaminants).

Cuts based on goodness of fit and the absence of multiple bumps are never enough to eliminate variable stars. Worryingly, for the subject of microlensing, there are light curves that are extremely well fitted by blended microlensing curves in two or more passbands, but that are almost certainly Miras or other long-period variables. Achromaticity tests are a conventional way of distinguishing stellar variability from microlensing. We find that achromaticity tests with the *r*- and *i*-band data are effective, but tests with the *r*- and *g*-band data are less powerful, mainly because most variables are red stars and so have low *g*-band flux. Cuts in an analogue of the colour–magnitude diagram can also be useful in excising contaminants from our data set, although care is needed to avoid eliminating too many possible sources.

It is also important that every stage in the selection of microlensing candidates is objective and reproducible by a computer, otherwise the efficiency of the process cannot be calculated by Monte Carlo simulations. The only permissible application of cuts by eye is to eliminate phenomena that are not included in the simulations – such as light curves associated with cosmic rays, satellite trails or CCD defects. In our work, any such light curves are eliminated (sometimes by visual examination of the original images) before the construction of the input catalogue of 44 554 light curves. We remark that the M31 microlensing survey carried out by de Jong et al. (2004) involved a final cull by eye at the end of candidate selection (reducing the number of candidates from 126 to 12). This is not reproducible by computer and so the efficiency of the de Jong et al. survey cannot be determined by a straightforward computer algorithm.

3.2 An automated algorithm for candidate selection

All light curves are automatically fitted with two models: (1) a Paczyński curve with a flat baseline and (2) a crude variable star light curve, using all the data points in all three passbands. The first fit has nine parameters (namely $f_{s,g}$, $f_{s,r}$, $f_{s,i}$, the source fluxes in *g*, *r* and *i*; $f_{b,g}$, $f_{b,r}$, $f_{b,i}$, the blending fluxes in *g*, *r* and *i*; u_0 the impact parameter; t_{max} the time of maximum and t_E, the Einstein radius crossing time). The fitted parameters are known to be degenerate in blended events (e.g. Woźniak & Paczyński 1997), but this is not important at such an early stage in the selection algorithm. There are some restrictions on the values of the fitted

parameters – namely, all the fluxes must be positive, the impact parameter is restricted to lie between 0 and 2, and the Einstein radius crossing time t_E must lie between 1.5 d and 1.5 yr. We refer to this fit as the global microlensing fit. The second model to be fitted is a single sinusoid. Of course, most variable stars have more complicated light curves than sinusoids, but this crude model is adequate for our purpose here. The fit has eight parameters (the baseline and the amplitude in g, r and i, the phase and the period). However, in practice, we do not fit the period as a freely selectable parameter, but fix it to one of 30 values and fit the rest of the parameters. The trial set of periods is determined as follows. The data are first 'robustified' (Alcock et al. 2000) by removing the five highest and lowest values. The robustified data for each passband are analysed with Lomb's periodogram (see, e.g., Press et al. 1992) and the values of the frequency corresponding to the highest 10 peaks in each of the three passbands are extracted. As before, all fluxes are restricted to be positive in the fit. We refer to this fit as the variable star fit.

From these two fits, we can compute the improvement of χ^2 for the Paczyński curve fit over a flat baseline ($\Delta\chi^2_{\mathrm{micro}} = \chi^2_{\mathrm{bl}} - \chi^2_{\mathrm{micro}}$) and the similar χ^2 improvement of the variable star light curve fit over a flat baseline ($\Delta\chi^2_{\mathrm{var}} = \chi^2_{\mathrm{bl}} - \chi^2_{\mathrm{var}}$). In other words, both $\Delta\chi^2_{\mathrm{micro}}$ and $\Delta\chi^2_{\mathrm{var}}$ are the χ^2 improvement (the absolute value of the decrease in χ^2) so that a larger number means a better fit. As shown in Fig. 1, there is an enormous cloud in the ($\Delta\chi^2_{\mathrm{micro}}$, $\Delta\chi^2_{\mathrm{var}}$) plane corresponding to variable stars. The *first cut* is chosen so as to excise this cloud by insisting that

$$\Delta\chi^2_{\mathrm{var}} < 0.75\,\Delta\chi^2_{\mathrm{micro}}. \tag{1}$$

This empirically derived cut is shown as the red line in Fig. 1. It ruthlessly reduces the candidate light curves by an order of magnitude from 44 554 to 4911.

It is instructive to compare this approach with the classical statistical method used for hypothesis testing, namely the F test (e.g. Lupton 1993, ch. 12). Suppose we wish to distinguish between two models (here variable stars and microlensing) that are both fitted to the same number of data points, but have different numbers of free parameters. Then, we construct $\Delta\chi^2 = \chi^2_{\mathrm{var}} - \chi^2_{\mathrm{micro}} = \Delta\chi^2_{\mathrm{micro}} - \Delta\chi^2_{\mathrm{var}}$ and compare this with the difference in the number of free parameters Δn. If the additional parameters are statistically redun-

dant, the χ^2 difference $\Delta\chi^2$ distributes as a χ^2 distribution with Δn degrees of freedom. Then, roughly speaking, provided that $\Delta\chi^2 \gg \Delta n$, one can conclude that the added free parameters are significant and thus that the model with the greater number of parameters is favoured over that with fewer. The exact criterion, which can be formally written as

$$\Delta\chi^2 > F(\Delta n), \tag{2}$$

is dependent on the chosen significance level. For the case in hand, $\Delta n = 1$. So, assuming the reported photometric errors are well determined and random, the 1σ (68 per cent) confidence corresponds to $F = 1$. However, under a conservative approach, $F = 50$ is probably more reasonable, given the characteristics of our data. The cuts corresponding to $F = 1$ and 50 are shown as blue and green lines in Fig. 1. They run through the cluster of variable stars, rather than just below it, so they are not optimal. The cut that we actually apply (equation 1) is both more stringent and more efficient, as it moves the intercept downwards to just below the cluster of variable stars.

The *second cut* demands that there are data points on both the rising and falling part of the light curve. Specifically, we require two data points within $1.5\,t_{1/2}$ of the peak on each side, and we require more than five data points within $6\,t_{1/2}$ in one passband, plus at least one further data point in a second passband. All these data points must be at least 3σ away from the baseline as judged by the global microlensing fit. This may lose some bona fide microlensing events, for which incomplete coverage means that we cannot be certain of the nature of the event.

The *third cut* proceeds by carrying out a local blended Paczyński fit to all data points within $6t_{1/2}$. This is performed in as many passbands as possible, usually two as the data are not generally available in all three bands during the event. The rationale for a local Paczyński fit is that the global fit may not return the proper baseline because of contamination by nearby variable stars. However, even if the baseline is bumpy, we expect that the maximum is a good fit to the theoretical curve. This is because the typical time-scale for the contaminating variable behaviour is much longer than the duration of microlensing, so that the baseline can be approximated as constant during the event. The fluxes are again restricted to be positive and the impact parameter must lie between 0 and 2. The Einstein radius crossing time t_E may take any value between 1.5 d and 1.5 yr. We refer to this fit as the local microlensing fit. If the χ^2_{local} per degree of freedom is greater than 2, then the event is rejected.

The *fourth cut* follows from insisting that the global microlensing fit is consistent with the fact that the source star must be unresolved. Recall that all the resolved stars were removed from our starting catalogue by our mask. Therefore, none of the candidates corresponds to a resolved star in the reference frame. In particular, global microlensing fits with low maximum magnification A_{max} may predict that the source is bright and resolved, which would make the microlensing interpretation inconsistent. It is straightforward to derive the detection limit for a resolved star, assuming (i) that it is limited by the Poisson noise of the 'sky', (ii) that the baseline counts of the superpixels (7×7 pixel^2) are dominated by the same 'sky' as that on the reference frame and (iii) that the point spread function is an isotropic Gaussian. The limit of the source flux to be $n\sigma$ above the sky noise is

$$f_{\mathrm{limit}} = \frac{na\sqrt{\pi}}{14\left[1 - 2^{-(a/s)^2}\right]}\sqrt{\frac{f_{\mathrm{bl}}}{E}}, \tag{3}$$

where E is the exposure time, f_{bl} is the baseline of the global microlensing fit, s is the seeing disc (FWHM) in pixels and a is the
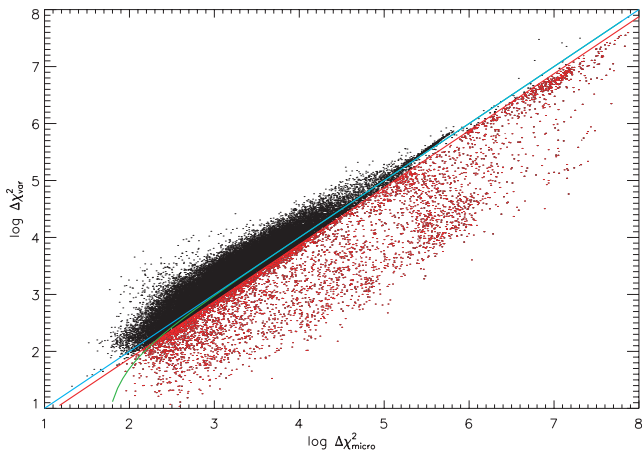

Figure 1. Each light curve is plotted in ($\log \Delta\chi^2_{\mathrm{micro}}$, $\log \Delta\chi^2_{\mathrm{var}}$) space. Note the dense black cloud representing variable stars. The first cut (red line) is given by equation (1) and is chosen to excise this cloud. Light curves passing the first cut are plotted as red dots. For comparison, the statistical $F = 1$ test of equation (2) is shown as a blue line, while the $F = 50$ test is shown as a green line.

aperture (diameter) size in pixels. The seeing and exposure times of the reference frames are listed in table 1 of An et al. (2004b). Assuming that $a \approx s$ then gives us values for all the physical quantities on the right-hand side of equation (3). Clearly, there is more noise than pure photon noise, and thus, to be conservative, we choose $n = 10$ for the detection limit;

$$f_{\mathrm{limit}} = \frac{10s\sqrt{\pi}}{7}\sqrt{\frac{f_{\mathrm{bl}}}{E}}, \tag{4}$$

where we convert the unit of fluxes into ADU s^{-1} so that, strictly speaking, the actual limit corresponds to n times the square root of the gain (in e ADU^{-1}). The fourth cut is to require that the 1σ lower bound of the r-band source flux[1] of the global microlensing fit be less than the flux limit of equation (4). There are well-known degeneracies in the fitting of blended microlensing events, with the consequence that the source flux from the fit is not wholly reliable. This is why such a cautious choice is made in the resolved flux limit in equation (4). The fitting routine can mimic flatter-topped light curves of a population of variable stars with low-magnification microlensing light curves. The effect of this cut is to remove such spurious, large impact parameter events.

The *fifth cut* is an achromaticity test. As microlensing is intrinsically achromatic, the ratio of the flux difference over the baseline (that is, the variable flux) between two different passbands must be constant over time (e.g. Ansari et al. 1997; Aurière et al. 2001). The same is not true for most pulsating stellar variables, which usually become bluer as they approach maximum (e.g. Sterken & Jaschek 1996). The data points in two passbands along the bump are binned, with the bin-size depending on the width of the bump and the number of available data points. A straight line and a quadratic polynomial are then fitted to the flux measurements in two passbands. If a linear fit is preferred (i.e. the quadratic coefficient is statistically redundant), then this suggests that the light curve is achromatic. The fifth cut is to insist that the linear fit is preferred, namely

$$\Delta\chi^2 = \chi^2_{\mathrm{linear}} - \chi^2_{\mathrm{quad}} < 2. \tag{5}$$

For events that peak in the first season, the r- and g-band data are used. For all other events, the r- and i-band data are used. The effect of the achromaticity cut is shown in Fig. 2.

The *sixth cut* is the elimination of the most troublesome long-period variable stars that show only one bump in the POINT-AGAPE data set, using an analogue of the colour–magnitude diagram. As a measure of the brightness of each variable, An et al. (2004b) introduced the 'pseudo-magnitude'. Because the fraction of light contributed by a variable object to the total superpixel flux at a given epoch is not known, the real magnitude of any variable cannot be determined. Under these circumstances, we measure the quantity Δf, which is the flux variation between minimum and maximum in ADU s^{-1}. We then convert Δf into magnitudes using $r(\Delta f) = r_0 - 2.5\log_{10}(\Delta f)$, where r_0 is the zero-point, i.e. the r-band magnitude of a star where the flux is 1 ADU s^{-1} in the reference image. This quantity $r(\Delta f)$ – or $i(\Delta f)$, $g(\Delta f)$ as appropriate – is called the pseudo-magnitude. Similarly, the difference between two pseudo-magnitudes, such as $r(\Delta f) - i(\Delta f)$, is a proxy for the colour. This allows us to build an analogue of the classical colour–magnitude diagram for pixellated data, as shown in Fig. 3. Note that,

---

[1] The source flux found from the fit is actually the fraction of the source light falling within the superpixel and hence is a lower limit to the true source flux. By ignoring this correction ($\approx$15 per cent for 1.5-arcsec seeing), our cut is somewhat looser in practice than reported in the text.
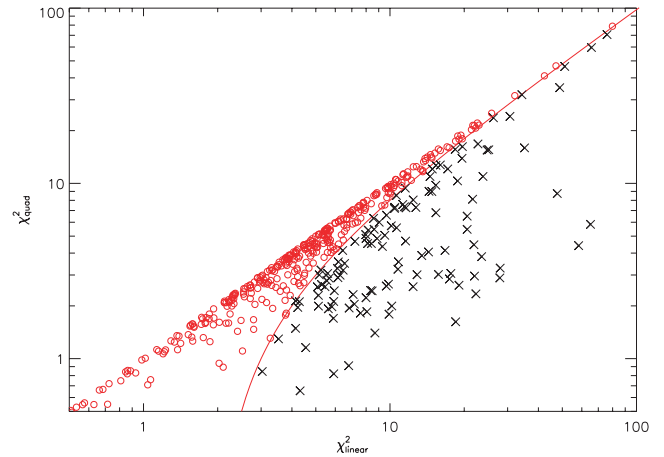
Figure 2. The 488 light curves that survive after the fourth cut are then subjected to an achromaticity test. The flux change in two passbands is fitted by a linear and a quadratic law and the light curves are plotted in the $(\chi^2_{\mathrm{linear}}, \chi^2_{\mathrm{quad}})$ space. The red circles represent the light curves that satisfy the achromaticity cut given in equation (5).
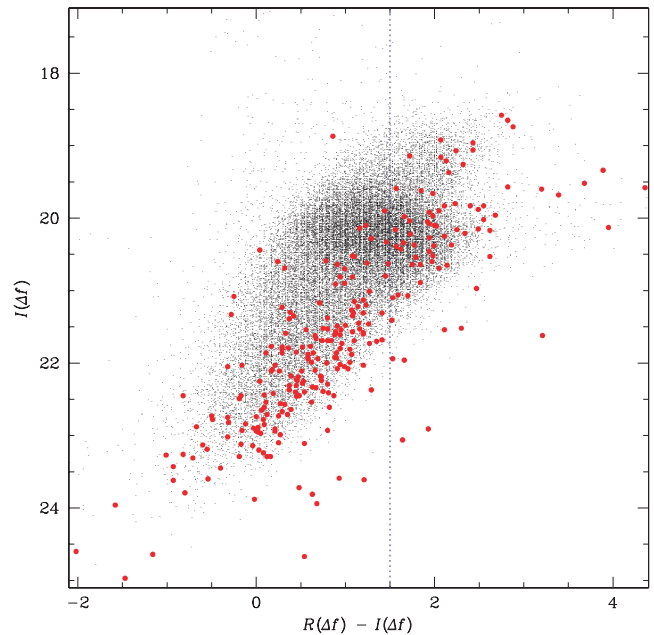

Figure 3. Variable stars are represented as grey dots on the analogue of the colour–magnitude diagram. The 369 survivors after the fifth cut are plotted as red dots. The sixth cut (dotted blue line) given in equation (6) is a colour cut that removes the redmost portion of the colour–magnitude diagram dominated by Miras. Note that $R(\Delta f) - I(\Delta f)$ and $I(\Delta f)$ are determined from the data using the transformations in Appendix A.

unlike the similar diagram presented in fig. 15 of An et al. (2004b), which is based on the instrumental Sloan system, here we have converted the Sloan magnitudes to the Cousins magnitudes using the flux calibration equation derived in Paulin-Henriksson (2002) and listed in Appendix A. Also plotted in Fig. 3 as red dots are the survivors after the fifth cut. Most of variable stars appearing in the diagram are long-period ($\gtrsim$100 d) Mira and semiregular variables, which are also the most insidious contaminants for the purpose of identifying microlensing events. We remove the redmost candidates which occupy the region of the Mira variables by imposing the

colour cut

$$R(\Delta f) - I(\Delta f) > 1.5, \qquad (6)$$

which is also shown as a straight line in Fig. 3. Note that this cut is not applied if there are only *r*- and *g*-band data. The Miras do not form a sufficiently distinct population in the $r(\Delta f)$ versus $g(\Delta f) - r(\Delta f)$ colour–magnitude diagram for us to contemplate using any such cut. The detailed form of the cut in equation (6) is chosen on the basis of numerical simulations of the source population, which suggest that the sources lie overwhelmingly in the band $0 < R(\Delta f) - I(\Delta f) < 1.5$ (see, for example, Fig. 6 discussed later).

The *seventh cut* is used to distinguish convincing, or first-level, candidates from more tentative, or second-level, candidates. First, an estimate of the signal-to-noise ratio in microlensing $(S/N)_{micro}$ is computed using

$$(S/N) = \sqrt{\frac{1}{N} \sum_i \left(\frac{f_i - f_{bl}}{\sigma_i} \right)^2}, \qquad (7)$$

where the sum is over all N data points within $t_{1/2}$, while f_i is the flux measurement, σ_i is its error and f_{bl} is the baseline of the global microlensing fit. Then, an estimate of the signal-to-noise ratio in the residuals $(S/N)_{res}$ is calculated as follows. The median of $|f_i - f_{bl}|$ is calculated for the set of flux measurements outside $3\,t_{1/2}$ of either side of the peak. Then, equation (7) is applied to compute $(S/N)_{res}$, with the sum running over all data points above the median (and outside the event). These signal-to-noise ratios are computed using all the data in each of the three passbands (g, r and i). As Fig. 4 shows, there is a cloud of data points in the $[(S/N)_{micro}, (S/N)_{res}]$ plane, together with a clear tail of events with high $(S/N)_{micro}$. The cut

$$(S/N)_{micro} > 15 + (S/N)_{res} \qquad (8)$$

is applied to leave three first-level microlensing candidates. The underlying meaning of this cut is that the signal-to-noise ratio in microlensing must exceed the signal-to-noise ratio in the residuals. Although the precise form of equation (8) is somewhat arbitrary,
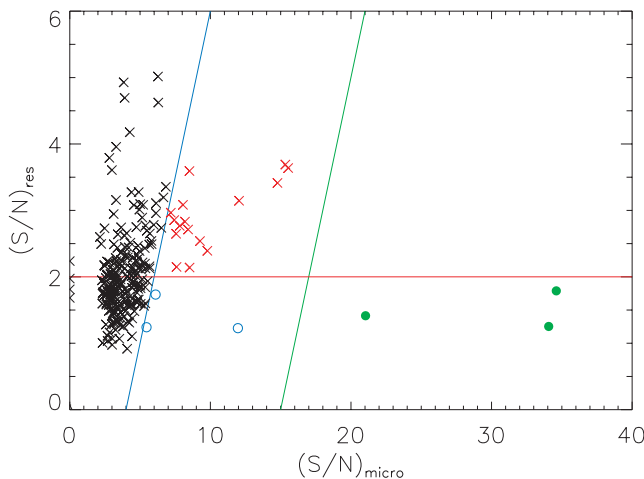

Figure 4. The 273 light curves that survive after the sixth cut are then plotted in the $(S/N)_{res}$–$(S/N)_{micro}$ plane. There is a tail of events with the microlensing signal-to-noise ratio dominating over the signal-to-noise ratio in the residuals. A set of three first-level candidates (filled green circles) are extracted by the stringent cut given in equation (8). Also shown are a looser set of three second-level candidates (open blue circles), which are extracted by the weaker cuts given in equation (9). Finally, a set of 16 third-level light curves (red crosses) are extracted by the cuts given in equation (10).

there is none the less a very clear separation between the three high signal-to-noise ratio events and the rest of the data points in Fig. 4. It is also useful to extract a looser set of candidates, the microlensing nature of which is possible but less certain. The cuts

$$(S/N)_{micro} > 4 + (S/N)_{res}, \qquad (S/N)_{res} < 2, \qquad (9)$$

are used to isolate a further three microlensing candidates, which we refer to as second level.

Finally, it also interesting to examine the light curves just below our threshold for designation as microlensing candidates. Accordingly, we extract a set of 16 third-level light curves with the cuts

$$(S/N)_{micro} > 4 + (S/N)_{res}, \qquad (S/N)_{res} > 2. \qquad (10)$$

As we will see, these are almost certainly mostly variable stars.

3.3 Properties of the candidate selection algorithm

Table 2 records the total number of light curves surviving each successive cut in the candidate selection algorithm, broken down according to the field and CCD combinations. Note that genuine microlensing events with very short $t_{1/2}$ may have failed to pass the first cut. For example, if a short time-scale event is superposed on a bumpy baseline (caused by variable stars contributions to the same superpixel), then the comparatively small number of data points in the microlensing event itself are outnumbered by the larger number of data points participating in the overall periodicity of the stellar variable. Such a light curve is better fitted by a variable star template rather than a microlensing one and so such short time-scale events fail the first cut. Similarly, the fifth cut (achromaticity) and the sixth cut (colour–magnitude) are especially needed for long time-scale events for which we cannot rely on periodicity to eliminate all repeating variables.

Fig. 5 shows the effect of each of the cuts applied singly to the starting catalogue of 44 554 light curves on the distributions of colour $R(\Delta f) - I(\Delta f)$ and magnitude $I(\Delta f)$. With the exception of the initial $\Delta \chi^2$ cut, which tends to promote brighter events, there is no obvious bias in magnitude introduced by the first six cuts. The same holds true for the colour distributions, with the exception of the explicit truncation introduced by the requirement $R(\Delta f) - I(\Delta f) < 1.5$ in the sixth cut.

This cut in the colour–magnitude diagram merits further discussion. By examination of Fig. 3, it is apparent that the cut is very loose. For example, another possibility is to excise the entire area in the colour–magnitude diagram that is occupied by Miras and long-period variables, by choosing cuts that tightly wrap around the densest regions. We elected not to do this, as this region of the colour–magnitude diagram is also where many of the potential microlensing sources are located.

A simple test of the plausibility of the magnitude and colour of our three first-level candidates is to see how representative they are of their local potential microlensing source population. By testing for local rather than global consistency, we avoid having to assume some underlying microlensing model. We define the potential source population by a simple S/N ratio criterion that reflects the quality of our candidate events without requiring a detailed simulation of the effect of each cut in our pipeline. If the cuts in our pipeline impose a strong bias upon the colour or magnitude of selected events, then we cannot expect a good agreement with a local source population selected by a simple S/N ratio cut. Similarly, there is no reason to expect agreement if the candidates are variable stars rather than microlensing events.

Table 2. The numbers of light curves surviving after each cut, broken down according to their occurrence in the CCDs and fields. F1 denotes the northern field and F2 the southern field. The original catalogue contains 44 554 variable light curves before the first cut is applied. The seventh cut is based on the signal-to-noise ratio in microlensing versus residuals and is used to extract first-, second- and third-level light curves. The first- and second-level light curves are microlensing candidates, the third-level ones are probably almost all variable stars.

| Cut	Description	F1, CCD1	F1, CCD2	F1, CCD3	F1, CCD4	F2, CCD1	F2, CCD2	F2, CCD3	F2, CCD4	Total
		6474	6518	4204	5498	2620	4705	8971	5564	44 554
1	Global fit	449	609	863	685	213	425	948	719	4911
2	Sampling	175	257	275	222	71	201	311	206	1718
3	Local fit	96	82	50	84	29	85	124	62	614
4	Unresolved source	77	60	35	64	23	70	105	54	488
5	Achromaticity	65	49	25	45	16	50	75	44	369
6	Colour–magnitude	44	33	21	34	12	32	61	36	273
7	Signal-to-noise: first level	0	0	0	0	0	1	1	1	3
7	Signal-to-noise: second level	0	0	0	0	1	0	2	0	3
7	Signal-to-noise: third level	2	2	0	4	0	2	3	3	16
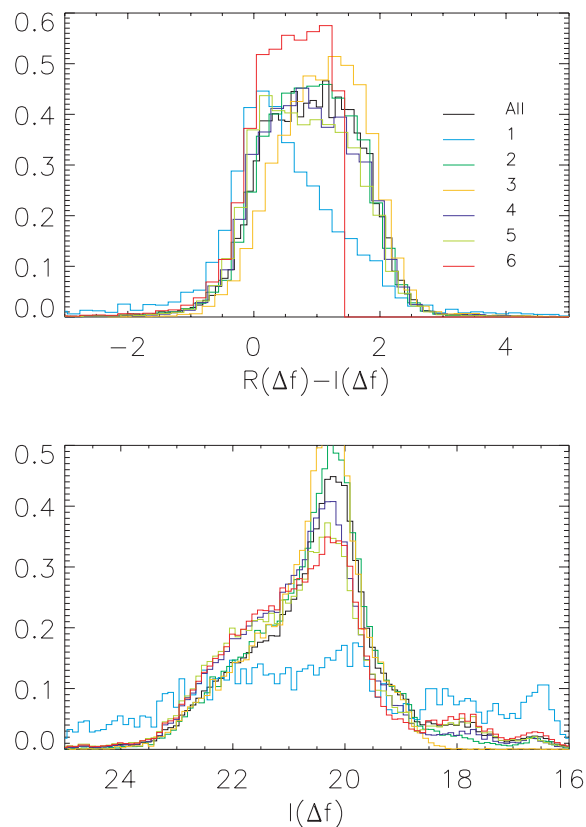
Figure 5. For each cut applied singly to the 44 554 light curves in the catalogue, the distributions of the survivors are shown as a function of the Cousins colour $R(\Delta f) - I(\Delta f)$ and magnitude $I(\Delta f)$. The cuts are: (1) $\Delta\chi^2$ (in cyan), (2) sampling (in green), (3) local fit (in brown), (4) resolved source (in blue), (5) achromaticity (in yellow) and (6) colour–magnitude (in red). The distributions in the starting catalogue are shown in black. There is no bias introduced via any of the cuts – except of course the colour–magnitude cut which introduces a cut-off at $R(\Delta f) - I(\Delta f) = 1.5$.

The sources are drawn from synthetic stellar population models (cf. Girardi & Salaris 2001), assuming similar mass, age and metallicity distributions to the Milky Way bulge and disc. For each synthetic star we compute the maximum impact parameter u_{max} that permits a detection of the magnification peak with S/N = 20. As u_{max} depends upon the local M31 surface brightness, it is computed

for the location of each of the three candidates using the surface photometry of Walterbos & Kennicutt (1987). As the microlensing rate scales with u_{max}, this parameter is essentially the statistical weight that the star may act as a potential microlensing source.

Fig. 6 shows the transformed Cousins $I(\Delta f)$ versus $R(\Delta f) - I(\Delta f)$ pseudo-magnitudes for each of the three first-level candidates, plotted together with their respective local source populations, as selected by their u_{max} weighting. The pseudo-magnitude for each synthetic star is computed from a randomly realized impact parameter u between 0 and u_{max}. For candidates 1 and 3 we employ a disc synthetic stellar population, whilst for candidate 2 we use a bulge stellar population model. The positions of candidates 1 and 3 appear consistent with main-sequence disc source stars, whilst candidate 2 is consistent with a lensed bulge giant source star. The consistency between the magnitudes and colours of the three events and their local source populations bolsters their interpretation as microlensing events.

4 THE THREE FIRST-LEVEL CANDIDATES

4.1 The light curves

Figs 7–9 show 3-yr light curves of the three first-level candidates in the r and i bands, together with close-ups of the peaks in g, r and i, if data are available. Overplotted on all the light curves are the best variable star fit (in blue), the best global microlensing fit (in red) and the best local microlensing fit (in green). The time of maximum magnification t_{max} is indicated by a red vertical line and the full-width at half-maximum time-scale $t_{1/2}$ by a red horizontal line. The lower three subpanels show blow-ups of the data points within $1.5 t_{1/2}$ on either side of the peak. These are the data points used in the local microlensing fit. The number of data points in each passband is recorded in the top right-hand corner. The χ^2 values calculated in the selection algorithm, together with the signal-to-noise ratios, are recorded in the top panel of each light curve, enabling the reader to locate the positions of the individual events in Figs 1 and 4. Fig. 10 shows the achromaticity test for the three first-level candidates. The red (green) curves show the straight line (quadratic) fit to the flux measurements in two passbands. In all cases, the quadratic coefficients are redundant, as judged by the χ^2 given in the top corners of the panels of Fig. 10.

Table 3 gives the locations of the three first-level candidates in right ascension and declination, together with the projected distance from the centre of M31. All lie in the southern field, one each on
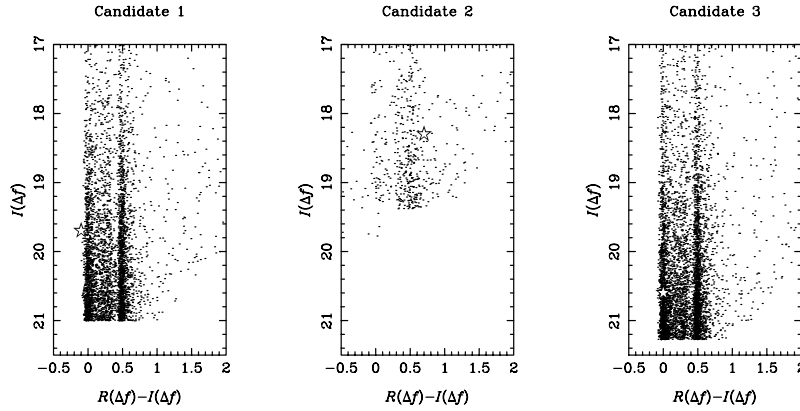
Figure 6. Synthetic Cousins $I(\Delta f)$ versus $R(\Delta f) - I(\Delta f)$ colour–magnitude diagrams showing the final three first-level candidates (star symbols) and their corresponding local source populations. For candidates 1 and 3 we assume a disc population, whilst for candidate 2 we assume the source belongs to the bulge. The local source populations are generated by assuming a peak S/N ratio threshold of 20 (see the main text). The size of the star symbols reflect the 1σ error of ~ 0.1 mag in the transformed photometry. (Note that the locations of the candidates are deduced using the microlensing fit, and so differ slightly from the locations in Fig. 3.)

CCDs 2–4. Table 4 shows the parameters of the microlensing fit for the three candidates. Nine microlensing parameters were fitted (namely the source flux in g, r and i, the blending flux in g, r and i, the impact parameter, the time of maximum and the Einstein radius crossing time) using all data points in all bands. However, the g-band data are patchy, so the source and blend fluxes in g are poorly constrained and not reported in Table 4. All three events have magnification A_{max} above the threshold value of $3/\sqrt{5} = 1.34$ and so have impact parameter smaller than unity. All three have very short time-scales ($t_{1/2} = 2.0$, 2.3 and 3.4 d, respectively). For each of the three candidates, thumbnail r- and i-band images of the reference frame (left) and the frame of the peak of the event (right) are shown in Fig. 11.

4.2 Candidate 1

Candidates 2 and 3 have been discussed previously by Paulin-Henriksson et al. (2002, 2003). However, candidate 1 is new and worthy of some detailed discussion here.

While candidate 1 has formally passed the second cut, we find that, from Fig. 7, the rising part of the light curve is not well constrained from the available data. This raises the question as to whether this light curve can be modelled by certain types of cataclysmic variable, notably a nova. This is particularly worrying as the calibrated colour of the candidate ($R - I = -0.1$) is close to that of the classical novae found in the POINT-AGAPE data set by Darnley et al. (2004). Furthermore, a fair fraction of the novae of Darnley et al. (2004) do not show very significant chromatic behaviour over a short time-scale. However, we find that the observed magnitude at the maximum of the light curve ($R_{max} \lesssim 19.5$) is rather faint compared with the other novae of Darnley et al. (2004) and that the decline rate is rather too fast ($dR/dt \gtrsim 0.5$ d^{-1}), although the sparse sampling may mean that we missed the actual maximum, which may have occurred earlier and been brighter. The possibility of this being a nearby dwarf nova is much less significant, as the light curve shows no sign of repeating bursts. It is also very unlikely to be a background (super-)nova considering its fairly blue colour. Its presence in the disc strongly suggests that, if it were a background event, it would be heavily reddened.

In fact, the microlensing fit result for candidate 1 indicates that the source should be relatively bright if this is indeed a microlensing

event. Based on equation (4), we find the nominal detection limit of the source in the r band at the location of candidate 1 to be 5.8 ADU s^{-1}, while the source flux from the fit is 8.8 ± 3.1 ADU s^{-1}. In other words, although this candidate has passed the fourth cut, it is still just possible that we may find a resolved source at the reference frame, if the microlensing fit is correct. [2]

After examining the frames when the event is at the baseline, including the reference frame, we find 'a resolved object' near the event ($\lesssim 1$ pixel). The fixed aperture (10-pixel aperture) photometry on this object indicates that, although the r-band flux (9.5 ± 1.7 ADU s^{-1}) is consistent with this being the microlensed source, the i-band flux (12.1 ± 2.0 ADU s^{-1}) is rather too bright; i.e. 'the colour' of this object is too red. However, we note that it is very likely that, for a given seeing, most 'resolved objects' in our images are blended mixtures of several stars and thus it is still possible that the object actually contains the microlensed source. This idea could be tested by examining an image with better resolution, such as a *Hubble Space Telescope* (*HST*) image. Unfortunately, the nearest *HST* archival pointing is 2.6 arcmin away. Another possibility is to examine the precise position of the centroid of this object and show that, while the r- and i-band centroid are shifted, the r-band position is closer to the centroid position of the event at maximum. This, though, is a difficult calculation to perform convincingly, as it requires precise modelling of the point spread function.

If we accept the fit, the magnitude and the colour of the source ($R = 21.9 \pm 0.4$, $R - I = -0.1$) are consistent with that of an M31 main-sequence B star under moderate extinction. This is quite reasonable as the event occurs on the disc near the dust lane, possibly along the star formation region or close to an OB association.

4.3 Comparison with other high signal-to-noise ratio surveys

Paulin-Henriksson et al. (2003) carried out a survey fine-tuned for short time-scale ($t_{1/2} < 25$ d), high signal-to-noise ratio

---

[2] This is not in contradiction with the fact that we masked the region around the 'resolved' stars as the mask was around *bright* resolved stars. The typical magnitude of the resolved stars used in building the mask is $R \sim 20.5$. The source magnitude corresponding to the fit result of candidate 1 is much fainter at $R = 21.9 \pm 0.4$.

Level 1 candidate 1

$\chi^2(\text{micro, global}) = 687.3$ (S/N) micro $= 34.61$

$\chi^2(\text{micro, local}) = 10.8$ (S/N) res $= 1.79$

$\chi^2(\text{const}) = 9483.6$ $A_{max} = 8.99 \pm 3.13$

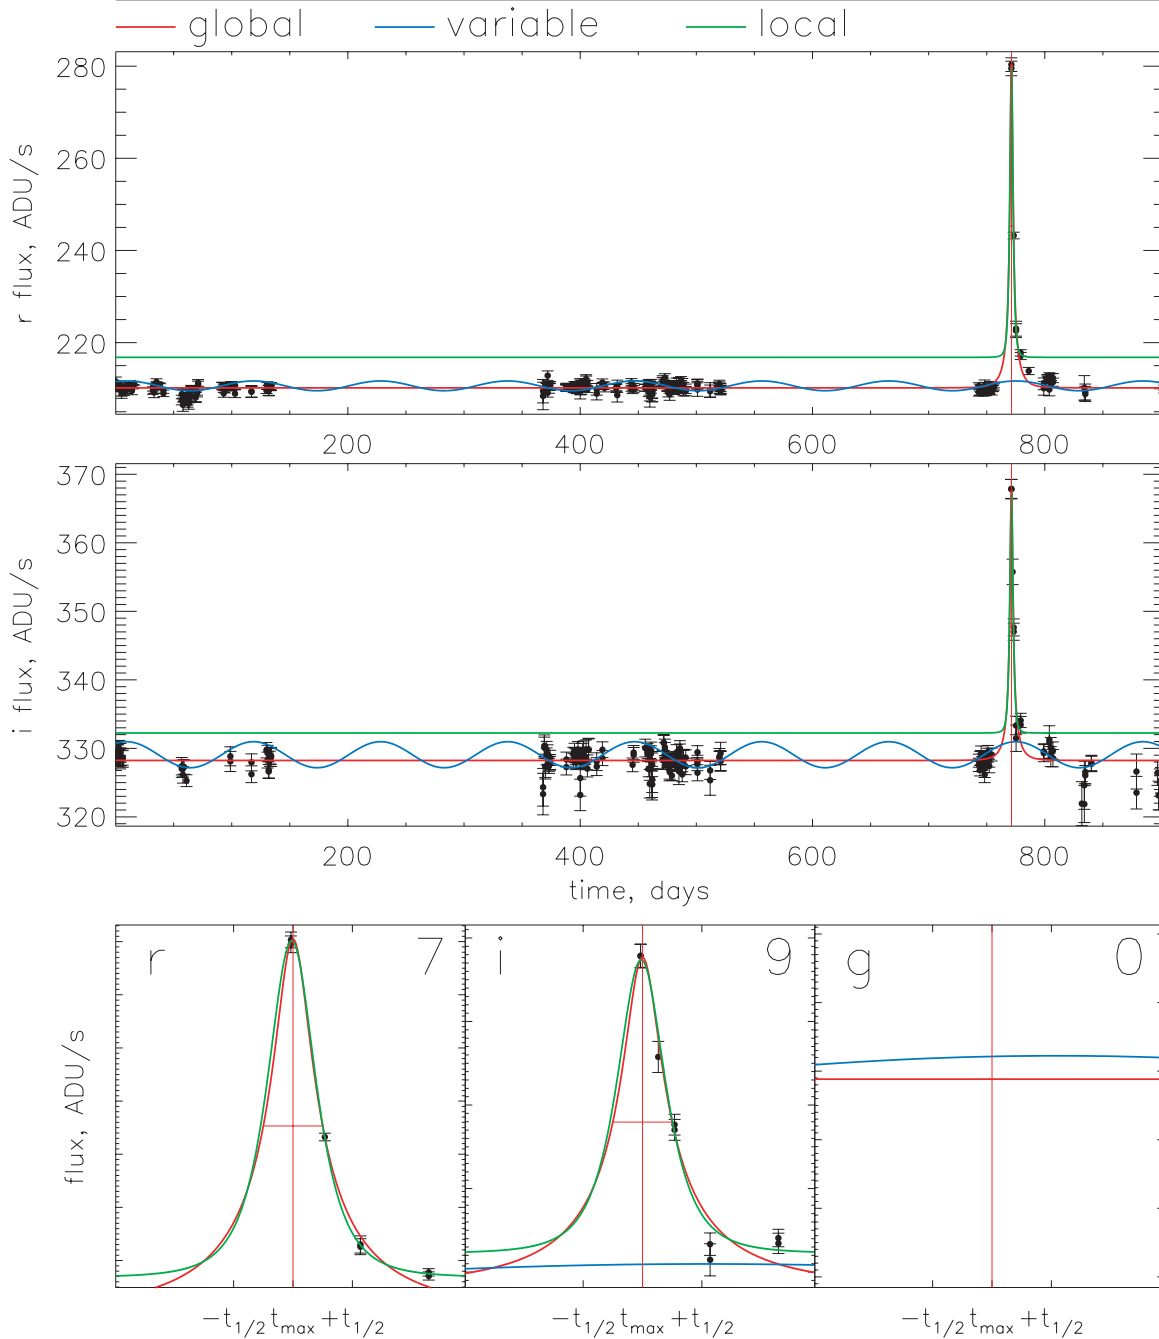$\chi^2(\text{var}) = 8969.9$ $t_{1/2} = 3.43 \pm 0.50$

Figure 7. Light curves in the *r* and *i* bands of first-level candidate 1 in Table 4 from the southern field, CCD 2. Overplotted in blue is the best bumpy variable fit, in red is the best global microlensing fit and in green is the best local microlensing fit. The location of the peak of the event is marked by a red vertical line. The lower subpanels show the peak in *g*, *r* and *i*, and the number of data points is recorded in the top corner of each subpanel.
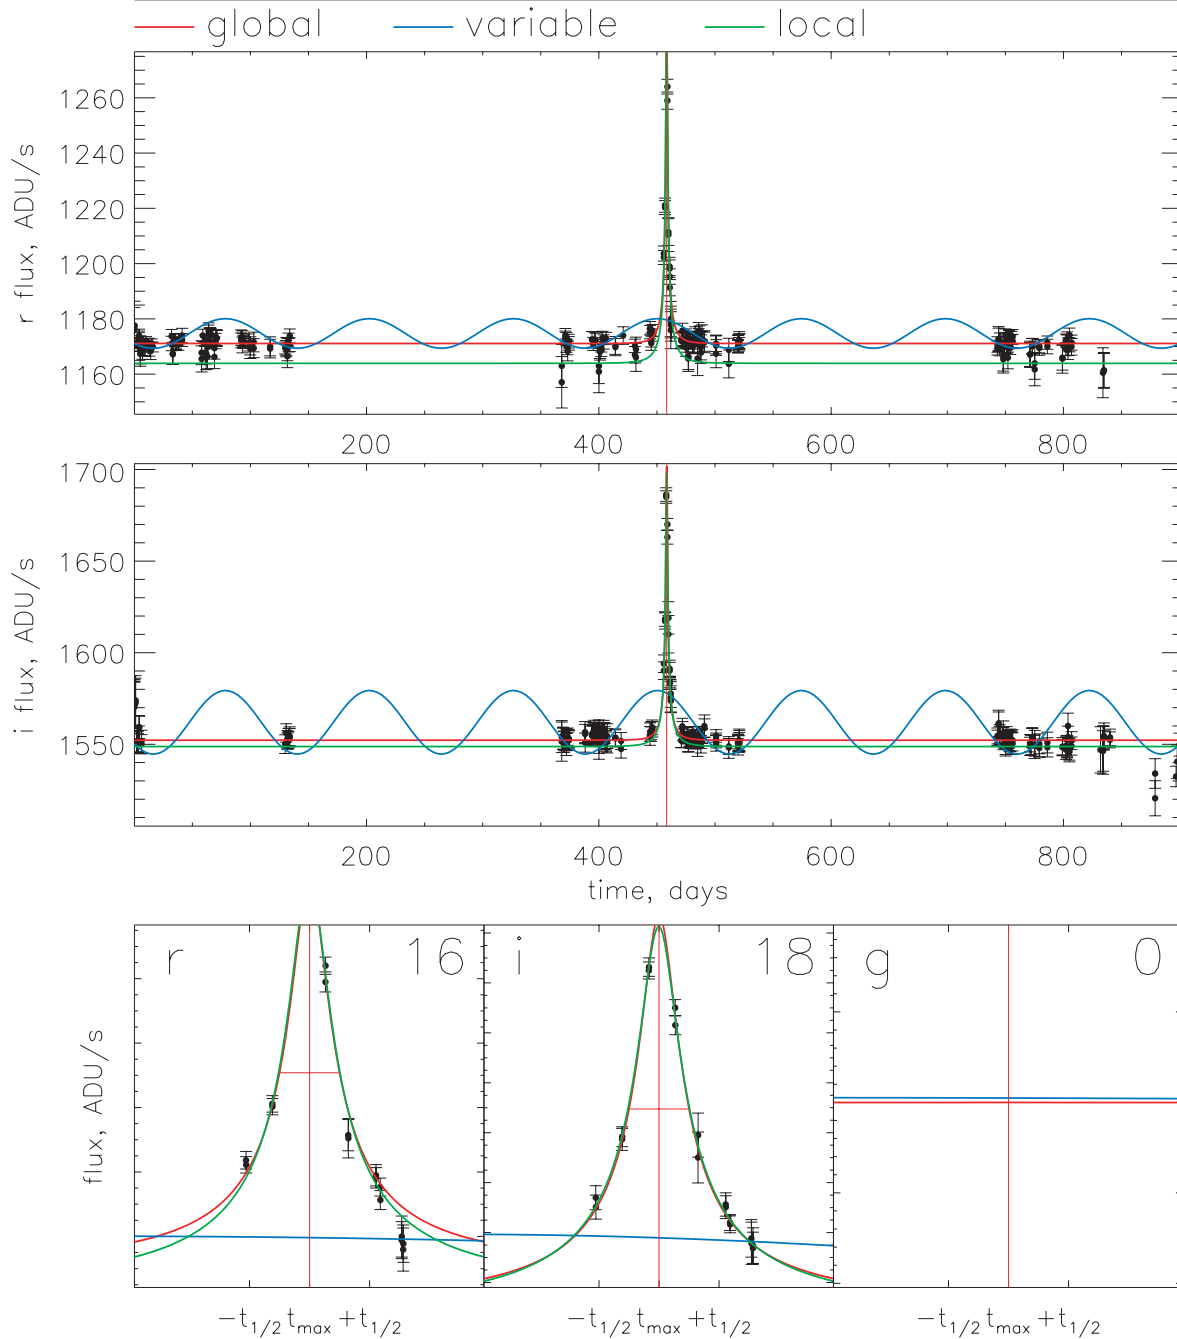
Figure 8. Light curve of first-level candidate 2 in Table 4 from the southern field, CCD 3. This is PA 00-S3 of Paulin-Henriksson et al. (2003) and WeCAPP-GL-1 of Riffeser et al. (2003).

Level 1 candidate 3

χ^2(micro, global) = 333.0   (S/N) micro = 21.04

χ^2(micro, local) =   7.8   (S/N) res = 1.41

χ^2(const) =2844.4 A_{max} =259.48±312.41

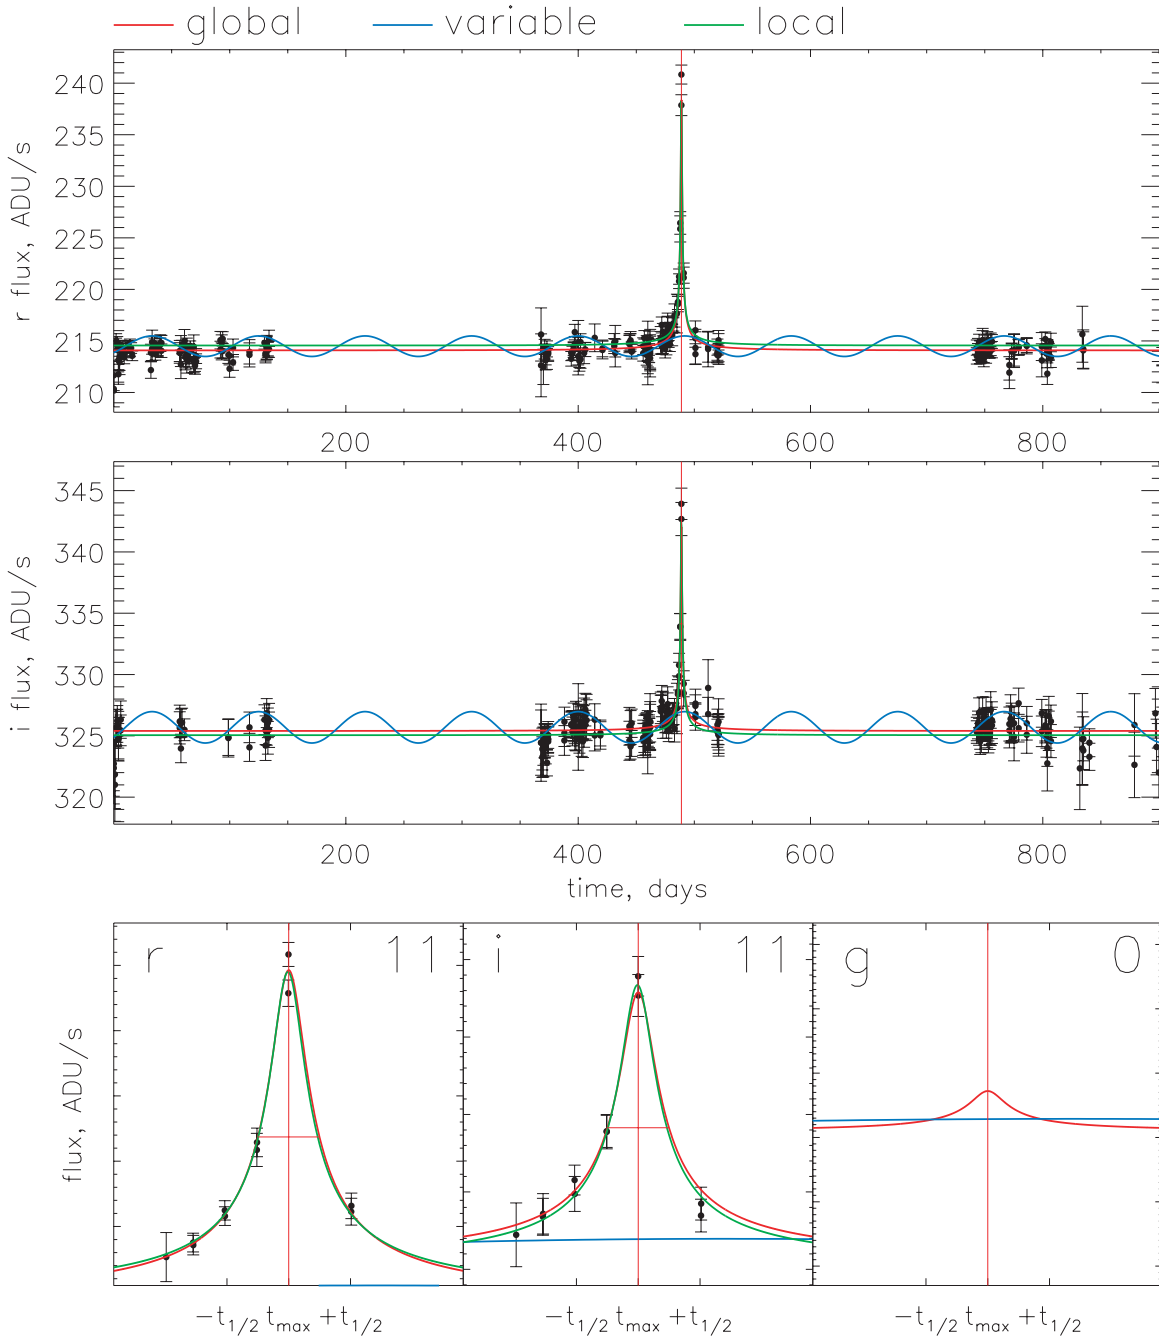χ^2(var) =2590.0 $t_{1/2}$ =   1.95±   0.10

Figure 9. Light curve of first-level candidate 3 in Table 4 from the southern field, CCD 4. This is PA 00-S4 of Paulin-Henriksson et al. (2003).
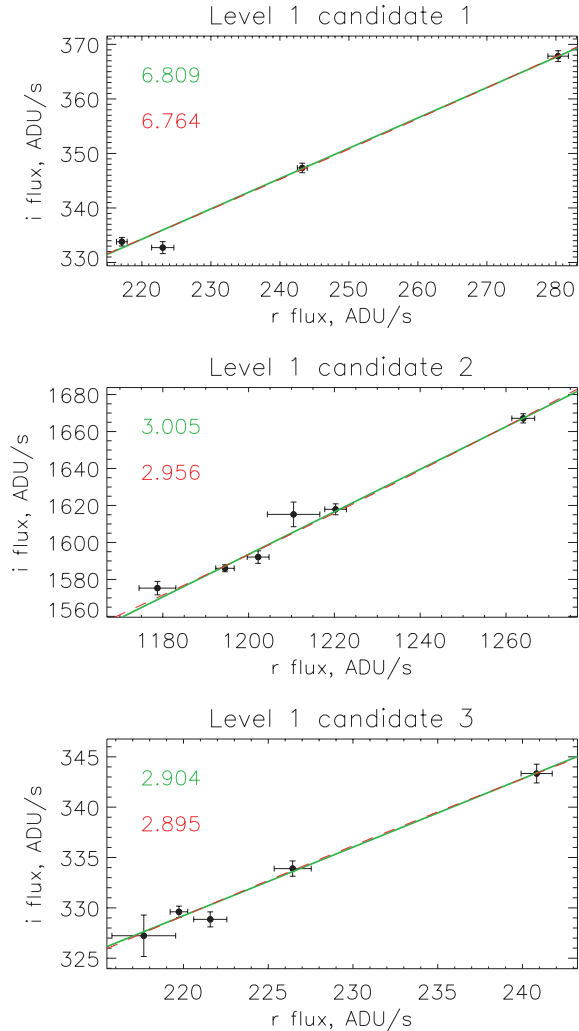
Figure 10. Achromaticity curves for the three first-level candidates. The flux in *r* is plotted against the flux in *i*, together with the best-fitting straight line (solid green) and quadratic curve (dashed red). The numbers in green (red) in the top corners of each panel are the χ^2 values for the linear (quadratic) fits.

$[R(\Delta f) < 21]$ microlensing events using the data taken in the 1999 and 2000 seasons. They found four events. One of these (PA 99-N2) is missing from our starting catalogue of 44 554 variable light curves as it has been masked out. This very bright event by chance occurs on the reference frame used to construct the resolved star catalogue (see An et al. 2004b). It is so bright that it is identified by our software as a resolved star! Hence, the area around PA 99-N2 is removed in the masking. An elaboration of our procedure – in which two resolved star catalogues are constructed at different epochs, and objects present in only one are identified as variables and so not deleted – would presumably allow us to detect PA 99-N2. A further one of the Paulin-Henriksson et al. (2003) candidates (PA 99-N1) is missing because it fails the first cut of our selection algorithm. It is a short time-scale candidate ($t_{1/2} = 1.8$) with a noisy baseline. It was originally selected based on a single significant peak in the 1999 data. Closer analysis over a longer timeframe, however, revealed that the baseline of the pixel light curve was bumpy, and that this was caused by contamination from a bright variable star located 1.1-arcsec south of the microlensed source (Aurière et al. 2001). The remaining two of the Paulin-Henriksson et al. (2003) candi-

Table 3. The locations of the three first-level and the three second-level candidates in right ascension and declination (J2000.0), together with the projected distance *R* from the centre of M31.

Candidate	Level	Field/CCD	*R* (arcmin)	RA	Dec.
1	1	F2,CCD2	22.94	$00^h\,42^m\,02^s.3$	$40°\,54'\,35''$
2	1	F2,CCD3	4.08	$00^h\,42^m\,30^s.3$	$41°\,13'\,01''$
3	1	F2,CCD4	22.54	$00^h\,42^m\,30^s.0$	$40°\,53'\,46''$
1	2	F2,CCD1	30.64	$00^h\,42^m\,57^s.7$	$40°\,45'\,37''$
2	2	F2,CCD3	3.40	$00^h\,42^m\,59^s.5$	$41°\,14'\,17''$
3	2	F2,CCD3	5.57	$00^h\,42^m\,23^s.9$	$41°\,12'\,06''$

dates (PA 00-S3 and 00-S4) are present in our sample of first-level candidates.

Very recently, a further three short-time-scale events have been detected by Paulin-Henriksson et al. (2004) using a different and independent candidate selection algorithm to this paper. A detailed analysis of these events is in preparation (Calchi Novati et al.).

The MEGA collaboration (de Jong et al. 2004) presented 14 candidate events based on the same INT imaging data as this paper, but using a different data analysis pipeline and candidate selection algorithm. Of these 14 events, 12 were new and two (PA 99-N2 and 00-S4) had already been discovered by Paulin-Henriksson et al. (2003). By matching the pixel positions, we find that (at most) six of the MEGA candidates are present in our catalogue of 44 554 light curves, namely MEGA 3, 5, 8, 9, 10 and 11. MEGA 1, 2, 4, 6, 12 and 13 are not present even in the catalogue before masking, as they are not judged to have three consecutive points above 3 σ. MEGA 7 (PA 99-N2) and 14 are removed by our mask. For the remaining six candidates, Table 5 gives the cuts in our pipeline that are failed. Note that MEGA 11 is the same as PA 00-S4, and so lies in our set of first level candidates. Apart from MEGA 11, the remaining candidates all fail at least one of our cuts. MEGA 10 only fails the final signal-to-noise ratio cut, and in fact lies very close to the second-level boundary in our Fig. 4. MEGA 3 would be adjudged a second-level candidate were it not for the fact that it failed our second cut on sampling.

4.4 The asymmetry signal

The spatial locations of our three first-level candidates, and the four candidates of Paulin-Henriksson et al. (2003), are plotted on a $g - r$ grey-scale map of M31 in Fig. 12. There are two events in common, so in total this gives five distinct candidates. Also shown is the projected major axis in blue. One of the Paulin-Henriksson et al. candidates (PA 99-N1) lies in the near side of the disc; all of the remaining events occur in the far side. Notwithstanding the claim of de Jong et al. (2004), it is hard to identify spatial gradients in the microlensing rate by a visual examination of such maps. As we have emphasized elsewhere (An et al. 2004b), the near–far asymmetry signal is confused by the effects of variable extinction and dust associated with the M31 spiral arms. This is a serious problem because the dust distribution is asymmetric over a large scale and, even worse, the dust asymmetry is maximized roughly along the major axis and so mimics the expected signal from microlensing by halo dark matter. This led An et al. (2004b) to propose searching for an east–west asymmetry, as the variable star distributions are roughly symmetric about the north–south line through the M31 centre. Considering the entire sample, the division is two events east and three events west of the north–south line. The very small number of convincing microlensing candidates that have unambiguously

Table 4. The parameters of the microlensing fit for the three first-level and the three second-level candidates: t_{max} is the time of the peak measured in days from JD 245 1392.5, $t_{1/2}$ is the full-width at half-maximum time-scale in days, t_E is the Einstein crossing time, A_{max} is the maximum amplification and f_s, f_b are the source and blend fluxes in ADU s^{-1} in the r and i bands.

Candidates	Level	t_{max}	$t_{1/2}$	t_E	A_{max}	$f_{s,r}$	$f_{b,r}$	$f_{s,i}$	$f_{b,i}$
1	1	771.3 ± 0.2	3.4 ± 0.5	10.1 ± 1.8	8.99 ± 3.128	8.8 ± 3.1	201.4 ± 3.1	4.9 ± 1.8	323.3 ± 1.8
2	1	458.4 ± 0.0	2.3 ± 0.1	11.2 ± 2.6	15.62 ± 4.350	8.2 ± 2.3	1162.9 ± 2.2	10.6 ± 3.0	1541.6 ± 2.9
3	1	488.9 ± 0.1	2.0 ± 0.1	146.6 ± 174.2	259.48 ± 312.411	0.1 ± 0.1	214.0 ± 0.1	0.1 ± 0.1	325.3 ± 0.1
1	2	753.6 ± 1.2	35.8 ± 3.8	36.3 ± 28.0	2.57 ± 2.407	3.2 ± 4.9	132.4 ± 4.9	2.2 ± 3.4	230.4 ± 3.4
2	2	46.6 ± 0.5	33.1 ± 2.1	11.1 ± 1.4	1.06 ± 0.032	257.9 ± 127.4	678.6 ± 127.4	–	–
3	2	130.7 ± 2.9	50.7 ± 5.4	21.4 ± 5.2	1.15 ± 0.122	102.3 ± 84.5	711.5 ± 84.5	295.8 ± 243.2	797.6 ± 243.2

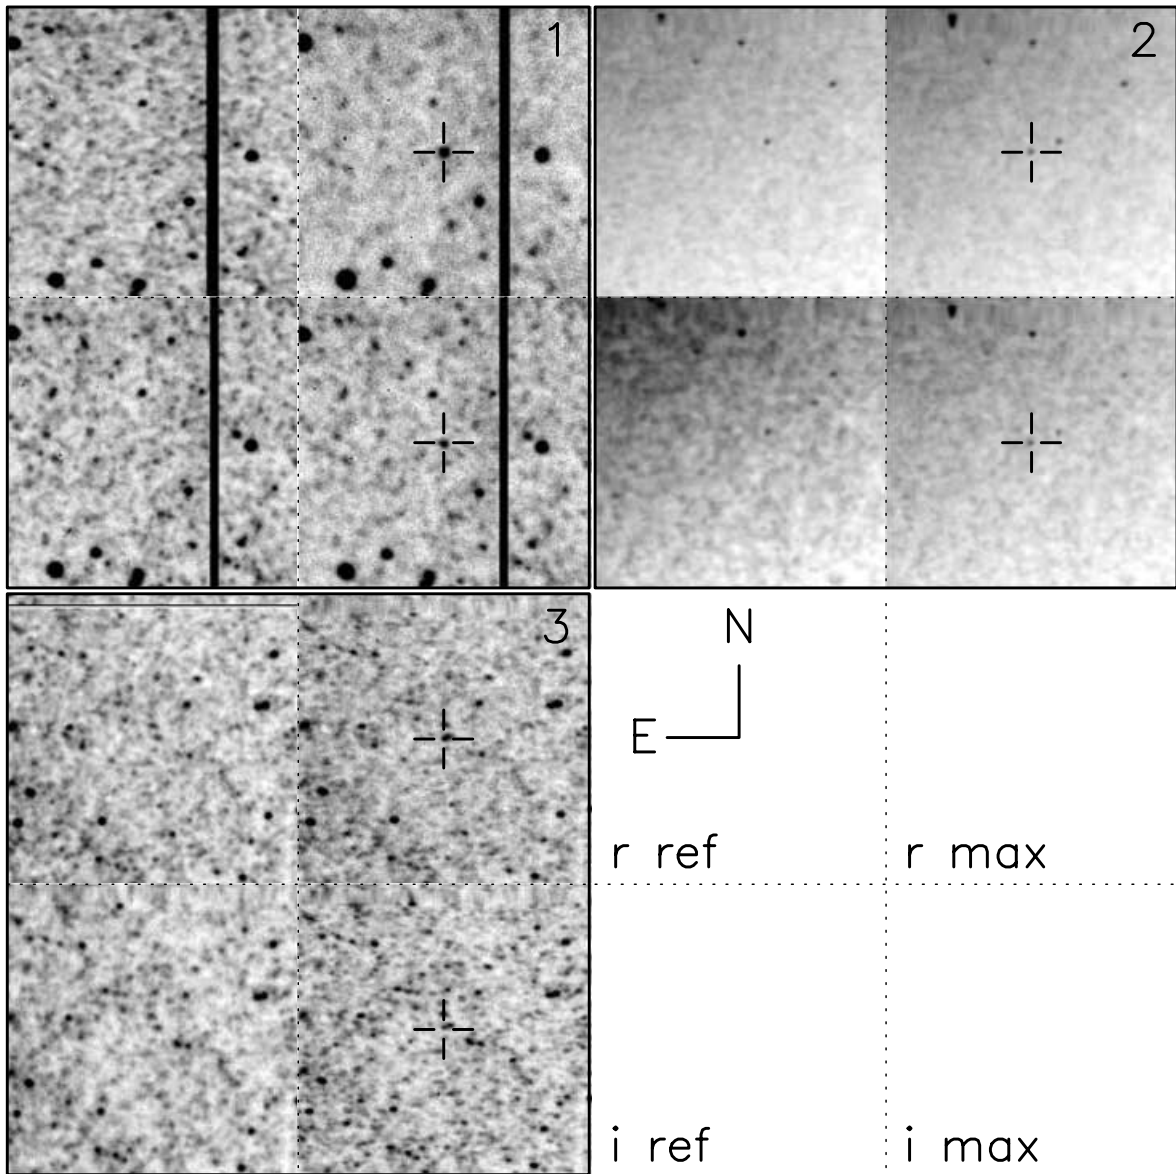Notes: the error bars assume symmetric Gaussian distributions; level 2, candidate 2 has no signal in the i band.

Figure 11. Thumbnail images (1×1 arcmin^2) of candidates 1 (top left), 2 (PA 00-S3, top right) and 3 (PA 00-S4, bottom). Each panel has four subpanels. In each panel, the r band is at the top, the i band is at the bottom, while the left-hand side corresponds to the reference frame and the right-hand side to the peak epoch. The location of the event is always at the centre of the frame and is marked with a cross in the right-hand panels.

Table 5. The MEGA candidates present in our starting catalogue and the cuts in our pipeline that they fail. Note that MEGA 11 is PA 00-S4 of Paulin-Henriksson et al. (2003) and is one of the first-level candidates in this paper.

MEGA candidate	Cuts failed
3	2, 7 (level 1)
5	1, 2, 3, 7 (all levels)
8	6, 7 (all levels)
9	2, 7 (all levels)
10	7 (all levels)
11	–
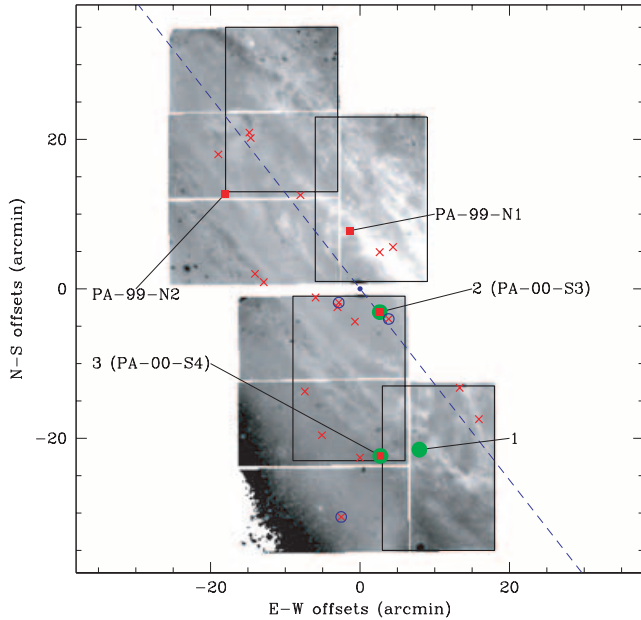
Figure 12. The locations of our three first-level candidates (green circles) are superposed on a $g - r$ grey-scale map of M31. Red squares are the locations of the four candidates given in Paulin-Henriksson et al. (2003). Also shown are the three second-level candidates (red crosses within blue circles) and the 16 third-level light curves (red crosses). The blue line marks the projected major axis and so separates the near and far disc. The optical centre of M31 is marked with a blue dot. The black rectangles indicate symmetric portions of M31 defined by An et al. (2004b). Combining all the first-, second- and third-level candidates, there are 15 events in the symmetric regions of the INT fields.

been determined from the 3 yr of M31 imaging implies that such ratios are dominated by Poisson noise.

Note that the absence of an evident asymmetry may still yield a strong result on the baryon fraction of the M31 halo, as many models predict asymmetries that should be easily detectable after 3-yr data collection (see Kerins et al. 2001). We will return to this matter in a later paper, in which the efficiency of our survey will be calculated.

5 THE THREE SECOND-LEVEL CANDIDATES AND THE 16 THIRD-LEVEL LIGHT CURVES

Bearing in mind that the second-level candidates are less securely established as microlensing, we briefly discuss some of their properties. The light curves of the three second-level candidates are shown in Figs 13–15. The parameters of the microlensing fits and the event locations are given in Tables 3 and

4. On scrutinizing the light curves, it is immediately apparent that – although the microlensing interpretation is possible – the light curves are in general noisier and sometimes show clear evidence of additional structure. The additional structure may be evidence of contamination from other nearby variable sources, or it may be evidence that some of the light curves do not correspond to genuine microlensing events at all. All three second-level candidates need to be examined over a longer baseline before they can be accepted or rejected as being genuine microlensing.

It is also interesting to see the light curves that lie just below our threshold for designation as microlensing candidates. These are the 16 third-level light curves, which are displayed compactly in Fig. B1 in Appendix B. The light curves are either those of variable stars or are highly contaminated by nearby variable stars. There are also indications that – here and there – our fitting software has not always found the best variable star fit to the data points. So, the light curves have survived to the final or seventh cut in our selection algorithm because of deficiencies in our comparison between the variable star and global microlensing templates. The locations of the events are listed in Table 6. Eight of the 16 third-level events occur in the first season, two in the second and six in the third. The second season, which has the largest number of epochs, contains the fewest events. This suggests very strongly that some (probably most) of the candidates are variables and that poor sampling is playing a role in allowing rogue candidates to survive the cuts. If the bump of a long-period variable occurs in the second season, then the end of the fall in the first season, and the beginning of the rise in the third season can be sampled, and so the event is discarded as being a possible microlensing event. If the bump occurs in the first season, then the second bump can lie in the ∼6 month gap in sampling after the second season, when M31 is not visible from La Palma, and so the light curve can survive. In the first season, data were taken predominantly in the g and r bands, whilst in the second and third season predominantly r- and i-band data were taken. For these red variables, a well-sampled light curve in the i band is best for uncovering discrepancies from the blended Paczyński fits. This pattern of sampling in different bands is almost certainly the explanation as to why half of all the third-level light curves peak in the first season. Although there may be one or two genuine microlensing candidates in the third-level sample, they are almost certainly mostly variable stars. We therefore do not classify them as microlensing candidates.

There are a total of 22 first-, second- and third-level light curves in all. Fig. 16 shows the entire sample on a colour–magnitude diagram (except for two events for which there is no i-band signal), while Fig. 12 shows their locations in M31. 17 of the 22 events lie within the symmetric regions of the INT fields. We can use the asymmetry signal-to-noise ratio definition of Kerins et al. (2003) on these 17 events to check the significance of their distribution. We find that the near–far and east–west signals are well within the Poisson error, whereas the north–south signal appears to be marginally significant. From An et al. (2004b), we recollect that a north–south signal is consistent with a population of objects affected by dust as the $g - r$ image of M31 indicates more dust in the north symmetric region. So, regardless of whether our second- and third-level events are predominantly microlenses or variable stars, the significant north–south asymmetry indicates that their distribution appears to be affected by the dust distribution across the M31 disc. The lack of far–near or east–west asymmetry again argues against a microlensing hypothesis for most of these events.

Alcock et al. (2000) have already established a precedent in microlensing studies by publishing a set of strong candidates (their 'set

Level 2 candidate 1

$\chi^2(\text{micro, global}) = 323.2$ $(S/N) \text{ micro} = 5.46$

$\chi^2(\text{micro, local}) = 57.1$ $(S/N) \text{ res} = 1.24$

$\chi^2(\text{const}) = 1377.4$ $A_{max} = 2.57 \pm 2.41$

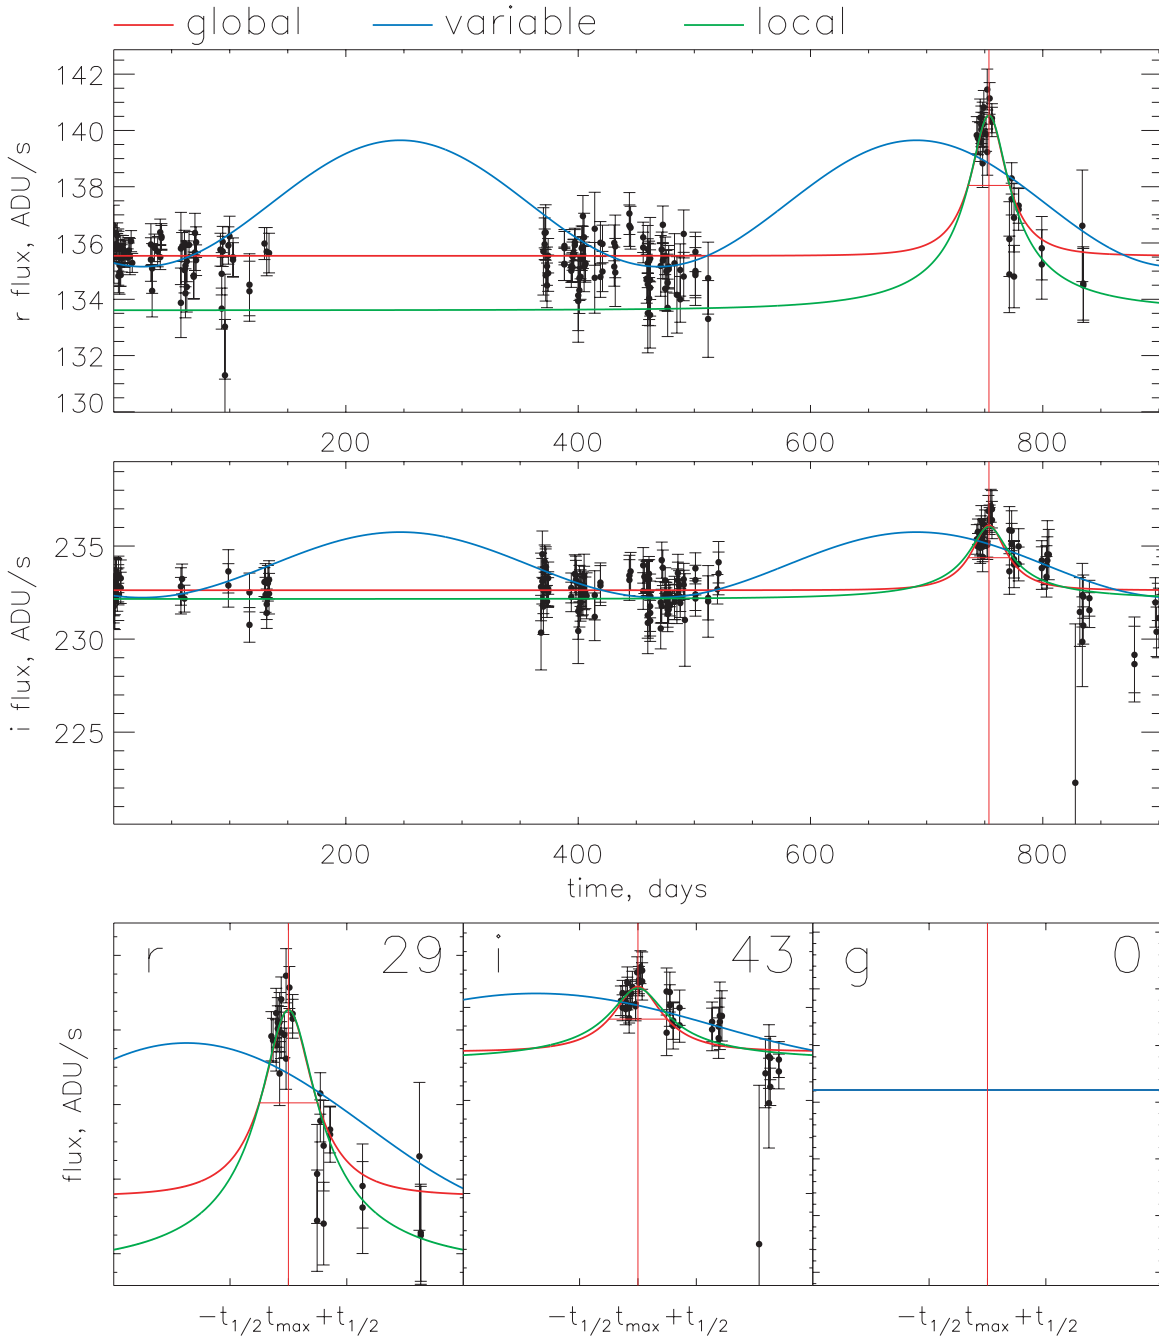$\chi^2(\text{var}) = 584.3$ $t_{1/2} = 35.83 \pm 3.80$

Figure 13. Light curves in the *r* and *i* bands of second-level candidate 1 in Table 4 from the southern field, CCD 1, Overplotted in blue is the best bumpy variable fit, in red is the best global microlensing fit and in green is the best local microlensing fit. The location of the peak of the event is marked by a red vertical line. The lower subpanels show the peak in *g*, *r* and *i*, and the number of data points is recorded in the top corner of each subpanel.

Level 2 candidate 2

$\chi^2(\text{micro, global}) = 551.3$ $(\text{S/N}) \text{ micro} = 6.10$

$\chi^2(\text{micro, local}) = 133.6$ $(\text{S/N}) \text{ res} = 1.73$

$\chi^2(\text{const}) = 1738.0$ $A_{max} = 1.06\pm 0.03$

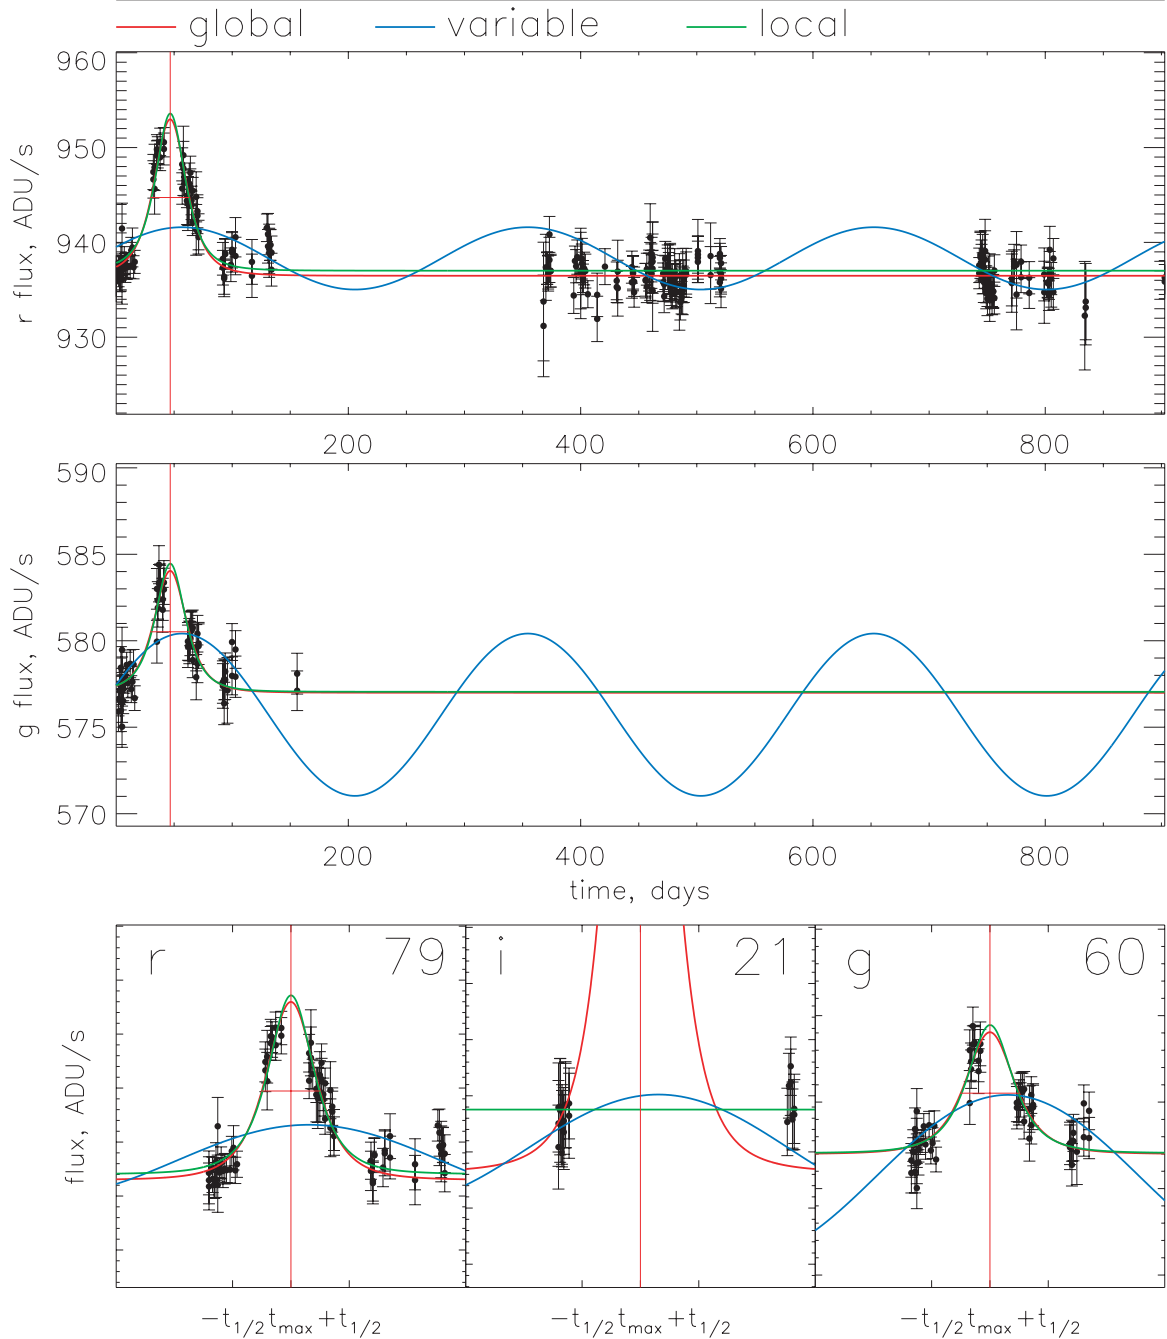$\chi^2(\text{var}) = 982.5$ $t_{1/2} = 33.08\pm 2.10$

Figure 14. Light curves in the *r* and *g* bands of second-level candidate 2 in Table 4 from the southern field, CCD 3.

Level 2 candidate 3

χ^2(micro, global) = 309.0   (S/N) micro = 11.96

χ^2(micro, local) =   69.5   (S/N) res = 1.23

χ^2(const) =2991.3 A_{max} =   1.15±   0.12

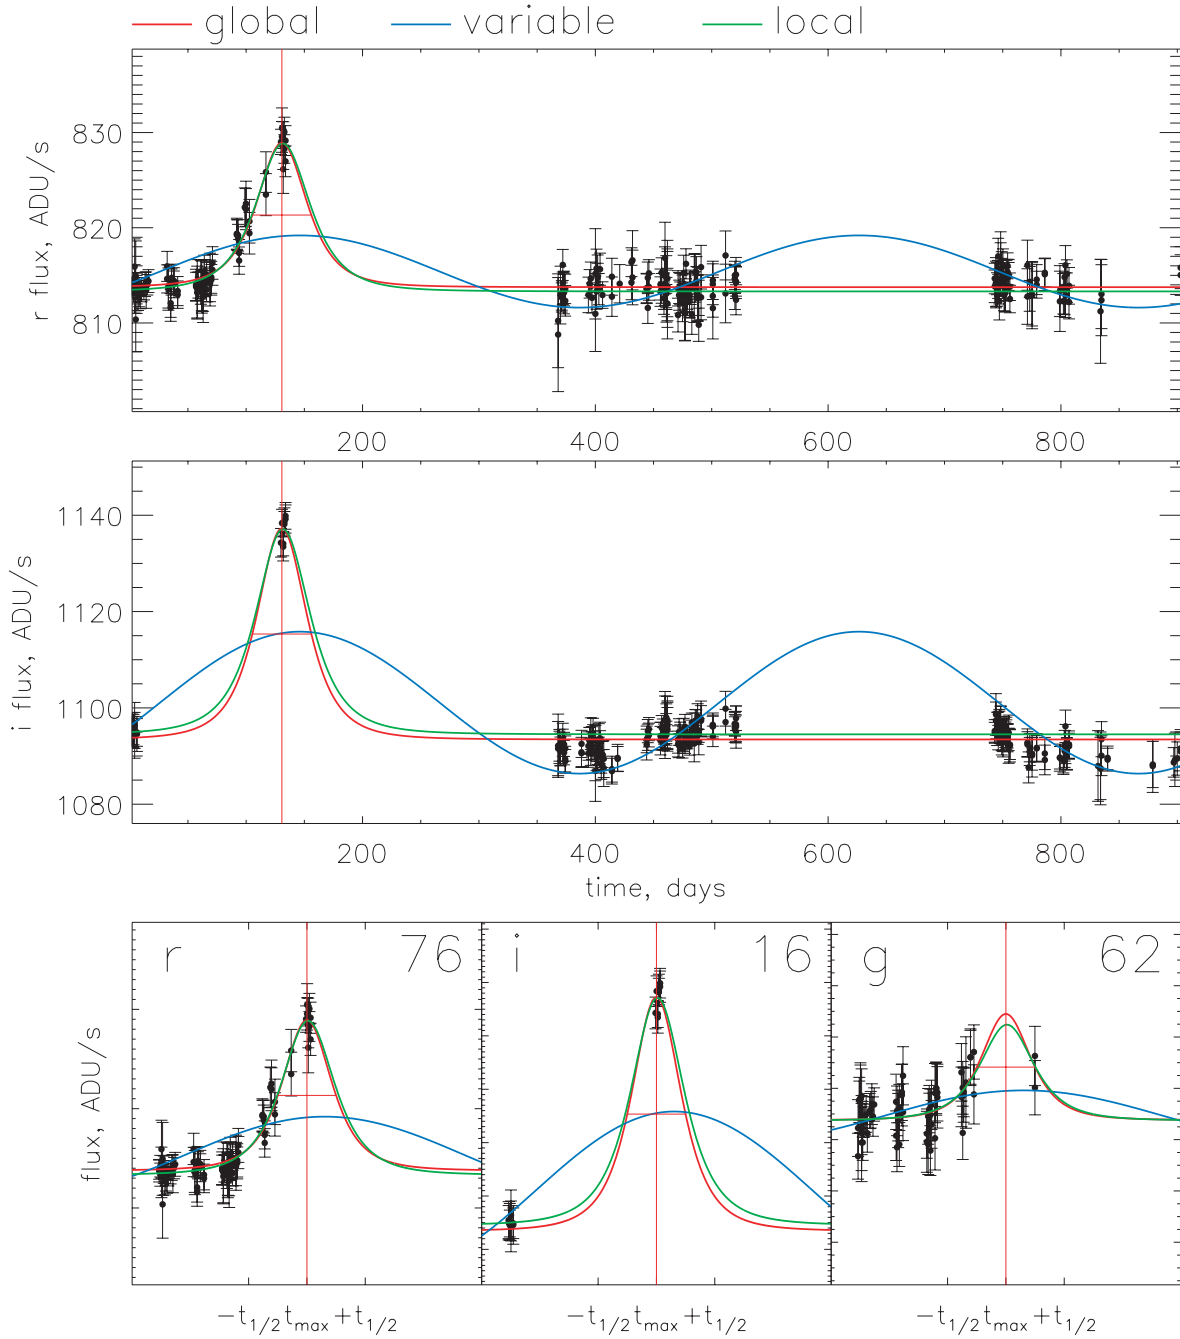χ^2(var) =1374.8 $t_{1/2}$ = 50.72±   5.40

Figure 15. Light curve of second-level candidate 3 in Table 4 from the southern field, CCD 3.

Table 6. The locations of the 16 third-level light curves in right ascension and declination (J2000.0), together with the projected distance R from the centre of M31.

Light curve	Field/CCD	R (arcmin)	RA	Dec.
1	F1,CCD1	12.91	$00^h\,43^m\,52\overset{s}{.}9$	$41°\,16'\,59''$
2	F1,CCD1	14.19	$00^h\,43^m\,59\overset{s}{.}2$	$41°\,18'\,06''$
3	F1,CCD2	7.12	$00^h\,42^m\,20\overset{s}{.}8$	$41°\,21'\,42''$
4	F1,CCD2	5.56	$00^h\,42^m\,30\overset{s}{.}3$	$41°\,21'\,01''$
5	F1,CCD4	26.14	$00^h\,44^m\,25\overset{s}{.}6$	$41°\,34'\,05''$
6	F1,CCD4	25.63	$00^h\,44^m\,03\overset{s}{.}6$	$41°\,37'\,01''$
7	F1,CCD4	24.93	$00^h\,44^m\,02\overset{s}{.}6$	$41°\,36'\,17''$
8	F1,CCD4	14.86	$00^h\,43^m\,26\overset{s}{.}9$	$41°\,28'\,40''$
9	F2,CCD2	18.79	$00^h\,41^m\,33\overset{s}{.}6$	$41°\,02'\,51''$
10	F2,CCD2	23.61	$00^h\,41^m\,20\overset{s}{.}1$	$40°\,58'\,39''$
11	F2,CCD3	6.06	$00^h\,43^m\,16\overset{s}{.}0$	$41°\,14'\,59''$
12	F2,CCD3	3.89	$00^h\,43^m\,00\overset{s}{.}4$	$41°\,13'\,41''$
13	F2,CCD3	4.44	$00^h\,42^m\,48\overset{s}{.}0$	$41°\,11'\,44''$
14	F2,CCD4	15.60	$00^h\,43^m\,23\overset{s}{.}6$	$41°\,02'\,24''$
15	F2,CCD4	20.25	$00^h\,43^m\,11\overset{s}{.}3$	$40°\,56'\,32''$
16	F2,CCD4	22.60	$00^h\,42^m\,44\overset{s}{.}4$	$40°\,53'\,32''$
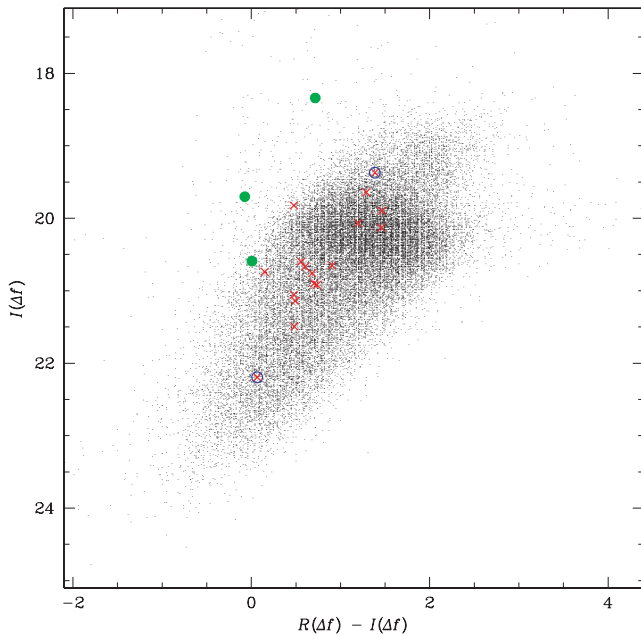
Figure 16. The locations of 20 events in the colour–magnitude diagram. The first-level candidates are denoted by green dots, the second-level candidates by red crosses within blue circles and third-level light curves by red crosses. (There are two light curves for which there is no *i*-band signal. Note also that the locations of the candidates are deduced using the microlensing fit, and so differ slightly from the corresponding locations in Fig. 3.)

A') and a set of weaker candidates (their 'set B'). The advantage of this is that the optical depth can then be studied as a function of the threshold of detection. We have followed their example by defining a set of first-level (or strong) candidates and second-level (or weaker) candidates. A detailed comparison of the numbers of first- and second-level microlensing candidates with theoretical predictions requires a careful calibration of our detection efficiency as a function of threshold and will be presented in a follow-up paper.

6 CONCLUSIONS

The POINT-AGAPE collaboration has presented the results of a fully automated search through a catalogue of 44 554 variable light curves built using the superpixel method from 3 yr of imaging data towards M31. In order to establish that the results obtained from microlensing experiments are robust, one should compare candidate samples derived from different sets of selection criteria. Our selection algorithm has no explicit restriction on time-scale. It is thus complementary to the work of Paulin-Henriksson et al. (2003, 2004) and Calchi Novati et al. (in preparation), who use a different algorithm that is fine-tuned for short, high-amplitude events. At the end of our automated search, we obtained a set of three first-level candidates, which are the most convincing. All of these events have short full-width at half-maximum time-scales. We also obtained a set of three second-level candidates, for which microlensing is a possible hypothesis.

We regard the three first-level light curves that pass all stages of the pipeline as excellent microlensing candidates. This includes the events PA 00-S3 and 00-S4 identified by Paulin-Henriksson et al. (2003), and a new candidate. We regard the three second-level candidates as less securely established. They need further analysis over a longer timeline before they can be convincingly accepted or rejected. We also showed a further 16 third-level light curves, which lie just below our threshold for designation as microlensing candidates.

The candidate selection algorithm proceeds by automatically performing fits to a blended microlensing event and a variable star template for every light curve. Only if the blended microlensing fit is preferred is the light curve retained for further cuts. This is a powerful way of discarding variable stars, but preserving long time-scale microlensing events. However, some short time-scale events with variable baselines (caused by contamination from nearby variables) are also thrown out. The selection algorithm also includes explicit cuts on achromaticity and on the location of events in an analogue of the colour–magnitude diagram. This is a useful way of discarding long-period variables and Miras, which given our sampling may fail to repeat during the 3-yr baseline. We note that some variable stars are excellent fits to blended microlensing events in two or three passbands. This is worrying for any experiment identifying long-time-scale microlensing events based on single passband data.

Our candidate selection algorithm is unrestricted, in the sense that it has no explicit threshold on time-scale or on flux variation. However, there is an implicit threshold on flux variation in the last cut, as the signal-to-noise ratio in microlensing must dominate over the signal-to-noise ratio in the residual variations. There are some obvious ways of extending the work in this paper. First, there is a fourth year of M31 imaging data taken by the MEGA collaboration, which is now publically available. It can provide additional discrimination against variable stars for our second-level candidates. Secondly, the efficiency of our survey must be calculated as a function of threshold to provide physically meaningful constraints on the baryon fraction of the M31 dark halo.

The great variety of light curves in pixel-lensing data implies that the optimum candidate selection algorithm may have a different structure, depending on the time-scale of the event. For example, for short time-scale events, the selection algorithm given above could be changed in the following manner. The data points associated with a candidate bump in the light curve are removed and a variable star template is fitted to the remainder. The first step in the selection then proceeds with a comparison between the goodness of fit of

a microlensing event with this bumpy baseline as compared with a variable star fit to the whole light curve. In other words, the two options are variable star and microlensing-contaminated by a nearby variable. If the latter is preferred, the light curve is retained and subjected to further analysis. It may also make sense to fine-tune the details of the cuts according to each field and CCD. The variable star contamination varies from CCD to CCD and therefore the optimum location of the cuts also changes.

The problem of identifying variable objects in pixel lensing is much more difficult than the corresponding problem in classical microlensing. Straight line cuts in the projected planes of multidimensional parameter spaces have been employed here and in other previous searches (e.g. Paulin-Henriksson et al. 2003; de Jong et al. 2004). However, the decision boundaries separating microlensing from non-microlensing events need not be as simple as hyper-planes in practice. It would be interesting to set more sophisticated pattern recognition algorithms, such as neural networks or self-organizing maps, on to this task (e.g. Belokurov, Evans & Le Du 2003, 2004). In this connection, we note that neural networks to hunt for the nova-like light curves in the POINT-AGAPE data set have already been successfully developed by Feeney (2004).

None the less, whatever microlensing candidate selection algorithm is used, the most important point is that the efficiency of the survey must be computable. Of course, different selection algorithms may have different efficiencies, but they must give consistent results for the characteristic mass and baryon fraction of the microlensing population. Here, the need for a well-defined efficiency has caused us to excise small portions of the data around bright resolved stars, which give rise to artefacts in the superpixel method. The fraction of the CCDs so masked is small, but this does mean that some genuine microlensing events have been excluded. In a future paper, we will compute the efficiency of this survey and present constraints on the dark matter content of the M31 halo.

REFERENCES

Alcock C. et al., 1997, ApJ, 486, 697
Alcock C. et al., 2000, ApJ, 542, 281
An J. H. et al., 2004a, ApJ, 601, 845
An J. H. et al., 2004b, MNRAS, 351, 1071
Ansari R. et al., 1997, A&A, 324, 843
Aubourg E. et al., 1995, A&A, 301, 1
Aurière M. et al., 2001, ApJ, 553, L137
Belokurov V., Evans N. W., Le Du Y., 2003, MNRAS, 341, 1373
Belokurov V., Evans N. W., Le Du Y., 2004, MNRAS, 352, 233
Calchi Novati S., Jetzer P., Scarpetta G., Giraud-Héraud Y., Kaplan J., Paulin-Henriksson S., Gould A., 2003, A&A, 405, 851
Crotts A. P. S., 1992, ApJ, 399, L43
Darnley M. et al., 2004, MNRAS, 353, 571
de Jong J. T. A. et al., 2004, A&A, 417, 461
Feeney S. M., 2004, M. Phil. thesis, Univ. Cambridge
Girardi L., Salaris M., 2001, MNRAS, 323, 109
Kerins E. et al., 2001, MNRAS, 323, 13
Kerins E. et al., 2003, ApJ, 598, 993
Lupton R., 1993, Statistics in Theory and Practice. Princeton Univ. Press Princeton
Paulin-Henriksson S., 2002, PhD thesis, College de France
Paulin-Henriksson S. et al., 2002, ApJ, 576, L121
Paulin-Henriksson S. et al., 2003, A&A, 405, 15
Paulin-Henriksson S. et al., 2004, Proc. of 'Rencontres de Moriond 2004' (astro-ph/0408204)
Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 1992, Numerical Recipes in Fortran 77, 2nd edn. Cambridge Univ. Press, Cambridge
Riffeser A. et al., 2001, A&A, 379, 362
Riffeser A., Fliri J., Bender R., Seitz S., Gössl C. A., 2003, ApJ, 599, L17
Sterken C., Jaschek C., 1996, Lightcurves of Variable Stars. Cambridge Univ. Press, Cambridge
Walterbos R., Kennicutt R., 1987, A&AS, 69, 311
Woźniak P. R., Paczyński B., 1997, ApJ, 487, 55

APPENDIX A: THE PHOTOMETRY TRANSFORMATION

In the main text, we frequently convert from Sloan to Cousins magnitudes, using equations derived by Paulin-Henriksson (2002). For the data taken in 1999, the transformations are

$$V = \Gamma_V + g + \beta_V(g - r)$$
$$R = \Gamma_{R_1} + r + \beta_{R_1}(g - r). \tag{A1}$$

For the data taken in 2000–2001, the transformations are

$$R = \Gamma_{R_2} + r + \beta_{R_2}(r - i)$$
$$I = \Gamma_I + i + \beta_I(r - i). \tag{A2}$$

The photometric calibration and colour transformation is worked out independently for each CCD. The coefficients are listed in Table A1.

Table A1. The coefficients in the transformations from Sloan to Cousins magnitudes according to each field and CCD.

| Field	CCD	Γ_V	Γ_{R_1}	Γ_{R_2}	Γ_I	β_V	β_{R_1}	β_{R_2}	β_I
Northern	1	24.64	24.33	24.30	23.90	−0.60	−0.18	−0.28	−0.20
Northern	2	24.63	24.32	24.27	23.88	−0.61	−0.21	−0.27	−0.15
Northern	3	24.81	24.61	24.35	23.60	−0.71	−0.29	−0.46	−0.23
Northern	4	24.74	24.49	24.22	23.53	−0.69	−0.30	−0.54	−0.40
Southern	1	24.61	24.21	24.26	24.04	−0.50	−0.06	−0.28	−0.25
Southern	2	24.59	24.26	24.24	23.88	−0.65	−0.19	−0.24	−0.12
Southern	3	24.59	24.29	24.26	23.98	−0.61	−0.17	−0.23	−0.17
Southern	4	24.61	24.29	24.27	23.96	−0.60	−0.20	−0.22	−0.13
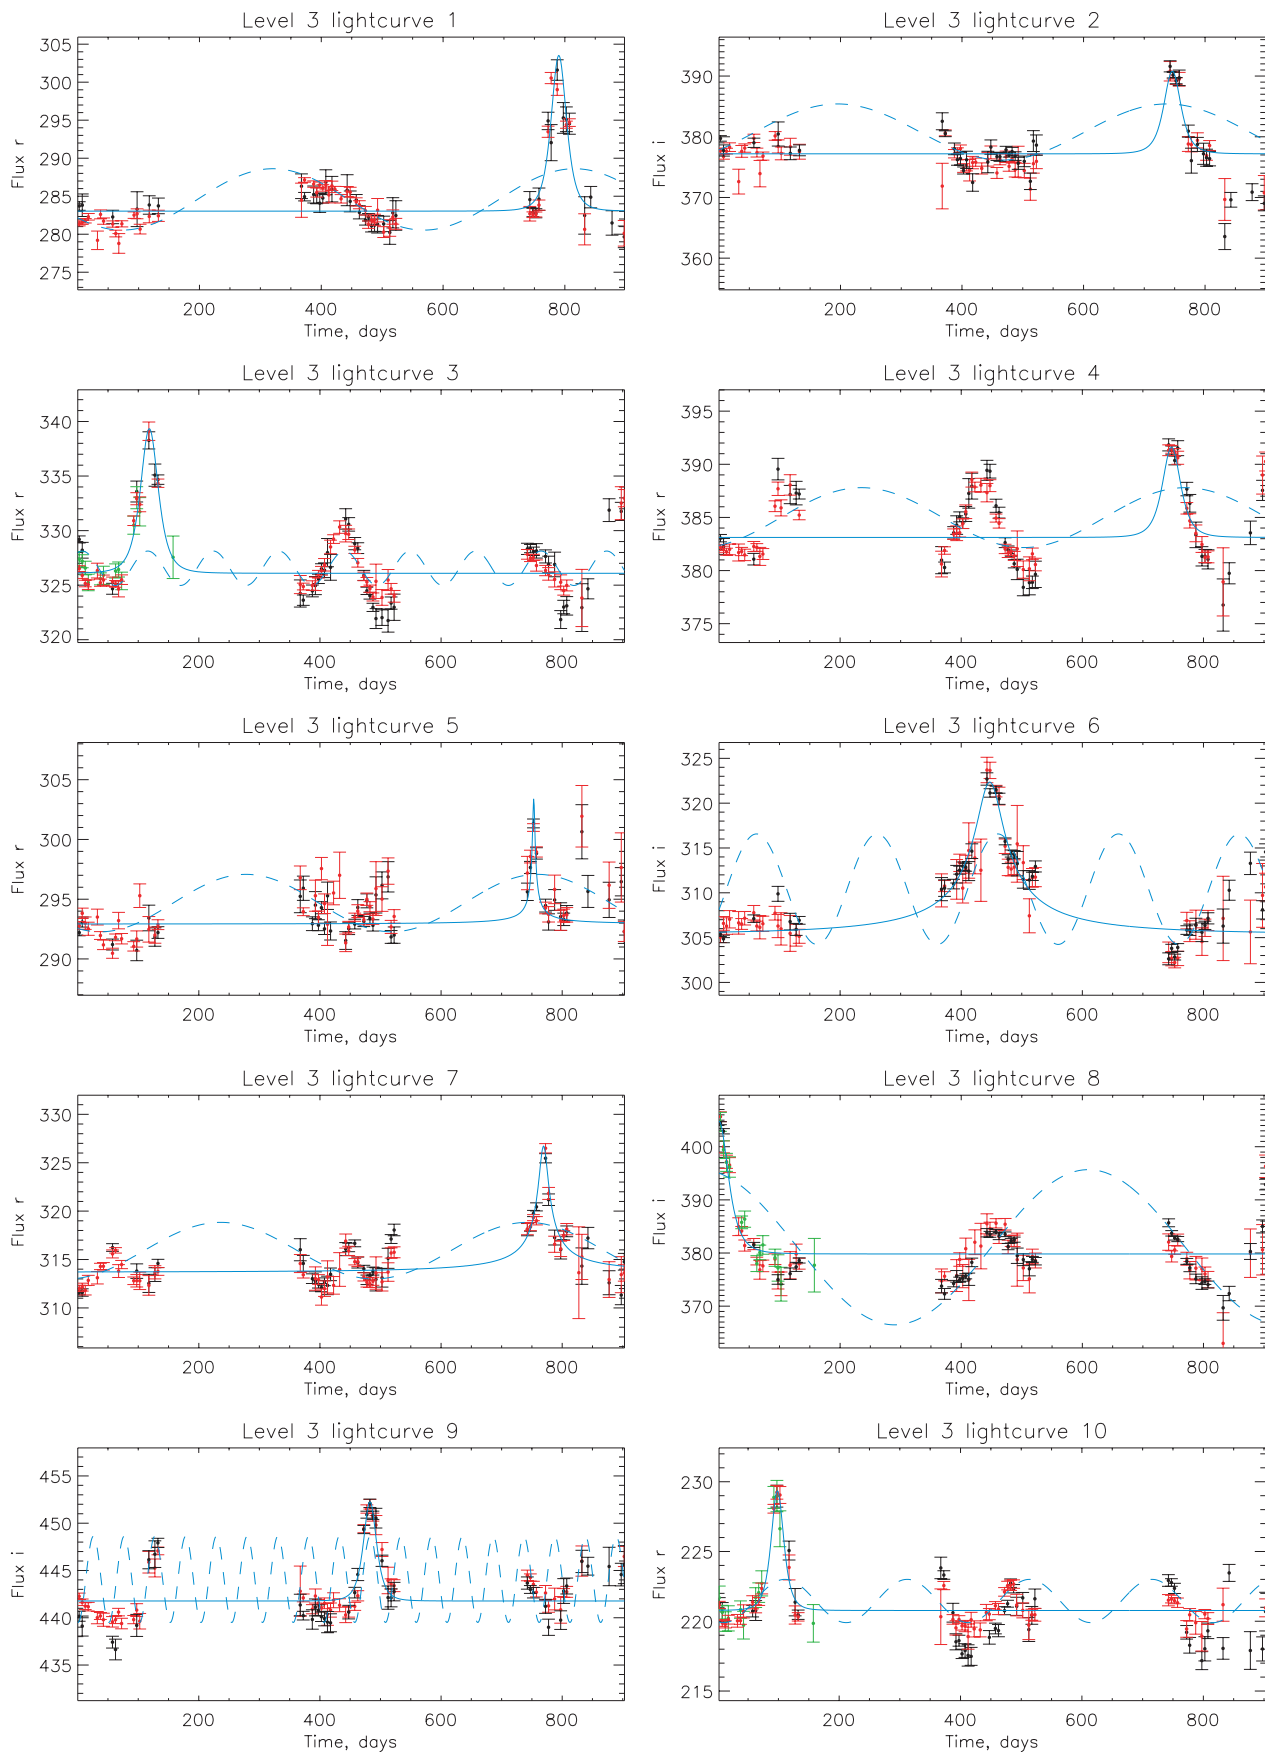
Figure B1. The third-level light curves. These just fail our threshold for designation as microlensing candidates. Data from all three bands are overplotted using the convention: *g* band as green, *r* band as red and *i* band as black. Also shown are the variable fit (blue-dashed) and the global microlensing fit (unbroken blue).
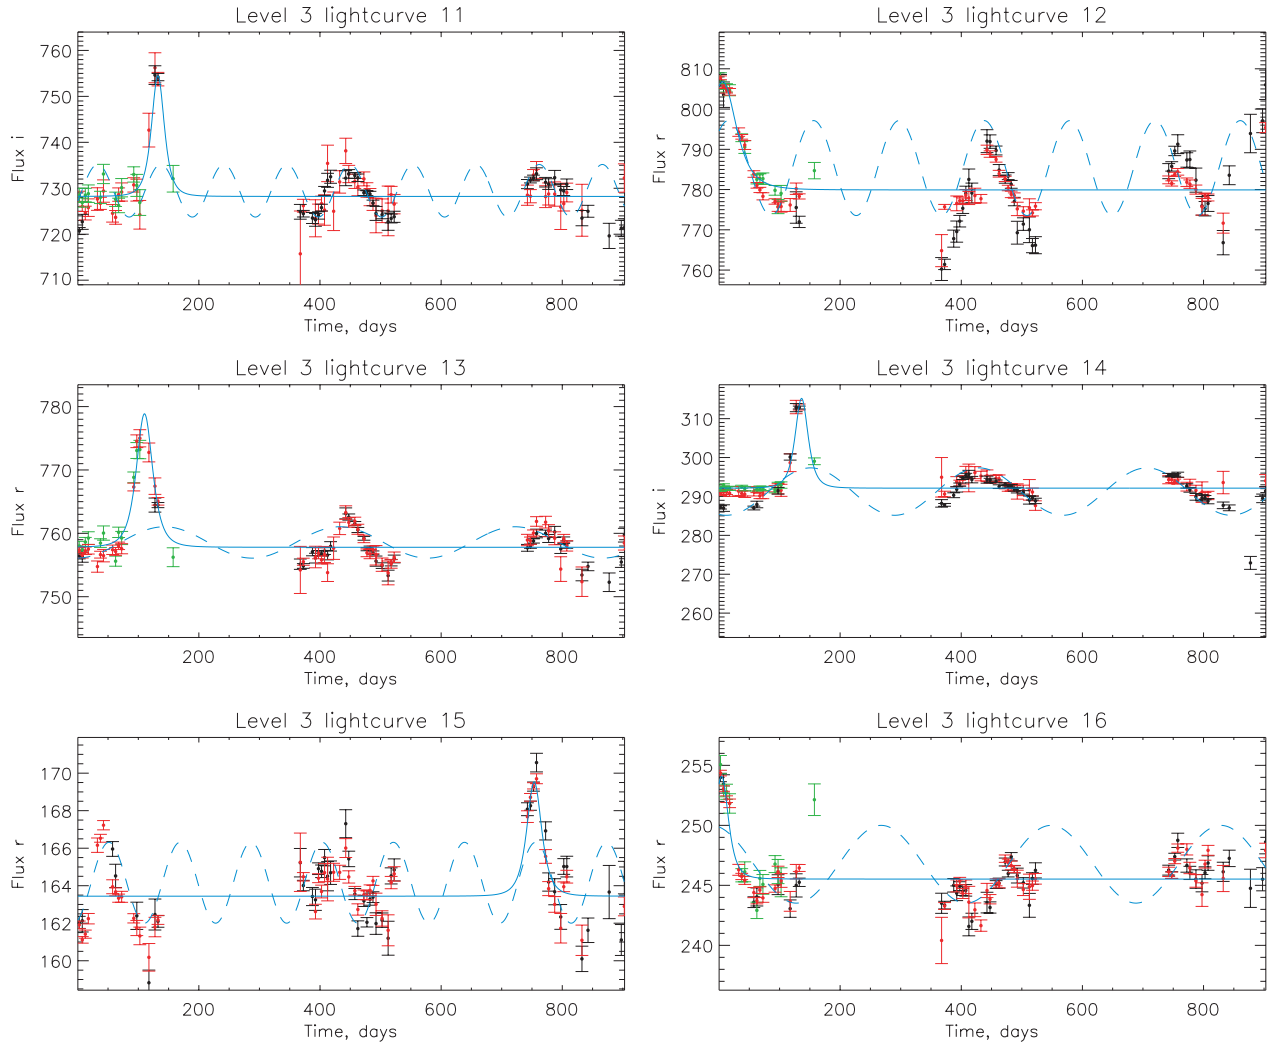
Figure B1 *– continued*

APPENDIX B: THE 16 THIRD-LEVEL LIGHT CURVES

The light curves of the 16 third-level events are displayed in Fig. B1.

This paper has been typeset from a TEX/LATEX file prepared by the author.