Realisation of a fully-deterministic microlensing observing strategy for inferring planet populations*

M. Dominik[1⋆⋆◦], **U. G. Jørgensen**[2,3], **N. J. Rattenbury**[4], **M. Mathiasen**[2], **T. C. Hinse**[2,5], **S. Calchi Novati**[6,7], **K. Harpsøe**[2], **V. Bozza**[6,7], **T. Anguita**[8], **M. J. Burgdorf**[9,10], **K. Horne**[1], **M. Hundertmark**[11], **E. Kerins**[12], **P. Kjærgaard**[2], **C. Liebig**[8], **L. Mancini**[6,7], **G. Masi**[13], **S. Rahvar**[14], **D. Ricci**[15], **G. Scarpetta**[6,7], **C. Snodgrass**[16], **J. Southworth**[17], **R. A. Street**[18], **J. Surdej**[14], **C. C. Thöne**[2,19], **Y. Tsapras**[18,20], **J. Wambsganss**[8], and **M. Zub**[8]

[1] SUPA, University of St Andrews, School of Physics & Astronomy, North Haugh, St Andrews, KY16 9SS, United Kingdom
[2] Niels Bohr Institutet, Københavns Universitet, Juliane Maries Vej 30, 2100 København Ø, Denmark
[3] Centre for Star and Planet Formation, Københavns Universitet, Øster Voldgade 5-7, 1350 København Ø, Denmark
[4] Department of Physics, The University of Auckland, Private Bag 92019, New Zealand
[5] Armagh Observatory, College Hill, Armagh, BT61 9DG, United Kingdom
[6] Università degli Studi di Salerno, Dipartimento di Fisica "E.R. Caianiello", Via Ponte Don Melillo, 84085 Fisciano (SA), Italy
[7] INFN, Gruppo Collegato di Salerno, Sezione di Napoli, Italy
[8] Astronomisches Rechen-Institut, Zentrum für Astronomie der Universität Heidelberg (ZAH), Mönchhofstr. 12-14, 69120 Heidelberg, Germany
[9] Deutsches SOFIA Institut, Universität Stuttgart, Pfaffenwaldring 31, 70569 Stuttgart, Germany
[10] SOFIA Science Center, NASA Ames Research Center, Mail Stop N211-3, Moffett Field CA 94035, United States of America
[11] Institut für Astrophysik, Georg-August-Universität, Friedrich-Hund-Platz 1, 37077 Göttingen, Germany
[12] Jodrell Bank Centre for Astrophysics , University of Manchester Alan Turing Building, Manchester, M13 9PL, United Kingdom
[13] Bellatrix Astronomical Observatory, Via Madonna de Loco 47, 03023 Ceccano (FR), Italy
[14] Department of Physics, Sharif University of Technology, P. O. Box 11155–9161, Tehran, Iran
[15] Institut d'Astrophysique et de Géophysique, Allée du 6 Août 17, Sart Tilman, Bât. B5c, 4000 Liège, Belgium
[16] Max Planck Institute for Solar System Research, Max-Planck-Str. 2, 37191 Katlenburg-Lindau, Germany
[17] Astrophysics Group, Keele University, Staffordshire, ST5 5BG, United Kingdom
[18] Las Cumbres Observatory Global Telescope Network, 6740B Cortona Dr, Goleta, CA 93117, United States of America
[19] INAF, Osservatorio Astronomico di Brera, 23806 Merate (LC), Italy
[20] School of Mathematical Sciences, Queen Mary, University of London, London, E1 4NS, United Kingdom

The dates of receipt and acceptance should be inserted later

Within less than 15 years, the count of known planets orbiting stars other than the Sun has risen from none to more than 400 with detections arising from four successfully applied techniques: Doppler-wobbles, planetary transits, gravitational microlensing, and direct imaging. While the hunt for twin Earths is on, a statistically well-defined sample of the population of planets in all their variety is required for probing models of planet formation and orbital evolution so that the origin of planets that harbour life, like and including ours, can be understood. Given the different characteristics of the detection techniques, a complete picture can only arise from a combination of their respective results. Microlensing observations are well-suited to reveal statistical properties of the population of planets orbiting stars in either the Galactic disk or bulge from microlensing observations, but a mandatory requirement is the adoption of strictly-deterministic criteria for selecting targets and identifying signals. Here, we describe a fully-deterministic strategy realised by means of the ARTEMiS (Automated Robotic Terrestrial Exoplanet Microlensing Search) system at the Danish 1.54m telescope at ESO La Silla between June and August 2008 as part of the MiNDSTEp (Microlensing Network for the Detection of Small Terrestrial Exoplanets) campaign, making use of immediate feedback on suspected anomalies recognized by the SIGNALMEN anomaly detector. We demonstrate for the first time the feasibility of such an approach, and thereby the readiness for studying planet populations down to Earth mass and even below, with ground-based observations. While the quality of the real-time photometry is a crucial factor on the efficiency of the campaign, an impairment of the target selection by data of bad quality can be successfully avoided. With a smaller slew time, smaller dead time, and higher through-put, modern robotic telescopes could significantly outperform the 1.54m Danish, whereas lucky-imaging cameras could set new standards for high-precision follow-up monitoring of microlensing events.

1 Introduction

The bending of light received from stars due to the gravity of intervening foreground stars has been discussed by Einstein already in 1912 (Renn, Sauer & Stachel 1979), three years before he properly described it by means of the theory of General Relativity (Einstein 1915), and shortly after he suggested to measure the bending of light rays grazing the limb of the Sun (Einstein 1911). Despite the success of Einstein's theory on the deflection by the Sun (Dyson et al. 1920), for other stars he concluded that "there is no great chance of observing this phenomenon" (Einstein 1936). In fact, it required several decades of advance in technology until the observation of such a *gravitational microlensing event* became a reality (Alcock et al. 1993).

Effects by intervening planets have already been considered by Liebes (1964), concluding that "the primary effect of planetary deflectors bound to stars other than the sun would be to slightly perturb the lens action of these stars". This has then been discussed by Mao & Pacyzyński (1991), and in 2004, gravitational microlensing joined radial-velocity Doppler-wobble surveys (Mayor & Queloz 1995) and planetary transit observations (Henry et al. 2000; Charbonneau et al. 2000; Udalski et al. 2002) as a successful technique for detecting planets orbiting stars other than the Sun (Bond et al. 2004).

An overwhelming majority of the more than 400 extra-solar planets[1] identified to date are quite unlike anything in the Solar system. The recent years, however, have seen the subsequent detections of more and more Earth-like planets (e.g. Rivera et al. 2005; Beaulieu et al. 2006; Udry et al. 2007, Marois et al. 2008) and first look-alikes of the the Solar system (e.g. Marcy et al. 2002; Gaudi et al. 2008; Marois et al. 2008) have also been discovered. This raises the chances for detecting a true sibling of our home planet, and one might even find evidence for life elsewhere. However, such a pathway towards habitable planets leaves an even deeper question unanswered: What is the origin of habitable planets and that of Earth in particular, with all their life forms? Rather than narrowing down the focus to habitable planets, it is the distribution of an as large as possible variety that provides a powerful test of models of planet formation and orbital evolution.

Amongst all the techniques that have been proposed for studying extra-solar planets, there is no single one superior to the others in every aspect, but instead, the different approaches each have a different focus, and complement each other nicely towards the full picture on planet populations, where some overlap in the sensitivity to regions of planet parameter space provides the opportunity to compare results and thereby check on whether the respective selection biases have been properly understood.

Gravitational microlensing favours the detection of planets in orbits of a few AU around M-, K-, and to a lesser extent G-dwarf stars at several kpc distance either in the Galactic disk and bulge – rather than in the Solar Neighbourhood – (e.g. Dominik et al. 2008c), and capable to even reveal planets orbiting stars in other galaxies, such as M31 (Covone et al. 2000; Chung et al. 2006; Ingrosso et al. 2009). Planet population statistics arising from microlensing observations are expected to shed light on a possible mass gap between super-Earths and gas giants, as well as on a predicted cut-off towards larger orbits, which provides a measure of the stellar disk surface density as well as the accretion rate of planetesimals, which constitute fundamental and crucial parameters for the underlying theories (Ida & Lin 2005). The technique of gravitational microlensing is a serious competitor in the race for the detection of an Earth-mass planet, and Earth mass does not constitute the limit of current efforts (Dominik et al. 2007). In fact, Paczyński (1996) already pointed out that in principle objects as small as the Moon could be detected, and he was not even talking about space-based observations in this context.

Given that the observed distribution of planets is the product of the underlying population and the detection efficiency of the experiment, simply counting the observed systems will not provide an appropriate picture of the planet statistics. Instead, the detection efficiency needs to be determined, quantifying what could have been detected and what has been missed. It is crucial that the detection criteria applied to the actual detections are identical to those used for determining the detection efficiency (e.g. Calchi Novati et al. 2009). First attempts to calculate the detection efficiency of microlensing observations (Gaudi & Sackett 2000) adopted χ^2 offsets as criterion for a detection, but this hardly relates to how the reported planets were actually revealed.

Studying the planet populations therefore calls for strictly-deterministic procedures for planet detection, while in contrast, human decisions by their unpredictability and irreproducibility are the enemy of deriving meaningful statistics. The ARTEMiS (Automated Robotic Terrestrial Exoplanet Microlensing Search) expert system[2] (Dominik et al. 2008a,b) enables to take the step from detecting planets to inferring their population statistics by not only providing an automated selection of targets, following earlier work by Horne, Snodgrass & Tsapras (2009), but also the automated identification of ongoing anomalies by means of the SIGNALMEN anomaly detector (Dominik et al. 2007).

Focussing on observations carried out by the MicroFUN team[3], Gould et al. (2010) have recently estimated the frequency of planetary systems with gas-giant planets beyond the snow line similar to the Solar system from high-magnification microlensing events. Looking at the planet detection efficiency as function of planet mass, they found that planets below $\sim$ 10 $M_\oplus$ are *not* best detected by focusing the ef-

⋆ Based on data collected by the MiNDSTEp consortium with the Danish 1.54m telescope at the ESO La Silla Observatory

⋆⋆ Royal Society University Research Fellow

° E-mail: md35@st-andrews.ac.uk

[1] http://exoplanet.eu

[2] http://www.artemis-uk.org

[3] http://www.astronomy.ohio-state.edu/~microfun

forts on the few high-magnification peaks, but rather by monitoring the lower-magnification wings of as many as possible events. Moreover, they stress the need for what they call a "controlled experiment", free from human intervention that would disturb the statistics. In fact, such campaigns, favourable to studying low-mass planets, are being piloted by the MiNDSTEp (Microlensing Network for the Detection of Small Terrestrial Exoplanets)[4] and RoboNet-II[5] (Tsapras et al. 2009) teams.

While a "controlled experiment" for inferring the statistics of an underlying population is a well-known concept in particicle physics, it is a rather unusual approach in astronomy. We are however approaching an era of large synoptic surveys, e.g. Gaia[6], Pan-STARRS[7], SkyMapper[8], or LSST[9], which will identify transient phenomena of different origin, and prompt follow-up observations on these. In particular, the Gaia Science Alerts Working Group[10] explicitly considers supernovae, microlensing events, exploding and eruptive stars. Whenever the opportunities exceed the available resources, an automated target prioritisation is required.

The framework for optimizing an observing campaign with regard to the return and required investment, as discussed in this paper, is a quite general concept, not at all restricted to the monitoring of gravitational microlensing events with the goal to infer planet population statistics. Moreover, it sets the demands for automated telescope scheduling within a heterogeneous, non-proprietary network (e.g. Steele et al. 2002; Allan et al. 2006; Hessman 2006), for which the study of planets by microlensing is a showcase application.

Here, we would like to discuss the specific implementation of a fully-deterministic ARTEMiS-assisted observing strategy at the 1.54m Danish telescope at ESO La Silla (Chile) as part of the 2008 MiNDSTEp campaign, and by providing a critical analysis of our operations, identifying challenges, and deriving suggestions for further improvements to share our experience. Moreover, there has been some scepticism on whether the theoretical foundations described by Dominik et al. (2007) for a strategy involving automated anomaly detection, which has been shown to extend the sensitivity to planets further down in mass by factors between 3 and 5 (Dominik 2008), can be turned into something that works in practice. We therefore address this request from the scientific community in substantial detail. For making scientific and technological progress, it is of crucial importance to properly document how experiments with a devised strategy turned out to actually work (or not to work).

[4] http://www.mindstep-science.org
[5] http://robonet.lcogt.net
[6] http://www.esa.int/science/gaia
[7] http://pan-starrs.ifa.hawaii.edu
[8] http://www.mso.anu.edu.au/skymapper
[9] http://www.lsst.org
[10] http://www.ast.cam.ac.uk/research/gsawg

In Sect. 2, we review the basics of revealing the presence of planets by gravitational microlensing, before we describe the adopted strategy for inferring planet population statistics by means of the MiNDSTEp observations in Sect. 3. Sect. 4 reports how this strategy is implemented by means of interaction of the observer with the ARTEMiS system, while Sect. 5 analyses how the strategy goals emerged into real acquired data. Finally, Sect. 6 presents conclusions and recommendations for improvements arising from the experience gained during 2008.

2 Planetary microlensing

The 'most curious' effect of gravitational microlensing (Einstein 1936) involves a single characteristic physical scale, namely the *angular Einstein radius*

$$\theta_{\rm E} = \sqrt{\frac{4GM}{c^2}\,\left(D_{\rm L}^{-1} - D_{\rm S}^{-1}\right)}\,, \tag{1}$$

characterized by the mass of the intervening foreground 'lens' star at distance $D_{\rm L}$, whose gravitational field yields a characteristic brightening of the observed background 'source' star at distance $D_{\rm S}$, which reads

$$A(u) = \frac{u^2+2}{u\,\sqrt{u^2+4}}\,, \tag{2}$$

where $u\,\theta_{\rm E}$ is the angular separation between lens and source star, which simply follows from geometry with the deflection angle

$$\alpha(\xi) = \frac{4GM}{c^2\,\xi} \tag{3}$$

for a light ray with impact parameter ξ. More precisely, $A(u)$ results as the combined magnification of two images at $x_\pm\,\theta_{\rm E}$, where

$$x_\pm = \frac{1}{2}\left(u \pm \sqrt{u^2+4}\right)\,, \tag{4}$$

which are too close to each other to be resolved (see the estimate on $\theta_{\rm E}$ below).

If one assumes a constant proper motion μ, the angular separation becomes a function of time (Refsdal 1964; Paczyński 1986)

$$u(t) = \sqrt{u_0^2 + \left(\frac{t-t_0}{t_{\rm E}}\right)^2}\,, \tag{5}$$

where the event time-scale is given by $t_{\rm E} = \theta_{\rm E}/\mu$, while $u_0\,\theta_{\rm E}$ is the closest angular approach occuring at epoch t_0. With $F_{\rm S}$ denoting the flux of the observed source star, and $F_{\rm B}$ denoting the flux of any background within the potentially unresolved target on the detector, the total observed flux reads

$$F(t) = F_{\rm S}\,A(t) + F_{\rm B} = F_{\rm base}\,\frac{A(t)+g}{1+g}\,, \tag{6}$$

with $F_{\rm base} = F_{\rm S} + F_{\rm B}$ and the blend ratio $g = F_{\rm B}/F_{\rm S}$. Consequently, the effective observed magnification becomes

$$A_{\rm obs}(t) = \frac{A(t)+g}{1+g}\,, \tag{7}$$

while the observed magnitude is given by

$$m(t) = m_{\rm base} - 2.5 \lg A_{\rm obs}(t) \,, \tag{8}$$

where $m_{\rm base}$ denotes the observed baseline magnitude and

$$m_{\rm S} = m_{\rm base} + 2.5 \lg(1+g) \tag{9}$$

is the magnitude of the source star.

If for observations towards the Galactic Bulge (Paczyński 1991; Kiraga & Paczyński 1994), one assumes typical distances of $D_{\rm S} = 8.5$ kpc and $D_{\rm L} = 6.5$ kpc for the source and lens star (of mass M), respectively,[11] the angular Einstein radius becomes

$$\theta_{\rm E} \sim 300\ \mu{\rm as} \left(\frac{M}{0.3\ M_\odot} \right)^{1/2} , \tag{10}$$

where $M \sim 0.3\ M_\odot$ corresponds to a rather typical M-dwarf star, which corresponds to a length

$$r_{\rm E} = D_{\rm L}\, \theta_{\rm E} \sim 2\ {\rm AU} \left(\frac{M}{0.3\ M_\odot} \right)^{1/2} , \tag{11}$$

at the distance of the lens star, while for a typical proper motion $\mu \sim 15\ \mu{\rm as\,d}^{-1}$, the event time-scale is

$$t_{\rm E} \sim 20\ {\rm d} \left(\frac{M}{0.3\ M_\odot} \right)^{1/2} . \tag{12}$$

With only one in a million observed stars being magnified by more than 30 per cent at a given time, the detection of microlensing events requires a substantial monitoring programme. Nowadays, the OGLE (Optical Gravitational Lensing Experiment)[12] and MOA (Microlensing Observations in Astrophysics)[13] surveys regularly monitor $\gtrsim 10^8$ stars each night, resulting in almost 1000 events being detected each year. Real-time data reduction systems allow for ongoing events, along with photometric data, being reported promptly to the scientific community (Udalski 2003; Bond et al. 2001).

Following the considerations by Liebes (1964), one expects planets to reveal their presence by causing small perturbations to the otherwise symmetric light curves of stars. If a planet were an isolated object of mass $M_{\rm p}$, the planetary deviation would last

$$t_{\rm p} = (M_{\rm p}/M)^{1/2}\, t_{\rm E} \,, \tag{13}$$

which evaluates to about a day for a Jupiter-mass and about 1.5 hrs for an Earth-mass planet. However, as Mao & Paczyński (1991) found, the tidal field of the planet's host star at the position of the planet not only substantially increases the planet detection probability, but also increases the signal duration by a fair factor. In fact, a 'resonance' occurs if the angular separation $\theta_{\rm p}$ of the planet from its host star becomes comparable to the angular Einstein radius $\theta_{\rm E}$, which makes gravitational microlensing most sensitive to planets orbiting at a few AU. In any case, due to the finite radius $R_\star$ of the observed source star, the signal duration cannot fall below

$$2\, t_\star = \frac{2\, R_\star}{D_{\rm S}\, \mu} \sim 2\ {\rm h} \left(\frac{R_\star}{R_\odot} \right) . \tag{14}$$

While the finite source size leads to spreading of the signal over a longer duration, it also reduces the signal amplitude and imposes a limit on it. As long as the source star can be approximated as point-like, the signal amplitude can reach any level regardless of the planet's mass, while just the signal duration and the probability for a signal to occur decreases towards smaller masses.

The need for both a huge number of stars to be monitored and a dense sampling to be achieved for being able to detect the planetary signatures lead to the adoption of a three-step strategy of survey, follow-up, and anomaly monitoring (Elachi et al. 1996; Dominik et al. 2007). A first microlensing follow-up network for planet detection that allows round-the-clock monitoring with hourly sampling was put into place in 1995 by the PLANET (Probing Lensing Anomalies NETwork) collaboration[14] (Albrow et al. 1998; Dominik et al. 2002). Fig. 1 shows the model light curve of the event OGLE-2005-BLG-390 together with data collected at 6 different sites. A planetary 'blip' on the night starting on 10 August 2005 showed a signature of a planet of 3-10 Earth masses (Beaulieu et al. 2006; Dominik, Horne & Bode 2006). Going for the next substantial science goals required an upgrade from the PLANET operations, and the case of event OGLE-2005-BLG-390 shows explicitly that with an immediate dense monitoring (of 10 min cadence) on the suspicion of a deviation from an ordinary light curve, an Earth-mass planet in the same spot could have been revealed. This pushed the timely development of the SIGNALMEN anomaly detector (Dominik et al. 2007), and subsequently its integration into the ARTEMiS system (Dominik et al. 2008a,b).

The number of ongoing microlensing events reported by the microlensing surveys at some point exceeded the number of targets that can be monitored by follow-up networks with the intended time sampling. This has prompted the need to select the most promising targets at any time. Griest & Safizadeh (1998) pointed out that the planet detection probability depends on the current angular separation between lens and source star, and thereby on the current event magnification. Based on this, Horne et al. (2009) have devised a priority algorithm, primarily for use with robotic telescopes. Vermaak (2000) and Han (2007) have also studied the roles of blending and the size of the source star on the target selection.

3 The 2008 MiNDSTEp strategy

The central aim of ARTEMiS is to enable the realization of an optimal planet search strategy by serving each telescope with the most promising target to be observed at any

[11] Full probability density distributions can be obtained by means of a proper model of the Milky Way (e.g. Dominik 2006).

[12] `http://ogle.astrouw.edu.pl`

[13] `http://www.phys.canterbury.ac.nz/moa/`

[14] `http://www.planet-legacy.org`

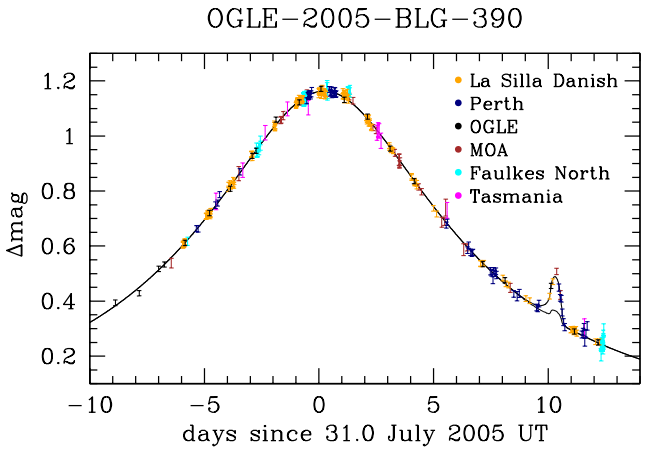

Fig. 1 Model light curve for microlensing event OGLE-2005-BLG-390 (Beaulieu et al. 2006; Dominik et al. 2006), along with the data acquired by the PLANET/Robonet, OGLE, and MOA collaborations. The $\sim$ 15 per cent blip, lasting about a day, revealed the 5-Earth-mass (with a factor 2 uncertainty) planet OGLE-2005-BLG-390Lb. The thinner line refers to the hypothetical detectable 3 per cent deviation that would have arisen from an Earth-mass planet in the same spot.

given time. However, a single global optimal strategy does not exist (Dominik 2008). Instead, the optimal choice of targets critically depends on 4 different kinds of input: the capabilities of potential observing sites, the science goals, the currently available data, and the current observability. Moreover, the science goals themselves can be manifold: one might want to (1) maximize the number of planets detected, (2) determine their abundance, or (3) obtain a census of their properties. In addition, the specific choice of strategy will depend on the kind of planets one would like to focus on. Actually, there is no need for any adopted strategy to be 'optimal' (whatever metric is used to measure this). Instead, the requirement for a successful strategy is to be reasonable and deterministic, where the determinism eliminates the interference of the planet abundance statistics with human judgement, and allows to derive proper results by means of Monte-Carlo simulations applying the defined criteria.

It is in fact favourable to miss out on some fraction of the optimally achievable detection efficiency (which could be gained by human intervention) if that would corrupt the statistics that are needed as ingredient for drawing the right conclusions about the acquired data set. The provision of alerts and monitoring by both the OGLE and MOA surveys however involve human decisions, but such can be modelled from the a-posteriori statistics of the vast number of events that comprise the respective sample, whereas such an approach is not viable for the small number of planetary events and the decision process that led to the identification of the planet.

Based on some simple thoughts and experience gained during the course of their observations, PLANET arrived at a rather pragmatic observing strategy that may not be optimal in a mathematical sense, but appears to be a workable match for achieving the science goals of the campaign (Albrow et al. 1998; Dominik et al. 2002). The fundamental principle behind the adopted strategy is in demanding a fixed photometric accuracy of 1–2 per cent, and selecting an exposure time $t_{\rm exp}$ for each of the targets, so that this can be met. For a target flux F, the number of collected photons becomes $N \propto F\,t_{\rm exp}$, and if the photon-count noise $\sigma_N = \sqrt{N}$ dominates the measurement uncertainties, $\sigma_F/F = \sigma_N/N$, so that

$$\frac{\sigma_F}{F} = \frac{\kappa}{\sqrt{F\,t_{\rm exp}}}\,, \tag{15}$$

where κ denotes a proportionality constant. Appendix A presents an extension to the presented formalism, taking into account a systematic uncertainty, which will dominate at some point as the exposure time increases, and prevent σ_F/F from approaching zero as $t_{\rm exp} \to \infty$.

Rather than from a flux contrast σ_F/F, planets are to be detected from a contrast in the magnification σ_A/A (e.g. Gaudi & Sackett 2000; Vermaak 2000; Dominik et al. 2007; Han 2007), which frequently makes a substantial difference, given that the observed targets are generally blended with other stars (e.g. Smith et al. 2007). With $F_{\rm S} > 0$ denoting the flux of the source star, whose light is magnified by a factor A due to the gravitational bending by the lens star, and $F_{\rm B}$ being a background flux, the total observed flux reads $F = F_{\rm S}\,A + F_{\rm B}$. Therefore, one finds

$$\frac{\sigma_A}{A} = \frac{\sigma_F}{F}\,\frac{A+g}{A}\,, \tag{16}$$

where $g = F_{\rm B}/F_{\rm S}$.

The effort to achieve a fixed σ_A/A is proportional to $t_{\rm exp}^{-1}$ required for the target at current flux $F(t)$, for which we find by combining Eqs. (15) and (16)

$$(t_{\rm exp})^{-1} = \kappa^{-2}\,F\,\left(\frac{A}{A+g}\right)^2\,\left(\frac{\sigma_A}{A}\right)^2\,. \tag{17}$$

Dominik et al. (2007) noted the need to invest data-acquisition efforts into the provision of the ability to sufficiently predict the underlying ordinary light curve of the event monitored, since deviations can only be detected in real time against a known model, whereas all deviations smaller than the uncertainty of model prediction cannot be properly assessed. The SIGNALMEN anomaly detector appears to fix this issue in practice to some extent by suspecting a deviation on a departure from the model expectations and requesting further data to be taken, which not only can lead to the detection of an anomaly, but also leads to a more appropriate model constraint if indicated. Imperfect event prediction has been ignored in the priority algorithm proposed by Horne et al. (2009), realised as web-PLOP (Snodgrass et al. 2008), and used by RoboNet (Burgdorf et al. 2007; Tsapras et al. 2009), which is based on the assumption that the model parameters are exactly known. However, the data-acquisition patterns of current campaigns do not make this a good approximation for practical purposes, and the degree of predictability of ongoing events is a quite relevant

issue (Dominik 2009). Rather than thinking about increasing the prospects for detecting planets, it was the need for proper event prediction that initially prompted PLANET to monitor events more densely over their peaks, where model uncertainties can result in rather severe misprediction of the current magnification. It is definitely worth studying how to optimally sample a light curve in order to make it sufficiently predictable. However, at this stage, we have just adopted a simple model that is expected to do the job reasonably well.

The ability to claim the detection of even Earth-mass planets does not require a follow-up sampling interval of more than 90 min, because this leaves sufficient opportunity to properly characterize arising anomalies by means of immediate anomaly monitoring activated after the first suspicion by the SIGNALMEN anomaly detector (Dominik et al. 2007). Such a moderate sampling on the other hand allows for a large enough number of targets to be followed for arriving at a fair chance to detect such in practice. The simple choice of a sampling interval

$$\tau = 90\ \mathrm{min}\ \sqrt{\frac{3/\sqrt{5}}{A}} \tag{18}$$

does not produce sampling overload as events get brighter, and for unblended events even leaves room to use exposure times that lead to a more accurate photometry over the peak region. We moreover require minimal and maximal values $\tau_{\rm min} = 2$ min and $\tau_{\rm max} = 120$ min, respectively, thereby avoiding over- or undersampling of events. In order to simplify the construction of an observing sequence of targets, we chose to further discretize the sampling intervals to values $\tau \in \{2, 3, 4, 5, 7.5, 10, 15, 20, 30, 45, 60, 90, 120\}$ min.

For the effects of planets orbiting lens stars, Gould & Loeb (1992) recovered the two-mass-scale formalism brought forward by Chang & Refsdal (1979) in a different context (lensing of quasars by stars within a galaxy), finding that these are maximal if the planet happens to be in the vicinity of one of the images due to light bending by the star. Horne et al. (2009) found that the area of the resulting planet detection zone on the sky $S_{\rm p}$ scales with the current event magnification A approximately as $S_{\rm p} = 2A - 1$, which moreover consists of zones around either of the images at x_+ or x_-, so that $S_{\rm p} = S_{\rm p}^+ + S_{\rm p}^-$ with $S_{\rm p}^+ = 2\,A_+ - 1$ and $S_{\rm p}^- = 2\,A_-$. However, this would directly yield the planet detection probability only if planets were distributed uniformly on the sky around their host star. The distortions of the detection zones are not far from those of the images, so that a tangential probability density proportional to $1/x$, where x is the image position, reflects an isotropic distribution. Very little however is known about a good radial prior, but rather than probing parameter space uniformly in the orbital axis, we adopt (with a degree of arbitrariness) a logarithmic distribution, leading to a further factor $1/x$. Thereby, one arrives at a planet detection probability proportional to

$$\Psi_{\rm p} = \frac{2\,A_+ - 1}{x_+^2} + \frac{2\,A_-}{x_-^2}\,. \tag{19}$$

With Eq. (4), one finds

$$x_\pm^2 = \frac{1}{2}\left(u^2 + 2 \pm u\,\sqrt{u^2+4}\right)\,, \tag{20}$$

so that with

$$A_+ = \frac{A+1}{2}\,,\quad A_- = \frac{A-1}{2} \tag{21}$$

and Eq. (2), $\Psi_{\rm p}$ as a function of the lens-source separation parameter u evaluates to

$$\Psi_{\rm p}(u) = \frac{4}{u\,\sqrt{u^2+4}} - \frac{2}{u^2+2+u\,\sqrt{u^2+4}}\,. \tag{22}$$

In general, the prioritisation of microlensing events will meaningfully follow a respective gain factor $\Omega_{\rm s} \propto R/I$, where R denotes the return and I the investment. In accordance with the previous considerations, we choose $R = \Psi_{\rm p}(u)$ and $I \propto t_{\rm exp}\,\tau^{-1}$, where the increased sampling effort $\tau^{-1} \propto \sqrt{A}$ is seen as a burden required to keep the event predictable rather than any further opportunity for planet detection. With Eq. (6), this leads to

$$\Omega_{\rm s} = 0.614\,\frac{F_{\rm base}}{F_{18}}\,\frac{A^{3/2}\,\Psi_{\rm p}(u)}{(A+g)\,(1+g)}\,, \tag{23}$$

where

$$\frac{F_{\rm base}}{F_{18}} = 10^{-0.4(m_{\rm base}-18)} \tag{24}$$

with $m_{\rm base}$ being the baseline magnitude. The normalization has been chosen so that $\Omega_{\rm s} = 1$ for an unblended target ($g = 0$) of 18th magnitude, separated from the lens star by the angular Einstein radius $\theta_{\rm E}$ (i.e. $u = 1$), which corresponds to a magnification $A = 3/\sqrt{5}$.

If one considers a finite effective slew time $t_{\rm s}$, which includes all immediate overheads (such as the read out, slew, starting guiding and setting up the new exposure, and the observer reaction time), the relative investment becomes $I = t_{\rm obs}\,\tau^{-1}$, where $t_{\rm obs} = t_{\rm exp} + t_{\rm s}$. Thereby, with defining t_{18} as the exposure time required for achieving the desired photometric accuracy of 1.5 per cent on an unblended target ($g = 0$) of 18th magnitude, one finds for the gain factor

$$\Omega_{\rm s} = \omega_{\rm s}\,\frac{\Psi_{\rm p}(u)}{\sqrt{A}\,\left[\frac{(A+g)\,(1+g)}{A^2}\,\left(\frac{F_{\rm base}}{F_{18}}\right)^{-1} + \frac{t_{\rm s}}{t_{18}}\right]}\,, \tag{25}$$

where $\omega_{\rm s}$ absorbs all constants, and we choose it so that $\Omega_{\rm s} = 1$ for an unblended target ($g = 0$) of 18th magnitude, separated from the lens star by the angular Einstein radius $\theta_{\rm E}$ (i.e. $u = 1$), which corresponds to a magnification $A = 3/\sqrt{5}$. For the 1.54m Danish telescope at ESO La Silla, we adopt $t_{\rm s} = 2$ min, and $t_{18} = 20$ min.

At some point, the invested time $t_{\rm obs}$ becomes dominated by the slew time, so that there is a minimal investment for any bright target. In particular, for large magnifications A, $\zeta = t_{\rm s}/t_{18}$ will rule over the denominator in the expression for the gain factor $\Omega_{\rm s}$, so that with $\Psi_{\rm p}(u) \simeq 2/u \simeq 2A$ in this case, one finds

$$\Omega_{\rm s} = \frac{2\,\omega_{\rm s}}{\zeta}\,\sqrt{A}\,. \tag{26}$$

We populate the target list in order of decreasing gain factor $\Omega_{\rm s}$, until the slew time $t_{\rm s}$ and the specific exposure times $t_{\rm obs}$ and sampling intervals τ do not allow any further observations to be accommodated. Events with ongoing or suspected anomalies thereby take precedence over ordinary events, as discussed in more detail in the next section. This procedure implies that if overheads turn out to be larger than expected or losses occur due to weather conditions or technical problems, the coverage of all events will be degraded in proportion, whereas no kind of redistribution of the lesser available resources will take place.

The adopted strategy explicitly forces the telescope to slew between targets rather than sticking on the same one for a longer time, and moreover we try to spread requested observations on the same target as equally as possible over the night. In contrast, Horne et al. (2009) considered the gain in total invested time by avoiding frequent slewing between targets. However, by not moving to another object, one runs into redundancy at some point as the achievable photometric precision approaches the systematic uncertainty which leads to lots of data points at high cadence not carrying substantially more information than a smaller number.[15] An optimal sampling strategy would not only consider the planet detection probability at any given time for all ongoing events, but also aim at: 1) maximizing the information, 2) avoiding redundancy for planet detection arising from overlapping detection zones, 3) avoiding gaps that would prevent a timely detection of an ongoing anomaly without investing more time than necessary. Given that these points are not sufficiently accounted for in the approach presented by Horne et al. (2009), we preliminarily adopted a hybrid strategy merging those concepts with the event selection criteria that evolved from experience with the PLANET campaign. Designing a strategy that is truly 'optimal' would require a proper assessment of the information content of acquired data with regard to the predictability of the event magnification. In fact, a strategy aiming at just observing the most favourable events, peaking at large magnifications, as pursued by the MicroFUN collaboration[16], has a similar requirement of being able to know the magnification sufficiently ahead of time before the light curve peak actually occurs, and thereby success of such a campaign also needs investment into collecting data for event predictability.

As a pilot setup to a more sophisticated system that we plan to put into place at a later stage, we took a simplified approach, and neglected for the time being some more complex dependencies on further parameters, such as various effects related to the finite angular size of the source stars (e.g. Vermaak 2000; Han 2007), the event time-scale $t_{\rm E}$, the exact relation between the light curve prediction and the required efforts, the sky background, the availability of information from observations at other sites, etc. In fact, ignoring the finite angular extent of the source star leads to an overprediction of the planet detection prospects for small u, but given that there are never more than a handful of events falling into that category, and the sampling rate does not depend on this, the amount of actually invested efforts remains unaffected. The adopted strategy also neglects any information about previous coverage of the events under consideration. In particular, there are no attempts to favour events that have received attention before, nor is it tried to explicitly keep coverage gaps to a minimum.

4 Implementation of strategy via ARTEMiS

Rather than wasting precious efforts on putting parallel infrastructure into place, MiNDSTEp has realized the implementation of its 2008 observing strategy via the ARTEMiS system, thereby profiting from the existing tools for real-time data assessment for anomalies (by means of the SIGNALMEN anomaly detector), as well as for the display of model light curves and data of ongoing events. It is a major design feature of ARTEMiS not to dictate on the adopted strategy that leads to the selection of targets, but to provide a general and freely-available framework able to cater for various needs and allow to choose between different options.

In 2008, data from the OGLE, MOA, MiNDSTEp, RoboNet-II, and PLANET-III campaigns were made available for assessment by SIGNALMEN shortly after the observations had taken place. We were short of time getting a data link to MicroFUN installed, but as of 2009, OGLE, MOA, MiNDSTEp, RoboNet-II, and MicroFUN data can be exchanged efficiently via rsync servers, which means that a substantial fraction of the total acquired data can be assessed within minutes of their acquisition. The rsync server running in St Andrews acts as a gateway to both the ARTEMiS and the RoboNet-II systems. We have realized a special direct internal link with privileged access between St Andrews and Santa Barbara in order to connect SIGNALMEN with the PLOP event prioritisation system (Snodgrass et al. 2008). To ensure availability of the provided services to the outside world, we plan to get all relevant software running in St Andrews and Santa Barbara (courtesy of Las Cumbres Observatory Global Telescope Network) in parallel, thereby gaining protection from network and/or power outages.

The event modelling by SIGNALMEN is immediately triggered by new incoming data, where the respective events are entered into a queue system, from which agents for each of the telescopes pick events to be processed. This guarantees that a large amount of data released for a specific site does not block the processing of events that have recently been monitored elsewhere. A live data monitoring system[17] allows one to see the acquired data for a selected set of telescopes together with model light curves, with the display automatically switching to the event that was most recently observed. Moreover, this live monitor also shows the SIGNALMEN assessment indicating potential deviations or

[15] see Appendix A for a modification of the gain factor that better matches these considerations

[16] `http://www.astronomy.ohio-state.edu/~microfun`

[17] `http://www.artemis-uk.org/livecurves.cgi`

identified ongoing anomalies. Thereby, one can directly see how the adopted strategy works out in practice. Since the SIGNALMEN anomaly assessment might not be able to keep track with the data acquisition rate, preliminary plots showing the most recent data without remodelling and assessments are being provided, so that observers can always see their data displayed in the sequence in which these have been collected.

While the data monitoring system provides plots with pre-defined ranges, full flexibility is given with an interactive plotter[18], using graphics routines derived and advanced from the PLANET light curve plotter, earlier developed by M. Dominik and M.D. Albrow.

A further tool provided by ARTEMiS is the event overview[19] which displays the event model parameters, namely u_0, t_0, t_E, m_base, g, and current properties at time t_now, namely $t_\mathrm{now} - t_0$, $A(t_\mathrm{now})$, $A_\mathrm{obs}(t_\mathrm{now})$, $m(t_\mathrm{now})$, along with links to the data and up-to-the-minute light curve plots. The event prioritisation is carried out just as an option to the event overview script, which optionally leads to a display of the gain factor Ω_s. In fact, the event overview allows to select events matching widely customizable criteria, where user-defined requests can be saved simply as a URL with a respective query string.

Given that ARTEMiS is a fully transparent and publicly accessible system, it offers the opportunity to share live results at the forefront of scientific research with the general public. In fact, an ARTEMiS-powered "Catch-a-planet" display formed part of an exhibit "Is there anybody out there? Looking for new worlds" designed for the 2008 Royal Society Summer Science Exhibition and shown thereafter at various other locations. An overview of the currently monitored targets, as distributed over the Galactic bulge, acts as entry point, which is also available on the web.[20]

The list of the selected targets to be monitored at the 1.54 Danish telescope at ESO La Silla by the MiNDSTEp consortium is obtained from running the ARTEMiS event overview script after the observer has set appropriate exposure times t_exp that match the data quality requirements via a password-protected area. The target list is converted into an observing sequence using a simple combinatorial approach that aims at spreading the observations around as much as possible. This observing sequence is accessed by the observer via the ARTEMiS webpages.

The SIGNALMEN anomaly detector assigns one of three possible status to an event: 'ordinary', 'check', or 'anomaly' (Dominik et al. 2007). The status 'ordinary' means that there is no hint for an anomaly in progress,'anomaly' means that there is evidence for an ongoing anomaly, and 'check' means that there is a fair chance for an anomaly in progress, but sufficient evidence has not been found yet. While 'anomaly' status justifies immediate scheduling of the respective event at high cadence, 'check' status is meant to turn either into 'anomaly' or to 'ordinary' status with the acquisition of further data. However, in order to avoid misestimates on crucial statistical properties due to small samples, data from a given telescope are only assessed once there are at least 6 data points in total spreading over at least 2 previous nights. Therefore, events in 'check' status are only put on the target list (at a fixed sampling interval of 10 min) if these criteria are met. Events in 'anomaly' status are initially added at the same sampling interval of 10 min, and once SIGNALMEN has put events into 'anomaly' status, these need to be controlled manually by revising the sampling interval or clearing 'anomaly' status and setting the event back to 'ordinary'. Events in 'anomaly' status are allocated the highest priority in order of sampling interval (shortest first), followed by events in 'check' status, in order of decreasing Ω_s. Manual control of events in 'anomaly' status does not break our paradigm of a fully-deterministic strategy for probing lens stars for the presence of planets, given that we only *increase* the sampling interval in order to allow for more time to be spent on other targets if a denser sampling is not required for properly characterising the anomaly, i.e. we only lose data that are redundant anyway. Moreover, any monitored events that are not selected by our prioritisation algorithm, or anomalies declared by means other than SIGNALMEN constitute a separate monitoring programme rather than being part of the primary MiNDSTEp strategy, and are therefore to be discarded from the statistical analysis.

The survey campaigns acquire a substantial amount of their data while the event is close to baseline, i.e. $A \sim 1$. Given that SIGNALMEN is designed to call checks on suspected anomalies on about 5 per cent of the incoming data if no anomalies are present, it will mean that 'check' status is invoked frequently at such event phases, whereas the probability that a real deviation is present is rather small. Moreover, the large number of events at small $A(t)$ as compared to a much smaller number at larger $A(t)$ means that extraordinary observing conditions or a failure in the data reduction preferentially affects the set of events with $A \sim 1$. We therefore decided to schedule follow-up observations only if $\Omega_\mathrm{s} \geq 1$ for invoked 'check' status or if $\Omega_\mathrm{s} \geq 0.5$ for invoked 'anomaly' status.

As a result, there are events in 'anomaly' or 'check' status that are effectively treated as 'ordinary' events, which is handled by assigning each event not only a SIGNALMEN status, but also a substatus (depending on the telescope) that indicates how to respond. Similarly, there is also the option to put an event at high priority by manually setting a sampling interval, which will set its substatus to 'anomaly'. Finally, uninteresting events can be dropped from the target list by setting the SIGNALMEN anomaly flag, while not assigning a sampling interval, which will lead these events to be marked as 'dropped' instead of 'anomaly'. According to the primary status of 'ordinary', 'check', 'anomaly', or 'dropped', the events are marked with the colours black, yellow, red, and white, respectively, in the ARTEMiS event overview.

[18] http://www.artemis-uk.org/plot.cgi

[19] http://www.artemis-uk.org/event_overview.cgi

[20] http://www.artemis-uk.org/catch-a-planet.html

5 The adopted strategy in practice

5.1 2008 season MiNDSTEp observing campaign

The main focus of the MiNDSTEp consortium is on dense follow-up of microlensing events in order to study planet populations within the Milky Way, but a number of other projects are also carried out in parallel (e.g. Southworth et al. 2009a,b,c), mainly while the Galactic Bulge is not visible. During 2008, observations at the Danish 1.54m telescope at ESO La Silla lasted from June 1 to October 5. While the first 3 weeks were devoted mainly to building up and testing new infrastructure, the last 6 weeks were dominated by projects not following our microlensing follow-up strategy. Therefore, the time-span of a systematic campaign with automated scheduling using the ARTEMiS system by the MiNDSTEp consortium at the Danish 1.54m telescope at ESO La Silla ranges from the night of 23 June (JD = 2454641.xxx) until the night of 25 August (JD = 2454704.xxx). Rather than scheduling observations on all ongoing events that were reported by the OGLE and MOA surveys, we avoided having to deal with a mixture of two different selection biases by only including events announced by OGLE in our target list, while we monitored events that were only alerted by MOA (but not by OGLE) solely in manual-selection mode whenever those were considered highly interesting prime targets. In fact, the OGLE-III campaign concluded on 3 May 2009 after having provided a total of 4057 events towards the Galactic bulge since 2002, comprising a well-defined sample for statistical studies.

Out of the 64 nights that comprise our observing window, data were acquired on 53 of them, making 11 nights (i.e. about 17 per cent) a total loss. Over the 53 nights with observations, we operated the telescope for a total of 313 hours (on average 5.9 hours per observing night), where 131 hours were used for exposures, 95 hours for overhead time in between, and 87 were further operational losses with regard to our microlensing programme. The latter include weather, technical, and operational issues, such as clouds, strong winds, high humidity, focus sequences, telescope tests, and also sequences devoted to other projects than microlensing monitoring. In particular, we adopted the simple statistical approach to count any break between exposures of 5 min or less as overhead time, and longer breaks as loss time (not to be confused with the usually reported down time loss). The total time of 226 hours for exposures and overheads corresponds to an average of 4.3 effectively used hours per observing night. Figure 2 shows a breakdown of the accumulated exposure, overhead, and loss time for each of the nights. Moreover, it shows the number of targets monitored (typically between 8 and 15 each night), and the number of images taken (3603 in total, or 68 on average per observing night). One clearly sees that losses and overheads were quite substantial. The overall structure reflects a combination of the typical weather pattern at La Silla (with July rather unstable) and the fact that the microlensing field (the Galactic bulge) is visible a smaller fraction of the night during August than in June-July. The lack of observations in mid-August results from a combination of moonlight close to the targets and bad weather. With the observations spreading over 68 monitored events (see next subsection), we were strikingly close to the rough prediction of the capabilities by Dominik et al. (2002), namely being able to monitor 20 targets per night or 75 per observing season, despite some smaller differences in the adopted strategy.

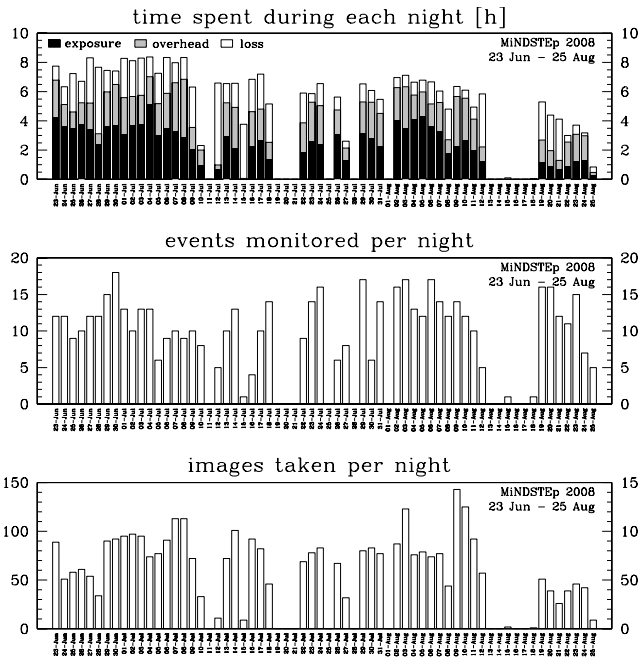

Fig. 2 Nightly breakdown of the exposure, overhead, and loss time, the events monitored, and the images taken during the systematic phase of the 2008 MiNDSTEp campaign at the 1.54m Danish telescope at ESO La Silla (Chile), extending from 23 June to 25 August.

It is hard to establish a correlation between the nightly overhead time averages, and other parameters that characterize the target selection. A strong hierarchy of sampling intervals, rather than events with comparably dense monitoring, would lead to observations getting clustered on the same target, thereby reducing the overhead times, but this does not imply that nights with a smaller number of targets see less slewing per exposure. In fact, we encounter the 'revenge of bright targets' which means that with their shorter exposure times, a smaller fractional amount of time is actually spent on the exposures themselves. For nights with more than 2 observations, Fig. 3 shows the overhead time average as a function of the number of targets and a spread coefficient $\kappa_{\rm s}$, which is defined as

$$\kappa_{\rm s} = \frac{N_{\rm slew}}{N_{\rm exp} - 1}\,, \tag{27}$$

where $N_{\rm exp}$ is the number of exposures (during a given night) and $N_{\rm slew}$ is the number of slew operations to a different target. This means that $\kappa_{\rm s} = 0$ corresponds to staying on the same target during all of the night, while $\kappa_{\rm s} = 1$ corresponds to always switching to a different object. As can

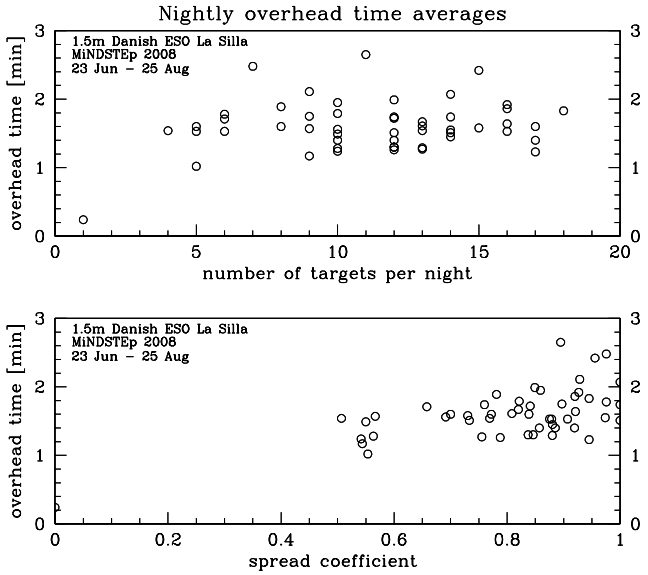

Fig. 3 Night averages of the overhead time (any break between exposures of 5 min or less) as a function of the number of targets observed during that night or the spread coefficient $\kappa_{\rm s} = N_{\rm slew}/(N_{\rm exp}-1)$, where $N_{\rm exp}$ denotes the nubmer of exposures and $N_{\rm slew}$ the number of slew operations.

be seen in the figure, the spread in the overhead time night average is substantial for $\kappa_{\rm s} = 1$, and the correlation with $\kappa_{\rm s}$ turns out not to be large in practice, unless one really sticks to a single target. In fact, it seems that operational delays on short time-scales significantly contribute to the statistics. We experienced an overall mean overhead time of 1.6 min. In theory, it would be possible to reduce the overhead time to something close to 30 s at the Danish 1.54m telescope, since this is the typical slew time between objects in the bulge as well as the typical read-out time (the read-out being done simultaneously). However, we find that adjusting the auto-guiding, adjusting for drifts, "guman overhead", etc. is the major factor of the delay time between the exposures in this type of projects.

5.2 Targets and observing mode

Any classification using the categories 'ordinary' and 'anomalous' for the events will lead to a substantial number of them being unclear or marginal, resulting in a large degree of arbitrariness. While no event will perfectly conform to being due to a single perfectly isolated lens star (given that such do not exist), the 'typical' event will show some weak deviation, frequently below the limit of characterizability. This means that 'ordinary' events exist only as approximation (no event can ever be proven not to involve any anomalies), and a comprehensive statistical analysis needs to take care of more general models for all events consistently. Consequently, we will avoid talking about 'ordinary' and 'anomalous' events here as far as possible, but instead refer to data being acquired in different SIGNALMEN status ('ordinary', 'check', 'anomaly'), which are well-defined.

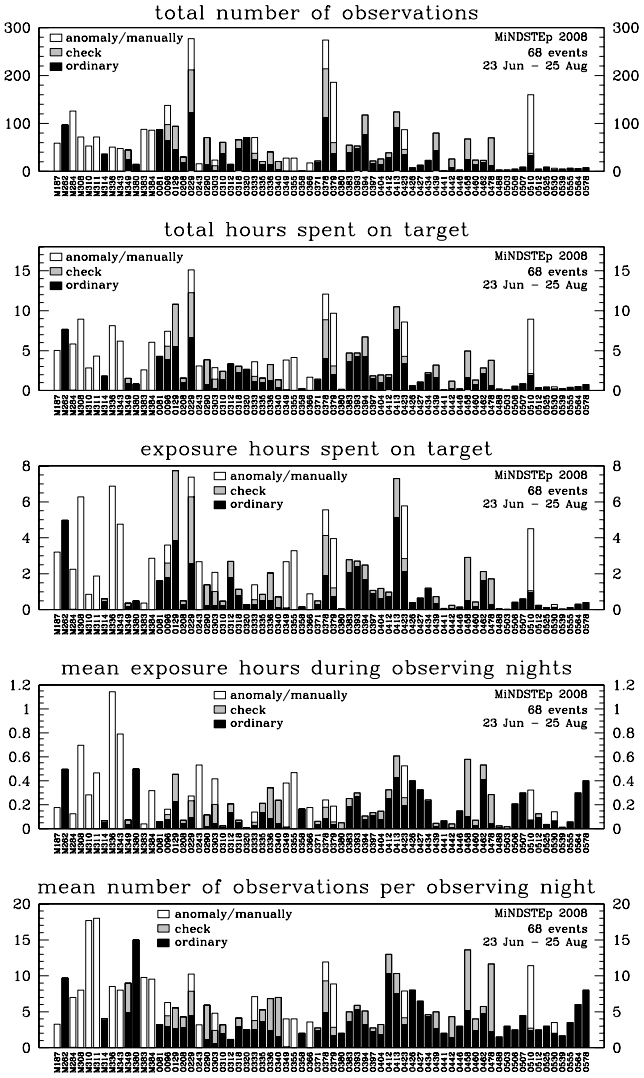

Fig. 4 Observations carried out per target in 'anomaly/manual', 'check' and 'ordinary' mode, respectively, during the systematic MiNDSTEp campaign at the 1.54m Danish telescope at ESO La Silla (Chile) from 23 June to 25 August 2008. In particular, the total number of observations, the total hours spent (exposure time plus overhead time), the exposure hours spent, the mean total exposure time (in hours) averaged over observing nights, and the mean number of observations during observing nights are shown. The event names OGLE-2008-BLG-xxx and MOA-2008-BLG-xxx have been abbreviated as Oxxx and Mxxx, respectively.

More precisely, we use the telescope-specific substatus (as explained in the previous section) for the classification, so that all manually-controlled events go under 'anomaly/manually', 'check' mode means that further dense observations have actually been requested, while all other data arose from 'ordinary' mode, prioritised according to the gain factor $\Omega_{\rm s}$.

Out of 3603 images acquired in total from 23 June to 25 August 2008, 1245 arose from 'anomaly/manually' mode, 784 from 'check' mode, and 1574 from 'ordinary' mode. Correspondingly, the 131 exposure hours distribute as 53 hours for 'anomaly/manually' (41 per cent), 28 hours for

Table 1 Final model parameters, as reported by SIGNALMEN, for the microlensing events observed in 2008 by the MiNDSTEp consortium at the Danish 1.54m telescope at ESO La Silla (Chile) during the systematic scheduling period covering the nights starting 23 June until 25 August.

event designation	alternative designation	$I_{\rm base}$		t_0 [UT (2008)]	t_0 [JD']	$t_{\rm E}$[d]	u_0	A_0	g	ΔI	$I_{\rm S}$	I_0
△ MOA-2008-BLG-187	—						*anomalous*					
MOA-2008-BLG-262	OGLE-2008-BLG-422	19.28	=	26 Jun, 16:55	4644.205	71.8	0.627	16.0	0.97	2.3	20.02	16.94
△ MOA-2008-BLG-284	—						*anomalous*					
△ MOA-2008-BLG-308	—						*anomalous*					
△ MOA-2008-BLG-310	—						*anomalous*					
△ MOA-2008-BLG-311	—						*anomalous*					
MOA-2008-BLG-314	OGLE-2008-BLG-459	15.69	=	25 Jul, 20:28	4673.353	10.6	0.249	4.1	0.03	1.51	15.72	14.18
▲ MOA-2008-BLG-336	OGLE-2008-BLG-524						*anomalous*					
▲ MOA-2008-BLG-343	OGLE-2008-BLG-493						*anomalous*					
MOA-2008-BLG-349	OGLE-2008-BLG-509	15.18	=	1 Aug, 1:07	4679.547	6.3	0.062	16.5	2.6	1.80	16.59	13.38
MOA-2008-BLG-380	OGLE-2008-BLG-543	19.79	=	10 Aug, 6:05	4688.754	7.5	0.045	22.4	0.33	3.1	20.10	16.71
△ MOA-2008-BLG-383	—						*anomalous*					
△ MOA-2008-BLG-384	—						*anomalous*					
OGLE-2008-BLG-081	—	15.44	=	15 Aug, 20:49	4694.368	97.3	0.615	1.85	0.37	0.52	15.78	14.92
■OGLE-2008-BLG-096	MOA-2008-BLG-166	16.45	+	11 Sep, 9:30	4720.896	104.4	0.539	2.1	0	0.78	16.45	15.67
OGLE-2008-BLG-129	MOA-2008-BLG-332	18.45	−	18 Jun, 2:42	4635.613	93.6	0.175	5.8	0.10	1.82	18.56	16.64
OGLE-2008-BLG-208	MOA-2008-BLG-212	16.63	−	8 Jun, 20:18	4626.346	25.0	0.030	34	0.13	3.7	16.76	12.93
■OGLE-2008-BLG-229	MOA-2008-BLG-272	17.19	=	18 Jul, 7:55	4665.830	51.0	0.151	6.7	0.44	1.73	17.59	15.46
□OGLE-2008-BLG-243	MOA-2008-BLG-184	19.44?					*anomalous*					
♦OGLE-2008-BLG-290	MOA-2008-BLG-241	16.94	−	(15 Jun, 0:27)	(4632.519)	(16.6)	(0.0006)		(0.006)			
■OGLE-2008-BLG-303	MOA-2008-BLG-267	19.38	−	17 Jun, 14:32	4635.106	34.1	0.022	46	0.22	3.9	19.60	15.43
OGLE-2008-BLG-310	MOA-2008-BLG-217	15.36	−	19 Jun, 18:23	4637.266	76.6	0.384	2.7	0.15	1.00	14.51	14.36
OGLE-2008-BLG-312	MOA-2008-BLG-247	19.42	−	8 Jun, 14:51	4626.119	30.7	0.588	17.0	0.33	2.8	19.74	16.64
OGLE-2008-BLG-318	MOA-2008-BLG-276	15.99	−	31 May, 9:51	4617.911	50.5	0.256	4.0	0	1.51	15.99	14.49
OGLE-2008-BLG-320	—	14.63	+	30 Aug, 13:07	4709.047	57.5	0.739	1.62	0	0.52	14.63	14.11
■OGLE-2008-BLG-333	MOA-2008-BLG-327	15.07	−	21 Jun, 8:08	4638.839	10.5	0.032	31	0	3.7	15.07	11.33
OGLE-2008-BLG-335	MOA-2008-BLG-296	17.36	+	17 Jul, 17:13	4665.218	75.6	0.176	5.7	3.8	0.74	19.07	16.62
OGLE-2008-BLG-336	MOA-2008-BLG-275	18.66	−	23 Jun, 13:42	4641.071	47.7	0.124	8.1	1.5	1.45	19.67	17.21
OGLE-2008-BLG-340	MOA-2008-BLG-256	16.65	−	14 Jun, 3:12	4631.634	7.1	0.112	9.0	0.11	2.3	16.77	14.37
▲OGLE-2008-BLG-349	MOA-2008-BLG-261	19.57	−	(15 Jun, 4:23)	(4632.683)	(24.4)	(0.028)		(0)			
□OGLE-2008-BLG-355	MOA-2008-BLG-288	18.4?					*anomalous*					
OGLE-2008-BLG-358	MOA-2008-BLG-264	18.79	−	10 Jun, 15:53	4628.162	21.4	0.059	17.0	6.0	1.3	20.90	17.49
□OGLE-2008-BLG-366	MOA-2008-BLG-321	17.70					*anomalous*					
OGLE-2008-BLG-371	—	17.53	=	21 Aug, 11:51	4699.994	81.2	0.120	8.4	3.3	1.09	19.12	16.45
■OGLE-2008-BLG-378	MOA-2008-BLG-299	15.81	=	11 Jul, 1:39	4658.569	14.6	0.311	3.3	0.06	1.26	15.88	14.55
■OGLE-2008-BLG-379	MOA-2008-BLG-293	17.52	=	(7 Jul, 22:36)	(4655.442)	(20.7)	(0.102)		(0.55)			
⋆OGLE-2008-BLG-380	—	17.47					*anomalous/unclear*					
OGLE-2008-BLG-383	—	17.70	−	23 Jun, 8:22	4640.849	28.0	0.095	10.5	0.99	1.91	18.45	15.79
OGLE-2008-BLG-393	MOA-2008-BLG-322	18.01	=	8 Jul, 14:16	4656.095	27.4	0.033	30	1.7	2.7	19.09	15.33
OGLE-2008-BLG-394	MOA-2008-BLG-300	16.22	=	29 Jul, 13:37	4677.068	23.0	0.324	3.2	0.01	1.26	16.23	14.96
OGLE-2008-BLG-397	MOA-2008-BLG-285	17.65	=	24 Jun, 14:13	4642.093	6.0	0.163	6.2	0.10	1.90	17.75	15.76
OGLE-2008-BLG-404	—	17.39	=	16 Aug, 23:25	4695.476	47.3	0.469	2.3	0.35	0.73	17.72	16.65
OGLE-2008-BLG-412	MOA-2008-BLG-292	18.60	=	30 Jun, 13:14	4648.052	15.5	0.206	4.9	0.38	1.46	18.95	17.14
OGLE-2008-BLG-413	MOA-2008-BLG-346	18.83	=	29 Jul, 1:30	4676.563	69.4	0.043	23.0	1.7	2.4	19.92	16.43
▲OGLE-2008-BLG-423	MOA-2008-BLG-302	16.69	=	(7 Jul, 8:49)	(4654.868)	(110.3)	(0.018)		(7.7)			
⋆OGLE-2008-BLG-426	—	19.54	=	(28 Jun, 18:53)	(4646.287)	(39.2)	(0.007)		(11)			
OGLE-2008-BLG-427	—	19.14	=	29 Jun, 8:29	4646.854	15.1	0.139	7.2	0.93	1.57	19.85	17.57
OGLE-2008-BLG-434	MOA-2008-BLG-309	19.41	=	6 Jul, 7:58	4653.832	16.6	0.046	21.8	0	3.3	19.41	16.06
OGLE-2008-BLG-439	MOA-2008-BLG-329	15.93	=	27 Jul, 23:08	4675.464	12.9	0.448	2.4	0	0.95	15.93	14.99
OGLE-2008-BLG-441	—	17.74	=	27 Jul, 21:17	4675.387	25.0	0.54	2.0	0	0.78	17.74	16.97
♦OGLE-2008-BLG-442	MOA-2008-BLG-317	15.99	+	(30 Aug, 12:44)	(4709.031)	(33.5)	(0.810)		(0)			
OGLE-2008-BLG-446	—	19.39	=	22 Jul, 19:52	4670.328	39.4	0.229	4.5	0.003	1.62	19.39	17.77
OGLE-2008-BLG-458	MOA-2008-BLG-331	18.95	=	5 Jul, 4:49	4652.701	20.3	0.109	9.3	0.84	1.85	19.51	17.01
OGLE-2008-BLG-460	MOA-2008-BLG-333	16.22	=	21 Jul, 21:00	4669.375	9.0	0.572	1.96	0.000	0.73	16.22	15.49
OGLE-2008-BLG-462	—	18.87	=	22 Jul, 8:44	4669.864	29.4	0.212	4.8	0.02	1.68	18.90	17.19
OGLE-2008-BLG-478	MOA-2008-BLG-392	16.96	=	11 Aug, 16:56	4690.206	37.2	0.174	5.8	0.54	1.54	17.43	15.42
OGLE-2008-BLG-488	MOA-2008-BLG-400	18.64	=	23 Aug, 11:18	4701.971	47.9	0.256	4.0	0.21	1.35	18.85	17.29
OGLE-2008-BLG-503	MOA-2008-BLG-354	14.72	=	2 Aug, 9:11	4680.883	5.8	0.784	1.55	0.06	0.45	14.78	14.26
OGLE-2008-BLG-506	—	18.58	=	25 Jul, 8:09	4672.840	32.7	0.023	44	20	1.23	21.86	17.35
OGLE-2008-BLG-507	—	17.92	=	2 Aug, 0:25	4680.518	8.8	0.457	2.4	0	0.93	17.92	16.99
■OGLE-2008-BLG-510	MOA-2008-BLG-369	19.23	=	(10 Aug, 2:01)	(4688.584)	(22.8)	(0.062)		(0.06)			
OGLE-2008-BLG-512	MOA-2008-BLG-363	17.85	=	5 Aug, 17:46	4684.240	12.0	0.179	5.6	2.5	0.92	19.21	16.93
OGLE-2008-BLG-525	MOA-2008-BLG-368	16.72	=	18 Aug, 21:13	4697.384	13.1	0.291	3.6	0.64	1.02	17.26	15.71
■OGLE-2008-BLG-530	MOA-2008-BLG-374	18.155	=	(8 Aug, 9:52)	(4686.911)	(15.4)	(0.084)		(4.5)			
OGLE-2008-BLG-539	—	16.60	=	24 Aug, 0:20	4702.514	12.6	0.456	2.4	0.68	0.65	17.16	15.96
OGLE-2008-BLG-555	MOA-2008-BLG-397	17.24	=	19 Aug, 2:18	4697.596	4.4	0.165	6.5	0.06	1.92	17.30	15.32
OGLE-2008-BLG-564	—	19.00	+	31 Aug, 6:53	4709.787	40.8	0.184	54	3.2	2.8	20.57	16.17
⋆OGLE-2008-BLG-578	—	18.44	=	(23 Aug, 9:33)	(4701.898)	(5000)	(6×10^{-5})		(1200)			

$I_{\rm base}$ baseline magnitude, t_0 epoch of peak, $t_{\rm E}$ event time-scale, u_0 , A_0 peak magnification, g blend ratio, $\Delta I = I_0 - I_{\rm base}$ observed brightening between peak and baseline, $I_{\rm S}$ intrinsinc source star
magnitude, I_0 peak magnitude
Indicators for events with data in 'anomaly/manually' mode: △ event without OGLE survey data, ▲ manual control gained, ■ SIGNALMEN/ARTEMiS-activated anomaly, □ anomaly reported by OGLE
before data release, ♦ anomaly outside the observing window, ⋆ event unclear, model estimate uncertain or erratic
Indicators of peak location: − before observing window, = within observing window, + after observing window
For some events with possible anomalies, indicative parameters are given in brackets, referring to a background ordinary model used for prioritising the 'ordinary' data that have been acquired

'check' (22 per cent), and 50 hours for 'ordinary'. With 10 to 15 per cent of all the events exhibiting apparent anomalies, one can expect ~ 10 ongoing anomalous events at any time. Given that the successful detection of anomalies caused by low-mass planets (10 $M_\oplus$ or below) requires the dense monitoring of a substantial number of ongoing (ordinary) microlensing events (Dominik et al. 2007), one needs to resist the temptation to put lots of efforts into anomalous events that are of less interest. This means to make the proper choice of balance between follow-up and anomaly monitoring. In fact, over the recent years, the balance between survey and follow-up monitoring has moved, in particular with MOA now monitoring several fields with sub-hourly sampling. While we were rather efficient in picking up ongoing anomalies (22 out of the 68 monitored events contain observations in 'anomaly/manually' mode), the total amount of time invested into these appears to be on the high side, and better mechanisms should be put into place that allow to assess the characteristics of ongoing anomalies, such as the use of simple indicators (e.g. Han & Gaudi 2008), so that efforts can be better focused on the major science goals.

Fig. 4 shows the efforts that went into the monitoring of each of the 68 targets, namely the number of observations, the exposure time and total time spent (exposure plus overhead), the mean exposure hours during the 53 observing nights, and the mean number of observations, whereas Table 1 lists the respective final event parameters. For the overhead time, we assigned the break between exposures to equal parts to the two images, as long as it is less than 5 min, and adopted a fixed value of 2 min otherwise. Further information, such as light curve plots, are available by means of the ARTEMiS system.[21] We discuss everything related to the prioritisation of ordinary events in the next subsection, the result of the interaction with the SIGNALMEN anomaly detector in the subsequent one, while first we keep to more general issues.

With 12 events monitored exclusively in 'anomaly/manually' mode, amongst those 7 that arose from the MOA survey only without a corresponding OGLE detection, and the unclear event OGLE-2008-BLG-380, one is left with 55 events on which data have been acquired in 'ordinary' mode. The events OGLE-2008-BLG-229, OGLE-2008-BLG-378, OGLE-2008-BLG-129, OGLE-2008-BLG-413, OGLE-2008-BLG-379, and OGLE-2008-BLG-510 attracted the largest investment of time effort within our two-month observing window, amongst which only OGLE-2008-BLG-129 and OGLE-2008-BLG-413 never had 'anomaly/manually' status invoked. Thereby, the lead investment is dominated by events that were considered to involve an ongoing anomaly. Despite the fact that a larger amount of time spent on a target does not necessarily imply a larger number of observations, given the vastly different exposure times (see Fig. 5), these events constitute 6 out of the 7 events with the largest number of observations, with MOA-2008-BLG-284 intervening, as the most demanding event with 'ordinary' observations only.

[21] http://www.artemis-uk.org/event_overview_2008.cgi

The left panel of Fig. 6 shows the trend between the number of observations and the time invested, and reveals as well that there are some events that do not follow it. The overhead time reduces substantially the contrast between brighter and fainter events, because the typical investment on an exposure does not fall below about 2.5 min, whereas exposure times are rarely in excess of 10 min.

The events OGLE-2008-BLG-426 and OGLE-2008-BLG-578 were alerted by OGLE due to a sudden peak, and were picked up promptly and immediately by the ARTEMiS system without any human intervention (thanks to the rsync link with the OGLE computers). Despite the fact that the respective model parameters for an ordinary microlensing light curve were meaningless, a fast reaction on potential short time-scale events is far more important than potentially wasting time on phenomena that have a different origin, because this harbours the potential for a study of isolated sub-stellar (including planetary) mass objects within the Milky Way.

Comparing a typical event-scale of 20 days (Eq. (12)) with the width of the observing window of our systematic operations in 2008, namely 2 months, it becomes obvious that there are events that are desired to be observed before or after. Moreover, we found that 1/3 of the monitored events have their peak magnification outside our observing window. In order for an event to be observed over a long time-span within the observing window (see Fig. 5), its time-scale should be long, the target should be bright, and the peak should occur near the centre of the observing window. In fact, OGLE-2008-BLG-081 and OGLE-2008-BLG-320 would have been monitored over more than the 2 month time-span, being 2 out of the three brightest events at baseline and having a time-scale in excess of 50 days. The other event with a top brightness, OGLE-2008-BLG-503, however has a short coverage time-span due to its event time-scale $t_\mathrm{E} \sim 6$ d. The rather small efforts that went into the events OGLE-2008-BLG-503 to OGLE-2008-BLG-564 are not the result of these being alerted near the end of the observing window, but due to a clustering of small event time-scales $t_\mathrm{E} \lesssim 15$ d. There are 3 exceptions: OGLE-2008-BLG-506 has $t_\mathrm{E} = 33$ d, but involves a faint target, OGLE-2008-BLG-510 involves a remarkable anomaly (Bozza et al. 2010) (and thereby is the only event in this group that received a huge amount of attention), and OGLE-2008-BLG-564 has $t_\mathrm{E} = 40$ d, is quite faint at baseline, and peaks after the end of the observing window. The fraction of nights with observations during the monitoring time-span usually ranges between 50 and 100 per cent, with events with a long time-span tending to have lower fractional coverage, in particular both OGLE-2008-BLG-081 and OGLE-2008-BLG-320 are around 50 per cent, i.e. a bit less than 30 nights.

Given that a large number of observing nights on a given event is likely to result from monitoring the wing regions, which comes with a smaller magnification and thereby longer sampling interval, many observations per night usually do not come in combination with many observing nights. In

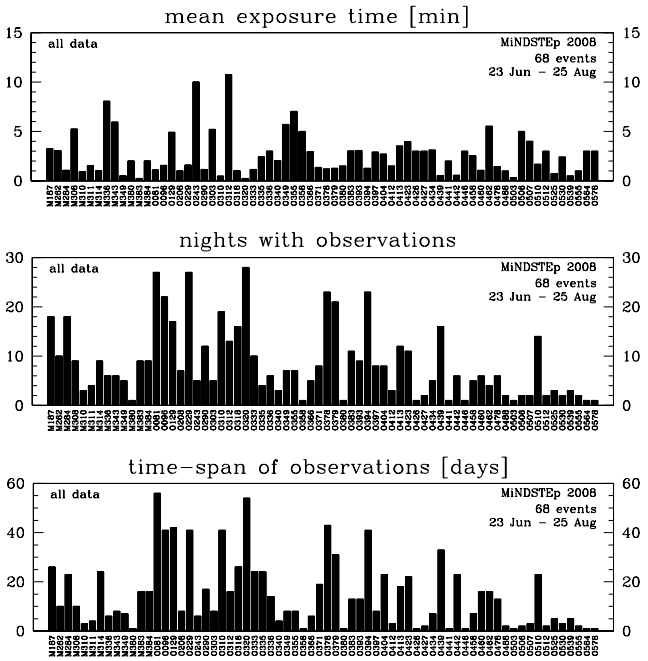

Fig. 5 Mean exposure time, nights with observations and the total time-span of the observations (last night − first night + 1) for each of the 68 events monitored within the MiNDSTEp campaign from 23 June until 25 August 2008. The event names OGLE-2008-BLG-xxx and MOA-2008-BLG-xxx have been abbreviated as Oxxx and Mxxx, respectively.

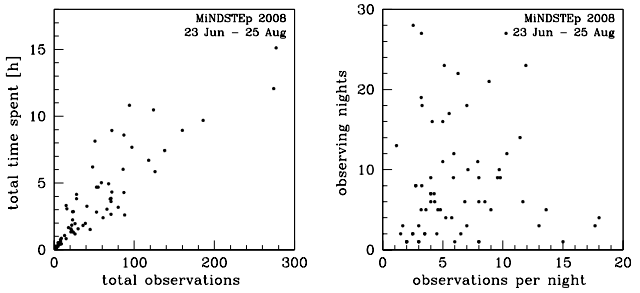

Fig. 6 (left) Total time spent on each of the 68 microlensing events monitored during the systematic phase of the 2008 MiNDSTEp campaign as a function of the total number of observations. (right) Observing nights as function of the mean number of observations per night for the same data set.

fact, only a small number of events (10) have more than 10 observations per night on average (see Fig. 6 right), and it is only OGLE-2008-BLG-229 and OGLE-2008-BLG-378 that due to observations in 'check' and 'anomaly/manually' mode have more than 10 observations per night on average and more than 20 observing nights (and thereby more than 200 observations). The densest average sampling has been obtained on three manually-controlled events without OGLE survey data, namely MOA-2008-BLG-311, MOA-2008-BLG-310, and MOA-2008-BLG-308 during 4, 3, or 1 observing nights, respectively, which on short time-scale reverted to being a faint target.

We aim at covering each of the monitored events at least once every 2 hours, which appears to have failed significantly if less than 3 images have been acquired per night. Looking at the averages shown in Fig. 4, 11 out of the 68 events fail on this criterion, however each of these under special circumstances. For OGLE-2008-BLG-312 and OGLE-2008-BLG-358, observations much after the peak that occured before the observing window of our systematic campaign appear to have been carried out by mistake. 5 events can be identified with those that just made it to the observing list as the event with the lowest gain factor Ω_s for 2 or 3 nights, respectively (OGLE-2008-BLG-441, OGLE-2008-BLG-488, OGLE-2008-BLG-506, OGLE-2008-BLG-512, and OGLE-2008-BLG-539), and therefore were just given the remaining available observing time without being able to fulfill the desired sampling rate. 2 events, namely OGLE-2008-BLG-371 and OGLE-2008-BLG-397 suffered from bad weather or priority given to anomaly monitoring for some of the observing nights out of a moderate total number (8). Similarly affected is the average long-term coverage (28 nights) of OGLE-2008-BLG-320 (while OGLE-2008-BLG-081 shows this trend to a lesser extent, keeping the average number of images per observing nights slightly above 3). Finally, we acquired just 2 data points on OGLE-2008-BLG-380 during a single night before dismissing this event on which the OGLE photometry indicated difficulties or weird behaviour.

5.3 Distribution of event characteristics

The statistics of event magnifications and observed magnitudes for the collected data are based on the rather small number of 52 events with observations in 'ordinary' monitoring mode while SIGNALMEN reported 'ordinary' status (1494 in total), after dismissing OGLE-2008-BLG-426 and OGLE-2008-BLG-578 which show a short strong rise due to anomalous behaviour, while there were no 'ordinary'-'ordinary' data on OGLE-2008-BLG-503 within the observing window. With the peak magnification and the baseline magnitude being specific to each event, the respective diagrams in Fig. 7 displaying the distribution of the event magnification, the observed event magnification, the current magnitude, the gain factor, and the event phase with respect to the number of observations and the time spent (exposure time plus overhead) therefore show a substantial scatter.

Given the strong correlation between the number of observations and the time spent (see also Fig. 6, left panel), the distributions for both ways of counting look quite similar, and for the gain factor Ω_s are even hard to distinguish, except for $\Omega_s < 2$, unless the current magnitude itself is considered, where the smaller efforts on brighter targets become apparent. Plotting the time spent on the target (exposure plus overhead time) as a function of the current magnitude (see Fig. 8) shows a large scatter with given magnitude over a rather small rise towards fainter targets. It is only for (the small number of) targets with $I > 18.5$ that the time invested into the observation becomes substantially larger.

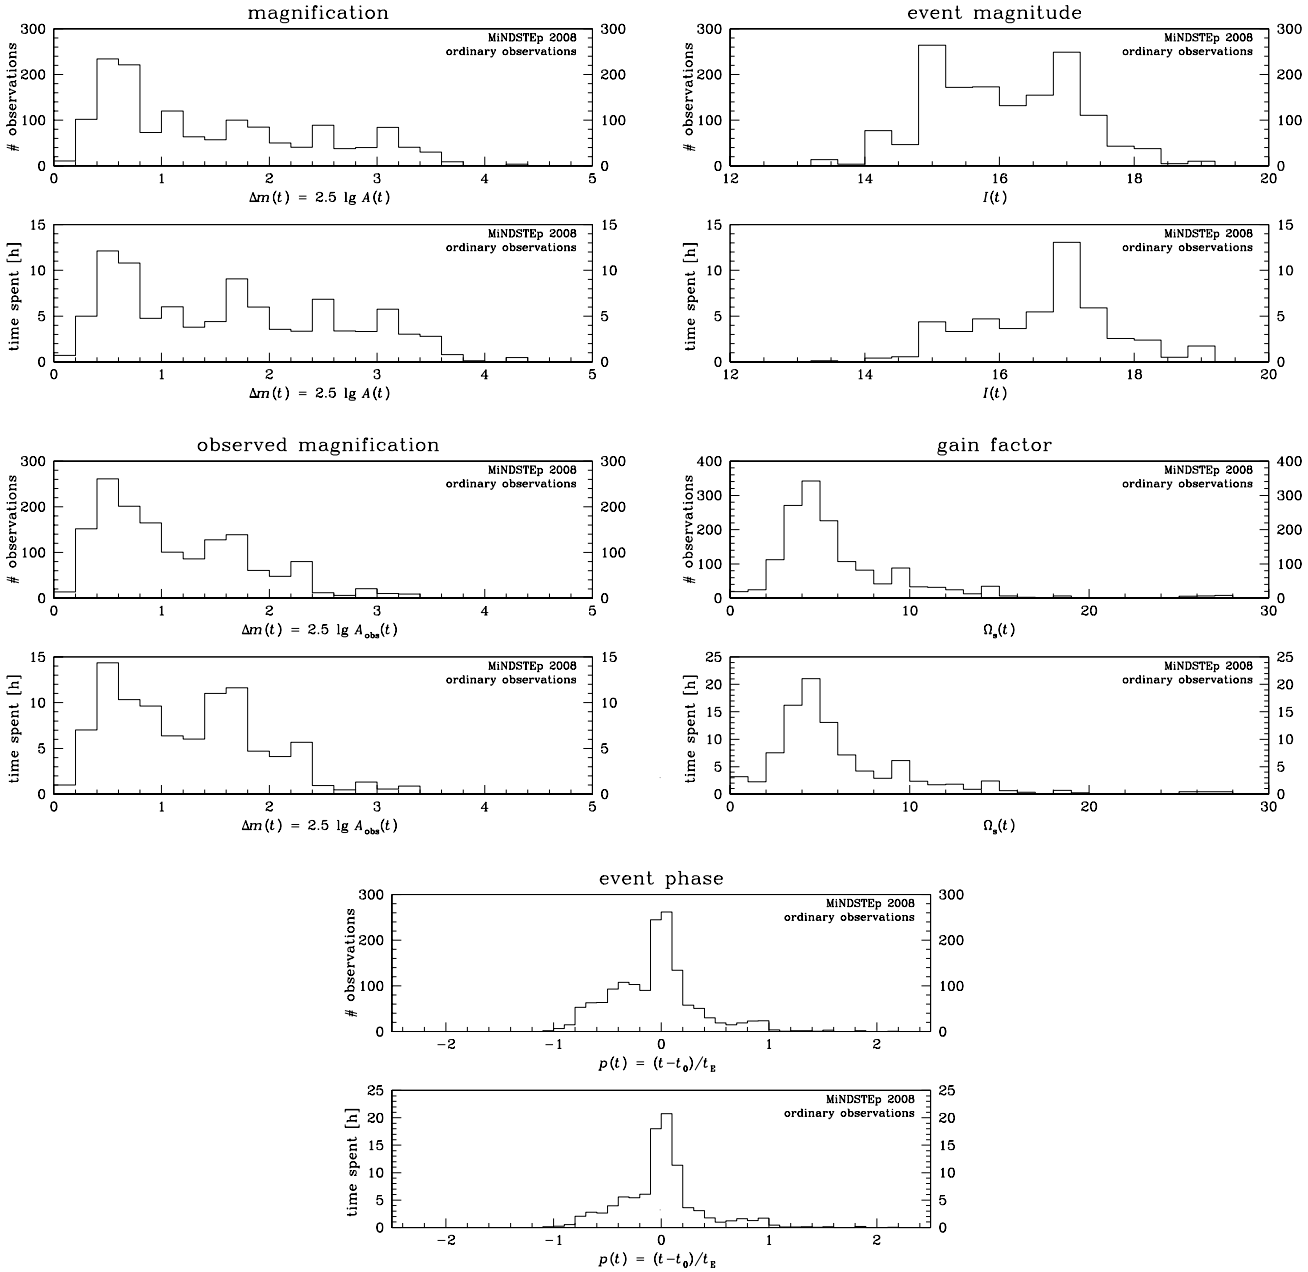

Fig. 7 Distribution of the current event magnification $A(t)$ (or respectively $\Delta m = 2.5 \lg A(t)$), the observed event magnification $A_{\rm obs}(t)$ (or respectively $\Delta m = 2.5 \lg A_{\rm obs}(t)$), the event magnitude $I(t)$, the gain factor $\Omega_{\rm s}(t)$, and the event phase $p(t) = (t - t_0)/t_{\rm E}$ with respect to the number of observations or the invested time into acquiring the frame for the observations in 'ordinary' mode at the 1.54m Danish telescope at ESO La Silla (Chile) during the systematic 2008 MiNDSTEp campaign between 23 June and 25 August.

The exposure time itself shows a stronger trend with target magnitude, but this is weakened by the overhead time. Fig. 9 allows to identify 'trajectories' of $I(t)$, $m(t)$ and $\Omega_{\rm s}$ for a given event, which themselves provide unambiguous relations: Δm decreasing with I, $\Omega_{\rm s}$ increasing with Δm, and $\Omega_{\rm s}$ decreasing with I. The spread resulting from the ensemble of events, with their various baseline magnitudes and peak magnifications, however destroys the one-to-one relationship. The most notable remaining feature indeed is the domination of $\Omega_{\rm s} < 2$ by faint events, while one also finds an excess of faint observed magnitudes in highly-magnified events (those common targets would not be monitored otherwise), whereas brighter events ($I < 15.5$) appear to be clusters at smaller magnifications ($\Delta m < 1$ mag).

There is a paucity of observations for a magnitude shift $\Delta m(t) = 2.5 \lg A(t)$ less than 0.2 mag, given that the respective target priority was not large enough to warrant selection, and also for $\Delta m > 3.6$ mag, reflecting the small chance for such a magnification to occur. The peak between 0.4 and 0.8 mag reflects the wing phases of the monitored

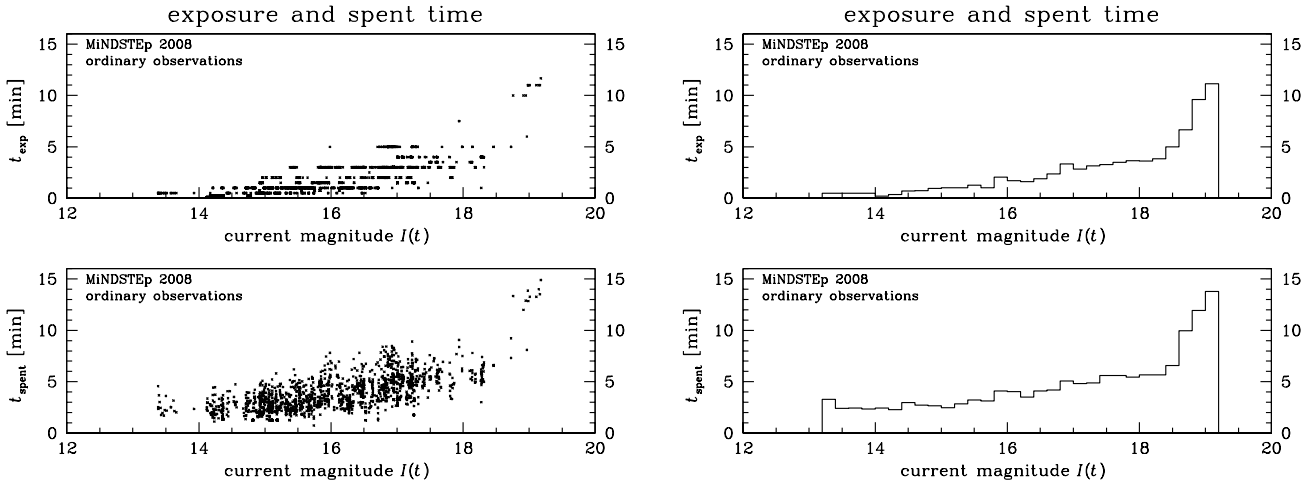

Fig. 8 Exposure time $t_{\rm exp}$ and spent time $t_{\rm spent}$ (including overhead time) as a function of the current target I-magnitude for the 1494 ordinary observations during the systematic phase of the 2008 MiNDSTEp campaign. While the left panels plot the respective quantities for each exposure, the binning applied in the right panels allows to reveal the underlying trend over the scatter.

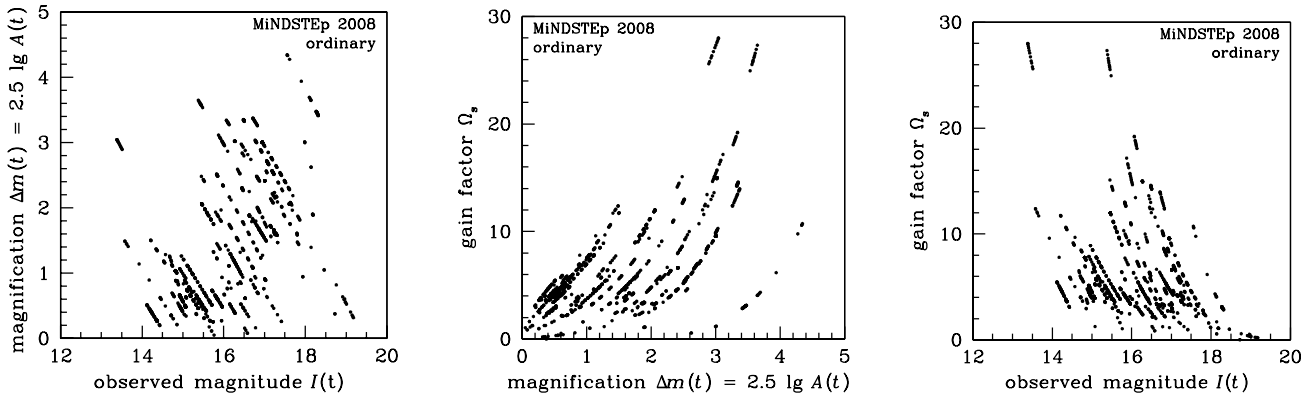

Fig. 9 (left) Magnitude shift $\Delta m = 2.5 \lg A(t)$ of the source star as a function of the current I magnitude for the 1494 'ordinary' observations during the systematic 2008 MiNDSTEp campaign at the 1.54m Danish telescope at ESO La Silla. (center) Gain factor $\Omega_{\rm s}$ as a function of Δm for the same observations. (right) Gain factor $\Omega_{\rm s}$ as a function of $I(t)$ for the same observations.

events, and *not* that large efforts are thrown at events with small peak magnifications. The underrepresentation of this peak on counting the invested time rather than the number of observations also indicates that extended wing coverage is more prominent with brighter targets. Source stars brightened between 0.8 mag and 3.6 mag received comparable amounts of attention with a decrease by about a factor 2 towards the higher magnifications, which is influenced by the $\tau \propto \sqrt{A}$ law for the sampling interval as well as the statistics of current event impact parameters $u(t)$.

The low-magnification peak is more prominent if the observed magnification $A_{\rm obs}(t) = [A(t) + g]/(1 + g)$ is considered, where the blend ratio g leads to $A_{\rm obs}(t) \leq A(t)$. Consequently, with a stronger concentration towards smaller magnitude shifts, the decrease of fractional investments towards actually brighter targets is stronger, and differences exceeding 2.4 mag already become rare.

The bimodality of the distribution of the current event magnitude amongst the observations with peaks around $I \sim 15$ and $I \sim 17$ appears to reflect the bimodality between either main-sequence or giant source stars. Only very few observations were taken on events that appeared brighter than 14 mag or fainter than 18.5 mag. Due to the smaller exposure time on bright targets, the distribution counting the invested time only retains a single peak at around $I \sim 17$, while fractional efforts were almost negligible for $I < 14.8$, whereas there was some significant (but small) fraction of time spent on targets down to $I \sim 19.2$ (actually less than 2.5 per cent went to $I > 18$).

The median gain factor $\Omega_{\rm s}$ of the 1494 'ordinary' observations is 4.3. With higher gains occuring less frequently and the sampling interval being a function of the magnification A rather than $\Omega_{\rm s}$, 56 per cent of the observations have moderate gain factors $3 < \Omega_{\rm s} < 6$, while 33 per cent of the observations have $\Omega_{\rm s} > 6$, and only 12 per cent have $\Omega_{\rm s} > 10$. This shows explicitly that strategies exclusively targeted at observations with large gain would occupy a rather small fraction of the time, whereas dedicated tele-

scopes collect information from a much larger number of targets with smaller gain factors. For a good scientific programme at a non-dedicated site with on-demand access, it would be meaningful to choose a proper cut-off in the gain factor that balances between the return-to-investment factor and the overall return by obtaining a larger data set. Given that the detection of OGLE-2005-BLG-390Lb (Beaulieu et al. 2006; Dominik et al. 2006), at a gain factor $\Omega_{\rm s} = 3.2$ with the parameters of the Danish telescope,[22] was the result that by far had the largest impact of all the Galactic gravitational microlensing observations, we are not in favour of advocating a strategy that would have missed it.

In any case, a further aspect that needs to be considered is that a proper characterisation of events is required both for the later scientific analysis and the real-time event prediction (without which it is not possible to identify the most favourable ones), calling for investment beyond that of the microlensing surveys.

The distribution of the event phases $p(t) = (t - t_0)/t_{\rm E}$ is an overlay of a narrower peak around the peak of the event with 34 per cent of the observations carried out while $|p| < 0.1$, and a broader distribution that 'kicks in' at about $p \sim -0.8$, which can be identified as alert jump, and basically ends at about $p = 1$, with very few observations taken further down the wing. The broader distribution is substantially skewed, leaning towards $p < 0$. This reflects the operations of the SIGNALMEN anomaly detector, sending events into 'check' and 'anomaly' mode, which cannot happen in early phases, given that there is a minimum requirement for the number of acquired data points, and deviations are difficult to assess with poorly-constrained models. This means that we also see the differential efficiency of SIGNALMEN on pre-peak and post-peak data.

Plotting the distribution of the maximal reached event magnification, observed event magnification, observed target brightness or gain factor for each night, while indicating the time spent on the respective event, as shown in Fig. 10, reveals that there is no need to devote a 1.54m telescope just to the monitoring of one or two events. No single 'ordinary' event ever dominated the observations, while even in the nights of 7 July or 31 July, comprising the events with the highest gain factors encountered, 10 or 14 events, respectively, were monitored. Few events are the result of, in first instance, huge losses during the night (compare Fig. 2), and in second instance of a (rare) huge load of anomaly monitoring. For the observations in 'ordinary' mode, prioritised by means of the gain factor $\Omega_{\rm s}$, it is the current ensemble that determines whether the number of monitored events turns out to be rather near 8 or rather near 15. One also sees that the time invested not just increases with the gain factor $\Omega_{\rm s}$, but the sampling interval τ is proportional to $\sqrt{A}$ (with A being the current magnification), while fainter targets require longer exposure times.

[22] A telescope slewing twice as fast would yield $\Omega_{\rm s} = 4.3$, while neglecting the slew time would mean $\Omega_{\rm s} = 7.5$.

5.4 Anomaly detection by means of immediate feedback

In total, 126 'check' requests on suspected deviations from an ordinary microlensing light curve, detected in 34 events, have been forwarded to the 1.54m Danish telescope by the SIGNALMEN anomaly detector. While the prompted dense monitoring caused SIGNALMEN to revert to 'ordinary' status in 120 cases, the remaining 6 cases led to flagging up an 'anomaly'.

By construction, SIGNALMEN should invoke 'check' status on about 5 per cent of the incoming data if no anomalies are present (Dominik et al. 2007). However, the 1.54m Danish telescope not only reacted to 'check' triggers arising from its own data, but also on those resulting from an assessment of data that were made available from other sites around the world where observations were carried out by OGLE (Las Campanas, Chile), MOA (Mt John, New Zealand), RoboNet-II (Faulkes North, Hawaii; Faulkes South, Australia; Liverpool Telescope, Canary Islands), and PLANET-III (Canopus Observatory, Tasmania; Perth Observatory, Western Australia; SAAO, South Africa). In fact, the latter provide a 2/3 majority on the 'check' triggers that the 1.54m Danish telescope reacted to, namely 85 in total (MOA: 44 , OGLE: 35, SAAO: 11, Canopus: 5, Perth: 2, Faulkes South: 1), whereas 41 'check' triggers arose from data taken with the 1.54m Danish at La Silla (Chile) itself. Given the longitudinal distribution of the sites, the Danish telescope will always pick up on 'check' triggers that arise from its own data, unless these occur very near the end of the night's observing window or the telescope needs to close due to weather or technical failures. In contrast, 'check' status arising from data acquired elsewhere can already revert to 'ordinary' before observations at ESO La Silla can commence. While it is explicitly desired for a follow-up telescope like the Danish 1.54m to invest into 'check' request that arise from data at a survey telescope, since this much increases the chances for detecting low-mass planets with the large number of events covered by the surveys, the 'check' load should be evenly distributed amongst a follow-up network, rather than being focused on the 1.54m Danish, as in 2008. Given that the main goal is a fast reaction to short-lived anomalies (on the time-scale of a few hours), one might also consider to let 'check' requests expire if these cannot be met with proper telescope resources. However, we can well accept 50 per cent of the 'check' trigger to arise from survey sites, favourably from the OGLE telescope at Las Campanas, because additional efforts at La Silla could be undertaken immediately and would allow the identification of planetary anomalies on just suspicious but not evident OGLE data, without severely impacting the OGLE survey itself by such an intensive follow-up effort. In a similar way, the Faulkes Telescope South (in Australia) could interact with the MOA telescope (in New Zealand), albeit with a slightly smaller time overlap.

Taking 5 further data points on 5 per cent of the acquired data would mean to take 20 per cent of total observations

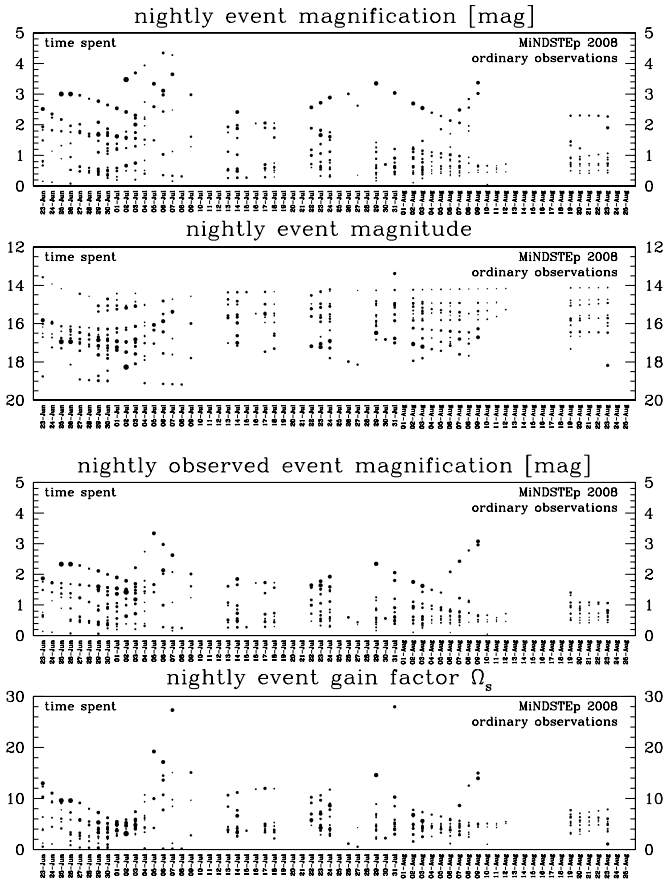

Fig. 10 Distribution of the event magnification $A(t)$ (respectively the source magnitude shift $\Delta m(t) = 2.5 \lg A(t)$), the observed event magnification $A_{\rm obs}(t)$ (respectively the observed magnitude shift $\Delta m_{\rm obs}(t) = 2.5 \lg A_{\rm obs}(t)$), the target magnitude, and the gain factor $\Omega_{\rm s}$ during each of the observing nights. Each event is represented by a single plot symbol, whose size corresponds to the time spent, while the reported value reflects the brightest exposure during the given night.

(except for anomaly monitoring) in 'check' mode. Now, for the systematic 2008 MiNDSTEp campaign, we find an investment of 33 per cent of the observations or 36 per cent of the exposure time. With 2/3 of the 'check' triggers not arising from data acquired with the 1.54m Danish, we would have expected a much larger investment. In fact, 44 'check' triggers from 1574 'ordinary' observations means triggering on ~ 2.8 per cent of the data on average. Looking more closely into this reveals that the 5 per cent level can only meaningfully be established based on the scatter of the data if more than 50 data points on the event under consideration have already been acquired, whereas SIGNALMEN is far more conservative at earlier stages, including to refuse assessing data unless observations from at least two previous nights have been carried out, which is reflected in the event-phase distribution for data acquired in 'ordinary' mode (see Fig. 7). With only 20 per cent of the 'ordinary' data taken when at least 50 previous points had already been acquired, the smaller number of 'check' triggers does not come as a complete surprise. Morevover, the small-number statistics well explain why we do not see any check requests on many of those events with less than 40 'ordinary' data points. An unrealistic, exceptional scatter can not only increase the number of 'check' triggers, namely when it first occurs, but also reduce it, namely every time thereafter by having increased the trigger threshold. By assessing the scatter of the received data, SIGNALMEN consequently does *not* trigger on improper data if those have already appeared earlier on the same event. However, if the data reduction fails for the first time during the night of consideration, or if the scatter substantially increases, there is no mechanism yet in place to properly pick up on this, whereas a respective additional criterion for judging about an anomaly has been described (Dominik et al. 2007). In fact, the large scatter arising from data-reduction problems explains the absence of any trigger on 1.5 Danish data on the events MOA-2008-BLG-262, OGLE-2008-BLG-081, and OGLE-2008-BLG-320.

On the other hand, there are 3 main factors that *increase* the number of 'check' triggers: the misestimation of model parameters leading to a discrepancy between the model light curve and the observed magnitude, a sudden increase in scatter due to failed data reduction, and finally, the events in fact contain real anomalies. Amongst all anomalies, those of planetary origin constitute a rather small minority, and by rough estimate 10 to 15 per cent of the events show apparent anomalies at some point. A particular problem arises from

events with intrinsic variability of the magnified blended source star, because this variation will increase over the course of the event. There are triggers on weak anomalies that never turn into sufficient evidence, and it is an established pattern that such occur repeatedly on the same event. As a relict from the inability to provide sufficient data for telling the difference between an anomalous and an ordinary event, further data on such marginal events are being requested, and this is being made through 'check' mode.

The vast majority ($\sim$ 95 per cent) of the incurred 'check' requests cause SIGNALMEN to revert to 'ordinary' status. Three different major routes to calling off an anomaly can be distinguished: First, a revision of the underlying ordinary light curve model or a reassessment of the scatter based on data points that are not suspected to indicate anomalous behaviour can mean the sudden death to a suspected anomaly with just one or two further data points (about 16 per cent of the encountered 'check requests'). Second, the main route (about 80 per cent) is to find evidence against a significant deviation, which with a reasonable data quality is achieved after 4 further data points (and this is also the mimimum number explicitly requested). Third, evidence for an anomaly might accumulate, but after 5 points have been acquired, SIGNALMEN decides in favour of revising the model rather than calling for an anomaly. This is typically the longest process with 5-10 further data points, and occured in 2.5 per cent of our 'check' requests.

We found 44 textbook examples (1/3 of all 'check' requests) of discarding a suspected anomaly by means of 4 data points as further evidence. Out of the 'check' requests dismissed by means of evidence against a real deviation, 45 per cent were called off by data from the 1.54m Danish only, 34 per cent together with data from other sites, and 22 per cent with other data only (see below for reasons why the data from La Silla were not sufficient).

While there were a large number of good performances, precisely as expected, we experienced some expected difficulties, and actually ran into one trap that we did not think about properly in advance. Namely, in order for a 'check' status being reverted to 'ordinary', data need to be reduced in real-time and made available to the ARTEMiS system. If either the timely data transfer or the data reduction itself fails, the event remained in 'check' status, where the 'check' request should have been cleared in such a case. It turned out that 14 per cent of the 'check' observations were wasted due to not having taken appropriate measures to account for this. A further loss of 8 per cent arose from a software bug that allowed a 'check' request to be submitted, while the SIGNALMEN criteria for data assessment (at least 6 previous points from at least 2 previous nights) were not met. Moreover, another 8 per cent of 'check' points were acquired unnecessarily, because the loop from data acquisition over the SIGNALMEN assessment and via the ARTEMiS tools back to the observer took too long, be it due to data not being uploaded pointwise, or the ARTEMiS machines being busy or the process itself requiring too much time. This means that by fixing these 3 issues, about 30 per cent of the time invested in 'check' mode could have been spent more efficiently otherwise. Given the rather small number of 'check' triggers per event, some events are particularly badly affected, namely OGLE-2008-BLG-478 (*all* of the 'check' data), OGLE-2008-BLG-442 (all), OGLE-2008-BLG-423 (all), OGLE-2008-BLG-336 (80 per cent), OGLE-2008-BLG-229 (60 per cent), and OGLE-2008-BLG-129 (50 per cent).

A timely call-off of a suspected (but not real) anomaly becomes difficult or impossible if the actual data quality is below the desired quality level. This effect was apparent for 4 per cent of the acquired 'check' mode data, and in particular on 40 per cent of the 'check' data for OGLE-2008-BLG-378. Moreover, there have been 4 events with excessive triggering on 1.54m Danish data, namely OGLE-2008-BLG-229, OGLE-2008-BLG-290, OGLE-2008-BLG-379, and OGLE-2008-BLG-458. However, together this accounts for less than 6 per cent of the 'check' observations, so that losses by bad data quality in total amount to less than 10 per cent. While the data quality is of course a crucial factor for the detectability of planetary (and other) deviations, sub-optimal photometry turns out not to be such an important factor for the amount of investment into confirming or rejecting suspected anomalies by means of providing immediate feedback to SIGNALMEN 'check' requests.

The assessment of data on potential anomalies by SIGNALMEN is a rather complex process, taking into account that data might arrive in blocks and out of time-sequence, might be altered at any time or withdrawn. Therefore, it is possible that a suspected anomaly is re-assessed due to data being received with a delay that were acquired before or during the suspected anomaly, which might alter the model parameters and more seriously, the conclusion about the presence of the anomaly.

Within our sample of 'check' requests, we found a few further oddities. First, the technical time delay of the data assessment not only led to data taken in 'check' mode while the event should have already been considered 'ordinary', but also to data taken in 'ordinary' mode while the 'check' flag had already been raised on earlier observations. Second, on OGLE-2008-BLG-439, a 'check' request resulted from an erratic point that was later removed. Third, one of the 'check' requests on OGLE-2008-BLG-229 disappeared too early due to a bug that restarted the assessment at a point that did not support the suspected anomaly. Fourth, and finally, on OGLE-2008-BLG-413, 4 out of 5 'check' triggers have been called off by model revision, reflecting the comparably poor guidance provided by the OGLE and MOA survey data.

Out of 22 events with data acquired in 'anomaly/manually' mode, there were 6 cases of 'check' requests turning into 'anomaly' status, namely on OGLE-2008-BLG-096, OGLE-2008-BLG-229, OGLE-2008-BLG-333, OGLE-2008-BLG-378, OGLE-2008-BLG-379, and OGLE-2008-BLG-510, whereas there were no 'check' requests for the 1.54m Danish between SIGNALMEN invoking 'check' status and turning

this into 'anomaly' for OGLE-2008-BLG-303 and OGLE-2008-BLG-530. Given a caustic exit in early MOA data of OGLE-2008-BLG-423, when realized, 'anomaly' status was set manually. 13 events were manually controlled from the beginning, amongst these 7 events without OGLE survey data, for which this was the mandatory mode according to our event selection criteria. Furthermore, 3 events were reported anomalous by OGLE before the respective data were released (OGLE-2008-BLG-243, OGLE-2008-BLG-355, and OGLE-2008-BLG-366), an anomaly was suspected in OGLE-2008-BLG-349 following SIGNALMEN assessment before the beginning of the systematic campaign, MOA-2008-BLG-343 was anomalous with the first release of data, and MOA-2008-BLG-336 was brought under manual control because the observed rise follows an earlier pair of caustic passages. During the 2 months of the systematic 2008 MiNDSTEp campaign, SIGNALMEN has not missed on flagging any present anomaly from assessing 'ordinary' or 'check' data in the monitored events that was reported by a third party.

In contrast to reverting from 'check' status to 'ordinary', clearing 'anomaly' status or revising the initial sampling interval of 10 min requires human intervention. In principle, we require that attention is given within 1 hour after the 'anomaly' status has been invoked. However, time appeared to be too short (and/or other commitments to large or too many) for getting a suitable web-interface in place before the beginning of the 2008 observing season that would have enabled this. Given that flagging up 'anomaly' status is a rather rare event, it is in any case a worthwhile investment of spending even a whole night with dense monitoring at 10 min intervals. While the obtained ~ 40 images are a substantial addition to the number of observations for the specific event (only about 10 per cent of the events had more than 100 exposures taken), we are talking about less than 10 per cent of all the observations in total.

Events with particularly striking anomalies monitored during the 2 months of the systematic 2008 MiNDSTEp campaign include MOA-2008-BLG-310 with a planetary anomaly partly covered by our data (Janczak et al. 2010), OGLE-2008-BLG-290 with a prominent finite-source peak (Fouqué et al. 2010), and OGLE-2008-BLG-510 where the otherwise undetected anomaly was revealed by MiNDSTEp 1.54m Danish data following a 'check' request of the SIGNALMEN anomaly detector (Bozza et al. 2010).

6 Conclusions and recommendations

While being far from claiming perfection, our setup appears to be technically ready for inferring the statistics of planetary populations. Several issues have been identified that should be addressed in order to improve the efficiency of our and related campaigns.

We have explicitly demonstrated that an automated selection of microlensing events for follow-up monitoring can show a reasonable performance in practice, without the need to devise an extremely detailed and sophisticated algorithm. With the return-to-investment ratio measured by a gain factor, well-balanced between the current event magnification, the current target magnitude, and the blend ratio, usually between 8 and 15 targets during a good observing night could have been monitoring in fulfillment of our scientific requirements. With a preference of event phases around the peak of the event, a fair fraction of the time has been devoted to the coverage of the wings, while omitting those regions that are not very sensitive to revealing planetary signals.

The amount of time devoted to acquiring further data in response of suspected ongoing anomalies as reported by SIGNALMEN is a quite valuable investment, given that it provides 1) the capability to detect signals of planets of Earth mass and below, 2) the opportunity to obtain a proper statistical sample, where however the achievable sample size for the least massive detectable planets poses a limit to the statistical significance. While spending 36 per cent of the exposure time on this during the 2 months of the systematic 2008 MiNDSTEp campaign was rather on the high side of the expectations, some simple measures could reduce this to 15–20 per cent without compromising on the efficiency. However, as compared to an 'optimal' performance, we did not lose 20 per cent of the coverage of average targets, but it is those with the lowest priority that got dropped off the list.

One should ensure to only observe in 'check' mode if real-time communication with the ARTEMiS system is possible. Moreover, 'check' mode observations should require that the real-time data reduction is producing data, and that the achieved data quality meets the requirements. Furthermore, the accumulation of requests for confirmation or rejection of suspected anomalies should not accumulate at a single telescope, but it should be tried to distribute these more evenly over a telescope network and/or let the requests expire after some time. One might also think about compensating for telescopes with extremely high data rate. With both the 2007 and 2008 microlensing seasons, SIGNALMEN has demonstrated its high sensitivity to ongoing anomalies (e.g. Dominik et al. 2008b; Han et al. 2009; Sumi et al. 2010), while keeping the number of false alerts at a low level. The balance between detection threshold and the number of false alerts could be improved with a more sophisticated noise model, considering correlations between subsequent data points, which appear to be particularly common in the MOA data set.

We also find that fast exchange of data and information is a determining factor in achieving our scientific goals. Any time delays, be it in the data reduction, data transfer, data assessment, or in communicating back to observers or robotic telescopes should be kept to an absolute minimum, given that such add up. Transmitting data in larger blocks should be avoided, while the rsync tool provides a very efficient means to keep information up to date at various locations. Currently, the ARTEMiS system runs on a rather small mini-cluster of 3 machines with 10 processor cores

in total, and to enhance the reliability and make the operations nearly technically failproof, mirrors would be required, while the performance needs to be improved both by optimizing the computer codes itself and extending the hardware pool.

While the efficiency of detecting planet rises and falls with the quality of the real-time photometry, this is not a crucial factor for the amount of time spent on 'check' requests, because SIGNALMEN performs an analysis of the real scatter. The data-reduction pipeline running at the Danish telescope during the 2008 MiNDSTEp campaign did not provide good results on strongly blended targets, thereby testing the ARTEMiS system near the worst-case scenario. This however, did not have a substantial effect on its performance, because data of bad quality were usually recognized as such.

Planetary anomalies make up only a small fraction of all deviations from an ordinary microlensing light curve that regularly show up in the data. An efficient strategy for probing planets would therefore need to sort out all 'uninteresting' anomalies as soon as possible and in particular avoid to invest a substantial amount of time on the respective events. There is an even higher need to reject events that are not due to microlensing, because by trying to fit their data to ordinary microlensing light curves, these enter the prioritisation algorithm with meaningless parameters.

Manually-operated telescopes that were built some decades ago, such as the 1.54m Danish telescope at ESO La Silla, have rather large slew times, and are therefore everything but optimal for continuous switching between targets. Moreover, the general time overheads of its operation with observers are of concern. Modern robotic telescopes explicitly designed for fast slewing can be expected to substantially increase the efficiency of a microlensing follow-up monitoring programme along our strategy. Still with the large slew time, we regularly managed to monitor 8 to 15 targets during a good observing night while matching all scientific requirements. Concentrating the efforts at a dedicated telescope of that diameter on just one or two targets thereby appears to be quite inefficient. As a precursor study for the upcoming SONG (Stellar Observations Network Group)[23] network (Grundahl et al. 2008), a lucky-imaging camera is being tested at the 1.54m Danish at ESO La Silla in 2009, which will advance the prospects for ground-based high-precision microlensing observations.

While our adopted monitoring strategy turned into a reasonable spread of the observing time over the ongoing microlensing targets, it still contains a consistency gap. Namely, we leave it to the observer to select an appropriate exposure time to match the data quality requirements, which is then used to determine how many events can be monitored, whereas the event prioritisation follows expected exposure times related to photon statistics, according to Eq. (17). For the upcoming operations with robotic telescopes, we therefore intend to adopt an iterative procedure where exposure times are automatically adapted according to the data quality returned. This would also allow for optimal exposure times in cases where the event magnification can hardly be predicted.

We also did not yet solve the problem on how frequently an event should be optimally sampled at a given time, i.e. how much time should be invested in achieving the ability to predict the light curve, against which deviations are detected.

Gaps in the coverage due to weather or technical failure (or engineering time) are a very serious issue. For signals of low-mass planets, lasting only a few hours, the instantaneous coverage at a given longitude is more important than a round-the-clock coverage. Therefore, a telescope network should consider several sites at comparable longitudes. Moreover, if one would like to start talking about 'evidence' for a detection, data from at least two sites are a scientific requirement for transient phenomena, whereas a one-off observation with a single instrument can never qualify as a scientific discovery.

[23] http://astro.phys.au.dk/SONG

Acknowledgments

We would like to thank the OGLE, MOA, RoboNet-II, PLANET-III, and MicroFUN collaborations for providing real-time data that gave useful guidance to our observations, and in particular substantially improved the event magnification predictions. Even more substantially, we are indebted to the OGLE and MOA survey programmes without which a real-time follow-up programme such as MiNDSTEp would not be possible. We also thank Scott Gaudi for his helpful comments and suggestions, and acknowledge support from the Danish Natural Science Research Council (FNU). Research at the Armagh Observatory is funded by the Department of Culture, Arts and Leisure (DCAL).

References

Albrow, M., Beaulieu, J.-P., Birch, P., et al.: 1998, ApJ 509, 687
Alcock, C., Akerlof, C. W., Allsman, R. A., et al.: 1993, Nature 365, 621
Allan, A., Naylor, T., Saunders, E. S.: 2006, AN 327, 767
Beaulieu, J.-P., Bennett, D. P., Fouqué, P., et al.: 2006, Nature 439, 437
Bond, I. A., Abe, F., Dodd, R. J., et al.: 2001, MNRAS 327, 868
Bond, I. A., Udalski, A., Jaroszyński, M., et al.: 2004, ApJ 606, L155
Bozza V., et al.: 2010, in preparation
Burgdorf, M. J., Bramich, D. M., Dominik, M., Bode, M. F., Horne, K. D., Steele, I. A., Rattenbury, N., Tsapras, Y.: 2007, P&SS 55, 582
Calchi Novati, S., Bozza, V., De Paolis, F., et al.: 2009, ApJ 695, 442
Chang, K., Refsdal, S.: 1979, Nature 282, 561
Charbonneau, D., Brown, T. M., Latham, D. W., Mayor, M.: 2000, ApJ 529, L45
Chung, S.-J., et al.: 2006, ApJ 650, 432

Covone, G., de Ritis, R., Dominik, M., Marino A. A.: 2000, A&A 357, 816
Dominik, M., et al.: 2002, P&SS 50, 299
Dominik, M., et al.: 2007, MNRAS 380, 792
Dominik, M., Horne, K., Allan, A., et al.: 2008a, AN 329, 248
Dominik, M., Horne, K., Allan A., et al.: 2008b, in Sun, Y.-S., Ferraz-Mello, S., Zhou, J.-L., eds, IAU Symposium Vol. 249 of IAU Symposium, ARTEMiS (Automated Robotic Terrestrial Exoplanet Microlensing Search) – Hunting for planets of Earth mass and below. pp 35–41
Dominik, M., Jørgensen, U. G., Horne, K., et al.: 2008c, Inferring statistics of planet populations by means of automated microlensing searches, White paper submitted to ESA's Exo-Planet Roadmap Advisory Team (EPR-AT), preprint arXiv:0808.0004
Dominik, M., Horne, K., Bode, M. F.: 2006, A&G 47, 3.25
Dominik, M.: 2006, MNRAS 367, 669
Dominik, M.: 2008, ARTEMIS, cooperative efforts, and optimal short-term strategies, in Kerins E., Mao S., Rattenbury N., Wyrzykowski Ł., eds., Proceedings of the Manchester Microlensing Conference: The 12th International Conference and ANGLES Microlensing Workshop, PoS(GMC8)048
Dominik, M.: 2009, MNRAS 393, 816
Dyson, F. W., Eddington, A. S., Davidson, C.: 1920, PhilTransA 220,291
Einstein, A.: 1911, Annalen der Physik, 340, 898
Einstein, A.: 1915, Sitzungsberichte der preußischen Akademie der Wissenschaften 47, 831
Einstein, A.: 1936, Science 84, 506
Elachi, C., et al.: 1996, A Road Map for the Exploration of Neighboring Planetary Systems (ExNPS), Technical report, Jet Propulsion Laboratory, California Institute of Technology Pasadena, CA United States
Fouqué, P., Heyrovský, D., Dong, S., et al.: 2009, OGLE 2008 BLG 290: An accurate measurement of the limb darkening of a Galactic Bulge K Giant spatially resolved by microlensing, A&A submitted
Gaudi, B. S., Bennett, D. P., Udalski, A., et al.: 2008, Science 319, 927
Gaudi, B. S., Sackett, P. D.: 2000, ApJ 528, 56
Gould, A., Dong, S., Gaudi, B. S., et al.: 2010, Frequency of Solar-Like Systems and Ice and Gas Giants Beyond the Snow Line from High-Magnification Microlensing Events, ApJ submitted, preprint arXiv:1001.0572
Gould, A., Loeb, A.: 1992, ApJ 396, 104
Gould, A.: 2008, MicroFUN 2007 in Kerins, E., Mao, S., Rattenbury, N., Wyrzykowski, Ł., eds., Proceedings of the Manchester Microlensing Conference: The 12th International Conference and ANGLES Microlensing Workshop, PoS(GMC8)038
Griest, K., Safizadeh, N.: 1998, ApJ 500, 37
Grundahl, F., Christensen-Dalsgaard, J., Kjeldsen, H., Frandsen, S., Arentoft, T., Kjaergaard, P., Jørgensen, U. G.: 2008, in Deng, L., Chan, K. L., eds., IAU Symposium Vol. 252 of IAU Symposium, SONG – Stellar Observations Network Group. pp 465–466
Han, C., Gaudi, B. S.: 2008, ApJ 689, 53
Han, C., et al.: 2009, Interpretation of Strong Short-Term Perturbations in the Light Curves of Moderate-Magnication Microlensing Events, in preparation
Han, C.: 2007, ApJ 661, 1202
Henry, G. W., Marcy, G. W., Butler, R. P., Vogt, S. S.: 2000, ApJ 529, L41
Hessman, F. V.: 2006, AN 327, 763
Horne, K., Snodgrass, C., Tsapras, Y.: 2009, MNRAS 396, 2087
Ida, S., Lin, D. N. C.: 2005, ApJ 626, 1045
Ingrosso, G., Calchi Novati N., De Paolis, F., Jetzer, P., Nucita, A. A., Zakharov, A. F.: 2009, MNRAS 399, 219
Janczak, J., Fukui, A., Dong, S., et al.: 2010, ApJ 711, 731
Kiraga, M., Paczyński, B.: 1994, ApJ 430, L101
Liebes, S.: 1964, PhysRev 133, 835
Mao, S., Paczyński, B.: 1991, ApJ 374, L37
Marcy, G. W., Butler, R. P., Fischer, D. A., Laughlin, G., Vogt, S. S., Henry, G. W., Pourbaix, D.: 2002, ApJ 581, 1375
Marois, C., Macintosh, B., Barman, T., Zuckerman, B., Song, I., Patience, J., Lafrenière, D., Doyon, R.: 2008, Science 322, 1348
Mayor, M., Queloz, D.: 1995, Nature 378, 355
Paczyński, B.: 1986, ApJ 304, 1
Paczyński, B.: 1991, ApJ 371, L63
Paczyński, B.: 1996, ARA&A 34, 419
Refsdal, S.: 1964, MNRAS 128, 295
Renn, J., Sauer, T., Stachel, J., 1979, Science 275, 184
Rivera, E. J., Lissauer, J. J., Butler, R. P., et al.: 2005, ApJ 634, 625
Smith, M. C., Woźniak, P., Mao, S., Sumi, T.: 2007, MNRAS 380, 805
Snodgrass, C., Tsapras, Y., Street, R., Bramich, D., Horne, K., Dominik, M., Allan, A.: 2008, The WEB-plop observation prioritisation system in Kerins, E., Mao, S., Rattenbury, N., Wyrzykowski, Ł., eds., Proceedings of the Manchester Microlensing Conference: The 12th International Conference and ANGLES Microlensing Workshop, PoS(GMC8)056
Southworth, J., Hinse, T. C., Jørgensen, U. G., et al.: 2009a, MNRAS 396, 1023
Southworth, J., Hinse, T. C., Burgdorf, M. J., et al.: 2009b, MNRAS 399, 287
Southworth, J., Hinse, T. C., Dominik, M. , et al.: 2009c, ApJ 707, 167
Steele, I. A., Naylor, T., Allan, A., Etherton, J., Mottram, C. J.: 2002, in Kibrick, R. I., ed., Advanced Global Communications Technologies for Astronomy II Vol. 4845 of Proc. SPIE, eSTAR: a distributed telescope network. pp 13–24
Sumi, T., Bennett, D. P., Bond, I. A., et al.: 2010, ApJ 710, 1641
Tsapras, Y., Street, R., Horne, K., et al.: 2009, AN 330, 4
Udalski, A., Żebrun, K., Szymański, M., Kubiak, M., Soszyński, I., Szewczyk, O., Wyrzykowski, Ł., Pietrzyński, G.: 2002, AcA 52, 115
Udalski, A.: 2003, AcA 53, 291
Udry, S., Bonfils, X., Delfosse, X., et al.: 2007, A&A 469, L43
Vermaak, P.: 2000, MNRAS 319, 1011

A Systematic measurement uncertainty and gain factor

The assumption that the photon-count noise $\sigma_{\rm N} = \sqrt{N}$ dominates the measurement uncertainties is actually equivalent to neglecting any systematic uncertainties, which in turn means that with sufficiently long exposures one can achieve arbitrarily precise measurement. Given that this is obviously not correct, this calls for a modification of the relations previously derived in Sect. 3.

In particularly, with a fractional systematic uncertainty σ_0, the relative flux uncertainty becomes

$$\left(\frac{\sigma_F}{F}\right)^2 = \sqrt{\frac{\kappa^2}{F\,t_{\rm exp}} + \sigma_0^2}\,, \tag{A1}$$

so that

$$(t_{\rm exp})^{-1} = \kappa^{-2}\,F\,\left[\left(\frac{A}{A+g}\right)^2\left(\frac{\sigma_A}{A}\right)^2 - \sigma_0^2\right]\,, \tag{A2}$$

where

$$\sigma_{\rm req} \equiv \frac{A}{A+g}\,\frac{\sigma_A}{A} > \sigma_0 \tag{A3}$$

is required, and $t_{\rm exp} \to \infty$ as $\sigma_{\rm req} \to \sigma_0$. Therefore, obtaining measurements with a desired relative precision $\sigma_{\rm req} \leq \sigma_0$ becomes impossible.

Consequently, with

$$\beta = \frac{\sigma_0}{\sigma_A/A}\,, \tag{A4}$$

the gain factor becomes

$$\Omega_{\rm s} = \frac{\omega_{\rm s}}{\sqrt{A}}\,\frac{\psi_{\rm p}(u)}{\frac{F_{18}}{F_{\rm base}}\,(1-\beta^2)\,\frac{1+g}{A+g}\,\left[\frac{A^2}{(A+g)^2} - \beta^2\right]^{-1} + \frac{t_{\rm s}}{t_{18}}} \tag{A5}$$

for

$$\frac{A}{A+g} > \beta \tag{A6}$$

and zero otherwise, which reproduces Eq. (25) for $\beta = 0$.

