

Library Declaration and Deposit Agreement

1. STUDENT DETAILS

Gregory Charles Brown
1258033

2. THESIS DEPOSIT

- 2.1 I understand that under my registration at the University, I am required to deposit my thesis with the University in BOTH hard copy and in digital format. The digital version should normally be saved as a single pdf file.
- 2.2 The hard copy will be housed in the University Library. The digital version will be deposited in the University's Institutional Repository (WRAP). Unless otherwise indicated (see 2.3 below) this will be made openly accessible on the Internet and will be supplied to the British Library to be made available online via its Electronic Theses Online Service (EThOS) service. [At present, theses submitted for a Masters degree by Research (MA, MSc, LLM, MS or MMedSci) are not being deposited in WRAP and not being made available via EThOS. This may change in future.]
- 2.3 In exceptional circumstances, the Chair of the Board of Graduate Studies may grant permission for an embargo to be placed on public access to the hard copy thesis for a limited period. It is also possible to apply separately for an embargo on the digital version. (Further information is available in the Guide to Examinations for Higher Degrees by Research.)
- 2.4 (a) Hard Copy I hereby deposit a hard copy of my thesis in the University Library to be made publicly available to readers immediately.
I agree that my thesis may be photocopied.
- (b) Digital Copy I hereby deposit a digital copy of my thesis to be held in WRAP and made available via EThOS.
My thesis can be made publicly available online.

3. GRANTING OF NON-EXCLUSIVE RIGHTS

Whether I deposit my Work personally or through an assistant or other agent, I agree to the following: Rights granted to the University of Warwick and the British Library and the user of the thesis through this agreement are non-exclusive. I retain all rights in the thesis in its present version or future versions. I agree that the institutional repository administrators and the British Library or their agents may, without changing content, digitise and migrate the thesis to any medium or format for the purpose of future preservation and accessibility.

4. DECLARATIONS

(a) I DECLARE THAT:

- I am the author and owner of the copyright in the thesis and/or I have the authority of the authors and owners of the copyright in the thesis to make this agreement. Reproduction of any part of this thesis for teaching or in academic or other forms of publication is subject to the normal limitations on the use of copyrighted materials and to the proper and full acknowledgement of its source.
- The digital version of the thesis I am supplying is the same version as the final, hardbound copy submitted in completion of my degree, once any minor corrections have been completed.
- I have exercised reasonable care to ensure that the thesis is original, and does not to the best of my knowledge break any UK law or other Intellectual Property Right, or contain any confidential material.
- I understand that, through the medium of the Internet, files will be available to automated agents, and may be searched and copied by, for example, text mining and plagiarism detection software.

(b) IF I HAVE AGREED (in Section 2 above) TO MAKE MY THESIS PUBLICLY AVAILABLE DIGITALLY, I ALSO DECLARE THAT:


- I grant the University of Warwick and the British Library a licence to make available on the Internet the thesis in digitised format through the Institutional Repository and through the British Library via the EThOS service.
- If my thesis does include any substantial subsidiary material owned by third-party copyright holders, I have sought and obtained permission to include it in any version of my thesis available in digital format and that this permission encompasses the rights that I have granted to the University of Warwick and to the British Library.

5. LEGAL INFRINGEMENTS

I understand that neither the University of Warwick nor the British Library have any obligation to take legal action on behalf of myself, or other rights holders, in the event of infringement of intellectual property rights, breach of contract or of any other right, in the thesis.

Please sign this agreement and return it to the Graduate School Office when you submit your thesis.

Student's signature:Date:



A Study of Exotic Nuclear Extragalactic Transients

by

Gregory Charles Brown

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Department of Physics

September 2016

THE UNIVERSITY OF
WARWICK

Contents

List of Tables	iv
List of Figures	v
Acknowledgments	vii
Declarations	viii
Abstract	ix
Chapter 1 Introduction	1
1.1 The Study of an Ever-changing Universe	1
1.2 An Introduction to Black Holes	3
1.2.1 Supermassive Black Holes	6
1.3 Active Galactic Nuclei	10
1.3.1 Background	11
1.3.2 AGN Variability	15
1.3.3 Identifying Active Hosts	19
1.4 Tidal Disruption Flares	20
1.5 Relativistic Tidal Disruption Flares	30
1.5.1 A connection between long-lived gamma-ray events?	34
1.6 Supernovae	36
1.6.1 Superluminous Supernovae	40
1.7 Summary of Transient Properties	41
Chapter 2 Methods and Observations	44
2.1 Detectors across the electromagnetic spectrum	44
2.1.1 UV, Optical and Near-Infrared detectors	44
2.1.2 X-ray and Gamma-ray detectors	47
2.2 Reduction and Analysis Techniques	49

2.2.1	Optical Image Reduction	49
2.2.2	Optical Spectroscopic Reduction	56
2.3	Astrometry	58
2.3.1	Astrometric Matching	58
2.3.2	Image Subtraction	60
2.3.3	Modelling Galaxy Morphology and Determining Centroid Positions	61
2.4	Magnitude Systems and Cosmology	62

Chapter 3 The Candidate Relativistic Tidal Disruption Flare Swift

J1112.2-8238		64
3.1	Introduction	64
3.2	Observations	65
3.2.1	Swift BAT data	65
3.2.2	X-ray data	66
3.2.3	Optical Imaging	67
3.2.4	Spectroscopy	75
3.3	Discussion	77
3.3.1	Physical properties	77
3.3.2	Comparison to other sources	78
3.4	Implications	83
3.5	Summary	87

Chapter 4 Late Time Observations of Swift J1112.2-8238

4.1	Introduction	88
4.2	Observations	89
4.2.1	HST Imaging	89
4.2.2	X-Shooter Spectroscopy	93
4.2.3	Late-time Radio Observations of the Host	101
4.3	Discussion	104
4.3.1	Host Morphology and Transient Position	104
4.3.2	Internal Extinction	106
4.3.3	Metallicity and Classification of the Host	107
4.3.4	Stellar Mass and Black Hole Mass	109
4.3.5	UV/Optical Star Formation Rate Estimates	110
4.3.6	Properties of the Radio Emission	111
4.4	Implications for the interpretation of the flare	115
4.5	Summary	118

Chapter 5	A Study of Four Unusual Nuclear Transients	119
5.1	Introduction	119
5.2	Observations	121
5.2.1	HST imaging	121
5.2.2	CSS100217: Optical Observations	130
5.2.3	CSS100217: Swift Observations	130
5.3	Discussion	135
5.3.1	Astrometry of the Transients	135
5.3.2	Evolution of the Transients	138
5.4	Implications for the origins of the flares	143
5.4.1	ASASSN14ae	143
5.4.2	ASASSN14li	144
5.4.3	CSS100217	145
5.4.4	ASASSN15lh	147
5.5	A Possible New Class of AGN flare or a Sub-Class of TDF?	148
5.6	Summary	151
Chapter 6	Conclusions	153
6.1	Summary of Results	153
6.1.1	The Relativistic Tidal Disruption Flare Candidate: Swift J1112.2-8238	153
6.1.2	The Late-time Analysis of Swift J1112-8238	155
6.1.3	An Expanding Variety of Nuclear Flares	156
6.2	Overview and Future Prospects	157

List of Tables

3.1	GMOS and FORS2 photometry of Swift J1112-8238	73
4.1	<i>HST</i> photometry of the host of Swift J1112-8238, broken down by component	94
4.2	The emission line fluxes determined from the X-Shooter spectrum of the host of Swift J1112-8238	101
4.3	The radio fluxes determined from the ATCA observations of Swift J1112-8238 and Swift J2058+05	104
4.4	The inferred star formation rate estimates for the components of the host of Swift J1112-8238	111
5.1	The key features of the <i>HST</i> observations	125
5.2	The <i>HST</i> photometry of the host and transient emission of ASASSN14ae, ASASSN14li, CSS100217 and ASASSN15lh	128
5.3	<i>Swift</i> -UVOT photometry of CSS100217	131
5.4	The angular distances and physical offsets of the host and transient centroids for ASASSN14ae, ASASSN14li and ASASSN15lh	137

List of Figures

1.1	Typical spectra of active galactic nuclei	12
1.2	The unified model of AGN	14
1.3	A typical multiwavelength AGN lightcurve	17
1.4	Baldwin Phillips & Terlevich diagrams for SDSS galaxies	21
1.5	Three possible configurations of the infall of a star towards a super-massive black hole	23
1.6	The lightcurves of <i>ROSAT</i> detected tidal disruption flares	26
1.7	Continuum subtracted spectra from a range of tidal disruption flare candidates showing a range of line strengths that indicate a variety of disrupted stars	29
1.8	The X-ray Luminosity - Optical Absolute Magnitude phase space for a selection of extragalactic transients and AGN	32
1.9	The T_{90} and hardness ratios for GRBs detected by <i>CGRO</i> , <i>BATSE</i> and <i>Swift</i>	35
1.10	A summary of the classes of supernovae	38
1.11	Typical lightcurves of the various classes of supernovae	40
2.1	Example of an optical image before reduction	49
2.2	Examples of the calibration frames used in the optical image reduction procedure	52
2.3	Example of a reduced optical image	53
3.1	The <i>Swift</i> -XRT spectrum of Swift J1112-8238	66
3.2	The hardness ratio evolution of <i>Swift</i> J1112-8238 as observed by <i>Swift</i> XRT	68
3.3	The finding chart for Swift J1112-8238	69
3.4	Examples of the methods used to subtract the contaminating star in the optical imaging of Swift J1112-8238	72

3.5	The multi-waveband lightcurve for Swift J1112-8238, with the X-ray lightcurves of Swift J1644+57 and Swift J2058+05 for comparison	74
3.6	The GMOS and FORS2 spectra of Swift J1112-8238	76
3.7	Swift J1112-8238 placed within the X-ray Luminosity - Optical Absolute Magnitude phase space	84
3.8	The optical image of Swift J1112-8238 with the positions of the transient and host centroid indicated	85
4.1	<i>HST</i> WFC3 images of the host of Swift J1112-8238 in the F606W and F160W bands	90
4.2	A composite image of the <i>HST</i> WFC3 images	91
4.3	The X-Shooter spectrum centred on the [OII] $\lambda\lambda$ 3726, 3729 emission doublet	96
4.4	The X-Shooter spectrum centred on the [OIII] λ 5007 emission doublet	97
4.5	The X-Shooter spectrum centred on the H α emission line	98
4.6	The X-Shooter spectrum centred on the H β emission line	99
4.7	The X-Shooter spectrum centred on the H γ emission line	100
4.8	The early epoch radio emission contours overlayed on the F160W <i>HST</i> image of Swift J1112-8238	103
4.9	The Baldwin Phillips & Terlevich (BPT) diagram of the bulge component of the host of Swift J1112-8238	108
4.10	The radio lightcurves of the three relativistic tidal disruption flare candidates	113
4.11	The spectral index evolution of the three relativistic tidal disruption flare candidates	114
5.1	The early and late epoch <i>HST</i> WFC3 images of ASASSN14ae and ASASSN14li and the corresponding subtractions	122
5.2	The <i>HST</i> WFC3 image of CSS100217 in the F763M filter and the point source subtraction	123
5.3	The early and late epoch <i>HST</i> WFC3 images of ASASSN15lh and the corresponding subtraction	124
5.4	The long-term, multiwaveband lightcurve of CSS100217	132
5.5	The <i>Swift</i> -XRT spectrum of CSS100217	134
5.6	The <i>Swift</i> -UVOT and <i>HST</i> lightcurve of ASASSN14ae	139
5.7	The <i>Swift</i> -UVOT and <i>HST</i> lightcurve of ASASSN14li	141
5.8	The lightcurves of the 5 most luminous supernovae to date	149

Acknowledgments

Through the four years of my PhD here at Warwick I have received a great deal of support from my colleagues, friends and family. Here I would like to briefly thank those who have made this work possible.

I would first and foremost like to thank my supervisors, Andrew Levan and Elizabeth Stanway for their constant support and bountiful patience throughout the last four years. Their time, knowledge and experience has been greatly appreciated, providing me the direction and confidence I needed to produce this thesis.

I would further like to thank the entire Warwick Astronomy group for their welcoming and helpful nature that has made my time here a pleasure. In just the few years I've been here, I've watched the group double in size, with each intake providing yet more wonderful people to the expanding Warwick family. In particular I should like to thank: Rachel for her patience in guiding a lowly first year away from some major pitfalls, while apologising to her for nonetheless falling into them; Penelope for her inextinguishable sense of humor and immoral support; Steph for her kind heart and fantastic conversation (whatever the topic...); Tom who's knowledge and advice has helped me deal with many a knotted Python (script); and Charlotte for her ability to take in vast quantities of caffeine in the pursuit of a (semi-)balanced life. From the original office and the cool office to the lame office and the new kids on the block, without any of you this experience would not have been half of what it was. I thank you all.

Finally, I thank my parents and brother for never giving up on me and always pushing me towards achieving my potential. Without their unwavering support, this would be nothing more than a book of blank pages.

I leave it to the reader to decide if that would have been preferable.

Declarations

Here I declare the extent of original work completed within this thesis and that of the work that has been published. This thesis has been composed by the author and has not been submitted in any previous application for any degree.

Chapter 3 is based on a paper published in the *Monthly Notices of the Royal Astronomical Society* titled “*Swift J1112.2-8238: A candidate relativistic tidal disruption flare*” (Brown et al., 2015). All reduction and analysis within was completed by the author, while Figure 3.7 was produced by Andrew Levan. The main collaborators on this work were Andrew Levan and Elizabeth Stanway.

The radio observations detailed in Chapter 4 were reduced by Elizabeth Stanway, who also provided Figures 4.8, 4.10 and 4.11. The X-Shooter spectroscopy was reduced by the author and independently by Thomas Krühler. The *HST* reduction and all other analysis was completed by the author.

In Chapter 5, the *Swift*-UVOT reduction for CSS100217 was completed by Samantha Oates. All other analysis was completed by the author.

Abstract

The analysis of nuclear flares is unfortunately difficult. Contamination of supernova surveys by active galactic nuclei (AGN) variability, and the difficulty in detecting transients in the high surface brightness nuclei of galaxies, has led to many surveys avoiding nuclear transients entirely. Even in cases where transients are detected, their identification and classification remains complex, with many possible progenitor pathways, overlapping models and wide ranges of observed properties to explain. Here I consider a sample of these events, placing them in the wider context of transient astronomy.

The detection of a class of relativistic tidal disruption flare, thought to be the capture and disruption of Sun-like stars that also powers a moderately relativistic jet, has prompted the search for more of these events. Within this work I analyse the properties of one such candidate, Swift J1112.2-8238, confirming its extragalactic origin and showing it came from a galaxy at a redshift of $z = 0.89$. Its high energy and optical properties are consistent with the previous candidates and its position, close to the centre of a likely star-forming host, continues to support the tidal disruption flare origin of these events. The rates of these events suggest that only a small fraction of tidal disruption flares launch similar jets.

Prompted by these findings, I proposed and obtained medium resolution spectroscopy and radio observations of the source, and analysed high-resolution *HST* imaging to further constrain the position of the transient within its host. The *HST* imaging shows that the host has a complex morphology, perhaps due to an interaction with another galaxy, with the transient loosely consistent with the centre of compact bulge-like component. I confirm the host's redshift and determine its nature as a star-forming galaxy. Radio emission detected coming from the host, that is too luminous to be associated with star-formation, shows evidence of variability, suggesting that it is associated with the transient flare and thus is perhaps confirmation of the jetted nature of the event. In almost all respects, Swift J1112.2-8238 remains an excellent candidate relativistic tidal disruption flare.

Finally, I analyse *HST* imaging of a number of flares with unusual properties. I greatly improve the astrometric tie of ASASSN14ae and ASASSN14li to the nuclear regions of their hosts and show that their properties are still most consistent with a tidal disruption flare origin. In the case of CSS100217 and ASASSN15lh however, the “accepted” classification of their origins as superluminous supernovae appears to be at odds with their host galaxies, with ASASSN15lh in particular coming from a massive host with minimal star-formation. I show that CSS100217 has undergone a significant drop in apparent quiescent-level emission following the flare, indicating the possibility that the flare may have directly impacted, or been caused by, a change in the accretion of the known AGN. I consider the possibility that both flares could be associated with unusual tidal disruption flares or AGN variability, though the current observations make it difficult to make strong claims about either flare's true origins.

Chapter 1

Introduction

1.1 The Study of an Ever-changing Universe

Our Universe is an ever-changing place. However, looking up at the night sky, it is easy to see why ancient civilizations often thought of the heavens as a constant realm. Naked eye events such as Galactic supernovae are rare, with only a small handful identified. Instead, as we now know, the sky is full of short-lived events. Nowhere is this more obvious than at γ -ray and X-ray wavelengths. In this regime, single short-lived events, lasting only a few seconds, can briefly outshine the rest of the Universe combined (Bloom, 2011). Even at optical wavelengths the sky is far from calm. Advances in technology have made possible the implementation of deep, high cadence surveys such as the Panoramic Survey Telescope and Rapid Response System (Pan-STARRS; Kaiser et al., 2002) and the Palomar Transient Factory (PTF; Law et al., 2009), scanning vast swathes of the sky to detect thousands of transients per year. These numbers are only set to increase with the coming Large Synoptic Survey Telescope (LSST; Ivezić and LSST Science Collaboration, 2013) era, when extragalactic transient rates alone are expected to be in excess of hundreds of thousands per year (LSST Science Collaboration et al.).

With the ever increasing number of transients comes the need for efficient classification and identification, all with a view to eventually understanding the cause of the event. As the detection rate of events has increased, and instrumentation and follow-up has improved, subtle differences in event and host properties have become apparent, prompting the definition of more and more classes and subclasses of event. This has forced theorists and observers alike to produce and test models to explain a growing variety of properties, from varying lightcurve decay rates, the presence, strengths and widths of spectral features and multi-wavelength detections

from γ -rays to radio. Major difficulties arise, however, whenever observed properties can be re-produced by a range of different progenitor mechanisms. This is of particular concern when transients are found in regions of parameter space where multiple pathways are likely and properties are expected to be peculiar.

Nowhere is this more likely to be the case, perhaps, than in the nuclear regions of galaxies. It is generally accepted that at the centre of most, if not all, large galaxies, there exists a supermassive black hole. Their presence makes an association of any nuclear flare with these extreme objects a strong possibility. Whether in one-off short-lived events from the shredding of stars (Rees, 1988) or from outbursts in long-lived active nuclei (e.g. Peterson, 2001, and references therein), transient behaviour is an unavoidable consequence of these behemoths.

Further transient behaviour that is not directly associated with the supermassive black hole can come from core collapse supernovae, a consequence of the considerable star formation that occurs within these regions. Indeed, $\sim 10\%$ of the star formation in our own Galaxy comes from a region only 500 pc in radius surrounding the Galactic centre and contains more high mass ($> 100M_{\odot}$) stars than anywhere else in the Galaxy (Figer, 2008). It is possible that this abundance of high mass stars is a result of an unusual initial mass function (IMF), the probability distribution (and thus, in large enough numbers, the population) of masses of stars upon beginning their main sequence lifetimes. While often thought to be relatively invariant throughout the local Universe at least, there is some evidence of variation of the form of the IMF with galactic environment (e.g. Geha et al., 2013). Indeed there is evidence that star formation in the extreme environments surrounding the central supermassive black hole can produce top heavy initial mass functions (e.g. Morris, 1993) as may be evidenced in our own Galaxy (Bartko et al., 2010), which, producing an over-abundance of massive stars, would naturally lead to a large number of core collapse events.

Unfortunately, while large samples of nuclear events would greatly help discern the properties and relative rates of separate populations, some surveys deliberately avoid the nuclear regions of galaxies due to the common confusion of black hole-related variability with the types of events (usually supernovae) that these surveys tend to focus on. Even when surveys are not designed such, the detection of these flares can still be very difficult. High line-of-sight gas and dust columns can produce significant extinction at optical and UV wavelengths, prompting attempts to look instead for nuclear supernovae in the infrared (e.g. Mattila et al., 2007). In addition, the luminous nature of the nuclear regions of galaxies can mask the presence of a transient through their high surface brightnesses and the resultant

large Poisson noise, making detection via photometry or image subtraction difficult. However, with recent and ongoing surveys such as the All Sky Automated Survey for Supernovae (ASASSN¹) and the *Swift* Gamma-Ray Burst Mission (Gehrels et al., 2005) producing a number of nuclear events, there is a great need to understand the natures of these transients.

Of particular interest are the tidal disruption flares (TDFs), outbursts that represent the disruption and accretion of a star by the central supermassive black hole (Rees, 1988). These events represent a possible route to the study of the supermassive black holes in distant, dwarf and inactive galaxies where other methods are impractical or observationally expensive. Further, the discovery of a possible variant of TDFs with relativistic jets (Levan et al., 2011; Bloom et al., 2011; Zauderer et al., 2011) opens a door to the study of jetted accretion analogous to long-lived blazar emission from onset to termination on human timescales.

This thesis covers the analysis of several exotic nuclear transients with a view to constraining the likely progenitor pathways in each case and bringing into focus the key difficulties in identifying distinguishing features in the properties of these extreme flares. I will also highlight the need for rapid and repeated multiwavelength follow-up observations of candidate flares and for an understanding of the properties of the underlying host as a key indicator of the event’s true origin. To this end I begin with an introduction to the topic of black holes (Section 1.2), focussing on the supermassive variant that is the source of tidal disruption flares that are of most interest in this study. I consider the properties and identification of a common source of confusion in both supernova surveys and the search for TDFs, active galactic nucleus variability (Section 1.3), before discussing the properties of the tidal disruption of stars (Section 1.4) and the discovery of the possible relativistic variant (Section 1.5). I discuss the properties of the nuclear supernovae that dwarf the rates of tidal disruption flares in all but the highest resolution imaging (Section 1.6) and finally I summarise the key properties of expected nuclear transients that form the basis for comparison throughout the rest of the study (Section 1.7).

1.2 An Introduction to Black Holes

The first suggestion of the existence of a “dark star” (the term “black hole” being coined in the mid 1960’s; Ewing, 1964) is purported to have occurred as early as 1783 by John Michell. He pointed out that, in the Newtonian model of gravity and light, the gravitational attraction of a star would slow down any light particles leaving

¹<http://www.astronomy.ohio-state.edu/~assassin/index.shtml>

it, with the extreme case of a sufficiently massive star causing light to halt and fall back, effectively making it impossible for light to escape (Hockey et al., 2007). This view, while missing the mark in its explanation of the effect of gravity on light, nonetheless correctly suggested that sufficiently massive objects would necessarily appear dark.

It was not until Einstein's famous theories of relativity in the early 20th century that the modern formalism of a black hole was produced. Within the scope of the theory of General Relativity, gravity is described as a curvature of spacetime, with objects of mass distorting space around them. The commonly used analogy for this is that of a taut rubber sheet, where placing objects at the centre causes a deformation proportional to the mass of the object. Any other object moving around the sheet then follows a trajectory that curves in towards the centre. Even massless particles, such as photons, are bound to this curvature. Over time it was realised that, in the case of a sufficiently dense object, there would be a radius within which the exceptional gravitational potential would produce a curvature too strong even for light to escape. The surface defined by this condition is known as the event horizon and within the special case of a non-rotating black hole, it occurs at the Schwarzschild radius, R_S . This Schwarzschild radius is defined in terms of the universal gravitational constant, G , the mass of the object, M , and the speed of light, c , as:

$$R_S = \frac{2GM}{c^2} \quad (1.1)$$

As there is no object or emission that can overcome this boundary, it is impossible to observe any object that crosses it. This includes the central singularity of the black hole, a point where physical properties such as density take on infeasible values under our current understanding of physics. Due to the peculiarities of the effects of such intense gravitational fields on an external observer's perception of time, objects crossing the event horizon will never be observed doing so, instead asymptotically closing in on the event horizon, with emission from them slowly redshifting out of existence. However, in reality, the object does cross in a finite proper time, that observed from the perspective of the object. Outside of the event horizon, stable orbits that are resistant to perturbations are possible up to the innermost stable circular orbit, which for a non-rotating black hole is at $3R_S$. Unstable freefall orbits are possible closer to the event horizon so long as they do not cross the surface defined by the photon sphere ($1.5R_S$ for a non-rotating black hole). Any closer approach requires that the object be capable of producing a balancing force that counters the intense gravitational pull of the black hole. Failing this, in all but

the fastest spinning black holes which are theoretically capable of supporting orbits within the event horizon, the object is fated to fall into the singularity.

Despite the physics of black holes and the effect on their environments being undeniably complicated, they can in some ways be thought of as some of the simplest objects in the Universe. Under the “no-hair” theorem, the properties of a black hole can be described in their entirety by only three parameters: the black hole’s mass, its spin or angular momentum, and its electrical charge. The last parameter is only important in exceptionally small black holes where Coulomb forces become comparable in magnitude to the black hole’s gravitational field and charge imbalances cannot easily be neutralised, a regime which is unlikely to be of astrophysical importance. Instead, astrophysical black holes can be considered only in terms of their mass (Schwarzschild black holes) or a combination of mass and spin (Kerr black holes).

Most astrophysical black holes are thought to be formed as a result of the deaths of massive stars. Low mass stars, which live extremely long main sequence life-times powered by the fusion of hydrogen into helium, eventually evolve to the point where they blow off their outer layers leaving behind a hot dense core of material that is supported against further collapse by electron degeneracy pressure. These objects, known as white dwarfs, will, barring any later mass transfer, slowly cool indefinitely. More massive stars have shorter main sequence lifetimes and are capable of fusing progressively more massive elements. Stars in excess of about $8M_{\odot}$ (though in reality the true limit is somewhat dependent on metallicity; Smartt, 2009) can continue to fuse material up to iron. The binding energy of iron is such that continued fusion is an endothermic process. This produces a sharp decrease in the temperature of the core, and so the thermal pressure holding the star from collapse ceases. This leads to a collapse of the inner core and an infalling of the outer envelope. These outer layers may then violently rebound producing a bright core-collapse supernova. The pressure on the central core is extreme enough to overcome electron degeneracy pressure, forcing protons to merge with electrons and forming an object with the density of an atomic nucleus completely composed of neutrons. This neutron star is once again held from collapsing further, this time through neutron degeneracy pressure and can maintain this state up to the Tolman-Oppenheimer-Volkoff (TOV) limit of $\sim 3M_{\odot}$ (Bombaci, 1996), though the exact conditions of this collapse are somewhat debated. At this point, barring the existence of currently purely theoretical quark (Ivanenko and Kurdgelaidze, 1965) or electroweak stars (Dai et al., 2010), there is no known force capable of holding up the remnant and thus its collapse into a singularity is inescapable. Neutron stars can also eventually

become black holes through mergers with other neutron stars or through accretion, with evidence suggesting that some of these events are the cause of the short variant of gamma-ray bursts (Bloom, 2011). In both the merger and massive star cases, the resulting black hole is of the order a few to a few tens of solar masses.

1.2.1 Supermassive Black Holes

However, at the centres of most, if not all, large galaxies exists an extremely massive black hole with a mass in the range $10^5 - 10^{10} M_{\odot}$. A number of mysteries continue to surround these extreme objects, not least of which concerns their origins. Multiple possible pathways have been formulated that may lead to the production of supermassive black holes, though none of them naturally lead to the extreme masses of SMBHs observed in the local Universe. Instead black hole seeds produced by these methods must have undergone considerable growth. These methods include:

- *Stellar Collapse*: The first population of stars, built from pristine material from the Big Bang with little to no trace of metals (here defined as any element of atomic number greater than that of helium) are known as Population III (Pop III) stars. Their extremely low metallicities enable the production of high initial mass stars and greatly reduce the mass loss due to primarily metal-driven stellar winds during their lifetimes (see Bromm and Larson, 2004, for a review). Though certain initial stellar mass ranges are expected to result in pair instability supernovae that leave no remnant, outside of these, the exceptional core masses produced in these early stars could produce far more massive black hole remnants upon core collapse than the stellar-mass black holes produced in higher metallicity environments later on in the Universe's lifetime. However, a lack of constraints on the initial mass functions (IMFs) of Pop III stars makes it difficult to determine the viability of this pathway. While some studies suggest that stellar masses of $10^2 - 10^3 M_{\odot}$ are likely to be relatively common (Hirano et al., 2014; Susa et al., 2014), others have suggested the IMF may be far less top heavy, with most stars inhabiting the tens of solar masses regime and few stars exceeding $100 M_{\odot}$ (Hosokawa et al., 2011). A more exotic pathway might involve stars held up by the annihilation of dark matter with initial masses of $500-1000 M_{\odot}$, producing even more massive black holes (e.g. Spolyar et al., 2008).
- *Direct collapse*: Under certain conditions, the gas clouds that would normally be the building blocks of star formation could collapse directly into massive black holes, perhaps through a supermassive star stage. This requires that

the gas not be subject to the usual fragmentation that occurs in collapsing molecular gas clouds. This in turn requires that the gas be exceptionally low metallicity, remain atomic and have a high virial temperature ($> 10^4\text{K}$), a situation that provides a far slower cooling rate than molecular hydrogen and suppresses star formation. The resulting supermassive star ($> 5 \times 10^4 M_\odot$) would eventually collapse, forming the supermassive black hole (Volonteri, 2010).

- *Stellar Mergers:* One possibility involves the merger of stars in regions of high stellar density such as compact nuclear clusters (Devecchi and Volonteri, 2009). This process could be important in regions that have already had Pop III formation, enriching the available gas above the level of the direct collapse mechanism.
- *Primordial Black Holes:* It is also possible that density fluctuations in the early Universe may have produced regions where the gravitational potential overcame the local pressure, forming primordial black holes over a large range of mass scales (Khlopov et al., 2005).

Certain mechanisms have predicted observable properties, such as the possibility of ultra-long gamma-ray bursts from the deaths of direct collapse black holes (Matsumoto et al., 2015). However, due to the individual final products of each being effectively identical, distinguishing these pathways will likely require an analysis of the SMBH population from the local universe back to extremely early times.

Growth of these black hole seeds can occur through a combination of mergers with other massive black holes, perhaps during the interactions and mergers of their host galaxies, and through accretion of diffuse material. This latter method of growth has an inherent limit, at least for the accretion of baryonic matter. Radiation pressure, the repulsive force of incident radiation on an absorbing or scattering material, will repel any material falling onto the emitting body. If the emission produced is allowed to become powerful enough for the radiation pressure to exceed the gravitational pull on the infalling material, the material can be pushed away. The luminosity at which this occurs is known as the Eddington luminosity, L_{Edd} , and, under the assumption of spherical symmetry both for the infalling material and the outgoing emission, is given by:

$$L_{\text{Edd}} = \frac{4\pi GMc}{\kappa} \quad (1.2)$$

where G is the Universal gravitational constant, M is the mass of the black hole, c is

the speed of light and κ is the opacity of the accreting material. While determining the opacity is difficult, if we assume that the infalling material is mostly formed of ionized hydrogen, and thus the opacity is dominated by Thompson scattering from electrons, we find:

$$L_{\text{Edd}} = \frac{4\pi GMcm_{\text{p}}}{\sigma_{\text{T}}} \quad (1.3)$$

where m_{p} is the mass of a proton and σ_{T} is the Thompson scattering cross-section for an electron. This becomes a problem when the emission is itself a product of the infalling material, as is the case for accretion onto a compact object, where a portion of the infalling material's gravitational potential energy is converted to radiation. Expressing this as a fraction, ε , of its rest mass energy, the luminosity from accretion, L_{acc} , is:

$$L_{\text{acc}} = \varepsilon \dot{M} c^2 \quad (1.4)$$

When the radiation reaches the Eddington luminosity stated above, the accreting matter will be blown away, accretion ceases and thus so does the emission. As a result, the accretion is intrinsically limited to a certain rate, \dot{M}_{Edd} , given by:

$$\dot{M}_{\text{Edd}} = \frac{4\pi GMm_{\text{p}}}{\varepsilon c \sigma_{\text{T}}} \quad (1.5)$$

Note that the maximal feeding rate scales with black hole mass.

From equation 1.5, assuming negligible contributions from merger events, a typical radiative efficiency $\varepsilon = 0.1$ and given an initial black hole seed mass, it is possible to determine the minimum time the seed can grow to any given mass². For a black hole seed of 10^2 (10^5) solar masses, it would take approximately 1 (0.5) Gyr to grow to $10^9 M_{\odot}$ (Volonteri, 2010).

However, SMBHs with masses of $10^9 M_{\odot}$ or more are known to exist at $z > 6$ from the detection of quasars in part by the Sloan Digital Sky Survey (see a review by Fan, 2006), with more recent surveys continuing to add to the list of examples (e.g. Venemans et al., 2015; Matsuoka et al., 2016). In fact one such object is known to exist at $z = 7.085$, only 0.77 billion years after the Big Bang (Mortlock et al., 2011), requiring that SMBH seeds already exist during or shortly after the cosmic ‘‘Dark Ages’’, the time of the first generation of stars. They must also exist in sufficient numbers to explain both the observed population of quasars and the rest

²though it should be noted that non-spherical geometries, where emission and accretion effectively occur apart from one another, can produce situations where super-Eddington accretion is possible, such as in tidal disruption flares (Section 1.4) and long gamma-ray bursts

of the proverbial iceberg, black holes that are too dim to be detected at present or are simply not accreting.

This problem is further exacerbated by the fact that the assumption of continuous feeding throughout the black hole’s lifetime requires a constant supply of material from the host. However observations do not support this with numerous studies placing accretion period lifetimes at between $10^6 - 10^8$ years (e.g. see Martini, 2003; Hopkins and Hernquist, 2009). It is possible that this may be caused by a simple lack of available material in the host (starvation), or due to a feedback process from the feeding itself. For example, the radiation from accretion, while insufficient to repel the material through radiation pressure, could nonetheless heat the gas clouds feeding it, the resultant expansion removing material from the black hole’s reach until it cools again (Alvarez et al., 2009). In any case, these “short”, intermittent periods of accretion would further lengthen the growing process making small but early black hole seeds more difficult to reconcile with observations. This situation could be circumvented with the possibility of super-Eddington accretion based on a spatial separation of the accretion and radiation (Shaviv, 1998), very low radiative efficiencies (less light, and thus radiation pressure, for the same accretion rate; Pacucci et al., 2015) or much more massive initial seeds with the last solution being strongly dependent on which of the above seed production mechanisms is viable/dominant in forming the SMBHs seen today.

While exceptionally massive objects, these supermassive black holes (SMBHs) still generally contribute only a tiny fraction of the total host galaxy mass. As such, the region over which they dominate the gravitational potential of their host galaxy, and thus the dynamics of their environments, is still only a meagre 1-100pc (Kormendy and Ho, 2013). Despite this small sphere-of-influence, the masses of SMBHs have been found to correlate with the wider properties of their hosts, far beyond the grasp of the black hole’s gravitational field. In many disc or spiral galaxies there exists a bulge component, a central section of the galaxy that shares a morphological similarity with elliptical galaxies (Renzini, 1999). Their similarity comes from their shared origin in the merger of smaller galaxies, the near sole distinguishing feature being the later development of a disk around what we now refer to as the bulge. The stellar masses of these bulges have been found to scale almost directly with the central SMBH mass, at a ratio of approximately 1000 to 1 (Haring and Rix, 2004; Kormendy and Ho, 2013). As would be expected assuming a constant stellar mass to light ratio, the luminosity of these bulges also follows this trend. More intriguingly, the galaxy stellar velocity dispersion (σ_*), the line of sight broadening of emission lines through the motion of stars in the galaxy, is also found to tightly

correlate with the SMBH mass. This $M - \sigma_*$ relation is found to be approximately:

$$M \sim 2 \times 10^7 M_\odot \sigma_{100}^\alpha \quad (1.6)$$

where M is the black hole mass, $\sigma_* = 100 \sigma_{100} \text{km s}^{-1}$ and $\alpha \sim 4.4 \pm 0.3$ (Ferrarese and Merritt, 2000; Gebhardt et al., 2000; Kormendy and Ho, 2013). This seems to indicate an element of coevolution, either with the host galaxy affecting the feeding, and thus the growth, of the black hole, or the black hole having influence over its host far beyond the range of its sphere of influence, or indeed a combination of the two. While the exact cause for these scaling relations is as yet unknown, it is thought to be due to a feedback process between the host and the emission produced during the feeding of these massive objects.

With clear links to galaxy formation and evolution, as well as the interest surrounding the potential for the detection of gravitational waves during SMBH binary mergers, these objects represent an important subject of study for a number of fields of astronomy. The characterization of the population of supermassive black holes, particularly at the low mass end and in the high redshift Universe, is important in order to understand their origins and evolution. However, a number of typical methods for their detection and characterization fail in small and distant galaxies. While the tracing of stellar orbits around an invisible compact source with a mass of $\sim 4 \times 10^6 M_\odot$ has provided by far the best evidence for a supermassive black hole in our own Galaxy (Ghez et al., 2008), this method is impossible to complete in any other host. Instead it is possible to observe stellar and gas dynamics within the sphere of influence of the black hole by looking for a characteristic rise in Doppler broadening of spectral features in spatially resolved spectroscopy across the face of the galaxy (Peterson, 2014). However, this method too is limited by distance to the host and by supermassive black hole mass, where if the sphere of influence has too small an angular size to be spatially resolved, confusion with other slower moving material further out from the SMBH will mask its influence. In galaxies where direct observations of the dynamics of the black hole's neighbourhood fail, attempts to characterize the SMBH population must rely on other mechanisms. These involve observing the emission that occurs when material is accreted onto the SMBH, the process that powers both active galactic nuclei and tidal disruption flares.

1.3 Active Galactic Nuclei

As the observational evidence of ongoing SMBH growth, active galactic nuclei (AGN) have provided extremely valuable insights into the central regions of galaxies

and added to our understanding of supermassive black holes themselves. However, the frequent variability of their emission makes them a significant source of contamination in transient surveys. Thus an understanding of the observational characteristics of AGN and their variability is key to the determination of the origins of nuclear transients.

1.3.1 Background

Our current understanding of active galactic nuclei (AGN) comes from an impressive effort to produce a unification of seemingly disparate and unrelated observations of what were at the time some of the most peculiar objects in the Universe (Antonucci, 1993). It has helped to explain a wide variety of phenomena, from the early 20th century observations of “spiral nebulae”³ with unusually bright and broad emission lines (Seyfert, 1943) to the post World War II detection of peculiar spatially-resolved two component radio emission (Jennison and Das Gupta, 1953). Eventually, as more pieces of the puzzle were found, such as the realisation of their extragalactic nature (e.g. Hazard et al., 1963; Schmidt, 1963; Oke, 1963; Greenstein and Matthews, 1963), it was found that each observation could be explained in the context of some form of central engine directly producing or illuminating various structures in the surrounding environment. The observation of short-term variability (see Subsection 1.3.2) required the engine be extremely compact ($< 10^{15}$ cm) and yet produce extreme luminosities (quasars often exceed $\sim 10^{46}$ erg s⁻¹). Attempts to explain this through dense stellar clusters inevitably resulted in systems where gravitational contraction dominated energy production (Lynden-Bell, 1969) and eventually a SMBH explanation was adopted that has survived until today. After briefly flirting with the concept of feeding the SMBH through a regular supply of stars for disruption (e.g. Hills, 1975), the current view of a quasi-steady flow of diffuse material was reached. It is this feeding that is thought to be the primary mechanism of SMBH growth throughout the history of the Universe and may in part provide the feedback mechanism responsible for the observed scaling relations of SMBHs and their hosts.

Broadly speaking, the unification of AGN has occurred on two fronts: one explaining much of the infrared through to the X-ray properties, the other focussing on the radio-loud/quiet dichotomy (see Antonucci, 1993; Shields, 1999). Many AGN have some combination of a bright, central point source with a non-stellar spectrum, broad emission lines with widths of $1000 - 20000$ km s⁻¹ and strong narrow emission lines with widths of $300 - 1000$ km s⁻¹, which is still wider than typical star-forming

³now known to be spiral galaxies, the debate over their Galactic or extragalactic origins not having been resolved at the time




Figure 1.1: Typical spectra for Type 1 and Type 2 AGN. Note the presence of narrow line emission in both spectra with the addition of much wider features in the Type 1 spectrum. This corresponds to the different inclinations at which each galaxy is observed, the Seyfert 2 galaxy’s broad line emission being obscured by a wide dusty torus. In both cases, the line emission is stronger and broader than a typical star-forming galaxy. Figure adapted from (Trump et al., 2009)

galaxies. In the simplest version of the unification model, these properties, or lack thereof, depend on the inclination of the AGN to the observer (see Figure 1.2). The central point source is a direct view of the accretion disk and the immediate environment of the black hole. The strong thermal emission and an infrared power-law from the accretion disk is accompanied by intense X-ray emission from a hot corona of electrons (Haardt and Maraschi, 1991). A group of high-density gas clouds, irradiated by the accretion disk and orbiting the SMBH at luminosity dependent distances of between 0.01 and 1 parsec, give rise to the broadest emission lines. Their high orbital velocities produce large Doppler shifts while their high densities preferentially produce the permitted and semi-forbidden emission lines. Beyond this broad-line region (BLR), a further group of lower density clouds make up the narrow-line region (NLR), orbiting at much greater distances and thus producing more moderate Doppler broadened lines. Their lower density also makes them more akin to the nebular gas clouds that produce the forbidden lines in star-forming galaxies. Finally, surrounding the central region is a dusty “torus”⁴ which is capable of obscuring much of the rest of the structure of the AGN. In this way, Seyfert 1

⁴though in reality its true morphology is difficult to determine

galaxies, with their broad emission lines and strong central point sources represent AGN viewed approximately face-on, bypassing the dusty torus. The Seyfert 2 category in contrast are AGN viewed at a higher inclination, missing most of the central point source except in X-rays and blocking the close in BLR from view, but still showing the narrow line emission from the NLR (see Figure 1.1). While more recent studies suggest additional luminosity dependent effects, the basic premise behind the unification of much of the mid to high frequency properties of AGN remains the same.

The picture is complicated further in the presence of a jet. Jets are a near ubiquitous feature of accretion with examples across all mass scales (e.g. De Gouveia Dal Pino, 2005). They are known to be present in a minority of AGN, though, when present, they can have a dramatic impact on their hosts. In the context of AGN they consist of collimated beams of fast moving material, often at large fractions of the speed of light, and stretching up to hundreds of kiloparsecs from the centres of their hosts. While the exact mechanism of their launching is not fully understood, in the case of an AGN it is thought to be produced through magnetic confinement of plasma by the magnetic field generated by the accretion disk, and may perhaps be partially powered by the angular momentum of the SMBH.

Jet emission covers a large portion of the electromagnetic spectrum, with strong emission from radio through to γ -ray wavelengths. Long-wavelength (radio to optical) emission comes from synchrotron emission (emission produced through the acceleration of electrons by magnetic fields) from within the core and the extended jet, while higher energy emission is a product of the upscattering (Inverse Compton scattering) of lower energy emission from the accretion disk. The kinetic power, the power contained within the bulk motion of the jet, can be vast, dwarfing the luminosity of the entire galaxy. This outflow can therefore have a great impact on the surrounding environment, potentially expelling molecular gas clouds and quenching star formation. However, the radiative efficiency of the jet is quite low, usually making the central regions of the AGN more intrinsically luminous (Marscher, 2010).

The jet can nonetheless appear to outshine the accretion disk in cases where the jet is oriented towards the observer. This is due to an effect known as relativistic beaming, where a combination of factors, including aberration⁵, the Doppler effect⁶ and time dilation⁷, result in an apparent increase in the luminosity of material

⁵otherwise known as the “headlight effect”, this is the tilting of emission toward an oncoming observer. It is analogous to the apparent angled descent of rainfall on a fast moving car when compared with a stationary one

⁶the shift of frequency of observed photons due to relative motion of the source and observer

⁷the apparent increase in the rate of photon production in the inertial reference frame of the observer due to time running faster in the reference frame of the fast-moving jet, a consequence of




Figure 1.2: The unified model of AGN. In this model, the majority of the UV and optical emission comes from the central accretion disk with X-ray emission from a surrounding corona. The close in broad-line region and more distant narrow-line region produce the broad and narrow line emission respectively, while the orientation of the observer with respect to an outer dusty torus dictates the observable properties of the central AGN. In the case of an AGN playing host to a jet, additional emission components are added and are considered to be separate from the whole in this simple unification model. Components associated with the jet can emit strongly across the electromagnetic spectrum but are particularly notable at radio and X/ γ -ray wavelengths. Figure from Urry and Padovani (1995).

moving towards the observer when compared with material that is stationary or receding. It is this effect that can result in situations where the anti-aligned double jet phenomena of AGN can be apparently reduced to a single forward facing jet, as while the first has been increased in apparent brightness, the second has been greatly reduced, each through beaming effects. The detection of both jets caused the peculiar, two-component radio detections of the 1950's and 1960's, with further advances since revealing additional structure including bright central point sources, shocks and cases where billowing lobes at the ends of each jet are punctuated by hotspots. The best studied jets of all are the most apparently luminous which tend to be those that benefit most from relativistic beaming, those that come within a few degrees of being observed on-axis. These cases, in addition to their high apparent luminosities, can also display some of the greatest variability in AGN (see Section 1.3.2).

While the various classes of AGN can be sometimes difficult to define, and are as a result defined somewhat arbitrarily, the fraction of AGN in the local Universe is undoubtedly quite low. About one in every thousand galaxies shows the properties of Seyfert galaxies, while the fraction of radio galaxies is a factor hundred smaller again (Osterbrock and Ferland, 2006; Mo et al., 2010). These numbers greatly increase at higher redshifts with studies showing an increase not only in the number of active nuclei, but also their intrinsic luminosity (Merloni and Heinz, 2013). However, with only a fraction of galaxies housing active nuclei, it is reasonable to ask the question how representative of the wider SMBH population these rare cases actually are. Nonetheless, their properties makes them useful objects to study on a number of fronts, with their high luminosities making them useful probes back to much earlier times in the Universe's history.

1.3.2 AGN Variability

Variability in AGN has been observed in every waveband they have been studied (Peterson, 2001). Even the earliest detections of quasars showed considerable variability (e.g. Matthews and Sandage, 1963). Some examples possess variability of a factor two in flux over extremely short timescales (days, hours or even minutes) whereas others may vary less than 10% over years of observations. Thus, while all AGN vary on some level, rapid and high amplitude variability is a common, but not ubiquitous, feature. While differences do exist from one object and waveband to another, broadly speaking any observed variability is stochastic (random) with little to no evidence of a dominant timescale or periodicity. In addition, multiple studies

special relativity

have shown correlation between variability in one waveband to another, sometimes with an associated time lag. However, the magnitude of this lag is not consistent between sources and even the order of the correlation (e.g. X-ray following optical or vice versa) has been shown to vary (e.g. Breedts et al., 2009, 2010).

For the moment putting to one side the extreme case of the most variable classes, AGN with significant observable central point source emission tend to show UV and optical variability on the order of 10-20% on timescales of a month, with smaller variability on day timescales (see Figure 1.3; Ulrich et al., 1997). The optical variability tends to be less pronounced in part due to a much larger, and unchanging, stellar continuum contribution. The variability is more dramatic at X-ray wavelengths where variability has been observed on timescales as short as hundreds of seconds (Mushotzky et al., 1993). This variability is thought to derive from the reprocessing of hard X-ray emission from the hot corona in the accretion disk. This is borne out by the shorter timescale of variability for the X-ray emission, with more smoothed out lower energy emission, and by observations of a time lag from X-ray to UV-optical variability.

This variability gives us considerable insight into the properties of the structure within AGN. For example, the shortest timescale of variability gives constraints on the emission region's size. Any change in emission coming from a region of radius x will inevitably be smeared out as viewed by an outside observer over a time $\frac{2x}{c}$, as the emission from the back of the region will arrive lagging behind that emitted from the front. As this assumes an instantaneous change across the whole region, and real variability is more likely to have intrinsic lag of its own, this therefore represents an upper bound on the emitting region size. For example, it was possible to determine that, in the case of an AGN without strong radio emission and thus no jet, evidence for a lack of variability in infrared emission on timescales shorter than a year supported the premise of a wide⁸ torus of dust that was re-emitting absorbed radiation. Similarly, the exceptionally small timescale of X-ray observations placed the strongest constraints on the size of the central engine, eventually leading to the SMBH explanation (Rees, 1977). However, this upper bound can be violated in situations of observer aligned motion, where the emitting region (e.g. the length of a jet) can seem to be shortened because the impetus for the change in emission (a shock or burst of energy) is coming from behind the emission region and is travelling with the resultant emission. The result is a compression of the signal as seen by the observer and thus appears to have a much smaller emission region size. However, this effect only significantly alters the apparent variability timescale

⁸when compared to the size of the central engine




Figure 1.3: A typical multiwavelength AGN lightcurve. Variability is clear at all wavelengths but is most evident at X-ray wavelengths and is considerably weaker at longer wavelengths. Also note the correlation between the variability in the different bands. Figure from Ulrich et al. (1997).

in situations where the motion is relativistic and thus the relative minority of cases where relativistic motion applies and the relatively small effect (one to two orders of magnitude) it has when it does, mean this does not greatly impact most timescale arguments (Krolik, 1999).

Variability can also be used indirectly to determine further properties of the system. The radius of the BLR can be determined through a method known as reverberation mapping (Peterson and Horne, 2006). Intrinsic variability in the emission of the accretion disk results in corresponding variability in the broad line emission from the BLR as its source of ionizing radiation fluctuates. The tracking and correlation of both the emission line and the continuum variability can be used to determine the time lag due to the light travel time and therefore the radius of the BLR. Then, from the velocity width of the emission lines of the BLR, it is possible to determine the mass of the central SMBH through the simple application of Kepler's laws. However, reverberation mapping is observationally expensive and difficult to perform and thus an observed scaling of the BLR radius to the AGN luminosity, originally determined via the reverberation mapping method, is generally used instead (Kaspi et al., 2005). This combination of methods has been used to determine the SMBH mass for a large number of AGN.

So far, I have focussed on the typical variability observed within AGN. However, one class of AGN, known as the blazars, show variability well above the scale of typical AGN. They exhibit the largest amplitude and smallest timescale variability of all AGN classes, with changes of a factor 100% or more in a matter of hours, even in the optical where other AGN are relatively calm (Krolik, 1999). The reasons for this vastly different behaviour are the very different emission mechanisms and regions involved. Instead of the X-ray reprocessing suspected in typical AGN, variability in blazars instead comes from bursts of injected relativistic electrons at the base of the jet, producing prompt UV and optical outbursts at early times and decaying into radio wavelengths as the fading burst propagates up the jet (Krolik, 1999). This jet interpretation is most strongly supported through the existence of large luminosity and short timescale γ -ray emission. The implied compactness of such a source would result in considerable electron-positron pair production, the annihilation of which would reprocess the γ -ray emission into lower energy bands. The only way to overcome this issue is for the timescale of variability to be shortened by relativistic beaming effects, as is the case in a jet. This argument is further strengthened through the association of γ -ray emission with radio-loud hosts (von Montigny et al., 1995) and through the presence of considerable polarisation, a hallmark of synchrotron emission (Angel and Stockman, 1980).

The brightest blazar flares can peak at extremely high luminosities across the electromagnetic spectrum, with absolute optical magnitudes as bright as $M \sim -30$ and isotropic X-ray luminosities of $\sim 10^{46} \text{ erg s}^{-1}$. They are often remarkably flat in νF_ν across infrared-X-ray emission, varying by only ~ 1 dex across the 5 decade range in frequency. However, several examples have shown the brightest flares are dominated by exceptionally luminous gamma-ray emission, peaking up to two orders of magnitude above the lower energy emission (Vercellone, 2012; Abdo et al., 2015).

It is worth noting that, while the Milky Way is far from being considered an “active” host, its supermassive black hole Sgr A*, and presumably other quiescent SMBHs like it, are nonetheless prone to some variability. Sgr A* is particularly low luminosity with steady X-ray emission at the $\sim 10^{33} \text{ erg s}^{-1}$ level (Markoff, 2005), equating to only one billionth of the SMBH’s Eddington luminosity (Melia and Falcke, 2001) and indeed cannot have approached its Eddington luminosity within the last few hundred years (Sunyaev et al., 1993). However it does still exhibit approximately daily flaring that temporarily increases its luminosity by a factor of a few (Baganoff et al., 2003). Similarly, radio, millimetre and sub-mm emission from the source has been seen to show short term variability on timescales down to ~ 1 hour (Yusef-Zadeh et al., 2011). However, in all cases these flares are many orders of magnitude fainter than the typical variability observed within active hosts and comparable flares from dormant hosts are not expected to be observable, let alone be a likely form of contaminant for extragalactic nuclear transient studies such as this.

1.3.3 Identifying Active Hosts

In order to determine the nature of nuclear transients, it is important to be able to determine the likelihood of the presence of significant AGN emission within the host. Here I briefly discuss the typical methods for identification that are used within this thesis.

The detection of strong X-ray emission is probably the most reliable method of identification with X-ray emission from AGN outshining all but the most star-forming of galaxies (Mendez et al., 2013). However, AGN can be obscured by large hydrogen column densities and the depth of X-ray observations required for identification at high redshift has been achieved across only a tiny fraction of the sky. Identification of heavily obscured AGN that are difficult to identify via optical or X-ray methods can be achieved at longer wavelengths with strong red colours indicating the presence of AGN emission (Stern et al., 2012). Surveys such as that performed by the *Wide-field Infrared Survey Explorer* (*WISE*; Wright et al., 2010) have been

particularly successful in this regard, though such surveys do suffer from issues with poor resolution due to the wide point-spread function at these wavelengths.

Detecting luminous radio emission is amongst the oldest methods for AGN identification, with most luminous extragalactic radio sources being associated with accreting supermassive black holes, as are sources with the classic double-lobed morphology that comes from jets (Mushotzky, 2004). Difficulties arise at low luminosities due to radio emission associated with star formation. However even in this regime it is sometimes possible to distinguish between accretion and star-formation emission with an accurate determination of the radio spectral index, α , defined such that the flux density F_ν at a frequency ν is $F_\nu \propto \nu^\alpha$. Radio emission from star formation has $\alpha \sim -0.9$ while AGN emission is somewhat flatter at $\alpha \sim -0.5$ (Mauch and Sadler, 2007) with a flat (i.e. $\alpha \sim 0$) spectral index being a good indicator of AGN emission (Mushotzky, 2004). However there is considerable spread in each and thus the spectral index is generally only suggestive of the emission’s origin.

Even in cases where the broad line region is masked, the spectral features produced by the narrow line region are nonetheless considerably broader than typical star-forming lines, a consequence of the larger Doppler broadening produced by the fast orbits deep in the black hole’s gravitational potential well. The identification of an AGN can also be achieved through comparison of certain emission line ratios. The power law continua produced by the AGN produces a distinct signature in the surrounding photoionized gas that can be seen by comparing the relative strengths of certain emission lines. As such, star-forming and active galaxies produce separate loci in plots such as the Baldwin, Phillips & Terlevich (BPT) diagram (Baldwin et al., 1981), a method of classification that has continued to be used to this day (e.g. Kewley et al., 2006). An example of this method is shown in Figure 1.4. However, such line emission can be masked in cases of high extinction.

1.4 Tidal Disruption Flares

With the vast gravitational potential well of a supermassive black hole, and the dense stellar environments present in the cores of galaxies, eventually and inevitably a star will stray too close. In these situations, the intense gravitational tides near the black hole can be strong enough to disrupt the star, the resulting diffuse material then accreting onto the black hole (Hills, 1975). This produces a short-term, months to years duration flare that is analogous to a short-lived AGN (Rees, 1988). As such a situation requires only the presence of a star and a SMBH, these tidal disruption flares (TDFs) can potentially occur in any galaxy; active or inactive; dwarf, grand




Figure 1.4: BPT diagram of SDSS galaxies. The differing photoionizing fluxes from star-forming and AGN processes produce considerable differences in certain observed emission line ratios, causing each group to follow different loci on the diagram. The solid red lines indicate the standard dividing line between AGN and star-forming hosts. Note the existence of an uncertain “composite” region, where neither process dominates line emission. Figure from Kewley et al. (2006)

spiral or massive elliptical. Their study and detection in large enough numbers could therefore lead to a tracing of the SMBH population that may be free from some of the biases present in the other methodologies discussed so far. At the very least, their correct identification proves the existence of a SMBH in the host, while opening avenues to the determination of the properties of the SMBH. In addition, their luminosity, comparable to core collapse supernovae, makes them visible well beyond local galaxies and thus may enable the tracking of SMBH evolution through time in similar ways to AGN observations.

A star can only be disrupted if it passes within the tidal disruption radius of the black hole. This radius corresponds to the distance from the black hole at which the gravitational potential of the black hole is equal to the gravitational potential of the star at its own half mass radius. This is therefore dependent both on the mass of the black hole and on the properties of the star, namely its mass and radius. The tidal disruption radius R_T , as determined in Rees (1988), is found to be:

$$R_T = 5 \times 10^{12} M_6^{\frac{1}{3}} \left(\frac{r_*}{r_\odot} \right) \left(\frac{m_*}{m_\odot} \right)^{-\frac{1}{3}} \text{ cm} \quad (1.7)$$

where M_6 is the mass of the black hole, M , in units of 10^6 solar masses, and r_* and m_* are the radius and mass of the infalling star respectively. Stars can fall within this limiting radius due to scattering processes, such as two-body relaxation in the

dense core that sends the star inwards (Frank and Rees, 1976). Alternatively, as R_T is dependent on a given star’s density, a star on an otherwise stable orbit around the black hole can become disruptable as its internal structure evolves (MacLeod et al., 2012).

Stars falling within this radius will be completely disrupted, with approximately half of the resulting material being flung out into deep space at a high velocity. The rest will be left in a set of bound orbits, falling back onto the black hole and accreting over the course of a few months to years (Figure 1.5a). Approaches that fall somewhat outside of this radius can still result in a stripping of some portion of the outer layers of the star. This results in a smaller, partial disruption that may repeat on the period of the stellar orbit (Figure 1.5b).

However, the above shows an important caveat for the ubiquity of TDFs. For a given star, R_T scales as $M^{\frac{1}{3}}$. However, as was seen in equation 1.1, the Schwarzschild radius (R_S) scales proportionally to M . As such there will be a black hole mass at which $R_T < R_S$. In such a case, the infalling star would not be disrupted until after it has crossed the event horizon and would therefore produce no observable flare (Figure 1.5c). Setting equation 1.1 and 1.7 equal to one another, the critical mass for the disruption of a star:

$$M_{\text{crit}} = \left(\frac{r_*^3 c^6}{8G^3 m_*} \right)^{\frac{1}{2}} \sim 1 \times 10^8 M_{\odot} \quad (1.8)$$

where the latter result is calculated under the assumption of the disruption of a solar mass star with a solar radius. As such, for a Sun-like star, supermassive black holes above this limit will swallow stars whole, though less dense giant stars would be disruptable up to much higher masses. This limit becomes even more important for more compact objects, with white dwarfs being disruptable only by black holes of $10^5 M_{\odot}$ or less, resigning them purely to the intermediate mass black hole regime.

While theoretical modelling of these flares is extremely difficult, in part due to the massive range of length scales important in these models (from the Schwarzschild radius out to the distance of the expelled material and associated winds), the general theoretical model is well established. Typically stars are pushed into the “loss cone”, the region of parameter space where each member (star) comes within its tidal disruption radius of the black hole through gravitational interactions with neighbours. A combination of the depletion rate and filling rate of this loss cone provides the expected tidal disruption flare rate. This is a complex function of the stellar dynamics of the central regions of the galaxy, but a number of studies suggest rates on the order of 10^{-5} per year per galaxy (e.g. Rees, 1988; Magorrian

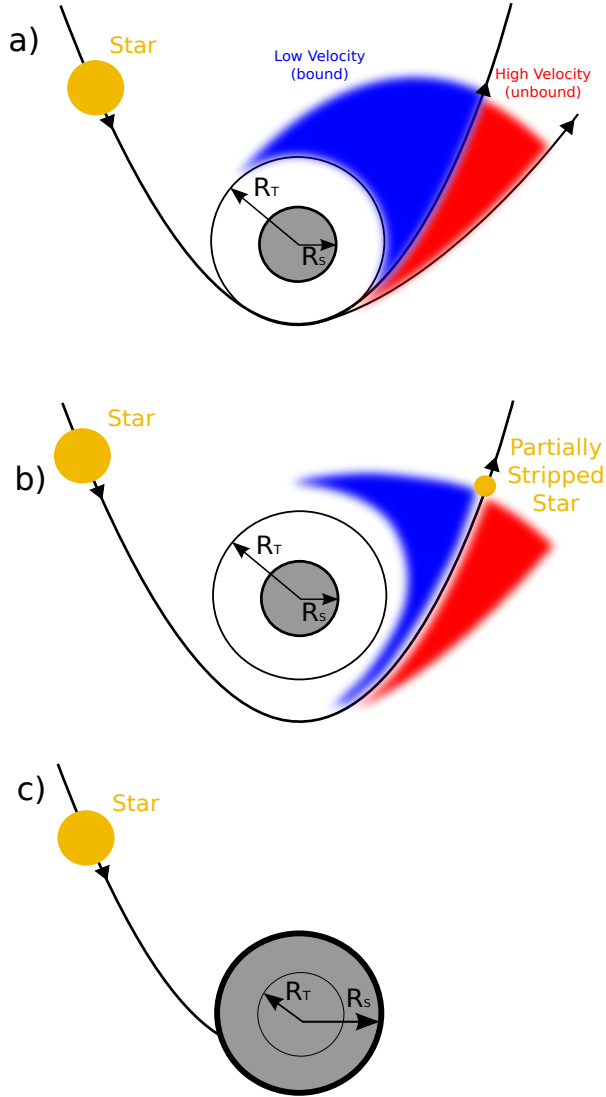


Figure 1.5: Three possible configurations of the infall of a star towards a super-massive black hole. a) The simple case of a star on a parabolic orbit that has its closest approach at the tidal disruption radius R_T . In this case the tidal forces of the SMBH causes a complete disruption of the star with about half of the remaining diffuse material falling back onto the SMBH producing a tidal disruption flare. b) When the closest approach is just outside of the tidal disruption radius a partial disruption is possible, where only the outer layers are stripped but the core remains intact. If the star is on a bound orbit, this configuration may produce repeated flares on the period of the orbit and could eventually lead to a full disruption. c) If the star's self-gravity is too strong (or equally the SMBH is too massive), the star will cross the event horizon before it can be disrupted. In this plunging star configuration, there is no observable flare. Based on a figure in Rees (1988)

and Tremaine, 1999) in keeping with current observed rates (Donley et al., 2002; van Velzen and Farrar, 2014). The rates could be enhanced in galaxies with hardening binary SMBHs, those that are slowly closing the gap between each other through a combination of dynamical friction (interactions with stars) and, at late times, gravitational wave emission (Ivanov et al., 2005; Chen et al., 2011). TDF rates can also be increased in cases where the number of potentially disruptable stars is high as might be the case in regions of enhanced number density of stars produced in recent starburst activity (Stone and van Velzen, 2016).

So far this discussion has not included the clearly important relativistic effects that occur in the extreme environments in the vicinity of a supermassive black hole. For example, high spin black holes are capable of disrupting more compact stars than would be possible if not spinning. This effect means that main sequence stars can potentially be disrupted by black holes up to $\sim 10^9 M_\odot$ as opposed to the normal $\sim 10^8 M_\odot$ limit (Kesden, 2012). As such, high spin black holes will have their TDF rates increased as the population of stars capable of being disrupted by them increases. However, relativistic effects also cause a fraction of the incoming stars to be directly captured by the event horizon before it can be disrupted, thus producing no visible flare and reducing the overall TDF rate (Kesden, 2012). These conflicting effects are highly dependant on the black hole’s mass and spin and thus affect TDF rates to varying extents around different black holes.

In cases where stars are disrupted, the diffuse material circularises forming an accretion disk, similar to that observed in AGN. These disks are expected to have high temperature ($\sim 10^5$ K) thermal spectral energy distributions and thus radiate much of their emission at ultraviolet and soft X-ray wavelengths (Rees, 1988; Strubbe and Quataert, 2009). Also, unlike the expanding ejecta of core collapse supernovae, the compact accretion disk is expected to retain its high temperature throughout the flare. As a result, many of the examples detected to date have been found at X-ray (e.g. Grupe et al., 1995; Komossa and Greiner, 1999; Greiner et al., 2000; Esquej et al., 2007; Cappelluti et al., 2009; Saxton et al., 2012b) and ultraviolet (e.g. Gezari et al., 2006, 2008) wavelengths, with a number of others detected as part of supernova surveys at optical wavelengths (e.g. van Velzen et al., 2011; Gezari et al., 2012; Chornock et al., 2014; Holoien et al., 2014; Holoien and Stanek, 2016). In total, a few dozen candidates of varying strength have been detected to date⁹. However finding and identifying examples is difficult, not least because the early blue colours of these sources are very similar to Type II supernovae (see Section 1.6).

⁹see <http://tde.space>

Observations following the long-term evolution of a TDF is an important route to its identification. With the associated emission typically assumed to be proportional to the accretion rate, we require a determination of the fallback rate onto the supermassive black hole. In the simple model put forth by Rees (1988), and corrected by Phinney (1989), for a star in hydrostatic equilibrium, the only unbalanced force is that of the gravity of the supermassive black hole, meaning that in effect each part of the star is on its own Keplerian orbit. While initially the orbits are tightly grouped around the centre of mass, the elliptical orbits bunch up near periapsis redistributing energy inside the star through internal pressure. With the wider energy distribution and under the assumption of a star on a parabolic orbit with closest approach at less than r_T , comes the result that the star is now split between material bound to the SMBH and that which is unbound. Following the paths of their newly attained orbits, the unbound material is flung off into space at high velocity. The bound material returns on a period T , based on its binding energy, E , given by:

$$E = -\frac{1}{2} \left(\frac{2\pi GM}{T} \right)^{-\frac{2}{3}} \quad (1.9)$$

Under the further assumption that the bound material loses its energy quickly upon reapproaching periapsis, the mass accretion rate is then the mass distribution of return times, given by:

$$\frac{dM}{dT} = \frac{dM}{dE} \frac{dE}{dT} = \frac{(2\pi GM)^{\frac{2}{3}}}{3} \frac{dM}{dE} T^{-\frac{5}{3}} \quad (1.10)$$

Finally, under the assumptions that the energy distribution $\frac{dE}{dT}$ is uniform, a finding supported by the numerical simulations of (Evans and Kochanek, 1989), and that the luminosity of the event is directly proportional to the accretion rate, the lightcurve would also be expected to decline as $t^{-\frac{5}{3}}$ (Lodato et al., 2009). The onset of the flare is expected to occur upon the return of the most bound material typically a few days to weeks (Lodato et al., 2015). However, more recent modelling shows that the energy distribution depends on the internal structure of the star, meaning this $t^{-\frac{5}{3}}$ decay rate is only applicable at late times (Lodato et al., 2009). Further modelling reveals that late time fallback of partial disruptions may become even steeper than the canonical value at late times (Guillochon and Ramirez-Ruiz, 2013). Thus the early description and often searched for decay rate may not accurately describe TDF evolution. It is also important to note that while the mass accretion rate (and therefore, under the assumption of a direct proportionality between accretion and




Figure 1.6: A selection of *ROSAT* discovered TDEs. Each is approximately consistent with the canonical $t^{-5/3}$ decay rate (red dashed line) with the possible exception of RXJ1624+75 which appears to decay much faster. Not plotted here is the very late time decay of RXJ1242-11 which decays faster than the canonical decay at late times. Figure from Komossa (2015)

luminosity, the bolometric lightcurve) is expected to follow this decay, the lightcurve in any given band may not follow this overall trend (Strubbe and Quataert, 2009; Lodato and Rossi, 2011). However, a number of events detected to date that have been attributed to tidal disruption events do indeed have the canonical decay rates over multiple years, particularly those discovered at X-ray wavelengths (see Figure 1.6; Komossa, 2015).

Beyond the high temperature thermal flares and general lightcurve properties, the expected properties of tidal disruption flares are somewhat difficult to define, for the most part due to the wide range of variables that can affect the resulting outburst. For example, the total energy released in the flare is in part dependent on the total mass accreted, which can be considerably less than the mass of the whole star, not least due to the portion of the star that is left unbound in the simple case considered above. A handful of events detected so far have had luminosities

considerably below the rest and with correspondingly low accreted mass estimates. For example, PS1-11af (Chornock et al., 2014) had an estimated minimum required mass for accretion of $\sim 0.002M_{\odot}$. As a result it was suggested to be the observation of a partial tidal disruption, where the star passes too far away from the black hole to be entirely disrupted but nonetheless has a fraction of its outer layers stripped off. This process may repeat on the period of the star, particularly if the star is an evolving giant star that expands after each partial disruption, resulting in the slow repeated feeding of the SMBH (MacLeod et al., 2013). One such case may have been observed in IC3599 (Grupe et al., 1995) which produced strong X-ray flares that appeared to repeat on a 10 year period. This may indeed be due to multiple partial disruptions (Campana et al., 2015), however the existence of optical emission preceding the X-ray flare is difficult to explain in the tidal disruption case and may be most readily explained as a flaring AGN (Grupe et al., 2015).

While the typical emission from an accretion disk is expected to be an effectively featureless, thermal continuum¹⁰, a number of TDF candidates have shown evidence for hydrogen and helium emission lines in their spectra. Perhaps the most striking was the detection of PS1-10jh, the spectrum of which exhibited strong HeII emission lines but placed limits on Balmer emission that meant the implied accreted hydrogen fraction was < 0.2 (Gezari et al., 2012). This has led to the suggestion that the star that had been disrupted had lost its hydrogen envelope, perhaps stripped off during an earlier encounter with the black hole, and thus the object had been a He-rich stellar core. Further candidates have since been identified with a range of hydrogen and helium line strengths, indicating a range of stellar progenitors have been observed (see Figure 1.7; Arcavi et al., 2014).

In a similar fashion to the effect exploited for reverberation mapping of AGN, a short, high-energy ultraviolet and X-ray flare is expected to produce considerable photoionization of nearby circum-nuclear material. This emission is then reprocessed into line emission. In addition to a number of sources that have been identified as continuum flares that have subsequently been shown to have variable emission lines (e.g. van Velzen et al., 2011; Gezari et al., 2012), a number of examples of otherwise quiescent galaxies with variable line emission have been detected (e.g. Komossa and Merritt, 2008; Komossa et al., 2009; Wang et al., 2011b; Wang and Cheng, 2012). These appear to have very different properties to emission line variability seen as a result of supernovae (Komossa et al., 2009), with some examples exhibiting extremely strong coronal line emission. With the potential for much longer timescales

¹⁰Indeed the early exclusion of the possibility of CSS100217 (see Chapter 5) as a TDF, was in part due to the detection of evolving Balmer emission (Drake et al., 2011)

of variability than the original flare, this represents an additional route to the identification of new candidate TDFs.

While TDFs are expected to be a near ubiquitous feature of galaxies hosting small enough black holes, a number of features are generally helpful for their conclusive identification. In particular hosts with minimal star formation helps eliminate core collapse supernova origins while the lack of AGN activity makes AGN variability a less likely explanation for a candidate's origin. An interesting recent trend in the hosts of TDF candidates has included an apparent preference for an extremely rare class of galaxy: the E+A galaxies. These galaxies, making up only between 0.01-1% of the galaxies in the local Universe (though the fraction is dependent on the exact definition, Quintero et al., 2004), have nonetheless played host to several of the tidal disruption flare candidates detected to date (Arcavi et al., 2014; French et al., 2016, 2017). These galaxies are characterised by unusually large numbers of A-type stars despite low current star formation rates, indicating a large change in star formation within the last Gyr (Quintero et al., 2004). One possible explanation for this behaviour resides in a galaxy recovering from a relatively recent merger. In such a situation, the combination of chaotic orbits and high central stellar number densities might produce an enhanced tidal disruption rate. Indeed, a recent analysis of the central region of NGC 3156, a nearby E+A type galaxy, has suggested the dominant cause might be stellar overdensities as the loss cone is rapidly refilled by two-body relaxation, producing a potential factor ~ 100 increase in TDF rates (Stone and van Velzen, 2016).

It has also been suggested that tidal disruption flare rates may be enhanced in the presence of a supermassive black hole binary (SMBHB, Ivanov et al., 2005; Chen et al., 2009), though the highest rates are expected to be achieved only during the short period preceding and following coalescence and thus likely does not contribute majorly to the overall TDF rate (Wegg and Nate Bode, 2011). In extremely tight binaries, signals may be evident in the tidal disruption flare lightcurve. In one case, the canonical power law lightcurve decay appears to be punctuated by substantial gaps (Saxton et al., 2012b) that may be indicative of a second SMBH periodically disrupting the stream of material falling back on to the first (Liu et al., 2009, 2014). While no other supermassive black hole binary is apparent in any of the other TDF candidate lightcurves to date, with the much larger expected rates of TDFs in the future, this may become an efficient method for the identification of further examples of SMBHBs that have overcome the final parsec problem of binary evolution. This has wide-ranging implications, not least for the detection of the powerful gravitational waves expected from the coalescence of such a binary in the current age of




Figure 1.7: Continuum subtracted spectra from a number of TDF candidates. The events show a range of observed line strengths, from the helium-dominated PS1-10jh (Gezari et al., 2012) to the hydrogen-dominated TDE2 (van Velzen et al., 2011). Figure from Arcavi et al. (2014)

multi-messenger astronomy.

1.5 Relativistic Tidal Disruption Flares

While all of the tidal disruption flare candidates considered so far have been detected at soft X-ray, UV and optical wavelengths, in part due to their high temperature thermal flares, a small subset of events have been detected by the *Swift* Burst Alert Telescope (BAT, Barthelmy et al., 2005) in hard X-rays. The first such event, Swift J164449.3+573451 (Swift J1644+57) was detected in March 2011 (Cummings et al., 2011; Levan et al., 2011; Burrows et al., 2011), while the second, Swift J2058.4+0516 (Swift J2058+05) was found only 2 months later in May 2011 (Cenko et al., 2012).

Each event had peculiar properties that set them apart from any γ -ray transients that had been detected at the time. In both cases, each was detected by *Swift*-BAT in unusually long triggers, with Swift J1644+57 requiring a long image trigger (~ 1000 s, Cummings et al., 2011) and Swift J2058+05 found in a 4-day average (Krimm et al., 2011a). Swift J1644+57 also produced repeat triggers of the satellite, with 4 triggers over the course of 48 hours (Barthelmy et al., 2011), indicating it was clearly not a standard gamma-ray burst (GRB), either short or long. While T_{90} ¹¹ is difficult to determine accurately in these cases, their visibility in γ -rays days after the original detection implied durations far in excess of normal GRBs (though see later for a discussion of possible exceptions to this).

X-ray observations of each flare revealed bright point source emission that remained visible for months after the initial flare. In the case of Swift J1644+57, a wild flaring phase was followed by a plateau of several days before dropping into a power law decline remarkably similar to the $t^{-5/3}$ evolution expected for tidal disruption flares (Levan et al., 2011; Burrows et al., 2011; Saxton et al., 2012a; Mangano et al., 2016). Swift J2058+05 followed a similar evolution with a long plateau and late power law decline, though with a steeper index than for Swift J1644+57 (-2.2 Cenko et al., 2012). Throughout, both events also showed considerable short term variability superimposed on their broad evolution, with Swift J1644+57 exhibiting order of magnitude changes on timescales of 100 s (Levan et al., 2011) and Swift J2058+05 showing factor of 1.5 variability on timescales of 10^4 s (Cenko et al., 2012).

Optical observations showed that the events came from the nuclei of dwarf galaxies with redshifts of $z = 0.353$ (Levan et al., 2011) and $z = 1.1853$ (Cenko

¹¹A typical measure for the duration of GRBs is the time within which 90% of the total fluence of the burst is received, known as T_{90} . It is this value that has (along with their spectral hardness) led to the distinguishing of the two main populations of GRBs, the short and long variants (Kouveliotou et al., 1993)

et al., 2012) for Swift J1644+57 and Swift J2058+05 respectively. Optical flares were visible in both cases, though for Swift J1644+57, the red colour of the flare and high inferred hydrogen column density implied the presence of moderate extinction, resulting in a flare that was most visible in the NIR (Levan et al., 2011; Bloom et al., 2011). Astrometric constraints placed the transients at < 150 pc (Levan et al., 2011) and < 400 pc (Pasham et al., 2015) from the centres of their hosts respectively, making an association with the central supermassive black hole plausible in each case.

Radio observations of Swift J1644+57 showed the presence of a rising unresolved source with an equipartition radius that implied moderately relativistic expansion with a Lorentz factor of ~ 2 and a formation epoch that coincided with the original detection (Zauderer et al., 2011). This radio flare has continued to be monitored to the present day with the observations showing it remains visible several years after its onset (Zauderer et al., 2011; Berger et al., 2012; Zauderer et al., 2013, and subsequent VLA observations PI:Zauderer). Radio monitoring of Swift J2058+05 has been far less diligent, though observations did show radio emission that appeared to be declining by 40 days post trigger (Cenko et al., 2012; Pasham et al., 2015).

The redshifts of the events implied that each flare had a peak X-ray luminosity in excess of $10^{47} - 10^{48} \text{ erg s}^{-1}$ while their optical emission was far more modest at $10^{42} - 10^{43} \text{ erg s}^{-1}$. This placed them both in a region of X-ray Luminosity - Optical Absolute Magnitude phase space that was devoid of other sources (see Figure 1.8, Levan et al., 2011; Cenko et al., 2012). Blazars with similar X-ray emission typically have optical luminosities several orders of magnitude more luminous than that observed in these flares, while GRBs, though sharing similar optical luminosities, have faded well below the level of these events in X-rays at late times. This indicated these flares belonged to a new class of event that had not been observed before.

Modelling of the emission of Swift J1644+57 across the electromagnetic spectrum indicated that it likely arose from a combination of the interaction of relativistic jets with the surrounding interstellar medium (IR-radio emission) and inverse compton scattering within a jet (X-ray and γ -ray emission, Bloom et al., 2011). This finding was supported by the relativistic expansion inferred from the radio emission and by the short-term variability of the high-energy emission. While their nuclear positions indicated a possible link with the central supermassive black hole, the star-forming nature of their host galaxies made an AGN origin for the flares unlikely. As a result, the current popular explanation for these events is in the tidal




Figure 1.8: The X-ray Luminosity - Optical Absolute Magnitude plot for a selection of sources. GRBs typically decay far faster than the rTDF candidates, and thus have far lower X-ray luminosities at late times. Blazar flares on the other hand are far more optically luminous. This places the rTDFs in a region otherwise devoid of sources, marking them as a separate form of transient. Figure from Cenko et al. (2012).

disruption of a Sun-like star that also launched a moderately relativistic jet. These events represent a chance alignment with the jet, enhancing the luminosity through a combination of collimation and relativistic beaming, naturally explaining how an event with a luminosity equivalent to the Eddington luminosity of a $10^{10} M_{\odot}$ black hole could originate from a host that likely houses a black hole of little more than a $10^7 M_{\odot}$ (Levan et al., 2011; Miller and Gültekin, 2011; Cenko et al., 2012).

These events represent an important new avenue in the studies of accretion onto massive compact objects, providing effectively miniature short-lived blazars that can be observed from onset of accretion to termination on human timescales. The increased apparent luminosity of these events also extends the range that tidal disruption flares could be observable to back to the early Universe.

However, their interpretation is not uncontested, with alternative explanations in the tidal shredding of a white dwarf (Krolik and Piran, 2011) or massive star core collapse (Quataert and Kasen, 2012b; Woosley and Heger, 2012). Recent observations of both flares have provided potential new evidence for the latter, as each underwent sizeable sharp drops in luminosity after a few hundred days (Sbarufatti et al., 2012; Levan and Tanvir, 2012; Zauderer et al., 2013; Levan et al., 2016; Pasham et al., 2015). Such a situation was explicitly expected in core collapse models of the flare (Quataert and Kasen, 2012a) as this would represent the point at which the entire star’s mass has been accreted onto the newly formed black hole. However, processes involving magnetically arrested accretion could also explain the drop in emission in the TDF case as the jet effectively shuts off once the accretion rate drops below a critical threshold (Tchekhovskoy et al., 2014). Further, in the core collapse case it is difficult to explain the near constant level of X-ray emission present for at least ~ 1000 days after the decline of Swift J1644+57 at a luminosity of $\sim 5 \times 10^{42} \text{ erg s}^{-1}$. This would require both high accretion rates and considerable beaming to be achievable for a stellar mass black hole, despite the original star’s matter having apparently been used up. Perhaps the most likely explanation for this late time emission is that of a low-level AGN separate to the transient emission, though the X-ray-optical emission does not fall within the normal locus of AGN (Levan et al., 2016).

A further peculiarity observed in the lightcurve of Swift J1644+57 is the presence of an optical/NIR rebrightening that peaked approximately 30 days post trigger (Levan et al., 2016). While explanations for this feature are numerous, including the detection of the thermal component of the TDF (Levan et al., 2016), the detection of the second synchrotron component (Berger et al., 2012) and reverberation and reprocessing of the X-ray emission (Levan et al., 2016), one possibility

is the detection of supernova emission given that the timescale and luminosities involved are approximately consistent with superluminous supernovae (SLSNe, see subsection 1.6.1; Gal-Yam, 2012). Luminous X-ray emission from superluminous supernovae would not be unprecedented, with at least one detection in the case of SCP06F6 (Levan et al., 2013), though the origin of this emission is not clear, and while the luminosity of the emission would make it amongst the most luminous SLSN observed to date, the two most luminous known examples, ASASSN15lh (Dong et al., 2016) and CSS100217 (Drake et al., 2011), are also located within the nuclei of their hosts (see Chapter 5. Perhaps more important is the possible connection to a further class of γ -ray transient, the ultra-long GRBs (ULGRBs).

1.5.1 A connection between long-lived gamma-ray events?

Whereas most long GRBs have T_{90} durations of up to a few hundred seconds, these ULGRBs can last thousands of seconds or more (see Figure 1.9, Levan et al., 2014). What is more, these flares tend to be dominated by longer-lived, lower luminosity tails as opposed to the short, high luminosity bursts that feature most strongly in classical GRBs. Indeed often these early peaks are not present in ULGRBs. As a result of this behaviour, with lower peak count rates than many GRBs, a number of these bursts have not been found using the usual rate triggers on *Swift*, instead relying on time-averaged methods such as image triggers. However these methods necessarily reduce the temporal resolution, making the duration of the burst difficult to determine precisely. In the case of ULGRBs, the short orbital period of *Swift* (~ 90 minutes) also causes issues with duration determinations, as bursts longer than the maximum uninterrupted pointing time will have their observation occulted by the Earth until the object becomes visible again. The high variability of GRBs makes a determination of the intervening emission difficult without simultaneous observations from other platforms. Also, in at least one case, GRB101225A, the flare was already in progress when *Swift* slewed in its direction (Palmer et al., 2010), making any estimate of T_{90} a lower limit.

At the forefront of these bursts are GRB101225A ($T_{90} > 7000$ s; Thöne et al., 2011; Levan et al., 2014), GRB110229A ($T_{90} \sim 10000$ s; Gendre et al., 2013; Stratta et al., 2013; Levan et al., 2014), GRB121027A ($T_{90} \sim 6000$ s; Levan et al., 2014) and GRB130925A ($T_{90} \sim 10000$ s; Evans et al., 2014). Of these, two (GRB101225A and GRB110229A) have been found to be consistent with the centres of their compact host galaxies, making an association with the central supermassive black hole plausible. However, the position of GRB130925A is offset significantly from the nucleus of its host, though it was noted that the distorted nature of the host may indicate




Figure 1.9: The T_{90} and hardness ratios (ratio of flux in 50-100keV band to flux in 20-50keV band) for *Compton Gamma Ray Observatory* (CGRO) and *BATSE* detected GRBs (red) and for *Swift* detected GRBs (blue). A further population of possible ULGRB candidates (yellow) have less certain T_{90} durations, determined instead through their X-ray, as opposed to γ -ray, lightcurves. The highlighted events (GRB101225A, GRB111209A and GRB121027A) are clear outliers to the population. All three are *Swift* detected, though GRB111209A was also detected by *Konus-WIND*. Note that the rTDF candidates, with T_{90} times of several days, are outliers even to these unusually long events. Also note the difference in observed GRB populations from the different instruments, with *Swift* preferentially detecting longer, softer bursts, a consequence of its lower energy wide-area burst detection instrument. Figure from Levan et al. (2014).

a recent merger and thus the existence of a second black hole (Tanvir et al., 2013). Indeed, while GRB101225A and GRB110229A were both consistent with the centres of their hosts, the compact nature of their hosts meant that the astrometric error bar encompassed the majority of the stellar light.

The true origin of these events may have now been determined given the detection of supernova emission in the lightcurve of GRB111209A, ruling out a tidal disruption origin for the flare (Greiner et al., 2015). Further, the low implied Nickel mass, insufficient to power the luminosity of the event, and the low likelihood of an interaction powered event suggests a central engine model involving the spin-down of a magnetar. While this detection has only been made for one event, the similarity between the ULGRBs suggest similar origins for the entire class.

The question then is: are the rTDF candidates an extension of this class or do they have a separate origin of their own? While the γ -ray emission of both classes are undeniably far longer than the bulk of GRBs, the rTDF candidates are still outliers to most of the ULGRBs, the only possible exceptions being those with lower limits on their true durations, such as GRB101225A (Palmer et al., 2010). Further, the late time X-ray emission of ULGRBs falls well short of the luminous emission of the rTDF candidates, with ULGRB afterglows being comparable, or perhaps somewhat fainter, than typical GRBs (Levan et al., 2014). Coupled with the strong astrometric ties to their host’s nuclei, it seems likely that the rTDF candidates are a separate class with a different origin to those of ULGRBs. Nonetheless, the properties of these flares, and the nuclear SLSNe, must continue to be compared and contrasted with the hope that further examples could answer the question definitively in the future.

1.6 Supernovae

While a flare’s position in the nuclear region of its host is generally seen as highly suggestive of an association with the central supermassive black hole, the high rates of supernovae inevitably mean a significant fraction of centrally located transients will be associated with them. This is particularly true in the absence of high resolution imaging of both the transient and the host. Even at the relatively nearby distance of $\sim 300\text{Mpc}$ ($z \sim 0.07$), the rate of supernovae expected to occur within $1''$ of the centre of a typical galaxy is of order 10^{-4}yr^{-1} , perhaps an order of magnitude greater than the rate of TDFs (Strubbe and Quataert, 2009). Thus it is important to consider the observable properties that might distinguish supernovae from other forms of nuclear event.

Supernovae are broadly split into two categories: those with hydrogen features in their spectra, the Type II's, and those without, the Type I's (Minkowski, 1941). Within these categories are multiple sub-categories each focusing on an observed feature of the transient, from the width of emission lines to the rate of its decline in luminosity. This multitude of supernova categories can be extremely complicated to navigate. An inexhaustive diagram of the major categories of supernovae under discussion is shown in Figure 1.10

The only major supernova class not thought to be due to massive star core collapse are the Type Ia's. This class of event, with characteristic lightcurves (Hamuy et al., 1996) and spectra with strong silicon absorption and no notable hydrogen features (Filippenko, 1997; Hillebrandt and Niemeyer, 2000), have long been known to be associated with the thermonuclear explosions of white dwarfs in binary systems. As mentioned in subsection 1.2, white dwarfs are limited to masses below $\sim 1.4M_{\odot}$, the Chandrasekhar limit. Above this value, the star is expected to collapse in upon itself, forming a neutron star or black hole. However, if approached slowly through long-term accretion from a binary companion, the result can instead be quite different. The rising temperature within the white dwarf eventually reaches the critical temperature for the fusion of carbon, and the star quickly goes through a massive chain reaction. The resultant outburst of energy is large enough to unbind the star and results in a supernova (Wheeler and Harkness, 1990; Hillebrandt and Niemeyer, 2000). The homogeneity of this process makes them useful for cosmologists as standard candles with a strong correlation between their decay timescale and peak absolute magnitude (see Branch and Tammann, 1992; Hamuy et al., 1996).

The core collapse supernovae, on the other hand, are extremely inhomogeneous. They vary greatly across a large range of properties (Smartt, 2009). However in each case, the typical model is the same. A massive star with a minimum mass of $7 - 10M_{\odot}$ begins to attempt the energetically undesirable fusing of iron. This endothermic process causes a sharp drop in the temperature of the core, greatly reducing the outwards thermal pressure which was holding up the star against its own immense gravitational pull. The resulting collapse produces a shock upon impacting the exceptionally dense proto-neutron star core and the star explodes in a bright supernova. The resulting supernova properties are greatly dependent on the progenitor and surrounding environment. Type II supernovae, with their strong hydrogen lines (at least at early times, with Type IIb's losing this feature later on) represent stars that had retained their hydrogen envelopes up to the point of the supernova occurring.




Figure 1.10: A brief summary of the various classes and sub-classes of supernovae and the typical observational features separating them. The diagram is not exhaustive, with a number of minor classes intentionally left out, instead covering the major categories of concern within this work. Figure adapted from <http://astro.matf.bg.ac.rs> (accessed 29/06/16)

Stars that have been stripped of their outer layers typically won't display hydrogen in their maximum light spectra, placing them in the Type I class of event. The presence or absence of helium features then further separates them into the Type Ib and Type Ic sub-classes respectively. Due to the difficulty in determining the difference between the two, members of both sub-classes can be known collectively as Type Ibc's. The differences between the two types has been interpreted as varying degrees of stripping, with Ib's retaining much of their helium-rich envelope (Filippenko, 1997). The progenitors of Type Ibc events are in part likely to be stars that are stripped of their outer envelopes through mass transfer to a binary companion (Podsiadlowski et al., 1992; NOMOTO et al., 1995). This is supported by the lack of direct progenitor detections in pre-explosion imaging (Van Dyk et al., 2003; Gal-Yam et al., 2005; Maund and Smartt, 2005; Maund et al., 2005; Crockett et al., 2007; Smartt, 2009) and the sometimes low ejecta masses when compared with more massive single star models (Mazzali et al., 2006; Valenti et al., 2008).

However, the coincidence of some Type Ic's with young star-forming regions and the more energetic events having high ejecta masses suggests a high mass progenitor (Smartt, 2009). In such a case, the progenitor is likely a Wolf-Rayet star, a massive star that loses much of its outer envelope due to strong stellar winds (Gaskell et al., 1986). It is worth noting that the only form of supernova so far found to be associated with long gamma-ray bursts are a variant of the Type Ic's with broad absorption lines in their spectra, the Type Ic-BL's.

In addition to the identification of supernovae via spectra taken throughout their evolution, clues as to the identity of a flare can come from their lightcurves. Each class of supernova has a typical evolution that can be identified through long-term monitoring of the transient (see Figure 1.11). For example, Type II-P supernovae are characterised by the presence of a long plateau phase of ~ 100 days starting about 20 days after peak before dropping sharply at late times (Sanders et al., 2014), the plateau being caused by the presence of an extended shock-ionised hydrogen envelope (Smartt, 2009; Kasen and Woosley, 2009). On the other hand Type II-L's have a long linear decay with time, likely due to the progenitor having a smaller hydrogen envelope.

One feature that is common amongst all supernovae that generally sets them apart from tidal disruption flares is their temperature evolution. The expanding and cooling outer layers of the supernova ejecta result in the transient's temperature cooling rapidly within days of peak light (e.g. Miller et al., 2009; Botticella et al., 2010; Inserra et al., 2013), an effect that can be seen both in the spectra and broadband photometric colours. Thus by comparing potential nuclear transients to




Figure 1.11: Typical lightcurves of the various classes of supernovae that can be used to identify new transients. For example, the Type II-L and Type II-P classes are characterised by the presence of a linear decay and a plateau phase respectively. Figure from <http://wiki.ivoa.net/twiki/bin/view/IVOA/CSPTTimeSeries> (accessed 27/04/17)

the spectral and lightcurve templates and models, and looking for evidence of temperature evolution, it is possible to determine the likelihood of a supernova origin for the flare.

1.6.1 Superluminous Supernovae

One class of supernova is of particular interest within this work. These rare and extremely luminous events, with an absolute magnitude of $M_V < -21$ (Quimby et al., 2011; Gal-Yam, 2012), approximately 100 times more luminous than typical core collapse supernovae and 10 times the peak of Type Ia supernovae, are known as superluminous supernovae (SLSNe). While thought to be core collapse supernovae of extremely massive stars, their longevity and luminosity pushes their energetics to the limits of theoretical models. Three progenitor paths have been suggested for the class: the pair instability model, where temperatures reach the point of spontaneous runaway production and annihilation of electron-positron pairs, unbinding the star and leaving no remnant (Rakavy and Shaviv, 1967; Heger et al., 2003); the

interaction model, where the supernova shock interacts with a cloud of circumstellar material, likely material blown off of the star during an earlier mass loss event (Chevalier and Irwin, 2011); and the central engine model, where the supernova ejecta is powered through the injection of energy from a central engine, possibly an accreting black hole or, more likely, the spin-down of a magnetar (Kasen and Bildsten, 2010; Dexter and Kasen, 2013). These transients have so far tended to be found in compact, star-forming hosts (Lunnan et al., 2014; Angus et al., 2016) with moderate to low metallicities (Leloudas et al., 2015) which have been suggested to be similar to those of long GRBs (Lunnan et al., 2014).

The general interest in this class of transient comes from a number of directions. If standardisable, the luminosity of these events would enable them to be used as cosmological probes over much greater distances than current standard candles. Also, particularly in the case of the pair instability mechanism, it has also been suggested that these could represent the deaths of stars made from pristine gas, the Population III stars (Gal-Yam et al., 2009).

SLSNe also have a number of links with the events studied within this work. Both CSS100217 and ASASSN15lh were originally suggested to be superluminous supernovae, a claim I analyse the likelihood of within Chapter 5. Further, given the currently favoured magnetar model of SLSNe (e.g. Nicholl and Smartt, 2016), a model also favoured for the powering of ultra-long GRBs (Levan et al., 2014), which themselves have a link with relativistic tidal disruption flare candidates, there is potential for a shared origin of the three classes of events. Indeed the prototype rTDF, Swift J1644+57 had an apparent re-brightening that could have been consistent with a superluminous supernova at late times (Levan et al., 2016). Also, at least one SLSN, SCP06F6 (Barbary et al., 2009), has had a strong detection of X-ray emission at late times that is approximately consistent with rTDF emission, though with a puzzling lack of emission at early times (Levan et al., 2013). As such the properties of these events are compared against when considering the identities of a number of the flares within this work.

1.7 Summary of Transient Properties

Here I briefly summarise the expected properties of the classes of transients that will be compared against within this thesis:

- *Tidal Disruption Flares:*

- Expected to exhibit a bolometric luminosity that falls as $t^{-5/3}$ at late times
 - High temperature thermal SED that remains approximately constant throughout its evolution
 - Typical limit on central supermassive black hole mass of $10^8 M_\odot$
 - Lack of AGN activity or intense star formation in the host can lead to the exclusion of alternative explanations
- *Relativistic Tidal Disruption Flares:*
 - Long-lived (days) γ -ray emission
 - Extremely luminous and long-lived X-ray emission with more moderate optical emission
 - Luminous radio emission with inferred moderately relativistic expansion
- *AGN flares:*
 - Typically stochastic variability across the electromagnetic spectrum
 - Presence of pre-flare X-ray and/or radio emission, or broad spectral lines in host spectrum
 - Optical emission typically correlated with X-ray, though possibly with time lag
- *Nuclear Supernovae:*
 - High temperature thermal spectrum that cools quickly after peak
 - Characteristic line emission and lightcurves across the various classes

Within this work, I shall compare the properties of a selection of nuclear flares against the properties of the classes of event considered here in an attempt to determine their natures.

In Chapter 3, I consider a new candidate relativistic tidal disruption flare, Swift J1112.2-8238, which shares many of the same properties of the previous rTDF candidates. This event, detected only 3 months after the prototype of the class, appeared to have similar duration and evolution at high energies as Swift J1644+57 and Swift J2058+05. However, the source had been mostly overlooked, leading to minimal follow-up observations and the lack even of a confirmed extragalactic origin. Within this work I analyse the high-energy properties along with optical imaging

and spectroscopy to determine whether or not this event is truly a member of the same class.

Building on this work, and requiring tighter constraints on the transient’s position within its host, in Chapter 4 I analyse high resolution *Hubble Space Telescope* (*HST*) imaging of the host of Swift J1112-8238. I combine this with medium resolution optical spectroscopy with X-Shooter in order to fully characterise the host galaxy, enabling a more complete comparison with those of the previous rTDF candidates. Importantly, I also present the first radio observations of the source, such observations being a key diagnostic of the origin of the other two flares.

Finally, in Chapter 5, I consider *HST* imaging of four nuclear flares: two likely thermal TDF candidates, ASASSN14ae and ASASSN14li, and two nuclear superluminous supernovae, CSS100217 and ASASSN15lh. These events all show properties that are atypical of their ascribed classification, such as the unusual lightcurves of the ASASSN TDF candidates or the presence of a superluminous supernova in a massive, low star formation host. I tighten the constraints on the positions on three of the four flares and discuss their properties within the context of the types of event considered above.

Chapter 2

Methods and Observations

The work within this thesis has been reliant on the data procured from a number of observing platforms across, and indeed above, the world; from the efficient detection of high-energy transients by the *Swift* Burst Alert Telescope (*Swift*-BAT), through the high resolution optical images produced by the *Hubble Space Telescope* (*HST*) to the sensitive spectroscopy of the Very Large Telescope's (VLT) X-shooter instrument. Each of these forms of observation require distinct methods for the obtaining, reduction and analysis of their associated data. Here I summarise the key detection techniques and analytical methods used extensively throughout this thesis.

2.1 Detectors across the electromagnetic spectrum

2.1.1 UV, Optical and Near-Infrared detectors

The most commonly analysed forms of emission within this work come from the visible band and its nearest neighbours, the ultraviolet and near-infrared. The properties of photons within these ranges make them suitable for the use of broadly similar detection techniques with the majority of modern observations across these three bands being made with the Charge Coupled Device or CCD (Martinez and Klotz, 1998).

These instruments are based on the principle of the photoelectric effect, whereby a photon is absorbed by a semiconducting material within each pixel, releasing an electron. This photoelectron is collected in a region of the pixel which is gated to ensure it does not recombine before it can be recorded, effectively turning each pixel into a light-sensitive capacitor. The size of the bandgap, that is the amount of energy required to promote an electron, is dependent on the choice of the semiconducting material and the level of doping. This means that CCDs can

be finetuned to be more efficient within certain frequency ranges. For example NIR CCDs, such as those mounted on the Gemini telescopes, are made from indium antimonide (InSb) or mercury cadmium telluride (HgCdTe) as the target photon energy is insufficient to excite electrons across the bandgap of silicon detectors that tend to be used for optical CCDs.

The ingenuity of the CCD is most evident during the read-out stage, at which point carefully controlled potential wells are manipulated to pass the charge from one pixel into the next in series such that each line of pixels is shifted down into the read-out line, which is itself shifted horizontally one pixel at a time into the read-out amplifier. This process enables the entire chip to be read-out systematically producing the resultant image.

The operational principles of a CCD have a number of advantages over the previously used photographic plates, not least of which is the quantum efficiency of detection; that is the fraction of photons that produce a detectable response within the medium used. In the case of photographic plates, the typical sensitivity is on the order of a few percent, while the most sensitive peaked at about 10% (Birney et al., 2006). The medium was also typically most sensitive at blue wavelengths. CCDs offer great improvement upon this with modern CCDs peaking at over 90% quantum efficiency with a typically wide range of wavelengths over which it remains sensitive. The ability to tune the bandgap by choosing the medium further improves the potential wavelength coverage afforded by CCDs. The linear response of a CCD, with each photon releasing a known number of photoelectrons, makes them far easier to use for source photometry than the non-linear photographic plates, while their regular, customisable pattern of pixels is preferable to the uneven and grainy appearance of plates.

While the invention and adoption of the CCD has revolutionised the field of astronomy, there are still a number of drawbacks with these devices. In addition to the remarkably complex electronics and manufacturing processes required to make the extremely sensitive CCDs used throughout professional astronomy, these devices usually require active cooling to temperatures well below 0°C (typical operating temperatures are about -100°C) to ensure that thermal noise is kept to a minimum. This is of particular importance in infrared astronomy, where the wavelengths of photons of interest are often within the range emitted by objects at room temperature. CCD's are also vulnerable to cosmic rays and the existence of malfunctioning (dead or hot) pixels that can make regions of individual images unusable. In addition, the read-out phase of a CCD's operation comes with it a new source of noise that increases with the number of exposures combined to make the final

image, though the extremely low read-out noise of many CCDs means this is rarely a concern. Read-out can also be a relatively lengthy process, adding to the overheads of an observation, though this can be reduced in situations where fast-read out is required through use of a second, unexposed chip that the charge is transferred into, enabling the main chip to start the next exposure while the second is read out.

A more commonly encountered issue is that of non-linearity. At low counts, CCDs are exceptionally accurate linear detectors. However, if an individual pixel is allowed to fill beyond a certain point, the Coulomb repulsion of the collected photoelectrons makes the collection of each new electron require more and more energy, thus changing the efficiency of the pixel. This has a considerable impact on the photometry as the lower efficiency results in lower implied count rates and thus bright objects look fainter than they actually are. In the case of very bright sources or very long exposures, pixels may become completely filled, or saturated, and any attempt to collect further photons may result in an overflow of charge, usually into the neighbouring pixels within the affected pixel's column. This effect can be avoided by using shorter exposures. Alternatively, if the effect is unavoidable, the damage to the required science goal can be minimised by the careful orientation of the chip to ensure any charge overflow does not affect the target of interest.

Finally, while the process of charge transfer in the read-out stage is extremely efficient, a small percentage of the charge in any one pixel is retained by traps in the CCD structure. While inherent in the design of CCDs, this effect can also become worse over time as energetic particles produce further defects in the CCD (Anderson and Bedin, 2010). While in ground-based systems, where sky emission is high, this has a negligible effect on the final image, low-background observations, such as those performed with *HST* WFC3 UVIS, exhibit clear streaking in the read direction as the counts from bright sources are retained each time charge is transferred. As this effect is cumulative upon transfer, it can be minimised by placing the target as close to the corner of the chip where read-out occurs. It can also be corrected for in the data reduction process (Anderson and Bedin, 2010) and is applied within this work on all *HST* WFC3 UVIS channel observations.

Most of the issues discussed above can be avoided or managed with careful calibration of the output data or specific observing parameters. As such, sensitive CCDs remain the instrument of choice for much of modern optical, infrared and ultraviolet astronomy. For most astronomical purposes, additional energy resolution is required to characterise the properties of sources. This is achieved in the camera design either by the application of a combination of filters designed to block all but the regime of light of interest to the observer or, in the case of spectroscopic

observations, by dispersing the observed light across the face of the CCD through use of a prism, diffraction grating or grism. Within this work, extensive use is made of observations made with *HST* WFC3, VLT FORS2 and X-Shooter, and Gemini GMOS all of which are instruments dependent on this technology.

2.1.2 X-ray and Gamma-ray detectors

With energies far above those of optical light, X-rays and γ -rays represent the emission generated by some of the most extreme environments the Universe has to offer, with emission mechanisms ranging from the peak energies of the highest temperature thermal flares to scattering from electrons in relativistic jets. However, the absorption of these high energy photons through the photoionization of single atoms is so efficient in the Earth’s atmosphere as to make their observation from the ground all but impossible. As such, the pioneers of X-ray observations made use of instruments mounted on balloons lifted into the upper atmosphere. In the current era, X-ray and γ -ray observations are generally performed by dedicated platforms launched into orbit. One of the earliest, a set of military satellites by the name of *Vela*, designed to police the then newly signed Nuclear Test-Ban Treaty, was instrumental in the discovery of the gamma-ray burst phenomenon (Bloom, 2011). Advancements in technology have enabled the application of more sophisticated detection techniques that are employed in the wide-area gamma/hard X-ray transient surveyors of *Swift* and *Fermi*, and in instruments dedicated to the observation of lower energy X-rays including those installed on *Swift*, *Chandra* and *XMM-Newton*.

Numerous methods for the detection of γ -rays and X-rays exist, from simple rate indicators like Geiger counters to more advanced methods capable of determining the direction and incident energy of the photon. As the latter properties are essential in imaging and spectroscopy, the methods employed in current observing platforms tend to have components capable of both spatial and spectral resolution.

While X-ray and gamma-ray emission is capable of penetrating conventional detectors used for longer wavelength emission, CCDs can nonetheless be used for X-ray observations. Unlike in the optical CCD case where in general a single photoelectron is produced by each incident (detected) photon, the much higher energies of X-ray and gamma-ray emission are capable of promoting numerous electrons through multiple secondary ionizations from the primary photoelectron. Thus while optical observations require integrated exposures in order to provide a significant detection, single X-ray detections promoting hundreds to thousands of electrons are significant in of themselves. This enables the operation of the CCD in “photon-counting” mode where exposure times are set short enough to detect individual

photons. Not only does this produce fine temporal resolution, enabling customisable binning of the data which is particularly useful in variability studies, the fact that each photoelectron represents the deposition of a known energy (typically a few eV depending on the detector properties), the number of counts received by the detector from each X-ray photon is effectively a measure of the energy of the photon. This means CCD observations of X-rays can act as both imaging and spectroscopy simultaneously making them extremely versatile. This approach is used in *Swift*'s Burst Alert Telescope (*Swift*-BAT) and X-ray Telescope (*Swift*-XRT), IBIS on *INTEGRAL* and the high energy instruments of *Chandra* and *XMM-Newton*.

A different approach was adopted for the *Fermi* gamma-ray burst detector, which instead makes use of the induced electron-positron production that occurs when a gamma-ray passes close to the nuclei of tungsten atoms in a series of thin foil layers. In this case, the energetics of the detected γ -rays are determined through the interaction of the electron-positron pairs with a cesium iodide calorimeter at the base of the instrument, the luminosity of the resulting scintillation being a measure of the energy of the incident photon.

The properties of γ -ray emission also makes it resistant to the conventional methods of focussing used in the optics of longer wavelength observing platforms. Instead, wide angle γ -ray detectors such as *Swift*'s Burst Alert Telescope tend to determine the direction of the incoming photon through use of a coded aperture mask. These components consist of a random, non-repeating pattern of lead tiles that strongly attenuate the gamma-ray emission attempting to pass through them. This array of tiles covers the telescope aperture such that any incident emission will be partially blocked. Spatial information can then be determined by modelling the shadow left on the detector.

The somewhat lower energies observed with X-ray telescopes makes it possible to use grazing incidence mirrors to focus the light in a similar way to the mirrors of optical telescopes. Hyperbolic and parabolic mirrors are arranged such that incoming light strikes at a low angle of incidence, below the critical angle for X-rays. The resulting mirror can only adequately focus X-rays approaching across a narrow range of angles and thus one mirror set has a very limited collecting area, so coaxial nested mirrors are used to expand the available collecting area. This technology is used in the *Swift*-XRT satellite and as a result is able to accurately determine the position of GRB X-ray afterglows to of order $1''$ precision.




Figure 2.1: An example of an optical ground-based image before the process of reduction is completed. The image has been inverted such that dark regions represent high counts. Note the uneven background and strong fringing which must be removed before the image can be analysed effectively. The image itself was taken with GMOS on Gemini in the i' -band.

2.2 Reduction and Analysis Techniques

As part of this work, I have reduced and analysed data from a number of different sources, each with a different science aim or purpose. Here I briefly discuss the key points of image and spectroscopic reduction.

2.2.1 Optical Image Reduction

Reduction of optical images is important in order to remove instrumental effects and contamination that plague the raw image (see Figure 2.1). The process of image reduction consists of 5 main stages:

- *Bias Removal:* Pixel values in CCDs do not begin at zero. CCD manufacturers add a small bias voltage to the operation of the CCD in order to fully sample the read-out noise (Birney et al., 2006). Thus, a zero exposure time image will result in counts in the resultant image (see Figure 2.2). This bias level is evident on every frame that is taken, be it science or calibration. To remove this bias, a zero-time exposure image is taken with the lens shutter (or aperture/telescope dome in the absence of a shutter) closed to ensure no stray light enters the CCD. The bias frame, which may be constructed from the median combination of a number of bias exposures, is subtracted from every image taken in the telescope run under identical conditions, including the raw science data. Alternatively, a small unexposed region of the chip, known as the overscan region, may suffice as an estimate of the bias level for an individual frame. However this region is incapable of accounting for variations in bias level across the chip, making the master bias image a generally superior method.
- *Dark Current Removal:* While CCDs are generally cooled to ensure their thermal noise is minimised, a certain level of “dark current” is always present. This effect is more prominent in devices intended to observe at near-infrared wavelengths which have smaller bandgaps designed to detect lower energy photons but also makes thermal electron excitation easier. The effect is also considerably stronger in hot pixels, defects in the CCD chip that may have existed since its manufacture or have developed since. Such defects can be caused through interactions with energetic particles, a situation of particular concern for space-based instruments. If significant enough to require it, the effects of dark current can be removed or reduced through the use of dark frames, exposures made with the CCD shutter closed or aperture covered in the case of shutterless CCDs. If the CCD is adequately temperature controlled, the effects of dark current can be effectively scaled from shorter exposure times and subtracted from the calibration and science products.
- *Flat Fielding:* The quantum efficiency of individual pixels can vary significantly across the CCD chip resulting in varying sensitivity. This causes issues when attempting to determine photometry of objects across the field. In order to correct for this effect, images are taken of a flatly illuminated surface, such as the interior of the dome of the telescope, or of the twilight sky before stars begin to dominate the image (see Figure 2.2). The combination of these images, normalised to the median value, is used to scale the response of each

pixel.

- *Fringe removal:* At red wavelengths, the superposition of incoming photons with reflected photons from the underside of the chip (thin film interference) produces an effect known as fringing. This is a consequence of the variable thickness of the chip meaning the reflected photons interact with differing path lengths and thus produce bands of varying intensity across the image. As the fringes are effectively fixed to the CCD and vary little image to image, they can be removed by taking multiple images of the target at small offsets between the frames, moving the objects in the field while keeping the fringes in the same position. A median average of the available frames fixed on the image coordinates (as opposed to world coordinates) produces an image with objects removed, leaving just the fringing which can then be subtracted from the individual science frames (see Figure 2.2).
- *Mosaicking and Combining:* The final step of image reduction consists of producing the final image by mosaicking and combining individual frames. Many instruments are made of multiple CCD chips with gaps between them. In order to produce a single image, the separate images can be mosaicked together with corresponding world coordinate system (WCS) offsets. The subsequent process of combining multiple frames serves a number of purposes in addition to the removal of fringing discussed above. Science objectives rely on achieving a sufficient signal to noise ratio on the target of interest. However, if the science target or a nearby object are particularly bright, the inherent non-linearity and even saturation that can occur as the pixel potential well fills may adversely affect the observations. Instead, repeat observations with exposure times short enough to not be affected by these effects can be combined to lengthen the effective exposure time of the observation. In addition, by applying small offsets in position to each exposure (dithering), cosmic rays can be removed when the images are combined based on WCS coordinates, applying either a median average or a sigma clip that ignores pixels from the average that show large deviations, indicative of cosmic rays affecting a single frame.

Each of these steps are generally present in the image reduction process, though dark current and fringe removal are not always required and additional steps can be required depending on the instrument and type of observation used. An example of the finished product following reduction is shown in Figure 2.3.




Figure 2.2: Examples of typically used calibration frames used in the optical image reduction procedure. Each image (Top, Middle, Bottom) consists of a slice of the three GMOS CCD chips (Left, Middle, Right) with clear gaps between the three chips. Again the image is inverted so that dark regions indicate high counts. The image scaling is not uniform between the three frames and instead has been chosen to emphasise the key features of each frame. (Top) An example of a bias image that is subtracted from all images in order to remove the bias level added to every frame. Note that the three CCD chips have subtly different bias levels. (Middle) An example of the flat frame, an image of a uniformly illuminated surface or the twilight sky, meaning any deviation in counts is due to the sensitivity of the chip. The chips have clearly differing sensitivity both within each chip and between the separate chips. (Bottom) An example of a fringe image, the median combination of the final science frames that leaves the fringing which is fixed to the CCD image coordinates.




Figure 2.3: The earlier GMOS image after it has undergone the complete reduction process. Note that, with the exception of a strong defect around the bright object in the top right corner of the image, the image now has a flat uniform background with little to no signs of contamination.

***HST* Sub-pixel Dithering and Drizzling**

In the case of many ground-based observing platforms that operate at optical, UV or NIR wavelengths, atmospheric effects dominate the resolution of astronomical observations. This regime, known as seeing-limited, is therefore generally a function of the specific atmospheric conditions at the time of the observation. In some cases, however, where adaptive optics are used on the ground or the telescope is situated above the atmosphere, the instrument can become diffraction-limited, where resolution is instead primarily governed by the optical power of the instrument. Such a situation can result in the point-spread function (PSF) of observations being so narrow as to be under-sampled by the pixel scale of the CCD, as is the case with the *Hubble Space Telescope*’s Wide Field Camera 3 (*HST* WFC3). This makes the modelling of the point source function difficult, reducing the precision of astrometric measurements and hampering morphological modelling.

With the ability to perform precision pointing of telescopes comes a solution to this undersampling issue. Whereas normal dither patterns are made with integer pixel offsets to enable simple combination into the final science image, the ability to offset subsequent frames at non-integer, sub-pixel values means it is possible to use each image as a separate sampling of the point spread function. Unfortunately the combining of such frames into a final, higher resolution frame, is non-trivial.

The simplest method, known as interlacing, takes each input pixel and places it in its corresponding place in a finer grid. Thus 2 images offset by half a pixel in the x direction would each fill in alternating pixels in the grid with half the pixel scale of the original in the x direction. However, this method requires that the pointing precision be extremely high and regular, and is incapable of dealing with distortions across the chip.

An alternative method, known as “shift-and-add”, works by producing a fine grid onto which each pixel from the input image is sub-divided, shifted and summed, making it possible to deal with arbitrary dither patterns. However, this effectively reconvolves the output image with the input pixel size, resulting in a reduction in resolution. In addition, as each input pixel contributes to multiple output pixels in the final image, the output suffers from correlated noise that can be difficult to accurately account for. Finally, this method is still incapable of dealing with geometric distortions across the chip and also cannot easily deal with missing data which result from cosmic ray removal or defects in the detector.

A technique that makes a compromise between these methods, borrowing on features of each, is Variable-Pixel Linear Reconstruction, or “drizzling” (Fruchter and Hook, 2002). This technique works by following a similar procedure to the

“shift-and-add” technique, but scales down the size of the input pixel before recombination, thus reducing the effect of the convolution by the input pixel’s footprint. This means the process works on a continuum between the above methods whereby a scaling of the input pixel by a factor 1 (i.e. no scaling) is equivalent to the “shift-and-add” technique, whereas a scale factor of 0 shrinks the pixel down to a delta function. As such, each input pixel contributes to only one output pixel, removing correlated noise and effectively recreating the interlacing technique. The key is in specifying an input pixel scaling that is small enough to retain the resolution of the image, but keeps the number of pixels contributing to each output pixel relatively uniform. This method is also capable of dealing with arbitrary dither patterns, rotations and geometric distortions that hamper the other methods (Gonzaga et al., 2012). This technique, written into the Python function `ASTRODRIZZLE` in the package `DRIZZLEPAC` (Fruchter and Hook, 2002; Fruchter, 2010), has now become a standard part of the *HST* pipeline.

While the effect of correlated noise can be minimised by careful choice of the input pixel scaling parameter, it is nonetheless a concern and must be corrected for in photometric uncertainties. It occurs because the naive assumption that the total variance in any output pixel is the arithmetic sum of the fraction of each input pixel that contributed to it multiplied by the variance of said input pixel ignores cross terms between adjacent fractions of the input pixel. Thus if an input pixel, with variance σ^2 contributes a fraction a and b to two output pixels respectively (such that $a + b = 1$, the total area of the input pixel), then the sum of the fractional variances added to each output pixel $(a^2 + b^2)\sigma^2$ is less than σ^2 by a factor $2ab\sigma^2$, the cross-term between the two. This effect becomes more difficult to correct for the more output pixels each input pixel contributes to. While on large scales the true RMS noise scales with the linear size of the region, as expected, on small scales, and particularly when approaching individual pixel level, the estimated single pixel noise must be corrected for the correlated noise term (Casertano et al., 2000). This is achieved by scaling by a factor R :

$$R = \frac{\sigma_C}{\sigma_P} \quad (2.1)$$

where σ_C is the standard deviation of the correlated noise and σ_P is the standard deviation of the effective single pixel noise. While accurate determination of this value for arbitrary dither patterns is extremely complex, under the assumption of a uniform dither pattern with many dithers, the effect can be approximated by the following relations (Gonzaga et al., 2012): If p is the input pixel scaling parameter used in the drizzle method and s is the output pixelscale as a fraction of the input

pixelscale and $r = p/s$, then for $r > 1$:

$$R = \frac{r}{1 - \frac{1}{3r}} \quad (2.2)$$

and for $r < 1$:

$$R = \frac{1}{1 - \frac{r}{3}} \quad (2.3)$$

As completing photometry in a large aperture of diameter N is effectively the same as summing into one large pixel of linear size N , s in the above equations can be replaced by Ns . This means that in the limit of large N , $R \rightarrow 1$ as expected. Similarly, if the input pixel scaling $p \rightarrow 0$ (that is, effectively approaching applying the interlacing method), $R \rightarrow 1$ again as expected.

2.2.2 Optical Spectroscopic Reduction

In addition to the same reduction techniques required for optical image reduction, the reduction of spectroscopic observations requires additional calibrations. Note that these steps refer primarily to long slit spectroscopy. Slitless, fibre-fed and echelle spectroscopy require additional steps that are not included here:

- *Rectification:* In order to accurately wavelength calibrate across the the entire chip, it is generally necessary to rectify the spectrum such that one axis of the chip represents spatial and the other spectral measurements. This involves modelling the curvature of sky emission lines and producing a transformation that maps the distorted spectrum onto orthogonal axes.
- *Wavelength Calibration:* Wavelength calibration can be achieved in one of two ways. The first involves taking observations of an arc lamp using an identical set-up to that used to take the science frames. The pattern of emission lines registered by the arc spectrum can then be fitted and interpolated to determine the wavelength calibration of the spectrum. Alternatively, sky line emission directly on the science frames can be used in the same way with the added bonus that its presence on the frame under study means there is no systematic error added by transforming the arc calibration to the science frame. However, sky line emission is most apparent at red wavelengths with little emission bluewards of 5500Å, meaning the calibration can become poor for spectra that extend into the blue or NUV, reliant on an extrapolation from the region with sky line coverage.

- *Background Subtraction:* While sky line emission can provide a useful wavelength calibration, it is also a strong source of contamination that must be removed before analysis of the underlying spectrum can be made. If the process of rectification has been adequately applied, this can often easily be achieved by selecting regions of the frame containing only sky line emission and subtracting the median contribution in the spatial direction. However, crowded fields, extended emission and poorly applied rectification can make this process more difficult. In addition, while the average contribution from the sky emission may be accurately determined, the Poisson noise may nonetheless drown out the target signal in regions of intense sky line emission.
- *Flux Calibration:* As the sensitivity of the CCD is a function of the frequency of the incident radiation, the flux calibration is non-uniform in the spectral direction. Typically, flux calibration is achieved through observations of a spectrophotometric standard star, an object with a known flux calibrated spectrum that has little to no variability in time. The same instrumental set-up and calibration is applied to the standard star and the ratio of the known flux of the star to the recorded counts gives the flux calibration as a function of wavelength which may then be applied to the science frame. In some cases this calibration must be altered to include correction for telluric absorption from the Earth's atmosphere, absorption features that cannot be removed via the same sky subtraction method that deals with emission sky lines. This can be achieved by fitting models to spectra of observations of telluric stars, with negative deviations from the model indicating the presence of telluric absorption.
- *Extraction:* Upon completion of the calibration, a one-dimensional spectrum may be extracted from the two-dimensional image. However, for observations taken off of the parallactic angle (that is, aligned with the target meridian), atmospheric dispersion causes a wavelength dependent deviation of the object's position on the chip that is not accounted for in image rectification. If present and significant enough, modelling of the locus of continuum emission (the trace) can be used to determine the position of the extraction aperture. However in cases where continuum emission is not detectable or the target has a complex morphology, it can be difficult to determine the optimal extraction method.

This method of analysis was applied to the GMOS and FORS2 spectroscopy within this thesis. The reduction of echelle spectroscopy, such as that produced by X-

Shooter, requires additional reduction steps to produce a single continuous spectrum from the oblique pattern of slightly overlapping wavelength ranges known as orders formed by the method. This involves identifying the positions of, and modelling the trace for, each order in preparation for extraction.

2.3 Astrometry

The location of a transient within its host is an important indication of its origin. Finding a strong preference for a particular type of transient to inhabit regions with strong UV emission might indicate an association with star formation, and thus young or massive stars, while a position far from the bulk of the galaxy could indicate the progenitor had received a kick prior to the event, ejecting it from its host. Indeed, in the case of transients considered within this work, one major requirement for inclusion is that they come from the nuclear regions of their hosts. Thus, both initially in order to be included in the study and again later if higher resolution imaging becomes available, the astrometric position of the transient must be determined. This position must also be compared to the morphology of the host and thus to the expected position of the supermassive black hole in the system which is generally assumed to be at the centroid position of the host light (though note that this assumption can fall apart in interacting hosts or on very small scales). In this section I discuss the various techniques which are used within this work to determine these properties.

2.3.1 Astrometric Matching

Several of the techniques discussed below require that the coordinate transformation between two images be known. This is important both in attempting to determine the relative positions of objects within the two images, as well as for cases where the images themselves need to be transformed to identical coordinates, perhaps in preparation for subtraction. While world coordinate system (WCS) information is encoded into many images taken on modern instruments, this is often insufficient for these purposes as it is subject to systematic offsets. Instead, it is preferable to determine the transformation between the two images using objects in the field.

Ideally the astrometric matching is completed using a number of point sources in the field, fitted using a simple Gaussian or Moffat profile, which are typically good fits in seeing-limited observations, or by using a central moment finding routine such as IRAF IMEXAM. Identifying point sources is done first by eye and then, having fitted each possible point source, a population of fitted profiles with a consistent

and small width is identified. This is because ideal point sources represent sources too small to be resolved and thus are merely scaled point spread functions, the narrowest profile a real¹ source can have. Any extended object, assuming the fitting process does not fail, should have a profile which is wider than point sources and so can be excluded. In cases where the point spread function is not an ideal Gaussian or Moffat profile, for example in cases where the pointing of an observation drifts during the exposure, elongating objects in the drift direction, or the observations are diffraction-limited and the optical system is not ideal (e.g. *HST* point spread functions are non-trivial), any systematics introduced should be equally felt by all point sources in the field so that the net relative systematic is zero.

With a sample of objects for use in the fitting process identified, the positions of each object in both images are used as variables in a fitting routine such as IRAF GEOMAP. This produces a functional form for the coordinate transformation which can include offsets in pointing, orientation and pixel scale, while the fitting polynomial's order can also be changed to allow for distortion of the image, assuming enough point sources have been identified to accommodate fitting all of the required variables. An estimate of the RMS of the fit can also be obtained which is useful for determining the astrometric uncertainty in later analysis. When required, anomalous positions, more often than not the result of misidentifying an extended source as a point source, can be removed from the fit. However on some occasions, as these point sources can be relatively local stars, they can have non-negligible proper motions meaning they can move across the field from one image to the next, requiring that they too be removed from the fitting process.

Most problems associated with astrometric matching can, for the most part, be overcome by having as large a number of well-detected point sources to fit the positions of as possible. However, in sparsely populated fields, with minimal overlap between the images, or with necessarily small fields of view, perhaps due to image buffer constraints on space-based instruments (e.g. *HST* sub-arrays), it can be difficult to find enough sources for comparison between the two. In cases where an insufficient number of point sources are available, extended sources can be used, though the systematics are more difficult to determine in this case and the fit is generally of poorer quality.

¹i.e. not a CCD defect, such as a hot pixel, or a cosmic ray hit, both of which should have been dramatically reduced in severity and frequency by the reduction process. They are also often identifiable by eye or fail to be fitted with simple profiles

2.3.2 Image Subtraction

While some transient phenomena are, at least for a short time, far more intrinsically luminous than their hosts, in most cases any determination of the position of the transient will necessarily be confused by the underlying host emission. In order to remove this host contribution and leave only the transient emission for modelling, a technique known as image subtraction can be used.

In its simplest and most ideal form, an image taken either before the transient occurred or after the transient has faded below the instrumental sensitivity is aligned and subtracted from an image taken while the transient was bright. As the host contribution is effectively unchanged over the short time between the two images, the subtraction should produce an image that includes only transient light with no appreciable contamination from the host. Ideally the two images should have been taken with the same instrument, in the same configuration and waveband, in identical observing conditions and with the the same pointing and orientation. If this is the case, the reduced images can be simply scaled by exposure time and subtracted from each other.

In reality though, many of the above requirements are not met, requiring a more in depth method of analysis. Typically, even if the observations were taken in near-identical conditions and were designed to have identical pointings and orientations, this will usually be confirmed first through astrometric matching (see subsection 2.3.1 above) and any required transformation applied. Large transformations are discouraged, however, as every transformation the data is put through adds systematic distortions and flux losses/inconsistencies that will affect the outputs. This effect can be minimised in cases where specialised software is available for the individual instrument, designed to deal with these problems, such as *ASTRODRIZZLE* for *HST* (Fruchter and Hook, 2002; Fruchter, 2010). Thus, if both of the intended images have been taken in identical instrumental configurations with *HST*, given the minimal changes in point spread function in this diffraction-limited case, the images can be merely aligned, scaled and subtracted. This method is used extensively in Chapter 5.

However, for most ground-based, seeing limited observations used in this work, changes in observing conditions will lead to an inevitable change in the point spread function, requiring one image be convolved to match the seeing of the other before the images may be subtracted. Within this work, the image subtraction software *ISIS* (Alard and Lupton, 1998; Alard, 2000) is used for this process. This software uses a reference image, typically the one with the best seeing, which is then convolved with a spatially variable kernel that allows for changes in the point

spread function across the image. The software is also capable of aligning images that have minor mismatches in pointing or orientation, though, once again, large translations should be avoided. Once complete, the convolved and aligned images can be subtracted from one another and the transient light analysed separately from the host emission. This process is used to subtract ground-based images with quite different seeing within Chapter 3.

2.3.3 Modelling Galaxy Morphology and Determining Centroid Positions

The final stage in determining the position of a transient within its host is to analyse the morphology of the host itself and determining the centroid position of both it and the transient. This can be completed in a number of ways depending on the signal to noise ratio of the detection, the resolution of the imaging, the presence or absence of clear structure within the host and any secondary science goals the study may have.

In the case of a low significance, low resolution (i.e. close to being unresolved) detection of a host galaxy, it is quite possible that no clear structure will be visible in the host, particularly if the host is at high redshift. In this case, there is often little benefit in completing a complicated morphological fit to the data. As such, in order to find the central position of the galaxy, a simple centroiding method, such as that used by IRAF IMEXAM can be used. The central moments of the object is calculated and returned as an estimate of the central position of the source. The uncertainty on this method can be reasonably approximated as the theoretical uncertainty on the centroid of a Gaussian with the FWHM and signal to noise (S/N) equal to that of the source, given by:

$$\sigma_{\text{cen}} = \frac{\text{FWHM}}{2.3 \times S/N} \quad (2.4)$$

This method is also used when determining the position of the centroid of the transient emission following subtraction or in higher signal to noise cases where an otherwise well-resolved galaxy has a clear, relatively narrow central peak that is unlikely to be affected by the more extended parts of the host.

In cases where the extent of the galaxy is larger and the signal to noise allows for it, more complicated methods can be applied and the precise morphology of the host can be determined. In particular, within this thesis, the galaxy morphology modelling code, GALFIT (Peng et al., 2002, 2010), is used extensively. This program enables the decomposition of galaxies into an arbitrary number of morphological

components that can include, for example, point sources, gaussians, sérsic profiles and edge-on disks. Each has a set of input variables that are then fitted to the 2-dimensional image. The best fit model is determined to be that with the minimum reduced chi-square when comparing the input model and actual image. The results include the best fit model parameters and estimates of the uncertainties on each, where possible.

However, a number of issues with this method exist. For example, in addition to the usual issues in fitting routines of local minima in the chi-square distribution and determining reasonable “first guesses” for the input values, the number of input components is based on a judgement call by the user. It is therefore difficult to tell when adding a new component to the model is a useful addition or an over-fit. Visual inspection of the residual image is usually the best way to attempt to determine whether another component has actually improved the fit or not. Further, because the sky background is also a fitted variable within the program and the input section of the image can be made arbitrarily large, a particular model’s goodness of fit can be skewed by choosing a larger region of sky. While experience with the program can help determine the “ideal” size of the fitting region to enable the accurate fitting of the sky without arbitrarily adding background pixels that improve the apparent quality of the fit, the uncertainties on the best fit parameters should always be viewed with caution. This method is used briefly in Chapters 3 and 4 and extensively in fitting the galaxies in Chapter 5.

2.4 Magnitude Systems and Cosmology

The apparent magnitudes, m , used throughout this thesis are defined as:

$$m = C - 2.5 \log_{10}(CR) \quad (2.5)$$

where CR is the count rate detected by the imaging device and C is an arbitrary zero-point offset. Unless otherwise stated, this constant will correct the magnitudes to the AB magnitude system, which is defined such that:

$$F_\nu = 10^{\frac{23.9-m}{2.5}} \quad (2.6)$$

where F_ν is the flux of the target in μJy .

In cases where absolute magnitudes, M , are quoted, they are defined:

$$M = m - 5(\log_{10}(d) - 1) + K \quad (2.7)$$

where d is the luminosity distance to the object in parsecs and K is the K-correction. While typically this correction is complex function of the bandpass and object's spectral energy distribution, in the absence of such information a first order approximation of:

$$K = 2.5 \log_{10}(1 + z) \quad (2.8)$$

is used, where z is the cosmological redshift of the source.

The cosmology adopted within this thesis is the standard Λ CDM cosmology with $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_M = 0.3$ and $\Omega_\Lambda = 0.7$.

Chapter 3

The Candidate Relativistic Tidal Disruption Flare Swift J1112.2-8238

3.1 Introduction

Whereas most TDFs are detected near the peak of their high temperature thermal spectral energy distributions, usually in the UV and soft X-ray, a handful of events have been discovered as part of the *Swift* gamma-ray transient monitoring program. Each had remarkably long-lived high energy emission (10^{47} erg s $^{-1}$ at 10^6 seconds after the trigger), rapid and high amplitude variability even at late times, moderate optical emission (10^{43} erg s $^{-1}$ Levan et al., 2011; Bloom et al., 2011) and evidence of a rising radio flare (Zauderer et al., 2011). Each also came from the central regions of galaxies with no evidence of ongoing AGN activity. Clearly atypical events, these flares have their most popular explanation in the tidal shredding of a star that launches a moderately relativistic jet. However other explanations, including a massive stellar core collapse event and the tidal shredding of a white dwarf, remain plausible.

The first such event, Swift J1644+57, was detected in April 2011 (Levan et al., 2011; Bloom et al., 2011). In the seven years of *Swift* operation preceding the detection of Swift J1644+57, no other event of its type had been detected. However, only two months later, Swift J2058+05 became the second candidate relativistic tidal disruption flare (Cenko et al., 2012). While changes to the conditions for event detection with the Burst Alert Telescope may have improved its ability to detect

longer-lived bursts, with such a short gap between the two events, the question was raised as to whether there were other events present in the *Swift* archive that had been overlooked (Levan et al., 2011). Attempts to find further examples failed to find any instances preceding the detection of Swift J1644+57 (Krimm et al., 2013).

However, in June 2011, only a month after Swift J2058+05, one further burst, Swift J1112.2-8238 (Krimm et al., 2011b) was discovered with properties that made it a potential new candidate. Its duration was far in excess of most GRBs, and its position far outside of the Galactic plane, made it a likely extragalactic event. While a short-lived multi-wavelength monitoring program was triggered, the source was mostly ignored, until it was reidentified in an archival search. In this chapter I cover the analysis of this data, including the high energy *Swift* observations and ground-based optical imaging and spectroscopy with Gemini and the VLT, and discuss the implications for the nature of the event.

3.2 Observations

3.2.1 Swift BAT data

The outburst of *Swift* J1112-8238 was originally discovered by the *Swift* satellite (Gehrels et al., 2005) in a four day integration¹ by the Burst Alert Telescope (BAT, Barthelmy et al., 2005) between 2011 June 16 and 19 (MJD 55728-55731, Krimm et al., 2011b). For the remainder of this work, the trigger time is taken to be the beginning of the first day of the integration, 2011 June 16 UT 00:01, although in practice a precise trigger time is poorly defined. The count rate in γ -rays across this period was $(2.9 \pm 0.7) \times 10^{-3} \text{ ph s}^{-1} \text{ cm}^{-2}$ (1σ uncertainty) with a peak daily average rate of $(1.9 \pm 0.5) \times 10^{-2} \text{ ph s}^{-1} \text{ cm}^{-2}$ (1σ uncertainty) recorded on the 16th, both in the 15-50 keV band (Krimm et al., 2011b). This peak, though high, was not in itself sufficient to produce a trigger of the automated transient monitor which requires a 5σ count rate for new sources (Krimm et al., 2011b, 2013).

In order to constrain pre or post flare activity coincident with the location of the flare, I used the available BAT daily average lightcurves² extending back as far as the launch of *Swift* in 2005. By computing the weighted average count rates across any four-day window within the archive (equivalent to the integration of the

¹The *Swift* Burst Alert Telescope has a number of methods for the identification of new transient flares. Rate triggers look for significant increases in count rates over short timescales of up to 64s, and it is this method that produces the majority of GRBs. Image triggers are based on images combined over timescales of up to 20 minutes (a full pointing). However, in the case of extremely long-lived transients as is the case here, daily mosaics are produced that can be combined on even longer timescales of up to 16 days (Krimm et al., 2013)

²<http://swift.gsfc.nasa.gov/results/transients/>




Figure 3.1: The *Swift*-XRT spectrum of Swift J1112-8238, produced by combining all available observations. The fitted model is indicated in red and the residuals presented in the lower panel. Figure produced with the Swift-XRT tools at the UK Swift Science Data Centre.

original detection), I found that there was no activity above a 4σ threshold, a finding that is also noted by Krimm et al. (2013). Given that *Swift*-BAT scans only 80-90% of the sky per day (Krimm et al., 2013), there are occasional gaps in the data within which a similar signal could have been missed. In the case of Swift J1112-8238, within the studied lightcurve the probability of any day being both observed and followed by 3 further observed days (i.e. four consecutive days) is $\sim 82\%$, making the first order probability of having missed a similar signal $\sim 18\%$ and is therefore unlikely.

3.2.2 X-ray data

Initial X-ray data was obtained by the *Swift* X-Ray Telescope (XRT; Burrows et al., 2005) in a 3000s target of opportunity observation approximately 10 days after the initial trigger (MJD 55741.7, Krimm et al., 2011b). The source was well detected with an observed count rate of $0.228 \pm 0.017 \text{ ph s}^{-1}$. An X-ray monitoring programme continued for a further 30 days with all observations obtained in photon

counting (PC) mode. The reduced XRT lightcurve and fitted spectrum were obtained from tools available at the UK *Swift* Science Data Centre³, and were created using the techniques outlined in Evans et al. (2007, 2009, 2010). The enhanced X-ray position derived from the UVOT boresight correction is RA= 11:11:47.32 DEC=-82:38:44.2 (J2000) with a 90% error radius of 1.4". The combined spectrum of all available PC-mode observations, setting the redshift of the event $z = 0.8901$ (as determined in Section 3.2.4), is well fit by an absorbed power-law, with a fit statistic $W_{\text{stat}} = 565.29$ (585). The output fit parameters are photon index $\Gamma_{\text{ph}} = 1.33 \pm 0.08$ and N_{H} (int) = $2.4^{+1.8}_{-1.6} \times 10^{21} \text{ cm}^{-2}$ while the Galactic value is $1.8 \times 10^{21} \text{ cm}^{-2}$ (Willingale et al., 2013), where all uncertainties are quoted as 90% confidence intervals. This produces a counts-to-flux ratio of $5.39 \times 10^{-11} \text{ erg cm}^{-2} \text{ ph}^{-1}$ observed ($6.12 \times 10^{-11} \text{ erg cm}^{-2} \text{ ph}^{-1}$ unabsorbed) in the 0.3-10 keV band. This means the first observation had an unabsorbed flux of $(1.4 \pm 0.1) \times 10^{-11} \text{ erg cm}^{-2} \text{ s}^{-1}$ (0.3-10 keV, 1σ uncertainty).

Splitting the available counts across all observations into bins of ~ 300 counts each, it is possible to look for evolution in the spectra of the event. There is evidence for a spectral hardening at late times with the hardness ratio of the first (~ 14 days post trigger) and last (~ 40 days post trigger) bins being 1.36 ± 0.11 and 2.17 ± 0.22 respectively, a 3.2σ difference (see Figure 3.2).

The light curve over the same period exhibits a gradual decay but with marked variability (a factor of ~ 2 in flux) between individual snapshots (uninterrupted pointings). Additional late time observations were obtained in April 2014, with a total XRT exposure time of 6960.3 s (in PC mode). This observation provides an upper limit on the source flux of $F_{\text{X}} < 4 \times 10^{-14} \text{ ergs s}^{-1} \text{ cm}^{-2}$ (99%, determined via the Bayesian method of Kraft et al., 1991) unabsorbed. This is a factor of ~ 250 fainter than the peak flux, confirming the source's transient nature.

3.2.3 Optical Imaging

Following a UVOT non-detection ($b > 22.0$ mag, Krimm et al., 2011b) made at the beginning of the X-ray monitoring programme, observations were obtained in the i' band with the Gemini Multiple Object Spectrograph on Gemini South (GMOS-S, Hook et al., 2004) at 2011 July 3 UT 00:58, 17 days after initial trigger (Berger and Chornock, 2011). Later follow-up was performed with GMOS-S in the r' and i' -bands at 1.5 years post-trigger (starting 2012 December 13 UT 06:50), and with the FOcal Reducer and low dispersion Spectrograph 2 (FORS2) on the Very Large Telescope (VLT) at 2 years post-trigger (starting 2013 August 31 UT 23:31) in I

³http://www.swift.ac.uk/user_objects




Figure 3.2: The *Swift*-XRT lightcurve in the soft (0.3–1.5 keV) and hard (1.5–10 keV) bands and the evolution of the ratio of the two when the available counts have been split into equal bins of ~ 300 counts a piece. The event shows evidence for spectral hardening between the first and last bins at the $\sim 3\sigma$ level. Figure produced with the Swift-XRT tools at the UK Swift Science Data Centre.




Figure 3.3: (Left) A GMOS-S i' -band finding chart for Swift J1112-8238. (Right) A comparison between the source at ~ 20 days (Top) and at ~ 1.5 years (Bottom) post trigger, each panel $15''$ across. The extended host's structure is far clearer in the later epoch, due in part to both the decline of the optical transient and the greatly improved seeing.

and z' . In addition, as part of the GMOS-S spectroscopic follow-up, a number of short exposure acquisition images were taken in r' (starting at 2012 December 16 UT 07:30 and 2012 December 23 UT 05:16) and in i' (starting at 2014 January 3 UT 07:01, ~ 3 years post trigger).

For each data set the reduction follows the basic outline presented in Chapter 2. Specifically, reduction of the Gemini-GMOS imaging was completed using the GMOS specific functions within the GEMINI package distributed for use with IRAF. The task GPREPARE was run on all raw images to add the various header information required by the later tasks. A master bias frame was built from a series of raw bias frames with the task GBIAS and was used in the processing of the raw flat and science frames. The master flat frame was built with the task GFLAT from multiple raw flat frames. The raw science images were then reduced with the task GIREDUCE and the i' -band images, which had significant fringing, were corrected with the tasks GFRINGE and GIRMFRINGE. Finally, the 3 CCDs of GMOS were mosaicked together (GMOSAIC) and combined with median-averaging (GEMCOMBINE).

The FORS2 imaging was reduced with the standard ESOREX pipeline which provides tools for the production of a master bias frame (FORS_BIAS), a master flat frame (FORS_IMG_SKY_FLAT) and the reduction of the science data (FORS_IMG_SCIENCE). The resulting reduced science frames were combined with median averaging using the IRAF function IMCOMBINE prior to analysis.

The presence of a nearby bright star ($R = 15.8$ mag at an angular distance of $\sim 5''$, Figure 3.3) complicated the analysis of this source. A number of methods were attempted in order to remove the contamination from this star prior to analysing the source of interest. The first involved rotating each image by $45^\circ/90^\circ$ centred on the bright star. The resulting image was then subtracted from the original. While this did appear to produce a good subtraction of the star's light as judged through visual inspection of the product, there was limited clean sky surrounding the source due to the already crowded field surrounding the star producing a number of negative depressions in the resulting image (see Figure 3.4). The second method attempted was to produce a model of the radial profile of the star using gaussian and moffat functions. However, the star is saturated in all of the images, meaning only a small and somewhat uncertain fraction of the wings of the star in each case could be used to fit the whole. The large uncertainty in the correct model choice and range of radii to fit over made it difficult to form an appropriate model and produced clear residuals in the image, so this method was also rejected (see Figure 3.4). The brightness of the star also hampered the use of other model point spread functions (PSF), produced with the IRAF function PSF from a combination of unsaturated

point sources in the field, as the subtraction of the model resulted in an extremely noisy background when scaled to the brightness of the contaminating star (see Figure 3.4). Finally, a model PSF was constructed as a median-averaged radial light profile and subtracted from the star. This produced the cleanest subtraction of the lights in the star’s wings and was adopted as the method used to remove the contamination from each image (see Figure 3.4).

Photometric calibration for the i'/I band was completed through comparison with observations of photometric standards analysed via the ESOREX FORS2 pipeline, the expected systematic offset between the GMOS i' and FORS2 I filters having been deemed negligible in this low signal to noise regime. The ESOREX pipeline analyses standard star images taken on the same night with known magnitudes to produce an airmass corrected zeropoint for the science frames with FORS_ZEROPOINT, enabling the direct calibration of the FORS2 photometry. To calibrate the GMOS images, photometry of a large number of objects in the field in both the GMOS and FORS2 science frames was taken and the raw (uncalibrated) magnitude values compared. The systematic offset between the two was determined and taken to be the zeropoint offset between the two images. Thus the zeropoint of the frames, and subsequently photometry of the target, could be determined in the GMOS images.

The non-standard filter z' was instead calibrated through comparison with the FORS2 standard star, Feige 110, which was observed within a few nights of our observations. Using a standard star observed on a separate night has the potential of suffering from various systematics produced by the different observing conditions. However, all observations were made during dark time, limiting the potential contamination from the Moon, and the atmospheric conditions on the different nights were very similar (seeing $\sim 1.1''$ at an airmass of 1.0, relative humidity $\sim 10\%$ and no evidence of clouds⁴) limiting any systematic offsets in the calibration. Feige 110 is a spectrophotometric standard star with a well established flux calibrated spectrum. The spectrum, accessed via the ESO archives⁵, was interpolated and convolved with the z' filter response curve and the CCD quantum efficiency using the ASTSED function in the python package ASTLIB. This provided the magnitude of Feige 110 in the z' band. By completing aperture photometry on the reduced image of Feige 110, the zeropoint of the CCD was determined and the target photometry calibrated. Finally

⁴based on the observing conditions reported by the ESO Differential Image Motion Monitor, DIMM, and the VAISALA Meteorological station, METEO; data accessed at <http://archive.eso.org/cms/eso-data/ambient-conditions/paranal-ambient-query-forms.html>

⁵<https://www.eso.org/sci/observing/tools/standards/spectra/feige110.html>




Figure 3.4: Examples of the methods used to subtract the contaminating star in the optical imaging of Swift J1112-8238. In each case, the contaminating star is indicated with a red cross, the source of interest with a green circle and any other subtracted sources with blue circles. (Top Left) An attempt to subtract a 45° rotated copy of the star centred on the red cross. The subtraction appears to be good, but the crowded field leaves numerous negative depressions that severely limit the available clear sky for background determination. (Top Right) Modelling the point spread function with simple gaussian/moffat profiles. In this case, additional objects in the field have also been subtracted with relatively small residual emission. However, the attempt leaves clear residuals surrounding the bright star, a consequence of the difficulty in modelling the whole star based on only a narrow part of the unsaturated wings of the point spread function. (Bottom Left) An attempt to use a point spread function produced from other point sources in the field. Again, this produces a clean subtraction of other sources in the field but the large scaling required to model the bright star greatly increases the noise in the background. (Bottom Right) A subtraction completed by producing a median-averaged radial profile of the star. At the radial distance of the source of interest, the contaminating star's emission has been well subtracted without the issues of the negative signals produced by the rotation method above. As such, this method was adopted to subtract the star's emission.

MJD	ΔT (d)	Instrument	Filter	Magnitude	Seeing (\prime)
55745.1	17.1	GMOS-S	i'	22.10 ± 0.10	1.4
55749.0	21.0	GMOS-S	i'	21.96 ± 0.10	1.4
56274.3	546.3	GMOS-S	r'	23.74 ± 0.17	0.7
56274.3	546.3	GMOS-S	i'	22.76 ± 0.12	0.7
56277.3	549.3	GMOS-S	r'	23.60 ± 0.26	0.8
56284.2	556.2	GMOS-S	r'	> 22.84	0.9
56536.0	808.0	FORS2	z'	> 22.10	1.6
56538.0	810.0	FORS2	I	23.28 ± 0.25	1.5
56538.0	810.0	FORS2	z'	23.29 ± 0.29	1.4
56660.3	932.3	GMOS-S	i'	> 22.37	1.4

Table 3.1: Swift J1112-8238 optical photometry. Limits are stated to 3σ . Photometry is presented without host subtraction, although it is likely that the late epochs represent the host; that is, not significantly contaminated by transient light. Note the i' GMOS-S magnitudes were calculated using relative photometry from the VLT I -band image and so have a minor systematic uncertainty not included here. All observation times are measured from the beginning of the first day of the 4 day *Swift* trigger observation (2011 June 16 UT 00:01). The seeing of each observation is included as it affects the contamination from the nearby bright star

the r' band was calibrated with reference to the Gemini standard zeropoints⁶.

The resultant photometry is detailed in Table 3.1 and plotted in Figure 3.5. The early time observations showed a point-like source while later observations (>1 yr) reveal emission with a flux a factor ~ 2 lower than recorded at early times. Modelled photometry of the late time emission with a Sérsic profile using GALFIT (Peng et al., 2002, 2010) was consistent with the aperture photometry detailed above, while PSF-matched point-source photometry (completed by scaling a PSF built from the image) yields results a magnitude dimmer, indicating the late time source is extended.

The photometry has been corrected for Galactic extinction, with $E(B-V) = 0.253 \pm 0.009$, based on values derived from Schlafly and Finkbeiner (2011) and accessed via the NASA/IPAC Infrared Science Archive⁷. The individual bandpass corrections were approximated from the corresponding SDSS filter corrections and thus have a minor systematic uncertainty not included in Table 3.1.

⁶<http://www.gemini.edu/sciops/instruments/gmos/calibration/photometric-stds>

⁷<http://irsa.ipac.caltech.edu/applications/DUST/>




Figure 3.5: The lightcurves for Swift J1112-8238 in various wavebands from a few hours to ~ 1000 days post-trigger. The time axis is displayed in both seconds and days in the observer frame. In all panels, the right hand axis indicates an isotropic equivalent luminosity or equivalent optical absolute magnitude. (Top) The median subtracted *Swift*-BAT daily average lightcurve in the 15-50 keV range, cut at 10^7 seconds post trigger for clarity. Note that the vertical scale is linear, and there are no significant detections beyond the first few days. (Middle) The *Swift*-XRT lightcurve in the 0.2-10 keV range. The black line indicates a $t^{-1.1}$ fit on the data preceding the sharp decline at ~ 30 days post trigger. In addition, the X-ray luminosities for Swift J1644+57 and Swift J2058+05 have been plotted. To allow for a direct comparison between the lightcurves, Swift J1644+57's and Swift J2058+05's lightcurves have had a cosmological time dilation correction to place them as though they had occurred at the same redshift as Swift J1112-8238. (Bottom) The optical lightcurves from GMOS-S and FORS2 photometry. There is considerable optical variability between the early and late time i'/I band magnitudes. The late time (>1.5 year) magnitudes are assumed to be at host level.

3.2.4 Spectroscopy

Optical longslit spectroscopy of Swift J1112.2-8238 was obtained on GMOS-S on 2012 December 16 and 23 using the R400_G5325 grating and independently on FORS2 using the 300I+11 grism on 2013 September 5. The GMOS-S spectra had a combined integration time of 2400 seconds (4×600) with spectral resolution of $\sim 7\text{\AA}$ and a spectral range of $3870 - 8170\text{\AA}$. The FORS2 spectrum also had an integration time of 2400 seconds (4×600) with spectral resolution of $\sim 12\text{\AA}$ and a spectral range of $5100-11000\text{\AA}$.

Following the basic outline of spectroscopic reduction outlined in Chapter 2, the GMOS spectroscopy was again reduced via the standard pipeline provided in IRAF. Namely, the master bias (GSBIAS), master dark (GSDARK) and master flats (GSFLAT) were produced with the standard tasks and the science frames reduced with GSREDUCE. Sky subtraction was completed with the task GSSKYSUB and wavelength calibration was completed with GSWAVELENGTH based on arc spectra taken contemporaneously with the observations. Finally the frame was rectified with GSTRANSFORM. One dimensional extraction using the normal trace-following methods was impossible due to the lack of any significant continuum emission.

Similarly, the standard reduction method was used for the FORS2 images using the package ESOREX. This involved the production of a master bias frame (FORS_BIAS) before the task FORS_CALIB was run to produce the wavelength calibration and master flat. Finally FORS_SCIENCE was run to apply the calibrations to the science frames. The resulting images were combined with median averaging using the IRAF task IMCOMBINE as before with offsets determined from the telescope due to the lack of a clear trace. As with the GMOS images, the lack of any significant continuum emission precluded the extraction of a one-dimensional spectrum through modelling the trace.

In all spectra, a single, weak emission feature was observed at $\sim 7045\text{\AA}$ (Figure 3.6), with a significance of $\sim 10\sigma$ in the GMOS spectrum. No continuum flux, or additional emission lines were seen. The line does not lie at the position of any common zero redshift features. It is offset by $\sim 600 \text{ km s}^{-1}$ from the He 7060\AA line that is sometimes seen in accreting binaries (Marsh et al., 1991). However, in these binaries the line is broad, and many other emission features are seen. In addition, the existence of an underlying extended source, interpreted as the host of the transient, greatly reduces the probability of a Galactic origin as nebulae are the only Galactic source likely to be resolvable, and these typically show multiple emission lines. This indicates that *Swift* J1112-8238 is not a Galactic source.

The non-detection of other lines proximate in wavelength disfavours the iden-




Figure 3.6: The (Top) GMOS and (Bottom) FORS2 spectra showing the clear emission line at a wavelength of $\sim 7045\text{\AA}$, interpreted as the OII emission line, though the low resolution and line signal to noise preclude the possibility of resolving its doublet nature. In each case the lines are unresolved in the dispersion direction with resolution FWHM of 7\AA and 12\AA for the GMOS and FORS2 spectroscopy respectively. The position of the removed sky lines are indicated as dashes between the two spectra. The scale is in units of Angstroms.

tification of this line as either [OIII]($\lambda 4959, 5007\text{\AA}$) or $H\beta$ at $z \sim 0.4$, since in either case the other lines would be expected to be observed. If the line were $H\alpha$ at $z = 0.07$, [NII] $\lambda 6584$ or $H\beta$ and [OIII] would likely be visible, since all lie within the spectral window covered by the GMOS observations. The expected $H\beta$ flux can be calculated directly under the assumption the observed line is $H\alpha$, and that the host galaxy extinction is minimal (though it should be noted that the requirement of an intrinsic absorbing material in the X-ray fit of $N_{\text{H}}(\text{int}) \sim 2.4 \times 10^{21} \text{ cm}^{-2}$ suggest a small extinction of $\sim A_V \sim 1 \text{ mag}$ if the hydrogen column density to optical extinction within the host follows the same scaling relation as within the Galactic environment; Güver and Özel, 2009). The combination of grating efficiency and Galactic reddening mean that $H\beta$ would likely be observed at a significance of $< 1.5\sigma$, and thus is not a strong argument against the observed lines identity as $H\alpha$. The [OIII] lines can frequently be substantially brighter than $H\beta$, and for a galaxy of metallicity $12 + \log(\text{O}/\text{H}) \sim 7.8$, consistent with the inferred absolute magnitude at $z = 0.07$ (rest frame $M_r \sim -14.8$, Sweet et al., 2014a,b), [OIII] ($\lambda 5007\text{\AA}$) would likely be a factor ~ 6 brighter than $H\beta$. Accounting for foreground extinction and grating efficiency as before, it is estimated that it would be visible at a significance of $\sim 8\sigma$, whereas no line is present at this location. Any emission at the location of [NII] ($\lambda 6584\text{\AA}$) would be well below the detection limit given this assumed metallicity. It is also noted that at $z = 0.07$ the absolute magnitude of the galaxy of $M_i > -15$ would be unusually faint. Given these combined constraints it is unlikely that the line's identity is $H\alpha$.

On balance, the most likely identification for the line is OII (λ 3727Å) at a redshift $z = 0.8901 \pm 0.0001$. In this case, the redward emission lines are beyond the range of the GMOS spectroscopy, and lie in bright sky lines in the FORS2 observations, precluding their detection. The low resolution of the spectra is unable to resolve the doublet nature of the observed line in this case. This interpretation is supported by the observed galaxy colours. After correction for foreground extinction they are relatively red in $r - i \sim 0.9 \pm 0.2$, and bluer $i - z \sim -0.5 \pm 0.3$ (based on the ~ 550 day Gemini i' band photometry). Although the errors are large, this is consistent with the presence of a Balmer break between the r - and i -bands, as might be expected for $z = 0.89$.

3.3 Discussion

The spectroscopic observations have confirmed the existence of a single emission line that is inconsistent with any zero redshift lines and consistent with the adopted redshift, $z = 0.8901 \pm 0.0001$. In this section I provide a short summary of the inferred rest-frame transient and host properties and then compare them to the properties of possible progenitors.

3.3.1 Physical properties

At the time of the first X-ray observations (10 days post trigger), the isotropic X-ray luminosity of the source was $\sim 6 \times 10^{46} \text{ erg s}^{-1}$. The source showed an approximate power law decay with time of $t^{-1.1}$. However this had considerable short-timescale variability superimposed upon it, with factor of 2 differences in flux on timescales of a few thousand seconds ($> 10^6$ s after the initial outburst). The X-ray spectrum was well fit by a power law spectrum with $\Gamma = 1.33$ but there was some evidence of spectral evolution with an apparent hardening at late times. While multi-component fits to the time series data involving a broken power law or a flare produce statistically better fits, this may simply be due to the intrinsic short-term variability of the source and the sparse sampling of the X-ray lightcurve, precluding the inference of more detailed information about the source.

By assuming the late-time optical epochs represent host level flux that is uncontaminated by transient light, this host contribution was subtracted from the earlier photometry. This was done through the use of the image subtraction software ISIS (see Chapter 2; Alard and Lupton, 1998). The images were aligned, convolved to the same seeing and then the late time (~ 1.5 years) image was subtracted from the early time (17 and 21 days) images in i' . The subtractions left a clear, point

source residual in each image with an inferred position that lay at $0.11'' \pm 0.12''$ and $0.22'' \pm 0.11''$ (1σ) from the centroid of the host galaxy. The host centroid position was determined using a Sérsic profile fit to the late-time image using the GALFIT software package (see Chapter 2), as shown in Figure 3.8. The error on the host centroid position is determined under the assumption of a Gaussian profile with FWHM based on the half light radius from the GALFIT Sérsic fit. The apparent asymmetry of the host means that this represents a lower limit on the true error in the centroid position. At the inferred redshift, our tightest constraint places the transient 0.85 ± 0.93 kpc from the centre of its host, for which the half light radius is ~ 6 kpc.

From the host subtracted image, the absolute magnitude of the flare was determined to be $M_{i'} = -20.0$, equating to a luminosity of $\sim 2 \times 10^{42} \text{ erg s}^{-1}$. The underlying host has a comparable absolute magnitude, $M_{i'} = -20.3$ (rest frame M_g at $z = 0.89$). Based on the luminosity function of galaxies from Gabasch et al. (2006) this places it somewhat below L_* at $z = 1$ (at a redshift of 0.89, the i' -band equates roughly to rest frame g' -band for which, in the redshift range $0.85 - 1.31$, the L_* magnitude is -21.7).

3.3.2 Comparison to other sources

GRBs

Swift detected gamma-ray bursts (GRBs) are typically detected on timescales much shorter than those for Swift J1112-8238. The majority arise from standard rate triggers, although a significant minority are longer-lived and trigger the detector via image triggers, sometimes on timescales of > 1000 s. However, even the ultra-long GRBs (Levan et al., 2014) that have durations of $\sim 10^4$ s are much shorter than Swift J1112-8238, whose several day long γ -ray emission would imply a duration (if defined as T_{90} as for GRBs) of closer to 10^6 s. Hence on the basis of the γ -ray properties alone, Swift J1112-8238 is a much closer analog with Swift J1644+57 and Swift J2058+05 than with any identified population of GRBs.

The X-ray properties are also apparently distinct, since the inferred isotropic X-ray luminosity lies an order of magnitude above GRBs at a similar epoch (see e.g. Nousek et al., 2006; Levan et al., 2014), and GRB afterglows at such late times seldom show such pronounced variability (likely due to the lack of engine activity). Despite the longevity of the gamma-ray emission in ULGRBs, their late time afterglows are generally consistent with, if not slightly fainter than those of normal long GRBs (Campana et al., 2011; Thöne et al., 2011; Evans et al., 2014), and so the

X-ray properties also would suggest a physically distinct system.

The optical properties of Swift J1112-8238 are rather less conclusive. The optical transient luminosity is comparable with the brightest end of the GRB afterglow distribution (e.g. Kann et al., 2011), although, given the X-ray brightness, the inferred X-ray to optical spectral slope is very flat ($\beta_{OX} \sim 0.14$). If the emission mechanisms were similar to GRBs this would identify the counterpart of Swift J1112-8238 as a dark burst, and would imply significant extinction (Fynbo et al., 2009; Perley et al., 2013), the correction for which would make the afterglow the brightest seen at similar epochs (e.g. Kann et al., 2011). Alternatively, one may ascribe rather different emission mechanisms to the counterpart to Swift J1112-8238, in which case little extinction may be needed. It is interesting to note in this regard that Swift J2058+0516 also has a very flat $\beta_{OX} \sim 0.11$, despite a strong UV-SED that implied little extinction (Cenko et al., 2012; Pasham et al., 2015). Thus, while in this case, without estimates of the internal extinction of the host and transient environment it is difficult to say for certain, this implies that the emission mechanism for GRBs and these rTDF candidates may be quite different and thus that the events do not have similar origins.

A further important diagnostic of extragalactic transients is their relative position within their hosts. However, the imaging analysed here is of too low resolution to be strongly diagnostic of the flare’s origin. In Chapter 4, I use high resolution imaging from the *Hubble Space Telescope* to investigate the location of the transient and compare its position to GRBs and other forms of extragalactic transient.

AGN

The apparent coincidence of the transient position and the host centroid makes an association with the central supermassive black hole of the galaxy plausible, and therefore possibly with ongoing AGN activity. No catalogued source is consistent with the position of Swift J1112-8238 in either the SUMSS 843GHz survey (60% complete down to 6mJy, 100% to 8mJy; Bock et al., 1999; Mauch et al., 2003) or the AT20G 20GHz survey (91% complete to 100mJy; Murphy et al., 2010). This places limits on the pre-flare underlying radio emission of the host to the $10^{32} - 10^{33} \text{ erg s}^{-1} \text{ Hz}^{-1}$ level, which is only capable of ruling out the most luminous BL Lac type objects (Marcha and Caccianiga, 2013). However, while the X-ray luminosity of the brightest blazar flares can reach the levels observed in Swift J1112-8238, this is generally accompanied by optical emission many magnitudes brighter than presented here, as seen in Figure 3.7. In addition, our late-time X-ray limit places constraints on any underlying activity to a limit of $L_X < 10^{44} \text{ erg s}^{-1}$, fainter than

the majority of quasars. For these reasons, it is unlikely the flare is associated with AGN activity.

Relativistic Tidal Disruption Flares

The association of the optical flare with the inferred location of the SMBH may indicate the discovery of a new tidal disruption flare. In order to determine if this is plausible, the mass of the black hole expected to occur within a galaxy of this size was estimated. Kauffmann et al. (2003b) measure mass to light ratios of galaxies for a given redshift and rest-frame $g-r$ colour. At a redshift of 0.89, this equates roughly to a $i-z$ band colour in the observer frame. Based on an $i-z$ colour of -0.5 the mass to light ratio is ~ 0.2 , and, coupled with the rest-frame g -band (observer frame i -band) absolute magnitude of -20.3 , implies a galaxy mass of $4 \times 10^8 M_\odot$. The stellar mass to black hole mass scaling relation of Bennert et al. (2011), produces an estimate for the SMBH mass of $\sim 2 \times 10^6 M_\odot$ (although there is considerable scatter in this relation and it is unclear whether the relation is applicable to such low masses). A similar result is obtained using the method from Haring and Rix (2004) of $\sim 3 \times 10^5 M_\odot$ (by assuming that the stellar mass estimate represents an upper limit on the bulge mass of the host). Both of these estimates are well within the $10^8 M_\odot$ limit for a Sun-like star to be disrupted by a SMBH and produce a visible TDF, making a TDF origin plausible.

From Figure 3.7, the optical absolute magnitude and X-ray luminosity of Swift J1112-8238 places it in a region of phase space that is devoid of any sources with the exception of the aforementioned relativistic TDF candidates Swift J1644+57 and Swift J2058+05. These candidates also match well with this flare in their late-time X-ray lightcurves as shown in Figure 3.5, particularly in the case of Swift J1644+57. The overall power law decay observed over the 30 days of *Swift*-XRT follow-up of Swift J1112-8238 is somewhat shallower than that of the other candidates with an index of ~ -1.1 . Swift J2058+05 had a much steeper decay at a similar epoch with an index of ~ -2.2 , while Swift J1644+57 had a late-time decay remarkably close to the $t^{-5/3}$ relation suggested to be a feature of TDF lightcurves (Rees, 1988; Phinney, 1989).

However, this decay index is somewhat sensitive to the choice of T_0 , which in this case is poorly defined, due in part to the unusual trigger method. Further, while often T_0 is taken to be the time at which the flare becomes observable, the true T_0 occurs some time earlier at the point of return of the most bound material which may precede visible emission by several days. In order to be consistent with a $t^{-5/3}$ decay, the “true” T_0 would have to have been 12^{+6}_{-4} days before the start

of the *Swift* detection image. This may not be unreasonable, since Swift J1644+57 was active at least 4 days prior to its first GRB trigger, and had a 3σ detection on a single day, 14 days earlier (Krimm and Barthelmy, 2011). Constraints on T_0 have been attempted in detailed models of previous flares (e.g. Guillochon et al., 2014), however the lack of comprehensive follow-up precludes that possibility in this case. Perhaps even more importantly, the short duration over which observations were made also makes it difficult to determine the behaviour of the lightcurve within the context of the longer term emission. Indeed, Swift J1644+57's lightcurve was relatively flat at a similar epoch. Calculations considering more detailed transport of material through the disc point to a more complex picture, in which the $t^{-5/3}$ decline is only present in certain bands and over a rather restricted range of time (Lodato and Rossi, 2011), while even more recent calculations suggest that the $t^{-2.2}$ decay seen in Swift J2058+0516 should be present in half of disruptions (Guillochon and Ramirez-Ruiz, 2013). These predictions show that the X-ray flux can plateau over a period of tens of days after the initial disruption meaning the shallow decay of Swift J1112-8238 cannot place strong constraints on its nature. However it should be noted that these simulations concern the disk emission, whereas, in relativistic TDFs, the X-ray emission is thought to be dominated by the jet. It is unclear if the assumption of a direct correlation between the jet and disc emission is reasonable.

Spectrally, the low number of counts recorded in Swift J1112-8238 restricts the information that can be extracted. However, the spectrum is well fit with a single, absorbed power-law with a relatively hard spectral index $\Gamma = 1.33 \pm 0.08$, (w-stat/dof = 574/586). This is somewhat harder than the late time power-law index in Swift J1644+57 ($\Gamma \sim 2$) or in Swift J2058+0516 ($\Gamma \sim 1.6$). One area in which previous rTDFs differ is in the apparent correlation between hardness and flux. Swift J1644+57 exhibits spectral softening as it fades (Levan et al., 2011), while Swift J2058+0516 appears to harden (Cenko et al., 2012). For Swift J1112-8238 there appears to be some evidence for a spectral hardening at late times, more akin to Swift J2058+05 than Swift J1644+57.

The rapid variability observed in the X-ray emission can place constraints on the nature of the emission region. While the variability is not as dramatic as that observed in Swift J1644+57, where factor of 100 changes in flux were observed on timescales of ~ 100 seconds, there is still evidence for factor of 2 variability on timescales of a few thousand seconds. Unfortunately the brightness of the source precludes timing at much higher resolution, and so light-travel time arguments would only place weak constraints on the size of the emitting region ($< 1 \times 10^{15}$ cm, or 100 R_S for a $\sim 10^7 M_\odot$ black hole). More compellingly, the gamma ray emission at

the time of the first XRT observations is close to the Eddington luminosity of a $10^9 M_\odot$ black hole, and an extrapolation to early times suggests it was brighter still. The expected black hole mass is a factor of several hundred small than this and it is unlikely that a black hole could accrete at such high super-Eddington rates, so while the constraints on beaming are weaker than for Swift J1644+57, this is still believed to be the most likely explanation for Swift J1112-8238.

The lack of more comprehensive optical follow-up precludes the building of an optical SED which would help distinguish between the thermal SEDs of previous TDFs, (e.g. ASASSN-14ae, Holoien et al. (2014); PS1-10jh, Gezari et al. (2012)), for which the peak absolute magnitudes are loosely consistent, and the differing, non-thermal emission mechanisms suggested in Burrows et al. (2011) and Bloom et al. (2011) for rTDF candidates. One of these models involves a blazar-analogue combination of inverse Compton emission at high frequencies (X-ray/ γ) with a second peak at low frequencies (optical etc.) from synchrotron emission. Alternatively the emission in different wavebands may come from spatially separate emission regions. In the lightcurves of Swift J1644+57, limits on optical/radio short-term variability set the emission apart from the rapidly varying high energy emission. Under the assumption of a spherical emitting region with a blackbody temperature of 10^5 K (10^4 K), the radius of the region emitting optical light in Swift J1112-8238 would be about $\sim 2 \times 10^{15}$ cm ($\sim 5 \times 10^{15}$ cm). This is approximately consistent with 10 (50) times the tidal radius of a Sun-like star around a $10^6 M_\odot$ black hole. This result is similar to those obtained from analysis of optical TDFs, perhaps unsurprisingly as the optical luminosity of Swift J1112-8238 is approximately equal to that seen in some other TDFs (e.g. van Velzen et al., 2011; Gezari et al., 2012). This may suggest a common mechanism for the optical emission from both relativistic and thermal TDFs.

It is also interesting to note Swift J1112-8238 shows a sharp decline in its X-ray flux at ~ 40 days post trigger. This may be indicative of dipping as seen in Swift J1644+57 (Levan et al., 2011; Bloom et al., 2011; Saxton et al., 2012a) at similar times, or perhaps of a longer term cessation of activity as identified in Swift J1644+57 at much later epochs of ~ 1.5 years (Sbarufatti et al., 2012; Levan and Tanvir, 2012; Berger et al., 2012; Levan et al., 2016) and similarly in Swift J2058+05 (Pasham et al., 2015). In any case, the final epoch of observations results in a limit which is significantly below the extrapolation of the early emission, requiring either a steepening of the decay or a rapid drop. This suggests broad similarities between the different events, although the sparse sampling of Swift J1112-8238, makes it difficult to rule out alternate interpretations.

With the previous rTDF candidates, a variable radio source with a measured Lorentz factor of ~ 2 or higher was detected (Zauderer et al., 2011; Cenko et al., 2012). In addition the inferred formation epoch from Swift J1644+57 implied a recently formed source, consistent with the start of the higher energy emission. This helped lead to the suggestion of a newly formed relativistic jet that accompanied the TDF. However, a Swift J1644+57-like radio lightcurve would be observable even several years after the flare (Zauderer et al., 2011; Berger et al., 2012; Zauderer et al., 2013, and subsequent VLA observations PI:Zauderer) and, in Chapter 4, I present results of a radio search and detection of Swift J1112-8238.

3.4 Implications

If Swift J1112-8238 is indeed a member of the same class of object as Swift J1644+57 and Swift J2058+0516 then it brings the total number of such events, as selected by the high energy emission, to three. Radio observations of thermal (non-relativistic) TDF candidates (e.g. Bower, 2011; Bower et al., 2013; van Velzen et al., 2013) have been used to attempt to determine the number of “off-axis” members and in the case of Bower et al. (2013), a few candidates may have been discovered. More recently, radio observations of ASASSN14li, have detected a low-level, non-relativistic jet implying jets of some level may be ubiquitous (Holoien and Stanek, 2016). However, it is clear that the detected population of high-energy jetted candidates is small. Nonetheless it is striking that these three outbursts were all discovered by *Swift* in the space of a 3 month window in 2011. At first sight it may be argued that the proximity, and consequent brightness, of Swift J1644+57 may have motivated the searches that led to the discoveries of the additional candidates. However, the lack of any further examples in the subsequent four years suggests that this is more likely a statistical fluke. It is possible to quantify this via an archival search of *Swift* GRBs and the BAT transient monitor (Krimm et al., 2013). Within the 6.5 years of data reported in Krimm et al. (2013) there are two events marked as TDFs (the previously identified bursts), while only a further three are marked as “unknown”. Two of these (Swift J1713.4-4219, IGR J17361-4441) lie close to the Galactic plane, and are most likely Galactic sources. This leaves only the source under discussion, Swift J1112-8238, as a candidate relativistic TDF. It is plausible, though, that some other sources within the catalogue have been misidentified. In particular, Swift J1644+57 was initially identified as a Galactic Fast X-ray Transient (Kennea et al., 2011). However, the population detected by the BAT transient monitor is necessarily small.




Figure 3.7: The X-ray luminosity and optical absolute magnitude plotted for a number of extragalactic transients, including AGN flares and GRBs at late times. Swift J1112-8238 is more X-ray luminous at $\sim 10^6$ seconds than GRBs at similar epochs and, while the brightest X-ray blazar flares can match it, Swift J1112-8238 is very optically underluminous in comparison. It instead occupies a region of the parameter space devoid of other sources except Swift J1644+57 and Swift J2058+05. At early times the luminosity of these flares exceeds $10^{48} \text{ erg s}^{-1}$, and is in excess of $10^{46} \text{ erg s}^{-1}$ at 10^6 s . This is more luminous than “classical” TDFs, which exhibit markedly lower X-ray luminosity ($10^{44} \text{ erg s}^{-1}$), although none of these have been observed close to peak. Adapted from Levan et al. (2011).




Figure 3.8: A zoomed in region ($\sim 5''$, North up, East left) around Swift J1112-8238, ~ 550 days after initial detection, smoothed via a 3 pixel Gaussian convolution for clarity. The positions of the optical transient centroid, as measured 17 (left) and 21 (right) days after trigger with the IRAF command IMEXAM, are displayed as red $1-\sigma$ error circles. Similarly, the host optical centroid (white) as measured via a Sérsic profile fit with GALFIT is plotted as a white $1-\sigma$ error circle. This error is a lower limit based on the assumption of a Gaussian profile with FWHM equal to twice the half-light radius of the Sérsic fit. The transient positions are thus coincident with the central position to 1σ and 2σ respectively. The consistency of these positions makes an association of the event with the SMBH in the galaxy plausible. It should be noted, though, that due to the low surface brightness and possible complex morphology of the host galaxy the host centroid is subject to substantial systematic uncertainty.

In total, therefore, it appears that at most a handful of such events have been recorded over *Swift*'s ~ 10 year lifetime. Similar to Cenko et al. (2012), it is possible to determine an implied analogous (i.e. similar *isotropic* luminosity) relativistic TDF rate based on the 3 events observed in ~ 10 years, using the volume bounded by the distance to Swift J2058+05, as the most distant yet observed at $z = 1.1853$ giving a comoving volume of 215 Gpc^3 , and assuming the local number density of $10^6\text{--}10^8 M_\odot$ SMBHs to be 10^{-2} Mpc^{-3} (Tundo et al., 2007). The resulting rate is found to be $\sim 3 \times 10^{-10}$ per galaxy per year, in stark contrast to the $\sim 10^{-5}$ inferred from thermal TDF detections (e.g. Donley et al., 2002; van Velzen and Farrar, 2014). Even if a significant fraction of the ULGRB population were related to similar phenomena this would be unlikely to constitute the majority of the factor of 3×10^4 required. To resolve this discrepancy likely requires a combination of tightly beamed high energy emission, such as that seen in GRBs, and that not all TDFs produce relativistic jets.

Recent late time radio surveys of thermal TDFs by Bower et al. (2013) suggest that up to $\sim 10\%$ of TDFs may have an associated relativistic jet, thus suggesting a intrinsic rate of rTDFs of $\sim 10^{-6}$ per galaxy per year. Given this, the required beaming angle for rTDF high energy emission would be of order 1° to match the observed rates with that suggested by Bower et al. (2013). At first sight this is not unreasonable, given that, for example, Swift J1644+57 produced an isotropic X-ray emission equivalent to the Eddington luminosity of a $10^{10} M_\odot$ black hole in a galaxy that is only expected to contain an SMBH of $\sim 10^6$ solar masses (Levan et al., 2011; Bloom et al., 2011). In this case beaming (relativistic and/or geometric) of a factor 10^4 , or highly super-Eddington accretion would seem to be necessary. However, radio observations of Swift J1644+57 point to a rather modest Lorentz factor, that would be unlikely to result in such strong collimation ($\Gamma \sim 2$; Zauderer et al., 2011), unless the radio and high-energy emission regions are spatially separate, each with their own Lorentz factors. This could suggest that observed rates of radio transients from rTDF candidates might be considerably higher than the observed rates detected via their high-energy emission (Zauderer et al., 2011). Indeed, based on the models of Mimica et al. (2015), the radio emission becomes isotropic within about a year of the initial outburst.

However, it is also possible that the survey of Bower et al. (2013) could be impacted by small number statistics and potential contaminants. One of the two detections made, RXJ1420.4+5334 has an uncertain host identification due to the large error in the X-ray flare position. The second, IC3599, may be an AGN (Grupe et al., 1995), and has recently exhibited repeated flares, either due to repeated partial

disruptions of the same star on an ~ 10 year orbit (Campana et al., 2015) or due to ongoing AGN activity (Grupe et al., 2015). Because of this, the suggested 10% jetted TDF fraction may be overestimated, which would explain the lack of detections in any of the other studies (e.g. Arcavi et al., 2014), thus further contributing to the apparent deficit of detected rTDFs. Clearly further observations of larger samples of sources across the electromagnetic spectrum are needed to resolve this question.

3.5 Summary

Here I summarise the key findings of this study:

- I have shown that Swift J1112-8238 likely does not originate from within the Galaxy given that it originates from an extended source that exhibits only a single emission line that is inconsistent with any likely zero-redshift features.
- The single emission line is likely the [OII](λ 3727Å) emission doublet, placing the host and transient at a probable redshift of 0.8901 ± 0.0001
- The inferred X-ray and optical luminosities and the X-ray/ γ -ray duration are all consistent with the properties of the previous rTDF candidates, making a shared origin plausible
- The nuclear position of the transient in a galaxy that hosts a relatively small SMBH, well below the 10^8 solar mass limit for Sun-like star disruption, makes a tidal disruption flare a plausible explanation for the outburst
- The continuing lack of additional candidate rTDFs in the *Swift* archive suggests that these flares are extremely rare, with an estimated beaming-corrected rate of less than 10% of their thermal cousins

The analysis of this event is continued in the next chapter with a view to improving the comparison between it and the previous rTDF candidates.

Chapter 4

Late Time Observations of Swift J1112.2-8238

4.1 Introduction

From my previous analysis, it is very likely the event Swift J1112-8238 has an extragalactic origin, being associated with an extended source that exhibits a single emission line in optical spectroscopy. The emission line is consistent with being the [OII]($\lambda 3727\text{\AA}$) emission doublet placing the host at a redshift of $z = 0.89$. If true, the broad properties of the flare are consistent with the previous relativistic tidal disruption flare candidates in luminosity, evolution and spectral energy distribution.

However this redshift could not be confirmed with the available data as the doublet nature of the line could not be resolved and no other emission lines were present. Further, key signals of the previous rTDF candidates had not yet been observed, namely the rising radio flare associated with the presumed relativistic jet and there was a need to strengthen the constraints on the position of the flare in its host, the current optical imaging being of too low resolution to resolve structure within the host. In this chapter I present further observations, including high resolution optical imaging from the *Hubble Space Telescope (HST)*, medium resolution spectroscopy from X-Shooter on the VLT, and radio observations with the Australian Telescope Compact Array, between them capable of answering many of the unresolved questions surrounding this source.

4.2 Observations

4.2.1 HST Imaging

Observations of the host of Swift J1112-8238 were obtained with Wide Field Camera 3 (WFC3, Dressel, 2016) on the *Hubble Space Telescope* on 2015 March 7 (MJD 57088), ~ 4 years after the initial trigger. Images were obtained in the F160W filter beginning at 19:39UT for an exposure time of 997s (4×249) and the F606W filter beginning at 20:21UT with an exposure time of 1568s (4×392).

For the optical observations, the target was placed near the lower-left corner of the CCD in order to reduce effects due to the inherent charge transfer efficiency (CTE) issues of the WFC3 UVIS chip and the images were CTE corrected via the method of Anderson and Bedin (2010). The position angle of the observations was chosen to ensure that diffraction spikes from a nearby bright ($R = 15.8$ mag) star would not interfere with the target. Each set of exposures was obtained with a sub-pixel dither pattern allowing the images to be redrizzled to half the native pixel scale (resulting in pixel scales of $0.065''/\text{pix}$ and $0.02''/\text{pix}$ in the IR and optical image respectively) and combined using the PYRAF routine ASTRODRIZZLE (Fruchter, 2010). A minor gradient in the background of the F606W imaging was removed through use of the source identification program SExtractor (Bertin and Arnouts, 1996).

The resulting images are shown in Figure 4.1 and Figure 4.2. In each, it is clear that the host has a highly complex, irregular morphology that is broadly split into two main components: one a compact, bulge-like component and the other a more diffuse, extended component which may constitute a disc. There is also evidence of a further extended component, a star-forming complex or “knot” on the south-western edge of the disc-like component that is most evident in the F606W imaging (see Figure 4.1). For this reason (and the findings of the analysis of the XShooter spectrum in subsection 4.2.2) we consider the photometry of each component separately as well as that of the whole system. From this point forwards, parts of the host’s morphology will be referred to separately as the bulge (the compact central component), the extended component (the bulk of the diffuse component) and the “knot” (an apparent extended star-forming complex on the south-western edge of the host).

Photometry was calibrated based on the standard *HST* zeropoints accessed via the Space Telescope Science Institute calibration pages¹ and was completed both via aperture photometry and through modelling with GALFIT where possible.

¹http://www.stsci.edu/hst/wfc3/phot_zp_lbn




Figure 4.1: *HST* WFC3 images of the host of Swift J1112-8238 in the wavebands (Upper) F606W and (Lower) F160W. The complex morphology hinted at in the lower resolution GMOS imaging (Chapter 3) is clear in the higher resolution *HST* imaging. The “host” as determined in Chapter 3 appears to be a combination of two separate components, a bulge-like feature and a more extended component. Also note the existence of a “knot” of emission, visible in both the F160W and F606W imaging, within the extended component on the south-west edge of the host complex. The complex morphology of the host may be indicative of a wide disk with a region of strong star formation, or perhaps an ongoing merger or interaction of two galaxies. The position of the transient as determined in the two early epoch GMOS images are indicated with 1σ error circles (red) while the centroid of the bulge component in each filter is indicated with a white cross. The minor North-South gradient and the feature to the West in the F160W image is caused by the bright star to the North of the visible field. The position angle of the X-Shooter slit is also indicated with a green arrow.




Figure 4.2: A composite image of the *HST* WFC3 images with F606W in blue and F160W in red. The image clearly shows the differences in the two components of the host complex, with a blue extended/knot component and a redder bulge.

In the F160W imaging, contamination from a nearby bright star, while small, was still detectable. This contamination was accounted for via two methods: with object and sky apertures placed at equal azimuthal distances from the star, thus including the star’s contribution in with the sky determination, and with the star subtracted through rotation of the original image by 90° centred on the star – the two methods produced statistically identical results. However, for the purposes of modelling the clean background produced by the rotation and subtraction method was preferable.

In this waveband, photometry of the entire system was completed in an aperture that covered the entire host complex (radius $1.6''$). In order to separate the bulge and extended components, modelling was completed using GALFIT (see Chapter 2; Peng et al., 2002, 2010). Attempts to model the bulge component as a point source (using a source from elsewhere in the image) left considerable residuals. An acceptable fit is produced by using 2 Sérsic profiles, one for each component. However, due to the low surface brightness and large extent of the extended component, coupled with the existence of a number of nearby features that had to be masked from the fit (nearby unrelated galaxy, diffraction spike from bright star etc.), the fit fails to determine accurate errors when required to fit its half-light radius. Instead, the fit is carried out with the half-light radius fixed to the best fit value ($1.3''$). The result ($\chi^2_{\text{red}}=1.073$) is a wide, very flat (Sérsic index $\ll 0.1$) extended disc, while the bulge component is much more compact (half-light radius of $0.22\pm0.04''$, Sérsic index of 2.1 ± 0.7). The centres of the two components are inconsistent with each other to approximately 10σ (distance of $0.18\pm0.02''$) which is consistent with the visible asymmetry of the disc-like component. It is also important to note that the knot of emission to the west of the host (indicated in Figure 4.1) was masked out to accommodate the fit. In order to determine its contribution to the host photometry, aperture photometry was used on the model subtracted image centred on the knot emission (radius $0.5''$).

In the F606W imaging, the morphology of the host is too complex to produce physically motivated models. Instead, photometry was completed with apertures fixed on the position of the bulge (radius $0.4''$) as determined in the F160W imaging. In addition, the photometry of the entire system was determined in two apertures, one that matched the entire host as determined in the F160W imaging (radius $1.6''$) and one that encompassed the visible emission in the F606W imaging (radius $1.1''$). Unlike in the F160W imaging, the extended component contributes minimally within the bulge aperture and thus we consider the components spatially separable with these apertures. In this way, the photometry of the extended component then

becomes simply the difference in flux between the total and bulge apertures. Finally, aperture photometry of the knot was determined through an aperture fixed to the position of the corresponding F160W aperture. However it should be noted that it is unclear what contribution, if any, comes instead from emission associated with the wide disc seen in the F160W imaging and as such the F606W knot emission is likely overestimated.

The resultant photometry is detailed in Table 4.1. Aperture photometry has had point source aperture corrections applied as a lower limit to the true aperture correction, though in practice the large aperture sizes when compared with the diffraction-limited seeing of *HST* make such corrections small. All photometry has been corrected for Galactic extinction with $E(B-V) = 0.253 \pm 0.009$ based on values derived from Schlafly and Finkbeiner (2011) and accessed via the NASA/IPAC Infrared Science Archive². The correction itself was determined at the effective wavelength of the filter using the Fitzpatrick extinction law (Fitzpatrick, 1999) with $R=3.07$, the values for which were accessed via the York Extinction Solver (YES, McCall, 2004).

4.2.2 X-Shooter Spectroscopy

Spectroscopy of the host of Swift J1112-8238 was taken using X-Shooter (Vernet et al., 2011) on the Very Large Telescope (VLT) on 2014 December 19 (MJD 57010) beginning at 04:44UT and on 2014 December 20 (MJD 57011) beginning at 04:43UT. In each observation the UVB and VIS arms were exposed for a total of 2720s (4×680) while the NIR arm exposure time was 2400s (4×600) due to inherent constraints in the X-Shooter observation parameters. The data was taken in NOD mode (ABBA pointing) with $5''$ offsets between the A and B nod positions in the spatial direction (i.e. along the slit).

Since 2012 August, the atmospheric dispersion correctors (ADCs) mounted on the UVB and VIS arms of X-Shooter have been offline due to malfunction, usually requiring observations to be oriented along the parallactic angle to reduce losses. Unfortunately, the presence of the aforementioned bright star due north of our source made this impossible and forced a position angle of $\sim 50^\circ$. Instead, wide slits ($1.6''$ and $1.5''$ in the VIS and UVB arms respectively) and maximised exposure times were used in the UVB and VIS arms to attempt to counter this effect. It is also worth noting that, given our previous identification of the emission line in GMOS and FORS2 spectroscopy, giving the host an inferred redshift of 0.89 as discussed in

²<http://irsa.ipac.caltech.edu/applications/DUST/>

Filter	Component	Aperture Magnitude	Modelled Magnitude
160W	Total	22.17 ± 0.07	22.11 ± 0.08
	Bulge	–	23.56 ± 0.13
	Extended	–	22.53 ± 0.04
	Knot	25.30 ± 0.29	–
606W	Total (1.1'')	23.61 ± 0.09	–
	Total (1.6'')	23.88 ± 0.21	–
	Bulge	25.06 ± 0.16	–
	Knot	24.77 ± 0.11	–
	Extended + Knot (1.1'')	23.94 ± 0.14	–
	Extended + Knot (1.6'')	24.33 ± 0.33	–

Table 4.1: *HST* photometry of the host of Swift J1112-8238. The component refers to the portion of the complex morphology being analysed, namely the total system, the position of the IR bulge and the extended remainder of the system. In the F606W imaging, the photometry for both the 1.1'' visible emission aperture and 1.6'' F160W matched aperture cases are shown. Where applicable, aperture and Galactic extinction corrections have been applied. The total photometry for the F160W modelled magnitude comes from the sum of the contributions of the Bulge and Extended modelled photometry and the Knot aperture photometry. Note that the total magnitude of the system is consistent in the modelled and aperture methods.

Chapter 3, standard strong star-formation lines were not expected to be observable in the UVB arm or the blue end of the VIS arm, where the effect is at its worst.

The data was reduced both via the standard REFLEX pipeline (Freudling et al., 2013) and independently via the method of Modigliani et al. (2010) by T. Krühler (priv. comm), producing consistent results³. Flux calibration was completed with respect to the spectrophotometric standard star Feige110. Telluric corrections were determined through use of the telluric star Hip058859, but were not found to be relevant in the wavelength range of interest.

In all 3 arms, there was insufficient flux to produce a detection of the continuum level emission in the unbinned data, though faint emission was discernible upon extreme binning of the VIS arm data. A strong, doublet emission feature, coincident with the unresolved line at $\sim 7045\text{\AA}$ from the GMOS and FORS2 spectra presented in Chapter 3, was clearly visible. Interestingly, the line is split both spatially and in velocity space into two clear components as can be seen in Figure 4.3, which is consistent with an asymmetry observed in the GMOS spectrum from Chapter 3. Other emission lines were also visible at wavelengths of $\sim 9465\text{\AA}$, $\sim 9190\text{\AA}$, and $\sim 8200\text{\AA}$ in the VIS arm and at $\sim 12405\text{\AA}$ in the NIR arm, all displaying the same two component nature except in the case of the $\sim 9465\text{\AA}$ line, where strong sky line emission masked the second component. No features were detected in the UVB arm. All of the detected features are consistent with the original interpretation of the single line discussed in Chapter 3 as being the [OII] emission doublet, coincident as they are with the relative positions of $H\gamma$, $H\beta$, [OIII] $\lambda 5007$ and $H\alpha$ respectively. The [OIII] $\lambda 4959$ line was covered by sky line emission for both spatially resolved components, while the [OIII] $\lambda 5007$ line was obscured for the southern component. Limits were also placed on the [NII] $\lambda 6583$ line emission.

One dimensional spectra of each of the two spatially resolved components were extracted and corrected to the heliocentric frame, and each line of sufficient significance was fitted with a gaussian profile. Based on the position angle of the slit, we determine that the two components (upper and lower) constitute separate detections of the bulge and extended/south-western knot components of the host system respectively. This results in measured redshifts of 0.8904 ± 0.0001 for the bulge and 0.8895 ± 0.0001 for the extended component. The line widths were resolved in velocity space and were deconvolved with the instrumental resolution to determine the velocity dispersion of the lines, found to be $105 \pm 15 \text{ km s}^{-1}$ and $120 \pm 10 \text{ km s}^{-1}$ for the bulge and extended components respectively.

³With the exception of potentially different user inputs (e.g. background/source aperture selection) and the automation of the REFLEX pipeline, these methods use the same reduction recipes




Figure 4.3: The X-Shooter spectrum centred on the feature at $\sim 7045\text{\AA}$ interpreted as the [OII] $\lambda\lambda 3726, 3729$ emission doublet. (Top Panel) The reduced 2D X-Shooter spectrum and (Second Panel) the same spectrum with a 3 pixel Gaussian smoothing. The doublet nature of the line is clearly visible, as is the existence of two spatially resolved components with a clear offset in velocity space. (Third Panel) The northern and (Fourth Panel) southern components extracted as 1D spectra with 1σ errorbars. The red line indicates the Gaussian line profiles fitted to the data. Brown bars indicate the fitted data while grey bars indicate data masked from the fit, often due to poor sky line subtraction. Black vertical dotted lines indicate the positions of the centres of the observed emission lines. (Bottom Panel) The relative intensity of the sky spectrum over the same wavelength range. This panel is intended to indicate the presence of emission features and give their relative strengths compared to features in Figures 4.3 through 4.7. It is not intended to give the absolute strengths of sky emission. Over this wavelength range the sky emission was minimal and did not greatly affect fitting.




Figure 4.4: The X-Shooter spectrum centred on the feature at $\sim 9465 \text{ \AA}$ interpreted as the [O III] $\lambda 5007$ emission line. The panels are as in Figure 4.3. Sky emission was sufficient to obscure the southern component of the source, making it impossible to determine the profile.




Figure 4.5: The X-Shooter spectrum centred on the feature at $\sim 12405 \text{ \AA}$ interpreted as the H α emission line. The panels are as in Figure 4.3. Sky line emission obscures a portion of the southern component of the source.




Figure 4.6: The X-Shooter spectrum centred on the feature at $\sim 9190\text{\AA}$ interpreted as the $\text{H}\beta$ emission line. The panels are as in Figure 4.3. Again, sky emission in this wavelength range was minimal.




Figure 4.7: The X-Shooter spectrum centred on the feature at $\sim 8205\text{Å}$ interpreted as the $\text{H}\gamma$ emission line. The panels are as in Figure 4.3. Sky emission in this wavelength range was minimal. The feature, while of low significance in both the northern and southern components, is nonetheless consistent with the expected position and morphology of $\text{H}\gamma$ based on the other detected emission features.

Line Identity	Bulge	Extended
[OII] λ 3727	0.76 ± 0.17	2.20 ± 0.19
[OII] λ 3729	1.50 ± 0.23	3.35 ± 0.19
H α	2.27 ± 0.24	3.68 ± 0.36
H β	0.59 ± 0.14	1.13 ± 0.14
H γ	0.32 ± 0.13	0.47 ± 0.12
[OIII] λ 5007	1.20 ± 0.18	*
[NII] λ 6583	< 0.72	< 1.02

Table 4.2: Inferred parameters of the host of Swift J1112-8238 determined from the XShooter spectrum. The properties are determined for each component separately. Fluxes are quoted in units of $10^{-17} \text{erg cm}^{-2} \text{s}^{-1}$, have been corrected for Galactic extinction and have no correction for slit losses. Limits on [NII] λ 6583 emission are given to 3σ . *The [OIII] λ 5007 extended component was contaminated by strong sky emission and was not recoverable

The line fluxes for each line were determined separately for each component and corrected for Galactic extinction. Due to the complex morphology of the host and inability to spatially resolve the continuum level of the separate components, no attempt was made at determining slit losses and thus the quoted fluxes are systematically underestimated by as much as a factor of a few. The line fluxes are presented in Table 4.2.

4.2.3 Late-time Radio Observations of the Host

Motivated by the possible identification of the outburst as a relativistic tidal disruption flare, radio continuum observations were made in an attempt to look for the radio flare observed in both of the previous candidates (Bloom et al., 2011; Cenko et al., 2012; Zauderer et al., 2013). While these observations were taken at very late times, the well-observed radio lightcurve of Swift J1644+57 indicates these events are capable of producing bright, long-term radio emission years after the initial event (Zauderer et al., 2013, and later observations with the VLA PI: Zauderer).

The initial observations were made using the Australia Telescope Compact Array (ATCA) on 2015 Jan 30 (~ 1300 days post trigger in the observer frame)⁴. The observations were taken simultaneously in two bands, each of bandwidth 2 GHz, centred at 5.5 and 9.0 GHz. The telescope was in its most elongated 6A configuration, with baselines between 5.938 and 0.337 km aligned East-West, and Earth rotation synthesis was used to improve coverage of the uv -plane. A total on-source integra-

⁴Observations associated with programme C3002, PI: Stanway

tion of 95 minutes was divided over hour angles spanning nearly 12 hours, allowing good reconstruction of the synthesised beam which had a $2.5'' \times 1.9''$ full-width at half-maximum at 5.5 GHz.

Secondary phase calibration was performed using regular observations of PKS 1057-797, and absolute flux and bandpass calibration were determined through observations of PKS 1934-638 (the standard calibrator for ATCA). The reduction of the data was achieved through the use of the standard software package MIRIAD (Sault et al., 1995). Each band comprised 2048 channels, each of 1 MHz bandwidth. Multi-frequency synthesis images were constructed using natural weighting and the full bandwidth between the flagged edges of each band. Radio-frequency interference (RFI) was flagged in the primary and secondary phase calibrators using the interactive flagging routine BLFLAG before being applied to the similarly flagged science data. Images were constructed with the function INVERT, and contaminating sources in the field removed with CLEAN. The result was fitted with the function IMFIT to determine the position, spatial extent and flux of the source, if present.

In both bands, a faint source was identified coincident with the coordinates of Swift J1112-8238 (RA=11:11:47.6 Dec=-82:38:44.44 in the 5.5GHz imaging with a positional uncertainty of $\sim 0.25''$; see Figure 4.8). Photometry was completed by fitting point sources to the emission (see Table 4.3). Attempts to fit an extended source to the 5.5 GHz data yields a flux estimate consistent within the point source estimate but with significant uncertainty on both the source size and resultant integrated flux, suggesting that the signal to noise ratio is insufficient to perform such an analysis.

Following the first detection, further observations were completed in an attempt to determine if the emission was variable. The observations were made with an identical instrumental configuration through 2016 May 11-16, approximately 1.3 years later in the observer frame⁵ with a total exposure time of 8.5 hours. As a comparison, observations were also completed on the second rTDF candidate Swift J2058+05 on 2016 May 14 with a total exposure time of 3.75 hours and using the secondary phase calibrator 2121+053. The data was reduced following the same procedure as before.

The determined fluxes are shown in Table 4.3. The observations of Swift J1112-8238 again revealed a point source in the 5.5 GHz band with a flux consistent with the previous epoch. However, the source had apparently declined in the 9.0 GHz band to a non-detection with a 3σ upper limit of $54\mu\text{Jy}$. Somewhat surprisingly given the greater distance to the event ($z = 1.12$ cf $z = 0.89$ for Swift J1112-8238),

⁵Observations associated with programme C3099, PI: Brown




Figure 4.8: The early epoch 5.5GHz radio contours overplotted on the F160W *HST* image. The solid contours represent 2, 3 and 4 times the sky RMS while the dashed contours represent 1σ fluctuations. The beam FWHM is indicated in the bottom right-hand corner. Note that, while astrometric ties between optical and radio images are inherently difficult and the radio point spread function is wide, the emission appears to be centred on the bulge component.

Source Name	Observation Date	MJD	5.5 GHz Flux / μ Jy	9.0 GHz Flux / μ Jy
Swift J1112-8238	2015 Jan 30	57052	76 ± 15	70 ± 29
	2016 May 11-16	57519-57524	70 ± 11	< 54
Swift J2058+05	2016 May 14	57522	225 ± 15	236 ± 13

Table 4.3: The observed radio flux from the ATCA observations. All photometry is based on fitting point sources to the images.

a source was well detected at the position of Swift J2058+05 in both bands. The northern position of the source and poor uv-coverage of the observation meant that the beam was strongly elongated along the north-south direction (FWHM $45''$). However, the lack of nearby sources in optical images makes confusion unlikely and thus it is likely this does constitute a detection associated with the host of Swift J2058+05.

4.3 Discussion

The mean of the redshifts of the two components as determined in the X-Shooter spectrum is adopted as the redshift of the host. This value, $z = 0.8900$, confirms the redshift determined in Chapter 3.

4.3.1 Host Morphology and Transient Position

In both bands of the *HST* imaging the host of Swift J1112-8238 shows clear evidence of a complex morphology with at least two main components. The first is a simple bulge-like structure visible in both the F160W and F606W imaging. In the F160W imaging, modelling shows the bulge is clearly extended with a half-light radius of $0.22\pm0.04''$. At this redshift, this constitutes a physical size of 1.7 ± 0.3 kpc. The Sérsic index of ~ 2.1 is also consistent with typical values for galactic bulges, placing it on the boundary between a classical bulge (or elliptical galaxy) and a disk-like bulge (pseudobulge; Fisher and Drory, 2008; Gadotti, 2009), though the large error on this value cannot distinguish between the two. The bulge component of the F606W imaging is clearly far more compact, but appears to be marginally more extended than point sources in the field, with FWHM $0.11''$ compared to $0.08''$ for other point sources based on measurements using the IRAF (Tody, 1986, 1993) function IMEXAM. Attempting a point source subtraction at the centroid position of the bulge does show evidence of residual emission, however, the low significance

of the feature makes this uncertain. If true, and assuming gaussian profiles for both the point source function of the imaging and the bulge morphology, the half-light radius of the bulge in the F606W imaging is $\sim 0.04'' = 0.3$ kpc. If true, this indicates the F606W emission comes from a small region within the IR bulge component. The second component visible in the imaging is much more diffuse, with modelling of the F160W imaging suggesting the presence of a wide flat disc with half-light radius of $\sim 1.3'' = 10$ kpc. However, the extended emission component displays considerable inhomogeneity with the most obvious feature being the knot of emission on the south-west edge of the system, most-likely a star-forming complex.

This duality is also apparent in the X-Shooter spectrum of the host with two clear components separated in both the spatial and dispersion directions, where the offset at the inferred redshift in velocity space constitutes a $\sim 100 \text{ km s}^{-1}$ shift. To have two sources coincident to such a degree both spatially and in velocity space indicates they must be interacting. The possibility then exists that the two sources may be separate galaxies that are in the process of a tidal interaction or merger, which would also explain the irregular nature of the second component, the knot of emission constituting the bulk of the second galaxy. However we also cannot rule out the presence of a strong star-forming region embedded in a disc surrounding the bulge component since the relative velocities are also consistent with the expectations of galactic rotation. Visual inspection of the images (see Figure 4.2) shows that the bulge and knot components appear to have quite different colours, the bulge exhibiting substantially more infrared emission than the bluer knot. Aperture photometry of the region in each band confirms this, with F606W-F160W ($\sim V-H$) colours of 1.5 ± 0.2 and -0.5 ± 0.3 for the bulge and knot component respectively. This indicates very different stellar populations in the two regions. However, it should be noted that if the knot photometry is expanded to include the entire extended component, the colours are remarkably similar (1.5 ± 0.2 and 1.5 ± 0.15 for the bulge and extended/knot components respectively).

Based on the transient position determined in Chapter 3 and matching coordinate systems to the new *HST* imaging through the use of the IRAF function GEOMAP, the position of the transient can be determined with respect to this newly discovered complex morphology. The results are shown in Figure 4.1 with the positions of the transient determined in the two early GMOS images represented as 1σ error circles. Note that the error is dominated by the poor seeing and signal to noise of the original transient images with the matching to the *HST* imaging contributing minimally to the final uncertainty. The transient is clearly associated with the bulge as opposed to the extended component. The transient position is

between 1σ and 2.5σ from the centroid position of the IR and optical bulge, depending on the GMOS observation and the HST band used. Combining the probability distributions of the two independently determined positions, the best position of the transient places it at an angular distance of $0.14 \pm 0.06''$ which corresponds to a projected physical distance of 1.1 ± 0.5 kpc equating to a 2.3σ correspondence with the determined centroid and thus is consistent with zero offset.

4.3.2 Internal Extinction

Emission line ratios and luminosities can provide useful diagnostics of the host galaxy properties. An estimate of the internal extinction of the host of Swift J1112-8238 can be determined through analysis of the Balmer decrements, the emission line ratios of $H\alpha/H\beta$ and $H\beta/H\gamma$. The theoretical values for these ratios based on Case B recombination of a gas with temperature of 10^4 K and electron density of 10^2 cm^{-3} , as is commonly used in the literature, are 2.86 and 2.14 respectively (Osterbrock and Ferland, 2006). Where the measured ratios are larger than this, dust extinction has reddened the emission and thus an estimate of the internal extinction can be made through correcting back to the theoretical values. However, the low signal to noise of the $H\gamma$ line in this case makes its use in this analysis limited and thus we base our results purely on the $H\alpha/H\beta$ ratio. We then convert the result to an inferred extinction for the line emission, $E(B-V)_{line}$ using the relation:

$$E(B - V)_{line} = \frac{2.5}{k(\lambda_{H\beta}) - k(\lambda_{H\alpha})} \log_{10} \left[\frac{(H\alpha/H\beta)_{obs}}{(H\alpha/H\beta)_{exp}} \right] \quad (4.1)$$

where $k(\lambda_{H\beta})$ and $k(\lambda_{H\alpha})$ are the reddening curves at the wavelengths of $H\alpha$ and $H\beta$ respectively (where here we assume the dust extinction law of Calzetti et al., 2000), and $(H\alpha/H\beta)_{obs}$ and $(H\alpha/H\beta)_{exp}$ are the observed and expected theoretical values of the Balmer decrement respectively (Domínguez et al., 2013).

From this, the bulge and extended components have $E(B-V)_{line}$ of 0.266 ± 0.175 and 0.119 ± 0.121 respectively. This correction is applied throughout the remainder of our analysis, but note that the extinction for continuum fluxes, $E(B-V)_{cont}$, is related to that derived from line fluxes, $E(B-V)_{line}$, by the relation of Calzetti et al. (2000):

$$E(B - V)_{cont} = (0.44 \pm 0.03) E(B - V)_{line} \quad (4.2)$$

However it should be noted that the large uncertainties on these $E(B-V)$ values means the uncertainties on the internal extinction corrections can dominate the errors on the absolute magnitudes of the hosts. This is particularly true in the F606W

filter where the rest frame wavelength is in the ultraviolet and the uncertainties become close to a factor of 2 in luminosity.

4.3.3 Metallicity and Classification of the Host

The procedures of Kewley and Ellison (2008) were used to determine the metallicity of the host. The four calibrations used were those of McGaugh (1991), Kobulnicky and Kewley (2004), Zaritsky et al. (1994), all of which are based on theoretical photoionization models of HII regions and are valid across a wide range of metallicities, and the calibration of Pilyugin (2001), which is based on an empirical electron temperature (T_e) method. While Kewley and Ellison (2008) has determined a calibration that include both a high metallicity ($12 + \log_{10}(\text{O}/\text{H}) > 8.25$) and a lower metallicity ($12 + \log_{10}(\text{O}/\text{H}) < 8.05$) branch for the Pilyugin (2001) method, the lack of galaxies over a metallicity of 8.5 that were used to fit these calibrations brings into question its validity at high metallicity. Indeed, T_e methods are known to underestimate high metallicities (e.g. Brown et al., 2016a).

Each calibration was applied to the Galactic and internal extinction corrected line fluxes of the bulge component of the host. The lack of an $[\text{OIII}]\lambda 5007$ detection due to coincident strong sky line emission precludes a similar analysis of the extended component. Due to the relatively low significances of many of the lines, the errors on the metallicity calibrations are sizeable. However, with only one exception, the calibrations are consistent with a metallicity of $12 + \log_{10}(\text{O}/\text{H}) \sim 8.5 \pm 0.2$, a metallicity that is approximately solar (8.69, Asplund et al., 2009). The only exception is the Pilyugin (2001) calibration, which given the valid range noted earlier is not a cause for concern.

Based on the available line fluxes, a Baldwin Phillips & Terlevich (BPT) diagram (Baldwin et al., 1981) was constructed in order to distinguish between a star-forming and AGN dominated classification for the host. The resulting diagram is plotted in Figure 4.9 along with a representative sample of SDSS galaxies and the dividing lines between the two regions of the plot as determined in Kewley et al. (2001) and Kauffmann et al. (2003a). The delineation determined in Kewley et al. (2001) is based on the extremal case that is designed to include all starburst galaxies at the expense of including some AGN, as opposed to the more conservative delineation determined in Kauffmann et al. (2003a). Lacking the detection of the $[\text{NII}]$ emission line, it is only possible to place limits on the position of the bulge component in this phase space, while the obscuring sky lines at the position of $[\text{OIII}]\lambda 5007$ in the extended/knot spectrum makes the determination of constraining limits impossible. Nonetheless, while the position of the bulge in the BPT diagram




Figure 4.9: The Baldwin Phillips & Terlevich (BPT) diagram of the bulge component of the host of Swift J1112-8238 plotted as a 3σ upper limit on the $[\text{NII}]/\text{H}\alpha$ line ratio. A sample of SDSS DR8 galaxies (Brinchmann et al., 2004; Kauffmann et al., 2003b; Tremonti et al., 2004) are plotted in grayscale. The extremal (Kewley et al., 2001) and more conservative (Kauffmann et al., 2003a) delineations between the AGN and star-forming regions of the diagram are plotted as solid and dashed lines respectively. The upper limit places the host well within the extremal delineation and is strongly suggestive of belonging to the star-forming locus.

is outside of the conservative delineation, the direction of the upper limit makes it likely the bulge belongs to the star-forming locus and makes it unlikely there is strong AGN activity within the host.

4.3.4 Stellar Mass and Black Hole Mass

At a redshift of 0.89, the F160W filter represents a rest frame wavelength of approximately 8000Å. While this is far from the K-band, the luminosity of which is commonly used as a proxy for stellar mass, under the assumption of a flat SED it can provide a useful estimate. Based on the K-band absolute magnitude to stellar mass conversion from Savaglio et al. (2009) and using the internal extinction corrected F160W absolute magnitudes, we find the stellar mass of the bulge and knot components are $(1.6 \pm 1.1) \times 10^9 M_\odot$ and $(3.2 \pm 1.5) \times 10^9 M_\odot$ respectively, and thus the combination of the extended and knot components produces a value of $\sim 5 \times 10^9 M_\odot$. Each value has a further $\sim 50\%$ uncertainty from the Savaglio et al. (2009) calibration that is not included here. The sum of these components gives a result that is somewhat larger than the stellar mass estimate for the whole system made in Chapter 3, due in part to the addition of an internal extinction correction. However it is likely that the inferred mass from Chapter 3 also suffers from the large error on the $i' - z'$ colour used to calibrate the conversion and from the application of a method used to determine the mass of a single galaxy to an unusual, interacting system.

The vastly better resolution of the *HST* imaging compared with the ground-based imaging analysed in Chapter 3 has revealed the presence of a possible bulge-disk system. Thus it is possible to apply the bulge mass scaling relation of Haring and Rix (2004) to determine the likely black hole mass such a system would have. This results in a mass of $1.5 \pm 1.1 \times 10^6 M_\odot$, though with few low mass black holes studied in Haring and Rix (2004) it is unclear how accurate this value is. Alternatively, under the assumption that each component (bulge and extended/knot) represents a separate galaxy, this in turn leads to estimates of the central SMBH masses of $(8 \pm 6) \times 10^6 M_\odot$ and $(1.5 \pm 0.7) \times 10^7 M_\odot$ based on the scaling relation of Bennert et al. (2011). As each are well within the $10^8 M_\odot$ limit for the disruption of a Sun-like star (Rees, 1988), either make plausible objects for being the site of a tidal disruption event. However, given the lack of association of the transient with the extended component, and particularly of the knot which could constitute the core of the second galaxy, it is unlikely that the latter black hole played any part in this event tidal disruption flare. It is possible that the extended component's SMBH has already been captured by the bulge component, a process that takes on the order

of tens to hundreds of Myrs following a merger (e.g. Just et al., 2011; Khan et al., 2016). However given that the interaction is apparently still ongoing and thus has not progressed to the point where the cores of the two hosts have coalesced, this seems unlikely.

4.3.5 UV/Optical Star Formation Rate Estimates

Given the detection of radio emission coming from the host which could come directly from the transient, it is important to be able to rule out a star formation driven origin. One method for doing this is to compare the star formation estimates determined through a number of other sources with that coming from the radio luminosity. Within this subsection, I determine the star formation rates for each component of the host as inferred from a number of UV and optical calibrations. Throughout this subsection, while internal extinction corrections have been made, the quoted uncertainties do not include the associated systematic uncertainty. This uncertainty can be very large, as much as 100% at rest-frame ultraviolet wavelengths where the correction is largest. As such the star formation rates determined here should be considered to be only representative of the true values.

Given that it has a rest-frame wavelength ($\sim 3100\text{\AA}$) bluewards of the 4000\AA break, it is possible to use the F606W absolute magnitude as a measure of the SFR within the host. Moustakas et al. (2006) derive a U -band ($\sim 3600\text{\AA}$) conversion of $(1.8 \pm 1.0) \times 10^{-43} L(U) M_{\odot} \text{yr}^{-1} (\text{erg s}^{-1})^{-1}$. It is also possible to determine star formation rates from certain line luminosities. From the relation given in Murphy et al. (2011), the star formation rate as determined from $H\alpha$ emission is defined as $\text{SFR} = 5.37 \times 10^{-42} L_{H\alpha} M_{\odot} \text{yr}^{-1} (\text{erg s}^{-1})^{-1}$. Finally, the $[\text{OII}]$ line emission can also be used, though it is subject to much greater uncertainty than $H\alpha$ due to the effects of metallicity, dust extinction and ionization, with the $H\alpha$ to $[\text{OII}]$ ratio exhibiting up to a 0.4 dex uncertainty (Moustakas et al., 2006). The $[\text{OII}]$ emission line star formation rates are determined by the conversion factor of Kewley et al. (2004) ($\text{SFR} = 6.58 \pm 1.65 \times 10^{-42} L_{[\text{OII}]} M_{\odot} \text{yr}^{-1} (\text{erg s}^{-1})^{-1}$). The results from the different calibrations are reported in Table 4.4.

While there is some disagreement between the various methods, in part perhaps due to the large uncertainties on the extinction corrections, it is clear that the UV and optical star formation rate indicators suggest a star formation rate no higher than a few solar masses per year for the combined system.

Method	Bulge SFR $/M_{\odot} \text{ yr}^{-1}$	Extended SFR $/M_{\odot} \text{ yr}^{-1}$
F606W ^a	0.17 ± 0.11 (0.14)	0.21 ± 0.13 (0.15)
H α ^b	0.71 ± 0.08 (0.48)	0.68 ± 0.06 (0.33)
[OII] ^c	1.6 ± 0.4 (1.3)	1.6 ± 0.5 (1.8)

Table 4.4: The inferred star formation rate estimates for the components of the host of Swift J1112-8238. Each estimate includes the estimated random uncertainty (a combination of the uncertainty in the calibration and the estimated uncertainty in the continuum magnitude/line flux) and, in brackets, the systematic uncertainty that results from the internal extinction determination. The methods are as follows: ^a The U -band continuum conversion from Moustakas et al. (2006) using the F606W magnitude as an approximation; ^b the H α emission line conversion from Murphy et al. (2011); and ^c the [OII] emission line conversion from Kewley et al. (2004).

4.3.6 Properties of the Radio Emission

While there is evidence for evolution of the source between the two radio epochs, it is possible that the emission from the position of Swift J1112-8238 may be associated with star formation. Condon et al. (2002) presents a calibration based on the total radio emission as measured at 1.4GHz. Accounting for the radio spectral slope, the rest frame 1.4GHz radio luminosity implied by the first epoch of radio observations suggests a star formation rate of a few hundred solar masses per year, though this is associated with considerable error, in part due to the uncertain spectral slope. This implies a very high star formation rate, between 1-2 orders of magnitude higher than even the largest estimates from the UV and optical measures. While it is true that estimates of radio emission are less affected by internal extinction within the host that could be systematically reducing the estimates made in the UV and optical, the moderate internal extinction determined in this case makes this unlikely to be the cause of the discrepancy, despite the sizeable error. Further, the radio spectral index (α defined as $S_{\nu} \propto \nu^{\alpha}$, where S_{ν} is the flux per unit frequency, ν) in both epochs is higher than would be expected for typical star formation associated synchrotron emission⁶, which tends to be within the range $\alpha = -0.5 - -1$ (e.g. Condon, 1992;

⁶Synchrotron emission is a non-thermal emission process that occurs when ultrarelativistic electrons interact with magnetic fields. The resulting acceleration of the charged particle produces emission with a characteristic power law spectrum that rises with increasing frequency. In the ideal case of a homogeneous, optically thin source, this has a spectral index of $\alpha = 2.5$, though astrophysical sources, being . However, at high frequencies, above the so-called synchrotron self-absorption frequency, the frequency at which the source becomes optically thick, the spectrum turns over, resulting in a power law index $\alpha = -0.7$ in the ideal case (Condon and Ransom, 2016)

Thompson et al., 2006; Seymour et al., 2008). Coupled with the apparent evolution in flux and spectral index, it is clear that the radio emission is inconsistent with coming solely from star formation.

It is interesting to note that, while it is difficult to reliably tie the absolute position of a radio source to the relative position on an optical image due to the inherent lack of comparison sources in the field, and with the large PSF and low significance of the radio detections the position’s measurement error is of order $0.25''$, the radio emission appears to be coincident with the bulge component (Figure 4.8). This is further tentative evidence for an association with the transient flare, a finding that is also backed up by the possible evolution of the source. It is also possible it is due to unrelated AGN activity from the bulge’s central supermassive black hole. However, the narrow emission lines visible in the X-Shooter spectrum with line ratios that are most consistent with a star formation origin suggests this is not the case.

In Figure 4.10, the radio lightcurves of all three rTDF candidates have been transformed to rest frame time and the fluxes plotted as though they had all occurred at the distance of the most well-studied of the three, Swift J1644+57, in order to compare their properties. The late time emission of all three candidates are within a factor few of each other at the same epoch. Further, the evolution of Swift J1112-8238 is consistent with the shallow rest-frame 10 GHz and steep 20 GHz evolution of Swift J1644+57. However, the lightcurve of Swift J2058+05 appears to be quite different to the other two, with an apparent sharp decline in flux by the rest-frame 40 day observation. While this measurement was made through VLBI observations and it is therefore conceivable that it is subject to a systematic offset in the flux calibration, it seems that there is some diversity in the evolution of the radio lightcurves of these events.

The spectral evolution of the three candidates (where contemporaneous observations in both bands allow) may show some evidence of similarity (see Figure 4.11). Both Swift J1644+57 and Swift J2058+05 begin with a clear negative spectral index that evolves as expected towards synchrotron self-absorption at late times ($\alpha \sim -1$; Metzger et al., 2012). The observations of Swift J1112-8238 are also consistent with following this broad evolution. However, given the large uncertainties on the observations of Swift J1112-8238 and the paucity of observations of both Swift J2058+05 and Swift J1112-8238, it is not possible to determine whether the radio emission has evolved consistently across all three candidates throughout. It should be also noted that, as mentioned before, high frequency evolution of synchrotron spectra toward ~ -1 is expected in ideal, homogeneous sources and is therefore quite




Figure 4.10: The radio lightcurves of the three rTDF candidates plotted in rest-frame time and scaled to the redshift of Swift J1644+57, the best studied of the three rTDF candidates. The lightcurves are plotted for Swift J1644+57 (black), Swift J2058+05 (blue) and Swift J1112-8238 (red and orange). The lightcurves are plotted at frequencies of ~ 10 and ~ 20 GHz rest-frame compiled from Zauderer et al. (2011), Berger et al. (2012), Zauderer et al. (2013) and from a VLA observation (associated with VLA/14A-423 PI: Zauderer) for Swift J1644+57, while the Swift J2058+05 observations are from Cenko et al. (2012) and Pasham et al. (2015). The luminosities of the three flares are consistent to within a factor few at late times.




Figure 4.11: The spectral index evolution of the three rTDF candidates, Swift J1644+57 (black), Swift J2058+05 (blue) and Swift J1112-8238 (red). The data for Swift J1644+57 is compiled from contemporaneous observations from Zauderer et al. (2011), Berger et al. (2012) and Zauderer et al. (2013), while the Swift J2058+05 observations are from Cenko et al. (2012) and Pasham et al. (2015). The vertical dashed lines indicate the times when the peak frequency of the modelled emission passes through the observed bands (Metzger et al., 2012). Note that all three candidates are consistent with following the same spectral evolution, though the large uncertainties and poor temporal coverage preclude stronger inferences.

common in astrophysical sources (Condon and Ransom, 2016).

It is also possible to use the fundamental plane of black hole activity (Merloni et al., 2003), an apparent correlation between the radio luminosity, X-ray luminosity and supermassive black hole mass, in order to estimate the black hole mass with the available X-ray and radio data. Taking the late time *Swift* X-ray observation at ~ 1000 days post-burst of $\sim 10^{44} \text{ erg s}^{-1}$ (see subsection 3.2.2) as an upper limit on the host X-ray emission, and given the inferred 5 GHz (10 GHz rest frame) radio luminosity of $\sim 3 \times 10^{40} \text{ erg s}^{-1}$, the source does appear to be approximately consistent with the correlation seen in Merloni et al. (2003). However this would imply a rather large supermassive black hole mass of $\sim 1 \times 10^8 M_{\odot}$. However, given that the true X-ray luminosity at the time of the radio emission is unknown, and that the host is in an unusual flaring state with likely differing timescales for evolution of the radio and X-ray emission, it is not clear how valid this estimate is.

4.4 Implications for the interpretation of the flare

The acquisition of high resolution imaging of the host with the *HST* has enabled better matching of the transient position to the host morphology. However, the loose tie to the bulge centroid makes it difficult to draw strong conclusions. As such, it remains possible that the flare is not associated with the centre of its host and thus could imply an unusual core collapse event origin for the flare. With the potential merging or interacting system comes the possibility that molecular gas clouds in either galaxy may be subject to shocks and tidal effects that may trigger star formation (e.g. Bournaud, 2011). However, while a considerable fraction of merger triggered star formation occurs in the nuclear region of the merging galaxy (Keel et al., 1985), in this asymmetrical and ongoing case, such triggered star formation might instead be expected to occur on the boundary of the interacting system, as has been seen in a number of examples such as the Antenna Galaxy (Wang et al., 2004). Indeed, this could explain the star-forming complex (knot) visible on the south-western edge of the system in the F606W imaging. Alternatively, the inconsistent centres of the bulge and disc component could be explained in light of an inclined disc with respect to the line of sight, where the under-side of the bulge has been obscured by the disc. This would in turn lead to a strong preference towards blue emission being visible on the leading edge of the disc, helping to explain the strong asymmetry observed in both *HST* bands and suggesting star formation may be present throughout the galaxy, but is only visible in a relatively small region. However, the relatively low internal extinction determined from the X-Shooter spec-

trum argues against this interpretation. In any case, the knot of star formation is far from the determined position of the transient and cannot therefore be associated with the flare.

Nonetheless, whether triggered through a merger or simply a result of ongoing normal star formation, the possibility of an unusual core-collapse event remains. It should be noted that recent analysis of a possibly related class of event, the Ultra-long GRB (see e.g. Levan et al., 2014), has been apparently been confirmed to be associated with massive core collapse through the detection of a supernova following the initial burst of a single event (Greiner et al., 2015). However, as noted in Chapter 3, a number of differences exist between the properties of the ultra-long GRBs and the three rTDF candidates, namely a much shorter gamma-ray flare duration, as defined by the T_{90} measure used in other GRBs (10^4 seconds for the ultra-long GRBs compared with the 10^6 s for the Swift J1644+57-like events), and far less luminous late-time X-ray emission. Because of this, the two classes could be physically distinct. However, the detection of an optical/NIR rebrightening in the long-term evolution of Swift J1644+57 (Levan et al., 2016) with an absolute magnitude, color and duration consistent with that of a superluminous supernova, may suggest all of the detected extreme duration gamma-ray flares originate from massive star collapse. However, this rebrightening has a number of alternative possible explanations including the reverberation of X-ray emission similar to the effect used to map the central regions of AGN, based on a possible time lag between the X-ray and optical lightcurves (Yoon et al., 2015). It could also be explained in terms of a late peaking component of thermal or synchrotron emission associated with a tidal disruption flare (Levan et al., 2016).

As determined in Chapter 3, with an absolute magnitude of $M_V = -20.3$, Swift J1112-8238 comes close to reaching the typical range of superluminous supernova luminosities. However, the peak optical observations made ~ 20 days post trigger (~ 10 days rest-frame) place it on a somewhat shorter timescale than the tens of days rise times seen in SLSNe (e.g. Gal-Yam, 2012). Unfortunately, the sparsely sampled optical lightcurve of Swift J1112-8238 makes further inference of the presence of an underlying supernova impossible in this case.

It is, however, possible to compare the observed properties of the host with those of typical GRB and superluminous supernova host galaxies. GRBs have been shown to occur within a wide range of hosts, with stellar masses, star formation rates and metallicities that easily include the properties inferred here (e.g. Fruchter et al., 2006; Savaglio et al., 2009; Krühler et al., 2015). However, recent work on superluminous supernova hosts shows a preference for a rather narrower range

of properties, at least where the hydrogen-poor variant is concerned. The bulge component of the host is consistent with being amongst the more massive and more star-forming of the superluminous supernova hosts (Lunnan et al., 2014; Angus et al., 2016). Meanwhile, the near solar metallicity of the host disfavours its inclusion with hydrogen-poor SLSNe (Leloudas et al., 2015), though the much wider range of galaxies playing host to the hydrogen-rich variant, with metallicities up to 8.9, makes it difficult to rule out a superluminous supernova association based on host properties.

The possible interacting nature of the host of Swift J1112-8238 also opens up a number of avenues whereby the rates of TDF events are enhanced. In such a system, the position of the SMBH becomes non-trivial and, in a sufficiently disrupted system, the natural assumption of the coincidence of the host centroid and the SMBH may be flawed and thus could potentially explain an off-centre flare. As a result, in such a system, the position of the transient may not be as strong a diagnostic of the flare’s origin as might be generally assumed. In addition, the possible existence of a second black hole that may be migrating in towards the dynamical centre raises the rate of TDFs. The perturbations of the stellar orbits caused by ongoing mergers can increase the rate of tidal disruption flares by as much as two orders of magnitude in both galaxies. This effect, while dependent on the mass of each galaxy, occurs when the perturber moves to within a few times the effective radius of host (Liu et al., 2009). Theoretical modelling of more evolved merging systems have shown that the presence of a SMBH binary could raise the rates of TDFs up to 1 yr^{-1} , though this effect occurs only within the short period preceding coalescence of the two (Chen et al., 2009).

In both of the previous rTDF candidates, a strong diagnostic of their nature was the detection of a rising radio flare with properties that indicated the production of a moderately relativistic jet (Zauderer et al., 2011; Berger et al., 2012; Cenko et al., 2012; Pasham et al., 2015). With the radio observations outlined in this work it becomes possible to search for a similar feature in the late-time evolution of Swift J1112-8238. With optical and UV indicators suggesting only moderate star formation rates while the inferred radio luminosity implies hundreds of solar masses per year, and with no indications of AGN activity in the optical spectrum of the host, the detection of radio emission that shows some signs of evolution may indicate just such a flare has been observed. However, the lack of long-term monitoring in two of the candidates, and the low significance of the detections of Swift J1112-8238, make more definite inferences impossible. It is important that any future candidates are followed up frequently and regularly in order to provide a useful

baseline for comparison with other events.

In any case, the detection of radio emission coming from the host that is likely consistent with nuclear activity, as opposed to being associated with star formation, is, though tentative, the first direct evidence of the presence of a supermassive black hole in the host. As such the TDF origin for the flare remains plausible, particularly with the likely SMBH mass hosted by such a galaxy falling well within the 10^8 solar mass limit for the disruption of a Sun-like star. In almost all regards, Swift J1112-8238 remains a close match to the two previous rTDF candidates, making a shared origin likely.

4.5 Summary

Here I summarise the key findings of this study:

- The X-Shooter spectroscopy has confirmed the previously determined redshift of the candidate as $z = 0.8900$
- Both the X-Shooter spectroscopy and *HST* imaging show the host has a complex morphology, consistent with an interaction or merger, with the transient associated with a bulge-like structure
- The transient is loosely associated with the bulge centroid, though strong inferences cannot be made based on the significance of the offset
- The detection of radio emission coming from the host of Swift J1112-8238 that shows signs of evolution indicates it may come from the transient flare. In addition, comparison of the radio lightcurves and spectral evolution of all three rTDF candidates suggest they may be consistent with being scaled versions of similarly evolving flares, indicating the transients share similar properties across the electromagnetic spectrum.

Based on these findings, it seems likely that the three rTDF candidates do indeed share a common origin, though the observations obtained here make it difficult to place strong constraints on the nature of that origin. Emphasis must therefore be placed on the obtaining of thorough follow-up of future candidates.

Chapter 5

A Study of Four Unusual Nuclear Transients

5.1 Introduction

The precision to which the relative position of a flare within its host can be determined is a strong function of the resolution of the observations used. Unfortunately, a trade off must be made between covering large areas of the sky in transient surveys and producing high resolution imaging. This means that when attempting to study rare nuclear transients, in particular tidal disruption flares, the contamination by non-nuclear supernovae can be of great concern. For example, Strubbe and Quataert (2011) estimate that, for ground-based surveys that can typically locate transients to an astrometric accuracy of $0.5''$, the supernova rate exceeds the tidal disruption flare rate by two orders of magnitude at a redshift of 0.1, with the effect growing rapidly worse at higher redshift. Only when closing in on diffraction-limited observations, such as with ground-based adaptive optics or space-based *Hubble Space Telescope* (*HST*) imaging do the rates become comparable (Strubbe and Quataert, 2011). Thus an important step in the analysis of nuclear transients is the obtaining of high resolution imaging that can both determine the position of the transient in its host and, as was demonstrated in the previous chapter, uncover interesting features in the host morphology that may be diagnostic of the flare’s origins. With this in mind, I present *HST* observations of four interesting nuclear transients. Each were optically detected and identified as coming from the nuclear regions of their galaxies. Each also was, with varying degrees of certainty, identified as a possible tidal disruption flare candidate. However each displays properties that, at least at the time of detection, made them atypical within that class.

The first two events, ASASSN14ae (Holoien et al., 2014) and ASASSN14li (Holoien et al., 2016) were both detected by the All Sky Automated Search for Supernovae (ASASSN) and each have been identified as tidal disruption flares. While many supernova surveys struggle to detect nuclear transients due to contrast issues with the bright central regions of their hosts, or intentionally avoid them to minimise contamination from inherently variable AGN, ASASSN has been unusually successful in detecting tidal disruption flare candidates (Holoien et al., 2016). This may be as a result of the large point spread function of the survey causing most objects to appear point-like, which may help make them avoid the bias against the central regions of galaxies. Interestingly, though, both of the candidates considered here have somewhat peculiar evolutions, with exponential declines producing superior fits to their bolometric lightcurves as opposed to the typically assumed power law decay. ASASSN14li is also possibly the first thermal TDF candidate to have been confirmed as having a weak radio jet (van Velzen et al., 2015; Alexander et al., 2016), barring the possible detections in the surveys of Bower (2011) and van Velzen et al. (2013). In these surveys, the presence of strong Swift J1644+57-like jets was ruled out in all but a handful of cases. These exceptions were themselves somewhat suspect, coming from possible active hosts or possibly unrelated galaxies that happened to lie within the large flare error box. However, the comparatively low luminosity and early peak of ASASSN14li’s radio emission meant that many of the limits on other thermal TDFs were insufficient to rule out a similar jet. This might then imply that jets are a ubiquitous feature of TDFs, perhaps with the rTDF candidates simply representing the strongest examples or a separate sub-class with similar yet distinct origins.

The other two events, CSS100217:102913+404220 (henceforth CSS100217 Drake et al., 2011) and ASASSN15lh (Dong et al., 2016), while each was discussed as a possible tidal disruption flare, were originally identified as superluminous supernovae, in fact making them the most luminous examples. For example, CSS100217 was excluded in part due to the presence of evolving line emission that had not been observed in TDF candidates before, while ASASSN15lh appeared to best fit the lightcurve models for superluminous supernovae. However, their extremely atypical host galaxies, CSS100217 coming from a known Seyfert 2 galaxy while ASASSN15lh was located within a unusually massive galaxy with minimal star formation, brought doubt on this interpretation. In the case of CSS100217 this was further compounded by the significant decrease in apparent quiescent level emission following the event, perhaps indicating a link between the flare and a decrease in the AGN emission from the host.

In this study, I utilise *HST* imaging in an effort to constrain the nuclearity of these events, a strong diagnostic of their nature, in essence determining the probability of the event originating from the supermassive black hole. I also analyse the available UV/optical/X-ray lightcurves, adding to them where possible with the new data reduced here, in order to determine their late-time evolution for comparison with the general properties of possible origin models. Throughout this analysis, the redshifts of the flares and their host galaxies as determined in previous work will be assumed. These redshifts are as follows: ASASSN14ae $z = 0.0436$ (Holoien et al., 2014), ASASSN14li $z = 0.0206$ (Holoien et al., 2016), CSS100217 $z = 0.147$ (Drake et al., 2011) and ASASSN15lh $z = 0.2326$ (Dong et al., 2016).

5.2 Observations

5.2.1 HST imaging

Observations of all four events were made with *HST* WFC3 in the UVIS channel. The observations of ASASSN14ae and ASASSN14li are new for this work¹ as are the observations of ASASSN15lh². The observations of CSS100217 were detailed in (Drake et al., 2011) and are reintroduced and reanalysed here. ASASSN14ae and ASASSN14li were observed in two epochs separated by approximately a year, while the ASASSN15lh epochs are separated by approximately 6 months. While ASASSN14ae was observed using the full array, the observations of ASASSN14li were made using the UVIS2-C1K1C-SUB sub-array and the observations of CSS100217 were made using the UVIS2-M1K1C-SUB sub-array. The early epoch of ASASSN15lh was also made with the UVIS2-M1K1C-SUB sub-array, while the late images used the full aperture. The main features of the observations are included in Table 5.1.

All observations were corrected for charge transfer efficiency by the method of (Anderson and Bedin, 2010). The images of ASASSN14ae and ASASSN14li were drizzled to native pixel scale ($0.04''/\text{pix}$) while the large number of individual exposures taken in each filter of CSS100217 (7 exposures per observation) and ASASSN15lh (20 and 6 exposures for the early and late observations respectively) enabled the redrizzling to smaller pixel scales of $0.02''/\text{pix}$ and $0.025''/\text{pix}$ respectively. All images were combined through use of the PYRAF routine ASTRODRIZZLE (Fruchter, 2010).

A sample of the final images are displayed in Figure 5.1, 5.2 and 5.3. Each presents a very different host morphology. ASASSN14ae’s host is an edge-on disk

¹Observations are associated with *HST* proposal 13026 PI:Levan

²Observations associated with *HST* proposal 14346 PI:Kochanek




Figure 5.1: The *HST* WFC3 images of (Left) ASASSN14ae in the F606W filter and (Right) ASASSN14li in the F621M filter. The images are tiled with the (Top) early epoch, (Middle) late epoch and (Bottom) subtraction of the two displayed. Both galaxies are well resolved with the host of ASASSN14ae being a clear edge on disc while the host of ASASSN14li is a somewhat simpler elliptical galaxy. In each case, the subtractions are clean leaving a clear point source consistent with coming from the very central regions of their hosts. The position of the host centroid is indicated in the lower panel with a green cross.




Figure 5.2: (Top) The *HST* WFC3 image of CSS100217 in the F763M filter. The emission is dominated by central point source emission, a combination of the transient and strong AGN emission. (Bottom) The image following the subtraction of a point source scaled to the peak of the emission. While the central regions are quite noisy, the outer regions show clear residual emission, evidence of the extended emission of the host.




Figure 5.3: The *HST* WFC3 images of ASASN15lh in the F606W filter. The images are tiled with the (Top) early epoch, (Middle) late epoch and (Bottom) subtraction of the two. The host is well resolved in both epochs and a point source is present in the subtracted image. However, the subtraction is not as clean as the attempts with the other ASASSN flares. This could be as a result of the somewhat poorer alignment between the images or due to the different roll angles of the two epochs. The position of the host centroid is indicated in the lower panel with a green cross

Object	Date	Observation Start	MJD	Filter	Exposure Time/s
ASASSN14ae	2014-06-13	06:28:47	56821.270	F606W	950
		06:51:27	56821.286	F275W	1020
	2015-05-15	06:23:22	57157.266	F606W	950
		06:51:27	57157.286	F275W	1020
ASASSN14li	2015-01-05	23:19:38	57027.972	F225W	600
		23:36:32	57027.984	F275W	560
		23:52:37	57027.995	F621M	300
	2015-12-26	05:04:35	57382.212	F225W	600
		05:25:44	57382.226	F275W	600
		05:42:29	57382.238	F621M	450
CSS100217	2010-05-31	07:18:04	55347.304	F390W	315
		07:34:33	55347.316	F555W	245
		07:49:43	55347.326	F763M	525
ASASSN15lh	2016-02-21	00:28:54	57439.020	F606W	1240
	2016-08-11	13:00:32	57611.542	F606W	2496

Table 5.1: The key features of the *HST* observations of the four events. Observation date and time are presented in UT.

galaxy and ASASSN14li comes from an apparently elliptical galaxy, while the only available epoch of imaging for CSS100217 shows a host dominated by a central point source with minimal extended emission. The host of ASASSN15lh also shows evidence of an edge-on disk structure, though its morphology is somewhat more difficult to determine (see below).

Photometry of each image was calibrated via the standard *HST* zeropoints accessed via the Space Telescope Science Institute calibration pages³. Galactic extinction was corrected for based on values of $E(B-V)$ derived from Schlafly and Finkbeiner (2011) and accessed via the NASA/IPAC Infrared Science Archive⁴. The resulting reddening values were $E(B-V) = 0.0155 \pm 0.0007$, 0.0217 ± 0.0013 , 0.0119 ± 0.0006 and 0.0304 ± 0.0011 for ASASSN14ae, ASASSN14li, CSS100217 and ASASSN15lh respectively. The corrections for each waveband were determined based on the Fitzpatrick (1999) reddening law.

In the cases of ASASSN14ae, ASASSN14li and ASASSN15lh, the large extents of the hosts in the F606W, F621M and F606W filters respectively made aperture photometry an impractical method for the determination of the photometry of the entire host. Instead, photometry of the host was completed through use of

³http://www.stsci.edu/hst/wfc3/phot_zp_lbn

⁴<http://irsa.ipac.caltech.edu/applications/DUST/>

SEXTRACTOR (using the elliptical aperture photometry of the MAG_AUTO option, Bertin and Arnouts, 1996). It should be noted, however, that the extremely low uncertainties reported by this method (~ 0.001 mag) are unlikely to be accurate. Millimag level precision photometry is extremely difficult to complete on extended sources. Systematic sources of uncertainty, such as variations in the background (which are difficult to determine accurately over such a large region of the CCD) and uncertainty in the low surface brightness wings of the galaxy, likely dominate the true uncertainty. This explains why the implied change in the source brightness between the early and late epochs of ASASSN14ae and ASASSN14li (in the F606W and F621M imaging respectively) is formally inconsistent with that measured in the subtraction image photometry. In an effort to estimate the true uncertainty in the total galaxy photometry, I first make the assumption that the subtraction image photometry is a more precise measure of the true difference between the early and late epochs. Scaling the total galaxy photometry uncertainties in the early and late epochs until the difference between the flux in the two epochs is consistent with the subtraction photometry to 1σ implies that the true uncertainty in the total galaxy photometry should be a factor ~ 5 higher (i.e. ~ 0.005 mag). This is still an extremely small uncertainty. Thus while using the early and late epoch photometry as an estimate for the change in transient emission is cautioned against, this small systematic effect is not enough to significantly affect the total photometry of the host for the rest of this analysis. Note that there is no similar inconsistency in the change in flux of the ASASSN15lh photometry compared with its subtraction photometry. This is in part due to the fainter host galaxy and the larger fractional change in flux, meaning other sources of error dominate over these small systematics.

In all other cases (ASASSN14ae F275W images, ASASSN14li F225W and F275W images and CSS100217 all bands) photometry was completed in large ($\sim 1''$) apertures to ensure the inclusion of all transient and extended host emission. All resultant photometry is detailed in Table 5.2.

Astrometric matching of the early and late F606W epochs of ASASSN14ae was completed through use of the IRAF function GEOMAP based on available point sources in the field (see Chapter 2). A total of 7 sources were used and, under the assumption of negligible difference in roll angle, an excellent match was achieved with a fit consisting of a simple shift in x and y coordinates (RMS of fit < 0.1 pixels, $< 0.004''$). The same offsets were assumed for the F275W epoch (a necessity due to the lack of bright UV sources in the field). All epochs were drizzled to the same position (determined by applying the transformation to the centroid of the transient in the early image with IRAF GEOXYTRAN) and the resulting images subtracted with

IRAF IMARITH. The subtraction produced clean, point source residuals in both bands and aperture photometry was completed as above and included in Table 5.2.

The small sub-array and sparsely populated field surrounding ASASSN14li precluded a similar method of astrometric matching. Only two objects were visible in the F621M imaging, an edge-on galaxy and a single bright point source. Again assuming minimal change in roll angle between the two epochs, attempts to use the positions of both (determined with IRAF IMEXAM and fitted with GEOMAP) to produce the transformation produced a poor fit ($\text{RMS} > 1$ pixel) and the subtractions produced strong negative residuals. It was determined that the point source was a nearby star that had undergone small but significant proper motion in the intervening time between the two epochs of observation. Instead, astrometric matching was made purely on the centroid of the edge-on galaxy. The uncertainty on the astrometric tie is naively determined based on the approximate centroid error on the galaxy, where the morphology of the central region of the galaxy used to determine the centroid position has been approximated as a gaussian with a corresponding centroid error. This results in an estimated uncertainty of 0.2 pixels ($0.008''$), though the true uncertainty may be somewhat larger. Again, the same offsets were assumed for the F275W and F225W images and each pair of images drizzled to the same position and subtracted as above. All three bands produced clean subtractions leaving only a point like transient in each upon which aperture photometry was completed as above and included in Table 5.2.

In the case of ASASSN15lh, astrometric matching was made using 6 sources in the field, limited by the small sub-array used in the early epoch imaging. The roll-angle of the telescope was different for the two epochs in this case. However, for the purposes of this study it is assumed that the error on the rotation determined through the WCS of each image is minimal, and thus the same method of astrometric matching is applicable upon rotation of the images to a common orientation. The fit produced a relatively poor RMS of ~ 0.2 pixels ($0.005''$) in part due to the necessity of using extended source centroids. The image subtraction revealed a point source though some additional structure was visible in the subtraction at the $\sim 1\%$ level, possibly a consequence of the poor alignment of the images or because the images were taken with different roll angles, thus making any asymmetries in the *HST* point spread function visible in the subtraction.

Modelling of the hosts of ASASSN14ae, ASASSN14li and ASASSN15lh was completed using GALFIT (see Chapter 2; Peng et al., 2002, 2010). With the exception of ASASSN15lh, the point source residual of the image subtraction was used to convolve the models as this provides an accurate measurement of the point source

Object	Date	Observation	AB Magnitude
ASASSN14ae	F606W	2014-06-13	16.784±0.001*
		2015-05-15	16.803±0.001*
		Subtract	20.858±0.013
	F275W	2014-06-13	20.061±0.021
		2015-05-15	20.884±0.077
		Subtract	20.754±0.056
ASASSN14li	F621M	2015-01-05	15.544±0.001*
		2015-12-26	15.676±0.001*
		Subtract	17.845±0.006
	F275W	2015-01-05	16.832±0.003
		2015-12-26	18.461±0.015
		Subtract	16.832±0.003
	F225W	2015-01-05	16.769±0.003
		2015-12-26	18.660±0.020
		Subtract	16.977±0.005
CSS100217	F763M	2010-05-31	16.099±0.001
	F555W	2010-05-31	16.711±0.001
	F390W	2010-05-31	17.081±0.001
ASASSN15lh	F606W	2016-02-21	18.166±0.007*
		2016-08-11	18.413±0.008*
		Subtract	19.837±0.060

Table 5.2: *HST* photometry of the hosts and transients of the four flares. All photometry has been corrected for Galactic extinction. The large aperture sizes and narrow point spread function (PSF) of *HST* make aperture corrections negligible. *Photometry determined through use of SEXTRACTOR instead of aperture photometry. The quoted uncertainties are those returned by the program and do not include the sources of systematic uncertainty discussed in the text, meaning they are likely slightly underestimated.

function at the position on the chip. In the case of ASASSN15lh, a superior fit was provided using a separate nearby point source in the field, possibly a consequence of the poorer subtraction.

The host of ASASSN14ae was best fit in both epochs by an edge-on disk (scale height $0.53''=455$ pc, scale length $1.15''=988$ pc) with a lower surface brightness wide halo-like feature (Sérsic index $n=1.24$, effective radius $R_{\text{eff}}=3.49''=3.00$ kpc). The central regions of the host included a compact bulge ($n = 1.55$, $R_{\text{eff}}=0.20'' = 172$ pc) and a further point source component was fitted for the transient in the early epoch and any remaining transient emission in the late epoch. Encouragingly, the decline in magnitude of the point source component (20.86 in the early and 21.89 in the late time epochs) is consistent with the subtraction determined decay of the transient and suggests that the transient has declined to less than a third of its flux in the early epoch. In both epochs, structured residuals were left near the central regions of the host with a morphology consistent with a dust ring surrounding the bulge, a feature which is difficult to model accurately. In any case, the residuals did not exceed more than 10% of the flux in any given pixel and so did not strongly alter the fit for the major components.

The host of ASASSN14li appears to be a simpler elliptical galaxy, well fit by two Sérsic profiles ($n=1.00$, $R_{\text{eff}}=1.98''=826$ pc and $n=1.98$, $R_{\text{eff}}=0.63''=263$ pc) and again a point source contribution was included for any transient emission in each epoch. Again the photometry of the point source (17.4 in the early and 18.4 in the late time epochs) agrees with the inferred decay of the transient and implies the transient emission has decayed by more than a factor two between the two epochs.

Modelling of the host of ASASSN15lh with GALFIT proved difficult. Fitting just a wide Sérsic profile and point source component produced residuals that indicated a truncated disk component. However attempts to also fit this component left strong asymmetrical residuals. This may be evidence of a strong dust lane or possibly indicate that the host is a post-merger galaxy with a complex morphology in its core. As a result, the host could not be fitted reliably.

The weak extended emission of the host of CSS100217, the emission being apparently dominated by a central point source, precludes accurate fitting of the host. In addition, the absence of other sources in the field, and no second epoch in order to produce a subtracted point source residual, makes the determination of the point source function for convolution difficult. Instead, suitably bright point source comparisons for each band were retrieved from other observations with similar instrumental configurations. Under the assumption that the central regions of the host are overwhelmingly dominated by point source emission, an assumption backed up

by the extremely similar radial profiles of the source and the point source comparisons in each band, the point source comparison was scaled to the peak determined in a Moffat profile fit of the source and subtracted. The results left noisy residuals but indicated the presence of extended emission in all three bands with similar spatial extent and magnitudes of $m_{F763M} \sim 19.5$, $m_{F555W} \sim 20.1$ and $m_{F275W} \sim 20.3$. While it is difficult to determine the systematic uncertainties on this method produced from subtle differences in the point spread function of the comparison and target images, this indicates the central point source emission, a combination of the transient flare and known AGN emission, constitutes $\sim 95\%$ of the observed source in all three bands, with only $\sim 5\%$ coming from extended emission.

5.2.2 CSS100217: Optical Observations

The transient CSS100217 was originally detected by the Catalina Real-time Transient Survey (CRTS, Drake et al., 2009) which covers a 33000 square degree region of the sky, performing unfiltered photometric monitoring of visible sources. The photometry is transformed to V-band magnitudes and is available as part of the Catalina Sky Survey (CSS).

The lightcurve for CSS100217 was retrieved from the archives⁵ and is plotted in Figure 5.4. The transient flare is clearly visible beginning at MJD ~ 55200 . However, most interestingly, when the flare decays, it appears to drop to a level ~ 0.4 mag fainter than before the flare began. In order to establish if this is due to a zeropoint offset in the local field, the long-term photometry of several other nearby objects of similar magnitude were also retrieved, one of which is plotted in Figure 5.4 for comparison. The source shows no signs of variability and does not display the same offset as seen in CSS100217, indicating it is an intrinsic property of CSS100217 and not a calibration error.

Further *R*-band observations were also retrieved from the Palomar Transient Factory (Law et al., 2009) lightcurve catalog⁶. These observations only covered two very short epochs at MJD ~ 56000 and MJD ~ 56750 and are therefore all after the flare has faded. They do however show that, assuming the *R*-band emission tracks the white-light of CRTS, the post-flare emission appears to have remained stable for a further 300 days after the end of the available CRTS lightcurve.

5.2.3 CSS100217: Swift Observations

Swift XRT (Burrows et al., 2005) observations of CSS100217 were completed both

⁵<http://nesssi.cacr.caltech.edu/DataRelease/>

⁶<http://irsa.ipac.caltech.edu/cgi-bin/Gator/nph-dd>

Filter	MJD	AB Magnitude
V	55292.586	16.10±0.05
	55311.695	16.25±0.05
	55326.009	16.27±0.05
	55339.192	16.53±0.06
B	55292.579	16.35±0.04
	55311.691	16.61±0.04
	55326.005	16.66±0.04
	55339.187	16.84±0.04
U	55292.578	16.73±0.04
	55311.690	16.90±0.04
	55326.004	17.01±0.04
	55339.187	17.05±0.04
UVW1	55292.576	17.48±0.03
	55311.689	17.66±0.04
	55326.003	17.83±0.04
	55339.184	17.86±0.04
	57107.631	18.91±0.24
	57109.968	19.13±0.11
	57110.962	19.17±0.07
	57114.684	19.09±0.18
	57116.580	19.09±0.10
	57118.144	19.00±0.10
	57123.566	19.15±0.12
UVM2	55292.590	17.77±0.03
	55311.696	18.00±0.04
	55326.012	18.14±0.04
	55339.196	18.24±0.04
	57107.630	19.58±0.34
	57109.965	19.38±0.16
	57110.960	19.31±0.11
	57114.683	19.40±0.24
	57116.579	19.50±0.17
	57118.143	19.35±0.14
	57123.565	19.68±0.20
UVW2	55292.582	18.00±0.03
	55311.693	18.22±0.03
	55326.007	18.32±0.04
	55339.190	18.42±0.03
	57107.629	19.29±0.21

Table 5.3: *Swift*-UVOT Photometry of CSS100217. Photometry is provided without host subtraction and is corrected for Galactic extinction.




Figure 5.4: The long term lightcurve of CSS100217. (Upper) *Swift* XRT observations during and after the flare. The late-time X-ray emission is clearly considerably brighter than the emission observed during the flare and also drops by a factor ~ 3 over ~ 10 days. (Middle) The optical/UV lightcurve including the long-term CRTS lightcurve (blue crosses), PTF lightcurve (green crosses) and *Swift* UVOT observations (squares): V (blue), B (green), U (red), UVW1 (yellow), UVM2 (cyan), UVW2 (magenta). The transient clearly decays to a fainter magnitude after the flare when compared with the level preceding it. A comparison source of similar magnitude is plotted in red, offset by $+2.5$ magnitudes. The lack of a similar apparent change in emission in this source indicates the decline is real and not caused by erroneous photometric calibration. (Lower Left) An enlarged portion of the optical lightcurve, covering the peak and decline of the flare. The ultraviolet bands decay along with, though more rapidly than, the optical flare, consistent with a cooling thermal SED. (Lower Right) An enlarged portion of the optical lightcurve covering the late time *Swift* observations. While the X-ray emission indicates the source is in an active and flaring state, the UV photometry shows no signs of variability and is consistent with the pre-flare *GALEX* photometry.

during the flare (as detailed in Drake et al., 2011) and ~ 5 years later. The early observations were made starting on 2010 April 6 12:57:00UT with several repeat visits lasting until 2010 May 23 with a total duration of 13.9 ks. During this time, a faint source was detected which, when combining all of the early observations, had an observed count rate of $(2.8 \pm 0.7) \times 10^{-3} \text{ ph s}^{-1}$ in the 0.3-10 keV band. The later observations began on 2015 March 26 15:01:59UT again with repeat visits over 15 days with a total exposure time of 4.7 ks. Interestingly, the source had clearly risen in X-ray intensity, peaking at ~ 10 times the count rate ($(3.4 \pm 0.4) \times 10^{-2} \text{ ph s}^{-1}$) it had during the optical flare. The reduced XRT lightcurve and fitted spectrum were obtained from tools available at the UK *Swift* Science Data Centre⁷, and were created using the techniques outlined in (Evans et al., 2007, 2009, 2010). The spectrum was based on the combination of all available observations (all made in PC-mode), making use of the known the redshift of the event ($z = 0.147$; Drake et al., 2011). The resulting spectrum is well fit by an absorbed power-law, with a fit statistic $W_{\text{stat}} = 65.77$ (67). The output fit parameters are photon index $\Gamma_{\text{ph}} = 3.2_{-0.4}^{+0.8}$ and N_{H} (int) $= 1_{-1}^{+11} \times 10^{20} \text{ cm}^{-2}$ while the Galactic value is $1.02 \times 10^{20} \text{ cm}^{-2}$ (Willingale et al., 2013), where all uncertainties are quoted as 90% confidence intervals. As such, the fit favours a fairly soft spectrum with large uncertainties on the intrinsic absorption (see Figure 5.5). This produces a counts-to-flux ratio of $2.55 \times 10^{-11} \text{ erg cm}^{-2} \text{ ph}^{-1}$ observed ($2.95 \times 10^{-11} \text{ erg cm}^{-2} \text{ ph}^{-1}$ unabsorbed) in the 0.3-10 keV band. This means the early observation had an unabsorbed flux of $(7.1 \pm 0.2) \times 10^{-14} \text{ erg cm}^{-2} \text{ s}^{-1}$ while the late time peak had an unabsorbed flux of $(8.7 \pm 1.0) \times 10^{-13} \text{ erg cm}^{-2} \text{ s}^{-1}$ (each in the 0.3-10 keV band, 1σ uncertainty). It was not possible to determine if the spectrum showed any signs of evolution in either a full spectral fit or a simple hardness ratio due to the low total count rate (particularly in the early epoch) and short exposure times (particularly in the late epoch).

Swift-UVOT observations were made concurrently with the XRT observations. Observations were taken in all 7 filters during 2010 and only the UV filters *UVW1*, *UVM2* and *UVW2*, during 2015. In order to obtain the photometry, the images taken around the same epoch were coadded, typically two to three exposures per observation. Source counts were initially extracted using a source region of $5''$ radius. As it is more accurate to use smaller source apertures when the count rate is low (Poole et al., 2007), source regions of $3''$ radius were used when the count rate dropped to below 0.5 counts per second. In order to be consistent with the UVOT calibration, these count rates were then corrected to $5''$ using the curve of growth contained in the calibration files. The background counts were extracted from a

⁷http://www.swift.ac.uk/user_objects




Figure 5.5: The *Swift*-XRT spectrum of CSS100217, produced by combining all available observations. The fitted model is indicated in red and the residuals presented in the lower panel. Figure produced with the Swift-XRT tools at the UK Swift Science Data Centre.

region $20''$ radius positioned in a blank portion of the sky close to the source. The count rates were obtained from the images using the *Swift* tool UVOTSOURCE. They were converted to magnitudes using the UVOT photometric zeropoints (Breeveld et al., 2011). The analysis pipeline used software HEADAS 6.17 and UVOT calibration 20150717. The UVOT data is provided in Table 5.3.

5.3 Discussion

5.3.1 Astrometry of the Transients

Based on the image subtractions of the three ASASSN candidate flares, it is possible to determine the positions of the flares relative to the centres of their hosts. The centroiding uncertainty on the point source residual in each subtracted image is determined under the assumption of a gaussian profile with signal to noise determined from its photometry and the FWHM of the point spread function.

The centroid of the host is determined in two ways. The first is based on the GALFIT models, using the position of the most compact extended component in each case, as it is assumed that any supermassive black hole would likely reside at the centre of this component. In the case of ASASSN14ae and ASASSN14li, this component is the compact bulge that dominates the central emission. In the case of ASASSN15lh, the difficulty in accurately modelling the host precluded this. However, by masking the entire central region of the host and fitting a simple Sérsic profile to the wings of the galaxy gives a first order estimate of the centroid position and is included in the analysis with due caution. The second method, which is less model dependent, makes use of the IRAF task IMEXAM to determine the centroid position based on a simple gaussian model of the central region of the host (8 pixel radius). This method works under the assumption that the remaining transient emission in the late-time epoch does not dominate the central region of the host, the case by case evidence for which is outlined below.

First it is assumed that all of the F275W/F225W emission in the images of ASASSN14ae and ASASSN14li is associated with the transient flare⁸ and that the flare has not undergone any considerable change in spectral energy distribution. Thus, the ratio of the late and early time F275W or F225W emission represents the fraction of transient emission that remains with respect to the early epoch. In the case of ASASSN14ae, the F275W emission has decayed to a flux of $16.1 \pm 1.1 \mu\text{Jy}$,

⁸Note that if this assumption is incorrect and host emission is present in either the F275W or F225W images of either candidate this would mean that the remaining transient emission constitutes a smaller fraction of the central region's late-time emission and would apply a smaller skew on the IMEXAM fit.

or $47 \pm 3\%$ of its early value. Assuming the thermal SED with a temperature of $\sim 20000 \pm 2000\text{K}$ determined in Holoien et al. (2014), this implies the remaining transient flux in the late F606W imaging is $8.7 \pm 2.6 \mu\text{Jy}$. Alternatively, assuming the F606W transient emission has decayed by the same fraction as the F275W transient emission, the remaining F606W transient would have a flux of $14.6 \pm 1.2 \mu\text{Jy}$. Note that the fact that these two values do not agree may indicate either that the spectral energy distribution of the flare has changed to become somewhat redder (i.e. cooler), or that the original fit did not accurately represent the SED of the flare. In any case, both implied transient emission fluxes constitutes less than 10% of the flux within the IRAF IMEXAM fitting region of the late time image (based on 8 pixel radius aperture photometry) and thus might not be expected to strongly skew the result. However, by adding fake point sources into the nuclear region of the late time image and determining the skew to the IMEXAM fitted region, it was found that the transient emission could produce sizeable skews to the fit, even when its total flux contributes less than 10% of the emission in the fitting region. As such the IMEXAM result for ASASSN14ae is presented with caution.

In the case of ASASSN14li, the F225W emission has decayed to a flux of $125 \pm 1 \mu\text{Jy}$, or $17.5 \pm 0.1\%$ of its original value. Assuming the admittedly poorly constrained thermal SED with a temperature of 35000K from Holoien et al. (2016), this implies the remaining F621M transient emission would have a flux of $\sim 35 \mu\text{Jy}$ or, by assuming again that the F621M transient has decayed by the same fraction as the F225W transient, the remaining F621M transient emission would have a flux of $\sim 56 \mu\text{Jy}$. In this case, the poor constraints on the SED of the flare mean it is not possible to determine if this difference is evidence of spectral evolution. Note also that both of these values are far below the remaining point source emission implied by the GALFIT modelling ($158 \mu\text{Jy}$). These estimates constitute less than 25% of the flux within the central region of the host ($229 \mu\text{Jy}$ within an 8 pixel radius). Again, modelling shows that this could produce a considerable skew in the IMEXAM fit and radial profile of the IMEXAM fit and thus again is considered with caution.

The radial profile of the central region of ASASSN15lh is extremely narrow, indicating point source emission may well be present in the late time image. Indeed, based on the model SED of the host produced by Leloudas et al. (2016), the V-band magnitude of the host emission is almost 0.5 mag fainter than that determined in the late-time F606W imaging, indicating the central region of the host may still be dominated by transient light. For completeness, however, the IMEXAM result is included in the astrometry of the host.

The determined offsets are detailed in Table 5.4. The IMEXAM offsets tend

Source	IMEXAM Offset ($''$)	IMEXAM Offset (pc)	GALFIT Offset($''$)	GALFIT Offset (pc)
ASASSN14ae	0.0034 ± 0.0028	2.9 ± 2.4	0.0130 ± 0.0030	11.2 ± 2.6
ASASSN14li	0.0148 ± 0.0080	6.2 ± 3.3	0.0260 ± 0.0080	10.8 ± 3.3
ASASSN15lh	0.0033 ± 0.0048	12 ± 18	0.0063 ± 0.0048	23 ± 18

Table 5.4: The angular distances and physical offsets of the two host centroid methods for the three ASASSN candidate flares. In the cases of ASASSN14ae and ASASSN14li, the transient is located within ~ 10 parsecs of the determined host centroid, making association with the supermassive black hole quite likely. The greater redshift of ASASSN15lh, and the consequently poorer physical resolution of the images of its host, make the astrometric ties somewhat poorer. Further the strong tie based on the IMEXAM position may be as a result of skewing from remaining transient emission. However, while it is unclear how applicable the GALFIT result is in the case of such a complex central morphology, the transient emission appears to be consistent with the host centroid.

to be somewhat smaller in all cases. This may indicate that the IMEXAM method is indeed skewed by transient emission, particularly in the case of ASASSN15lh. While ASASSN14ae and ASASSN14li are technically inconsistent with their host centroids as determined by GALFIT to the $3\text{--}4\sigma$ level, even those poorer matches are still within ~ 10 pc of the host centroid. This is an extremely close tie, especially when compared with their inner bulge effective radii of 170 pc and 260 pc respectively. ASASSN15lh, being at a somewhat higher redshift, has a consequently poorer physical resolution. However, based on the wider galaxy model the transient appears to be consistent with the host centroid to ~ 20 pc.

In the case of CSS100217, as in Drake et al. (2011), it was not possible to distinguish between the position of the transient and the underlying AGN point source emission in the HST imaging. In addition, while the CRTS lightcurve of the event implies that the transient constitutes approximately two thirds of the emission at the time of the *HST* imaging, the true fraction of the light coming from the transient is somewhat uncertain. As a conservative estimate on the minimum resolvable separation of the two, the WFC3 UVIS minimum separation of faint companion stars determined in Gilliland and Rajan (2011) is extrapolated to a magnitude difference of ~ 1 mag resulting in an estimated separation of $\sim 0.05''$. This is equivalent to ~ 100 pc at the redshift of CSS100217, comparable to the ~ 150 pc determined in Drake et al. (2011).

5.3.2 Evolution of the Transients

The *Swift* UVW1 and UVM2 filters have similar effective wavelengths to the *HST* F275W and F225W filters respectively, making possible the production of long-term lightcurves of each event. The resulting lightcurves for ASASSN14ae and ASASSN14li are plotted in Figure 5.6 and 5.7 respectively.

ASASSN14ae

As noted in Holoien et al. (2014), ASASSN14ae has a somewhat peculiar evolution compared to other tidal disruption flare candidates. Based on the bolometric luminosity of the flare determined from fits to the *Swift*-UVOT observed SED, they were able to determine that an exponential decline produced a better fit to the early time data than the normal power law models with indices between $-5/3$ and $-5/12$, as are commonly seen in other flares (Strubbe and Quataert, 2009; Lodato and Rossi, 2011). Using these new *HST* observations, and including further *Swift* data from Brown et al. (2016b), it is possible to extend the lightcurve to a potentially host dominated epoch and thus the F275W/UVW1 data was fitted with power law ($F = (t - T_0)^\alpha + \text{const}$) and exponential ($F = a \times e^{t/t_{\text{decay}}} + \text{const}$) models, each with an additional constant component to represent the underlying host emission. The power law fits were completed both with free fits of T_0 and α and by holding α to the canonical value of $-5/3$ varying only T_0 .

If permitted to vary freely, the power law fits prefer early values of T_0 (40+ days before the ASASSN trigger on 2014 January 25, MJD 56682.5) and steep power law indices ($\alpha \sim -3$). Unfortunately, pre-flare imaging can only place a weak limit on the beginning of the flare at a factor of ~ 2 in flux below peak 24 days before the flare’s detection (Holoien et al., 2014). Further, some leeway does exist in the “true” T_0 given that it represents the time of the return of the most bound material (Rees, 1988; Lodato and Rossi, 2011) and not necessarily the time of the first observable emission. As such it is difficult to exclude a slow rising transient. If on the other hand the flare was detected within a few days of T_0 , this simple power law fit would imply that the flare still dominates the ultraviolet emission, making the flat late time observations somewhat at odds with the fit. If α is held to the canonical value of $-5/3$, the best fit T_0 occurs ~ 17 days before the ASASSN trigger, consistent with the 18 days determined in Brown et al. (2016b) who instead fitted the *Swift* UVW2 photometry at early times. The exponential fit also shows a reasonable match to the data, particularly at late times, where the fit implies the flare has now reached host level.




Figure 5.6: The *Swift* UVW1 (crosses) and *HST* F275W (squares) lightcurve for ASASSN14ae. All photometry is presented without host subtraction. Also displayed are the canonical power law fit with a index of $-5/3$ (solid), upon which the value of T_0 is determined, and the best fit exponential model (dashed), each model with a constant factor included to represent host level emission.

However, none of the fits adequately reproduce the lightcurve in its entirety, with the power law fits producing reduced chi-squared values (χ^2_{red}) of between 3 and 9, depending on the value of T_0 used, and the exponential fit having a $\chi^2_{\text{red}} = 5.3$, implying that these simple fits do not fully represent the evolution of the flare. A more complicated fit with a broken power law improves the fit greatly ($\chi^2_{\text{red}} = 1.7$) with an early power law index of -0.9 ± 0.1 breaking to a much steeper decay of -2.5 ± 0.2 after ~ 45 days. This is qualitatively similar to the models of Lodato and Rossi (2011) that predict that tidal disruption flare lightcurves follow a shallow decay at early times and steepen at late times.

Alternatively, a possible cause for the poor fits could be systematic offsets in the photometric calibration of the photometry, or the possibility that the small offset in the central wavelengths of the filters used produces a larger than expected offset in the output photometry. Visual inspection of the lightcurve shows the F275W and UVW1 emission exhibits little sign of any systematic offset except perhaps at late times where the UVW1 observation at ~ 800 days from Brown et al. (2016b) is brighter (3.8σ) than the late F275W photometry at ~ 500 days. However, given the existence of a red leak in *Swift* UVW1 filter⁹, it is likely that this is due to contamination from optical host emission in the presence of a reddening SED as the blue transient fades. Given that the exponential model fit agrees well with the ~ 800 days UVW1 emission, whether or not it is included in the fitting process, and that the photometry is approximately consistent with the implied magnitude of the host fitting in Holoien et al. (2014), it is likely that the late time photometry represents host level emission.

ASASSN14li

As with ASASSN14ae, the lightcurve of ASASSN14li had previously been noted as being best fit by an unusual exponential fit (Holoien et al., 2016). Again, the combined *Swift*-UVOT observations from Holoien et al. (2016) and the *HST* observations from this study are fitted with the same power law and exponential models as ASASSN14ae, each with a constant factor to account for the underlying host emission.

As before, the free power law fits favour exceptionally early values of T_0 , hundreds of days prior to the trigger on 2014 November 22 (MJD 56983.5). Due to the host being behind the Sun for a considerable period of time preceding the flare, the limit on early emission of 2014 July 13 (MJD 56851, 132 days before the ASASSN trigger Holoien et al., 2016) is not as constraining as for ASASSN14ae. Fixing

⁹http://swift.gsfc.nasa.gov/analysis/uvot_digest/redleak.html




Figure 5.7: The *Swift* UVM2 and *HST* F225W lightcurve of ASASSN14li. Again, the photometry is presented without host subtraction. Also plotted are the $t^{-5/3}$ power law (solid) and exponential (dashed) fits, again with a constant component included to represent host level emission. The exponential model produces the superior fit to the data and shows that the flare has now reached host level.

the power law index to the canonical $-5/3$, however, produces a likely T_0 around MJD 56940, ~ 45 days before the ASASSN detection. This fit would also imply that the flare still contributes considerably to the remaining ultraviolet emission. However, an exponential model with $t_{\text{decay}} = 52$ days produces a far superior fit ($\chi^2_{\text{red}} = 2.3$ cf 5.9 for the canonical power law case). The same goodness of fit is only achieved in single power law fits with unfeasibly early values of T_0 , more than 300 days before the trigger, which are precluded by the limits on early detection. The constant component of the exponential fit also agrees well with the late epoch *HST* photometry, indicating the flare has now decayed to host level. More recent *Swift* UVOT observations that are not analysed in this work appears to support this, having plateaued at a magnitude of ~ 18.7 , consistent with the late-time F225W photometry¹⁰. Alternatively, a broken power law fit does produce a comparable fit to the data ($\chi^2_{\text{red}} = 2.3$) but indicates a very steep late time decay (~ 3.2).

CSS100217

A number of features of interest exist the long-term lightcurve of CSS100217 (Figure 5.4). The first is, as mentioned above, a clear offset in the pre- and post-flare photometry. The most logical explanation for this, particularly given the known Seyfert nature of the host, is a change in output of the central AGN, either through a change in the accretion rate or through a phase change in the emission mechanism. The temporal coincidence of the flare and this decline, along with the spatial coincidence of the flare and the centre of the host, makes it plausible that there is causal connection between the two events, perhaps with the flare resulting in the decline of the AGN or some mechanism causing both emission features.

The second point of interest is the apparent lack of correlation between the X-ray and optical emission. During the flare, while the UV and optical observations showed a clear linear decline in magnitude with time, the combined *Swift* X-ray observations were only able to determine the presence of a weak X-ray source with an inferred luminosity of $(4.1 \pm 1.2) \times 10^{42} \text{ erg s}^{-1}$, placing it at the low end of the X-ray luminosity function of AGN at low redshift, well below L_* at $1.5 \times 10^{44} \text{ erg s}^{-1}$ (Aird et al., 2015). However, the late *Swift* observations showed the source was going through a possible X-ray flare with bright, X-ray emission up to a factor of 10 more luminous than during the high optical state and showing factor ~ 3 variability over ~ 10 days. The *Swift*-UVOT observations obtained during the same epoch showed no evidence for evolution and were consistent with *GALEX* NUV (Bianchi et al., 2011) observations of the host made on 2004 January 24, 6 years before the flare

¹⁰Accessed via ASI Science Data Center at <http://swift.asdc.asi.it/>

($UVW2 = 19.4 \pm 0.2$ cf $NUV = 19.3 \pm 0.1$) and thus consistent with host level.

5.4 Implications for the origins of the flares

Each of these four flares exhibit properties that are somewhat atypical of their formally ascribed origins, from the low temperatures and exponential lightcurves of the ASASSN candidates to the decline in host level emission of CSS100217 and the extremely low star formation rate in the host ASASSN15lh. Here I consider the implications of these new observations on the mechanisms producing the three flares.

5.4.1 ASASSN14ae

The *HST* imaging has been able to further constrain the location of the flare to within ~ 10 parsecs of the centroid of the galaxy, the high resolution imaging and relatively local distance to the host making this one of the tightest constraints yet determined for a tidal disruption flare candidate. While the small uncertainties on the host and transient centroids make the two positions technically inconsistent to almost 3σ , the minimal offset between the two positions makes it entirely plausible that the transient remains associated with the central supermassive black hole.

With such an association comes the possibility that the flare is of AGN origin. Analysis of SDSS spectra of the host by Arcavi et al. (2014) and Holoiu et al. (2014) show some evidence of AGN activity in its $[OIII]/H\beta$ and $[NII]/H\alpha$ emission line ratios. However the blue *WISE* (Wright et al., 2010) W1-W2 colour of the host noted in Holoiu et al. (2014) disfavors strong AGN activity. It is interesting to note that, while the fitting of the F275W lightcurve implies that the flare has decayed to host level, the late F275W image remains consistent with being dominated by unresolved, point source emission. With an absolute magnitude of $M_{F275W} = -15.496 \pm 0.077$, the star formation rate implied by the final F275W epoch, under the assumption that the remaining emission is from stellar light, is $0.07 M_{\odot} \text{yr}^{-1}$, based on the scaling relation of Calzetti (2012). This is a factor of 2.5 higher than that determined from the SDSS spectrum (0.02 ± 0.01 Arcavi et al., 2014), and if applied instead to the final *Swift* epoch, would be almost a factor of two higher again. While determinations of star formation rates are inherently difficult and systematic errors between methods are common, this seems to imply that either the flare continues to contribute to the observed emission or that the remaining ultraviolet emission has a component coming from an active nucleus. In addition, while emission associated with star-formation would be expected to

come from an extended region, a point source subtraction completed on the late-time imaging shows the emission is consistent with being unresolved. Finally, it is possible that the apparent inconsistency of the late UVW1 and F275W photometry could be indicative of a change in the AGN output. This effect would be most noticeable in the ultraviolet where AGN variability is typically most pronounced (10-20% on timescales of months Ulrich et al., 1997). Given the low significance of the inconsistency and the known red-leak of the *Swift* UVW1 filter, it is more likely this represents a systematic calibration error and not an intrinsic property of the emission. In any case, for the original flare to be produced by AGN emission would require an extremely large outburst to have come from an otherwise relatively low luminosity AGN, an oddity even for the selection of events considered in this chapter.

In light of the extremely tight constraint on the nuclear position and only marginal evidence for a low luminosity AGN, the likely explanation for ASASSN14ae remains a tidal disruption flare. The slowly cooling thermal SED is most reminiscent of TDF behaviour and inconsistent with the rapidly cooling SED of SNe. Further, while the possible exponential decline of the flare is somewhat unusual for tidal disruption flares, the wide range of evolutionary paths determined to exist in theoretical modelling of these events does not currently exclude this situation (e.g. Lodato et al., 2009). As such, the popular explanation for the flare as a TDF remains the most likely origin for ASASSN14ae.

5.4.2 ASASSN14li

A similar situation exists in the case of ASASSN14li. The astrometry of the flare has tied its position to within ~ 11 parsecs of the centroid of the galaxy. While the nominal uncertainty places it just over 3σ from the determined centre of its host, the true uncertainty on this position is likely to be somewhat higher given the lack of sources available for astrometric matching in the field. Therefore it is likely that the source is consistent with the centre of its host and therefore an association with the central supermassive black hole remains plausible.

While the SDSS spectrum of the host of ASASSN14li shows little evidence of AGN activity, radio emission and [OIII] line emission coming from the host cannot be easily explained in terms of a star formation origin, making the galaxy likely host to a low-level AGN (Holoien et al., 2016). As in the case of ASASSN14ae, the radial profile and point source subtractions of the remaining F225W and F275W emission shows the majority of the remaining flux is consistent with coming from unresolved point emission. From fits to the broadband photometry, Holoien et al.

(2016) were able to place limits on the host’s star formation rate of $< 0.009 M_{\odot} \text{yr}^{-1}$ (1σ). As the lightcurve fitting suggests the emission has returned to host level, the remaining F225W absolute magnitude of $M_{\text{F225W}} = -16.080 \pm 0.020$, if assumed to come entirely from stellar emission, suggests a star formation rate of $0.11 M_{\odot} \text{yr}^{-1}$, based on the scaling relation of Calzetti (2012). This is a factor of 10 higher than the limit placed by the SED fitting, implying that either the flare still contributes to the late time emission or that the emission has considerable AGN contribution. However, given that the flare comes from at most a common, low-luminosity AGN, it is unlikely it could have caused the unusual bright flare.

With little constraint on the temperature of the flare it is difficult to determine if the transient has shown any spectral evolution, though the colour of the transient F621M-F275W = 1.013 ± 0.007 mag is quite similar to the ~ 1.25 expected for a blackbody temperature of $\sim 35000\text{K}$ as determined in previous analysis (Holoien and Stanek, 2016). In addition, while the lightcurve is again somewhat peculiar with an apparent exponential decay instead of the typical power law, the deviation is somewhat smaller in this case when compared to ASASSN14ae. With minimal spectral cooling and a strong nuclear tie, this flare remains a strong candidate for a tidal disruption event.

5.4.3 CSS100217

The origins of CSS100217 have been a mystery since its discovery. As with the other events considered here, the likely nuclear position of the flare suggests a possible connection to the central supermassive black hole. However the possibility of a supernova coming from a core collapse event cannot easily be ruled out. Indeed, in previous analysis, the outburst was considered most likely to be a member of the class of superluminous supernovae (Drake et al., 2011), based on its relatively low temperature (15000K), variable Balmer line emission and lightcurve shape. In addition the host, while determined to be a Seyfert galaxy based on certain emission line ratios, shows evidence of starburst properties in others (Drake et al., 2011). The flare itself is the second brightest event identified as a superluminous supernova ($M_V = -22.7$), the brightest being the more recent ASASSN15lh.

However, the observations noted here, along with analysis of other events and theoretical models, puts further doubt on this interpretation. While tidal disruption flares are generally expected to have particularly high temperature thermal SEDs of $\sim 10^5\text{K}$, possible cooler candidates do exist including PS1-11af (19000K Chornock et al., 2014), PS1-10jh ($2\text{-}5 \times 10^4\text{K}$ based on fits Gezari et al., 2012, though this may be a lower limit) and both ASASSN14ae and ASASSN14li as discussed in this

chapter. In addition, recently the existence of a continuum of events with hydrogen and helium emission lines in TDF candidates could explain the existence of the variable Balmer emission (Arcavi et al., 2014). While typical tidal disruption flares have much fainter optical emission, the relativistic tidal disruption flare candidates have all been accompanied by more luminous than normal optical emission ($M \sim -21$; Levan et al., 2011; Cenko et al., 2012; Brown et al., 2015), though it is not clear whether this emission is dominated by the same mechanism as in the thermal TDF candidates. Further, a possible interpretation of ASASSN15lh suggests such luminosities are possible in the presence of a rapidly spinning black hole, which increases the accretion efficiency compared with the static case (Leloudas et al., 2016).

Perhaps the most interesting feature of this new analysis is the apparent drop in emission between the pre- and post-flare host level. The drop of ~ 0.4 mag can only be logically explained in a phase change in the emission from the AGN. The simultaneity of this change and the outburst, while certainly possible to be a simple coincidence, may also indicate a direct connection between the two events. However, with any explosive event occurring outside of the accretion disk, such as an extremely nearby supernova, it is difficult to conceive of a situation that would produce a decrease in the accretion rate. It is plausible that strong emission from an AGN flare or TDF might have dispersed the accretion disk, in which case fallback from such an event might be expected to occur some time in the future, potentially causing a rebrightening of the AGN. However, without careful modelling, the true effect of either external supernovae or intense flares from the SMBH on the accretion disk is difficult to determine.

If CSS100217 were produced directly by the AGN, it would be extremely atypical for the generally stochastic variability observed in other examples (e.g. MacLeod et al., 2010). However it is possible that the host of CSS100217 represents an unusual AGN. The apparent lack of correlation between optical and X-ray emission is itself peculiar. Many AGN show strong correlations between their X-ray and optical variability, albeit with a commonly observed time lag of a few hours to days (e.g. Breedts et al., 2009, 2010), the correlation and lag both a consequence of the reprocessing of X-ray emission from the high temperature corona by the cooler accretion disk. In addition, based on the modelling of the three *HST* images, the central point source makes up an extremely large fraction of the total emission. Based on the CRTS lightcurve, the transient-host system is approximately one magnitude (or a factor 3) brighter than the post-flare host level. Thus, when the transient has faded, the extended stellar emission would contribute approximately 15% to the

remaining emission based on the point source subtraction in section 5.3.1. While this is not uncharacteristic for an AGN of luminosity $\sim 10^{44} \text{ erg s}^{-1}$ as exists here, it is nonetheless at the extreme end of the relation for Seyfert galaxies (Kotilainen et al., 1993). Further, the host of CSS100217 is unusually small for a galaxy hosting an AGN. Based on the point source subtracted F763M emission, corrected for the expected transient emission, and assuming a flat SED for the host, the F763M emission implies a host mass of $10^9 M_{\odot}$ based on the scaling relation of Savaglio et al. (2009). However, as noted by Kauffmann et al. (2003b) very few galaxies less than $10^{10} M_{\odot}$ host AGN. As such, it is conceivable that AGN from such a galaxy could give rise to unusual behaviour.

5.4.4 ASASSN15lh

As with all of the sources considered here, ASASSN15lh remains strongly associated with the central region of its host. While it is difficult to determine how accurate this method is in the case of such a complex central morphology, the GALFIT analysis suggests the transient and host centroids are offset by only $23 \pm 18 \text{ pc}$. This is consistent with the recent analysis of (Leloudas et al., 2016) which compares *HST* ACS images with a pre-flare image taken with the CTIO 4m+DECam, finding that the transient is offset by $131 \pm 192 \text{ pc}$ from the host centroid. The host of ASASSN15lh has been shown to be a massive red and dead galaxy with very little star-formation (Dong et al., 2016; Leloudas et al., 2016). Thus the combination of the low star formation rate and the nuclear position in the host, where massive stars are unlikely to exist, makes a SN explanation for the flare quite unlikely.

The nuclear position of the transient makes a TDF origin plausible. Indeed, the analysis of Leloudas et al. (2016) comes to the conclusion that the flare comes from a stellar disruption around a rapidly spinning black hole. In such a situation, the energy budget for the flare increases considerably due to a greatly increased radiative efficiency, explaining the extreme luminosity. Further, the rapid spin enables a star's disruption outside the Schwarzschild radius of a black hole that would otherwise be too massive, naturally explaining its presence in a host with an expected SMBH mass well above the $10^8 M_{\odot}$ limit (Leloudas et al., 2016). The temperature evolution of the flare is also qualitatively similar to ASASSN14ae, though the timescales involved are somewhat longer. Finally, this model naturally explains the ultraviolet rebrightening observed at late times, as the emission associated with shocks between streams of material during circularisation is expected to be visible somewhat before the reprocessed emission from the accretion disk.

It is also possible the transient may be associated with an AGN. The likely

existence of point source emission at the centre of the host could indicate AGN emission. If the point source emission is assumed to come from an active nucleus, it is conceivable that the flare is also associated with AGN emission in a situation similar to that suggested above for CSS100217. Indeed, the lightcurves of both flares show some striking similarities in terms of both peak luminosity and duration (see Figure 5.8). However, it should be noted that Leloudas et al. (2016) place limits on radio emission from the host during the transient of $\nu L_\nu < 10^{38} \text{ erg s}^{-1}$ which places it somewhat below the region where AGN hosts dominate the local radio luminosity function (Mauch and Sadler, 2007) and similarly that X-ray observations, also made during the flare with *XMM-Newton*, place limits on the X-ray luminosity of $< 3 \times 10^{40} \text{ erg s}^{-1}$, again at the low end of the X-ray luminosity function in the appropriate redshift range (Aird et al., 2015). As such, it is unlikely ASASSN15lh is associated with AGN emission.

5.5 A Possible New Class of AGN flare or a Sub-Class of TDF?

While the possibility of low-level AGN activity in both ASASSN14ae and ASASSN14li is interesting, the flare and host properties remain most consistent with a tidal disruption flare origin. Further, based on the analysis of ASASSN15lh in (Leloudas et al., 2016) and the continuing nuclear position of the flare, it seems most likely that ASASSN15lh is also associated with a tidal disruption flare. However, its massive host (and likely correspondingly massive SMBH) and extreme luminosity would likely require a rapidly spinning black hole for this interpretation to be viable. The peculiar asymmetric morphology of the centre of the host may be indicative of a post-merger host, possibly providing a further link with the tidal disruption flare interpretation given the recent discovery of a preference for E+A galaxies, a rare galaxy class that may have undergone a recent merger (Arcavi et al., 2014; French et al., 2016, 2017).

The change in the host emission in CSS100217 is evidence that the flare may have directly impacted the AGN emission. A change in the luminosity of AGN can occur in so-called changing-look AGN (e.g. Matt et al., 2003; Puccetti et al., 2007) where the observed X-ray emission can be seen to change on relatively short timescales of months to years, though these situations have for the most part been explained as absorption effects from intervening gas clouds along the line of sight. However a number of sources have also been shown to transition between Seyfert 1 and Seyfert 2 classifications, or vice versa, based on optical line emission. Such




Figure 5.8: The lightcurves of the 5 most luminous supernovae detected to date: iPTF13ajg (R-band, cyan squares; Vreeswijk et al., 2014), SN2008es (V-band, magenta triangles; Miller et al., 2009), CSS121015 (V-band, green stars; Benetti et al., 2014), CSS100217 (V-band, blue triangles; Drake et al., 2011) and ASASSN15lh (V-band, red circles; Dong et al., 2016). CSS100217 and ASASSN15lh peak somewhat above the other candidates and exhibit similar durations and late time decay rates while the other flares decline at a somewhat faster rate.

situations may be explained in a similar absorption case due to patchy coverage of the torus (Elitzur, 2012), or be due to changes in accretion rate (Elitzur et al., 2014). Perhaps most interestingly for this study, it has also been suggested that such changes could be triggered by short-lived transient events, such as tidal disruption flares. In the case of SDSS J015957.64+003310.5, for example, the AGN has evolved from Type 1 to Type 1.9 within the space of ten years, a change that accompanied a drop in AGN emission of a factor 6 (LaMassa et al., 2015), though it has also been suggested that this change is the result of long-term evolution of a tidal disruption event (Merloni et al., 2015). Importantly, this change in optical flux was accompanied by a corresponding drop in X-ray emission, a situation that does not appear to have occurred in the case of CSS100217 considering the apparent rise in X-ray emission at late times. Nonetheless, it is possible that the host of CSS100217 is undergoing a similar change that may become apparent in its spectrum over time, thus providing motivation for spectroscopic observations of the host over the coming years.

Strong AGN flares reminiscent of TDFs have been observed in at least one other event. One early interpretation of a flare from the known Seyfert 1.9 galaxy IC3599 was that of tidal disruption flare, in part due to its apparent decline similar to the expected $t^{-5/3}$ for TDFs (Grupe et al., 1995). Approximately 10 and 20 years later, the host appeared to display similarly evolving flares, and it was suggested that this may be evidence of repeat partial disruptions from a star on a 10 year orbit of its host’s supermassive black hole (Campana et al., 2015). However, analysis of optical emission based on CSS observations showed that optical flaring with a long duration and slow rise preceded the third X-ray flare by hundreds of days, a situation that was difficult to reconcile with the quickly rising nature of tidal disruption flares (Grupe et al., 2015). Thus it seems most likely that the flaring is as a result of AGN variability, possibly due to episodic feeding through interaction in a black hole binary, or due to disk instability. While the magnitude of the change in this case (~ 0.2 mag increase in the optical emission) is considerably smaller than that observed in CSS100217, and the timescale of optical emission was longer by several hundred days, it is possible that the host has undergone a similar flare. It is therefore also possible that each may undergo quasiperiodic flaring activity as IC3599 may have done, in which case further significant flares may be expected in the future.

A further possibility, suggested by (Wrzykowski et al., 2017), is that of a new class of tidal disruption flare that originates from within active hosts. Along with CSS100217 and ASASSN14li being likely candidates for this class examined within this work, other examples include OGLE16aaa (an optical flare with a re-

markably similar temperature evolution to both ASASSN14ae and ASASSN14li; Wyrzykowski et al., 2017), SDSS J095209.56+214313.3 (SDSS J0952+21; Komossa et al., 2008), SDSS J074820.67+471214.3 (SDSS J0748+47; Wang et al., 2011a) and IGR J12580+0134 (IGR J1258+01, a likely disruption of a super-Jupiter; Nikolajuk and Walter, 2013). While the observations of each flare are quite different, with SDSS J0952+21 and SDSS J0748+47 identified through transient coronal line emission that exceeds that expected from supernovae, IGR J1258+01 detected in high energy X-rays with *INTEGRAL* and OGLE16aaa (and of course CSS100217 and ASASSN14li) found through the more typical optical imaging survey method, all have been identified as being most likely associated with a tidal disruption event despite residing in hosts with evidence of AGN activity. The suggestion then is that these events could represent a missing link between the more typical TDFs from quiescent hosts and a fraction of “changing-look” QSOs with transient broad emission lines and an increase in blue continuum emission (MacLeod et al., 2016). It is therefore possible that the bias towards searching for TDFs in inactive hosts could be resulting in an underestimate of the true TDF rate.

The possible detection of this new class of AGN flare and possible sub-class of TDF further complicates an already complex situation in the classification of nuclear flares. While continued observations of each host may yet shed light on their as yet uncertain origins, the possible existence of these new classes can only be confirmed in the continued search for further examples.

5.6 Summary

Here I summarise the key findings of this study:

- Both ASASSN14ae and ASASSN14li have very strong ties to the centres of their hosts, with offsets of less than ~ 10 pc, making their association with the central supermassive black hole quite likely.
- While the lightcurves of ASASSN14ae and ASASSN14li are somewhat unusual, their temperature evolution, host properties and central position means they remain strong candidate for tidal disruption flares.
- CSS100217 has exhibited a marked change in its host level emission, likely a change in accretion rate of the AGN, that occurred contemporaneously with the flare, indicating a possible link.
- Given the similarities between the events, it is conceivable that both CSS100217 and ASASSN15lh belong to a new class of AGN flare. However the lack of

obvious signs of AGN activity in the host of ASASSN15lh may make a TDF flare origin more likely. Alternatively, ASASSN14li and CSS100217 could represent members of a new sub-class of TDF that fills a possible gap between inactive host TDFs and “changing-look” QSOs.

Chapter 6

Conclusions

The main aim of this work has been to attempt to constrain the physical origins of a number of nuclear transients. This has involved both analysis of the flares themselves including their lightcurve properties and spectral evolution, and of their host galaxies, all with a view to comparing the observed properties with those expected based on theoretical modelling and previous examples. In so doing, this work has highlighted the large range of properties that can be exhibited by such nuclear flares, spanning orders of magnitude in optical luminosity, with vastly different X-ray, optical and radio properties and originating from a wide range of differing host galaxies. Even within relatively similar transients, subtle differences in evolution may indicate levels of complexity that are only now coming to light with the advent of systematic deep transient surveys. Below I summarise the key findings of the individual studies described in this thesis before placing them in the wider context of the field and discussing the future of the field as instrumentation and observing strategies continue to evolve and improve.

6.1 Summary of Results

6.1.1 The Relativistic Tidal Disruption Flare Candidate: Swift J1112.2-8238

The detection of a strikingly unusual flare in the form of Swift J164449.3+573451 (Swift J1644+57, Levan et al., 2011; Bloom et al., 2011; Burrows et al., 2011), and subsequently of a possible similar event in Swift J2058.4+0516 (Swift J2058+05, Cenko et al., 2012), indicated the potential discovery of an exciting new class of transient. These flares exhibited long-lived γ -ray and X-ray emission with durations far in excess of typical long GRBs. Their association with the central regions of

seemingly inactive hosts made them prime candidates for tidal disruption flares. However their detection at γ -ray wavelengths was unique among the class, while their high-energy and radio properties suggested the formation of collimated relativistic outflows similar to GRB jets (Bloom et al., 2011; Zauderer et al., 2011). The popular explanation for these events was in the tidal disruption of a star that had also launched a relativistic jet (Levan et al., 2011), though other explanations involving stellar collapse (Quataert and Kasen, 2012b; Woosley and Heger, 2012) or the disruption of a white dwarf (Krolik and Piran, 2011) have also been suggested. With their discovery, came the search for further candidates detected by *Swift*, a search that has so far yielded only one viable candidate, Swift J1112.2-8238.

The early analysis of this flare conducted within this work confirmed the flare’s suspected extragalactic origin from a host with a likely redshift of $z = 0.89$, eliminating the possibility of a local or Galactic phenomenon. In addition, the astrometry afforded by the ground-based optical imaging indicated the flare’s position was coincident with the centroid of the host to within 0.85 ± 0.93 kpc, indicating a plausible association with the central supermassive black hole. The flare displayed a number of similarities to the previous rTDF candidates including similar X-ray longevity and luminosity ($\sim 10^{46}$ erg s $^{-1}$ at 10^6 s after trigger), particularly when compared with Swift J1644+57. This fact, coupled with its more modest optical emission, places all three events in the same region of X-ray luminosity - Optical absolute magnitude phase space that is devoid of other sources. The flare even showed some signs of a rapid cessation of activity that was common to both of the previous flares (Zauderer et al., 2013; Levan et al., 2016; Pasham et al., 2015). Thus, while the somewhat flatter X-ray lightcurve and harder X-ray spectrum remained factors that set the events apart, the properties were nonetheless suggestive of belonging to the same class of event as the preceding candidates.

The lack of new candidates in the 5 years preceding Swift J1644+57 and the 5 years following Swift J1112-8238 suggests the clustering of these events in a three month period in the first half of 2011 is a simple statistical fluke and the rate for such events must necessarily be low. The determined rate of 3×10^{-10} per galaxy per year is considerably below the estimated rate for thermal TDFs of $\sim 10^{-5}$ per galaxy per year. Evidence for a collimated beam may suggest a narrow beaming angle for emission, thus limiting the detected rate to those few that happen to be orientated towards us and making the true rate somewhat higher. However, even with a quite narrow GRB-like jet, the estimates still fall short of the thermal TDF rate, indicating it is likely that only a small fraction of TDFs launch Swift J1644+57-like jets, a finding that is supported by radio observations of thermal

TDF candidates (e.g. Bower et al., 2013; van Velzen et al., 2013).

6.1.2 The Late-time Analysis of Swift J1112-8238

With a likely link to the previous rTDF candidates, further observations of the host of Swift J1112-8238 were warranted and this work recounts the findings of an analysis of high-resolution *Hubble Space Telescope* imaging, medium-resolution optical spectroscopy and radio observations of the source. As was hinted at in the lower resolution ground based imaging, the host of Swift J1112-8238 displays a rather complex morphology with an indication of at least two major components. While a bright star-forming region embedded in a wide disc cannot be entirely ruled out, the strong asymmetry of the host indicates an interacting system made up of two galaxies in the process of a merger or disruption. The X-Shooter spectrum, which confirmed the previous redshift determination, also shows signs of this complex morphology, with a clear two component structure in each detected emission line separated both spatially and in velocity space. The transient position is located 1.1 ± 0.5 kpc from the centroid of a compact bulge component, making it consistent with the centroid of the host, though the weak significance of this offset precludes strong inferences. It is also conceivable that the event represents a TDF by a migrating black hole within the interacting system.

Perhaps the most interesting development is in the detection of luminous radio emission from the host. If attributed to star formation, the luminosity implies a rate that is considerably above that determined from star-formation lines in the X-Shooter spectrum. In addition, the possible declining nature of the emission with time indicates the emission may be coming from the fading transient as opposed to steady ongoing AGN activity. The radio lightcurves of all three rTDF candidates show broad similarities, such that each may represent a scaled version of the others and the spectral indices of the events are approximately consistent throughout their evolution, though the radio follow-up of Swift J2058+05 and Swift J1112-8238 is too sparse to be certain.

Nonetheless, all three events show remarkable similarities with only minor differences between them, strongly suggesting a common origin. While these events remain the only rTDF candidates currently detected through *Swift*, evidence of low-level jet activity has been detected in observations of ASASSN14li (van Velzen et al., 2015; Holoien et al., 2016; Alexander et al., 2016), indicating some form of jet may be a ubiquitous feature of TDFs. However, it is currently unclear as to whether this represents a continuum of events with varying strengths of jetted emission or two separate sub-classes with substantially different progenitor systems.

6.1.3 An Expanding Variety of Nuclear Flares

Unlike many previous supernova surveys, ASASSN has shown an impressive ability to detect flares coincident with their nuclei (Holoien and Stanek, 2016). ASASSN14ae (Holoien et al., 2014) and ASASSN14li (Holoien and Stanek, 2016) represent two such flares whose origins have been suggested to be tidal disruption flares. *HST* imaging of these events analysed in this work continues to support this interpretation with the transient positions within ~ 10 parsecs of the centres of their hosts, placing them amongst the most tightly constrained nuclear transients detected to date. Interestingly though, they exhibit somewhat unusual lightcurves, somewhat removed from the typical power law declines seen in other candidates, and, at least in the case of ASASSN14li, may come from hosts with active nuclei. However despite these abnormalities, a tidal disruption flare origin remains the most likely explanation for these events.

CSS100217 (Drake et al., 2011) and ASASSN15lh (Dong et al., 2016), both also detected in the nuclear regions of their hosts, are two exceptionally luminous events. Indeed, their interpretation as superluminous supernovae make them the two brightest such events detected to date. However, the events have unusually bright and long-lived lightcurves, while the hosts of each show peculiar properties. Interestingly, CSS100217, a known Seyfert 2 galaxy (Drake et al., 2011), shows signs of a decline in host emission occurring contemporaneously with the flare, indicative of a change in the AGN emission that may be connected with the transient, and flaring X-ray emission that shows no correlation to optical emission. The host of ASASSN15lh is very massive and exhibits extremely minimal star formation (Dong et al., 2016; Leloudas et al., 2016), a peculiarity for a class of event inextricably linked with massive stars and typically residing in small, star-forming hosts (Lunnan et al., 2014; Leloudas et al., 2015; Angus et al., 2016). Given the properties of CSS100217 and the similarities between the two events, it is therefore possible that these outbursts represent a new class of AGN flare with structured lightcurves that contrast strongly with the typical stochastic behaviour of AGN variability. However, the lack of evidence for AGN emission in the host of ASASSN15lh argues against such an interpretation. Alternatively, either may be associated with an unusual tidal disruption flare associated with a rapidly spinning black hole (Leloudas et al., 2016), both explaining each event’s luminosity and, in the case of ASASSN15lh, its occurrence within a massive galaxy with a likely black hole mass in excess of the $10^8 M_{\odot}$ limit for the disruption of a Sun-like star (Rees, 1988). The difficulty in separating these flares from nuclear supernovae may further complicate the classification of these flares in the future. For now at least, the origins of these events

remains a mystery.

6.2 Overview and Future Prospects

From γ -ray bright rTDF candidates, through thermal TDFs to luminous atypical AGN flares and nuclear supernovae, the range of properties of nuclear transients is extremely wide. Taking just the flares discussed within this work as an (extremely incomplete) sampling of the population, these outbursts span over six orders of magnitude in X-ray luminosity and a further two orders of magnitude in peak optical luminosity. Patterns are, however, beginning to emerge as sample sizes grow. For example, recent work on tidal disruption flares has shown a continuum of events ranging from hydrogen to helium dominated and that these events show a significant preference for post-merger hosts (Arcavi et al., 2014; French et al., 2016, 2017). Despite this, the classification of individual flares remains difficult.

Part of this difficulty comes from the often sparse and inconsistent nature of multi-wavelength follow-up. The radio observations of the rTDF candidates are an excellent example of this. The first rTDF candidate, Swift J1644+57 (Levan et al., 2011; Bloom et al., 2011; Burrows et al., 2011), was comparatively nearby, luminous and particularly unusual. As a result numerous radio observations were made across a wide range of frequencies which resulted in an extremely detailed radio lightcurve (Zauderer et al., 2011; Berger et al., 2012; Zauderer et al., 2013), a portion of which is plotted in Figure 4.10. However, the subsequent rTDF candidates, Swift J2058+05 (Cenko et al., 2012; Pasham et al., 2015) and Swift J1112-8238, were more distant and less immediately spectacular, receiving far less attention and resulting in large gaps in their radio lightcurves that makes comparison to the prototype event difficult. This is, however, unsurprising, as telescope time is already an extremely precious commodity and follow-up facilities are already straining to deal with the rate of new transients produced by modern surveys.

Unfortunately this problem is only likely to increase. When the Large Synoptic Survey Telescope (LSST, Ivezić and LSST Science Collaboration, 2013) begins science operations in the early 2020's, transient rates are expected to run to tens of thousands per night. As an example, the rate of TDF candidates has greatly improved from one every few years in the 1990's to one every few months thanks to the efforts of the Palomar Transient Factory, PanSTARRS and ASASSN to name just a few. In the LSST era, this number is expected to rise steeply to several thousand per year (LSST Science Collaboration et al.). These rates are far in excess of the ability of current facilities to perform detailed and consistent follow-up observations, even

if limited to just a representative fraction of the whole. Thus considerable emphasis must be placed on the fast and accurate identification of interesting transients that would then be prioritised for further study. For example, the identification of tidal disruption flares has centred on a nuclear high temperature flare with minimal cooling, implying a blue transient with minimal change in colours for a protracted period. However, as has been shown in this work, even with extensive follow-up, multiple progenitor pathways remain possible for each flare. In many cases, theoretical modelling shows that many of these routes have overlapping properties that make it difficult to differentiate them based on the available observations.

However, such determinations are being made easier through the wealth of archival observations that continues to expand, providing baseline comparisons for many of the host galaxies of newly discovered flares. Photometric and spectroscopic surveys like SDSS, *WISE* and *GALEX* cover vast swathes of the electromagnetic spectrum, enabling the comparison of host properties to those of other candidates, even while transient emission still dominates the host. Photometric monitoring campaigns like CSS (and eventually LSST) will provide limits on host variability over baselines of potentially decades, a feature of particular interest in the study of possible AGN flares.

With the discovery of the first gravitational waves comes the possibility for multi-messenger astronomy that could one day make all the difference in determining the origins of some nuclear flares. While such observations would necessarily need to be carried out in space due to the limitations of low frequency signal detection on the ground, the detection of a gravitational wave signal from the position of a nuclear transient flare could finally provide the evidence for the categorical identification of a star being disrupted by a supermassive black hole. Unfortunately, the small signal produced by a tidally disrupted star around a massive black hole would have limited detection by *Evolved Laser Interferometer Space Antenna* (*eLISA*; Amaro-Seoane et al., 2013) out to ~ 20 Mpc (Kobayashi et al., 2004; East, 2014), a distance that only contains most of the Virgo cluster with only about 10^3 member galaxies. While the TDF rates for low mass black holes are poorly constrained, even if each of those galaxies hosted a $10^6 M_\odot$ black hole or greater, the expected detectable TDF rate would then be about 1 per hundred years, making a detection unlikely in the 5 year *eLISA* mission length. Thus the gravitational wave detection of a disrupted star will likely be beyond our capability for some time to come.

However, extreme mass ratio inspiral (EMRI) rates for compact objects that would be swallowed whole by the SMBH would stand a much better chance of detection with expected rates of a few tens within the mission’s lifetime (Gair and

Porter, 2012). Such rates would provide constraints on the SMBH population that could be contrasted with those determined by AGN and TDF methods. It has also been suggested that the recoil on the merged result of supermassive black hole binary (SMBHB) coalescence could result in the formation of a new loss cone of stars, resulting in extremely high TDF rates of up to $0.1 - 1 \text{ yr}^{-1}$ (Stone and Loeb, 2011). This suggests that the detection of a coalescing SMBHB signal from *eLISA* could precede a string of tidal disruption flares, making them prime targets for targeted follow-up.

Identification of nuclear flares remains a difficult process with considerable overlap between observed properties and limited constraints on theoretical models. However the correct identification of these flares are essential in developing our understanding of galaxy formation and evolution. With the as yet poorly understood links between supermassive black holes and their host galaxies, evidenced by the the well-documented scaling relations between SMBH mass and their host's masses (e.g. Haring and Rix, 2004; Kormendy and Ho, 2013) and stellar velocity dispersions (e.g. Ferrarese and Merritt, 2000; Gebhardt et al., 2000; Kormendy and Ho, 2013), comes the need to carefully study the population and evolution of SMBHs. Tidal disruption flares in particular offer a useful view into the cores of otherwise unreachable galaxies and with the possibility of a variant of these flares with relativistic jets that are observable to much greater distances comes the chance to study the SMBH population back to the early Universe. Such observations could help determine the origins of SMBH seeds by placing limits on early black hole masses, help determine the importance of galaxy mergers in SMBH growth by tracking their evolution across cosmic time and finally fill in the gap between stellar mass and supermassive black holes by finding the elusive intermediate mass black hole population through observations of tidal disruption flares in tiny dwarf galaxies beyond the reach of other methods. Without a clear and effective method of identification, however, such insights could be missed, prompting the need for thorough and extensive study of the few candidates we have and the multitude expected to come.

Bibliography

- Abdo, A. A., Ackermann, M., Ajello, M., et al. 2015, *The Astrophysical Journal*, 799(2):143.
- Aird, J., Coil, A. L., Georgakakis, A., et al. 2015, *Monthly Notices of the Royal Astronomical Society*, 451(2):1892–1927.
- Alard, C. 2000, *Astronomy and Astrophysics Supplement Series*, 144:363–370.
- Alard, C. and Lupton, R. H. 1998, *The Astrophysical Journal*, 503(1):325–331.
- Alexander, K. D., Berger, E., Guillochon, J., Zauderer, B. A., and Williams, P. K. G. 2016, *The Astrophysical Journal Letters*, 819(2):L25.
- Alvarez, M. A., Wise, J. H., and Abel, T. 2009, *The Astrophysical Journal*, 701(2):L133–L137.
- Amaro-Seoane, P., Aoudia, S., Babak, S., et al. 2013, *GW Notes*, 6:4–110.
- Anderson, J. and Bedin, L. R. 2010, *Publications of the Astronomical Society of the Pacific*, 122:1035–1064.
- Angel, J. R. P. and Stockman, H. S. 1980, *Annual Review of Astronomy and Astrophysics*, 18(1):321–361.
- Angus, C. R., Levan, A. J., Perley, D. A., et al. 2016, *Monthly Notices of the Royal Astronomical Society*, 458(1):84–104.
- Antonucci, R. 1993, *Annual Review of Astronomy and Astrophysics*, 31(1):473–521.
- Arcavi, I., Gal-Yam, A., Sullivan, M., et al. 2014, *The Astrophysical Journal*, 793(38).
- Asplund, M., Grevesse, N., Sauval, A. J., and Scott, P. 2009, *Annual Review of Astronomy and Astrophysics*, 47(1):481–522.

- Baganoff, F. K., Maeda, Y., Morris, M., et al. 2003, *The Astrophysical Journal*, 591: 891–915.
- Baldwin, A., Phillips, M. M., and Terlevich, R. 1981, *Publications of the Astronomical Society of the Pacific*, 93(551):817.
- Barbary, K., Dawson, K. S., Tokita, K., et al. 2009, *The Astrophysical Journal*, 690 (2):1358–1362.
- Barthelmy, S. D., Baumgartner, W. H., Beardmore, A. P., et al. 2011, *GRB Coordinates Network*, 11824(1).
- Barthelmy, S. D., Barbier, L. M., Cummings, J. R., et al. 2005, *Space Science Reviews*, 120(3-4):143–164.
- Bartko, H., Martins, F., Trippe, S., et al. 2010, *The Astrophysical Journal*, 708(1): 834–840.
- Benetti, S., Nicholl, M., Cappellaro, E., et al. 2014, *Monthly Notices of the Royal Astronomical Society*, 441(1):289–303.
- Bennert, V. N., Auger, M. W., Treu, T., Woo, J.-H., and Malkan, M. A. 2011, *The Astrophysical Journal*, 742(2):107.
- Berger, E. and Chornock, R. 2011, *ATel*, 3469.
- Berger, E., Zauderer, A., Pooley, G. G., et al. 2012, *The Astrophysical Journal*, 748 (1):36.
- Bertin, E. and Arnouts, S. 1996, *Astronomy and Astrophysics Supplement Series*, 117(2):393–404.
- Bianchi, L., Herald, J., Efremova, B., et al. 2011, *Astrophysics and Space Science*, 335(1):161–169.
- Birney, D. S., Gonzalez, G., and Oesper, D. *Observational Astronomy*. Cambridge University Press, second edi edition, 2006. ISBN 978-0-521-85370-5.
- Bloom, J. S. *What Are Gamma-Ray Bursts?* Princeton University Press, Princeton, 2011. ISBN 9781400837007.
- Bloom, J. S., Giannios, D., Metzger, B. D., et al. 2011, *Science (New York, N.Y.)*, 333(6039):203–6.

- Bock, D. C.-J., Large, M. I., and Sadler, E. M. 1999, *The Astronomical Journal*, 117(3):1578–1593.
- Bombaci, I. 1996, *Astronomy and Astrophysics*, 305:871–877.
- Botticella, M. T., Trundle, C., Pastorello, a., et al. 2010, *The Astrophysical Journal*, 717:L52–L56.
- Bournaud, F. 2011, *EAS Publications Series*, 51:107–131.
- Bower, G. C. 2011, *The Astrophysical Journal*, 732(1):L12.
- Bower, G. C., Metzger, B. D., Cenko, S. B., Silverman, J. M., and Bloom, J. S. 2013, *The Astrophysical Journal*, 763(2):84.
- Branch, D. and Tammann, G. A. 1992, *Annual Review of Astronomy and Astrophysics*, 30(1):359–389.
- Breedt, E., Arévalo, P., McHardy, I. M., et al. 2009, *Monthly Notices of the Royal Astronomical Society*, 394(1):427–437.
- Breedt, E., McHardy, I. M., Arévalo, P., et al. 2010, *Monthly Notices of the Royal Astronomical Society*, 403(2):605–619.
- Breeveld, A. A., Landsman, W., Holland, S. T., et al. 2011, *AIP Conference Proceedings*, 1358:373–376.
- Brinchmann, J., Charlot, S., White, S. D. M., et al. 2004, *Monthly Notices of the Royal Astronomical Society*, 351(4):1151–1179.
- Bromm, V. and Larson, R. 2004, *Annual Review of Astronomy and Astrophysics*, 42(1953):79–118.
- Brown, G. C., Levan, A. J., Stanway, E. R., et al. 2015, *Monthly Notices of the Royal Astronomical Society*, 452(4):4297–4306.
- Brown, J. S., Martini, P., and Andrews, B. H. 2016, *Monthly Notices of the Royal Astronomical Society*, 458(2):1529–1547.
- Brown, J. S., Shappee, B. J., Holoién, T. W. S., et al. 2016, *Monthly Notices of the Royal Astronomical Society*, 462(4):3993–4000.
- Burrows, D. N., Kennea, J. a., Ghisellini, G., et al. 2011, *Nature*, 476(7361):421–424.

- Burrows, D. N., Hill, J. E., Nousek, J. A., et al. 2005, *Space Science Reviews*, 120 (3-4):165–195.
- Calzetti, D. 2012, *XXIII Canary Islands Winter School of Astrophysics*, page 41.
- Calzetti, D., Armus, L., Bohlin, R. C., et al. 2000, *The Astrophysical Journal*, 533 (2):682–695.
- Campana, S., Lodato, G., D’Avanzo, P., et al. 2011, *Nature*, 480(7375):69–71.
- Campana, S., Mainetti, D., Colpi, M., et al. 2015, *Astronomy & Astrophysics*, 581: A17.
- Cappelluti, N., Ajello, M., Rebusco, P., et al. 2009, *Astronomy and Astrophysics*, 495(2):L9–L12.
- Casertano, S., Mello, D.de, Dickinson, M., et al. 2000, *The Astronomical Journal*, 120(6):2747–2824.
- Cenko, S. B., Krimm, H. a., Horesh, A., et al. 2012, *The Astrophysical Journal*, 753 (1):77.
- Chen, X., Madau, P., Sesana, A., and Liu, F. K. 2009, *The Astrophysical Journal*, 697(2):L149–L152.
- Chen, X., Sesana, A., Madau, P., and Liu, F. K. 2011, *The Astrophysical Journal*, 729(13).
- Chevalier, R. a. and Irwin, C. M. 2011, *The Astrophysical Journal Letters*, 729:4–7.
- Chornock, R., Berger, E., Gezari, S., et al. 2014, *The Astrophysical Journal*, 780 (1):44.
- Condon, J. J. 1992, *Annual Review of Astronomy and Astrophysics*, 30(1):575–611.
- Condon, J. J. and Ransom, S. *Essential Radio Astronomy*. Princeton University Press, Princeton, 2016.
- Condon, J. J., Cotton, W. D., and Broderick, J. J. 2002, *The Astronomical Journal*, 124(2):675–689.
- Crockett, R. M., Smartt, S. J., Eldridge, J. J., et al. 2007, *Monthly Notices of the Royal Astronomical Society*, 381(2):835–850.

- Cummings, J. R., Barthelmy, S. D., Beardmore, A. P., et al. 2011, *GRB Coordinates Network*, 11823.
- Dai, D.-c., Lue, A., Starkman, G., and Stojkovic, D. 2010, *Journal of Cosmology and Astroparticle Physics*, 2010(12):4.
- De Gouveia Dal Pino, E. M. 2005, *Advances in Space Research*, 35(5):908–924.
- Devecchi, B. and Volonteri, M. 2009, *The Astrophysical Journal*, 694(1):302–313.
- Dexter, J. and Kasen, D. 2013, *The Astrophysical Journal*, 772(1):30.
- Domínguez, A., Siana, B., Henry, A. L., et al. 2013, *The Astrophysical Journal*, 763(2):145.
- Dong, S., Shappee, B. J., Prieto, J. L., et al. 2016, *Science*, 351(6270):257–260.
- Donley, J., Brandt, W., Eracleous, M., and Boller, T. 2002, *The Astronomical Journal*, 124:1308–1321.
- Drake, A. J., Djorgovski, S. G., Mahabal, A., et al. 2009, *The Astrophysical Journal*, 696(1):870–884.
- Drake, A. J., Djorgovski, S. G., Mahabal, A., et al. 2011, *The Astrophysical Journal*, 735(2):106.
- Dressel, L. 2016, *Wide Field Camera 3, HST Instrument Handbook*, (January).
- East, W. E. 2014, *The Astrophysical Journal*, 795:135.
- Elitzur, M. 2012, *The Astrophysical Journal Letters*, 33(1991):2007–2009.
- Elitzur, M., Ho, L. C., and Trump, J. R. 2014, *Monthly Notices of the Royal Astronomical Society*, 438(4):3340–3351.
- Esquej, P., Saxton, R. D., Freyberg, M. J., et al. 2007, *Astronomy and Astrophysics*, 462(3):L49–L52.
- Evans, C. R. and Kochanek, C. S. 1989, *The Astrophysical Journal*, 346:L13.
- Evans, P. A., Beardmore, A. P., Page, K. L., et al. 2007, *Astronomy and Astrophysics*, 469:379–385.
- Evans, P. A., Beardmore, A. P., Page, K. L., et al. 2009, *Monthly Notices of the Royal Astronomical Society*, 397(3):1177–1201.

- Evans, P. A., Willingale, R., Osborne, J. P., et al. 2014, *Monthly Notices of the Royal Astronomical Society*, 444(1):250–267.
- Evans, P., Willingale, R., Osborne, J. P., et al. 2010, *Astronomy & Astrophysics*, 102:7.
- Ewing, A. 1964, *The Science News-Letter*, 85(3):39.
- Fan, X. 2006, *New Astronomy Reviews*, 50(9-10):665–671.
- Ferrarese, L. and Merritt, D. 2000, *The Astrophysical Journal*, 539(1):L9–L12.
- Figer, D. F. 2008, *Proceedings of the International Astronomical Union*, 250:247–256.
- Filippenko, A. V. 1997, *Annual Review of Astronomy and Astrophysics*, 35(1):309–355.
- Fisher, D. B. and Drory, N. 2008, *The Astronomical Journal*, 136(2):773–839.
- Fitzpatrick, E. L. 1999, *Publications of the Astronomical Society of the Pacific*, 111(755):63–75.
- Frank, J. and Rees, M. J. 1976, *Monthly Notices of the Royal Astronomical Society*, 176:633–647.
- French, K. D., Arcavi, I., and Zabludoff, A. 2016, *The Astrophysical Journal*, 818(1):L21.
- French, K. D., Arcavi, I., and Zabludoff, A. 2017, *The Astrophysical Journal*, 835(2):176.
- Freudling, W., Romaniello, M., Bramich, D., et al. 2013, *Astronomy and Astrophysics*, 559:1–12.
- Fruchter, A. S. and Hook, R. N. 2002, *Publications of the Astronomical Society of the Pacific*, 114(792):144–152.
- Fruchter, A. S., Levan, A. J., Strolger, L., et al. 2006, *Nature*, 441(7092):463–468.
- Fruchter, A. S. 2010, *2010 Space Telescope Science Institute Calibration Workshop*.
- Fynbo, J. P. U., Jakobsson, P., Prochaska, J. X., et al. 2009, *The Astrophysical Journal Supplement Series*, 185(2):526–573.

- Gabasch, A., Bender, R., Seitz, S., et al. 2006, *Astronomy & Astrophysics*, 448(1): 101–121.
- Gadotti, D. A. 2009, *Monthly Notices of the Royal Astronomical Society*, 393(4): 1531–1552.
- Gair, J. R. and Porter, E. K. 2012, *arXiv:1210.8066*.
- Gal-Yam, A. 2012, *Science*, 337(6097):927–932.
- Gal-Yam, A., Mazzali, P., Ofek, E. O., et al. 2009, *Nature*, 462(7273):624–7.
- Gal-Yam, A., Fox, D. B., Kulkarni, S. R., et al. 2005, *The Astrophysical Journal*, 630(1):L29–L32.
- Gaskell, C. M., Cappellaro, E., Dinerstein, H. L., et al. 1986, *The Astrophysical Journal*, 306:L77–L80.
- Gebhardt, K., Bender, R., Bower, G., et al. 2000, *The Astrophysical Journal*, 539(1):L13–L16.
- Geha, M., Brown, T. M., Tumlinson, J., et al. 2013, *The Astrophysical Journal*, 771:29–9.
- Gehrels, N., Chincarini, G., Giommi, P., et al. 2005, *The Astrophysical Journal*, 621(1):558.
- Gendre, B., Stratta, G., Atteia, J. L., et al. 2013, *The Astrophysical Journal*, 766(1):30.
- Gezari, S., Martin, D. C., Milliard, B., et al. 2006, *The Astrophysical Journal*, 653(1):L25–L28.
- Gezari, S., Basa, S., Martin, D. C., et al. 2008, *The Astrophysical Journal*, 676(2): 28.
- Gezari, S., Chornock, R., Rest, A., et al. 2012, *Nature*, 485(7397):217–20.
- Ghez, A. M., Salim, S., Weinberg, N. N., et al. 2008, *The Astrophysical Journal*, 689(2):1044–1062.
- Gilliland, R. L. and Rajan, A. 2011, (Instrument Science Report WFC3):1–18.
- Gonzaga, S., Hack, W., Fruchter, A., and Mack, J. 2012, *The DrizzlePac Handbook*, (June).

- Greenstein, J. L. and Matthews, T. A. 1963, *Nature*, 197(4872):1041–1042.
- Greiner, J., Schwarz, R., Zharikov, S., and Orio, M. 2000, *Astronomy and Astrophysics*, 362:25–28.
- Greiner, J., Mazzali, P. a., Kann, D. A., et al. 2015, *Nature*, 523(7559):189–192.
- Grupe, D., Beuermann, K., and Mannheim, K. 1995, *Astronomy & Astrophysics*, 299:L5–L8.
- Grupe, D., Komossa, S., and Saxton, R. 2015, *The Astrophysical Journal*, 803(2):L28.
- Guillochon, J. and Ramirez-Ruiz, E. 2013, *The Astrophysical Journal*, 767(1):25.
- Guillochon, J., Manukian, H., and Ramirez-Ruiz, E. 2014, *The Astrophysical Journal*, 783(1):23.
- Güver, T. and Özel, F. 2009, *Monthly Notices of the Royal Astronomical Society*, 400:2050–2053.
- Haardt, F. and Maraschi, L. 1991, *The Astrophysical Journal*, 380:L51.
- Hamuy, M., Phillips, M. M., Suntzeff, N. B., et al. 1996, *The Astronomical Journal*, 112(9):2438.
- Haring, N. and Rix, H.-W. 2004, *The Astrophysical Journal*, 604(2):L89–L92.
- Hazard, C., Mackey, M., and Shimmins, A. 1963, *Nature*, 197(4872):1037–1039.
- Heger, A., Fryer, C. L., Woosley, S. E., Langer, N., and Hartmann, D. H. 2003, *The Astrophysical Journal*, 591(1):288–300.
- Hillebrandt, W. and Niemeyer, J. C. 2000, *Annual Review of Astronomy and Astrophysics*, 38(1):191–230.
- Hills, J. G. 1975, *Nature*, 254:295–298.
- Hirano, S., Hosokawa, T., Yoshida, N., et al. 2014, *The Astrophysical Journal*, 781(2):60.
- Hockey, T. A., Trimble, V., Bracher, K., et al. *The Biographical Encyclopedia of Astronomers*. Springer New York, New York, NY, 2007. ISBN 978-0-387-31022-0.
- Holoien, T. W. S., Prieto, J. L., Bersier, D., et al. 2014, *Monthly Notices of the Royal Astronomical Society*, 445(3):3263–3277.

- Holoien, T. W. S., Kochanek, C. S., Prieto, J. L., et al. 2016, *Monthly Notices of the Royal Astronomical Society*, 455(3):2918–2935.
- Holoien, T. and Stanek, K. 2016.
- Hook, I. M., Jørgensen, I., Allington-Smith, J. R., et al. 2004, *Publications of the Astronomical Society of the Pacific*, 116:425–440.
- Hopkins, P. F. and Hernquist, L. 2009, *The Astrophysical Journal*, 698(2):1550–1569.
- Hosokawa, T., Omukai, K., Yoshida, N., and Yorke, H. W. 2011, *Science*, 334(6):1250.
- Insera, C., Pastorello, a., Turatto, M., et al. 2013, *Astronomy & Astrophysics*, 555:A142.
- Ivanenko, D. D. and Kurdgelaidze, D. F. 1965, *Astrophysics*, (5):251–252.
- Ivanov, P. B., Polnarev, A. G., and Saha, P. 2005, *Monthly Notices of the Royal Astronomical Society*, 358(4):1361–1378.
- Ivezić, v. and LSST Science Collaboration, . 2013, *Large Synoptic Survey Telescope (LSST) Science Requirements Document*.
- Jennison, R. C. and Das Gupta, M. K. 1953, *Nature*, 172(4387):996–997.
- Just, A., Khan, F. M., Berczik, P., Ernst, A., and Spurzem, R. 2011, *Monthly Notices of the Royal Astronomical Society*, 411:653–674.
- Kaiser, N., Kaiser, N., Aussel, H., et al. 2002, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 4836:154–164.
- Kann, D. A., Klose, S., Zhang, B., et al. 2011, *The Astrophysical Journal*, 734(2):96.
- Kasen, D. and Bildsten, L. 2010, *The Astrophysical Journal*, 717(2007):245–249.
- Kasen, D. and Woosley, S. E. 2009, *The Astrophysical Journal*, 703(2):2205–2216.
- Kaspi, S., Maoz, D., Netzer, H., et al. 2005, *The Astrophysical Journal*, 629(1):61–71.
- Kauffmann, G., Heckman, T. M., Tremonti, C., et al. 2003, *Monthly Notices of the Royal Astronomical Society*, 346(4):1055–1077.

- Kauffmann, G., Heckman, T. M., White, S. D. M., et al. 2003, *Monthly Notices of the Royal Astronomical Society*, 341(1):33–53.
- Keel, W. C., Kennicutt, R. C., J., Hummel, E., and Hulst, J. M. van der. 1985, *The Astronomical Journal*, 90(5):708.
- Kennea, J., Romano, P., Krimm, H., et al. 2011, *ATel*, 3242.
- Kesden, M. 2012, *Physical Review D - Particles, Fields, Gravitation and Cosmology*, 85(2):1–9.
- Kewley, L. J., Dopita, M. A., Sutherland, R. S., Heisler, C. A., and Trevena, J. 2001, *The Astrophysical Journal*, 556(1):121–140.
- Kewley, L. J. and Ellison, S. L. 2008, *The Astrophysical Journal*, 681(2):1183–1204.
- Kewley, L. J., Geller, M. J., and Jansen, R. A. 2004, *The Astronomical Journal*, 127(4):2002–2030.
- Kewley, L. J., Groves, B., Kauffmann, G., and Heckman, T. 2006, *Monthly Notices of the Royal Astronomical Society*, 372(3):961–976.
- Khan, F. M., Fiacconi, D., Mayer, L., Berczik, P., and Just, A. 2016, *The Astrophysical Journal*, 828(2):73.
- Khlopov, M. Y., Rubin, S. G., and Sakharov, A. S. 2005, *Astroparticle Physics*, 23(2):265–277.
- Kobayashi, S., Laguna, P., Phinney, E. S., and Meszaros, P. 2004, *The Astrophysical Journal*, 615(2):855–865.
- Kobulnicky, H. A. and Kewley, L. 2004, *The Astrophysical Journal*, 617(1):240–261.
- Komossa, S. 2015, *Journal of High Energy Astrophysics*, 7:148–157.
- Komossa, S. and Merritt, D. 2008, *The Astrophysical Journal Letters*, 3:1–6.
- Komossa, S., Zhou, H., and Wang, T. 2008, *The Astrophysical Journal*, 678:13–16.
- Komossa, S., Zhou, H., Rau, A., et al. 2009, *The Astrophysical Journal*, 701(1):105–121.
- Komossa, S. and Greiner, J. 1999, *Astronomy & Astrophysics*, 48:45–48.
- Kormendy, J. and Ho, L. C. 2013, *Annual Review of Astronomy and Astrophysics*, 51(1):511–653.

- Kotilainen, J. K., Ward, M. J., and Williger, G. M. 1993, *Monthly Notices of the Royal Astronomical Society*, 263(3):655–674.
- Kouveliotou, C., Meegan, C. A., Fishman, G. J., et al. 1993, *The Astrophysical Journal*, 413:L101.
- Kraft, R. P., Burrows, D. N., and Nousek, J. A. 1991, *The Astrophysical Journal*, 374(9):344.
- Krimm, H. A., Kennea, J. A., Holland, S. T., et al. 2011, *ATel*, 3384.
- Krimm, H. A., Holland, S. T., Corbet, R. H. D., et al. 2013, *The Astrophysical Journal Supplement Series*, 209(1):14.
- Krimm, H. and Barthelmy, S. 2011, *GRB Coordinates Network*, 11891.
- Krimm, H., Kennea, J., Holland, S., et al. 2011, *ATel*, 3463.
- Krolik, J. H. *Active Galactic Nuclei: From the Central Black Hole to the Galactic Environment*. Princeton University Press, 1999.
- Krolik, J. H. and Piran, T. 2011, *The Astrophysical Journal*, 743(2):134.
- Krühler, T., Malesani, D., Fynbo, J. P. U., et al. 2015, *Astronomy & Astrophysics*, 581:A125.
- LaMassa, S. M., Cales, S., Moran, E. C., et al. 2015, *The Astrophysical Journal*, 800(2):144.
- Law, N. M., Kulkarni, S. R., Dekany, R. G., et al. 2009, *Publications of the Astronomical Society of the Pacific*, 121(886):1395–1408.
- Leloudas, G., Schulze, S., Krühler, T., et al. 2015, *Monthly Notices of the Royal Astronomical Society*, 449(1):917–932.
- Leloudas, G., Fraser, M., Stone, N. C., et al. 2016, *Nature Astronomy*, 1.
- Levan, A. J., Tanvir, N. R., Cenko, S. B., et al. 2011, *Science (New York, N.Y.)*, 333(6039):199–202.
- Levan, A. J., Read, A. M., Metzger, B. D., Wheatley, P. J., and Tanvir, N. R. 2013, *The Astrophysical Journal*, 771(2):136.
- Levan, A. J., Tanvir, N. R., Starling, R. L. C., et al. 2014, *The Astrophysical Journal*, 781(1):13.

- Levan, A. J., Tanvir, N. R., Brown, G. C., et al. 2016, *The Astrophysical Journal*, 819(1):51.
- Levan, A. and Tanvir, N. 2012, *ATel*, 4610.
- Liu, F. K., Li, S., and Chen, X. 2009, *The Astrophysical Journal*, 706(1):L133–L137.
- Liu, F. K., Li, S., and Komossa, S. 2014, *The Astrophysical Journal*, 786(2):103.
- Lodato, G., King, A. R., and Pringle, J. E. 2009, *Monthly Notices of the Royal Astronomical Society*, 392(1):332–340.
- Lodato, G. and Rossi, E. M. 2011, *Monthly Notices of the Royal Astronomical Society*, 410(1):359–367.
- Lodato, G., Franchini, A., Bonnerot, C., and Rossi, E. M. 2015, *Journal of High Energy Astrophysics*, 7:158–162.
- LSST Science Collaboration, , Abell, P. A., Allison, J., et al. Technical report.
- Lunnan, R., Chornock, R., Berger, E., et al. 2014, *The Astrophysical Journal*, 787(2):138.
- Lynden-Bell, D. 1969, *Nature*, 223(5207):690–694.
- MacLeod, C. L., Ivezić, v., Kochanek, C. S., et al. 2010, *The Astrophysical Journal*, 721(2):1014–1033.
- MacLeod, C. L., Ross, N. P., Lawrence, A., et al. 2016, *Monthly Notices of the Royal Astronomical Society*, 457(1):389–404.
- MacLeod, M., Guillochon, J., and Ramirez-Ruiz, E. 2012, *The Astrophysical Journal*, 757(2):134.
- MacLeod, M., Ramirez-Ruiz, E., Grady, S., and Guillochon, J. 2013, *The Astrophysical Journal*, 777(2):133.
- Magorrian, J. and Tremaine, S. 1999, *Monthly Notices of the Royal Astronomical Society*, 309:447.
- Mangano, V., Burrows, D. N., Sbarufatti, B., and Cannizzo, J. K. 2016, *The Astrophysical Journal*, 817(2):103.
- Marcha, M. J. M. and Caccianiga, A. 2013, *Monthly Notices of the Royal Astronomical Society*, 430(3):2464–2475.

- Markoff, S. 2005, *The Astrophysical Journal*, 618:L103–L106.
- Marscher, A. P. 2010, *The Jet Paradigm: From Microquasars to Quasars*, 794: 173–201.
- Marsh, T. R., Horne, K., and Rosen, S. 1991, *The Astrophysical Journal*, 366: 535–543.
- Martinez, P. and Klotz, A. *A Practical Guide to CCD Astronomy*. Cambridge University Press, Cambridge, 1998.
- Martini, P. 2003, *Coevolution of Black Holes and Galaxies*, 1:16.
- Matsumoto, T., Nakauchi, D., Ioka, K., Heger, A., and Nakamura, T. 2015, *The Astrophysical Journal*, 810(1):64.
- Matsuoka, Y., Onoue, M., Kashikawa, N., et al. 2016, *arXiv:1603.02281 [astro-ph]*.
- Matt, G., Guainazzi, M., and Maiolino, R. 2003, *Monthly Notices of the Royal Astronomical Society*, 342(2):422–426.
- Matthews, T. A. and Sandage, A. R. 1963, *The Astrophysical Journal*, 138:30.
- Mattila, S., Väisänen, P., Farrah, D., et al. 2007, *The Astrophysical Journal*, 659 (1):L9–L12.
- Mauch, T., Murphy, T., Buttery, H. J., et al. 2003, *Monthly Notices of the Royal Astronomical Society*, 342(4):1117–1130.
- Mauch, T. and Sadler, E. M. 2007, *Monthly Notices of the Royal Astronomical Society*, 375(3):931–950.
- Maund, J. R. and Smartt, S. J. 2005, *Monthly Notices of the Royal Astronomical Society*, 360(1):288–304.
- Maund, J. R., Smartt, S. J., and Schweizer, F. 2005, *The Astrophysical Journal*, 630(1):L33–L36.
- Mazzali, P. a., Deng, J., Nomoto, K., et al. 2006, *Nature*, 442(7106):1018–1020.
- McCall, M. L. 2004, *The Astronomical Journal*, 128:2144.
- McGaugh, S. S. 1991, *The Astrophysical Journal*, 380:140–150.
- Melia, F. and Falcke, H. 2001, *Annual Review of Astronomy and Astrophysics*, 39: 309–352.

- Mendez, A. J., Coil, A. L., Aird, J., et al. 2013, *The Astrophysical Journal*, 770(1): 40.
- Merloni, A., Dwelly, T., Salvato, M., et al. 2015, *Monthly Notices of the Royal Astronomical Society*, 452(1):69–87.
- Merloni, A. and Heinz, S. In *Planets, Stars and Stellar Systems*, pages 503–566. Springer Netherlands, Dordrecht, 2013.
- Merloni, A., Heinz, S., and Matteo, T. D. 2003, *Monthly Notices of the Royal Astronomical Society*, 345:1057–1076.
- Metzger, B. D., Giannios, D., and Mimica, P. 2012, *Monthly Notices of the Royal Astronomical Society*, 420(4):3528–3537.
- Miller, A. A., Chornock, R., Perley, D. A., et al. 2009, *The Astrophysical Journal*, 690:1303.
- Miller, J. M. and Gültekin, K. 2011, *The Astrophysical Journal*, 738(1):L13.
- Mimica, P., Giannios, D., Metzger, B. D., and Aloy, M. A. 2015, *Monthly Notices of the Royal Astronomical Society*, 450:2824–2841.
- Minkowski, R. 1941, *Publications of the Astronomical Society of the Pacific*, 53:224.
- Mo, H., Bosch, F. van den, and White, S. *Galaxy Formation and Evolution*. Cambridge University Press, Cambridge, 2010. ISBN 9780511807244.
- Modigliani, A., Goldoni, P., Royer, F., et al. 2010, *Proceedings of SPIE*, 7737: 773728.
- Morris, M. 1993, *The Astrophysical Journal*, 408:496.
- Mortlock, D. J., Warren, S. J., Venemans, B. P., et al. 2011, *Nature*, 474(7353): 616–619.
- Moustakas, J., Kennicutt, Jr., R. C., and Tremonti, C. A. 2006, *The Astrophysical Journal*, 642(2):775–796.
- Murphy, E. J., Condon, J. J., Schinnerer, E., et al. 2011, *The Astrophysical Journal*, 737(2):67.
- Murphy, T., Sadler, E. M., Ekers, R. D., et al. 2010, *Monthly Notices of the Royal Astronomical Society*, 402(4):2403–2423.

- Mushotzky, R. In Barger, A. J., editor, *Supermassive Black Holes in the Distant Universe*, chapter 2, pages 53–87. Kluwer Academic Publishers, Dordrecht, 2004.
- Mushotzky, R. F., Done, C., and Pounds, K. A. 1993, *Annual Review of Astronomy and Astrophysics*, 31(1):717–761.
- Nicholl, M. and Smartt, S. J. 2016, *Monthly Notices of the Royal Astronomical Society: Letters*, 457:L79–L83.
- Nikolajuk, M. and Walter, R. 2013, *Astronomy & Astrophysics*, 552:A75.
- NOMOTO, K., IWAMOTO, K., YAMAOKA, H., et al. 1995, *Annals of the New York Academy of Sciences*, 759(1):360–367.
- Nousek, J. A., Kouveliotou, C., Grupe, D., et al. 2006, *The Astrophysical Journal*, 642(1):389–400.
- Oke, J. 1963, *Nature*, 197(4862):1040–1041.
- Osterbrock, D. E. and Ferland, G. J. *Astrophys. of Gaseous Nebulae and Active Galactic Nuclei*. University Science Books, 2nd editio edition, 2006. ISBN 1-891389-34-3.
- Pacucci, F., Volonteri, M., and Ferrara, A. 2015, *Monthly Notices of the Royal Astronomical Society*, 452(2):1922–1933.
- Palmer, D., Bethelmy, S., Baumgartner, W., et al. 2010, *GRB Coordinates Network*, 11500.
- Pasham, D. R., Cenko, S. B., Levan, A. J., et al. 2015, *The Astrophysical Journal*, 805(1):68.
- Peng, C. Y., Ho, L. C., Impey, C. D., and Rix, H.-W. 2002, *The Astronomical Journal*, 124:266–293.
- Peng, C. Y., Ho, L. C., Impey, C. D., and Rix, H.-W. 2010, *The Astronomical Journal*, 139(6):2097–2129.
- Perley, D. a., Levan, a. J., Tanvir, N. R., et al. 2013, *The Astrophysical Journal*, 778(2):128.
- Peterson, B. and Horne, K. In Livio, M. and Casertano, S., editors, *Planets to cosmology : essential science in the final years of the Hubble Space Telescope, Proceedings of the Space Telescope Science Institute Symposium*, page 89, Cambridge, 2006. Cambridge University Press.

- Peterson, B. M. 2001, *The Starburst-AGN Connection*, page 63.
- Peterson, B. M. 2014, *Space Science Reviews*, 183(1-4):253–275.
- Phinney, E. 1989, *The Center of the Galaxy*.
- Pilyugin, L. S. 2001, *Astronomy and Astrophysics*, 374(2):412–420.
- Podsiadlowski, P., Joss, P. C., and Hsu, J. J. L. 1992, *The Astrophysical Journal*, 391:246.
- Poole, T. S., Breeveld, A. A., Page, M. J., et al. 2007, *Monthly Notices of the Royal Astronomical Society*, 383(2):627–645.
- Puccetti, S., Fiore, F., Risaliti, G., et al. 2007, *Monthly Notices of the Royal Astronomical Society*, 377(2):607–616.
- Quataert, E. and Kasen, D. 2012, *Monthly Notices of the Royal Astronomical Society: Letters*, 419(1):L1–L5.
- Quataert, E. and Kasen, D. 2012, *Monthly Notices of the Royal Astronomical Society: Letters*, 419(1):L1–L5.
- Quimby, R. M., Kulkarni, S. R., Kasliwal, M. M., et al. 2011, *Nature*, 474(7352):487–9.
- Quintero, A. D., Hogg, D. W., Blanton, M. R., et al. 2004, *The Astrophysical Journal*, 602(1):190–199.
- Rakavy, G. and Shaviv, G. 1967, *The Astrophysical Journal*, 148:803–816.
- Rees, M. J. 1977, *Quarterly Journal of the Royal Astronomical Society*, 18:429–442.
- Rees, M. J. 1988, *Nature*, 333(6173):523–528.
- Renzini, A. In *The Formation of Galactic Bulges*, page 9. Cambridge Contemporary Astrophysics, 1999. ISBN 9780521663342.
- Sanders, N. E., Soderberg, A. M., Gezari, S., et al. 2014, *American Astronomical Society*, 223:216.
- Sault, R., Teuben, P., and Wright, M. 1995, *Astronomical Data Analysis Software and Systems IV, A.S.P. Conference Series*, 77:433.
- Savaglio, S., Glazebrook, K., and Le Borgne, D. 2009, *The Astrophysical Journal*, 691(1):182–211.

- Saxton, C. J., Soria, R., Wu, K., and Kuin, N. P. M. 2012, *Monthly Notices of the Royal Astronomical Society*, 422(2):1625–1639.
- Saxton, R., Read, A., and Esquej, P. 2012, *Astronomy & Astrophysics*, 541:A106.
- Sbarufatti, B., Burrows, D., Gehrels, N., and Kennea, J. 2012, *ATel*, 4398.
- Schlafly, E. F. and Finkbeiner, D. P. 2011, *The Astrophysical Journal*, 737(2):103.
- Schmidt, M. 1963, *Nature*, 197(4872):1040.
- Seyfert, C. K. 1943, *The Astrophysical Journal*, 97(2):28.
- Seymour, N., Dwelly, T., Moss, D., et al. 2008, *Monthly Notices of the Royal Astronomical Society*, 386(3):1695–1708.
- Shaviv, N. J. 1998, *The Astrophysical Journal*, 494(2):L193–L197.
- Shields, G. a. 1999, *Proceedings of the Astronomical Society of the Pacific*, 111:661.
- Smartt, S. J. 2009, *Annual Review of Astronomy and Astrophysics*, 47(1):63–106.
- Spolyar, D., Freese, K., and Gondolo, P. 2008, *Physical Review Letters*, 100(5):051101.
- Stern, D., Assef, R. J., Benford, D. J., et al. 2012, *Astrophysical Journal*, 7530:30.
- Stone, N. and Loeb, A. 2011, *Bulletin of the American Astronomical Society*, 218:23504.
- Stone, N. C. and Velzen, S.van. 2016, *The Astrophysical Journal*, 825(1):L14.
- Stratta, G., Gendre, B., Atteia, J. L., et al. 2013, *The Astrophysical Journal*, 779(1):66.
- Strubbe, L. E. and Quataert, E. 2009, *Monthly Notices of the Royal Astronomical Society*, 400(4):2070–2084.
- Strubbe, L. E. and Quataert, E. 2011, *Monthly Notices of the Royal Astronomical Society*, 415(1):168–180.
- Sunyaev, R. a., Markevitch, M., and Pavlinsky, M. 1993, *Astrophysical Journal*, 407:606–610.
- Susa, H., Hasegawa, K., and Tominaga, N. 2014, *The Astrophysical Journal*, 792(1):32.

- Sweet, S. M., Drinkwater, M. J., Meurer, G., et al. 2014, *The Astrophysical Journal*, 782(1):35.
- Sweet, S. M., Drinkwater, M. J., Meurer, G., et al. 2014, *The Astrophysical Journal*, 782(1):35.
- Tanvir, N., Levan, A., Hounsell, R., et al. 2013, *GRB Coordinates Network*, 15489.
- Tchekhovskoy, A., Metzger, B. D., Giannios, D., and Kelley, L. Z. 2014, *Monthly Notices of the Royal Astronomical Society*, 437(3):2744–2760.
- Thompson, T. A., Quataert, E., Waxman, E., Murray, N., and Martin, C. L. 2006, *The Astrophysical Journal*, 645(100):186.
- Thöne, C. C., de Ugarte Postigo, A., Fryer, C. L., et al. 2011, *Nature*, 480(7375):72–74.
- Tody, D. 1986, *Instrumentation in astronomy VI*, 627:733.
- Tody, D. 1993, *Astronomical Data Analysis Software and Systems II, A.S.P. Conference Series*, 52:173.
- Tremonti, C. A., Heckman, T. M., Kauffmann, G., et al. 2004, *The Astrophysical Journal*, 613(2):898–913.
- Trump, J. R., Impey, C. D., Elvis, M., et al. 2009, *The Astrophysical Journal*, 696:1195–1212.
- Tundo, E., Bernardi, M., Hyde, J. B., Sheth, R. K., and Pizzella, A. 2007, *ApJ*, 663(1):53–60.
- Ulrich, M.-H., Maraschi, L., and Urry, C. M. 1997, *Annual Review of Astronomy and Astrophysics*, 35(1):445–502.
- Urry, C. M. and Padovani, P. 1995, *Publications of the Astronomical Society of the Pacific*, 107:803.
- Valenti, S., Benetti, S., Cappellaro, E., et al. 2008, *Monthly Notices of the Royal Astronomical Society*, 383(4):1485–1500.
- Van Dyk, S. D., Li, W., and Filippenko, A. V. 2003, *Publications of the Astronomical Society of the Pacific*, 115(803):1–20.
- Velzen, S. van, Frail, D. a., Körding, E., and Falcke, H. 2013, *Astronomy & Astrophysics*, 552:A5.

- Velzen, S.van, Anderson, G. E., Stone, N. C., et al. 2015, *Science*, 351:62.
- Velzen, S.van and Farrar, G. R. 2014, *The Astrophysical Journal*, 792(1):53.
- Velzen, S.van, Farrar, G. R., Gezari, S., et al. 2011, *The Astrophysical Journal*, 741 (2):73.
- Venemans, B. P., Bañados, E., Decarli, R., et al. 2015, *The Astrophysical Journal*, 801(1):L11.
- Vercellone, S. 2012, *arXiv:1205.5510*.
- Vernet, J., Dekker, H., D’Odorico, S., et al. 2011, *Astronomy & Astrophysics*, 536: A105.
- Volonteri, M. 2010, *Astronomy and Astrophysics Review*, 18(3):279–315.
- Montigny, C.von, Bertsch, D. L., Chiang, J., et al. 1995, *The Astrophysical Journal*, 440:525.
- Vreeswijk, P. M., Savaglio, S., De Cia, A., et al. 2014, *The Astrophysical Journal*, pages 206–213.
- Wang, F. Y. and Cheng, K. S. 2012, *Monthly Notices of the Royal Astronomical Society*, 912:no–no.
- Wang, T.-G., Zhou, H.-Y., Wang, L.-F., Lu, H.-L., and Xu, D. 2011, *The Astrophysical Journal*, 740:85.
- Wang, X.-Y., Liu, R.-Y., Dai, Z.-G., and Cheng, K. S. 2011, *Physical Review D*, 84 (8):081301.
- Wang, Z., Fazio, G. G., Ashby, M. L. N., et al. 2004, *Astrophysical Journal*, 154(1): 193.
- Wegg, C. and Nate Bode, J. 2011, *The Astrophysical Journal*, 738(1):L8.
- Wheeler, J. C. and Harkness, R. P. 1990, *Reports on Progress in Physics*, 53(12): 1467–1557.
- Willingale, R., Starling, R. L. C., Beardmore, A. P., Tanvir, N. R., and O’Brien, P. T. 2013, *Monthly Notices of the Royal Astronomical Society*, 431(1):394–404.
- Woosley, S. E. and Heger, A. 2012, *The Astrophysical Journal*, 752(1):32.

- Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, *The Astronomical Journal*, 140(6):1868–1881.
- Wyrzykowski, ., Zieliski, M., Kostrzewa-Rutkowska, Z., et al. 2017, *Monthly Notices of the Royal Astronomical Society: Letters*, 465:L114–L118.
- Yoon, Y., Im, M., Jeon, Y., et al. 2015, *The Astrophysical Journal*, 808(1):96.
- Yusef-Zadeh, F., Wardle, M., Miller-Jones, J., et al. 2011, *The Astrophysical Journal*, pages 1–26.
- Zaritsky, D., Kennicutt, R. C., and Huchra, J. P. 1994, *Astrophysical Journal*, 420: 87.
- Zauderer, B. A., Berger, E., Soderberg, A. M., et al. 2011, *Nature*, 476(7361): 425–428.
- Zauderer, B. A., Berger, E., Margutti, R., et al. 2013, *The Astrophysical Journal*, 767(2):152.