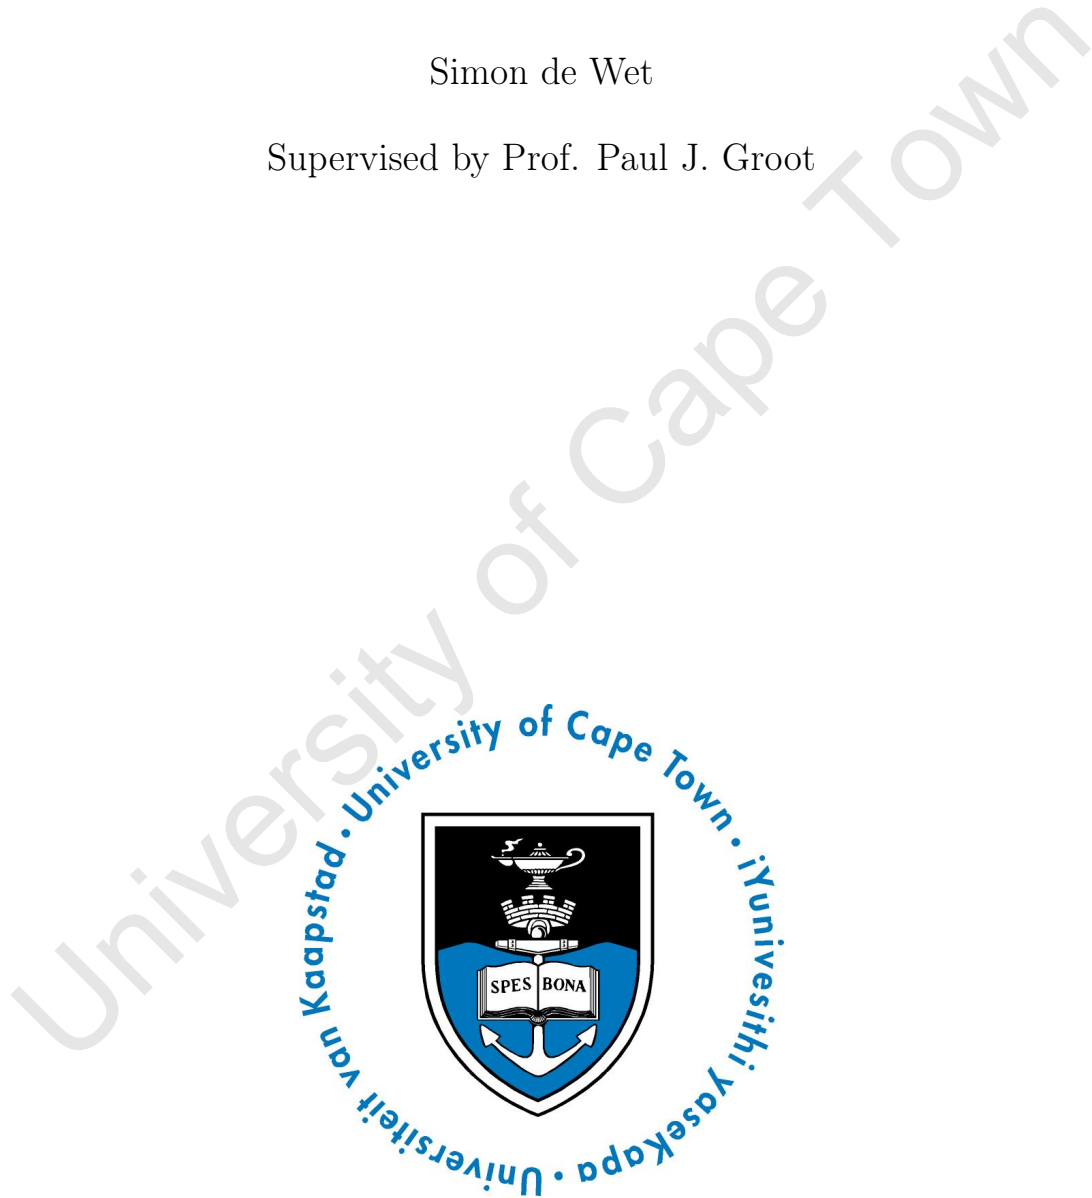


A thesis presented for the degree of
Master of Science

Transients in the errorbox of GW190814

Simon de Wet

Supervised by Prof. Paul J. Groot



Department of Astronomy
University of Cape Town
July 2020

The copyright of this thesis vests in the author. No quotation from it or information derived from it is to be published without full acknowledgement of the source. The thesis is to be used for private study or non-commercial research purposes only.

Published by the University of Cape Town (UCT) in terms of the non-exclusive license granted to UCT by the author.

Contents

1	Introduction	1
1.1	GW190814	2
1.2	NSBH binaries	3
1.2.1	Scenarios for NSBH merger and post-merger	4
1.2.2	Mass ejection and outflows	6
1.2.3	Kilonovae	7
1.2.4	NSBH KN simulations	9
2	MeerLICHT	12
2.1	Telescope	12
2.2	Software pipeline	14
2.2.1	Pre-processing	14
2.2.2	Astrometry and photometry	16
2.2.3	Image subtraction and transient detection	16
2.2.4	Data products	17
2.3	Science goals	18
3	The data	19
3.1	GCN alert	19
3.2	Observations with MeerLICHT	19
3.3	The gradient problem	20
3.3.1	Re-reducing the raw data with good flats	23
3.3.2	Re-reducing science images without gradients	25
3.3.3	The gradient problem solved	26
3.4	The (good) reduced data	28
3.5	Finding transients	29
3.5.1	Identifying matching pairs	30
3.5.2	Vetting of matching pairs	32
3.5.3	Identifying moving objects	34
3.5.4	Light curves	34
3.5.5	Determining if transients were unreported	35
3.6	Transients reported by other groups	38
4	Limits on the electromagnetic counterpart to GW190814	40
4.1	Comparison with AT2017gfo-like KN	40
4.2	POSSIS kilonova simulations	43
4.2.1	Covered probability in mass ejecta-viewing angle parameter space	45
4.2.2	Colour-evolution of KN models	47
4.2.3	Colour-colour evolution of KN models	49

5 Conclusion	52
5.1 Future work	53
A	58
A.1 Light curves of transients in table 3.3	58

Chapter 1

Introduction

We are now firmly in the era of multi-messenger astronomy. The detection of the first binary black hole (BBH) merger in GW150914 [1] opened up the era of gravitational wave astronomy, with a further 9 such mergers being detected during the first two observing runs (O1 and O2) of the LIGO Scientific and Virgo Collaborations (LVC). The first – and currently only – multi-messenger source was detected during O2 and was caused by the merger of two neutron stars in a binary system (BNS) [2]. The electromagnetic (EM) counterparts to GW170817 [3] were observed across the EM spectrum by numerous observing facilities, with implications across a vast range of scientific disciplines. Optical/near-infrared observations demonstrated that the emission was due to a kilonova powered by the radioactive decay of r-process material produced during the merger. For the first time short gamma-ray bursts were convincingly linked to BNS mergers, as observed in GRB170817A [3]. The third LVC observing run (O3) began 2019 April 1 and concluded 2020 March 27. The signal from GW190425 [4] was likely caused by the coalescence of two neutron stars, with the system having a larger total mass than any currently known BNS system. Furthermore, the detection of GW190412 revealed the first BBH merger with a clearly unequal mass ratio of $q = m_2/m_1 = 0.28$ along with significant higher-multipole gravitational radiation [5].

The subject of this thesis is an event that occurred during this same observing run – GW190814. Its initial significance was that it had an extremely low false-alarm rate and was localised to a relatively small region of the sky. The source classification of a neutron star black hole (NSBH) merger made the existence of an EM counterpart particularly promising. Among many other facilities, the MeerLICHT optical telescope in Sutherland performed an intense follow-up campaign, observing more than 95% of the probability localisation each night for over a week in three optical bands (u , g , and i), with our initial observations being some of the earliest of any group – the first observation beginning almost exactly 2 hours after the GW detection. No potential counterpart was found during follow-up observations by us nor any other group.

In the following sections of this chapter an overview of GW190814 and NSBH mergers and their expected optical emission is given. In chapter 2 the optical telescope MeerLICHT – our data-collecting instrument – is introduced. In chapter 3 the search for new transients in MeerLICHT data is described, while in chapter 4 it is demonstrated how our limiting magnitudes resulting from the non-detection of any counterpart could be used to constrain possible kilonova and merger models.

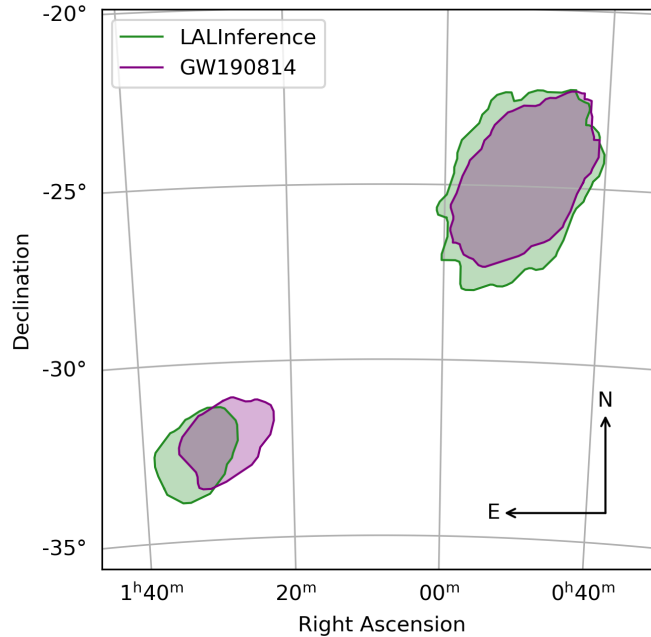


Figure 1.1: Until the recent LVC publication on GW190814 [6], the LALInference sky map (green) was used for planning follow-up observations, as detailed in section 3.1 and section 3.2. The 90% credible region encompassed a sky-area of 23 deg^2 at a distance of $267 \pm 52 \text{ Mpc}$. The sky map accompanying the recent paper (purple) encompassed an area of 18.5 deg^2 at a distance of $241_{-45}^{+41} \text{ Mpc}$. The two regions are shown for comparison.

1.1 GW190814

On 2019 August 14 at 21:11:00 UTC, GW190814 was detected in data from LIGO Livingston and Virgo with a signal-to-noise (S/N) ratio of $\simeq 25$. Initially classified as a MassGap¹ merger, reanalysis of the data from all three detectors updated the source classification to that of a NSBH merger at a distance of $267 \pm 52 \text{ Mpc}$ with a sky-area of 23 deg^2 at the 90% credible level (see fig. 1.1). Notices and Circulars through NASA’s Gamma-ray Coordinates Network (GCN) provided the electromagnetic community with sky maps from which follow-up observations could be planned. Numerous groups undertook surveys – both galaxy-targeted and wide-field – with no viable counterpart being found.

Further analysis of the GW190814 data in the recent publication by the LVC (see [6]) revealed that the 90% probability region encompassed 18.5 deg^2 at a distance of $241_{-45}^{+41} \text{ Mpc}$ and was caused by a compact binary coalescence with the most unequal mass ratio yet measured in gravitational waves, with $q = 0.112_{-0.009}^{+0.008}$. Furthermore, the secondary component of the binary had a mass that would make it either the lightest BH or heaviest NS ever discovered, with $m_2 = 2.59_{-0.09}^{+0.08} M_{\odot}$. The primary component was confidently declared to be a BH with a mass $m_1 = 23.2_{-1.0}^{+1.1} M_{\odot}$ and dimensionless spin tightly constrained to $\chi_1 \leq 0.07$, while the merger-rate density for such systems was calculated as $7_{-6}^{+16} \text{ Gpc}^{-3} \text{ yr}^{-1}$ [6]. In their paper the mass of the secondary component is compared to estimates of the maximum NS mass from a number of studies that consider a

¹A MassGap merger is a classification by the LVC that at least one of the components has a mass between 3 and 5 solar masses.

variety of NS equations of state (EOS), mass measurements of the known NS population, and constraints on the mass of the merger remnant of GW170817. From these comparisons they infer that GW190814 was most likely not due to the merger of a NS and BH, though importantly they do not rule it out entirely due to the current uncertainties in the maximum NS mass. Furthermore, they leave open the possibility that the secondary component is an unknown exotic compact object such as a boson star [7] or gravastar [8].

There is much uncertainty surrounding the origin of systems having such an unequal mass ratio. It is unknown if the lack of observed NSs and BHs in the MassGap range is physical or due to selection biases. Theoretical models used in calculating the masses of compact remnants are subject to challenges arising from the complex physics of supernova explosions and the details of stellar evolution [6]. Some of them have even been developed to reproduce this perceived mass gap. For this reason the authors recommend the adjusting of models to allow for such masses. Population synthesis simulations predict that most BBH mergers have $q > 0.5$, while for NSBH mergers the mass-ratio distribution peaks at $q \approx 0.2$ with more massive NSs ($\gtrsim 1.3M_{\odot}$) paired with fairly low mass BHs ($\lesssim 15M_{\odot}$) being favoured [6]. Additionally, these models tend to favour more symmetric masses in isolated binaries as a result of common envelope and mass transfer. The highly unequal mass ratio of GW190814 therefore poses a challenge to current population synthesis models.

Another suggestion is that such systems may be of dynamical origin, particularly in young star clusters where the number of compact-object binaries having unequal mass ratios may be much higher than in globular clusters. The reasoning is that compact objects of similar mass are paired up more favourably in dense stellar environments like globular clusters, with models showing that most BBH systems have $q \simeq 1$ (as in [9]), and that the merger rate of NSBH systems is low since the BHs dominate the core of the clusters, preventing NSs and BHs from interacting. There is also the possibility that GW190814-like binaries were formed from hierarchical systems where a higher mass BH or supermassive BH (in galactic centres) acts as a perturber that induces the merging of two GW190814-like components. In fact, along with young star clusters, disks of gas surrounding supermassive BHs in active galactic nuclei are seen as the most promising sites for the formation of systems akin to GW190814 [6].

In summary, the novel discovery by the LVC of a compact binary merger having such an unequal mass ratio with the secondary component squarely within the MassGap has challenged population synthesis models and existing assumptions on the lightest BH or heaviest NS. Though it is unlikely that the secondary component was a NS due to existing estimates on the maximum NS mass, the uncertainties in these estimates mean that it is still possible that GW190814 was caused by the coalescence of a NS and BH. For this reason, we continue subsequent analysis assuming that the source of GW190814 was a NSBH merger.

1.2 NSBH binaries

Mergers of binaries containing two neutron stars (NSNS) or a neutron star and black hole (NSBH) are exciting sources of gravitational waves for ground-based GW detectors such as Advanced LIGO/Virgo [10][11] and KAGRA [12]. Additionally, mergers of these binaries are promising sites of r-process nucleosynthesis due to the expected abundance of neutron-rich material ejected during the merger. The subsequent radioactive decay of the heavy r-process nuclei is predicted to power strong EM emission in the form of a kilonova (KN) – as seen in the multi-wavelength observations of GW170817 [2]. Numerical-relativity




Figure 1.2: The two scenarios for a NSBH merger. Either the NS is completely swallowed by the BH resulting in no EM emission, or the NS is tidally disrupted. Figure from [13].

simulations taking into account detailed microphysical processes, neutrino radiation and magnetohydrodynamics (MHD) have been used to study the merger phenomena and predict potential observational features of these mergers, and have shown that the mass of dynamically ejected matter relies heavily on the neutron star equation of state (EOS), total mass and mass ratio of the binary, and the BH spin (for NSBH binaries) [13].

There are important differences between NSNS and NSBH mergers. It is beyond the scope of this work to describe NSNS mergers in detail, hence we will focus only on NSBH mergers in the following subsections. In general, the trends are that higher BH spins and more equal binary mass ratios produce larger ejecta masses and therefore more luminous kilonovae, while ejecta with smaller electron fractions (greater neutron richness) synthesise more heavy elements, resulting in redder kilonovae due to the higher opacity of these heavy elements.

1.2.1 Scenarios for NSBH merger and post-merger

There are essentially two possible merger scenarios for a NS and BH, as depicted in fig. 1.2. Either the NS is completely swallowed by the BH resulting in little to no mass remaining outside the BH, or the NS is tidally disrupted with the BH remnant surrounded by matter consisting of two components: a bound disk, and unbound dynamical ejecta. No EM emission is expected in the first scenario (a direct plunge) while a variety of emission mechanisms are expected in the latter [13].

Tidal disruption occurs when the tidal force of the NS due to the BH is greater than the self-gravity of the NS given by the approximate condition

$$\frac{GM_{\text{BH}}R_{\text{NS}}}{r^3} > \frac{GM_{\text{NS}}}{R_{\text{NS}}^2}, \quad (1.1)$$

where G is the gravitational constant, M_{BH} and M_{NS} are the BH and NS masses respectively, R_{NS} is the NS radius, and r is the orbital separation. The tidal and self-gravity forces are equal at the “tidal distance” where the orbital radius

$$r = d_{\text{tidal}} \sim \left(\frac{M_{\text{BH}}}{M_{\text{NS}}} \right)^{1/3} R_{\text{NS}}. \quad (1.2)$$

The innermost stable circular orbit (ISCO) of the BH is important here as it determines whether the orbiting NS is swallowed by the BH before tidal disruption occurs. The




Figure 1.3: The location of the ISCO (red line) and other special orbits are shown as a function of the dimensionless spin parameter a for a $10M_{\odot}$ BH. For a generic BH having mass M_{BH} , $R_{\text{ISCO}} = 6GM_{\text{BH}}/c^2 = 3R_{\text{S}}$ when the spin is zero ($a = 0$), and $R_{\text{ISCO}} = 3GM_{\text{BH}}/c^2 = 3R_{\text{S}}/2$ for a maximally spinning BH ($a = 1$). R_{S} is the Schwarzschild radius. Prograde spins – when the BH angular momentum and orbital angular momentum are aligned – correspond to positive a values and are in the region to the right of the vertical dashed line. Figure from [15].

ISCO is the closest distance from the centre of a Kerr BH at which a test particle will orbit the BH in a circular orbit [14]. Within this radius, orbits are unstable and particles spiral rapidly into the centre of the BH. The radius of the ISCO depends only on two properties – the BH mass and the BH spin. The Kerr solution to Einstein’s General Theory of Relativity shows that spacetime rotates like a vortex outside of a spinning BH, dragging matter and light into rotation around the BH [15]. Figure 1.3 demonstrates how the radius of the ISCO varies as a function of the dimensionless spin parameter $a = cJ/GM_{\text{BH}}^2$ (where J is the BH angular momentum). We see that BHs with larger spins have smaller R_{ISCO} . Therefore tidal disruption will occur if

$$d_{\text{tidal}} > R_{\text{ISCO}}, \quad (1.3)$$

or using the approximate relation from eq. (1.2)

$$\left(\frac{M_{\text{BH}}}{M_{\text{NS}}}\right)^{1/3} R_{\text{NS}} > R_{\text{ISCO}}. \quad (1.4)$$

For NSBH mergers we consider only prograde orbits – that is, the orbital angular momentum is aligned with the BH spin angular momentum – as anti-aligned spins result in larger R_{ISCO} (see fig. 1.3), making it more likely that the NS is swallowed before tidal disruption. Following eq. (1.4), we notice that fixing the BH mass and increasing the NS mass requires larger BH spins to ensure there is leftover mass to produce EM emission.

Higher mass NSs are more compact and therefore have smaller tidal radii, meaning larger BH spins are required to satisfy the inequality. Similarly, increasing the BH mass results in larger gravitational radii, requiring higher spins to satisfy the inequality [16]. Equation 6 in the Shibata and Hotokezaka review (2019) [13] summarises these competing factors in the conditional inequality

$$\left(\frac{\xi}{6}\right)^{-3/2} \left(\frac{Q}{7}\right)^{-1} \left(\frac{R_1}{10Gc^{-2}M_{\text{NS}}}\right)^{3/2} > 3.25, \quad (1.5)$$

where ξ is a dimensionless measure of the BH spin ranging from 1 for maximally spinning BHs (where $a = 1$ in fig. 1.3) to 6 for non-spinning BHs (where $a = 0$ in fig. 1.3), Q is the binary mass ratio (with $Q = M_{\text{BH}}/M_{\text{NS}} > 1$), and R_1 is the semi-major axis of the NS. Systems with larger BH spin (smaller ξ) and smaller mass ratios (smaller Q) are favoured for tidal disruption to occur with a greater mass left outside of the BH post-merger.

1.2.2 Mass ejection and outflows

Numerical relativity simulations have shown that the bulk of the remnant NS material forms a gravitationally bound accretion disk (up to $0.5M_{\odot}$) surrounding the spinning BH, while some of the material is dynamically ejected – typically $\sim 20\%$ of the total mass left outside of the BH, which amounts to $0 - 0.1M_{\odot}$ of dynamical ejecta [13]. The average velocity of the dynamical ejecta is typically around $0.3c$ with the high-velocity matter coming from near the BH event horizon. Shock heating and neutrino irradiation can increase the electron fraction Y_e of the ejecta through weak interaction processes: electron-positron pair production and subsequent neutron capture (via $n + e^+ \rightarrow p + \bar{\nu}_e$) in the case of shock heating, and neutron and proton absorption of neutrinos (via $n + \nu_e \rightarrow p + e^-$ and $p + \bar{\nu}_e \rightarrow n + e^+$) with neutrino irradiation. For NSBH systems the ejecta is mainly launched by tidal torque instead of shock wave-induced ejection, while strong neutrino irradiation only plays a significant role in the case of massive neutron star formation (a strong neutrino emitter) for some NSNS mergers. The result is that weak interaction processes are largely absent in NSBH mergers, and thus the Y_e is always low ($\lesssim 0.1$) [13].

The accretion disk has differential rotation and should possess magnetic fields from the NSs magnetic field, leading to magnetohydrodynamic (MHD) turbulence induced by magnetorotational instability (MRI) [13]. Turbulent viscosity is enhanced and thus angular momentum is transported outwards, causing matter in the inner part of the disk to fall into the BH and matter in the outer part to expand equatorially. In the early stages of the disk evolution, thermal energy from viscous heating leads to increased neutrino emission. If the disk density is high enough, neutrinos are not allowed to escape freely due to the increased optical depth to neutrinos, reducing their emissivity. In the later stage of disk evolution, the decreasing disk density caused by mass infall into the BH and disk expansion from viscous angular momentum transport leads to a decline in the optical depth to neutrinos. This allows the viscous heating to be used for the adiabatic expansion of the disk, leading to the eventual ejection of disk matter. Simulations show that $\sim 20\%$ of the disk mass is ejected with a characteristic velocity $\lesssim 0.1c$, slower than the dynamical ejecta. As for the Y_e , a wider range of values (0.05–0.3) are plausible due to the viscous heating and neutrino irradiation [13].

Accretion of mass in the disk onto the central BH is also expected to launch a relativistic jet via the Blandford-Znajek process [17], in which poloidal magnetic fields penetrating the BHs event horizon lead to energy extraction of the spinning BHs rotational kinetic

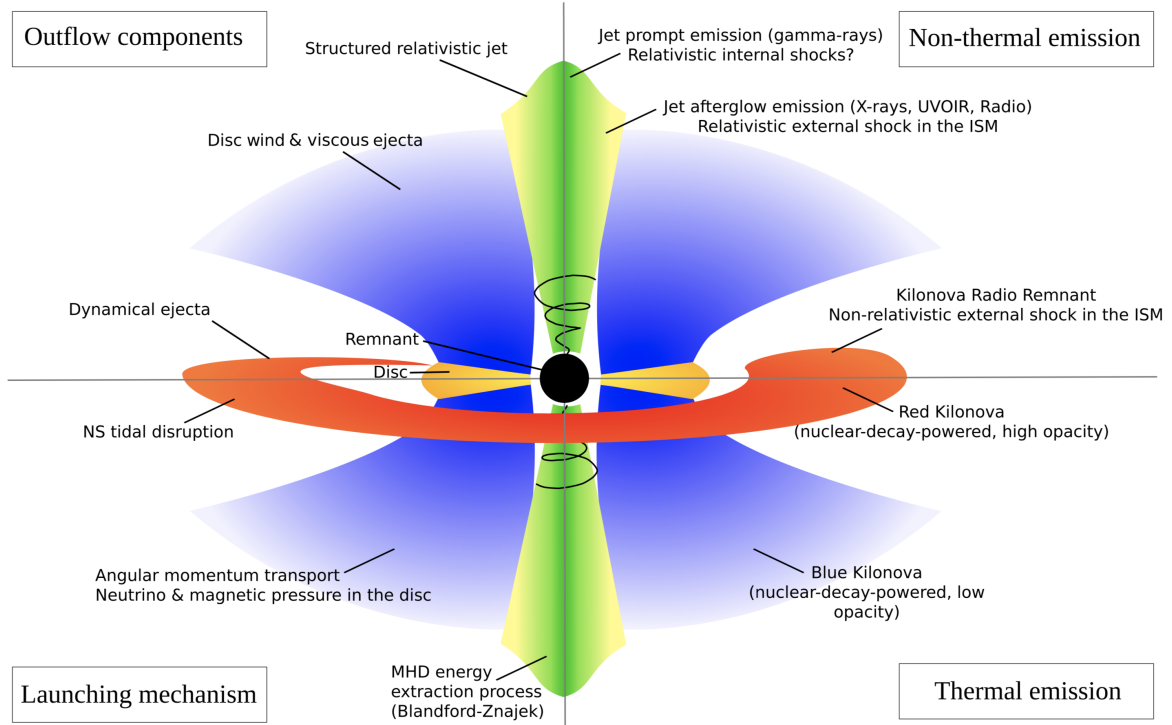


Figure 1.4: Figure showing the various outflows and emission mechanisms for a NSBH merger, taken from [16]. The dynamical ejecta (red) and disk wind ejecta (blue) power the KN, which should be observable in optical and infrared wavelengths.

energy. Additionally, the annihilation of high-energy neutrinos from the disk could lead to γ -rays being produced via the annihilation of electron-positron pairs resulting from neutrino pair annihilation. Both these mechanisms could be responsible for powering the relativistic jet and generating a short γ -ray burst (sGRB). The interaction of the jet with the interstellar-medium (ISM) is likely to produce a shock that generates synchrotron radiation from accelerated electrons. This afterglow could be observable in the X-ray to radio bands.

Figure 1.4 provides a schematic of the outflow components resulting from the merger, along with their launching mechanism and expected emission. The dynamical ejecta (in red) and disk ejecta (in blue) are together responsible for the KN, observed as thermal emission. The relativistic jet (green) could be observed as a sGRB along with its afterglow.

1.2.3 Kilonovae

The type of emission from a NSBH merger that is most relevant to this thesis is a kilonova (KN), since it is expected to be visible at optical and infrared wavelengths and thus could be observable with MeerLICHT. The emission from a KN is powered by the radioactive decay of heavy r-process nuclei synthesised in the neutron-rich merger ejecta. R-process nucleosynthesis refers to the rapid neutron-capture process in which heavy seed nuclei capture successive neutrons before radioactive β^- decay ($n \rightarrow p + e^- + \bar{\nu}_e$) can occur, and is thought to be responsible for the creation of about half of the atomic nuclei heavier than iron [18]. Core-collapse supernovae and neutron star mergers are seen as the most likely sites of the r-process, though it is believed by some researchers [18] that the latter site could be more dominant.

The neutron richness (characterised by the electron fraction Y_e) is a key quantity in determining the abundance patterns of r-process elements, and therefore has a strong influence on the opacity of the EM emission. Research [19] has shown that neutron-rich ($Y_e \lesssim 0.1$) ejecta are responsible for synthesising heavy nuclei ($A \gtrsim 120$), while higher electron fractions ($Y_e \gtrsim 0.25$) are associated with the synthesis of the lighter ($70 \lesssim A \lesssim 120$) elements. Since the ejecta (dynamical and disk) of NSBH mergers are expected to have low electron fractions, heavier elements are dominantly synthesised.

The radioactive decay of r-process nuclei proceeds through three decay channels: β decay, α decay, and spontaneous fission. The kinetic energy of the decay products is transferred to the merger ejecta as thermal energy. It is the thermal emission of the radioactively heated merger ejecta which is visible as a KN. During β decay, unstable nuclei move towards the nuclear valley of stability while maintaining their atomic mass number and releasing an electron with an amount of kinetic energy. The energy-releasing rate for these electrons per unit mass is usually [13] given by

$$\dot{q}_e(t) \approx 3 \times 10^9 \text{ erg s}^{-1} \text{g}^{-1} \left(\frac{t}{1 \text{ day}} \right)^{-4/3}. \quad (1.6)$$

During α decay, an α particle (helium nucleus) having 5–10 MeV of kinetic energy is emitted by an initially neutron-rich nucleus that has increased its proton fraction through successive β decays. A number of α and β decays can occur before a stable nucleus ($A < 210$) is reached. It is possible for α heating to dominate the β heating if the mass of α -unstable elements is above $10^{-3}M_\odot$. Nuclides with atomic numbers greater than uranium's are all unstable to radioactive decay. Those with $A \gtrsim 240$ can undergo spontaneous fission in which the fission fragments can have 100–200 MeV of kinetic energy. Though the energy released during an α decay and a spontaneous fission is much higher than for a β decay, it is β decay and subsequent γ -ray emission that are responsible for most of the thermal energy of the KN.

The opacity plays an important role in determining the spectrum and light curves of a KN. The composition of the ejecta determines the opacity as heavier r-process nuclei (lanthanides and actinides) are open f -shell elements with many more excited states having low excitation energies than lighter d -shell elements. Lanthanide-rich ejecta thus have higher opacities than lanthanide-free ejecta. The effect of the composition on the opacity and therefore the light curves and spectra are demonstrated in fig. 1.5. Light r-process materials (where $X_{lan} \leq 10^{-4}$) have low opacities ($\kappa \leq 1 \text{ cm}^2 \text{ g}^{-1}$) and so they radiate primarily in optical light that fades quickly – within days. The heavy r-process ejecta (with $X_{lan} \geq 10^{-2}$) have higher opacities ($\kappa \approx 10 \text{ cm}^2 \text{ g}^{-1}$) and thus a longer diffusion time. Optical light is obscured by the lanthanide lines and so most of the radiation is emitted in the infrared, with a longer duration bolometric light curve [20].

The effect is summarised in the following equations [13] for a typical light curve of a KN produced by merger ejecta having mass M_{ej} , velocity v_{ej} , and opacity κ . The peak occurs at time² t_p

$$t_p \approx \sqrt{\frac{\kappa M_{ej}}{4\pi c v_{ej}}} \approx 10 \text{ days} \left(\frac{\kappa}{10 \text{ cm}^2 \text{ g}^{-1}} \right)^{1/2} \left(\frac{M_{ej}}{0.04 M_\odot} \right)^{1/2} \left(\frac{v_{ej}}{0.1c} \right)^{-1/2}, \quad (1.7)$$

²Derived by setting the photon diffusion time equal to the expansion time as in [21].

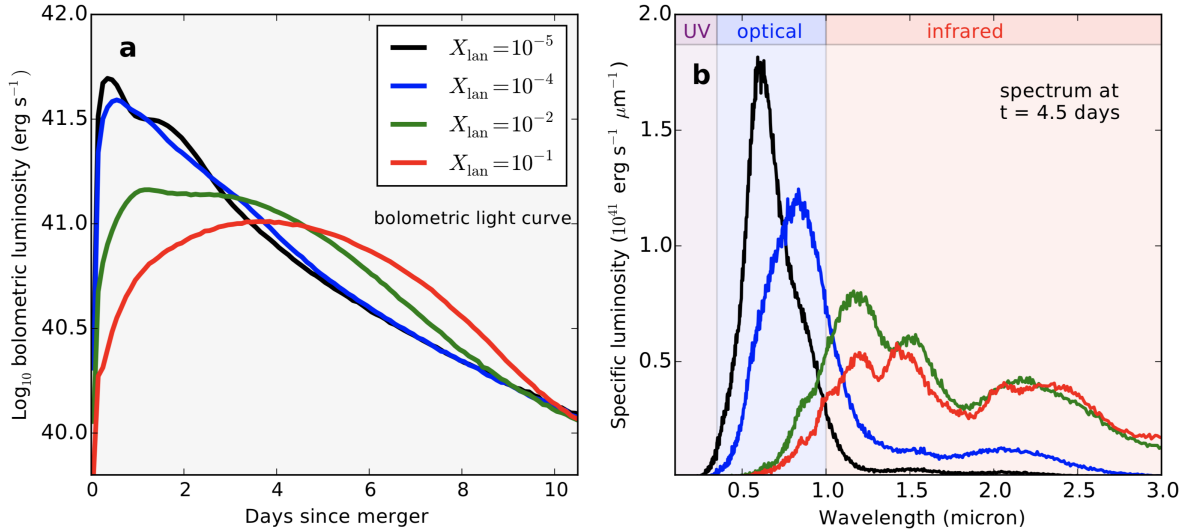


Figure 1.5: The effects of the lanthanide fraction on the bolometric luminosity and spectra of model KNe assuming an ejecta mass of $0.05M_{\odot}$ with velocity of $0.2c$. Figure from [20].

while the bolometric luminosity is approximately equal to the thermal heating rate at t_p

$$L_{\text{bol}}(t_p) \approx \dot{Q}_{\text{th}}(t_p) = M_{\text{ej}} \times \dot{q}_{\text{th}}(t_p) \approx 4 \times 10^{40} \text{ erg s}^{-1} \left(\frac{t_p}{10 \text{ days}} \right)^{-1.3} \left(\frac{M_{\text{ej}}}{0.04M_{\odot}} \right) \quad (1.8)$$

The effective temperature at peak luminosity is approximately derived from the Stefan-Boltzmann law as

$$T_{\text{eff}}(t_p) \approx \left(\frac{L_{\text{bol}}(t_p)}{4\pi\sigma v_{\text{ej}}^2 t_p^2} \right)^{1/4} \left(\frac{v_{\text{ej}}}{0.1c} \right)^{-1/2} \left(\frac{t_p}{10 \text{ days}} \right)^{-1/2}, \quad (1.9)$$

where σ is the Stefan-Boltzmann constant. NSBH mergers – usually having a low Y_e – are more likely to have lanthanide-rich ejecta and thus “redder” spectra that peak later than lanthanide-free ejecta, provided the mass, velocity and specific heating rate are held constant.

1.2.4 NSBH KN simulations

In section 1.2.1 we presented the general conditions necessary for tidal disruption to occur with mass left outside of the BH. Along with the binary mass ratio and BH spin, the tidal deformability of the NS (Λ_{NS}) plays an important role in its fate, given by

$$\Lambda_{\text{NS}} = \frac{2}{3} k_2 C_{\text{NS}}^{-5}, \quad (1.10)$$

where $C_{\text{NS}} = GM_{\text{NS}}/(R_{\text{NS}}c^2)$ is the NS compactness and k_2 is the tidal Love number (as defined in [16]). Tidal disruption is favoured for NSs with higher Λ_{NS} . The EOS sets the radius and therefore compactness of the NS, meaning that “stiffer” EOSs result in NSs with a higher Λ_{NS} , and so tidal disruption is more likely than for “softer” EOSs.

Using their semi-analytical models, Barbieri *et al.* [16] investigate the dependence of the binary mass ratio and BH spin on the EM emission of a NSBH merger while keeping the NS mass and Λ_{NS} constant – with $M_{\text{NS}} = 1.4M_{\odot}$ and $\Lambda_{\text{NS}} = 330$. The expected disk




Figure 1.6: The disk (left panel) and dynamical (right panel) ejecta mass as a function of BH mass and spin combinations is shown, from [16]. The squares represent KN models having constant spin and varying mass, while circles represent models with constant mass ($5.3M_{\odot}$) and varying spin. Light curves corresponding to each model are shown in fig. 1.7. For BH mass and spin combinations in the white region, the NS plunges directly into the BH with no NS material left outside the BH.

and dynamical ejecta masses are shown as a function of the binary mass ratio and BH spin in fig. 1.6. We see confirmation of the general conditions for tidal disruption that were presented in section 1.2.1 – smaller binary mass ratios (smaller BH mass) and greater BH spins produce more ejecta. Figure 1.7 presents light curves in the r (657 nm) and K (2143 nm) bands at a luminosity distance of 230 Mpc and viewing angle of 30° for two scenarios: varying the BH mass and keeping the spin constant (left panel), and varying the spin while keeping the BH mass constant (right panel). The light curves demonstrate that lower BH masses and higher spins result in brighter KNe, with the peak shifting to later times. Importantly, the light curves show a high level of degeneracy. This means that inferring the intrinsic parameters (M_{BH} and χ_{BH}) from the light curves alone is impossible. Knowledge of the BH and NS masses and the luminosity distance from the GW signal could be used to break this degeneracy as demonstrated in [16].

The dependence of the NS properties on the EM emission of NSBH mergers is explored in the paper of Barbieri *et al.* [22]. Similar to fig. 1.6 and fig. 1.7, they compute the ejecta mass and produce light curves for models with varying NS tidal deformability and NS mass, but fixing the BH mass and spin. For a fixed NS mass, the trend is that increasing Λ_{NS} leads to more massive ejecta and therefore brighter KNe. For a fixed Λ_{NS} , increasing the NS mass produces the same result. There are exceptions, however, for certain combinations of Λ_{NS} and M_{NS} , where the dynamical ejecta mass decreases instead of increases for increasing NS mass. Since the disc mass always increases with increasing NS mass at a fixed Λ_{NS} , it is possible that the light curves with higher NS mass are dimmer than those with lower NS mass at early times, when the dynamical ejecta dominates. Again, degeneracies in the light curves exist which could be broken by a combined EM and GW analysis.

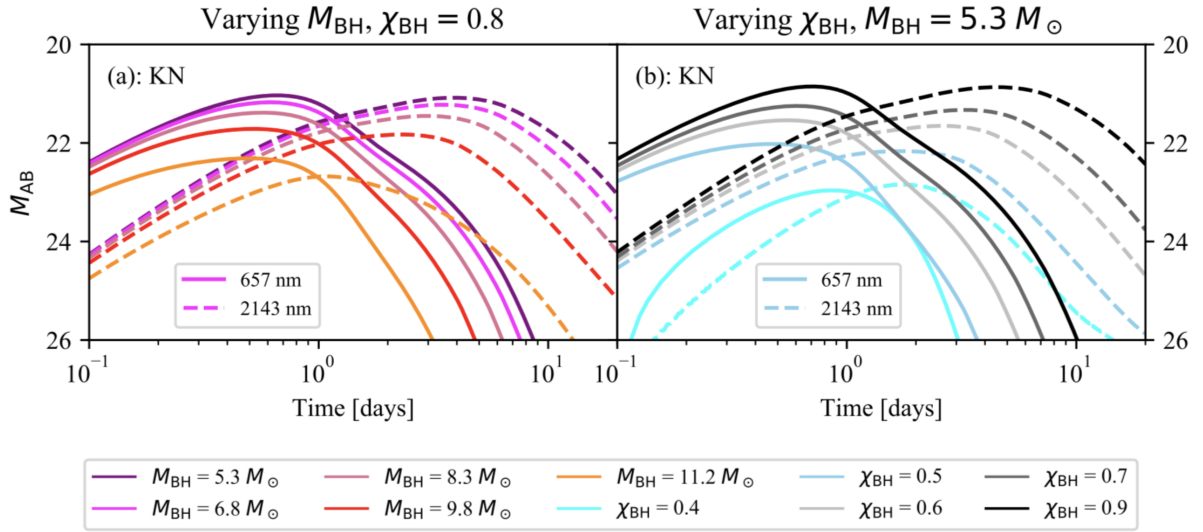


Figure 1.7: Light curves in the r (smooth lines) and K (dashed lines) bands for varying BH mass (left panel) and varying BH spin (right panel). The colours of each line correspond to a particular $M_{BH} - \chi_{BH}$ combination shown in fig. 1.6. The figure and legend are from figure 6 in [16].

Chapter 2

MeerLICHT

Answering fundamental science questions depends on having an instrument through which the sky can be measured. For astronomers such an instrument is usually a telescope, and for this thesis in particular the MeerLICHT optical telescope was our data-collecting instrument.

MeerLICHT is a wide-field and fully robotic optical telescope located at the South African Astronomical Observatory site near Sutherland. It was designed and built with the intention of being a prototype for the BlackGEM array and being paired with the MeerKAT radio telescope to provide a simultaneous optical and radio view of the transient sky. Inaugurated in May of 2018, the name MeerLICHT is Dutch for *more light* and refers both to its pairing with MeerKAT and its wide field of view.

2.1 Telescope

Though they have different science goals, the MeerLICHT and BlackGEM telescopes are identical in design [23], using a modified Dall-Kirkham Cassegrain optical system where the spherical secondary mirror is suspended above the 65 cm parabolic primary mirror as shown in fig. 2.1. The carbon fibre frame housing the optical system is fixed to an equatorial mount, with the counterweight housing the glycol-cooled electronics. Baffling below the secondary mirror shields the detector from stray light. A set of three lenses in the lens barrel correct for atmospheric dispersion and an off-axis coma introduced by the telescope design. The resulting light passes through one of six filters before arriving at the cryostat-cooled CCD detector – a single large chip consisting of 10560 by 10560 pixels. This results in a 2.7 deg^2 field of view with a pixel scale of $0.56''/\text{pixel}$. The exceedingly stable point spread function from the optical design means that the image quality degrades by less than $0.3''$ from the optical axis to the corner of the field [24].

The filter wheel consists of six filters which can be changed in less than 3 seconds: five SDSS-type filters – u , g , r , i , z – and an additional wide filter – q – similar to $g+r$. The wavelength range for each filter is shown in table 2.1. MeerLICHT is seeing limited due to the atmospheric conditions at the site, such that a default integration time of 60 seconds per exposure reaches a sufficient depth. The wide field of view, consistently high sensitivity across the image plane, and the ability to be operated fully robotically makes these telescopes perfectly suited for rapid follow-up of gravitational wave events in addition to exploring the multi-colour explosive Universe [24].




Figure 2.1: Model of the MeerLICHT telescope demonstrating the main components of its design. Figure from [23].

Filter	Wavelength range (nm)
u	350 – 410
g	410 – 550
r	563 – 690
i	690 – 840
z	840 – 990
q	440 – 720

Table 2.1: MeerLICHT filters and their approximate wavelength ranges.

2.2 Software pipeline

The software pipeline for reducing raw images from the telescope is largely based on SkyMapper’s [25], with modifications and additions. Written in Python the software consists of two components: the first is known as BlackBOX¹ which performs standard CCD reduction tasks on the raw science images, including gain-, crosstalk-, overscan-, bias-, and flatfield-corrections; the second is ZOGY² which is used for identifying sources, performing astrometry, and finding transients through image subtraction.

2.2.1 Pre-processing

In order to accurately measure the flux (photons per unit time interval) arriving at the CCD, a series of pre-processing steps must be undertaken to correct for the various sources of noise introduced by the surroundings and the CCD itself. MeerLICHT’s CCD is divided into 16 equal-sized channels (see fig. 2.2) for reading out the charges generated by photons falling onto the semi-conductor pixels. Because MeerLICHT’s CCD is placed in a 170 K cryostat, the thermal noise is relatively low. For this reason, dark frames are not used during processing though they are still taken at the beginning and end of each night for monitoring. The other processing steps are explained below:

Gain

The charge produced by photons striking a pixel is measured by the CCD in analog to digital units (ADUs). In order to determine the number of electrons created by incident photons, it is necessary to determine the conversion factor between the CCD-measured counts (ADUs), and the actual number of electrons produced by incident photons. This factor is known as the gain and it is determined separately for each of the 16 channels.

Crosstalk

A very bright star can leave ghost images at the same relative pixel location in horizontally neighbouring channels – a phenomenon known as cross-talk. During commissioning, measurements of the amount of cross-talk between channels as a function of exposure levels were taken. This information is used to correct for cross-talk in science images.

Overscan and bias

To prevent negative counts being detected, an offset is introduced by the electronics which can vary across the CCD. Overscan is the bulk offset of the pixel values. Since it can vary according to temperature and time, it is calculated for each channel by calculating the mean level in the artificially added overscan regions (visible in fig. 2.2). Once subtracted from the pixels, the overscan regions are trimmed out of the image. Bias images are then used to correct for the remaining pixel-to-pixel offset across the image plane. These are taken with a closed shutter and null exposure time, and are also corrected for overscan. Multiple images are stacked to create a master bias, though they are currently not used on science images since overscan correction appears to eliminate most of the bias.

¹Developed by Paul Vreeswijk and Kerry Patterson, the source code is available at <https://github.com/pmvreeswijk/BlackBOX>.

²A python implementation of proper image subtraction based on the paper by Zackay *et al.* [26]. Source




Figure 2.2: A raw q band image taken by MeerLICHT on the 24th of August, 2019. The field has ID 4878 and contains the photogenic Sculptor Galaxy (NGC 253) visible towards the left edge. The horizontal overscan regions can be seen between the two rows of channels, while the vertical overscan regions are located immediately to the right of each channel.

Flatfielding

Flatfields are used to correct for differences in sensitivity between pixels and also possible sources of uneven illumination – for instance, vignetting. Observing a uniformly lit surface ensures that each pixel receives the same amount of light. MeerLICHT takes flatfields of the sky during twilight (dawn and dusk) in each of its filters. The individual flatfields are stacked to create a master flatfield with an improved signal-to-noise ratio. Usually, flatfields in a particular filter that were taken between 3 days before and 3 days after the particular observing night are combined to form the master flat for that night. All science images from that night are then divided by the master flat. Figure 3.4 provides examples of flatfield images in multiple filters.

2.2.2 Astrometry and photometry

Before sources can be identified in an image, a mask containing pixels that should not be used for science processing is created. Cosmic rays, satellite trails, bad pixels, edge pixels, and saturated pixels are identified in the image and added to its mask. SExtractor [27] is used for source identification in MeerLICHT images. Those sources with at least a 5σ significance above the background will end up in the full source catalogue.

MeerLICHT uses a calibration catalogue consisting of a selection³ of Gaia DR2 [30] stars to perform astrometric and photometric calibration on images. To convert pixel coordinates to sky coordinates, Astrometry.net [31] is used to compare and match the pixel coordinates of the 1000 brightest sources in an image with their positions in MeerLICHT’s calibration catalogue. The resulting WCS solution uses the International Celestial Reference System (ICRS) frame with a reference epoch of J2015.5 – the same used in the Gaia DR2 data.

MeerLICHT utilises the AB magnitude system, where the apparent magnitude m of a source with flux F is given by

$$m = -2.5 \log\left(\frac{F}{F_{AB}}\right), \quad (2.1)$$

with F_{AB} equal to 3631 Jy. The MeerLICHT magnitudes of stars in the calibration catalogue are determined by fitting spectral models to the photometric data from a large number of UV, optical and NIR surveys and calculating the flux MeerLICHT would observe in each of its filters. Since the point-spread function (PSF) generally varies across an image, the PSFEx [32] software is used to calculate the PSF at any position, allowing for accurate measurement of the fluxes of individual sources. The magnitudes in the calibration catalogue are then used to calculate the zero point for the image. Thereafter, all fluxes can be photometrically calibrated.

2.2.3 Image subtraction and transient detection

For MeerLICHT to find transients, difference image analysis is performed on two images of the same field taken at different epochs. In most cases, a deep composite reference image is subtracted from a new image, though at the time of writing MeerLICHT’s current procedure is to set the first good quality image of a field (in a particular filter) as the reference image to be used. The image subtraction routine formulated by Zackay *et al.*

code available at <https://github.com/pmvreeswijk/ZOGY>.

³The criteria for selection are outlined in section 4.8 of Daniëlle Pieterse’s Master’s thesis [28], a very

[26] is used for performing difference image analysis. Known as ZOGY, the method uses statistical principles to derive the optimal statistic for transient detection, taking into account the PSFs of both the new and reference images to produce the difference image [29]. The significance image represents the significance of pixels in the difference image, while the corrected significance (S_{corr}) image takes into account the source noise and astrometric error, meaning that errors due to bad subtractions are less likely to show up.

Transients are identified from the S_{corr} images. All sources having $|S/N| > 6$ are included in a transient catalogue file associated with the new image. A positive S_{corr} value for a source indicates that the source is new or has brightened with respect to the reference image, while negative values indicate that the source has faded. A machine learning classifier is being developed by Zafirah Hosenie (PhD student at the University of Manchester) to identify the real sources from the many bogus sources that end up in the transient catalogues.

2.2.4 Data products

The headers of the image files contain important information about the observing conditions, the telescope and the observation itself, all of which are important for image reduction. Keywords are added to the headers of files during reduction giving information on the image and the reduction process itself. The software pipeline performs quality control during reduction by examining the header keywords and determining if they are within the acceptable range of values. Quality control flags are assigned to each image based on the worst flag assigned to a keyword. Red flagged files are unusable since they are not reduced further, while orange, yellow and green- flagged files are usable. The following data products are produced for each reduced raw file:

1. A log file containing information about the reduction process
2. The reduced image
3. A mask file indicating pixels that should not be used for science
4. A full source catalogue file containing all sources detected in the image along with their important parameters
5. Background and background standard deviation images created by SExtractor – currently, these are not used
6. A difference image created with ZOGY
7. A corrected significance (S_{corr}) image
8. A catalogue of transient sources extracted from the S_{corr} image
9. A transient limiting magnitude image, providing the limiting magnitude of the observation for each pixel

The last four images are only created if a reference image of the observed field already exists. The data products are stored on the Inter-University Institute for Data-Intensive Astronomy (IDIA) research cloud in the folder `/ceph/meerlicht/red/`, ordered by date. In most cases, the reduced image, full source and transient catalogue files are sufficient for scientific analysis.

2.3 Science goals

The science goals of MeerLICHT centre around its novel pairing with the 64-antennae MeerKAT radio array – the precursor to the Square Kilometre Array (SKA). Each antenna is 13.5 m in diameter with an offset Gregorian design, together comprising a radio interferometer with a maximum baseline of 8 km. Currently, two receivers are in use: the UHF-band (580-1015 MHz), and L-band (900-1670 MHz) [33]. Approximately two thirds of observing time is dedicated to a number of Large Survey Projects (LSPs). Two of these focus on the detection and studying of radio transients, namely *ThunderKAT* [34] and *TRAPUM* [35]. *ThunderKAT* aims to study Galactic and extragalactic synchrotron transients associated with high energy accretion and explosive events in compact object binaries (consisting of neutron stars, black holes and white dwarfs), gamma-ray bursts, supernovae, and gravitational wave events. *TRAPUM* will focus on fast (<1 s) radio transients with the aim of discovering new pulsars and fast radio bursts (FRBs). Both projects will rely on the commensal use of MeerKAT data from the other LSPs.

Their similar geographic locations in the Karoo region of South Africa and alike fields of view – 5 deg^2 at 580 MHz and 0.85 deg^2 at 1400 MHz for MeerKAT compared to 2.7 deg^2 for MeerLICHT – makes their pairing particularly suitable for the simultaneous detection of short-timescale transients. Triggering optical follow-up observations for a radio transient detected by MeerKAT could be time-consuming, meaning that the possible optical counterpart to any short-timescale transient might be missed. For this reason, MeerLICHT aims to provide continuous night-time, multi-filter coverage matching MeerKAT’s.

In addition to co-observing with MeerKAT, MeerLICHT has a local transients program to search for transients in nearby (<100 Mpc) galaxies. Currently, MeerLICHT is completing a multi-filter full southern sky ($\text{dec} < +30^\circ$) survey, as it is vital to have a complete set of reference images for transient detection. The whole southern sky has been divided into a fixed grid of 15946 fields, each with a unique ID and a sky area matching MeerLICHT’s field of view. Given a set of coordinates, MeerLICHT observes the field having the field ID that contains the desired coordinates, and uses the existing reference image for that field to perform transient analysis. There is an additional small set of “special fields” covering areas of high scientific interest including galaxies, galaxy clusters, and globular clusters.

In the context of this thesis it is important to mention the BlackGEM array situated at the ESO La Silla site in Chile. The array currently consists of 3 optical telescopes identical to MeerLICHT in optical design, but with each telescope housed on top of a 5 m double-cylinder tower to reduce night-time air turbulence, resulting in sub-arcsecond image quality across the image plane [24]. BlackGEM’s highest priority science goal is the rapid follow-up of gravitational wave events detected by the LIGO/Virgo Collaboration (LVC). The next phase of the project will extend the array to 15 such telescopes at La Silla, with the option available for further expansion to southern Africa and Australasia providing all-day southern sky coverage. Additionally, BlackGEM will undertake a number of surveys aimed at finding sub-day timescale transients related to high-energy explosive events including supernovae shock breakouts, white dwarf binaries, fast radio bursts and fast X-ray transients [24].

Since BlackGEM is not yet operational, it made sense to utilise MeerLICHT in the search for the possible electromagnetic counterpart to the gravitational wave event GW190814 – the subject of this thesis.

useful reference for this chapter, together with Kerry Paterson’s PhD thesis [29].

Chapter 3

The data

3.1 GCN alert

A preliminary GCN Notice was sent out by the LVC at 21:31:40 UT on the 2019 August 14 indicating that a gravitational wave event had been detected. The event was given the superevent ID S190814bv. The initial BAYESTAR [36] sky map – sent out 21 minutes after the merger – gave the 90% probability region a sky-area of 772 deg². About two hours later at 23:39:31 UT the first GCN circular [37] was sent out by the LVC along with an updated BAYESTAR sky map, giving the 90% probability region a much-reduced area of 38 deg² at a luminosity distance of 276 ± 56 Mpc. Additionally, the false-alarm rate was very low at one event per 1.559×10^{25} years.

The following day at 10:41:01 UTC, a further GCN Circular [38] was sent out by the LVC with a final sky map produced by the LALInference [39] software library. The 90% region had further reduced to an area of 23 deg² at a luminosity distance of 267 ± 52 Mpc. The previous circular had classified the compact binary coalescence as a MassGap merger but this was revised to a NSBH merger with a probability greater than 99%. However, it also had a low probability (<1%) of having a remnant. If there is neutron star material remaining outside the final remnant compact object, the likelihood of there being observable EM emission is greater.

3.2 Observations with MeerLICHT

Having received the GCN notices along with the GW localisation sky map, the MeerLICHT team decided to trigger a follow-up observing campaign of the GW localisation region. The ranked-tiling method [40] was used to determine which fields would be observed over the course of the follow-up campaign. The ranked tiling method improves the localisation coverage over the contour-covering method, and is particularly suitable for MeerLICHT's fixed sky-grid. The integrated probability in each of the fields covering the localisation area is calculated, allowing the fields to be ranked by integrated probability. The ranked fields which cumulatively encompass a probability of 95% are then chosen to be observed. Having used this method on MeerLICHT's fixed sky-grid along with the updated BAYESTAR sky map (from GCN 25324), 24 fields were identified. These are shown in the left-hand plot of fig. 3.1.

The first observation began at 23:11:33 UTC on the 14th of August, with a total of 191 60-second observations of the 24 fields taken over the course of the night, in the u , g and i bands. With the release of the LALInference sky map the next day, implementation of

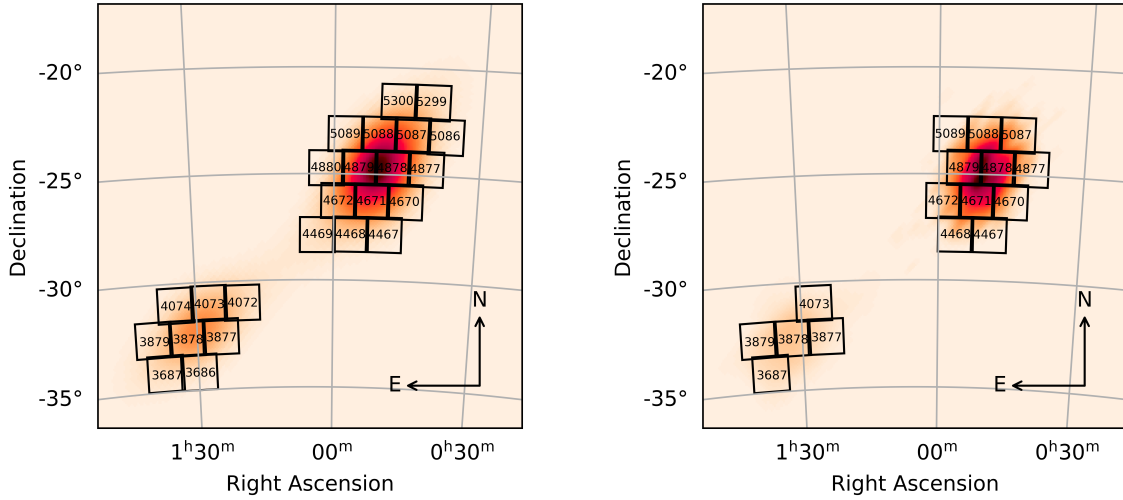


Figure 3.1: The 24 MeerLICHT fields observed over the course of the first night of observations are shown in relation to the second BAYESTAR sky map, in the figure on the left. Darker regions have a higher probability of being the source of the gravitational wave. The numbers are each field’s respective ID. Following the release of the LALInference sky map with a reduced localisation area, 16 fields were required to cover a probability of at least 95% according to the ranked-tiling method, as shown in the figure on the right.

the ranked-tiling method to the reduced localisation area resulted in 16 fields covering a probability of at least 95%, as shown in fig. 3.1. For the remainder of the observing time spent on the follow-up of GW190814, these 16 fields were observed each night. A total of 1484 observations were taken from the 14th to the 24th of August. The breakdown according to field, filter and night is shown in table 3.1.

3.3 The gradient problem

The raw files were reduced using MeerLICHT’s software pipeline as outlined in section 2.2. These files were transferred to a local drive for analysis and were sorted by field, filter, and date. A preliminary investigation of the 778 non-red flagged transient files revealed that many of them contained abnormally large numbers of transients, particularly in the q and i bands as demonstrated in fig. 3.2. Even taking into account large numbers of bogus transient detections, the massive number of detected transients alluded to a deeper problem.

A brief look at the transient file containing the most sources – 990 – revealed that the distribution of transients in the image plane was heavily correlated with the corrected significance S_{corr} values of the sources. Figure 3.3 demonstrates how the positive S_{corr} values are distributed towards the top right corner of the image, with negative S_{corr} values towards the lower left.

The reduced image associated with this transient file possessed a strong positive gradient from the bottom left to the top right of the image. The reference image also possessed a gradient, but it was positive in the opposite direction (top right to bottom left) and not as strong. The gradients in the reduced image and its reference image was

Field ID	Date of night in August										Total
	14	15	16	17	18	19	20	21	22	24	
3686	(2,2,2)										6
3687	(2,2,2)	(2,2,2)	(1,1,1)	(4,4,4)	(3,3,3)	(4,4,4)	(4,4,4)	(4,4,4)	(2,2,2)	(1,1,1)	81
3877	(3,3,3)	(2,2,2)	(1,1,1)	(4,4,4)	(3,3,3)	(4,4,4)	(4,4,4)	(4,3,3)	(2,2,2)	(1,1,1)	82
3878	(3,3,3)	(2,2,2)	(1,1,1)	(4,4,4)	(3,3,3)	(4,4,4)	(4,4,4)	(4,4,4)	(2,2,2)	(1,1,1)	84
3879	(2,2,2)	(2,2,2)	(1,1,1)	(4,4,4)	(3,3,3)	(4,4,4)	(4,4,4)	(4,4,4)	(2,2,2)	(1,1,1)	81
4072	(2,2,2)										6
4073	(3,3,3)	(2,2,2)	(1,1,1)	(4,4,4)	(3,3,3)	(4,4,4)	(4,4,4)	(4,4,4)	(2,2,2)	(1,1,1)	84
4074	(2,2,2)										6
4467	(2,2,2)	(2,2,2)	(1,1,1)	(4,4,4)	(3,3,3)	(4,4,4)	(4,4,4)	(4,4,4)	(2,2,2)	(1,1,1)	81
4468	(3,3,3)	(2,2,2)	(1,1,1)	(4,4,4)	(3,3,3)	(4,4,4)	(4,4,4)	(4,4,4)	(2,2,2)	(1,1,1)	84
4469	(2,2,2)										6
4670	(3,3,3)	(2,2,2)	(2,2,2)	(5,5,4)	(4,4,4)	(5,5,5)	(5,5,5)	(4,4,4)	(3,3,3)	(1,1,1)	101
4671	(3,3,3)	(2,2,2)	(2,2,2)	(4,4,4)	(4,4,4)	(4,4,4)	(4,4,4)	(4,4,4)	(3,3,3)	(2,2,2)	96
4672	(3,3,3)	(2,2,2)	(1,1,1)	(4,4,4)	(3,3,3)	(4,4,4)	(4,4,4)	(4,4,4)	(3,3,3)	(1,1,1)	87
4877	(3,3,3)	(2,2,2)	(2,2,2)	(5,5,5)	(8,8,7)	(5,5,5)	(5,5,5)	(4,4,4)	(3,3,3)	(1,1,1)	113
4878	(3,3,3)	(2,2,2)	(2,2,2)	(4,4,4)	(4,4,4)	(5,4,4)	(4,4,4)	(4,4,4)	(3,3,3)	(2,2,2)	97
4879	(3,3,3)	(2,2,2)	(2,2,1)	(4,4,4)	(4,4,4)	(4,4,4)	(4,4,4)	(4,4,4)	(3,3,3)	(2,2,2)	95
4880	(3,3,3)										9
5086	(2,2,2)										6
5087	(3,3,3)	(2,2,2)	(1,1,1)	(4,4,4)	(4,4,3)	(4,4,4)	(4,4,4)	(4,4,4)	(3,3,3)	(1,1,1)	89
5088	(3,3,3)	(2,2,2)	(1,1,1)	(4,4,4)	(4,4,4)	(4,4,4)	(4,4,4)	(4,4,4)	(3,3,3)	(2,2,1)	92
5089	(3,3,3)	(2,2,2)	(1,1,1)	(4,4,4)	(3,3,3)	(4,4,4)	(4,4,4)	(3,3,3)	(2,2,2)	(1,1,1)	81
5299	(3,3,2)										8
5300	(3,3,3)										9
Total	191	96	62	197	175	199	198	187	120	59	1484

Table 3.1: Table of GW190814 follow-up observations taken with MeerLICHT. The numbers in parentheses are the numbers of (u, q, i) observations taken in a particular field per night. It is evident that 8 of the 24 fields were only observed on the night of the 14th, due to the release of the updated LALInference sky map the next day.




Figure 3.2: Histogram of the number of transient sources in each of the 778 non-red flagged transient catalogue files, with a bin width of 20 sources. The maximum number and minimum number of sources in a file was 990 and 20 respectively. A median of 247 transients per file was abnormally large, alluding to a deeper underlying problem.




Figure 3.3: The distribution of the 990 transient sources in the file `ML1_20190815_013946_red_trans.fits`. The field was observed on the night of the 14th of August in the q -band and had ID 3878.

believed to be the cause of the uneven distribution in S_{corr} values in the transient file. At the time it was known that some of the flatfields possessed similarly strong gradients. Analysis of the flatfields taken in 2019 showed that the gradient only showed up in flats taken in the evenings before observing. Furthermore, the strength of the gradient would decrease as night approached. The flatfield used in the reduction of the image associated with fig. 3.3 did indeed possess a noticeable gradient. It was also surprising to notice that all of the q and i -band images had been reduced with their respective master flats from the 16th of July. This was unusual as an observation is usually reduced with the master flat produced on the day of the observation, with the master flat being made by combining all non-red flagged flats in that filter from the three days preceding (and succeeding, if they exist) the observation. Clearly, no non-red flagged flatfields in the q and i bands existed after the 16th of July.

It was believed that the flatfields with gradients were introducing gradients into the reduced and reference images, leading to the very large numbers of transients in the transient files. Consequently, it was decided that the raw data would be re-reduced using different master flats – flats without a gradient.

3.3.1 Re-reducing the raw data with good flats

First, flatfields without a strong gradient had to be identified. To do this, a python script written by Daniëlle Pieterse was used. The script divided a reduced flatfield image into a 3×3 grid and calculated the median pixel count level in each grid element. Since the gradients always appeared diagonally across the images from the bottom left to top right corners, the relative median count difference between the elements in the two opposing corners was calculated. For both the q and i bands we chose to produce their master flats using a flat window of three around the date of the 17th of June, and excluded all those files that had a relative median count difference greater than $|0.005|$. The reason this date was chosen was because this was the first time period of a few consecutive days before August where the flatfields were consistently good, while the relative median count difference of $|0.005|$ was chosen as a quality threshold as it was sufficiently low to ensure a minimal gradient, and sufficiently large to include a reasonable sample size. The u -band flats taken during the week of GW follow-up were good enough to be used, so the master flat was created by combining all the flats around the date of the 17th of August with a flat window of three and again excluding those flats with a relative median count difference greater than $|0.005|$. The old and newly created master flats for each filter are shown in fig. 3.4 for comparison.

With the newly created master flats, the raw GW data were re-reduced by the MeerLICHT pipeline. Analysis of the numbers of transients per file revealed a clear improvement from the original data reduced without these good master flats. Figure 3.5 shows a reduction from 247 to 126 in the median number of transients per file – a substantial reduction of almost 50%. Despite the improvement the long tail in the histogram showing many files having greater than 100 sources alluded to yet another effect in play. A cursory glance at the reduced images of some of these files with large numbers of transients revealed the same trend as before – a gradient. The conclusion was therefore that some of the raw science images had such a gradient which must have been introduced during observing and not by the flatfields.




(a) *u*-band flats(b) *q*-band flats(c) *i*-band flats

Figure 3.4: Comparison of old master flats (left) with the newly created master flats (right), for each of the three filters. Particularly noticeable is the strong gradient in the old *q* and *i*-band flats. The old and new *u*-band flats are almost identical.




Figure 3.5: Histogram of the number of transient sources in each of the 743 non-red flagged transient catalogue files that were reduced using good master flats. The bin width is 10 sources per file. A median of 126 transients per file was a great improvement on the median of 247 for the original data reduced without good master flats.

3.3.2 Re-reducing science images without gradients

To identify science images containing a gradient a similar approach was taken as for the flats, but instead using the background image files – those ending in `red_bkg_mini.fits.fz`. These files are produced by the MeerLICHT pipeline during image processing, and are created by masking the sources found by SExtractor in the SExtractor background image. The background is then re-determined by calculating the sigma-clipped median and standard deviation per 30×30 pixel section. Currently the files are not used though they may be used in the future. A strong gradient is clearly visible in the background image associated with the transient file having the most sources (see fig. 3.6), thereby justifying our use of these files to determine which of the reduced images possessed a gradient.

The relative median count difference between the top right and bottom left corners was calculated for the background images of non-red flagged reduced images, following the same procedure as for the flatfields. A plot of the relative median count difference versus the time of the observation is shown for all 10 days’ worth of data in fig. 3.7, while the same plot with data only from the night of the 14th is shown in fig. 3.8. The plot in fig. 3.8 is more interesting as it reveals that there was a strong gradient across all three bands until around 01:00:00 UTC, whereupon it weakened and was virtually nonexistent an hour later. This trend was apparent on most nights.

Because many of the reference images possessed a gradient, we could not use all of the transient files – even for those reduced images that did not possess a gradient. It was decided that the best way to completely eliminate the problem would be to select all those reduced images without gradients and re-reduce their raw files, thereby ensuring that all the reference images would be good. All those files that had a relative median




Figure 3.6: The reduced image `ML1_20190823_002949_red.fits.fz` (left) is shown alongside its background image. The scaling in the reduced image hides the background gradient. These files were associated with the transient file having 573 sources – the most in the data reduced using good master flats.

count difference greater than $|0.02|$ were excluded. The threshold of $|0.02|$ was chosen because the files without a gradient in fig. 3.8 (those in the lower right along a somewhat uniform line) appeared to lie within this range. Of the 958 non-red flagged reduced images, 681 passed this quality threshold test. Their raw files were re-reduced by the MeerLICHT pipeline, again with the good master flats.

3.3.3 The gradient problem solved

The underlying cause of the gradient problem in both the flatfields and science images was discovered by the MeerLICHT team in February 2020 after a trip to monitor the telescope’s operations. It appeared that the gradients were caused by vignetting of the telescope’s dome. The dome of the telescope has a slit of fixed orientation through which the telescope observes. During normal operations, the dome rotates in azimuth such that the telescope is pointing directly through (or near) the centre of the slit. The dome of the telescope has four sensors – one at each of the four cardinal points – that record the dome’s position so that the dome software knows into which position the dome should be moved to accommodate the telescope for observing. It appeared that the East sensor was malfunctioning and thus reporting that the dome was in the correct position to accommodate the telescope, while in reality it wasn’t. This meant that the telescope was not observing directly through the centre of the slit, leading (in many cases) to vignetting of the dome and thus a non-even distribution of light intensity across the image plane. This was the cause of the gradient in brightness across some of our images.

Once the dome moved through North, the software would reset meaning the actual position of the dome was correctly being recorded by the sensors. Figure 3.9 demonstrates how stronger gradients – and therefore more dome vignetting – occurred when the telescope was pointed towards the East and the malfunctioning dome. This also explained the trend of a gradient in the evening flatfields and in fig. 3.8 where at the beginning of the night




Figure 3.7: A plot of relative median count difference in the background images versus the time of the observation is shown for the 958 non-red flagged reduced images from the 10 days' worth of data. In general, the gradient was strongest in the observations from the beginning of each night. C_{ur} and C_{ll} are the median count rates in the upper right and lower left sections of a background image, respectively.




Figure 3.8: The plot of relative median count difference versus time for the night of the 14th of August revealed that the strong gradient started to decrease around 01:00:00 UTC and was all but gone an hour later.




Figure 3.9: A plot of the relative median count difference as a function of the telescope’s azimuth revealed how the strongest gradients occurred when the telescope was pointed towards the East, and thus in the direction of the malfunctioning dome.

the telescope would start observing in the East and thus be affected by the faulty dome. As the night progressed and the telescope moved further towards the zenith, the gradient problem would weaken until vanishing completely.

3.4 The (good) reduced data

The 681 images without a gradient that were reduced with good master flats became the final dataset on which subsequent analysis took place. None of the 681 reduced files were flagged red. Each of the 681 reduced files had a corresponding catalogue file, of which 47 were flagged red, mainly during astrometry with certain parameters not falling within their acceptable ranges. Because an early observation in a particular field and filter would become the reference image for that field and filter, every reduced file did not necessarily have a corresponding transient file. This explained why there were 618 transient files. Of these, 47 were flagged red, the same 47 that were flagged red as catalogue files. This resulted in a total of 571 usable transient files.

Figure 3.10 demonstrates the vast reduction in the numbers of transients per file. Originally the median value was 247 sources per file. This was reduced to 126 using good master flats, and when excluding all those files that had a gradient, was further reduced to a median of 46 sources per file, a much more realistic – and manageable – number. It is important to notice, however, that the q -band transient files still possessed a reasonably large number of transients per file. There is a clear peak between 200 and 250 transients per file in fig. 3.10. The distribution of transients across all of the q -band transient files revealed an overdensity of sources in the top left edge of the image plane (see fig. 3.11). Upon examination of the reduced images, it was noticed that this region of the CCD had




Figure 3.10: Histogram of the number of transient sources in each of the 571 non-red flagged transient catalogue files that formed part of the final dataset. The bin width is 10 sources per file. A median of 46 transients per file was an even bigger improvement on the numbers of transients in files that were reduced with good flats, however the q -band files still possessed reasonably large numbers of transients.

been affected by etching during manufacture. The combination of the high night time background due to the full moon along with the master flatfield having an insufficient SNR to correct for the etching meant that the reduced images had artefacts left in them that were detected as transients. This explained the overdensity of transients in the top left edge of the q -band files.

The limiting magnitudes in all three bands as well as the seeing for all the observations are shown in fig. 3.12. The rise and fall of the limiting magnitudes was a result of the moon which was full on the 15th of August. The median seeing value was approximately 3.27 arcseconds over the course of the observing run.

3.5 Finding transients

One of the principal aims of this project was to perform an independent search for transients in our MeerLICHT data on GW190814. To do this, only the good reduced data (see section 3.4) of the 16 fields visible in the right of fig. 3.1 was analysed. Table 3.2 shows the numbers of usable transient files per field and filter for each observing night. A total of 564 transient files were searched for transients. The following procedure was drawn up for finding real transients in the data, with each stage being explained in detail in the subsequent subsections:

1. In a particular field on a particular night, perform a crossmatch of all unique combinations of pairs of transients files, regardless of filter.




Figure 3.11: A 2-D histogram of the distribution of transients in all of the non-red flagged q -band transient files. Each bin was 330×330 pixels in size. There is a clear overdensity of transients in the top left edge of the image plane, linked to the etching of the CCD.





2. Manually vet each matching pair using the reduced, reference, difference and S_{corr} image cutouts.
3. Identify moving objects that made it through the manual vetting stage.
4. Produce light curves of remaining astrophysical sources.
5. Search the Simbad, Gaia, and TNS databases in order to determine if any of the transients were unreported.

3.5.1 Identifying matching pairs

In order to reduce the number of bogus sources to be vetted, two steps were taken. The first was to only consider those sources in the transient files which had an $|S_{\text{corr}}|$ value greater than 10. Currently, the MeerLICHT pipeline automatically excludes those sources with $|S_{\text{corr}}| < 6$. A higher S_{corr} value indicates a more significant transient detection. Sources with a lower significance are more likely to be bogus detections caused by artefacts such as cosmic rays, bad pixels, or poor flatfielding (see fig. 3.11).

For the second step it was decided that any source must have been detected in at least two transient files on a particular night, regardless of the filter of the observation. To do this, a list of all unique combinations of transient file pairs from a particular field and night was created. The number of unique pairs is given by the simple combinatorial formula

$$C(n, r) = \frac{n!}{r!(n-r)!}, \quad (3.1)$$

(a) u -band limiting magnitudes(b) q -band limiting magnitudes(c) i -band limiting magnitudes

(d) Seeing conditions

Figure 3.12: Limiting magnitudes across all observations and bands are shown in (a), (b), and (c), while seeing is shown in (d). The median seeing value was approximately 3.27 arcseconds over the course of all the observations.

Field ID	Date of night in August							Total	
	14	15	17	18	19	20	21		22
3687		(1,2,2)	(1,1,2)	(1,1,1)	(1,2,2)	(2,3,3)	(0,4,4)	(1,1,1)	36
3877	(0,1,0)	(0,2,2)	(2,2,2)	(1,1,1)	(2,1,2)	(3,3,2)	(3, 3,3)	(1,1,1)	39
3878	(0,1,1)	(0,1,1)	(0,1,1)	(1,1,1)	(1,2,2)	(3,3,3)	(4,4,4)	(1,1,1)	38
3879	(0,1,0)	(0,1,2)	(1,1,1)	(1,1,1)	(1,2,2)	(3,3,3)	(1,4,4)	(1,1,1)	36
4073		(2,2,2)	(1,1,2)	(1,1,1)	(2,1,2)	(3,3,2)	(3,4,4)	(1,1,1)	40
4467		(1,2,2)	(1,1,1)	(1,1,1)	(2,1,3)	(3,3,2)	(3,4,4)	(1,1,1)	39
4468	(0,1,1)	(1,2,2)	(0,2,2)	(0,1,1)	(2,2,2)	(2,3,3)	(2,4,4)	(1,1,1)	40
4670	(0,1,1)	(0,1,1)	(1,1,1)		(3,4,4)	(3,3,3)	(3,4,4)	(1,1,1)	41
4671		(0,1,1)	(1,1,1)	(1,0,0)	(2,3,3)	(3,3,3)	(2,4,4)	(1,1,1)	36
4672						(2,0,0)	(3,0,0)	(1,0,0)	6
4877	(1,1,1)	(0,1,1)	(1,1,1)		(4,4,3)	(3,3,3)	(4,4,4)	(1,1,1)	43
4878		(0,1,1)	(1,1,1)		(2,3,2)	(3,3,3)	(3,4,4)	(1,1,0)	34
4879		(0,1,1)	(0,1,0)		(1,3,3)	(2,3,3)	(3,4,4)	(1,1,1)	32
5087		(0,2,2)			(1,3,3)	(3,3,3)	(4,4,4)	(1,1,1)	35
5088		(0,1,2)	(1,1,1)		(3,3,3)	(3,3,1)	(4,4,4)	(1,1,1)	37
5089	(0,1,1)	(0,2,2)	(0,1,1)		(2,2,3)	(3,3,3)	(2,3,3)		32
								Total	564

Table 3.2: Table of the non-red flagged transient files that were used in our search for transients, organised by field ID. The numbers in parentheses are the numbers of (u, q, i) transient files per field per night.

with n equal to the number of transient files and $r = 2$ since we were looking for pairs of transient files. Each pair of files was then cross-matched with a one arcsecond sky radius to identify sources catalogued in both files, as a two arcsecond matching radius did not appear to produce a substantial increase in matches. These pairs of sources occurring in both files – known hereon as a matching pair – would be manually vetted.

A bash script was written to automate the identification of these matching pairs for each night per field. The STILTS [41] command-line tool was used to filter out all sources having $|S_{\text{corr}}| < 10$ while performing the cross-matching of the pairs of transient files. STILTS is used for the processing of tabular data, and is especially suitable for large astronomical datasets. For each night per field, a table of matching pairs was generated showing the coordinates and names of the two files in which each source was co-identified. A total of 2768 matching pairs across all fields were identified for vetting.

3.5.2 Vetting of matching pairs

With the tables of matching pairs, the question arose as to how the pairs should be vetted manually. After much consideration, it was decided that a web interface¹ would be an effective way to perform the vetting and keep track of selected candidates through a

¹A number of reasons influenced my decision to use a web interface rather than something like a local GUI. Firstly, I felt that learning the basics of web development would be useful for possible future projects and would add another skillset to my quiver of skills. Secondly, a website allows other users to view and interact with the data more easily. Lastly, a web framework like Django lends itself to using databases, making storage and organisation of data relatively simple.




Figure 3.13: An example vetting image for the matching pair having coordinates (13.04404, -28.49332) in field with ID 4467. This source was identified in Simbad as an RR Lyrae variable.

database. The free and open-source Python web framework [Django](#) was used to develop the vetting interface. Being completely new to web development, a sizeable time investment had to be made initially to learn the basics, but thereafter, things went more smoothly. In essence, the website contained two databases: a vetting database containing important parameters and vetting images for each of the 2768 matching pairs; and a candidates database containing parameters, links to external databases, and light curves for sources which were identified as real transients during vetting. The website was initially developed and hosted on a local computer with the option of deploying it online in the future.

For the vetting database, a figure containing 8 cutouts was created for each matching pair. As can be seen in [fig. 3.13](#), the four cutouts on the top row of the figure are all related to the transient detection from the transient file with the prefix given above those cutouts, while the bottom four cutouts are related to the second transient file in which the source was co-identified (prefix also given above). For both rows, the cutouts are (from left to right) from the reduced image, reference image, difference image, and the S_{corr} image, with each cutout being 1×1 arcminute in dimension. The cutouts are displayed with [DS9](#)'s Zscale normalisation where darker pixels indicate a lower count level than lighter pixels.

Using the web interface, all of the matching pairs were manually vetted, with pairs which appeared to be real being added to a list of candidate transients. A large number of star-like sources – in most cases probably variables – had duplicate matching pairs. In this case, the first matching pair for the source at those coordinates was added to the list of candidates. The majority of matching pairs were bogus detections (see [fig. 3.14](#) for an example).




Figure 3.14: An example of a bogus matching pair, likely due to a bad pixel in the reference image.

3.5.3 Identifying moving objects

The list of candidate transients in the candidates web database contained approximately 58 sources. Many of the sources had detections in both their reduced and reference images, meaning they were almost definitely astrophysical in origin. Those sources which were only visible in one of either their reduced or reference images would have to be checked for moving objects as it may have been the case that a moving object was detected in two images observed close to each other temporally, resulting in a matching pair. The course of action was simple – produce cutouts centred at the source’s coordinates from observations taken on the same night, and examine if the source appears to move across the field. Figure 3.15 demonstrates this for two matching pairs. For both series of images, q -band images are shown on the left, and i -band images on the right. Earlier observations are shown towards the top, with later ones towards the bottom. For both examples, a source appears to move across the image field, identified with a green circle. If a matching pair source was only visible in its reference image and at no later epochs, it was discarded as a transient candidate. A total of 12 candidate transients were excluded from the list using this method.

3.5.4 Light curves

The remaining 46 candidates were likely all astrophysical in origin. Light curves were produced for each candidate by searching each catalogue file for a one arcsecond match within the candidate’s coordinates. If a match existed, the magnitude and its error would form a datapoint on the candidates’ light curve at that epoch. If no match existed, the limiting magnitude would be shown at that epoch. The light curves were added to the candidates database in the web interface to allow for easy perusal. Figure 3.16 provides




Figure 3.15: An example of two moving objects which were identified using image cutouts from the night of their matching pair detection. The left and right figure cutouts are centred on coordinates $(13.21366, -27.05262)$ and $(13.31326, -26.96465)$ respectively.

an example of such a light curve.

3.5.5 Determining if transients were unreported

The [Simbad](#), [Gaia DR2](#), and Transient Name Server ([TNS](#)) databases were queried to determine if any of our candidates were new or unreported. Every single source had a match within one arcsecond in the Gaia DR2 database, while a large number of sources were also found in Simbad, providing a classification for the source. A single source – the quasar QSO B0035-252 – had been reported to the TNS on the 19th of August by the MASTER group [42] and given name [AT2019nvx](#), though from the light curve it simply appeared to be showing variable behaviour. Table 3.3 displays the transient candidates of astrophysical origin that were found through this analysis.

RA (deg)	Dec (deg)	Field ID	Simbad classification	G mag	\bar{q} mag
24.03886	-34.50995	3687		14.22	14.52
25.14406	-34.29417	3687		13.37	13.54
23.98727	-33.64103	3687		14.32	14.35
24.49559	-33.68419	3687		16.85	17.08
21.20969	-32.68997	3877	Eclipsing binary	11.31	11.43
20.47796	-31.84143	3877	RR Lyrae variable	14.42	14.34
21.79382	-32.23449	3877	RR Lyrae variable	15.58	15.85
21.66897	-33.30406	3877	Possible RR Lyrae variable	16.81	17.02
23.72503	-33.04030	3878		12.45	12.52

23.67084	-33.32235	3878		14.14	14.28
23.91714	-33.37142	3878		15.15	15.29
24.84324	-33.07292	3879	RR Lyrae variable	12.80	12.91
25.77454	-32.71110	3879		14.06	14.24
24.33520	-32.94701	3879		14.43	14.56
20.85303	-31.30931	4073		13.97	14.24
22.49961	-30.67453	4073		12.84	13.10
12.37119	-27.38717	4467	RR Lyrae variable	13.32	13.44
11.77327	-27.73729	4467	Star	13.76	13.94
11.98638	-27.75949	4467	Star	15.36	15.48
13.04404	-28.49332	4467	RR Lyrae variable	17.44	17.49
13.67282	-28.23188	4468	Star	13.64	13.71
13.32939	-27.55863	4468	Star	15.39	15.54
11.07424,	-26.07290	4670	Possible RR Lyrae variable	16.51	16.64
13.19452	-26.56870	4671	RR Lyrae variable	15.18	14.89
13.16781	-26.54159	4671	RR Lyrae variable	15.15	14.91
10.25079	-24.88181	4877		14.05	14.14
9.56137	-24.98390	4877	Quasar	19.01	15.69
10.60015	-24.52526	4877	RR Lyrae variable	15.63	15.77
11.03642	-23.93112	4877		14.68	14.81
11.14006	-24.39554	4877		14.52	14.65
11.22108	-24.71085	4878	RR Lyrae variable	13.43	13.49
11.43848	-24.75447	4878	Eclipsing binary	14.68	14.94
12.73859	-23.91721	4878		14.00	14.16
11.38689	-24.33887	4878		14.82	14.97
12.55806	-24.03788	4878	RR Lyrae variable	17.39	17.56
12.95900	-24.65356	4879		12.94	13.22
12.90751	-23.95032	4879		13.64	13.83
13.06334	-25.10326	4879		14.64	14.74
14.42013	-24.02047	4879	Horizontal branch star	17.21	17.40
11.82490	-22.75212	5087	Rotationally variable star	10.93	11.17
10.84540	-23.28073	5087		14.19	14.50
11.25533	-22.39226	5087	Possible RR Lyrae variable	15.37	15.47
10.32200	-23.03056	5087		See footnote ²	
11.98874	-22.35416	5088	RR Lyrae variable	16.54	16.66
12.16165	-22.41826	5088		15.10	15.24
13.97614	-23.25184	5089	RR Lyrae variable	16.67	16.97

Table 3.3: Table of astrophysical transient candidates, with coordinates given in the ICRS system. Simbad classifications have been given where they exist while the G -band magnitudes are from the Gaia DR2 catalogue. The column labelled \bar{q} gives the average q -band magnitude from our observations. The locations of the sources with respect to the GW localisation are shown in fig. 3.17. Appendix A.1 presents the light curves of all 46 candidates.

²Gaia DR2 has two sources which are 2.29" and 2.90" away from these coordinates with G -band magnitudes of 19.70 and 17.57 respectively. It is likely that these two sources were not resolved by MeerLICHT and are possibly a double star or binary.




Figure 3.16: An example light curve for the transient candidate at coordinates (13.04404, -28.49332) in field 4467. The source was identified as an [RR Lyrae variable](#) in Simbad. Non-detections are shown as vertical dashes – note that these are distinct from the thicker error bars for the light curve data points.

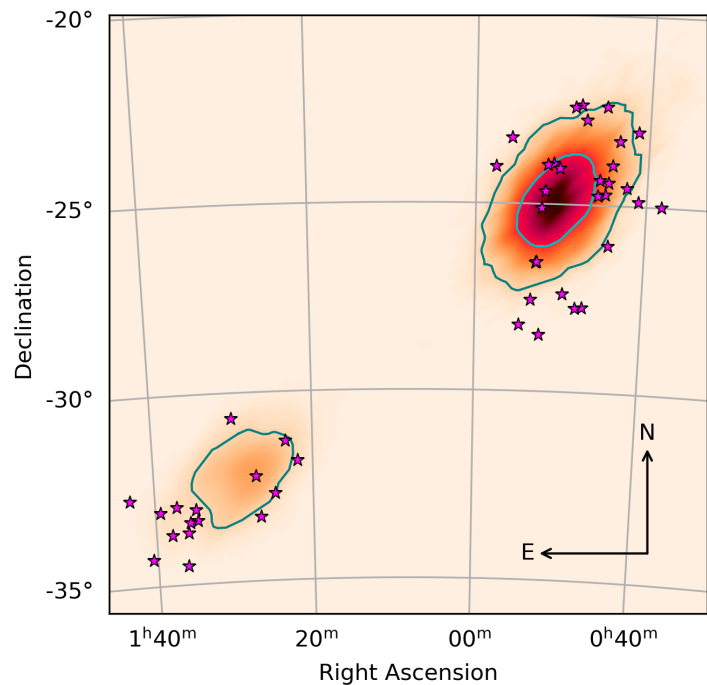


Figure 3.17: The sky positions of astrophysical transient candidates found during our search for transients in MeerLICHT data are shown with respect to the 50% and 90% contained probability contours for the GW localisation. All sources were found in Gaia data, excluding them from being associated with GW190814.

Any source in Gaia data was unlikely to be associated with GW190814 as a new kilonova-like transient would not have been expected to be visible prior to the GW event, especially given the expected luminosity distance of 241_{-45}^{+41} Mpc. Since we found no new transients in the data, we can conclude that none of our candidates were connected with the GW event.

3.6 Transients reported by other groups

Table C. 1. from Ackley *et al.* [43] lists all reported candidate counterparts associated with GW190814 that were reported to TNS from a week before up until two weeks after the GW event. Of the 78 candidates, 7 would not have been visible in any of the 16 MeerLICHT fields due to not falling within their fields-of-view. The transient and catalogue files from our data were searched to determine if any of the candidates had been observed by MeerLICHT. None of the candidates were detected in the transient files from the good reduced data (see section 3.4). A total of 22 sources had at least one detection within one arcsecond of the TNS-published coordinates. Table 3.4 compares the discovery magnitudes of these sources with the earliest magnitude detected by MeerLICHT. The discovery magnitudes are generally much fainter than the ML magnitudes, indicating that in some cases MeerLICHT may have detected a host galaxy rather than the actual transient. The source AT2019nqr was detected as a transient three times in the original reduced data, however, we refrain from reporting any magnitude changes due to the gradient problem described in section 3.3. The reason that all but one of these sources were not detected as transients by the MeerLICHT pipeline is most likely due to the fact that a faint transient source (in the new or reference image) will have a very low SNR in the difference image and thus an $|S_{\text{corr}}|$ value well below the threshold of 10 we set for finding transients.

IAU name	RA (deg)	Dec (deg)	Discovery mag	ML mag	Epochs	Comment
AT2019nbp	11.73914	-24.36186	$w = 20.0$	$q = 18.30$	24	Pre.Det.
AT2019noq	12.19951	-25.30652	$i = 19.9$	$i = 18.61$	2	SN
AT2019nor	12.46663	-24.27159	$i = 19.7$	$q = 18.03$	35	SN?
AT2019npe	10.38887	-23.74221	$z = 20.1$	$i = 18.16$	15	Non.det
AT2019npz	13.27317	-24.36075	$z = 20.6$	$i = 18.09$	22	AGN?
AT2019nqp	14.21001	-24.34721	$z = 20.7$	$q = 19.40$	5	Pre.Det.
AT2019nqr	23.57354	-32.74178	$i = 18.3$	$i = 15.64$	40	SN
AT2019nqs	23.39652	-31.78013	$i = 20.4$	$i = 17.73$	16	Spec.Host-z
AT2019nqz	11.69340	-24.33668	$z = 21.1$	$i = 17.89$	23	Spec.Host-z
AT2019nra	13.79460	-25.94920	$i = 20.5$	$q = 19.37$	18	?
AT2019ntp	12.55025	-26.19800	$z = 21.1$	$q = 19.67$	1	SN
AT2019nul	13.81851	-26.94294	$z = 20.3$	$q = 18.26$	21	Spec.Host-z
AT2019nun	14.20250	-24.90847	$i = 21.7$	$q = 18.00$	27	Spec.Host-z
AT2019nve	14.02296	-24.64067	$z = 21.3$	$q = 19.99$	1	SN?
AT2019nvr	12.06703	-25.47083	$z = 20.9$	$q = 19.68$	2	SN?
AT2019nvs	13.15733	-26.19483	$z = 21.4$	$q = 19.57$	5	SN?
AT2019nzs	11.83921	-24.57683	$i = 21.8$	$q = 19.38$	17	?
AT2019odc	11.50704	-25.45915	$i = 21.6$	$q = 17.48$	25	?
AT2019ofb	13.83015	-26.19744	$z = 21.4$	$q = 19.57$	8	?
AT2019omu	23.49537	-34.33889	$i = 21.3$	$q = 19.24$	8	Spec.Host-z
AT2019qcb	11.57932	-26.14545	$r = 21.6$	$q = 19.14$	23	SN?
AT2019tii	12.39668	-27.03593	$i = 21.9$	$q = 19.52$	1	Photo.evol.*

Table 3.4: Table of transients with at least one detection in MeerLICHT data using a cross-match radius of one arcsecond. The ML magnitude is the earliest observed magnitude with MeerLICHT, regardless of filter. Comments indicate the reasons given in table C. 1. of Ackley *et al.* for ruling out the transients as counterparts.

Chapter 4

Limits on the electromagnetic counterpart to GW190814

The absence of any new transients in our data meant that no electromagnetic counterpart to GW190814 was detected by MeerLICHT over the course of our observations. This non-detection allows us to place limits on any potential counterpart and constrain possible KN models.

4.1 Comparison with AT2017gfo-like KN

It is instructive to compare the limiting magnitudes from our observations with the light curve of the only observed kilonova to date – AT2017gfo. Ackley *et al.* [43] perform phenomenological fits to the AT2017gfo light curve in each of the relevant bands so as to compare their limits with a possible AT2017gfo-like KN.

Our MeerLICHT limiting magnitudes were obtained from the headers of the non-red flagged reduced images in each of the u , q and i bands (shown in fig. 3.12). Photometric data on AT2017gfo was obtained from Villar *et al.*'s [44] compilation of light curves associated with that event. Combined U - and u -band data was used for our u -band fit; V -band data was regarded as a proxy for MeerLICHT's q -band as they both have similar central wavelengths, though our q -band is wider; and the i -band data was naturally used for our i -band. Gompertz *et al.* [45] fit either an exponential or Bazin function depending on if a clear peak is visible in the light curve post-merger. Exponentials were fitted to the three light curves in flux space and converted back to magnitudes, as shown in fig. 4.1. The normalisation constant and characteristic time scale for each fit is shown in table 4.1.

The model fits for each band were then converted to apparent magnitudes at the recently published [6] distance of GW190814. Using distances of 40 Mpc and 241^{+41}_{-45} Mpc to AT2017gfo and GW190814 respectively, the conversion amounted to a shift of 4.12 magnitudes fainter. Our limiting magnitudes are shown in relation to these model fits

Data used	Normalisation (μJy)	τ_{fall} (days)
$U + u$	731.01	0.48
V	685.67	1.11
i	610.69	1.85

Table 4.1: Parameters of exponential fits to AT2017gfo photometric data, shown in fig. 4.1.




Figure 4.1: Photometric data on AT2017gfo from [45] fitted by exponentials along with residuals (in magnitudes). All magnitudes are in the AB system. The x -axis times are post-merger. The exponential fits start to noticeably under-predict the data at 1.4 days post-merger.




Figure 4.2: The phenomenological fits to the AT2017gfo light curve (with u dashed, q smooth and i dotted) are shown in relation to MeerLICHT’s limiting magnitudes. The shaded regions indicate the 1σ distance uncertainties. Apparent magnitudes were converted to absolute magnitudes using $M = m - 5\log_{10}(d_L/10\text{pc})$ at the distance of GW190814.

in fig. 4.2. Our most sensitive limits are naturally in the wider q band, though even the deepest limits are insensitive to an AT2017gfo-like KN by at least a magnitude.

Using our limiting magnitudes and the 3D probability distribution for GW190814, we can calculate the covered probability of our observations in each band assuming an AT2017gfo-like KN. Appendix A in Ackley *et al.* [43] outlines the approach for converting limiting magnitudes to limits in the model parameter space. Given an EM counterpart model defined by the intrinsic light curve $dL/d\nu(\nu, t)$, we can exclude such emission out to our limiting flux levels. For a wide-field search the exclusion confidence – commonly known as the “covered probability” – is defined as

$$1 - P(\xi) = \sum_{i=1}^{N_{\text{tiles}}} P_{\text{tile},i} \int_0^\infty dr \frac{dP}{dr} \varepsilon_i(r, \xi), \quad (4.1)$$

with

$$\varepsilon_i(r, \xi) = 1 - \prod_{j=1}^{N_{\text{obs},i}} H(F_{\text{lim},i,j} - F_{\text{model}}(r, \xi, \nu_{i,j}, t_{i,j})), \quad (4.2)$$

where H is the Heaviside function. The sum runs across all non-overlapping observed tiles where $P_{\text{tile},i}$ is the 2D probability density of the GW event across the sky for tile i , and dP/dr describes how the probability varies with luminosity distance. The quantity $\varepsilon_i(r, \xi)$ equals one up until the distance that the counterpart model is too faint to be detected by our most sensitive observations for a particular tile. These distances are presented in table 4.2, assuming an AT2017gfo-like KN. ξ is the set of parameters and assumptions that define our counterpart model, and for this analysis, we scale the flux of our AT2017gfo KN model by a range of multiplicative factors. Figure 4.3 demonstrates how the covered probability of our survey varies as a function of the scaling factor applied to the model

Field ID	u	q	i
3687	79	142	100
3877	25	131	91
3878	25	136	105
3879	75	139	104
4073	72	133	96
4467	72	134	94
4468	77	139	102
4670	74	142	101
4671	22	121	86
4672	68	128	100
4877	76	138	101
4878	22	121	81
4879	<1	115	83
5087	67	125	94
5088	67	124	85
5089	77	139	99
Mean	56	132	95

Table 4.2: Distances in Mpc at which an AT2017gfo-like KN would have been visible based on our most sensitive observations per field and filter. The mean values per filter have been rounded to the nearest Mpc.

flux. Assuming an AT2017gfo-like KN model, the covered probability of our observations is very low – in the q -band it is only 8.75×10^{-5} , and even lower in u and i . We see that in our most sensitive band – q – the covered probability is close to unity for a KN 5 times brighter than AT2017gfo.

4.2 POSSIS kilonova simulations

Andreoni *et al.* [46] use their observations of GW190814 to constrain the mass ejecta-viewing angle parameter space of KN models produced by the time-dependent 3D Monte Carlo code POSSIS outlined in the paper by Bulla [47]. The code models radiation transport in supernovae and kilonovae (KNe) using wavelength and time-dependent opacities. For KNe, it assumes a spherical two component ejecta model consisting of a lanthanide-rich component distributed around the equatorial plane with half-opening angle Φ , and a lanthanide-free component at higher latitudes (see fig. 4.4). The lanthanide-rich component can be thought of as the dynamical ejecta, and the lanthanide-free component as the disk wind ejecta.

The KN models produced by POSSIS depend on three parameters – the total ejecta mass (M_{ej}), the half-opening angle Φ , and the temperature T of the ejecta at 1 day post-merger. A number of these models appear in the papers of Bulla [47] and Dhawan *et al.* [48], and have been made [publicly available](#). Models are computed for ejecta mass values in the range $[0.01, 0.10]$ in steps of $0.01M_{\odot}$, Φ in $[15, 75]$ in steps of 15° , and T in $[3000, 9000]$ in steps of 2000K. For our analysis we only considered models with temperatures $T = 5000$ K, as in the work of Bulla. For each model, spectral energy distributions (SEDs) in the




Figure 4.3: The covered probability of our observations as a function of scaled AT2017gfo-like KN flux.

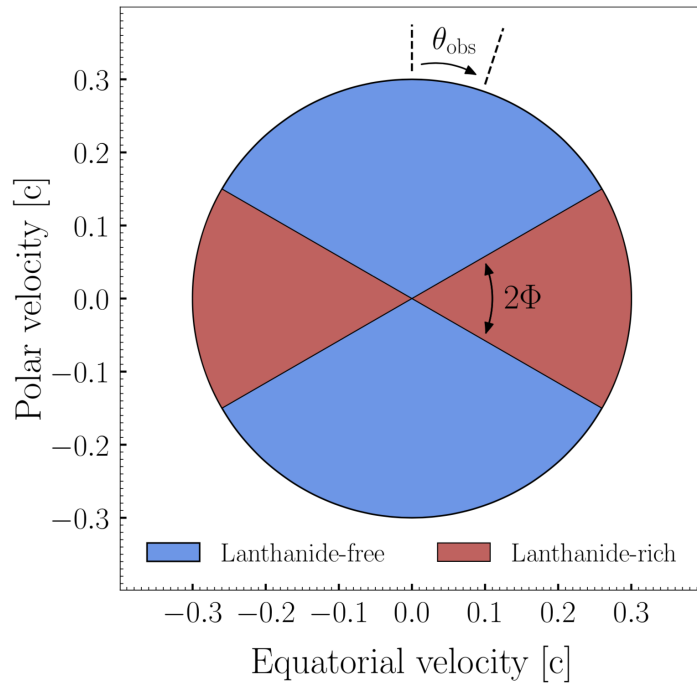


Figure 4.4: A meridional cross-section depicting the two component KN model from the Bulla paper [47] (figure 1 in that work) that is characterised by a lanthanide-rich equatorial component (red) where opacities are high, and a lanthanide-free component (blue) with lower opacities. The half-opening angle Φ determines the relative fraction of mass in each ejecta component, while θ_{obs} is the angle between the observer’s line of sight and the total angular momentum axis of the merger.




Figure 4.5: Spectral energy distribution for a KN model with $M_{\text{ej}} = 0.10M_{\odot}$, $\Phi = 15^{\circ}$, and $\cos(\theta_{\text{obs}}) = 1$ at $t = 0.5$ days post-merger is shown alongside the approximate wavelength ranges for MeerLICHT’s u , q and i bands.

wavelength range $0.1 - 2.3 \mu\text{m}$ ($\Delta\lambda = 0.022 \mu\text{m}$) were available at times ranging from 0.5 to 15 days post-merger in time steps of 0.5 days, for 11 viewing angles equally spaced in $\cos(\theta_{\text{obs}})$ in the range $[0, 1]$. Viewing angles varied from face-on ($\cos(\theta_{\text{obs}}) = 1$) to edge-on ($\cos(\theta_{\text{obs}}) = 0$). An example SED for a KN model having $M_{\text{ej}} = 0.10M_{\odot}$ and $\Phi = 15^{\circ}$ is shown in fig. 4.5 for a face-on viewing angle at $t = 0.5$ days post-merger.

Using the SEDs – where the fluxes are given at a distance of 10 pc – along with MeerLICHT’s filter transmission curves, absolute AB magnitudes in a particular filter could be calculated by integrating the flux according to

$$M_{AB} = -2.5 \log \left(\frac{\int F_{\lambda} S_{\lambda} d\lambda}{\int S_0 S_{\nu} d\nu} \right), \quad (4.3)$$

where F_{λ} and S_{λ} are the flux and filter sensitivity at wavelength λ , and $S_0 = 3631 \text{ Jy}$ is the zero-point of the AB scale. The term in the denominator sets the zero-point in that filter in the AB magnitude system. Light curves can thus be easily extracted and converted to apparent magnitudes at the desired distance.

4.2.1 Covered probability in mass ejecta-viewing angle parameter space

In order to see if our limiting magnitudes were able to constrain any of the KN models, it made sense to start off by comparing them with the model that would produce the brightest KN. As shown in the Bulla paper [47], smaller half-opening angles, higher ejecta masses, and more polar viewing angles produce brighter KNe. Figure 4.6 presents light curves extracted from the KN model having an ejecta mass of $0.10M_{\odot}$, $\Phi = 15^{\circ}$, and




Figure 4.6: Extracted u (dashed), q (smooth) and i -band (dotted) light curves from the KN model with $M_{\text{ej}} = 0.10$, $\Phi = 15^\circ$, and $\cos(\theta_{\text{obs}}) = 1$ are shown scaled at the most recent distance estimate for GW190814 – 241_{-45}^{+41} Mpc, where the shaded regions indicate the uncertainty in distance. The dashed-dotted lines indicate where the model light curves have been extrapolated. Limiting magnitudes taken with MeerLICHT are shown for comparison.

$\cos(\theta_{\text{obs}}) = 1$, and scaled to the most recent distance estimate for GW190814 – 241_{-45}^{+41} Mpc. Furthermore, the light curves have been extrapolated to times earlier than 0.5 days post-merger using a cubic spline¹, so that our earliest limiting magnitudes can be compared with the model. Comparing our limiting magnitudes taken with MeerLICHT, we see that at all times the model light curves are at least a magnitude too faint to have been observed.

In the same way as in section 4.1, we can calculate the covered probability of our observations using eq. (4.1) and eq. (4.2), assuming a given KN model. Setting $\Phi = 30^\circ$, the covered probability of our observations was calculated for a variety of KN models, varying the ejecta mass and viewing angle. Light curves were extrapolated at early times as shown in fig. 4.6. Figure 4.7 presents the covered probability in the q band according to a number of given KN models. The figure shows two trends – the covered probability increases for more polar viewing angles and is probably due to the fact that more polar viewing angles result in brighter KNe. A more interesting trend is that the covered probability peaks at an ejecta mass of $0.05M_\odot$. This is unexpected since increasing ejecta mass generally results in brighter KNe (as shown in the Bulla paper). The effect is likely due to the influence of our limiting magnitudes at early ($t < 0.5$ days) times post-merger, where the extrapolated light curve is brighter for an ejecta mass of $0.05M_\odot$ than for $0.10M_\odot$, thereby increasing the covered probability. We caution for an over-interpretation of this result as

¹Using the SciPy function `interpolate.interp1d`.




Figure 4.7: Covered probability of our q band observations for KN models varying with ejecta mass and viewing angle.

the extrapolation to earlier times is not based on the models which have their earliest data point at 0.5 days post merger. Calculation of earlier-time models is therefore highly encouraged.

4.2.2 Colour-evolution of KN models

Figures 5 and 6 in the Bulla paper [47] compare the light curves of KN models by varying a single parameter while keeping the others constant. The parameters investigated are the viewing angle, ejecta mass and half-opening angle. We thought it would be instructive to perform a similar analysis, but instead investigate the temporal evolution of colours using various combinations of MeerLICHT’s 6 bands. A total of 7 colour combinations were investigated: $u - g$, $g - r$, $r - i$, $i - z$, $u - q$, $q - i$, and $u - i$. In fig. 4.8, the viewing angle is varied while setting $M_{\text{ej}} = 0.05M_{\odot}$ and $\Phi = 30^{\circ}$; in fig. 4.9, the ejecta mass is varied while keeping $\Phi = 30^{\circ}$ and $\cos(\theta_{\text{obs}}) = 1$; while in fig. 4.10, Φ is varied keeping $M_{\text{ej}} = 0.05M_{\odot}$ and $\cos(\theta_{\text{obs}}) = 1$.

The obvious trend for all the colour-evolution plots is the general redward evolution of the KNe with increasing time post-merger. There are subtle differences in how the colours evolve as a parameter is varied. In fig. 4.8 we see that for $r - i$, $i - z$, and $q - i$, more polar viewing angles often start out bluer than for equatorial angles, but around 1 day post-merger they become redder than equatorial viewing angles. In contrast, colour combinations involving the bluer bands (u and g) are bluer for more equatorial viewing angles at times until 5 days post-merger where we see that more polar viewing angles become bluer and remain so. Varying the ejecta mass (in fig. 4.9) shows that lower ejecta masses are generally bluer between roughly 1 and 7 days post-merger, with convergence at later times. For colours containing the i -band data, there is a faster redward evolution

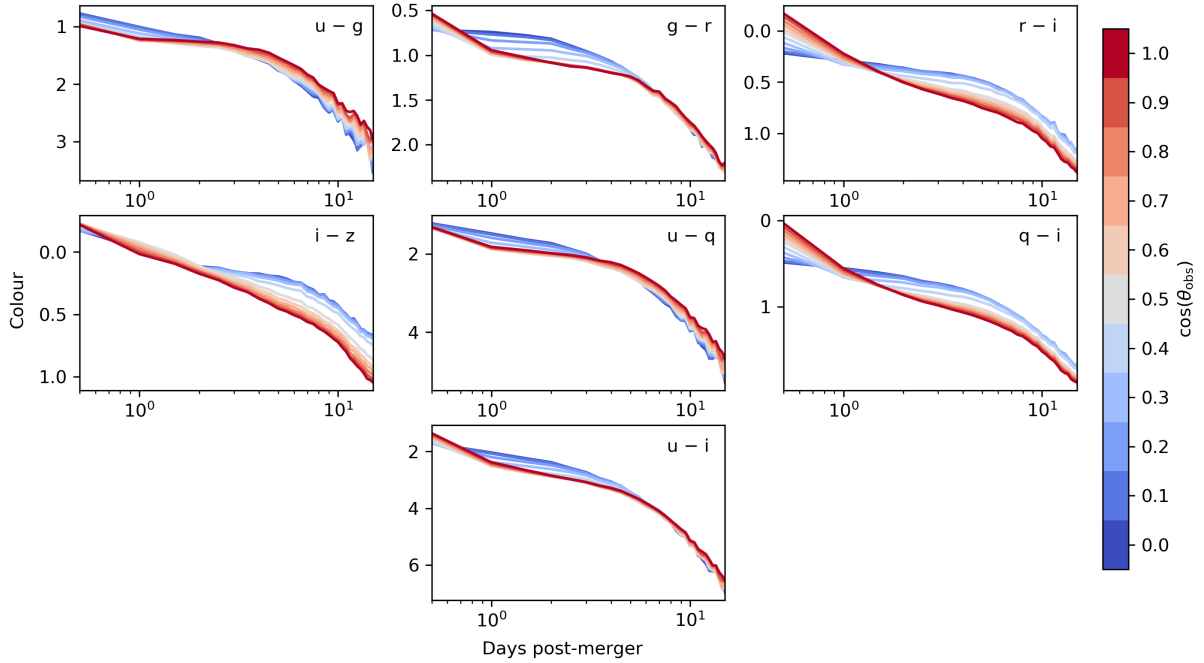


Figure 4.8: Colour evolution for KN models having $M_{\text{ej}} = 0.05M_{\odot}$ and $\Phi = 30^{\circ}$ while varying the viewing angle.

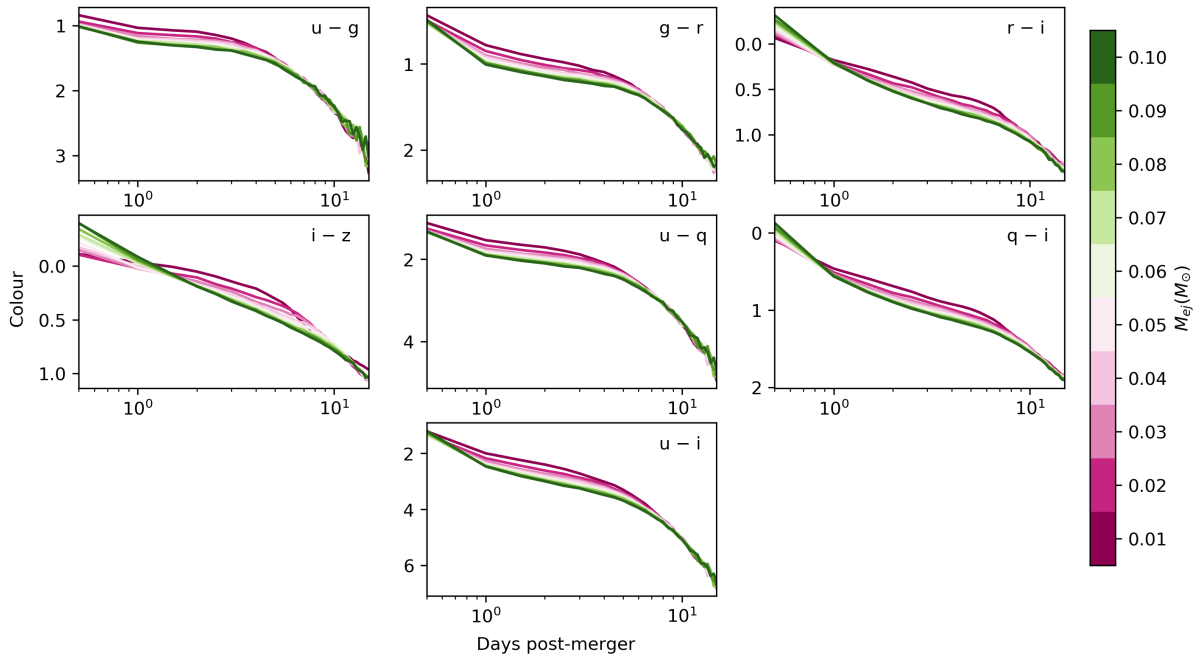


Figure 4.9: Same as previous figure but varying the ejecta mass with $\Phi = 30^{\circ}$ and $\cos(\theta_{\text{obs}}) = 1$.




Figure 4.10: Same as previous figure but varying Φ with $M_{\text{ej}} = 0.05M_{\odot}$ and $\cos(\theta_{\text{obs}}) = 1$.

for larger ejecta masses at early times (< 1 day). Smaller half-opening angles result in a reduced proportion of lanthanide-rich ejecta and thus produce bluer SEDs. This effect is clearly visible in fig. 4.10, particularly at times later than 2 days post-merger. For the blue colour combinations (containing u and g), we see that more lanthanide-rich ejecta (higher Φ) actually produce bluer KNe during early evolution (until 1 day).

4.2.3 Colour-colour evolution of KN models

A similar analysis but for colour-colour evolution is demonstrated in fig. 4.11, fig. 4.12 and fig. 4.13, comparing $u - q$ with $q - i$, and $u - g$ with $g - r$. Each dotted point along a line is associated with a time post-merger, in 0.5 day intervals. In fig. 4.11 we see in both plots that more equatorial viewing angles follow a more linear evolution in the colour-colour space, starting out bluer and ending redder in $u - g$ and $g - r$ than for polar viewing angles. More polar viewing angles start out bluer and end redder in $q - i$ and $r - i$, while evolving more quickly in these colours at earlier times as evidenced by their less steep downward evolution in the colour-colour space at early times. It appears as if all the models converge at the same point around $t = 1.5$ days in both plots (where the blue and red lines intersect), indicating that all viewing angles may have a similar SED at 1.5 days post-merger. In fig. 4.12 the general trend for all models is a more rapid evolution in $q - i$ and $r - i$ at early times. A cusp is apparent around $t = 3$ days for all models where the redward evolution in the colour-colour space drops more steeply. Prior to this, models with smaller ejecta masses are bluer in $u - q$ and $u - g$. There is general convergence across models at late times. Varying Φ as in fig. 4.13, we see that models with less lanthanide-rich ejecta have more prominent cusps around 2.5 days post-merger, where most of the models appear to converge. Lanthanide-rich models are redder at the start and finish of their $q - i$ and $r - i$ evolution, but bluer in $u - q$ and $u - g$ at their start of their evolution.

Having these colour and colour-colour evolution tracks available may be useful for a prompt KN parameter estimation as well as for distinguishing possible KN signals from




Figure 4.11: Colour-colour evolution for KN models having $M_{\text{ej}} = 0.05M_{\odot}$ and $\Phi = 30^{\circ}$ while varying the viewing angle. Each dotted point along a line is associated with a time post-merger, in 0.5 day intervals with the first ($t = 0.5$ days) and last ($t = 15$ days) times shown.

other astrophysical transients as shown in [49] and chapter 3 in [50].




Figure 4.12: Same as previous figure but varying ejecta mass with $\Phi = 30^\circ$ and $\cos(\theta_{\text{obs}}) = 1$.




Figure 4.13: Same as previous figure but varying Φ with $M_{ej} = 0.05M_{\odot}$ and $\cos(\theta_{\text{obs}}) = 1$.

Chapter 5

Conclusion

The prospect of finding an EM counterpart to a high significance GW event detected by the LVC on 2019 August 14 was made particularly promising given its early classification as a NSBH merger. Numerous groups conducted follow-up observations which were facilitated by the small sky-localisation of approximately 23 deg^2 at the 90% credible level. The MeerLICHT optical telescope in Sutherland observed the GW localisation each night for more than a week, covering at least 95% of the localisation probability in three bands (u , q , and i), often 3 or more times per night per band. Altogether 1484 one-minute exposures were taken over the course of our follow-up campaign. Analysis of the data (detailed in chapter 3) uncovered a gradient across the image plane in a large number of observations, meaning only $\sim 46\%$ of the total dataset could be used for scientific analysis. The median depth per exposure of our observations (in the AB magnitude system) was 18.98 in u , 20.02 in q , and 19.09 in i . A total of 46 astrophysical transients were found in the good reduced data – none of them new, with each source having prior detections in Gaia data. Our non-detection of any counterpart confirmed the findings of other groups.

The detection of GW190814 was a scientifically rich discovery in light of the recent LVC paper which stated that the source of the GW was a binary merger with the most unequal mass ratio yet discovered in gravitational waves, and that the lighter component sat squarely within the “MassGap”, making it either the lightest BH or heaviest NS ever discovered or even an as yet unknown exotic compact object. The mass ratio and component masses of such a system have challenged existing population synthesis models. Though it is unlikely that the source of the GW was a NSBH merger due to its high mass ratio, uncertainties in estimates of the maximum NS mass left open the possibility that the lighter component may have been a NS.

Continuing our analysis assuming a NSBH merger, in chapter 4 we used MeerLICHT limiting magnitudes to calculate the covered probability of our observations assuming an AT2017gfo-like KN at the distance of GW190814. Our covered probability in all 3 bands was negligible ($< 10^{-4}$), however it is highly probable that we would have been able to detect a KN approximately 5 times brighter than AT2017gfo. Furthermore, we used our limiting magnitudes to investigate the mass ejecta-viewing angle parameter space of KN models produced by the time-dependent 3D Monte Carlo code POSSIS. For KNe with a half-opening angle of 30° we found that an ejecta mass of $0.05M_\odot$ and edge-on viewing angle had the greatest probability of being observed, though it is still very low ($p \sim 10^{-3}$). The POSSIS models were also used to explore the colour and colour-colour evolution of KNe in MeerLICHT’s 6 optical bands. Both MeerLICHT and BlackGEM will be used extensively to find transients and observe their multi-colour evolution. Our exploration of the KN

colour and colour-colour space could be a valuable diagnostic tool in constraining the parameters of possible KNe discovered with these (and other) telescopes.

5.1 Future work

The loss of a significant fraction of the raw data taken with MeerLICHT was an unfortunate occurrence. Given the efforts of other groups it was highly unlikely that a counterpart to the GW would be found, but nevertheless, the possibility remains that there are transients “lurking” in the lost data, waiting to be found. Paul Vreeswijk (of Radboud University) has recently been working on background correction techniques for MeerLICHT images, and there is every possibility that the lost data could be salvaged and re-analysed for transients. Analysis using nightly co-additions of images taken in a particular filter and of a particular field could allow us to probe even deeper.

Though motivated by the strong science case of finding a GW counterpart, this project also served as an opportunity to test the transient detection capabilities of MeerLICHT. Going forward, the detailed analysis of such a dataset will be rendered obsolete by the addition of a machine learning component to the software pipeline, freeing the astronomer to focus on the science rather than the instrumentation and imaging aspects. This project enabled us to scrutinise MeerLICHT’s data products and evaluate its performance in preparation for the machine learning pipeline addition.

With respect to GW follow-up, we were fortunate that the sky localisation of the GW was so small, especially since the majority of O3 detections had much larger localisations. Though MeerLICHT does have a large field-of-view when compared to most optical telescopes, it is wholly insufficient for these larger localisations. The increase in sensitivity of the current generation of ground-based GW detectors and the addition of KAGRA will greatly reduce the median size of the 90% credibility regions, making follow-up more efficient. The BlackGEM array is particularly well-poised to take advantage of these gains in detector sensitivity, especially since it can cover three times the sky area as MeerLICHT, with the possibility that the array may be expanded to 15 such telescopes.

A paper on some of the findings of this work is in preparation.

Bibliography

- [1] B. P. Abbott et al. “Observation of Gravitational Waves from a Binary Black Hole Merger”. In: *Physical Review Letters* 116.6, 061102 (Feb. 2016), p. 061102. DOI: [10.1103/PhysRevLett.116.061102](https://doi.org/10.1103/PhysRevLett.116.061102). arXiv: [1602.03837](https://arxiv.org/abs/1602.03837) [gr-qc].
- [2] B. P. Abbott et al. “GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral”. In: *Physical Review Letters* 119.16, 161101 (Oct. 2017), p. 161101. DOI: [10.1103/PhysRevLett.119.161101](https://doi.org/10.1103/PhysRevLett.119.161101). arXiv: [1710.05832](https://arxiv.org/abs/1710.05832) [gr-qc].
- [3] B. P. Abbott, R. Abbott, and T. D. Abbott. “Multi-messenger Observations of a Binary Neutron Star Merger”. In: *The Astrophysical Journal Letters* 848.2, L12 (Oct. 2017), p. L12. DOI: [10.3847/2041-8213/aa91c9](https://doi.org/10.3847/2041-8213/aa91c9). arXiv: [1710.05833](https://arxiv.org/abs/1710.05833) [astro-ph.HE].
- [4] B. P. Abbott, R. Abbott, and T. D. Abbott. “GW190425: Observation of a Compact Binary Coalescence with Total Mass $\sim 3.4 M_{\odot}$ ”. In: *The Astrophysical Journal Letters* 892.1, L3 (Mar. 2020), p. L3. DOI: [10.3847/2041-8213/ab75f5](https://doi.org/10.3847/2041-8213/ab75f5). arXiv: [2001.01761](https://arxiv.org/abs/2001.01761) [astro-ph.HE].
- [5] The LIGO Scientific Collaboration et al. “GW190412: Observation of a Binary-Black-Hole Coalescence with Asymmetric Masses”. In: *arXiv e-prints*, arXiv:2004.08342 (Apr. 2020), arXiv:2004.08342. arXiv: [2004.08342](https://arxiv.org/abs/2004.08342) [astro-ph.HE].
- [6] R. Abbott et al. “GW190814: Gravitational Waves from the Coalescence of a 23 Solar Mass Black Hole with a 2.6 Solar Mass Compact Object”. In: *The Astrophysical Journal Letters* 896.2, L44 (June 2020), p. L44. DOI: [10.3847/2041-8213/ab960f](https://doi.org/10.3847/2041-8213/ab960f).
- [7] David J. Kaup. “Klein-Gordon Geon”. In: *Physical Review* 172.5 (Aug. 1968), pp. 1331–1342. DOI: [10.1103/PhysRev.172.1331](https://doi.org/10.1103/PhysRev.172.1331).
- [8] Pawel O. Mazur and Emil Mottola. “Gravitational vacuum condensate stars”. In: *Proceedings of the National Academy of Science* 101.26 (June 2004), pp. 9545–9550. DOI: [10.1073/pnas.0402717101](https://doi.org/10.1073/pnas.0402717101). arXiv: [gr-qc/0407075](https://arxiv.org/abs/gr-qc/0407075) [gr-qc].
- [9] Carl L. Rodriguez, Sourav Chatterjee, and Frederic A. Rasio. “Binary black hole mergers from globular clusters: Masses, merger rates, and the impact of stellar evolution”. In: *Physical Review D* 93.8, 084029 (Apr. 2016), p. 084029. DOI: [10.1103/PhysRevD.93.084029](https://doi.org/10.1103/PhysRevD.93.084029). arXiv: [1602.02444](https://arxiv.org/abs/1602.02444) [astro-ph.HE].
- [10] LIGO Scientific Collaboration et al. “Advanced LIGO”. In: *Classical and Quantum Gravity* 32.7, 074001 (Apr. 2015), p. 074001. DOI: [10.1088/0264-9381/32/7/074001](https://doi.org/10.1088/0264-9381/32/7/074001). arXiv: [1411.4547](https://arxiv.org/abs/1411.4547) [gr-qc].
- [11] F. Acernese et al. “Advanced Virgo: a second-generation interferometric gravitational wave detector”. In: *Classical and Quantum Gravity* 32.2, 024001 (Jan. 2015), p. 024001. DOI: [10.1088/0264-9381/32/2/024001](https://doi.org/10.1088/0264-9381/32/2/024001). arXiv: [1408.3978](https://arxiv.org/abs/1408.3978) [gr-qc].

- [12] Kentaro Somiya. “Detector configuration of KAGRA—the Japanese cryogenic gravitational-wave detector”. In: *Classical and Quantum Gravity* 29.12, 124007 (June 2012), p. 124007. DOI: [10.1088/0264-9381/29/12/124007](https://doi.org/10.1088/0264-9381/29/12/124007). arXiv: [1111.7185](https://arxiv.org/abs/1111.7185) [gr-qc].
- [13] Masaru Shibata and Kenta Hotokezaka. “Merger and Mass Ejection of Neutron Star Binaries”. In: *Annual Review of Nuclear and Particle Science* 69.1 (2019), pp. 41–64. DOI: [10.1146/annurev-nucl-101918-023625](https://doi.org/10.1146/annurev-nucl-101918-023625). eprint: <https://doi.org/10.1146/annurev-nucl-101918-023625>. URL: <https://doi.org/10.1146/annurev-nucl-101918-023625>.
- [14] James M. Bardeen, William H. Press, and Saul A. Teukolsky. “Rotating Black Holes: Locally Nonrotating Frames, Energy Extraction, and Scalar Synchrotron Radiation”. In: *The Astrophysical Journal* 178 (Dec. 1972), pp. 347–370. DOI: [10.1086/151796](https://doi.org/10.1086/151796).
- [15] Christopher S. Reynolds. “Observing black holes spin”. In: *Nature Astronomy* 3 (Jan. 2019), pp. 41–47. DOI: [10.1038/s41550-018-0665-z](https://doi.org/10.1038/s41550-018-0665-z). arXiv: [1903.11704](https://arxiv.org/abs/1903.11704) [astro-ph.HE].
- [16] C. Barbieri et al. “Light-curve models of black hole - neutron star mergers: steps towards a multi-messenger parameter estimation”. In: *Astronomy & Astrophysics* 625, A152 (May 2019), A152. DOI: [10.1051/0004-6361/201935443](https://doi.org/10.1051/0004-6361/201935443). arXiv: [1903.04543](https://arxiv.org/abs/1903.04543) [astro-ph.HE].
- [17] R. D. Blandford and R. L. Znajek. “Electromagnetic extraction of energy from Kerr black holes.” In: *Monthly Notices of the Royal Astronomical Society* 179 (May 1977), pp. 433–456. DOI: [10.1093/mnras/179.3.433](https://doi.org/10.1093/mnras/179.3.433).
- [18] F. K. Thielemann et al. “Neutron Star Mergers and Nucleosynthesis of Heavy Elements”. In: *Annual Review of Nuclear and Particle Science* 67 (Oct. 2017), pp. 253–274. DOI: [10.1146/annurev-nucl-101916-123246](https://doi.org/10.1146/annurev-nucl-101916-123246). arXiv: [1710.02142](https://arxiv.org/abs/1710.02142) [astro-ph.HE].
- [19] Stephane Goriely, Andreas Bauswein, and Hans-Thomas Janka. “r-process Nucleosynthesis in Dynamically Ejected Matter of Neutron Star Mergers”. In: *The Astrophysical Journal Letters* 738.2, L32 (Sept. 2011), p. L32. DOI: [10.1088/2041-8205/738/2/L32](https://doi.org/10.1088/2041-8205/738/2/L32). arXiv: [1107.0899](https://arxiv.org/abs/1107.0899) [astro-ph.SR].
- [20] Daniel Kasen et al. “Origin of the heavy elements in binary neutron-star mergers from a gravitational-wave event”. In: *Nature* 551.7678 (Nov. 2017), pp. 80–84. DOI: [10.1038/nature24453](https://doi.org/10.1038/nature24453). arXiv: [1710.05463](https://arxiv.org/abs/1710.05463) [astro-ph.HE].
- [21] B. D. Metzger et al. “Electromagnetic counterparts of compact object mergers powered by the radioactive decay of r-process nuclei”. In: *Monthly Notices of the Royal Astronomical Society* 406.4 (Aug. 2010), pp. 2650–2662. DOI: [10.1111/j.1365-2966.2010.16864.x](https://doi.org/10.1111/j.1365-2966.2010.16864.x). arXiv: [1001.5029](https://arxiv.org/abs/1001.5029) [astro-ph.HE].
- [22] C. Barbieri et al. “Electromagnetic counterparts of black hole-neutron star mergers: dependence on the neutron star properties”. In: *European Physical Journal A* 56.1, 8 (Jan. 2020), p. 8. DOI: [10.1140/epja/s10050-019-00013-x](https://doi.org/10.1140/epja/s10050-019-00013-x). arXiv: [1908.08822](https://arxiv.org/abs/1908.08822) [astro-ph.HE].
- [23] Steven Bloemen et al. “MeerLICHT and BlackGEM: custom-built telescopes to detect faint optical transients”. In: *Ground-based and Airborne Telescopes VI*. Ed. by Helen J. Hall, Roberto Gilmozzi, and Heather K. Marshall. Vol. 9906. International Society for Optics and Photonics. SPIE, 2016, pp. 2118–2126. DOI: [10.1117/12.2232522](https://doi.org/10.1117/12.2232522). URL: <https://doi.org/10.1117/12.2232522>.


- [24] Paul J. Groot. “The multi-colour dynamic Universe explored”. In: *Nature Astronomy* 3 (Dec. 2019), pp. 1160–1160. DOI: [10.1038/s41550-019-0964-z](https://doi.org/10.1038/s41550-019-0964-z).
- [25] R. A. Scalzo et al. “The SkyMapper Transient Survey”. In: *Publications of the Astronomical Society of Australia* 34, e030 (July 2017), e030. DOI: [10.1017/pasa.2017.24](https://doi.org/10.1017/pasa.2017.24). arXiv: [1702.05585](https://arxiv.org/abs/1702.05585) [[astro-ph.IM](#)].
- [26] Barak Zackay, Eran O. Ofek, and Avishay Gal-Yam. “Proper Image Subtraction—Optimal Transient Detection, Photometry, and Hypothesis Testing”. In: *The Astrophysical Journal* 830.1, 27 (Oct. 2016), p. 27. DOI: [10.3847/0004-637X/830/1/27](https://doi.org/10.3847/0004-637X/830/1/27). arXiv: [1601.02655](https://arxiv.org/abs/1601.02655) [[astro-ph.IM](#)].
- [27] E. Bertin and S. Arnouts. “SExtractor: Software for source extraction.” In: *Astronomy and Astrophysics Supplement* 117 (June 1996), pp. 393–404. DOI: [10.1051/aas:1996164](https://doi.org/10.1051/aas:1996164).
- [28] Daniëlle Pieterse. “Telescope characterisation of the wide-field optical telescope MeerLICHT”. MA thesis. Radboud University Nijmegen, 2019.
- [29] Kerry Paterson. “Detecting optical transients and variables with MeerLICHT”. PhD thesis. Astronomy department, University of Cape Town, 2018.
- [30] Gaia Collaboration et al. “Gaia Data Release 2. Summary of the contents and survey properties”. In: *Astronomy & Astrophysics* 616, A1 (Aug. 2018), A1. DOI: [10.1051/0004-6361/201833051](https://doi.org/10.1051/0004-6361/201833051). arXiv: [1804.09365](https://arxiv.org/abs/1804.09365) [[astro-ph.GA](#)].
- [31] Dustin Lang et al. “Astrometry.net: Blind Astrometric Calibration of Arbitrary Astronomical Images”. In: *The Astronomical Journal* 139.5 (May 2010), pp. 1782–1800. DOI: [10.1088/0004-6256/139/5/1782](https://doi.org/10.1088/0004-6256/139/5/1782). arXiv: [0910.2233](https://arxiv.org/abs/0910.2233) [[astro-ph.IM](#)].
- [32] E. Bertin. “Automated Morphometry with SExtractor and PSFEx”. In: *Astronomical Data Analysis Software and Systems XX*. Ed. by I. N. Evans et al. Vol. 442. Astronomical Society of the Pacific Conference Series. July 2011, p. 435.
- [33] J. Jonas and MeerKAT Team. “The MeerKAT Radio Telescope”. In: *MeerKAT Science: On the Pathway to the SKA*. Jan. 2016, p. 1.
- [34] R. Fender et al. “ThunderKAT: The MeerKAT Large Survey Project for Image-Plane Radio Transients”. In: *MeerKAT Science: On the Pathway to the SKA*. Jan. 2016, p. 13. arXiv: [1711.04132](https://arxiv.org/abs/1711.04132) [[astro-ph.HE](#)].
- [35] B. Stappers and M. Kramer. “An Update on TRAPUM”. In: *MeerKAT Science: On the Pathway to the SKA*. Jan. 2016, p. 9.
- [36] Leo P. Singer and Larry R. Price. “Rapid Bayesian position reconstruction for gravitational-wave transients”. In: *Physical Review D* 93.2 (2016). ISSN: 2470-0029. DOI: [10.1103/physrevd.93.024013](https://doi.org/10.1103/physrevd.93.024013). URL: <http://dx.doi.org/10.1103/PhysRevD.93.024013>.
- [37] The LIGO Scientific Collaboration and the Virgo Collaboration. *GCN Circular 25324*. URL: <https://gcn.gsfc.nasa.gov/gcn3/25324.gcn3>.
- [38] The LIGO Scientific Collaboration and the Virgo Collaboration. *GCN Circular 25333*. URL: <https://gcn.gsfc.nasa.gov/gcn3/25333.gcn3>.
- [39] J. Veitch et al. “Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library”. In: *Physical Review D* 91.4 (2015). ISSN: 1550-2368. DOI: [10.1103/physrevd.91.042003](https://doi.org/10.1103/physrevd.91.042003). URL: <http://dx.doi.org/10.1103/PhysRevD.91.042003>.

- [40] Shaon Ghosh et al. “Tiling strategies for optical follow-up of gravitational-wave triggers by telescopes with a wide field of view”. In: *Astronomy & Astrophysics* 592 (2016), A82. ISSN: 1432-0746. DOI: [10.1051/0004-6361/201527712](https://doi.org/10.1051/0004-6361/201527712). URL: <http://dx.doi.org/10.1051/0004-6361/201527712>.
- [41] M. B. Taylor. “STILTS - A Package for Command-Line Processing of Tabular Data”. In: *Astronomical Data Analysis Software and Systems XV*. Ed. by C. Gabriel et al. Vol. 351. Astronomical Society of the Pacific Conference Series. 2006, p. 666.
- [42] V. Lipunov et al. “MASTER Transient Discovery Report for 2019-08-19”. In: *Transient Name Server Discovery Report 2019-1541* (Aug. 2019), p. 1.
- [43] K. Ackley et al. “Observational constraints on the optical and near-infrared emission from the neutron star-black hole binary merger S190814bv”. In: *arXiv e-prints*, arXiv:2002.01950 (Feb. 2020), arXiv:2002.01950. arXiv: [2002.01950](https://arxiv.org/abs/2002.01950) [[astro-ph.SR](#)].
- [44] V. A. Villar et al. “The Combined Ultraviolet, Optical, and Near-infrared Light Curves of the Kilonova Associated with the Binary Neutron Star Merger GW170817: Unified Data Set, Analytic Models, and Physical Implications”. In: *The Astrophysical Journal Letters* 851.1, L21 (Dec. 2017), p. L21. DOI: [10.3847/2041-8213/aa9c84](https://doi.org/10.3847/2041-8213/aa9c84). arXiv: [1710.11576](https://arxiv.org/abs/1710.11576) [[astro-ph.HE](#)].
- [45] B. P. Gompertz et al. “The Diversity of Kilonova Emission in Short Gamma-Ray Bursts”. In: *The Astrophysical Journal* 860.1, 62 (June 2018), p. 62. DOI: [10.3847/1538-4357/aac206](https://doi.org/10.3847/1538-4357/aac206). arXiv: [1710.05442](https://arxiv.org/abs/1710.05442) [[astro-ph.HE](#)].
- [46] Igor Andreoni et al. “GROWTH on S190814bv: Deep Synoptic Limits on the Optical/Near-infrared Counterpart to a Neutron Star?Black Hole Merger”. In: *The Astrophysical Journal* 890.2, 131 (Feb. 2020), p. 131. DOI: [10.3847/1538-4357/ab6a1b](https://doi.org/10.3847/1538-4357/ab6a1b). arXiv: [1910.13409](https://arxiv.org/abs/1910.13409) [[astro-ph.HE](#)].
- [47] M. Bulla. “POSSIS: predicting spectra, light curves, and polarization for multidimensional models of supernovae and kilonovae”. In: *Monthly Notices of the Royal Astronomical Society* 489.4 (Nov. 2019), pp. 5037–5045. DOI: [10.1093/mnras/stz2495](https://doi.org/10.1093/mnras/stz2495). arXiv: [1906.04205](https://arxiv.org/abs/1906.04205) [[astro-ph.HE](#)].
- [48] S. Dhawan et al. “Constraining the Observer Angle of the Kilonova AT2017gfo Associated with GW170817: Implications for the Hubble Constant”. In: *The Astrophysical Journal* 888.2, 67 (Jan. 2020), p. 67. DOI: [10.3847/1538-4357/ab5799](https://doi.org/10.3847/1538-4357/ab5799). arXiv: [1909.13810](https://arxiv.org/abs/1909.13810) [[astro-ph.HE](#)].
- [49] J. van Roestel et al. “The Palomar Transient Factory Sky2Night programme”. In: *Monthly Notices of the Royal Astronomical Society* 484.4 (Apr. 2019), pp. 4507–4528. DOI: [10.1093/mnras/stz241](https://doi.org/10.1093/mnras/stz241). arXiv: [1902.03768](https://arxiv.org/abs/1902.03768) [[astro-ph.HE](#)].
- [50] Jan van Roestel. “Optical Variability with the Palomar Transient Factory”. PhD thesis. Radboud University, Nov. 2018.


Appendix A

A.1 Light curves of transients in table 3.3


The coordinates (RA, Dec) of each source is given below its respective light curve, in the ICRS system.




(24.03886, -34.50995)




(25.14406, -34.29417)




(23.98727, -33.64103)




(24.49559, -33.68419)




(21.20969, -32.68997)




(20.47796, -31.84143)




(21.79382, -32.23449)




(21.66897, -33.30406)




(23.72503, -33.04030)




(23.67084, -33.32235)




(23.91714, -33.37142)




(24.84324, -33.07292)




(25.77454, -32.71110)




(24.33520, -32.94701)




(20.85303, -31.30931)




(22.49961, -30.67453)




(12.37119, -27.38717)




(11.77327, -27.73729)




(11.98638, -27.75949)




(13.04404, -28.49332)




(13.67282, -28.23188)




(13.32939, -27.55863)




(11.07424, -26.07290)




(13.19452, -26.56870)




(13.16781, -26.54159)




(10.25079, -24.88181)




(9.56137, -24.98390)




(10.60015, -24.52526)




(11.03642, -23.93112)




(11.14006, -24.39554)




(11.22108, -24.71085)




(11.43848, -24.75447)




(12.73859, -23.91721)




(11.38689, -24.33887)




(12.55806, -24.03788)




(12.95900, -24.65356)




(12.90751, -23.95032)




(13.06334, -25.10326)




(14.42013, -24.02047)




(11.82490, -22.75212)




(10.84540, -23.28073)




(11.25533, -22.39226)




(10.32200, -23.03056)



(11.98874, -22.35416)



(12.16165, -22.41826)



(13.97614, -23.25184)