The multiplicity of 250-μm *Herschel* sources in the COSMOS field

Jillian M. Scudder,[1][*] Seb Oliver,[1] Peter D. Hurley,[1] Matt Griffin,[2] Mark T. Sargent,[1] Douglas Scott,[3] Lingyu Wang[4] and Julie L. Wardlow[5]

[1]*Astronomy Centre, Department of Physics and Astronomy, University of Sussex, Brighton BN1 9QH, England*
[2]*Cardiff School of Physics and Astronomy, Cardiff University, Queens Buildings, The Parade, Cardiff CF24 3AA, Wales*
[3]*Department of Physics and Astronomy, University of British Columbia, 6224 Agricultural Road, Vancouver, BC V6T 1Z1, Canada*
[4]*SRON Netherlands Institute for Space Research, Landleven 12, NL-9747 AD Groningen, the Netherlands*
[5]*Department of Physics, Centre for Extragalactic Astronomy, Durham University, South Road, Durham DH1 3LE, England*

ABSTRACT

We investigate the multiplicity of extragalactic sources detected by the *Herschel Space Observatory* in the COSMOS field. Using 3.6- and 24-μm catalogues, in conjunction with 250-μm data from *Herschel*, we seek to determine if a significant fraction of *Herschel* sources are composed of multiple components emitting at 250 μm. We use the XID+ code, using Bayesian inference methods to produce probability distributions of the possible contributions to the observed 250-μm flux for each potential component. The fraction of *Herschel* flux assigned to the brightest component is highest for sources with total 250-μm fluxes <45 mJy; however, the flux in the brightest component is still highest in the brightest *Herschel* sources. The faintest 250-μm sources (30–45 mJy) have the majority of their flux assigned to a single bright component; the second brightest component is typically significantly weaker, and contains the remainder of the 250-μm source flux. At the highest 250-μm fluxes (45–110 mJy), the brightest and second brightest components are assigned roughly equal fluxes, and together are insufficient to reach 100 per cent of the 250-μm source flux. This indicates that additional components are required, beyond the brightest two components, to reproduce the observed flux. 95 per cent of the sources in our sample have a second component that contains more than 10 per cent of the total source flux. Particularly for the brightest *Herschel* sources, assigning the total flux to a single source may overestimate the flux contributed by around 150 per cent.

Key words: galaxies: high redshift – galaxies: photometry – galaxies: statistics.

1 INTRODUCTION

As part of the drive to understand the evolution of galaxies over cosmic history, we look back at systems which existed at early times. One of the more puzzling sub-sets of galaxies which have been identified is the population of very dusty, infrared (IR) luminous, presumably highly star-forming galaxies at high redshifts. (For a review, see Blain et al. 2002; Casey, Narayanan & Cooray 2014.) As they contain a large amount of heated dust, these objects are typically very bright at far-infrared (FIR) or sub-mm wavelengths. Thus, the majority of the discoveries of these objects have come from FIR and/or sub-mm facilities. Attempts to identify optical counterparts to these sub-mm sources have historically been stymied by that same dust, as these objects are usually heavily extincted in the optical regime. The 'typical' galaxy of this type has so far been found to reside at redshifts of 1–3 (Eales et al. 2000; Smail et al.

2002; Chapman et al. 2005), though this is likely to partially depend on the selection wavelength (e.g. Zavala, Aretxaga & Hughes 2014; Béthermin et al. 2015).

This population of dusty, strongly star-forming sources at high redshift has drawn numerous comparisons to the local ultraluminous infrared galaxy (ULIRG) population (e.g. Lonsdale, Farrah & Smith 2006), as ULIRGs are the closest local analogues to the high-redshift systems in terms of their unusually high star formation rates for their redshift. ULIRGs, defined as having a total IR luminosity above 10^{12} L$_\odot$,[1] are similarly a rare sub-set of the local galaxy

---

[1] ULIRGs are defined by a fixed luminosity threshold; in the local Universe, this selects a rare sub-set of galaxies above the star formation rate (SFR)–stellar mass relation, but at high redshifts, the same luminosity threshold permits the inclusion of galaxies with less extreme SFR for their stellar mass (Lin et al. 2007; Ellison et al. 2013). The dusty, strongly star-forming galaxies at high redshift are comparable to local ULIRGs in that they are both populations which are significantly offset from the SFR–stellar mass relation.

[*] E-mail: j.scudder@sussex.ac.uk

population, typified by unusually high star formation rates for their stellar mass, and contain a large volume of molecular gas (Combes et al. 2013). If ULIRGs are truly local analogues of the high-redshift systems, one might expect to find a substantial number of the high-redshift galaxies are involved in some form of close galaxy–galaxy encounter, as suggested by several authors (e.g. Tacconi et al. 2008; Younger et al. 2010; Zamojski et al. 2011; Wiklind et al. 2014). Unfortunately, it is difficult to unambiguously identify physical pairs at high redshift, and the fraction of the total samples which are ascribed to merger events therefore varies significantly. Both Tacconi et al. (2008) and Younger et al. (2010) find that their sources can all be explained with mergers, but sample only a few galaxies at the very brightest end. Zamojski et al. (2011) has a much larger sample size, and estimates that their merger fraction should be 80 per cent, though this is extrapolated and could vary between 60 per cent and 91 per cent. If the galaxies are not involved in an interaction to trigger their star formation, it remains unclear what secular mechanisms could drive a galaxy to such high star formation rates at early times in the universe.

It is possible that the extremely high star formation rates implied for these FIR/sub-mm sources are boosted by the presence of a second, nearby source (either in projection or physically associated) which is also contributing to the FIR flux. Single-dish observatories at the sub-mm and FIR wavelengths (whether space or ground based) typically have very large beam sizes [full width at half-maximum (FWHM) of the order of ~ 20 arcsec, which at a redshift of 2 is ~ 170 kpc], making it difficult to unambiguously identify a single optically detected galaxy as the source of the dust emission; the high space density of optical sources makes it likely that many potential counterparts will fall within the beam. Hughes et al. (1998) and Downes et al. (1999) both find 3–5 optical sources at the approximate location (i.e. within the positional error) of a sub-mm source (see also Dunlop et al. 2004). If multiple optical components are contributing roughly equally to the FIR flux, then a single FIR source may well be comprised of multiple, potentially independent physical sources. It is difficult to determine whether the optical counterparts are indeed contributing to the sub-mm flux, and often a single counterpart is assigned responsibility for the entirety of the sub-mm flux (typically the counterpart with the smallest projected separation), though in some cases, other methods, such as a Bayes factor, have been used (e.g. Roseboom et al. 2009). The issue of multiple possible sources will always be present in some fraction of the data, as there will be line-of-sight projection effects which will layer unrelated foreground or background sources in the same region of sky as the intrinsic FIR source.

To overcome the challenge of identifying the correct source of the FIR flux, several methodologies have been developed. One approach is to use data from other wavelengths, such as the radio, where higher resolution imaging of the same area of the sky can be obtained, making use of the strong radio–FIR correlation (e.g. Helou, Soifer & Rowan-Robinson 1985). Radio observations provided some of the earliest hints that a significant fraction of the galaxy population might be composed of blended components (e.g. Ivison et al. 1998, 2000; Smail et al. 2000; Ivison et al. 2002). More recently, it has become possible to make use of interferometric facilities such as the Sub-Millimetre Array (SMA; e.g. Younger et al. 2007, 2008; Wang et al. 2011), the Plateau de Bure Interferometer (e.g. Smolčić et al. 2012), or the Atacama Large Millimetre/sub-millimetre Array (ALMA; e.g. Hodge et al. 2013), which have the power to observe some of these systems with sub-arcsecond resolution.

Recent ALMA results have added to a growing pool of bright FIR-selected and sub-mm sources which resolve into multiple compo-

nents. Whereas previous work had focused on deep optical imaging or small samples (e.g. Younger et al. 2010; Wang et al. 2011; Chen et al. 2013) to find potential counterparts, ALMA can observe in the FIR with high resolution, avoiding complications which arise when comparing with other wavelengths (e.g. Hodge et al. 2013). However, the fraction of sources which divide into multiple components varies substantially between studies. Multiplicity fractions (defined as the fraction of total sources which resolve into multiple components) appear to range between 10 per cent (Chen et al. 2013) and 50 per cent (Hodge et al. 2013), with other studies finding values within that range (e.g. Barger et al. 2012; Hezaveh et al. 2013, measuring 25 per cent and > 20 per cent fractions respectively).[2] This variation is likely due in part to different sample selection between studies, which select their targets based on flux measurements based anywhere between 450 μm (e.g. Chen et al. 2013) and 1.2 mm (e.g. Hezaveh et al. 2013), though largely the systems are selected close to 850 μm (Hodge et al. 2013; Karim et al. 2013; Simpson et al. 2015). Another possible source of variation is the indication that the multiplicity fraction may be dependent on the FIR flux observed, with brighter sources more likely to have multiple components (Karim et al. 2013). While ALMA has the advantage of directly observing with high resolution in the FIR, it is limited by sky coverage. In this work, we will make use of the availability of multi-wavelength coverage of large areas of the sky, which allows us to make use of higher resolution data to determine how many FIR sources divide into multiple flux-emitting components.

If a significant fraction of the FIR sources are in fact multiple components, then the individual galaxies are much less extreme systems than previously thought, regardless of whether or not the component galaxies are physically interacting with each other. In order to develop a better understanding of how the population of star-forming galaxies has evolved over time, we must determine the exact nature of these seemingly extreme systems. We wish to undertake a comprehensive survey of the multiplicity of these dusty star-forming galaxies using a new approach: Bayesian statistical analysis of multiwavelength coverage of a given field. This study will be the first to undertake such a statistical survey of bright *Herschel* sources.

In Section 2, we describe the data used for this study, including a brief description of the multi-wavelength catalogues. In Section 3, we describe the selection criteria used for the data, and in Section 4, we describe the methods used to analyse this data. In Section 5, we discuss the implications of our results in the context of previous works, and in Section 6, we summarize our conclusions.

2 THE DATA

We base our study on data from the *Herschel Space Observatory* (Pilbratt et al. 2010). *Herschel* was equipped with two continuum imaging instruments, the Photodetector Array Camera and Spectrometer (PACS; Poglitsch et al. 2010), and the Spectral and Photometric Imaging Receiver (SPIRE; Griffin et al. 2010). Both instruments observed a number of legacy fields as part of the *Herschel* MultiTiered Extragalactic Survey (HerMES;[3] Her-MES Collaboration et al. 2012), which provided FIR coverage at

---

[2] For the works cited here, the beam size is not substantially different between studies, varying between 3 arcsec (Chen et al. 2013) and 1.5 arcsec (Hodge et al. 2013), but may play a role when comparing to other analyses.
[3] hermes.sussex.ac.uk

100 and 160 μm in PACS, and at 250, 350, and 500 μm from SPIRE.

A new EU-funded project, the *Herschel* Extragalactic Legacy Project (HELP;[4] e.g. Vaccari & HELP Consortium 2015), aims to collect and reprocess data products across multiple wavelengths for all regions contained within the *Herschel* survey fields, including fields in surveys not contained in HerMES, such as *Herschel*-Astrophysical Terahertz Large Area Survey (H-ATLAS; Eales et al. 2010) and *Herschel* Stripe 82 Survey (HerS; Viero et al. 2014). The main goal of HELP is to produce homogeneously calibrated, value-added catalogues covering roughly 1000 deg², which matches individual sources across wavelengths, allowing for statistical, multiwavelength studies of the local-to-intermediate redshift galaxy population. In this work, we benefit from the multi-wavelength ancillary catalogues already compiled by HELP. For the current work, we limit our study to the COSMOS field (Scoville et al. 2007), a field observed by both PACS and SPIRE instruments as part of the PACS Evolutionary Probe program (Lutz et al. 2011) and HerMES (HerMES Collaboration et al. 2012), and which has excellent multiwavelength coverage over the 2 deg² field. We focus on the 250-μm data catalogue from the SPIRE instrument, as it has the highest SPIRE resolution, and of the three SPIRE bands will be the least affected by blending. Our *Herschel* source list is produced by the STARFINDER software and is a completely blind catalogue (Smith et al. 2012; Wang et al. 2014). The *Herschel* maps are generated by the SMAP/SHIM software and are a standard HerMES Data Release 2 product (Levenson et al. 2010).

Traditionally, studies of multiplicity within the *Herschel* catalogues have been tackled using the positions of 24-μm sources from the *Spitzer Space Telescope* (Werner et al. 2004) as more precise indicators of the position of likely sources within a 250-μm data set (e.g. Brisbin et al. 2010; Roseboom et al. 2010). A similar method was used by Smith et al. (2011), employing positional information to assign H-ATLAS source flux to a counterpart in the Sloan Digital Sky Survey Data Release 7 (Abazajian et al. 2009). For this study, we use the positions of sources from the 24-μm *Spitzer* Multiband Imaging Photometer catalogue (Le Floc'h et al. 2009) and the 3.6-μm *Spitzer* Infrared Array Camera (IRAC) catalogue, which were produced as part of S-COSMOS (Sanders et al. 2007). The 24-μm catalogue was extracted from the maps using the SEXTRACTOR software (Bertin & Arnouts 1996) to identify sources, with fluxes determined through the use of DAOPHOT (Stetson 1987). The 24-μm catalogue is estimated to be over 90 per cent complete above a flux of 80 μJy (Le Floc'h et al. 2009); our catalogue is flux-limited to contain only sources above 150 μJy, so our sample should be complete.

The 3.6-μm catalogue contains 5σ detections down to a flux of 0.9 μJy, and is also extracted from the map using the SEXTRACTOR software (Bertin & Arnouts 1996; Sanders et al. 2007). The 3.6-μm data provide a much higher number of potential counterparts than the 24-μm data, since it is much higher resolution (1.7 arcsec at 3.6 μm versus 6 arcsec at 24 μm). By using the 3.6-μm data, our resolution information has improved by a factor of 10, relative to the 250-μm data. The 3.6- and 24-μm sources will be used to help us probe the question of how frequently FIR sources divide into multiple components; by using the higher resolution information from shorter wavelength data, we can use the positions of these sources as prior information for our analysis of the SPIRE photometry.

3 SAMPLE SELECTION

We wish to select a sample of *Herschel* sources that have higher resolution counterparts in order to examine the frequency with which they divide into multiple components. We therefore identify a sample of 250-μm-detected sources that have at least one 3.6-μm source and at least one 24-μm source within the size of the *Herschel* beam at 250 μm (18.1 arcsec; Griffin et al. 2010).

We begin our sample selection by including in our master catalogues only reasonably secure detections. We include all 250-μm sources within the COSMOS field above a flux limit of 30 mJy, which corresponds to approximately 5σ above the confusion noise limit in the *Herschel* observations (Nguyen et al. 2010; HerMES Collaboration et al. 2012). Confusion noise dominates over all other sources of noise in the SPIRE COSMOS data; see Nguyen et al. (2010). As a first quality control measure on both the 24-μm catalogue (Le Floc'h et al. 2009) and the 3.6-μm catalogue, we require that the flux values recorded are above zero, since negative flux values are unphysical. For the 3.6-μm catalogue (Sanders et al. 2007), we additionally require that the source flux flag is set to 0, which indicates that the source flux is not contaminated by a nearby bright source in the IRAC catalogue.[5]

For each 250-μm source in the COSMOS field, we search for those that have at least one 24-μm source that passes our quality control measures within the 18.1-arcsec FWHM of *Herschel* at 250 μm. We require a 24-μm detection for our sample, as many previous works (including standard HerMES products) have required 24-μm detections at the locations of a *Herschel* source to ensure that their *Herschel* detections are reliable astrophysical sources. For consistency with these methodologies, we impose the same criterion on our sample. However, the 24-μm data are of limited spatial resolution, so to increase the resolution information, we also require a 3.6-μm detection. We use the same procedure to select out all 250-μm sources that have at least one 3.6-μm source within the 250-μm beam. We then cross-match the 3.6-μm matches with the 24-μm matches, to select those 250-μm sources that have potential components in both the 24-μm and 3.6-μm catalogues. This cross-matching marks the selection of the final sample of 250-μm sources.

For each of our final 250-μm sources, if a 24-μm source is within 6.0 arcsec of a 3.6-μm source matched to the same 250-μm source (6 arcsec is the resolution of *Spitzer* at 24 μm), we remove the 24-μm location, making the assumption that the 3.6-μm source will be a more accurate position for the same physical object. This process does not exclude single-counterpart *Herschel* sources from our sample, as we do not require more than one possible counterpart per detected *Herschel* source. In some cases, the 24-μm object is farther than 6.0 arcsec away from the nearest 3.6-μm source. In these cases, we keep both the 3.6- and the 24-μm object positions for our list of potential contributing components for the 250-μm source.[6]

Our final sample consists of 360 *Herschel* 250-μm sources within the COSMOS field, each having at least one possible contributing source at 24 and at 3.6 μm. The limiting criterion is requiring a 24-μm source within the 250-μm beam, which restricts us to 14 per cent of the total 250-μm catalogue in this field. However, this sample represents 98 per cent of the 250-μm sources with 24-μm detections within the *Herschel* beam; the inclusion of the 3.6-μm criterion only reduces the sample by 2 per cent. Although

---

[5] In the IRAC catalogue, flag = 0 is shortened to 'good data'.
[6] Within our sample, this is true in 25/360 of our 250-μm sources.

we did not explicitly require that each 250-μm source have more than one potential component at 24 and 3.6 μm, the fewest number of potential components matched to each *Herschel* source is 3. The maximum number of potential components within the 250-μm beam is 26; the median number of candidate sources is 14 per 250-μm source.

3.1 Using XID+ to assign component fluxes

With many potential contributing sources per 250-μm source, we must find a way to assign flux to each individual component if we are to investigate the properties of the underlying components. To do so, we make use of the XID+ software (Hurley et al. 2016), a prior-driven source-extraction tool, and a significant improvement upon the previous XID software, also known as DESPHOT (Roseboom et al. 2010; Wang et al. 2014). This new software's methodology and validation are fully described in Hurley et al. (2016), so we refer the reader to that work for a complete explanation of the methods, with a brief explanation here.

The XID+ algorithm uses the 250-, 350-, and 500-μm maps from *Herschel*, along with the positional information of known sources to find the most likely distribution of flux among these sources. XID+ assigns the appropriate *Herschel* flux to each of the sources at the known input positions. XID+ places a positive prior on the fluxes; in other words, it cannot assign negative flux to a source to improve the fit, as this has no physical meaning for observations. The code has been tested on synthetic maps of the COSMOS field used here and is found to recover injected source fluxes accurately and with a reliable error estimate. The previous code DESPHOT is found to less reliably recover the true flux of injected sources, and also systematically underestimates the errors on those measurements (Hurley et al. 2016). This difference is largely due to DESPHOT's use of a maximum likelihood algorithm which gradually switches on sources to increase the goodness of fit; sources which do not improve the fit quality are assigned a flux of 0 mJy. We consider XID+ to be a considerable improvement over the methods used in DESPHOT.

For this work, since we have already selected a sub-sample of the 250-μm catalogue, we do not need to run XID+ on the full COSMOS map, which would require us to simultaneously fit several tens of thousands of sources. Instead, we select a 180 arcsec × 180 arcsec box surrounding each 250-μm source.[7] We then find all of the 3.6-μm sources within that field (typically several hundred in number), in addition to those identified as potential counterparts, in order to ensure a good fit to the entire region, which will ensure that we do not neglect to assign flux to other sources near our 250-μm source of interest. We neglect the 24-μm positions here as the typical field contains no additional 24-μm sources which are not coincident with a 3.6-μm source, so the inclusion of the 24-μm sources is unlikely to improve the fit. This full list is then used as the 'known sources' for the positional prior information required by XID+. XID+ is then run on each of the 360 unique fields surrounding our 250-μm sources, and the resulting flux probability distributions are stored for all of the fitted sources, not just those which fall within the beam of the strong 250-μm source. However, for the analysis which follows, we investigate only those potential counterparts which were flagged as within 18.1 arcsec of the *Herschel* source.

XID+ produces a full Bayesian statistical posterior probability distribution through Markov Chain Monte Carlo sampling, which allows strongly correlated sources to be identified. In cases where more than one component is likely to be producing FIR flux, the posterior allows us to examine all of the possible permutations of flux assignment between these components. The full posterior also provides accurate error estimates on the flux estimates. In instances where the final probability distribution is Gaussian, or nearly Gaussian, taking the median, 16th and 84th percentiles of the probability distribution will provide an accurate summary of the typical flux and its error. Where the fluxes assigned are very low (they are constrained to remain above zero), the median will tend to be biased towards slightly higher values than are truly typical (i.e. the calculated median may not coincide with the mode of the distribution). The raw probability distributions can be viewed through a corner plot (Foreman-Mackey et al. 2014, Fig. 1). Fig. 1 shows the probability distributions for each of the 12 components of the 250-μm source in a typical field. The majority of the components are given fluxes that are very close to zero, with the bulk of the flux assigned to component 5 and component 8. Of the two sources with significant flux, component 5 dominates the flux budget, with a median flux of 46.7 mJy, and the 16th and 84th percentiles falling at 44.8 and 48.6 mJy, respectively. 16 and 84th percentiles for component 8's flux set a range between 10.7 and 13.5 mJy, with component 2 receiving a median of 2.0 mJy, with the 84th percentile setting an upper limit of 3.5 mJy of flux. The contours which populate the off-diagonal show the correlations between components. There is a slight correlation between component 5 and component 8, but the probability distributions for these two components are relatively narrow, indicating that the possible error due to degeneracy in the solution is minimal. In the case where two components are completely degenerate, the probability between sources will appear as a tightly correlated diagonal across a wide range of possible fluxes. This poor constraint on parameter space would be reflected in the width between the 16th and 84th percentile values being very large. We have inspected all the corner visualisations for the 360 fields in our sample, and find that flux is typically assigned to 2–4 components out of a median number of 14 potential components.[8]

We note that XID+ uses a prior on the fluxes which is flat in parameter space. While this prior has the effect of preferring solutions with a larger number of lower flux objects relative to higher flux objects, this is a physically motivated prior, as very bright sources should be rarer objects than objects which are intrinsically fainter. We also note that XID+ is limited to assigning fluxes to objects which are provided as known sources. However, if the known source list was very incomplete, XID+ would find a poorer solution to the map, which would be reflected in a broader width of the posterior distribution. We account for the uncertainty in the posterior distribution throughout this work, and do not believe that the 3.6-μm catalogue is a significantly incomplete representation of the FIR-emitting population (Duivenvoorden et al. 2016); these effects should not bias our results.

XID+ also includes statistical methods to test for convergence, allowing us to be certain that the parameter space has been appropriately explored. We require that for each chain, the individual runs have completely converged. Following the guidelines presented in

---

[7] The exact size of this box is somewhat arbitrary, but this ensures that the entire 250-μm source, along with a buffer region surrounding it, are fit by XID+ for each 250-μm source.

[8] The full set of triangle plots is available for public download at the following: http://jmscudder.github.io/XID-figures.
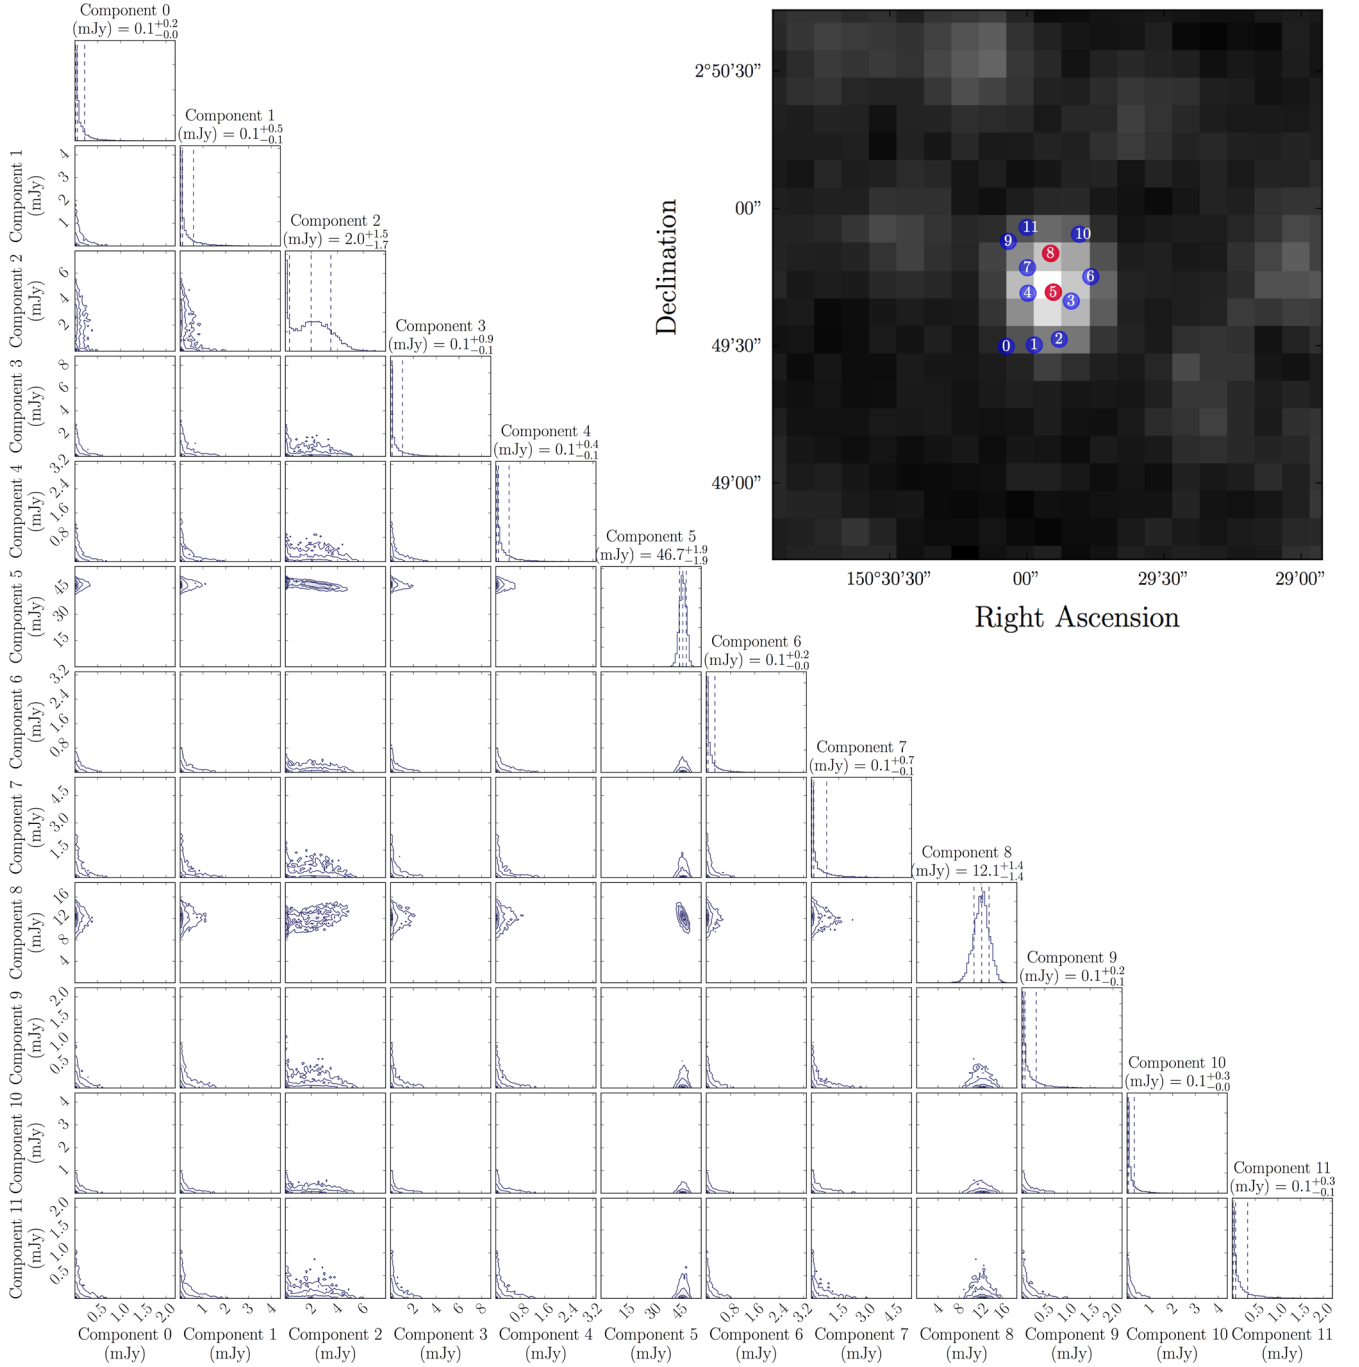
Figure 1. A triangle plot of the correlations between the 12 potential components for a single 250-μm source (ID 10009). Histograms along the top diagonal edge indicate the probability distribution of flux for the component in that column. Above each histogram, the 50th percentile value for the component is labelled, plus/minus the distance to 16th and 84th percentiles. Vertical dashed lines indicate the 16th, median, and 84th percentiles (from left to right in each panel). The interior panels give the contours of correlation between pairs of components. In the top-right corner, the 250-μm map is shown in grey scale, with the positions of the known potential counterparts shown as coloured circles. Each potential component is labelled as in the triangle plot. We can see that only two sources are assigned significant flux, with the majority of the flux assigned to Source 5. We note that the axis ranges are not consistent across all components, though all axes begin at 0 as their lower limit. With the exception of components 5 and 8, and to a lesser extent component 2, all other sources have distributions peaking very close to 0 flux. Sources 5 and 8 are therefore marked in red in the inset figure.
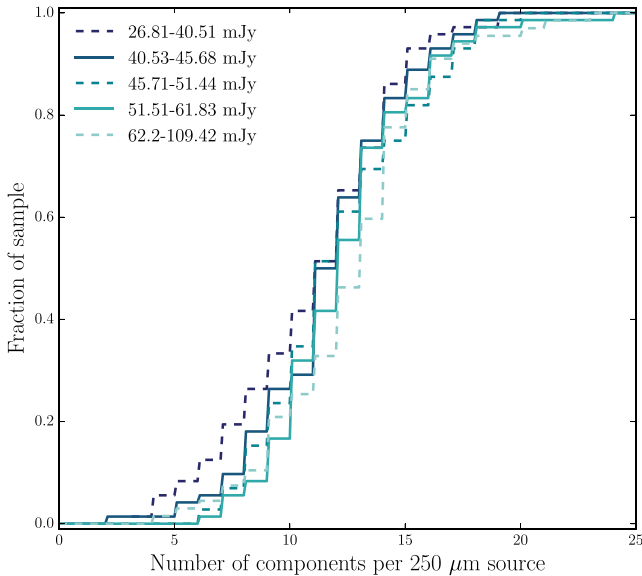
Figure 2. The cumulative distribution of the number of all potential components per 250-µm source, binned as a function of total 250-µm source flux. Each bin contains an equal fraction of the total sample. Line style alternates between dashed and solid for added visual clarity. A KS test indicates that all flux bins are consistent with being drawn from the same parent population as all other bins.

Hurley et al. (2016), we constrain $\hat{R}^9$ to be <1.2, and $n_{\rm eff}{}^{10} >40$. For the work done here, we use four MCMC chains of 2000 iterations each, which is sufficient to meet the convergence criteria for 99.5 per cent of all runs.

4 ANALYSIS

Our ultimate goal in analysing the results of XID+ is to determine what fraction of these 250-µm sources have strong multiple components, and if there are systematic differences within the population that we have selected, which depend on the brightness of the 250-µm source. To investigate whether we have a fair base of comparison within the sample, we first wish to determine if there are systematic variations in the number of potential counterparts surrounding the selected 250-µm sources as a function of 250-µm flux. If there is a systematic change in the number of components, then any properties which are plotted as a function of the fraction of total flux or as a fraction of the number of components will also be systematically biased. If we find that there is no strong systematic shift with total 250-µm source flux, then the raw number of potential components will not be a source of bias in our analysis of trends with total 250-µm flux.[11] In Fig. 2, we plot the cumulative distribution of the number of potential components found per 250-µm source, divided

into bins of total flux across all potential components. Bins here are defined such that each bin has an equal fraction (20 per cent) of the total sample. We have also binned the data evenly in total flux, and the results are broadly unchanged, but at high 250-µm source flux, bins suffer from low number statistics.

We note that the lowest flux limit in Fig. 2 is slightly below 30 mJy, which was our original flux limit described in Section 3. Here, and throughout this work, we use a total 250-µm calculated from the output of XID+. For each iteration of XID+, we sum the fluxes of all components identified as within 18.1 arcsec. We then use the median of the distribution of summed fluxes as our total flux throughout this analysis. In the case of the lowest fluxes plotted here, 4 mJy of flux has been assigned to a component outside of the beam of the 250-µm source, lowering the total flux below our nominal 30 mJy threshold.

To determine if the bins are statistically consistent with each other, we run a Kolgolmorov–Smirnoff (KS) test on the distributions. A KS test finds the maximum vertical separation of the cumulative distributions of two populations to determine if the two populations are consistent with being drawn from the same parent population. If the maximum distance is too large, the distributions are inconsistent. Each of the five bins is compared with the remaining four in turn. We find that the KS test indicates that all bins are consistent with being drawn from the same parent sample as all other bins.[12] The 250-µm sources are therefore not populated with a significantly different number of sources in the field as a function of total flux. If the 250-µm sources have consistent numbers of potential components as a function of total flux, we should be able to make robust relative comparisons as a function of total flux for other properties of the components.

In Fig. 3, we plot the fraction of the total 250-µm flux assigned to each component of our full sample. Each data point indicates the median of the posterior distribution of possible fluxes. We plot the data as a density histogram, where colour indicates the log of the number of components within that bin, indicated by the colour bar on the far right of the figure. Note that, as Fig. 1 implied, there are a large number of components with near-zero fluxes, which results in a concentration of potential components along the x-axis (visible as the dark blue stripe). On the right-hand panel of Fig. 3, we show the 1D histogram of component flux fractions (with the horizontal axis plotted in log), collapsed along the total 250-µm source flux axis. This panel emphasizes the abundance of components assigned very little flux. Curiously, we note that the upper-right-hand corner of this diagram appears unpopulated by data points. We also note that in our density histogram, the low 250-µm source flux end appears to have a much higher concentration of points with zero flux than is present at the high 250-µm source flux end. To verify that this is not simply due to the flux distribution of 250-µm sources, we divide the sample into the same set of equally populated bins as Fig. 2 and compare the distribution of flux fractions in each bin to all other bins with a KS test. For the full sample, the KS test between the three lowest total flux bins and the highest total flux bin exclude the null hypothesis of being drawn from the same parent population at much more than 5σ. The second highest flux bin is also inconsistent with the highest flux bin at $>4\sigma$. This indicates that the distributions of flux fraction assigned to components are statistically distinct from each other as a function of total 250-µm flux. A priori, this might lead us to expect that a larger fraction of

---

[9] $\hat{R}$ is the potential scale variation as estimated from a measurement of the posterior variance within chains of the MCMC method, and measures how well each run has found the global minimum.

[10] $n_{\rm eff}$ is the effective number of samples, and checks if the global minimum has been well traversed. Both convergence metrics are fully explained in (Hurley et al. 2016). The limits used here are found to ensure that the code has converged.

[11] We note that a uniform underrepresentation of potential components across all 250-µm sources will still have an effect on, e.g. number counts, where the absolute number of sources is important. For the current work, we are not sensitive to this issue.

[12] Our most discrepant value (between the lowest and highest 250-µm flux bins) give p-values of 0.144, or less than 1.5σ.
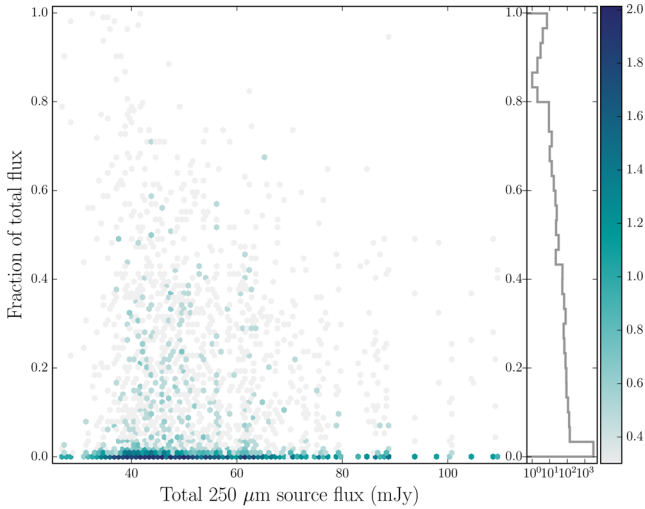
Figure 3. Left-hand panel: the unbinned data points for all potential components in our sample, plotted as a density histogram. We plot the fraction of the total 250-µm source flux assigned to each potential component, as a function of the total 250-µm flux. Colour indicates the log of the number of points contained within that bin, as indicated by the colour bar on the far right. Points close to 1 on the *y*-axis indicate that they were assigned almost 100 per cent of the total flux, meaning that they are the only source assigned significant flux. For high total flux 250-µm sources, there is a deficit of components at large flux fractions (i.e. high-flux objects tend to have that flux divided amongst more components, explored in more detail in text). The right-hand panel shows a histogram of the distribution of flux fractions across all 250-µm source fluxes. The *x*-axis in the right-hand panel is in log space.

the low 250-µm flux components are being assigned very low-flux values, and that at high 250-µm flux values, significant flux is being divided among more components. We investigate this further in the following sections.

4.1 The brightest components

As our sample has a large number of components which have been assigned extremely low flux, it may be that components assigned low-flux values are swamping any interesting trends within the fraction of the sample which has been assigned significant flux. To investigate, we use the brightest component from each of the 360 250-µm sources in our sample.

It is possible that some of the 250-µm sources have two (or more) nearly equal brightness components. In this case, identifying a physical component as the brightest component is potentially unwise, as the posterior distribution may indicate that it is equally likely for the second component to be the brightest. The same physical component will not necessarily always be the brightest in all of the iterations. To circumvent this problem, we select the brightest component on a sample by sample basis. We therefore accumulate a distribution of brightest component fluxes for each of the 2000 iterations. We use the median of this distribution as the 'brightest component flux' for that 250-µm source. We place these brightest components into equally populated bins of 250-µm flux, and plot them as a function of 250-µm source flux in Fig. 4. The horizontal error bars in Fig. 4 reflect the width of each bin. Below each point, we label the number of 250-µm sources contained within that bin.

To make an estimate of the error on the median flux of the brightest components in each bin, we calculate the errors in two ways. The
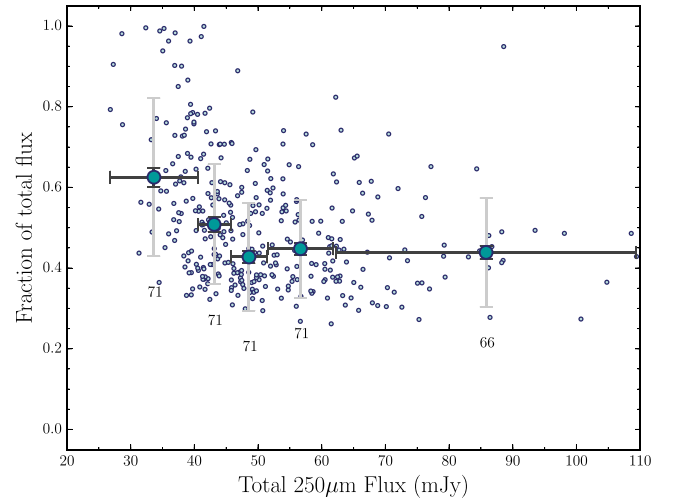

Figure 4. For each 250-µm source, we select the component with the highest assigned flux fraction. We plot these brightest components as a function of the total 250-µm flux in small background points. Horizontal error bars on the larger, binned points indicate the width of the bin, and the numbers below each point indicate the number of brightest components within each bin. Pale vertical error bars show the standard deviation on the distribution of all possible medians, including both the flux uncertainty and the dispersion of brightest component contained within the bin. Black vertical error bars (typically of comparable size to the data point) are calculated by dividing the pale error bars by $\sqrt{n}$, applicable as each data point is an independent measurement and measurement error should not introduce a systematic bias. We refer to the black $\sqrt{n}$ error bars for the remainder of this work. The brightest components show a trend towards higher fractions of 250-µm flux as total 250-µm flux decreases. A Spearman Rank correlation test on the unbinned data gives a *p*-value of 1.86×10^{-12}, which excludes the null hypothesis of no correlation at $>6\sigma$.

first takes into account the full posterior distribution of the brightest component flux measurement itself. To calculate the error, we resample the posterior of each brightest component contained within a given 250-µm bin. If the brightest component flux varies strongly across iterations, we will capture this variation with the resampling. For each sampling of the distribution of each brightest component, we recalculate the median of that bin. This is repeated for all 2000 iterations, to fully sample the distribution of each brightest component. The standard deviation on the distribution of resampled median binned points is plotted as the pale grey vertical error bars in Fig. 4. This error bar reflects not only the dispersion within the population contained within that bin, but also the possible measurement error due to variations in how well constrained the fluxes are.

However, each of the brightest components within the bin is given a completely independent flux measurement, as they are fitted by XID+ separately. We do not expect the measurement error to have any systematic effect on the medians, so we plot the above described error bars divided by the square root of the number of brightest components contained within the bin (shown in Fig. 4 as the number plotted below the error bars). This $\sqrt{n}$ error bar is plotted in black. Typically, the $\sqrt{n}$ set of error bars is approximately the size of the plotted median point or smaller. While we plot both error bars for the remainder of this work, we expect that for trends within the sample, the $\sqrt{n}$ error bars are a better reflection of our expectation of the uncertainty in the median itself, and refer to those error bars for the remainder of this work.
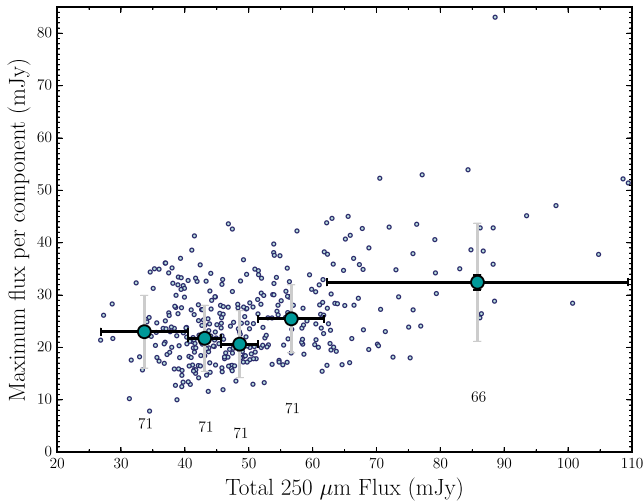
Figure 5. The same as Fig. 4, but now plotting the flux value assigned to the brightest component instead of the fraction of the total flux. The brightest components at high 250-μm fluxes are still brighter than the low 250-μm flux components. A Spearman Rank correlation test on the unbinned data gives a *p*-value of 1.08×10^{-13}, indicating an exclusion of the null hypothesis of no correlation at $>7\sigma$.

In the binned values, it is clear that the median flux fraction in the brightest component increases at the low total 250-μm flux end of the diagram. To determine if this is statistically significant, we run a Spearman Rank correlation test on the unbinned data. This tests if two properties are correlated by checking if both quantities increase with each other, but has the advantage of not requiring that this correlation be linear, as (e.g.) the Pearson test does. By running the correlation on the unbinned data, we are also insensitive to our choice of bin. The Spearman Rank correlation test returns a *p*-value of 1.86×10^{-12}, which excludes the null hypothesis of no correlation at $>6\sigma$. The Spearman Rank test also returns a correlation coefficient of -0.36.

We also investigate the raw flux assigned to each component. The exact relationship between brightest component flux and 250-μm flux can be inferred from Fig. 4. As the flux fraction dependence on 250-μm flux is not linear, we plot the dependence of brightest component flux on 250-μm flux explicitly in Fig. 5 for clarity. This figure is otherwise constructed identically to Fig. 4. Here we see a strong increase in the amount of flux assigned to the brightest component as a function of total flux, with the Spearman Rank test applied to the unbinned data indicating a *p*-value of 1.08×10^{-13}, an exclusion of the null hypothesis (no correlation) at $>7\sigma$. The Spearman Rank correlation test indicates a correlation coefficient of 0.37.

Together with Fig. 4, Fig. 5 shows that while the brightest components at lower total 250-μm flux often have a single component with a significant fraction of the total flux of the 250-μm source, those components are still assigned significantly lower fluxes than the components at high total 250-μm fluxes, which are sharing flux amongst multiple objects. While bright 250-μm sources divide their flux into multiple objects, the brightest of the high-flux components are still brighter than the lower 250-μm flux brightest components.[13]

---

[13] We note that our binned points at fluxes lower than 50 mJy have a poor positive correlation. However, this is due to our choice of bins, and disappears if we use bins which are equal in width. The full population (plotted as background small points) does not have a strong break at 50 mJy.
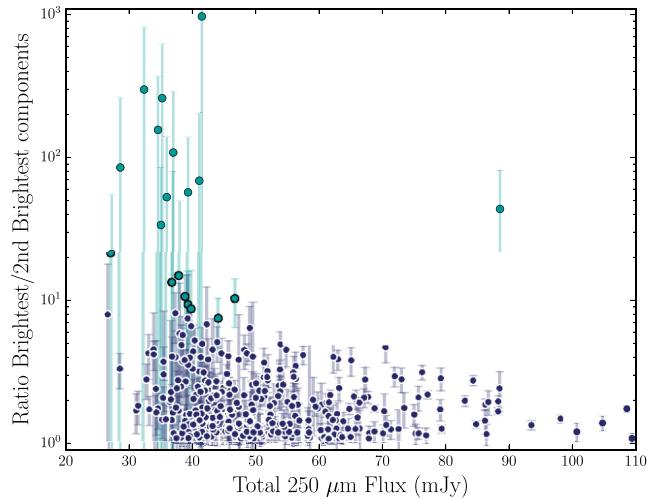

Figure 6. The ratio of the brightest component to the second brightest component in log, as a function of the total 250-μm source flux. Values close to 1 indicate that the brightest component is nearly matched by the second brightest component. Conversely, values much higher than 1 indicate that the brightest component is significantly brighter than the second brightest component. Dark blue points indicate that the second brightest component contains more than 10 per cent of the total 250-μm source flux. Cyan points indicate that the second brightest component contains <10 per cent of the total 250-μm source flux. Vertical error bars on both sets of points are the standard deviation on the distribution of all possible ratios for the full 2000 samplings of brightest and second brightest components. A Spearman rank correlation test on the full sample of points plotted gives a *p*-value of 4.54×10^{-9}, which excludes the null hypothesis of no correlation at $>5\sigma$. A Spearman Rank correlation test on the sub-sample with second brightest components containing >10 per cent of the total flux gives a *p*-value of 1.34×10^{-5}, which excludes the null hypothesis of no correlation at $>4\sigma$.

4.2 Brightest component versus second brightest component

As an additional check into the behaviour of these brightest components, we investigate the relationship between the fraction of flux assigned to the brightest component and the fraction assigned to the second brightest component. We identify the second brightest component in an identical fashion to identifying the brightest component. In Fig. 6, the ratio of the brightest component's flux to the second brightest component's flux is plotted in log against the total 250-μm source flux for each 250-μm source.

There is an additional consideration at this stage, which is that some fraction of our sample will have the majority of their flux assigned only to a single source, with extremely low fluxes found in the second brightest component. In order to consider more physically meaningful second components, we set an 'insignificant flux' threshold to 10 per cent of the total flux, which gives sources of >3 mJy. In Fig. 6, we plot the full sample, but colour code according to the fraction of the total 250-μm source flux contained within the second brightest component. The dark blue points are those where the second brightest component has more than 10 per cent of the total source flux. The cyan points indicate the 250-μm sources where the second brightest component contains less than 10 per cent of the total flux. Our vertical error bars are the standard deviation on the full distribution of ratios, calculated by iterating through the full distribution of brightest and second brightest component fluxes. These error bars therefore represent our full measurement uncertainty on each point. We note that those data points plotted in dark blue typically have significantly smaller error bars, indicating that all iterations of the code have arrived at very similar flux values
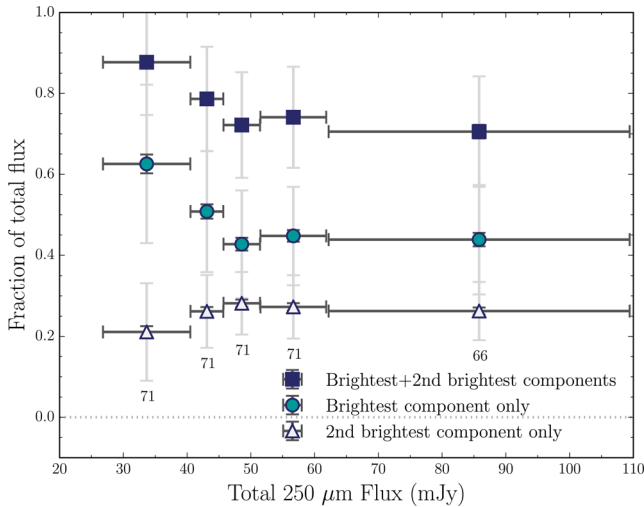
Figure 7. This figure is constructed identically to Fig. 4. We plot the brightest component's fraction of the total flux, as a function of the total flux, in cyan circles. In grey triangles, we plot the second brightest component's contribution to the total flux. In dark blue squares, we plot the fraction of the total flux contained in the brightest and the second brightest components combined. At the low total 250-μm flux end, the combination of the top two components accounts for approximately 90 per cent of the flux in the 250-μm source. At the high-total flux end, where we expect the brightest and second brightest components to be of nearly equal brightness, the top two sources account for approximately 70 per cent of the total flux. At intermediate total fluxes, the sum of the brightest two components is reasonably constant at 70 per cent of the total flux.

for the brightest and second brightest components, and that their probability distributions (as shown in Fig. 1) are narrow. We emphasize that the points plotted are the medians of each distribution, for each 250-μm source.

As implied by Figs 4 and 5, the majority of sources with very high brightest/second brightest component ratios are found at low 250-μm flux. As the total flux of the 250-μm source increases, the maximal ratio between brightest and second brightest components decreases and approaches unity (i.e. the brightest and second brightest components are approximately equal in brightness). At the low-total source flux end, some of the brightest components are 1000 times brighter than their next brightest component; however, these extremes are present only for those 250-μm sources where the second component has insignificant flux. This extreme difference between first and second brightest component also seems to dominate only where the total flux is very low (i.e. <40 mJy); above this level, the brightest component tends to be no more than 10 times brighter than the next brightest component. In this flux range (>40 mJy), components span the range from equally bright as their secondary component to several times brighter. A Spearman Rank correlation test (which is not sensitive to the specific shape of the correlation) on the total sample finds a correlation coefficient of -0.3, with a p-value of 4.54×10^{-9}, which excludes the null hypothesis at $>5\sigma$. The same test run on only the sub-sample with strong second brightest components (those points in dark blue) gives a p-value of 1.34×10^{-5}, which excludes the null hypothesis of no correlation at $>4\sigma$.

We plot the fraction of the total 250-μm source flux accounted for by the brightest two companions explicitly in Fig. 7. As previously plotted in Fig. 5, the fraction of the total flux due to the brightest component is shown in cyan circles; the second brightest component is plotted on its own in grey triangles. The sum of the

brightest two components is plotted in dark blue squares. The top two brightest components contribute approximately 80 per cent of the total 250-μm source flux at the high 250-μm total flux end, with only 20 per cent of the total flux remaining to be produced by additional companions. The typical contribution of the brightest two components is remarkably consistent between 70 and 80 per cent. The exception to this is the lowest total flux bin, which is consistent with the brightest two components containing 100 per cent of the flux. The second brightest component is also remarkably consistent, producing 20–25 per cent of the total 250-μm flux across the full 250-μm flux range.

We examine the fraction of all 250-μm sources where the second brightest component contains a very small fraction of the total flux. If the second component contains very little flux, we can consider this source to be a single component. If we continue to use our limit of the second brightest component containing at least 10 per cent of the total 250-μm flux, we find that only 5 per cent of all 250-μm sources are excluded. We therefore conservatively conclude that our sample has a multiplicity fraction of 95 per cent, although the exact value will depend on the fractional threshold imposed on the secondary component.

5 DISCUSSION

It is difficult to compare the results presented here directly to other studies, since the majority of studies in the FIR are selected (and observed) at different wavelengths than the current work. However, our results may be more directly compared in the context of the conclusions drawn from other studies, keeping in mind the variable sample size between studies. In general, our results are consistent with previous studies of multiplicity in FIR sources.

Problems with correctly identifying a shorter wavelength counterpart of a source detected in the FIR or sub-mm regime have been longstanding issues in this field. Both Hughes et al. (1998) and Downes et al. (1999) struggled with the optical counterpart selection of their sub-mm observations, finding multiple possible options within a reasonable radius. Hughes et al. (1998) had a set of five sub-mm detected sources, with optical counterparts either invisible, or with up to five potential counterparts within the beam. Downes et al. (1999) conducted deep follow-up imaging of the brightest of the Hughes et al. (1998) sources and proposed that the sub-mm source was associated with four components, but noted the presence of nine other sources within 6 arcsec of the sub-mm source. This difficulty then pushed the need for high-resolution observations in the FIR or nearby wavelengths; observing at the same wavelength would help determine the correct optical counterpart.

Early attempts to do so were somewhat hampered by low number statistics. One of the earlier high-resolution observations of sub-mm source multiplicity was undertaken at the SMA by Wang et al. (2011), who observed two sub-mm sources and found that they each source divided into two or three components. The conclusion drawn was that it was highly likely that sub-mm sources were being misidentified by being assigned to the nearest optical counterpart, and that all sub-mm sources would need to be re-observed with high-resolution capabilities to ensure that the correct counterpart was identified at other wavelengths.

Subsequent studies have been highly varied in sample size, selection wavelength, and in how the multiplicity fraction was determined. Early studies of multiplicity were necessarily of only a few objects at a time (e.g. Wang et al. 2011), but the majority of recent studies have been of samples of the order of 20–30 objects, much smaller than the sample of 360 presented in the current work.

For instance, Chen et al. (2013) had a sample of 20 sources observed by SCUBA-2 at 450-µm or 850-µm;[14] they determine that 2–3 (10–15 per cent) may be multiple components, given multiple nearby IRAC or 20 cm radio counterparts. This estimate was made by investigating offsets between the centroid of the FIR flux and the radio counterpart. However, one might expect this method to underestimate the true multiplicity fraction, as any multicomponent structure which had a strong component associated with the radio component, and a second component not associated with the radio flux, would not be counted. Hezaveh et al. (2013) have a similar sample size to Chen et al. (2013), finding four galaxies (originally detected at 1.4 mm at the South Pole Telescope[15]), which divide into directly imaged multiple components (in this case, strong gravitational lenses) out of a sample of 20, when observed with ALMA at 860 µm, implying a multiplicity fraction of at least 20 per cent.

The majority of recent studies have relied on direct detection of bright components with interferometric observations. However, this method also has its limitations. Smolčić et al. (2012) calculated that at least 15 per cent of their 28 objects detected with 870-µm observations with The Large Apex BOlometer CAmera (LABOCA) on the APEX 12-m single-dish telescope[16] were blends of multiple components, as the multiple components were directly imaged by the IRAM Plateau de Bure facility. A number of the objects from the single-dish data were completely undetected with the higher resolution data. Smolčić et al. (2012) concludes that it is very plausible that these non-detections are the result of the single dish flux being subdivided, so that the individual components are below the flux limit of their survey, which has a 3σ limit of $\sim$1.4 mJy. We note that this 3σ limit is approximately half as bright as the conservative threshold we impose in Section 4.2 to denote very low-flux levels. By counting the non-detections as having multiple components, the multiplicity fraction of Smolčić et al. (2012) increases to 40 per cent. This flux threshold limitation is in place for all such direct imaging campaigns to determine the multiplicity fraction. The majority of such campaigns have similar number statistics to the Smolčić et al. (2012) work; Bussmann et al. (2015) finds that of their 29 FIR-selected sources, 20 (69 per cent) are multicomponent in nature, down to a 5σ limit of 1 mJy. Similarly, Simpson et al. (2015) investigate a sample of 30 galaxies detected by the SCUBA-2 Cosmology Legacy Survey at 850 µm with ALMA, and find that 60 per cent of their sample is resolved into more than one component above a 4σ limit of $\sim$1 mJy, with four sources resolved into three or four components. Chen et al. (2016) similarly finds that bright 850-µm sources (>4 mJy) are split into multiple components at least 40 per cent of the time. While these smaller studies have multiplicity fractions that are significantly higher, they are still prone to the loss of low-flux counterparts, which may partially explain the difference between these results of 60–70 per cent and our result of a multiplicity fraction of >95 per cent.

With the advent of ALMA, larger surveys such as the ALMA LABOCA ECDFS Submillimeter Survey (ALESS) have been able to study samples of over a hundred sources (Hodge et al. 2013; Karim et al. 2013) in high resolution, which brings much larger number statistics than have previously been available. Hodge et al. (2013) find a multiplicity fraction of 35 per cent in a sample of 126 galaxies selected with LABOCA. This 35 per cent reflects the fraction which are directly detected above a 3σ threshold of

$\sim$1.5 mJy; considering the non-detections, this value increases to 50 per cent, similar to the results of Smolčić et al. (2012). This is, however, a significantly lower fraction than found by Simpson et al. (2015) and Bussmann et al. (2015). Simpson et al. (2015) point out that this discrepancy could well be due to the lower resolution of the ALESS sample relative to the Simpson et al. (2015) work, and that typically the single-dish fluxes of ALESS sources are lower than those of the sample probed in Simpson et al. (2015). Both of these aspects would make it more difficult to find lower flux companions near a bright source, and explain the lower multiplicity fraction relative to that which is reported in the current work.

An additional complication in comparing different studies is the increasing number of studies which suggest that the multiplicity fraction is dependent on the FIR or sub-mm source flux itself. A number of studies have now shown evidence that the higher the total source flux, the more likely the source is to be divided into multiple components (Hodge et al. 2013; Karim et al. 2013; Swinbank et al. 2014; Simpson et al. 2015). As Simpson et al. (2015) looked at higher flux objects than the ALESS team, one might expect them to find a higher multiplicity fraction. In particular, Karim et al. (2013) find that above a given flux limit, all sources are multiple components. Cowley et al. (2015) present theoretical work which also suggests that bright sources (>5 mJy) have the majority of their flux (90 per cent) contained within three to six components, but that this multiplicity declines as source flux declines. Similarly, another theoretical work (Hayward et al. 2013) finds that the vast majority of sub-mm sources should be blends of multiple components; above a synthesized 850-µm flux of 3 mJy, the sub-mm sources are dominated by at least two physical components. The findings of Hayward et al. (2013) are consistent with our results; the 250-µm sources at the faintest fluxes are most likely to be dominated by a single component, with the highest fluxes more likely to be divided into roughly equal flux components.

Of the existing studies, only a few have begun to examine the properties of the resolved components. Simpson et al. (2015) is one such, finding that the brightest component accounts for approximately 80 per cent of the total sub-mm flux, and that where there are two or more components, the secondary source tends to add approximately 25 per cent to the total flux. The brightest component in the simulations presented in Cowley et al. (2015) is 70 per cent of the total flux on average. This is significantly higher than we see here; typically our brightest component is 45 per cent of the total source flux, except at the lowest flux end, where the brightest component contains 60 per cent of the total flux. Consistent with Simpson et al. (2015) observational study, we find that the secondary component has a fairly consistent contribution to the total flux, of about 20–25 per cent across all 250-µm source fluxes. However, we typically reach fractions of 70–80 per cent only when considering the brightest two components jointly. The Simpson et al. (2015) sources are selected at 850 µm, so it is unclear how bright these sources would appear at 250 µm. If the Simpson et al. (2015) study had investigated sources similar to those found within our lowest 250-µm flux bins, a brightest component fraction of 80 per cent would be close to what we observe.

With the measurements of individual 250-µm fluxes associated with each 3.6-µm source, we can naturally extend this work to an analysis of whether or not the bright components are truly physically associated or not. Existing theoretical works suggest that the bright components are unlikely to be physically associated (Hayward et al. 2013; Cowley et al. 2015; Muñoz Arancibia et al. 2015) except in a small fraction of cases. However, as the goal of this first paper is to determine how many of the 250-µm sources might be expected

---

[14] SCUBA-2 has an FWHM beam size of 14.5 arcsec at 850 µm.
[15] The South Pole Telescope has an FWHM beam size of 1 arcmin at 1.4 mm.
[16] LABOCA has an FWHM beam size of 19 arcsec at 870 µm.

to divide into at least two components, a detailed observational analysis is beyond the scope of the current paper, and we defer it to a forthcoming work. Furthermore, our sample is typically divided into more counterparts than are traced by the 24-μm component (see also Hodge et al. 2013). Number counts derived only from 24-μm positions (e.g. Béthermin et al. 2012) will likely be affected by the findings presented in this work. Given that multiplicity changes as a function of total 250-μm flux density, it is not straightforward to estimate the exact influence; we also leave this to a future work.

6 CONCLUSIONS

We have undertaken an analysis of a sample of 360 different 250-μm sources above 30 mJy, found within the COSMOS field, which are selected to have at least one 24-μm counterpart and at least one 3.6-μm counterpart within the *Herschel* beam. We use the XID+ code, which statistically determines all possible allocations of 250-μm flux based on the positions of the potential 24- and 3.6-μm components. Our conclusions can be summarized as follows.

(i) We find that while there is no statistically significant difference in the raw number of potential counterparts at 3.6- and 24-μm as a function of 250-μm source flux, the brightest and faintest 250-μm sources have statistically distinct distributions of flux to their respective components.

(ii) At low 250-μm source fluxes (30–45 mJy), the brightest components are assigned a larger fraction of the total 250-μm source flux. Above a 250-μm source flux of 45 mJy, the brightest components only contain ∼45 per cent of the total 250-μm source flux.

(iii) The brightest components of the most luminous 250-μm sources are typically assigned more flux than the brightest components of faint 250-μm sources.

(iv) The most luminous 250-μm sources typically have brightest and second brightest components which are similar in flux; fainter 250-μm sources more often contain components which are very unequal in flux.

(v) If both the brightest and second brightest components are considered, at the low 250-μm flux end (<45 mJy), 90 per cent of the 250-μm source flux is accounted for, indicating that only one or two sources are required to assemble the observed 250-μm flux. At larger 250-μm fluxes, the two brightest components contain only 70 per cent of the total flux, indicating that further components are contributing significant flux.

(vi) Only 20/360 (5.6 per cent) have a secondary component that contains less than 10 per cent of the 250-μm flux, indicating that ∼95 per cent of our sample can be considered to be divided into multiple strong components.

REFERENCES

Abazajian K. N. et al., 2009, ApJS, 182, 543
Barger A. J., Wang W.-H., Cowie L. L., Owen F. N., Chen C.-C., Williams J. P., 2012, ApJ, 761, 89
Bertin E., Arnouts S., 1996, A&AS, 117, 393
Béthermin M. et al., 2012, A&A, 542, A58
Béthermin M., De Breuck C., Sargent M., Daddi E., 2015, A&A, 576, L9
Blain A. W., Smail I., Ivison R. J., Kneib J.-P., Frayer D. T., 2002, Phys. Rep., 369, 111
Brisbin D. et al., 2010, MNRAS, 409, 66
Bussmann R. S. et al., 2015, ApJ, 812, 43
Casey C. M., Narayanan D., Cooray A., 2014, Phys. Rev., 541, 45
Chapman S. C., Blain A. W., Smail I., Ivison R. J., 2005, ApJ, 622, 772
Chen C.-C., Cowie L. L., Barger A. J., Casey C. M., Lee N., Sanders D. B., Wang W.-H., Williams J. P., 2013, ApJ, 762, 81
Chen C.-C. et al., 2016, ApJ, 820, 82
Combes F., García-Burillo S., Braine J., Schinnerer E., Walter F., Colina L., 2013, A&A, 550, A41
Cowley W. I., Lacey C. G., Baugh C. M., Cole S., 2015, MNRAS, 446, 1784
Downes D. et al., 1999, A&A, 347, 809
Duivenvoorden S. et al., 2016, MNRAS, submitted
Dunlop J. S. et al., 2004, MNRAS, 350, 769
Eales S., Lilly S., Webb T., Dunne L., Gear W., Clements D., Yun M., 2000, AJ, 120, 2244
Eales S. et al., 2010, PASP, 122, 499
Ellison S. L., Mendel J. T., Scudder J. M., Patton D. R., Palmer M. J. D., 2013, MNRAS, 430, 3128

Foreman-Mackey D., Price-Whelan A., Ryan G., Emily Smith M., Barbary K., Hogg D. W., Brewer B. J., 2014, triangle.py v0.1.1, (available at http://dx.doi.org/10.5281/zenodo.11020)
Griffin M. J. et al., 2010, A&A, 518, L3
Hayward C. C., Behroozi P. S., Somerville R. S., Primack J. R., Moreno J., Wechsler R. H., 2013, MNRAS, 434, 2572
Helou G., Soifer B. T., Rowan-Robinson M., 1985, ApJ, 298, L7
HerMES Collaboration et al., 2012, MNRAS, 424, 1614
Hezaveh Y. D. et al., 2013, ApJ, 767, 132
Hodge J. A. et al., 2013, ApJ, 768, 91
Hughes D. H. et al., 1998, Nature, 394, 241
Hurley P. et al., 2016, MNRAS, submitted
Ivison R. J., Smail I., Le Borgne J.-F., Blain A. W., Kneib J.-P., Bezecourt J., Kerr T. H., Davies J. K., 1998, MNRAS, 298, 583
Ivison R. J., Dunlop J. S., Smail I., Dey A., Liu M. C., Graham J. R., 2000, ApJ, 542, 27
Ivison R. J. et al., 2002, MNRAS, 337, 1
Karim A. et al., 2013, MNRAS, 432, 2
Le Floc'h E. et al., 2009, ApJ, 703, 222
Levenson L. et al., 2010, MNRAS, 409, 83
Lin L. et al., 2007, ApJ, 660, L51
Lonsdale C. J., Farrah D., Smith H. E., 2006, Astrophysics Update 2, Springer, Praxis Publ., Chichester, p. 285
Lutz D. et al., 2011, A&A, 532, A90
Muñoz Arancibia A. M., Navarrete F. P., Padilla N. D., Cora S. A., Gawiser E., Kurczynski P., Ruiz A. N., 2015, MNRAS, 446, 2291
Nguyen H. T. et al., 2010, A&A, 518, L5
Pilbratt G. L. et al., 2010, A&A, 518, L1
Poglitsch A. et al., 2010, A&A, 518, L2
Roseboom I. G., Oliver S., Parkinson D., Vaccari M., 2009, MNRAS, 400, 1062
Roseboom I. G. et al., 2010, MNRAS, 409, 48
Sanders D. B. et al., 2007, ApJS, 172, 86
Scoville N. et al., 2007, ApJS, 172, 1
Simpson J. et al., 2015, ApJ, 807, 128
Smail I., Ivison R. J., Owen F. N., Blain A. W., Kneib J.-P., 2000, ApJ, 528, 612
Smail I., Ivison R. J., Blain A. W., Kneib J.-P., 2002, MNRAS, 331, 495
Smith D. J. B. et al., 2011, MNRAS, 416, 857
Smith A. J. et al., 2012, MNRAS, 419, 377
Smolčić V. et al., 2012, A&A, 548, A4
Stetson P. B., 1987, PASP, 99, 191
Swinbank A. M. et al., 2014, MNRAS, 438, 1267
Tacconi L. J. et al., 2008, ApJ, 680, 246
Vaccari M., HELP Consortium T., 2015, preprint (arXiv:1508.06444)
Viero M. P. et al., 2014, ApJS, 210, 22
Wang L. et al., 2014, MNRAS, 444, 2870
Wang W.-H., Cowie L. L., Barger A. J., Williams J. P., 2011, ApJ, 726, L18
Werner M. W. et al., 2004, ApJS, 154, 1
Wiklind T. et al., 2014, ApJ, 785, 111
Younger J. D. et al., 2007, ApJ, 671, 1531
Younger J. D. et al., 2008, ApJ, 688, 59
Younger J. D. et al., 2010, MNRAS, 407, 1268
Zamojski M., Yan L., Dasyra K., Sajina A., Surace J., Heckman T., Helou G., 2011, ApJ, 730, 125
Zavala J. A., Aretxaga I., Hughes D. H., 2014, MNRAS, 443, 2384

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