Mem. S.A.It. Vol. 86, 598

Memorie della

A 3D dust map from Pan-STARRS 1 photometry

Gregory M. Green

Harvard-Smithsonian Center for Astrophysics, 60 Garden St., Cambridge, MA 02138, USA

Abstract. We have constructed a three-dimensional map of dust in the Milky Way, tracing reddening on $\sim 7'$ scales out to a distance of several kiloparsecs. We trace reddening using stars embedded in the dust, by simultaneously inferring stellar distance, stellar type, and the reddening along the line of sight. We use 5-band *grizy* Pan-STARRS 1 photometry of 800 million stars, augmented by 3-band 2MASS JHK_s photometry when available. The full map is available at `http://argonaut.skymaps.info`. An online version of this talk is available at `http://http://greg.ory.gr/present/ewass2015`.

1. Introduction

Dust pervades interstellar space, extinguishing and reddening near-infrared, optical and ultraviolet light, and emmitting far-infrared light. Many astronomical observations require precise corrections for the effects of dust extinction, reddening and emission. Dust emission is, for example, the most significant foreground to the Cosmic Microwave Background (CMB) in the far-infrared. Supernova Ia modeling meanwhile requires careful correction for dust reddening and extinction, both as the light travels out of the host galaxy and through the Milky Way. In the plane of the Milky Way, where most of the stars in our Galaxy lie, inference about stellar populations requires accurate 3D reddening corrections and knowledge of the survey selection function, itself dependent of dust reddening (see, e.g., Bovy et al. 2015).

Two-dimensional dust maps, giving the projected line-of-sight dust reddening over the sky, have been assembled using a number of techniques. Schlegel et al. (1998, hereafter SFD) maps dust reddening across the entire sky, using far-infrared emission to trace the dust. By modeling $100\,\mu$m emission and dust temperature, a dust optical depth can be obtained, which can be calibrated against color standards behind the dust, such as galaxies of known color, to obtain a map of dust reddening across the sky. Dust reddening has also been mapped more directly by estimating the reddening of large numbers of stars across the sky. If the stars are behind all (or most) of the dust, they trace projected dust reddening. A third technique relies on stellar number counts to trace dust. If one has a model of the stellar number density and luminosity function throughout the Galaxy, one can predict the observed number density of stars across the sky, as a function of stellar magnitude. Dust extinction reduces the number of observed stars, so comparison between observed and predicted number counts, does, in principle, give information on the density of dust throughout the Galaxy.

There are limitations, however, to the use of 2D dust maps, which can be overcome by 3D dust maps incorporating distance information. Two-dimensional dust maps give projected dust reddening or extinction, but objects embedded within our Galaxy may not lie behind all of the Galactic dust. This is especially true for studies of the plane of our Galaxy, where most of the observed stars are behind only a fraction of the dust, and extinction can increase steeply with distance. When studying objects within the Milky Way, or studying the structure of the Milky Way itself, one really requires a dust map that includes distance information.

Far-infrared emission-based dust maps, in particular, can suffer from systematic effects, as they translate from observed emission to dust optical depth, to dust reddening. Uncertainties in these translations, as well as variations in the emission and scattering properties of the dust, can introduce systematic errors into such maps which are difficult to quantify. Dust maps which directly measure reddening avoid these particular systematic effects. Dust maps that rely on stellar reddening can therefore complement two-dimensional emission-based maps, by providing an independent calibrator against which systematics can be explored.

Finally, 3D dust maps are of use beyond simple corrections to NIR, optical and UV observations. Three-dimensional dust maps allow distance determinations to objects of known reddening within the Milky Way. If the line-of-sight distance-reddening relation towards the object in question is known, and the object has a measured reddening, one can place the object at the correct distance in the map to yield the measured reddening. One can also measure distances to interstellar clouds, which are the sites of star formation (Schlafly et al. 2014). Accurate distance determinations to these clouds are important to determining their physical properties, such as their physical size and mass. Three-dimensional dust maps also give a direct window into the structure of our Galaxy. Not only do they directly trace the interstellar medium, but inference of the stellar number density throughout the Galaxy requires either simultaneous inference of the dust density, or a pre-computed 3D dust map (Bovy et al. 2015).

2. What the map looks like

Before discussing how we go about mapping dust in 3D, we first provide a look at the finished product. Video flythroughs of our map, which can be accessed at the project website, `http://argonaut.skymaps.info`, give a much better sense of what our map contains than still frames. Nevertheless, in Fig. 1, we provide a volume rendering of interstellar dust. The Sun is at the center of the frame, with various famous clouds labeled.

As can be seen in Fig. 1, our map reveals a wealth of structure in the ISM. We recover well-known cloud complexes, such as Orion A & B, California, and Monoceros R2 at distinct distances from the Sun, even though each sightline is fit separately.

The voxelization of the map is also apparent in Fig. 1. As each sightline is divided up into distance bins, and as our angular resolution is much finer than our distance resolution, the voxels are pencilbeams centered on the Sun's position in the Galaxy. The distance resolution of our map is $\sim 25\%$. Close to the Sun (within the inner few hundred parsecs), where there are fewer stars per angular pixel, our ability to constrain distance degrades, increasing the distance uncertainty of the nearest clouds. This can be seen particularly in the stretched-out appearance of the Taurus and Perseus cloud complexes.

3. How is the map constructed?

In order to trace dust with stars, we have to make a basic assumption about the smoothness of the dust density in space. If dust density were to vary wildly on very small scales, then each star could be behind a completely different column of dust, and combining information from many stars would not help us build up a 3D map of dust. We need the dust density to be smooth on small spatial scales, so that stars along nearly the same line of sight trace roughly the same dust column.

We begin by dividing up the sky into angular pixels of a few arcminutes on a side. In detail, we use the HEALPix pixelization, and subdivide pixels until all the pixels have ap-

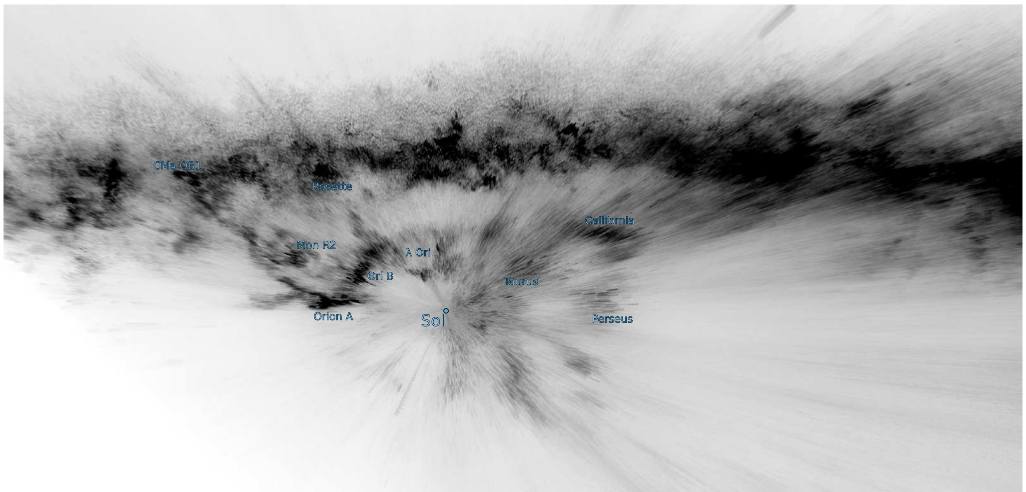

Fig. 1. A volume rendering of the 3D dust map. The camera has a 120° field-of-view, is placed 165 pc away from the Sun, and is pointed directly at the Sun, with the Galactic anticenter roughly ahead. The plane of the Milky Way is viewed nearly edge-on in this perspective. The Sun's position is labeled ("Sol"), as are various famous clouds.

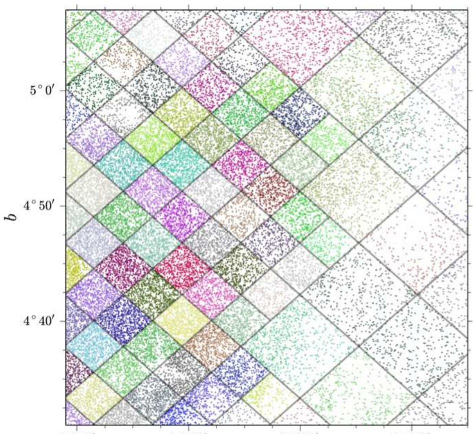

Fig. 2. An example of our pixelization on a small patch of sky. We split the sky into HEALPix pixels of varying resolution, so that each pixel contains a similar number of stars (typically a few hundred stars per pixel).

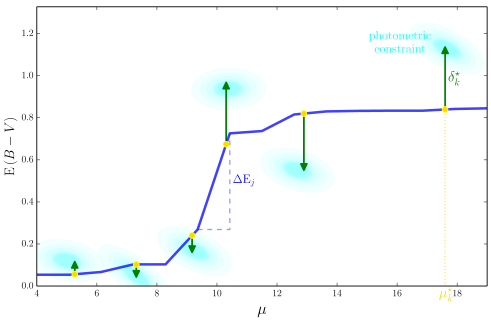

Fig. 3. Schematic picture of the variables that describe the dust and stars along one line of sight. The mean dust column across the pixel is described by the jump in reddening in each distance bin, ΔE_j. Each star k has a distance modulus, $\mu_k^\star$, and a stellar type, $\Theta_k^\star$ (not shown here). The fractional offset in the reddening from the mean reddening across the pixel at the angular position (and distance) of star k is given by $\delta_k^\star$. The distance and reddening of each star is determined by these parameters, as can be seen. In this figure, it lies at the tip of each arrow. This distance and reddening must be consistent with the constraint from the stellar photometry. Thus, the photometry from each star loosely constrains the overall dust column. Loose constraints from hundreds of stars along the line of sight combine to give a strong constraint on the distance-reddening relationship.

proximately the same number of stars. Our angular representation of the sky is therefore a multi-resolution HEALPix map. Fig. 2 gives an example of our pixelization on a small patch of sky.

We divide each angular pixel into distance bins, spaced at equal intervals in distance modulus. The mean (across the pixel) increase in reddening in distance bin j is given by ΔE_j. Each star (labeled by k) has a distance modulus, $\mu_k^\star$, which places it somewhere along the reddening column. However, we allow each star to deviate slightly from the mean reddening at its distance modulus. In effect, we are treating the reddening across the angular pixel not as a perfectly uniform column that increases with distance, but rather as a white noise process whose mean increases with distance. We call star k's fractional offset from the mean reddening (at the star's distance) $\delta_k^\star$. Given ΔE_j for each distance bin j, and $\mu_k^\star$ and $\delta_k^\star$ for each star k, we fully specify the reddening and distance for each star. In order to predict the appearance of the star, we also need the type of each star, which we label $\Theta_k^\star$. In our model, $\Theta_k^\star$ contains a luminosity and metallicity. The interplay of these variables is sketched out in Fig. 3.

Assuming that we knew all these variables, we could predict the flux of each star in each passband. By comparing the observed flux with the predicted flux, we obtain a likelihood. We also apply priors on the distribution of stars of different types throughout the Galaxy, and on the distribution of dust throughout the Galaxy. We can therefore compute a posterior probability density for any choice of ΔE_j (for all distance bins j), and $\mu_k^\star$, $\delta_k^\star$ and $\Theta_k^\star$ (for all stars k in the pixel). Markov Chain Monte Carlo (MCMC) sampling then yields possible line-of-sight dust distance-reddening curves, stellar distances and types.

As described in Green et al. (2015), we do not sample all the variables at once. It is possible to factorize the full line-of-sight problem so that one first samples from a simplified model for each star individually, and then uses the results for all the stars to sample the line of sight dust profile, described by ΔE_j. Readers interested in the details of how we factorize the problem are encouraged to read §3 of Green et al. (2014).

The end result of this process is a probabilistic map of reddening vs. distance, as well as stellar distances, reddenings and types. For each pixel, we draw hundreds of samples of the distance–reddening curve from the posterior density on ΔE_j.

4. The data

The method we describe here would not be possible without copious amounts of stellar photometry. Luckily, in the present era, digital surveys such as the Sloan Digital Sky Survey, 2MASS and Pan-STARRS 1 have mapped large areas of the sky in multiple passbands, producing large, very well calibrated catalogs of stellar magnitudes.

We make use of multi-band Pan-STARRS 1 photometry of 800 million stars, 200 million of which were also observed by 2MASS. Pan-STARRS 1 is a 1.8 m optical and NIR telescope located on Mt. Haleakala, Hawaii (Hodapp et al. 2004; Kaiser et al. 2010), which between May 2010 and April 2014 devoted the majority of its observing time to a multi-epoch survey of the sky north of $\delta = -30°$, called the 3π survey. The 3π survey was conducted in five passbands (Tonry et al. 2012; Stubbs et al. 2010; Magnier 2007), $g_{\rm P1}$, $r_{\rm P1}$, $i_{\rm P1}$, $z_{\rm P1}$ and $y_{\rm P1}$, similar to the SDSS passbands (York et al. 2000), but without a u-band filter and with the addition of the NIR $y_{\rm P1}$ filter. Pan-STARRS 1 observed to a typical depth ranging from 22 mag (AB) in $g_{\rm P1}$ to 20 mag in $y_{\rm P1}$. 2MASS is a uniform all-sky survey in three NIR passbands, J, H, and K_s, conducted by two 1.3 m telescopes – one at Mt. Hopkins, Arizona, and one at Cerro Tololo, Chile (Skrutskie et al. 2006). 2MASS achieves a typical point-source depth from 15.8 mag (Vega) in J to 14.3 mag in K_s. Together, these two surveys provide homogeneous optical and NIR coverage of three-quarters of the sky.

Although Pan-STARRS 1 observed more deeply than 2MASS, the latter complements the former well, extending our photometry farther into the infrared, where dust reddening has less of an effect. In regions of high reddening, the effective depth of 2MASS is closer to that of Pan-STARRS 1 than in relatively unobscured regions, and a substantial fraction of our sources are observed in 2MASS.

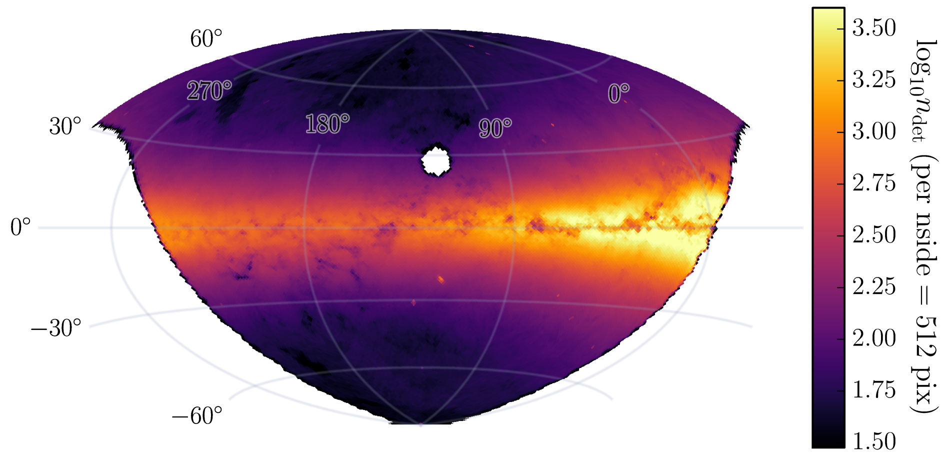

Fig. 4. Sky surface density (on a logarithmic scale) of point sources used to generate our 3D dust map. Each point source is detected in at least four passbands, at least two of which are PS1 passbands. A Hammer projection is used, with the $\ell = 0°$ meridian on the right-hand side of the image. The missing chunk of sky corresponds to $\delta < -30°$, which is not observable from Hawaii. The plane of the Milky Way is clearly visible, as are dust lanes and nearby cloud complexes that dramatically lower the observed source density in certain regions of the Galaxy.

In order to use a source, we require detection in at least four passbands, at least two of which must be PS1 passbands. In addition, we require that the source pass a cut on aperture magnitude minus point-spread-function magnitude, meant to weed out extended sources. If the source is detected in 2MASS, we require that it not be marked as an extended source. The resulting point-source surface density across the sky is plotted in Fig. 4. The typical source is detected in four or five passbands, with 20% of sources having detections in six or more passbands.

5. Comparison with 2D maps

One way to test the accuracy of our 3D reddening map is to project out distance, so that we are left with a 2D map, and then to compare that map with existing 2D maps. Because our map is based on the modeling of stellar colors, it has very different types of systematics than 2D dust maps based on the modeling of far-infrared dust emission.

We first compare our map (Green et al. 2015), integrated to 8 kpc, with the Planck 353 GHz optical depth-based dust reddening map (Planck Collaboration et al. 2014), which we will refer to as τ_{353}. Fig. 5 shows our map integrated to 8 kpc map, while Fig. 5 shows the difference between the our map and τ_{353}. Both images are stretched to a maximum reddening of 0.25 mag in $E(B-V)$, in order to accentuate the residuals.

At high Galactic latitudes, we agree well with the Planck τ_{353}-based dust map. Deep in the plane of the Galaxy, and through the cores of dense clouds, Planck τ_{353} (and other FIR emission-based dust maps) see more dust. This is because at high enough reddening, we do not observe stars behind the clouds. Our method essentially brackets the dust by measuring the reddening of stars in both the foreground and background of clouds. At distances and extinctions where we cannot observe the main sequence (for which our method fairs best), we can no longer trace reddening.

We do see systematic differences in reddening between our map on the one hand, and

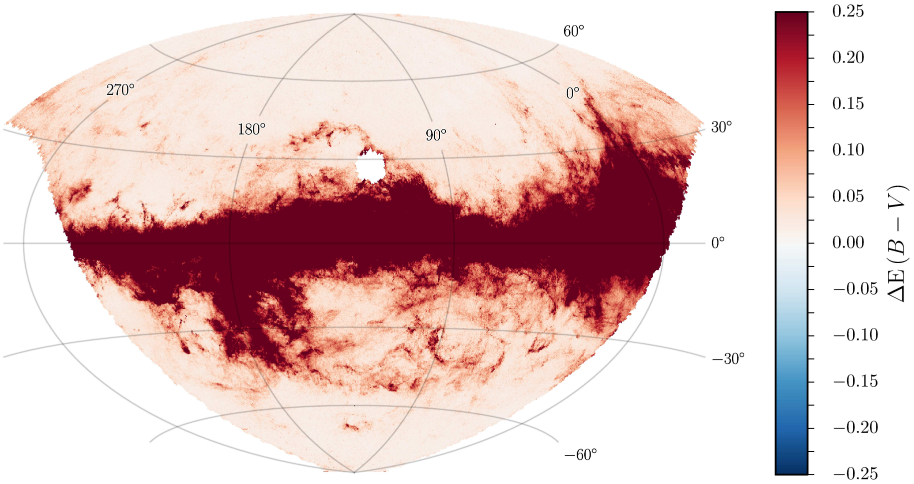

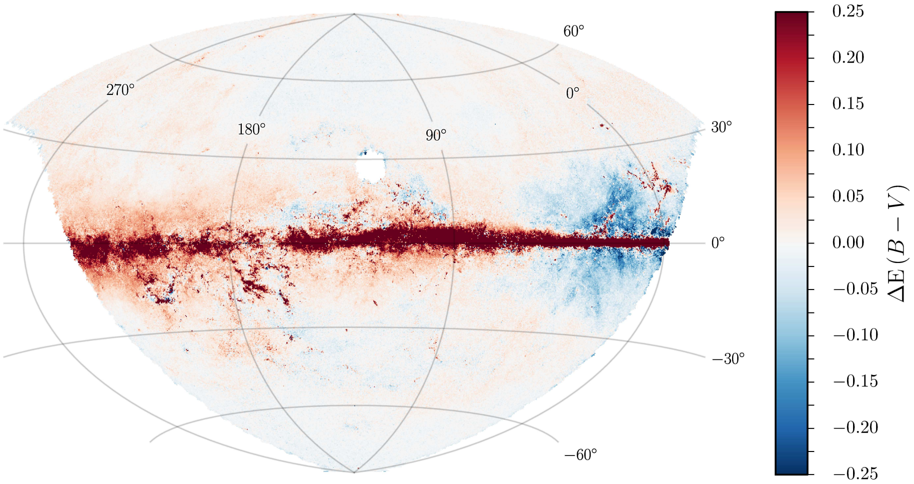

Fig. 5. Up: The Green et al. (2015) dust map, integrated to 8 kpc, Down: The Planck τ_{353} map minus the Green et al. (2015) map at 8 kpc.

both τ_{353} and SFD on the other hand, in the approximate region $0° < \ell < 45°$, $|b| < 20°$ (but excluding the immediate Galactic plane). In this region, our map infers more reddening than the FIR emission-based maps. As the excess is correlated with the reddening, one possible explanation for the discrepancy is that the reddening vector in nearby clouds in this direction (e.g., the Aquila Rift) deviates slightly from the reddening vector we use in Green et al. (2015). Inaccuracies in the reddening vector should manifest themselves as residuals

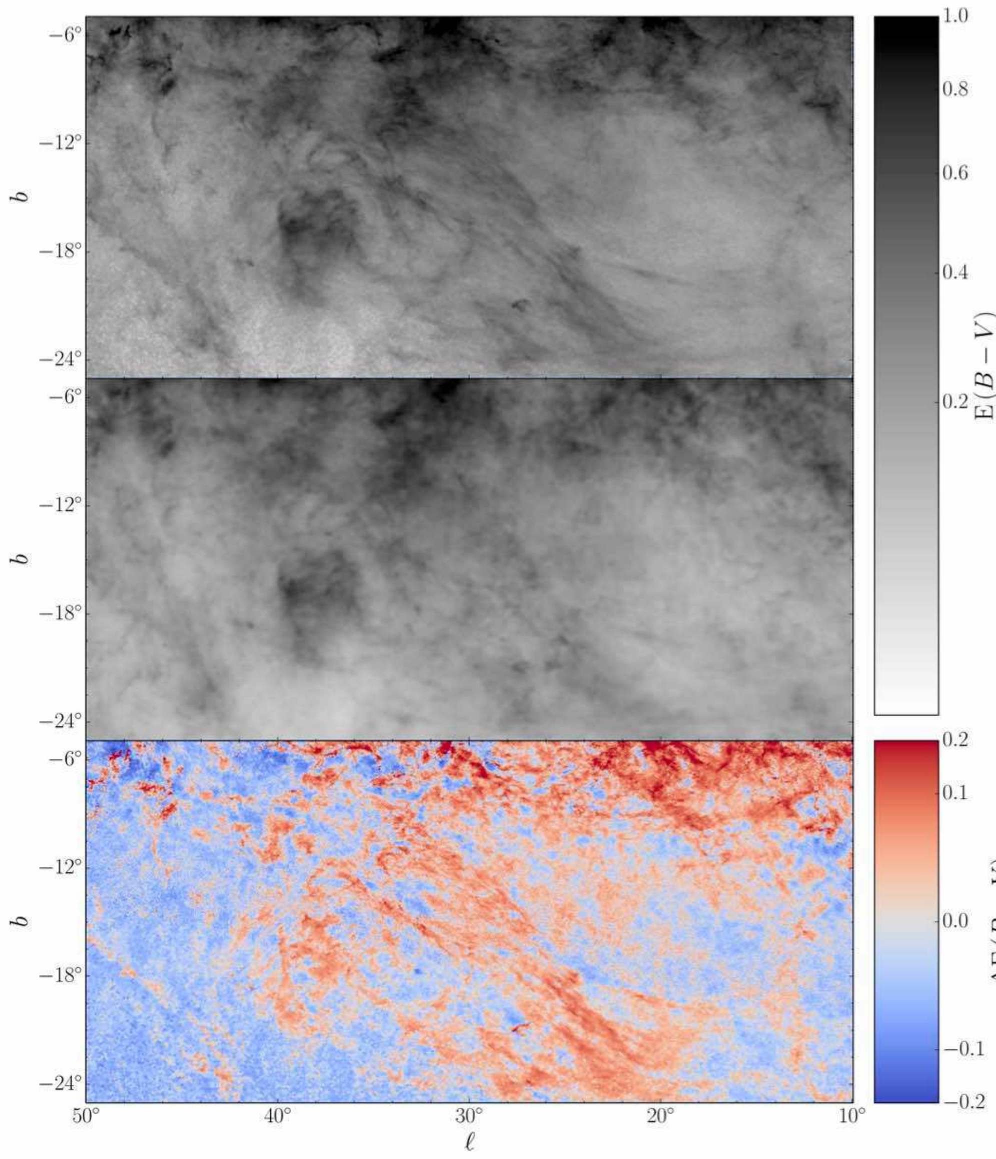

Fig. 6. Detailed comparison between our 3D dust map, projected to 8 kpc, and SFD, in the vicinity of the Aquila South cloud complex. The top panel shows the Green et al. (2015) map integrated to 8 kpc, and the middle panel shows SFD on the same scale. The bottom panel shows the difference (the 3D map minus SFD), on a stretched scale. All three panels use a squareroot color scale. Although Green et al. (2015) and SFD have similar resolution in this part of the sky, Green et al. (2015) shows fine filamentary structures that are absent in SFD.

in the reddening map that are correlated with reddening, just as the observed residuals are.

Next, we compare our map, again integrated to 8 kpc, with the SFD dust map in the vicinity of the Aquila South cloud complex. The results are plotted in Fig. 6. Even though the nominal resolution of the two maps is approximately the same in this region of the sky, we observe much more defined filamentary structure than SFD. This is likely due to the temperature correction in SFD, which is based on the DIRBE $100\,\mu$m and $240\,\mu$m maps, and which has $\sim 1°$ resolution (Schlegel et al. 1998). Due to the lower resolution of the temperature map, filaments that are colder than their surroundings would have an overestimated temperature. To compensate, SFD would infer a shallower optical depth, and thus lower inferred reddening.

The effect of this would be to smooth out fine filamentary structure in SFD. The Planck Collaboration et al. (2014) reddening maps, which infer temperature and optical depth simultaneously from the same emission maps, and therefore do not suffer from the temperature-map resolution effect seen in SFD, do indeed see the same fine filamentary structure as our 3D dust map.

Additional tests, particularly of the distance scale of our map, are possible. See, for example, Schlafly et al. (2014), which uses a closely related method to Green et al. (2014, 2015) to determine distances to a large number of dust clouds. Schlafly et al. (2014) finds good agreement with literature distances, as well as determining distances to many clouds with no reliable previous distance measurement.

6. Accessing the map

The Green et al. (2015) dust map is available at `http://argonaut.skymaps.info`. There, the map can be downloaded in full, or queried through an interactive web form. We also provide an API, whereby users can access small portions of the map remotely as a function call from within their own code.

References

Bovy, J., et al. 2015, ArXiv e-prints, arXiv:1509.06751
Green, G. M., Schlafly, E. F., Finkbeiner, D. P., et al. 2014, ApJ, 783, 114
Green, G. M., Schlafly, E. F., Finkbeiner, D. P., et al. 2015, ApJ, 810, 25
Hodapp, K. W., Kaiser, N., Aussel, H., et al. 2004, Astron. Nachr., 325, 636
Kaiser, N., Burgett, W., Chambers, K., et al. 2010, Proc. SPIE, 7733, 77330E
Magnier, E. 2007, in The Future od Photometric, Spectrophotometric and Polarimetric Standardization, ed. C. Sterken (ASP, San Francisco), ASP Conf. Ser., 364, 153
Planck Collaboration, Abergel, A., Ade, P. A. R., et al. 2014, A&A, 571, A11
Schlafly, E. F., Green, G., Finkbeiner, D. P., et al. 2014, ApJ, 786, 29
Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, ApJ, 500, 525
Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, AJ, 131, 1163
Stubbs, C. W., Doherty, P., Cramer, C., et al. 2010, ApJS, 191, 376
Tonry, J. L., Stubbs, C. W., Lykke, K. R., et al. 2012, ApJ, 750, 99
York, D. G., Adelman, J., Anderson, Jr., J. E., et al. 2000, ApJ, 120, 1579