represents anywhere between 10^3 to 10^5 stars) into physical stars. To generate the phase-space density, we used the code *EnBid* (Sharma & Steinmetz 2011), which is unique in that it allows the density to be approximated in an arbitrary number of dimensions.

During the second visit, an entirely different approach was considered. Instead of parametrizing the simulation in position-velocity space, we considered the possibility of extracting the distribution function in integral (action) space. This approach is promising because it reduces the number of dimensions needed to parametrize the distribution function to just three (instead of 6). Furthermore, recent work by e.g. Bovy et al. (2013) and Sanders & Binney (2013) has shown that fitting such action-space DFs to Milky Way data is a promising avenue for obtaining physically-meaningful analytic descriptions of the Galactic mass distribution. If we use the same techniques on the simulations, comparisons between data and observation will be greatly simplified.

3) Description of the main results obtained

This project was technically very demanding given the short amount of time we had at our disposal thus far. Although quite a lot of progress was made, these are largely only qualitative at the moment, but at the same time they provide us with a robust direction to pursue in the future.

The main issue with approach #1 described above was that obtaining the local phase-space density at the position of each particle very quickly erodes the resolution of the model. Our initial idea was to sample the distribution function in position, velocity, metallicity space and sample this DF to get physical stars and color-magnitude diagrams. However, when using anything more than 3 dimensions (i.e. more than just position), it is impossible to retain the model locality because the smoothing lengths simply become too large. That is, instead of each particle sampling a volume of a few parsecs, they start to sample a region of many hundreds of parsecs or even a kiloparsec. The following figure illustrates this problem:

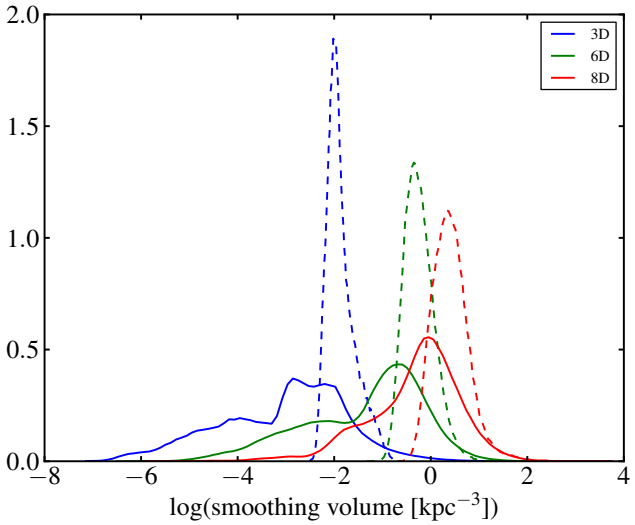

When translating this into an actual mock survey, the problem becomes even more apparent.
In this example, I used *Galaxia* to create a mock survey looking from the solar position direectly down through the plane of the Galaxy. The figure below shows the “parent” particles which are used to create the actual catalog. In the leftmost column, the locality is satisfactory, clearly

showing the “cone” of the survey in the particle selection. However, in the remaining two columns, using 6D and 8D information respectively, we see that because of the large smoothing lenghts, particles from vastly different regions of the model are scattering into the survey cone.

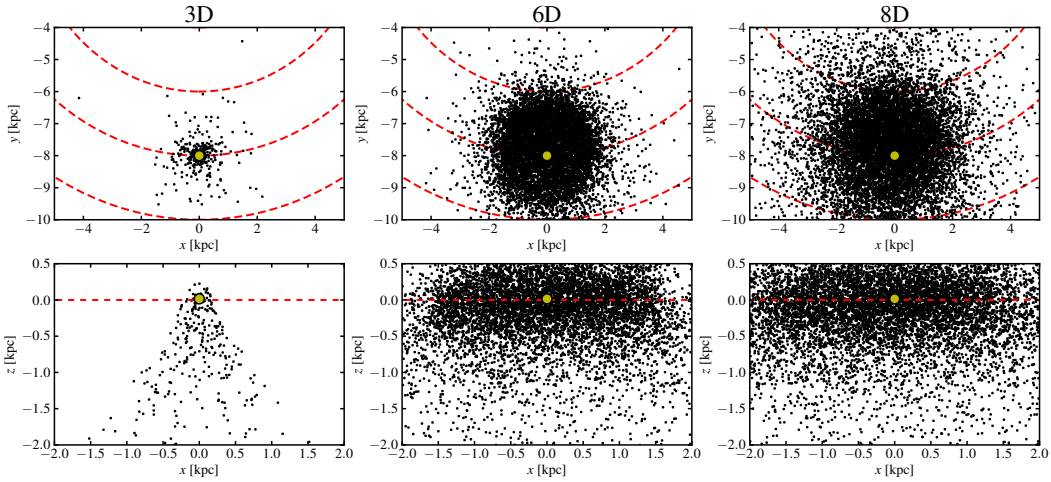

The simulation used for this example has 2 million star particles. It is reasonable to assume that in the next several years, simulations will employ 10 or 20 times this number of particles, but it is clear that this would not be sufficient to yield significantly better results. Thus, using this first approach, we are limited to only 3-dimensions, which eliminates the use of kinematics or chemistry. Given the wealth of data in upcoming surveys, this is clearly unacceptable.

For this reason, we decided in the second part of the project to try an entirely new approach: deriving actions for particles in the model and fitting them with quasi-isothermal distribution functions (Binney 2012). The figures below show the action distributions for the solar neighborhood in the model -- in the second panel of Jz vs. Jr, the evenly spaced log-space contours mean that the DF is indeed close to isothermal, which allows us to use the formalisms developed by Binney and collaborators for observational data sets and apply them to the simulations.

To generate the action distributions, I adapted the code *Galpy* developed by Jo Bovy (https://github.com/jobovy/galpy) which has been used in recent papers on MW structure (Bovy et al. 2012, 2013). The additions I made to handle arbitrary simulations will be publicly available in conjunction with the open-source analysis code *Pynbody* of which I am a co-author (https://github.com/pynbody/pynbody).

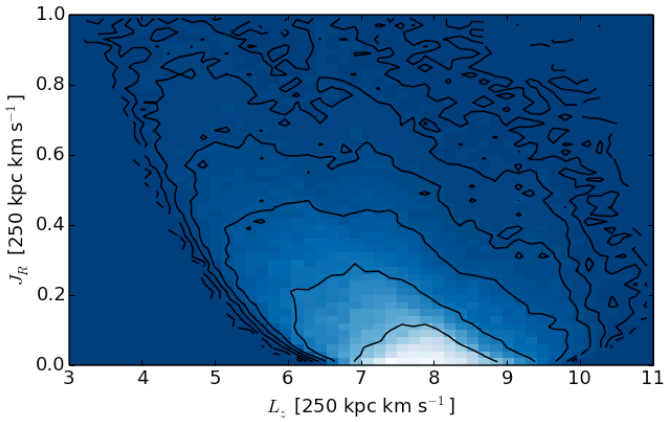

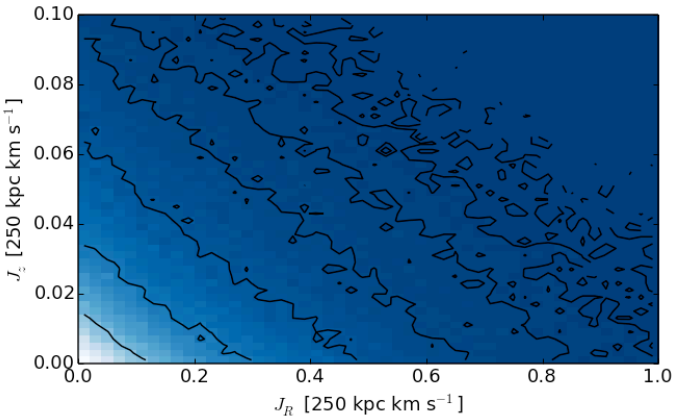

4) Future collaboration with host institution (if applicable)

The current plan is for me to finalise the action code and obtain DF fits. Once the DFs are in place, it should be possible to adapt *Galaxia* to spawn mock catalogs using this information. Then we plan on using these mocks to compare to Gaia-ESO survey data with the team at the Observatoire de la Cote d'Azur.

5) Projected publications / articles resulting or to result from the grant *(ESF must be acknowledged in publications resulting from the grantee's work in relation with the grant)*

While no publications have yet been started as a result of this work, I presented the early results at the Gaia-ESO early science meeting in April 2013 in Nice. Once the method is operational, it will certainly be publishable as it is, to our knowledge, the first attempt at such a simulation - observation comparison to-date.