1. Getting a pulsar timing solution

Aims and objectives

- Obtaining data from the Parkes Pulsar Data Archive
- Learning the basic usage of the psrchive software package
- Getting a pulsar ephemeris from the pulsar catalogue
- Forming an analytic standard template
- Using tempo2
- Improving a pulsar timing model
- Combining data sets from different observatories

Background reading

- TEMPO2, a new pulsar-timing package - I. An overview, Hobbs, Edwards and Manchester (2006; MNRAS)
- TEMPO2, a new pulsar timing package - II. The timing model and precision estimates, Edwards, Hobbs and Manchester (2006; MNRAS)
- PSRCHIVE and PSRFITS: An Open Approach to Radio Pulsar Data Storage and Analysis, Hotan, Van Straten and Manchester (2004; PASA)
- The Parkes Observatory Pulsar Data Archive, Hobbs et al. (2011; PASA)
- Tutorial on the Pulsar Data Archive (J. Khoo 2012, arXiv)
- Tutorial on tempo2 (G. Hobbs, arXiv:1205.6273)
- Tutorial on PSRCHIVE (W. van Straten et al. 2012, arXiv:1205.6276)

Who to find if you get stuck during the student week

David Champion, George Hobbs, Mike Keith, Duncan Lorimer, Dick Manchester, Maura McLaughlin, Ryan Shannon, Ingrid Stairs

Experts to discuss these topics with during the science week

The Parkes Pulsar Data Archive and the TEMPO2 software package have been developed by George Hobbs. For more information on PSRCHIVE talk to Paul Demorest and Stefan Oslowski. For information on the various IPTA data sets speak with Dick Manchester and he'll point you at the correct person for your query. There are many experts on pulsar timing at the meeting. These include David Nice, Ingrid Stairs, Joris Verbiest, You Xiaopeng and Gemma Janssen.

1 Getting data from the archive

A large number of pulsar observations from the Parkes radio telescope can be obtained from the Parkes Pulsar Data Archive. The Parkes Pulsar Data Archive is part of the CSIRO data access portal that enables users to access data related to numerous research topics. To access the archive navigate to:

`data.csiro.au`

To access pulsar data you must:

- Click on 'specific search'
- Find the "ATNF PULSAR OBSERVATION SEARCH" box and click on "START".

Now you should search for observations of PSR J1713+0747 obtained using the 10 cm receiver and the PDFB4 backend instrument between MJDs of 54750 and 55320. The archive provides raw and pre-processed files (have a look at the background material to understand the difference). For now, let's select all 29 of the pre-processed files and download them to disk. The resulting tar file should be ∼1.4 MB in size.

What is the difference between the raw files and the pre-processed files?

____
____
____

(If the internet has failed these files can be obtained in data/module1/1713_data_archive/)

2 Measuring pulse arrival times

For PTA applications the most fundamental measurement is the arrival time of a pulse at one of the IPTA observatories. Hereafter we label such arrival times as "site arrival times" (or times-of-arrival). In order to measure the site arrival times, a "standard template" of the pulse profile is initially formed. This standard template is subsequently cross-correlated with each observation to measure the arrival time and its corresponding uncertainty.

How do we measure the arrival time of a single pulse? What is the typical observation length for an IPTA observation?

____
____
____

We will make use of the psrchive software package to view the observations and measure the site arrival times. You can view the profiles using

```
$ pav -CD *.FTp
```

(experts may wish to try `psrplot`). Note: throughout these worksheets we are purposely not providing much detail on how the software works or what key-stokes to press. Often you can use a `-h` option to obtain helpful information on how to use the software package (note that you have to focus on an appropriate window, by clicking on it to get the help to work).

Do all the observations look the same? If no, how do the observations differ?

____
____
____

What are the other display options for pav? Can they be used with these files?

In order to form pulse arrival times you need to create an analytic standard template. Choose a high signal-to-noise observation (e.g., t090211_215434.FTp) and run

```
$ paas -R 0.5 -i t090211_215434.FTp
```

Follow the onscreen instructions to make a template (note that you need to add components and then press ‘f’ to fit the component to the data, then repeat for another component). Press ‘q’ to quit. You will now have a file “paas.std” that is your standard template. You can view this template using `pav` or `psrplot`.

How many components did you need to use? How many are typically used to model this pulsar (see, M. Kramer et al. 1998)?

Does this template make use of all the available information about the pulse shape?

Bonus question: These data were taken using a backend that uses *incoherent* dedispersion. How would the profile differ with observations using coherent dedispersion ? (hint: compare the DM smearing with the pulse width). Can you find an observation on the data archive to justify your conclusion? (Hint: The APSR machine uses coherent dedispersion)

Now you need to form pulse arrival times:

```
$ pat -f tempo2 -s paas.std *.FTp
```

This output goes to the screen. You need to write this output to a file (e.g., arrivalTimes.tim). This file contains 5 columns. The first column is the filename, the second column is the observing frequency (MHz), the third column is the site-arrival-time (as a Modified Julian Date), the fourth column is the site-arrival-time uncertainty (in microseconds) and the fifth column gives the telescope identification code.

What is the typical uncertainty on the arrival time measurements? Why is there a variation?

Instead of using an analytic template, you can form a template using the average of the observations or simply the brightest observation. Please form a second set of TOAs using the brightest observation as the template.

Can we use the brightest observation as the template?

3 Getting the timing residuals

In order to obtain pulsar timing residuals you need to obtain an ephemeris for the pulsar. Ephemerides for all published pulsars can be obtained from the pulsar catalogue:

`http://www.atnf.csiro.au/research/pulsar/psrcat`

Find out how to obtain an ephemeris for PSR J1713+0747 and save this to disk (e.g. in a file called ephemeris.par).

(If there is no internet connection then you can find a suitable ephemeris in /data/module1/psrcat/)

Give a description for all the parameters in the ephemeris. What do they correspond to and where did they come from?

Now you can plot the timing residuals:

```
$ tempo2 -gr plk -f ephemeris.par arrivalTimes.tim
```

The resulting plot has multiple sections. There are a set of boxes and parameter labels across the top. Buttons such a RE-FIT near the top, a panel to the left, a graph and a set of white buttons near the bottom. Experiment or read the tutorial to determine what is the basic purpose of these different sections?

Note that some text has been written out on the terminal. Give a brief description of what this text contains

It will be necessary to improve the timing ephemeris. You may find it useful to plot the timing residuals versus orbital phase - using the menu bar on the left of the plot.

Are there any "bad" points? If so, you can first view them by clicking with the middle mouse button and then they can be deleted by clicking on them with the right mouse button.

Why do bad points exist? What phenomena could lead to such outliers?

Why is the ephemeris that you obtained from the pulsar catalogue not perfect? Where did it come from?

Tempo2 carries out a least-squares-fit to minimise the timing residuals by changing the parameter values. By default all the parameters are held fixed, but you can choose to let various parameters be included in the fit by selecting them from the top bar (select e.g., PB, T0, A1 etc.) and then by clicking on re-fit. To see the post-fit residuals ensure that the "post-fit" box has been selected for the y-scale using the left-hand menu bar. Initially fit for the pulse frequency, its first time derivative and the five Keplerian binary parameters.

What is the post-fit, non-weighted rms timing residual that you are getting?

By default the tempo2 fit is unweighted. If you press 'w' and then click on 're-fit' you will carry out a weighted fit. Does the rms value significantly change? If so, why?

Do some observations have much larger error bars than other observations? If yes, click on an observation with a large error bar with the central mouse button. This should display the pulse profile. Now repeat by clicking on one with a small error bar.

Why do the error bars vary in size?

You should now have a new timing ephemeris which can be written to disk by clicking on the "New par" button. If this is the first time that you have made a timing solution then please check with one of the experts that this has been completed correctly.

What is your measurement of PSR J1713+0747's orbital period? How does it compare to the value published in the pulsar catalogue? How does it compare to the value published in Verbiest et al. (2009)?

4 More advanced analysis

This section demonstrates a few more features of TEMPO2 and describes issues that are important for the IPTA.

4.1 Dealing with instrumental time offsets

For this part of the project we need to access more data. Download pre-processed 10 cm, J1713+0747 observations with the PDFB4 backend using the Parkes Data Archive for observations after MJD 55320 and before 55510. You should download 11 files.

(If the internet is not working then you can find a suitable ephemeris in (If the internet has failed these files can be obtained in /data/module1/1713_data_archive2/))

Form site arrival times using the same standard template as before and append these to the original set of arrival times. Now re-run tempo2 and re-do the fitting as before. Do you notice a jump of $\sim 8\mu$s around MJD 55328? This occurred when the backend instrumentation was upgraded.

Let's assume that the size of the offset is unknown and we need to measure it from the data. Open the arrival time file in your favourite editor (emacs or vim). For all the observations before the upgrade we need to add a "flag" to the end of each arrival time line, e.g.,

```
t081025_043858.FTp 3100.500 54764.216087999808628 0.105 7 -section 1
...
```

and for all observations after MJD 55327 we should change the flag value:

```
t101109_043534.FTp 3100.500 55328.769560199533500 0.105 7 -section 2
...
```

When running tempo2 we can highlight the two regions using:

```
$ tempo2 -gr plk -f ephemeris.par arrivalTimes.tim -colour -section
```

To fit for an arbitrary time "jump" between the early and late sections we need to include the jump parameter in the pulsar timing model. Edit the ephemeris.par file to add the line:

```
JUMP -section 2 0 0
```

(The first "0" gives the initial guess for this time offset and the second "0" tells tempo2 not, for now, to include this in the fit.)

Why can't we include a fit for -section 1 and also for -section 2?

____

____

____

Re-run tempo2 and fit as normal for the pulse and binary parameters. When you have the best fit possible then click on the red '2' in the bottom-left corner. It should turn green indicating that the time offset will now be included in the fit. Click on 're-fit' and check that the post-fit residuals do not show any evidence of the time offset.

What was the time offset caused by the instrumental upgrade? What length of cable would need to be added to cause the change?

____

____

____

Save your final ephemeris by clicking on "New par" and saving with a file name such as `ppta_final.par`.

4.2 Combining Parkes and NANOGrav data

The IPTA will allow team members to access data from telescopes situated in Australia, Europe and North America. We have made site arrival times for NANOGrav data available for PSR J1713+0747 during the same time range as the Parkes data that you have processed above. These data can be obtained from:

```
/data/module1/nanograv/nano.1713.par
/data/module1/nanograv/nano.1713.tim
```

We have only selected a subset of all the available NANOGrav data.

Which observing frequency was selected for the NANOGrav data?

____

____

____

The goal here is to combine the two data sets, form post-fit timing residuals, highlight each observatory individually and to search for any unexplained offsets between the two data sets.

To do this you should

- Combine your arrival time file with the NANOGrav data.
- Add a flag on the end of each arrival time line to indicate whether the arrival time was from NANOGrav or PPTA.
- Add a jump parameter in your ephemeris to fit for the arbitrary time offset between the two observatories.
- Study your pre- and post-fit timing residuals (use the -colour option to highlight the different observatories.)

Why is there an offset between the two data sets and what is the size of the offset?

What is the typical sampling rate for the Parkes data compared with the NANOGrav data?

Why does the NANOGrav data have larger error bars than the Parkes data?

What is the mean and median arrival time uncertainty for both telescopes? Why are they different?

Congratulations you have successfully combined IPTA data from the PPTA and from NANOGrav. Show somebody your final post-fit timing residuals!

5 To think about

Do you know the answers to the following questions? If not, perhaps you should speak with one of the experts.

- Are the resulting timing residuals consistent with white noise and do the arrival time uncertainties represent that white noise?
- What happens if you change the parameter values in the ephemeris file?
- Can you think of any way to decrease the rms timing residuals?
- What would be the typical amplitude of a gravitational wave signal in these data?
- You started with an ephemeris from the pulsar catalogue? What would you do if you had just discovered the pulsar and no such ephemeris exists?
- Do tempo1 and psrtime give the same answers as tempo2?
- What is the effect of a poorly made standard template?

The following problems are not fully solved. Any solution would be important for the IPTA.

- Why are the timing residuals not consistent with white noise?
- What would the effect of fitting for the pulsar parameters be on a gravitational wave signal?
- Are there are physical phenomena that are not being included in the tempo2 timing model?
- Can we improve on the cross-correlation technique to obtain site-arrival-times?
- How do we know that our timing software (e.g., tempo2) is giving the correct result?

6 Publication

Find a publication on ADS that is related to this worksheet and summarise the basic conclusions of that paper.

7 Mistakes

How many mistakes did you find in this worksheet?