

Introduction

The busgrap pre-processor is intended to let you embed business type graphics within the flow of your groff document. All the graphics are produced using groff's native drawing primitives, so the results are pure vector based drawings, no bit mapped graphics. Each graph quires you to embed a call to macro ".BGS" followed by a set of parameters, described below, which are terminated by a call to ".BGE". Between these two macro calls only busgrap parameters are expected, any normal groff language will cause an error. The format of the ".BGS" parameter is:-

```
.BGS name
```

Where 'name' can be "pi" or "graph". The difference determines whether the graph will be a pie chart or a line/bar chart.

In order to place a graph in the document you need to specify a "frame" which will contain the actual graph when it is generated.

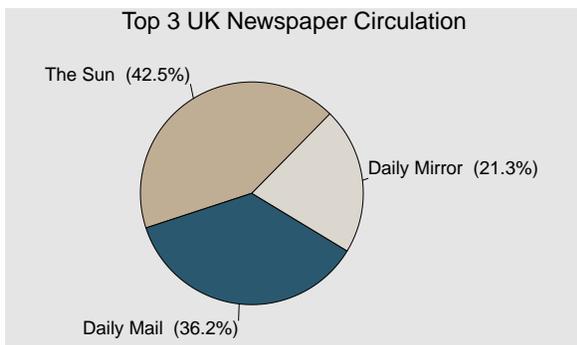
The Frame

In order to place a graph in the document you need to specify a "frame" which will contain the actual graph when it is generated. Busgrap will try to fit your graph within the given frame, but it is not very intelligent so it may not be successful in keeping within the frame, if this happens, increase the size of the frame until it is successful, or reduce the size of labels, etc..

The series labels, (the name associated with the series of data you are plotting), can either be shown on the graph or in a separate "key box" which is a separate box holding the legend of labels. If you define a key box it must be included in the area defined as the frame.

These are the parameters which can affect a frame:-

```
Frame:width ^ height           # default: 6c ^ 5c
Origin:relx ^ rely
Just:{left|centre|right}      # default: left
Flow:{yes|no}                  # default: no
Border:size                     # default: zero (no border)
BRDcolour:                      # default: black border
BGcol:colour                    # default: no frame background
```

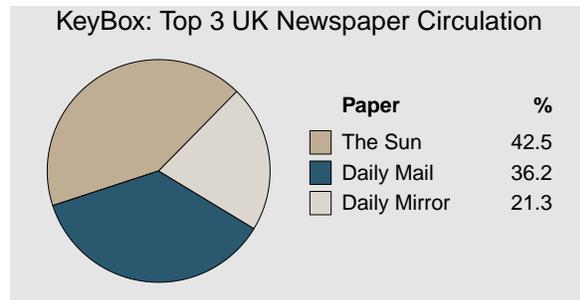


When dealing with pie charts the Origin parameter moves the centre of the pie, the x and y offsets are relative to the top left corner of the frame. The default value is to centre the pie within the frame. For line or bar charts it has a slightly different format which is described later.

The "Flow" parameter causes the text following the graph to be flowed down the side of the graph. It is only active if the "Just" parameter is left or right, you cannot flow test if the graph is centred.

This example has Flow turned on and the text is flowing down the right side of the graph. Along with these parameters there are a set which affect any key box required:-

```
KeyBox:{yes:no}                # default: no (series names shown within the graph)
BoxFrame:xoffset ^ yoffset { ^ width ^ height}
BoxHeads:Col1 ^ Col2 ^ Col3   # default: Name ^ Value ^ %
BoxLabels:template             # default: $text ^ $value ^ $percent
BoxTabs:groff tab string      # default: 0cL 2.8cR 4cR
```



We will now show the same graph but with a KeyBox instead of labelling the graph on the pie itself. We will need to move the origin to the left to leave room for the KeyBox on the right. The contents of the box will be set so that the BoxHeads and BoxLabels only show the text and percent.

Using a KeyBox with a legend for the pie can often produce a clearer graph for the reader, and you can use a slightly larger font size.

Graphing Data

The data for graphing is entered as series of data. The format is slightly different, depending on whether you wish to produce a pie chart or a line/bar chart. For pie chart each series is a wedge of pie and just consists of the name and a value, so the format is simply:-

```
Series:name ^ value
Label:template          # default: $text ($percent%)
```

The Series parameter is the only graph parameter which may appear multiple times in the graph definition, once for each wedge. The software will add up the value in each series and calculate the percentage for each pie.

The label parameter controls how the data is labelled, when you choose to have inline labelling rather than a separate KeyBox. It performs a similar purpose to the BoxLabel parameter shown above, but produces a string, rather than a tab separated table. In the template there are three inbuilt variables which can be used in the template string: \$text is replaced by the name of the series, \$value is replaced with the value passed with a particular name, and \$percent is replaced with the value as a percentage of the total of all series values added together. Also, note that the value for each Series can be a "formatted" number, i.e. may have currency symbol or thousand separators, it will be this string which is passed back as \$value, rather than a pure number.

If you are creating a line or bar chart, the following parameters control the data:-

```
Series:type ^ name ^ value1 ^ value2 ^ ...
Steps:number
Floor:{C}value
```

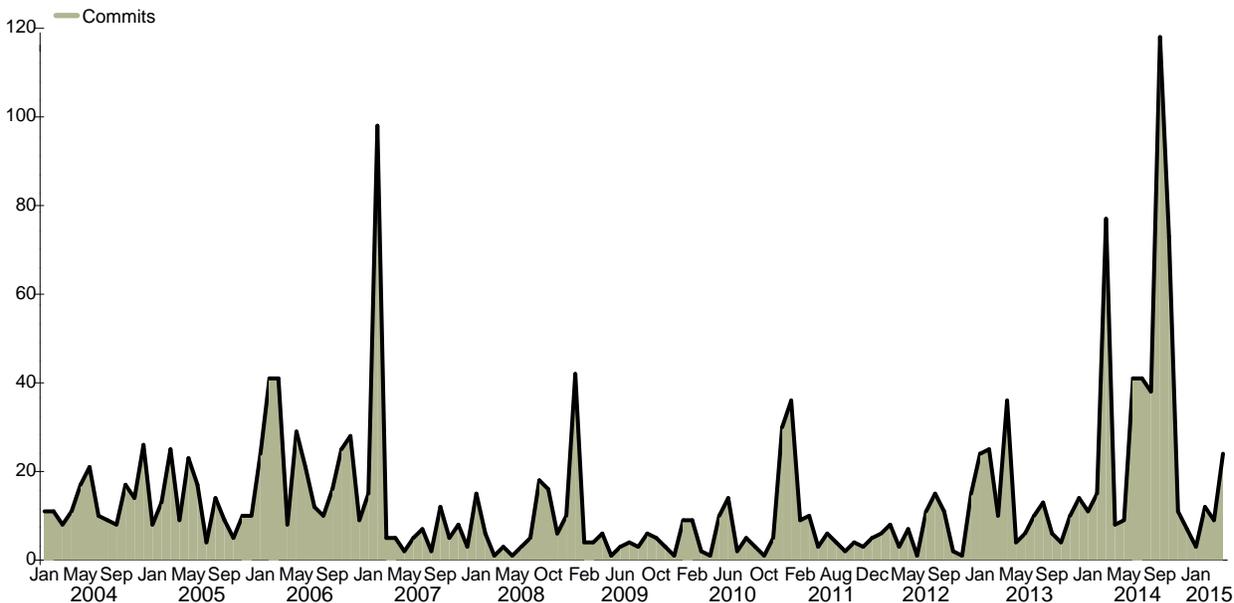
The series parameter now has an extra field called "type" and each series can contain multiple values. There are two types of series you must supply: the "X" type (which labels the X-axis of the graph) and at least 1 numerical type, i.e. the type is 1-n (each numbered series contains the values for one line in the graph). Sometimes the X axis will be time based. If the dates are given as mmm-yy{yy} or mm/yy{yy} then the value is recognised as a valid date and the X-labels will be formatted on two lines for the year and month. The X-type series should hold a label for each data point you intend to graph, so if your data is monthly and you have 12 months worth of data there should be 12 values in the X-type (and 12 in each of the numbered series parameters).

Another Series type is "A" (anchor) this should hold one value per numbered series and is the starting value for that series, so lines are drawn from the Y-axis to the first data point. If, for example, you have two a data set which holds the value of a portfolio, and one for a financial benchmark. Both data sets start in January 2014. You rebase both data sets so that they start at 100 and the following figures will show gain/loss for each following month. In this case you would remove the first data point from both series, and the X-series and add two values of 100 to the A-series parameter. The effect of this would be for the lines for both series starting at the 100 point on the Y-axis to the first data point which would be February 2014.

The final series type is "B". This can be used instead of Series 1 (you must not have a series 1 if you have a series "B". The only difference is that the data in a Series B is plotted as an area graph, i.e. the area under the line is filled with colour. Since this Base series is a replacement for Series 1 it plotted first so the lines of any subsequent series are plotted on top of the area.

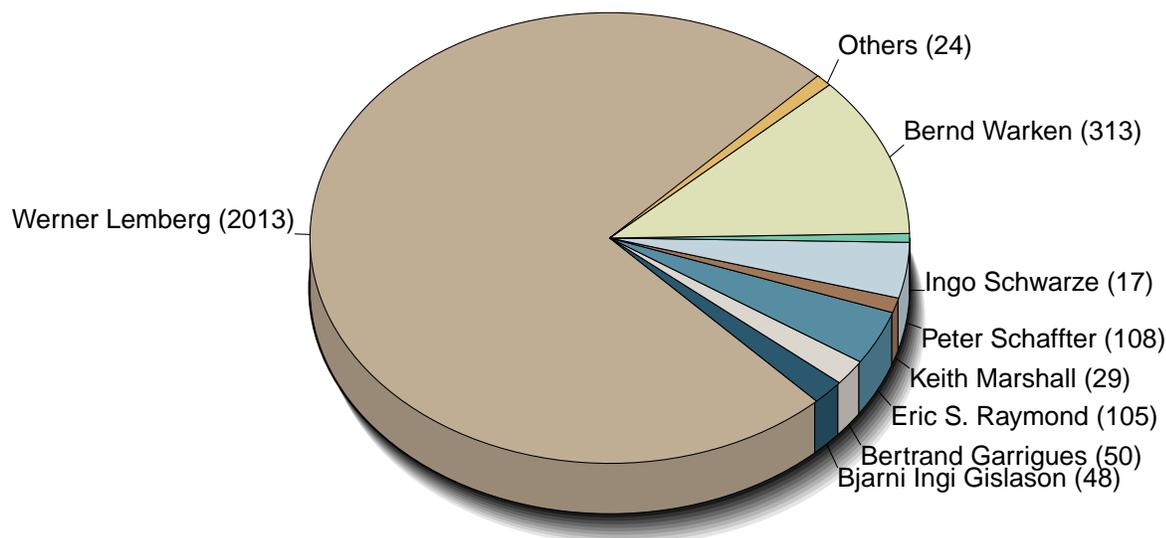
Using all the values in the numbered series, busgrap will calculate the appropriate Y-scale labels to use, it tries hard to settle on a scale which is "natural". The parameter Steps influences the algorithm, by telling it how many steps you would prefer labelled in the Y-axis, but it is quite stubborn in using what it thinks is best! The Floor of the graph is normally the bottom of the graph but the Floor parameter can be used to change this, the value given will be used as the new floor. If the floor value is prefixed by "C" the floor will be centred in the graph. This is probably only useful when doing bar chart since bars are drawn upwards or downwards from the floor, so if you are graphing both positive and negative data as a bar chart it can help clarity if you use Floor:0. The other occasion to use Floor, for all types of line graph is when the Y-axis does not start from zero but you would like it to start from zero.

In this example, which shows the number of commits to groff by month since January 2004, only one series has been graphed, but it has been specified as a Series:B, so it produced an area graph.



The data comes from the git log and the next graph comes from the same source.

Top group of committers to Groff



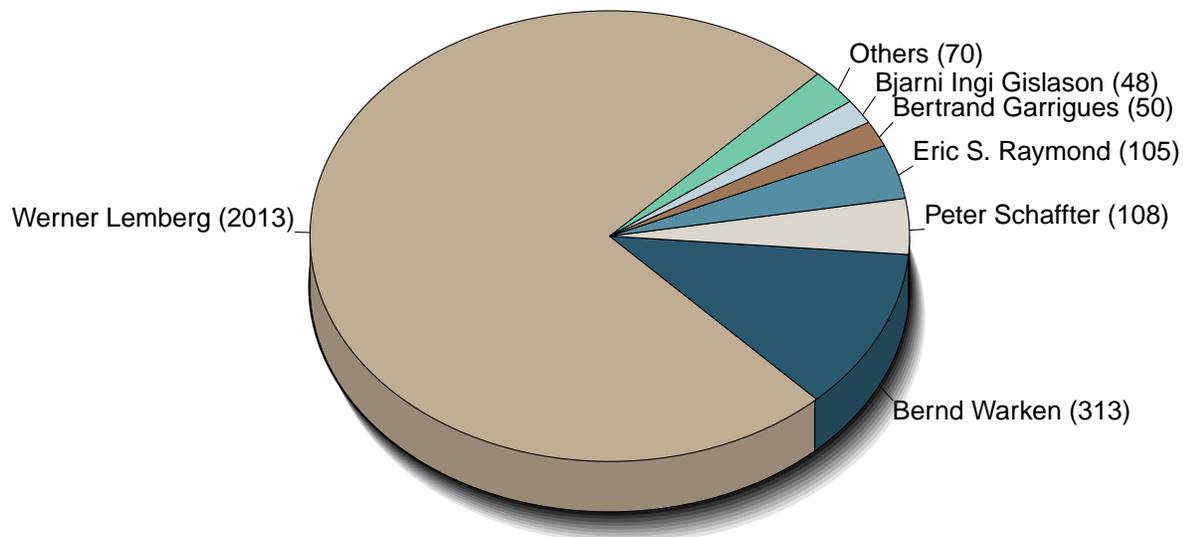
Others includes: Steffen Nurpmeso (14), James Clark (10)

In this pie chart you can see that two of the series have been automatically combined into one segment, called "Others", this is because the segments would be too small to have individual inline labels. This is controlled by parameters:-

```
Others:percentage      # default: 2.05
Sort:yes|no           # default: yes
```

The Others parameter specifies the minimum percentage, if a particular segment is below this size then it is a candidate for combination. Whether it will be combined depends on the size of its neighbours, since if both neighbours are large enough there will be room for the small segment to be labelled. This is why the Sort parameter has an effect, since when Sort is turned on the segments are sorted decreasing from largest to smallest, so the smallest segments will clump together, so more will be put into Others. With Sort turned off, and careful ordering of the series, you can arrange so the smallest are interspersed with the largest. This is the same graph with Sort turned on and Others increased slightly to avoid labels overwriting:-

(Sort) Top group of committers to Groff

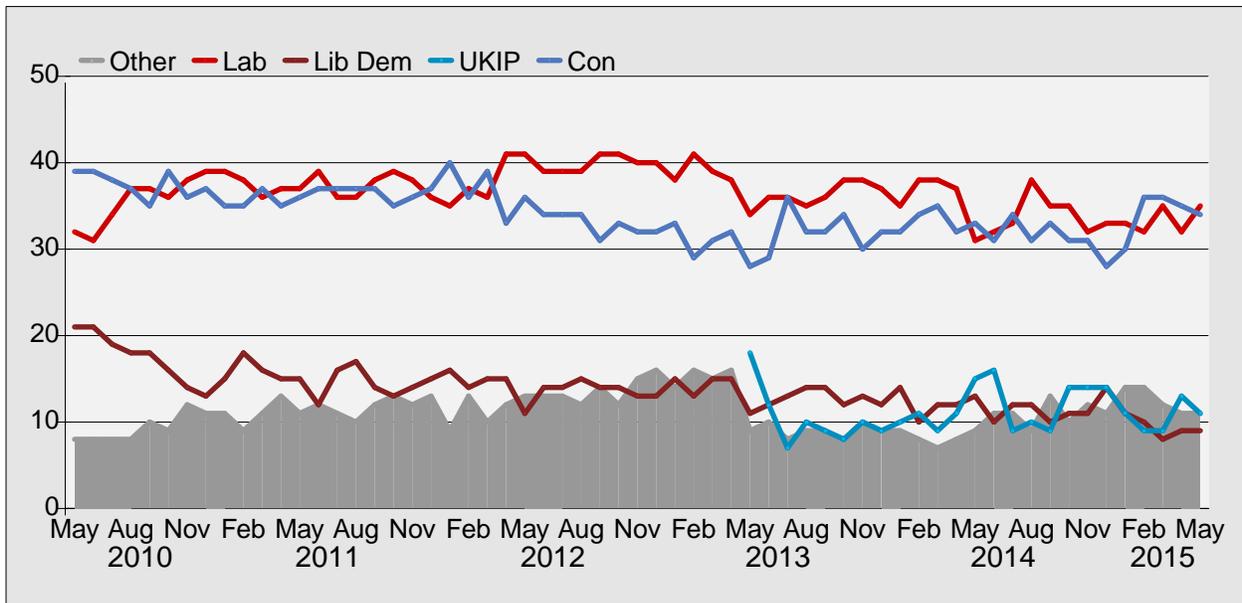


Others includes: Keith Marshall (29), Ingo Schwarze (17), Steffen Nurpmeso (14), James Clark (10)

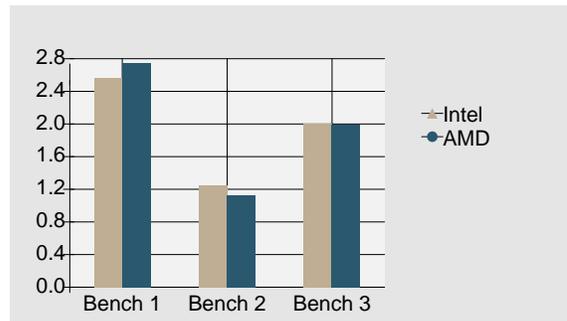
Other methods of avoiding labels overwriting each other are to adjust the point size of the labels and/or increase the size of the pie, see affecting pie style below.

In any statistical work it is useful to have a "missing value". Busgrap uses a single "." (dot) to signify a value is missing, so rather than have a blank value, which equates to zero, use the dot instead.

In the following graph which shows the polling figures for UK political parties from May 2010, the UKIP party was not established until May 2013, so the missing values have been replaced with dots in the Series parameter for the UKIP party.



When the input data includes decimal places the algorithm which decides on the Y-labels may or may not use decimals. The decision is based on how close the range of numbers is between the lowest and the highest. When the range is very close together the algorithm will choose an appropriate number of decimals. The choice can be overridden by setting parameter PDecimals to the number of places required. In this example the algorithm has chosen 1 decimal place, which is probably the correct choice in this case.

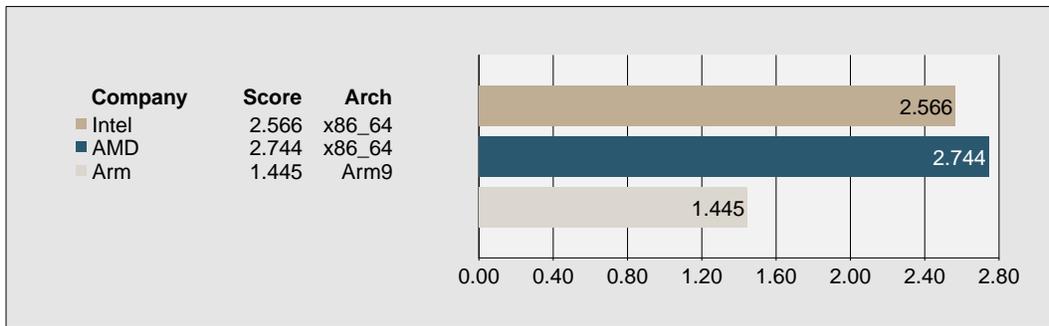


To graph a particular series as bars, instead of lines you just put "BAR:" in front of the series name, so two series in this graph were called "BAR:Intel" and "BAR:AMD". It is perfectly permissible to mix bars and lines in the same graph. Any bars will always be drawn first then any lines appear over the top.

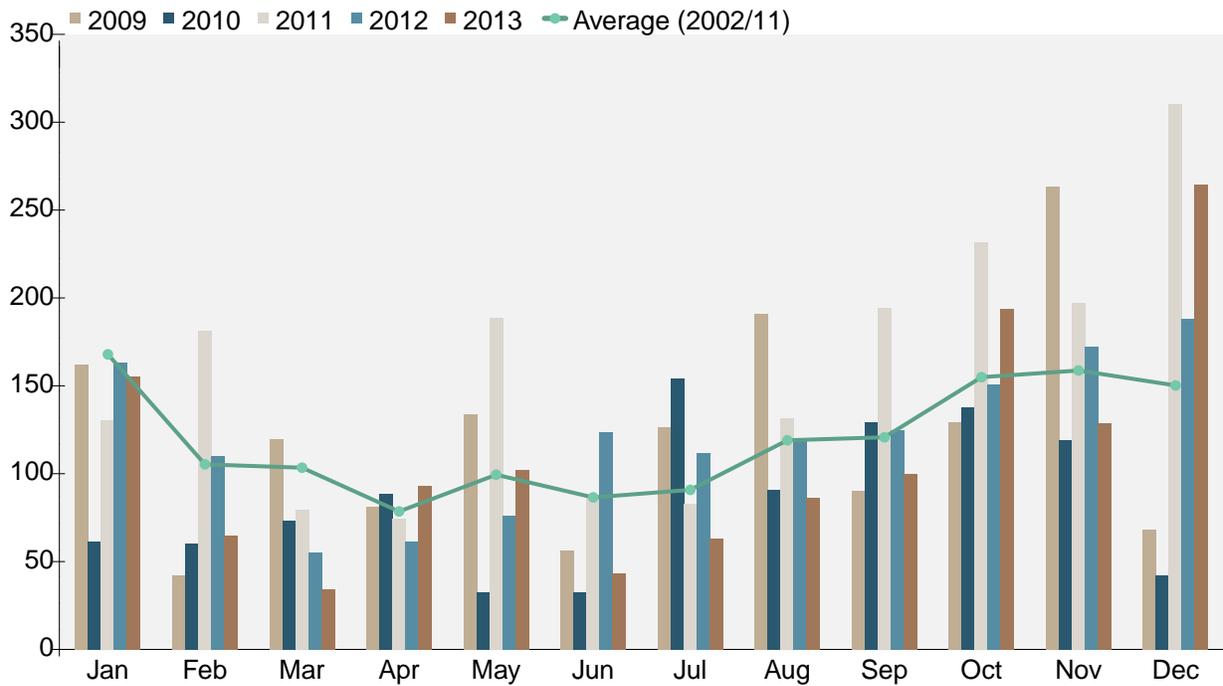
The actually data to produce this bar chart was:-

```
Series:X ^ ^ Bench 1 ^ Bench 2 ^ Bench 3
Series:1 ^ BAR:Intel ^ 2.566 ^ 1.253 ^ 2.000
Series:2 ^ BAR:AMD ^ 2.744 ^ 1.122 ^ 1.990
```

So, even though the data was to 3 decimals, only 1 place of decimals was relevant to the graph labels.



UK Rainfall in mm



Affecting the pie style