

What's in a Name?

The Importance of Naming grid Grobs When Drawing Plots in R

by Paul Murrell

Abstract Any shape that is drawn using the **grid** graphics package can have a name associated with it. If a name is provided, it is possible to access, query, and modify the shape after it has been drawn. These facilities allow for very detailed customisations of plots and also for very general transformations of plots that are drawn by packages based on **grid**.

Introduction

When a scene is drawn using the **grid** graphics package in R, a record is kept of each shape that was used to draw the scene. This record is called a *display list* and it consists of a list of R objects, one for each shape in the scene. For example, the following code draws several simple shapes: some text, a circle, and a rectangle (see Figure 1).

```
> library(grid)
> grid.text(c("text", "circle", "rect"),
+          x=1:3/4, gp=gpar(cex=c(3, 1, 1)))
> grid.circle(r=.25)
> grid.rect(x=3/4, width=.2, height=.5)
```




Figure 1: Some simple shapes drawn with **grid**.

The following code uses the `grid.ls()` function to show the contents of the display list for this scene. There is an object, called a *grob* (short for “graphical object”), for each shape that we drew. The output below shows what sort of shape each grob represents and it shows a name for each grob (within square brackets). In the example above, we did not specify any names, so **grid** made some up.

```
> grid.ls(fullNames=TRUE)
text [GRID.text.5]
circle [GRID.circle.6]
rect [GRID.rect.7]
```

It is also possible to explicitly name each shape that we draw. The following code does this by specifying the `name` argument in each function call (the

resulting scene is the same as in Figure 1) and call `grid.ls()` again to show that the grobs on the display list now have the names that we specified.

```
> grid.text(c("text", "circle", "rect"),
+          x=1:3/4, gp=gpar(cex=c(3, 1, 1)),
+          name="leftText")
> grid.circle(r=.25, name="middleCircle")
> grid.rect(x=3/4, width=.2, height=.5,
+          name="rightRect")
> grid.ls(fullNames=TRUE)
text[leftText]
circle[middleCircle]
rect[rightRect]
```

The **grid** package also provides functions that allow us to access and modify the grobs on the display list. For example, the following code modifies the circle in the middle of Figure 1 so that its background becomes grey (see Figure 2). We select the grob to modify by specifying its name as the first argument. The second argument describes a new value for the `gp` component of the circle (in this case we are modifying the `fill` graphical parameter).

```
> grid.edit("middleCircle", gp=gpar(fill="grey"))
```




Figure 2: The simple shapes from Figure 1 with the middle circle modified so that its background is grey.

The purpose of this article is to discuss why it is useful to provide explicit names for the grobs on the **grid** display list. We will see that several positive consequences arise from being able to identify and modify the grobs on the display list.

Too many arguments

This section discusses how naming the individual shapes within a plot can help to avoid the problem of having a huge number of arguments or parameters in a high-level plotting function.

The plot in Figure 3 shows a *forest plot*, a type of plot that is commonly used to display the results of a meta-analysis. This plot was produced using the `forest()` function from the **metafor** package (Viechtbauer, 2010).

This sort of plot provides a good example of how statistical plots can be composed of a very large number of simple shapes. The plot in Figure 3 consists of




Figure 3: A forest plot produced by the `forest()` function from the **metafor** package.

many different pieces of text, rectangles, lines, and polygons.

High-level functions like `forest()` are extremely useful because, from a single function call, we can produce many individual shapes and arrange them in a meaningful fashion to produce an overall plot. However, a problem often arises when we want to *customise* individual shapes within the plot.

For example, a post to the R-help mailing list in August 2011 asked for a way to change the colour of the squares in a forest plot because none of the (thirty-three) existing arguments to `forest()` allowed this sort of control. The reply from Wolfgang Viechtbauer (author of `metafor`) states the problem succinctly:

“The thing is, there are so many different elements to a forest plot (squares, lines, polygons, text, axes, axis labels, etc.), if I would add arguments to set the color of each element, things would really get out of hand ...

... what if somebody wants to have a different color for *one* of the squares and a different color for the other squares?”

The reality is that it is impossible to provide enough arguments in a high-level plotting function to allow for all possible modifications to the low-level shapes that make up the plot. Fortunately, an alternative is possible through the simple mechanism of providing names for all of the low-level shapes.

In order to demonstrate this idea, consider the **lattice** plot (Sarkar, 2008) that is produced by the following code and shown in Figure 4.

```
> library(lattice)
> xyplot(mpg ~ disp, mtcars)
```

This plot is simpler than the forest plot in Figure 3, but it still contains numerous individual shapes. Anyone familiar with the **lattice** package will also know that it can produce plots of much greater complexity; in general, the **lattice** package faces a very difficult problem if it wants to provide an argument in its high-level functions to control every single shape within any of its plots.

However, the **lattice** package also provides names for everything that it draws. The following code shows the contents of the **grid** display list after drawing the plot in Figure 4.

```
> grid.ls(fullNames=TRUE)
rect[plot_01.background]
text[plot_01.xlab]
text[plot_01.ylab]
segments[plot_01.ticks.top.panel.1.1]
```

```
segments[plot_01.ticks.left.panel.1.1]
text[plot_01.ticklabels.left.panel.1.1]
segments[plot_01.ticks.bottom.panel.1.1]
text[plot_01.ticklabels.bottom.panel.1.1]
segments[plot_01.ticks.right.panel.1.1]
points[plot_01.xyplot.points.panel.1.1]
rect[plot_01.border.panel.1.1]
```




Figure 4: A simple **lattice** scatterplot.




Figure 5: The **lattice** plot from Figure 4 with the x-axis modified using low-level **grid** functions.

Because everything is named, it is possible to access any component of the plot using the low-level **grid** functions. For example, the following code modifies the x-axis label of the plot (see Figure 5). We specify the component of the scene that we want to modify by giving its name as the first argument to `grid.edit()`. The other arguments describe the

changes that we want to make (a new label and a new `gp` setting to change the `fontface`).

```
> grid.edit("plot_01.xlab",
+           label="Displacement",
+           gp=gpar(fontface="bold.italic"))
```

That particular modification of a **lattice** plot could easily be achieved using arguments to the high-level `xyplot()` function, but the direct access to low-level shapes allows for a much wider range of modifications. For example, figure 6 shows a more complex multipanel **lattice** barchart.

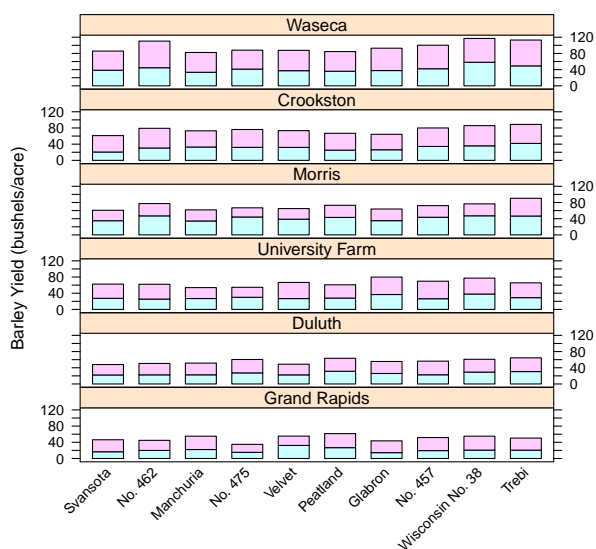


Figure 6: A complex multipanel **lattice** barchart.

This is generated by the following code

```
> barchart(yield ~ variety | site, data = barley,
+          groups = year, layout = c(1,6),
+          stack = TRUE,
+          ylab = "Barley Yield (bushels/acre)",
+          scales = list(x = list(rot = 45)))
```

There are too many individual shapes in this plot to show the full display list here, but all of the shapes have names and the following code makes use of those names to perform a more sophisticated plot modification: highlighting the sixth set of bars in each panel of the barchart (see Figure 7).

```
> grid.edit("barchart.pos.6.rect",
+          grep=TRUE, global=TRUE,
+          gp=gpar(lwd=3))
```




Figure 7: The barchart from Figure 6 with the sixth set of bars in each panel highlighted.

The first argument to `grid.edit()` this time is *not* the name of a specific grob. This time we have given a name *pattern*. This is indicated by the use of the `grep` argument; `grep=TRUE` means that the change will be made to a component that matches the name pattern (that was given as the first argument). The `global` argument is also set to `TRUE`, which means that this change will be made to not just the first component that matches the name pattern, but to all components that match. The `gp` argument specifies the change that we want to make (make the lines nice and thick).

It would not be reasonable to expect the high-level `barchart()` function to provide an argument that allows for this sort of customisation, but, because **lattice** has named everything that it draws, `barchart()` does not need to cater for every possible customisation. Low-level access to individual shapes can be used instead because individual shapes can be identified by name.

Post-processing graphics

This section discusses how naming the individual shapes within a plot allows not just minor customisations, but general transformations to be applied to a plot.

The R graphics system has always encouraged the philosophy that a high-level plotting function is only a starting point. Low-level functions have always been provided so that a plot can be customised by *adding some new drawing* to the plot.

The previous section demonstrated that, if every shape within a plot has a label, it is also possible to customise a plot by *modifying the existing shapes*

within a plot.

However, we can go even further than just modifying the existing parameters of a shape. In theory, we can think of the existing shapes within a picture as a basis for more general post-processing of the image.

As an example, one thing that we can do is to query the existing components of a plot to determine the position or size of an existing component. This means that we can position or size new drawing in relation to the existing plot. The following code uses this idea to add a rectangle around the x-axis label of the plot in Figure 4 (see Figure 8). The `grobWidth()` function is used to calculate the width of the rectangle from the width of the x-axis label. The first argument to `grobWidth()` is the name of the x-axis label `grob`. The `downViewport()` function is used to make sure that we draw the rectangle in the right area on the page.¹

```
> xyplot(mpg ~ disp, mtcars)

> rectWidth <- grobWidth("plot_01.xlab")

> downViewport("plot_01.xlab.vp")
> grid.rect(width=rectWidth + unit(2, "mm"),
+           height=unit(1, "lines"),
+           gp=gpar(lwd=2),
+           name="xlabRect")
```

The display list now contains an new rectangle `grob`, as shown below.

```
> grid.ls(fullNames=TRUE)

rect [plot_01.background]
text [plot_01.xlab]
text [plot_01.ylab]
segments [plot_01.ticks.top.panel.1.1]
segments [plot_01.ticks.left.panel.1.1]
text [plot_01.ticklabels.left.panel.1.1]
segments [plot_01.ticks.bottom.panel.1.1]
text [plot_01.ticklabels.bottom.panel.1.1]
segments [plot_01.ticks.right.panel.1.1]
points [plot_01.xyplot.points.panel.1.1]
rect [plot_01.border.panel.1.1]
rect [xlabRect]
```




Figure 8: The **lattice** plot from Figure 4 with a rectangle added around the x-axis label.

Importantly, the new `grob` depends on the size of the existing x-axis label `grob` within the scene. For example, if we edit the x-axis label again, as below, the rectangle will grow to accommodate the new label (see Figure 9).

```
> grid.edit("plot_01.xlab",
+           label="Displacement",
+           gp=gpar(fontface="bold.italic"))
```




Figure 9: The **lattice** plot from Figure 4 with a rectangle added around the *modified* x-axis label.

A more extreme example of post-processing is demonstrated in the code below. In this case, we again query the existing x-axis label to determine its

¹This `downViewport()` works because the `grid` viewports that `lattice` creates to draw its plots all have names too!

width, but this time, rather than adding a rectangle, we *replace* the label with a rectangle (in effect, we “redact” the x-axis label; see Figure 10).

```
> xyplot(mpg ~ disp, mtcars)

> xaxisLabel <- grid.get("plot_01.xlab")
> grid.set("plot_01.xlab",
+         rectGrob(width=grobWidth(xaxisLabel) +
+                 unit(2, "mm"),
+                 height=unit(1, "lines"),
+                 gp=gpar(fill="black"),
+                 name="plot_01.xlab"))
```




Figure 10: The **lattice** plot from Figure 4 with the x-axis label redacted (replaced with a black rectangle).

The display list now consists of the same number of grobs as in the original plot, but now the grob named "plot_01.xlab" is a *rectangle* instead of text (see the second line of the output below).

```
> grid.ls(fullNames=TRUE)

rect[plot_01.background]
rect[plot_01.xlab]
text[plot_01.ylab]
segments[plot_01.ticks.top.panel.1.1]
segments[plot_01.ticks.left.panel.1.1]
text[plot_01.ticklabels.left.panel.1.1]
segments[plot_01.ticks.bottom.panel.1.1]
text[plot_01.ticklabels.bottom.panel.1.1]
segments[plot_01.ticks.right.panel.1.1]
points[plot_01.xyplot.points.panel.1.1]
rect[plot_01.border.panel.1.1]
```

The artificial examples shown in this section so far have been deliberately simple in an attempt to make the basic concepts clear, but the ideas can be applied on a much larger scale and to greater effect. For example, the **gridSVG** package (Murrell, 2011) uses these techniques to transform static R plots into

dynamic and interactive plots for use in web pages. It has functions that modify existing grobs on the **grid** display list to add extra information, like hyperlinks and animation, and it has functions that transform each grob on the **grid** display list to SVG code. The following code shows a simple demonstration where the original **lattice** plot is converted to an SVG document with a hyperlink on the x-axis label. Figure 11 shows the SVG document in a web browser.

```
> xyplot(mpg ~ disp, mtcars)

> library(gridSVG)

> url <-
+   "http://www.mortality.org/INdb/2008/02/12/8/document.pdf"

> grid.hyperlink("plot_01.xlab", href=url)
> gridToSVG("xyplot.svg")
```

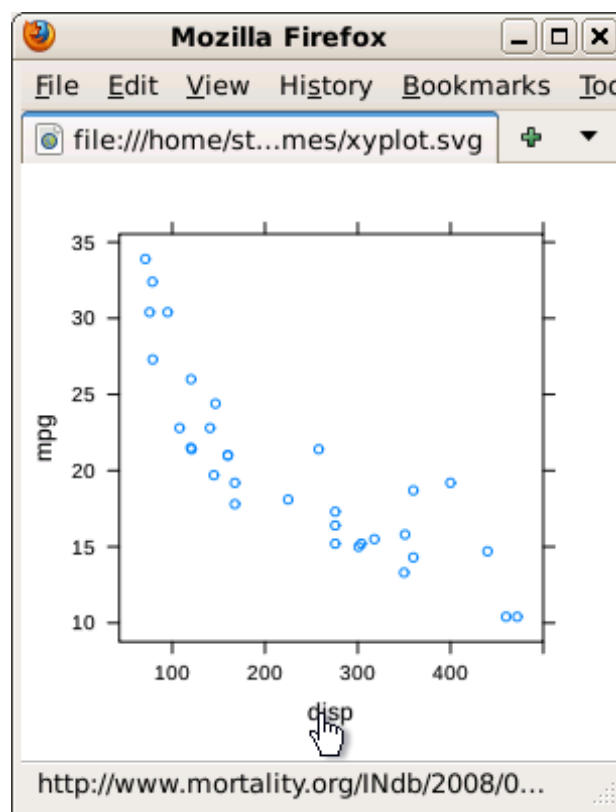


Figure 11: The **lattice** plot from Figure 4 transformed into an SVG document with a hyperlink on the x-axis label.

The significant part of that code is the first argument in the call to the `grid.hyperlink()` function, which demonstrates the ability to specify a plot component by name.

More sophisticated embellishments are also possible with **gridSVG** because the names of plot components are exported to SVG code as `id` attributes of the corresponding SVG elements. This facilitates the development of javascript code to allow user interaction with the SVG plot and allows for the possibility of CSS styling of the SVG plot.

Naming schemes

The basic message of this article is straightforward: name everything that you draw with **grid**. However, deciding what names to use—deciding on a *naming scheme*—is not necessarily so easy.

The approach taken in the **lattice** package is to attempt to reflect the structure of the plot in the naming scheme. For example, everything that is drawn within a panel region has the word "panel" in its name, along with a suffix of the form *i.j* to identify the panel row and column.

The decision may be made a lot easier if a plot is drawn from **gTrees** rather than simple grobs, because the **gTrees** reflect the plot structure already and names for individual components can be chosen to reflect just the "local" role of each plot component. The naming scheme in the **ggplot2** package (Wickham, 2009) is an example of this approach.

In addition to the code developer deciding on a naming scheme, the code user also faces the problem of how to "discover" the names of the components of a plot.

From the developer side, there is a responsibility to document the naming scheme (for example, the **lattice** naming scheme is described on the package's R-Forge web site²). It may also be possible to provide a function interface to assist in constructing the names of grobs (for example, the `trellis.grobname()` function in **lattice**).

From the user side, there are tools that help to display the names of grobs in the current scene. This article has demonstrated the `grid.ls()` function, but there is also a `showGrob()` function, and the **gridDebug** package (Murrell and Ly., 2011) provides some more tools.

Caveats

The examples used for demonstrations in this article are deliberately simplified to make explanations clearer. This section addresses two complications that have not been raised previously.

One issue is that, while each call to a **grid** drawing function produces exactly one grob, a single call to a drawing function may produce *more than one* shape in the scene. In the very first example in this article (Figure 1), the call to `grid.circle()` creates one circle grob and draws one circle.

```
> grid.circle(r=.25, name="middleCircle")
```

The call to `grid.text()` also creates only one text grob, but it draws *three* pieces of text.

```
> grid.text(c("text", "circle", "rect"),
+          x=1:3/4, gp=gpar(cex=c(3, 1, 1)),
+          name="leftText")
```

²<http://lattice.r-forge.r-project.org/documentation.php>

Modifying this text grob is slightly more complex because there are three locations and three sets of graphical parameter settings for this single grob. For example, if we modify the text grob and supply a single `cex` setting, that is applied to all pieces of text (see Figure 12).

```
> grid.edit("leftText", gp=gpar(cex=2))
```




Figure 12: The simple shapes from Figure 1 with the text grob modified using a single `cex` value.

If we want to control the `cex` for each piece of text separately, we must provide three new settings (see Figure 13).

```
> grid.edit("leftText", gp=gpar(cex=c(1, 2, 3)))
```




Figure 13: The simple shapes from Figure 1 with the text grob modified using three distinct `cex` values.

Another topic that has not been mentioned is **grid** viewports. This is because, although **grid** viewports can also be named, they cannot be edited in the same way as grobs (the names are only used for navigation between viewports). Furthermore, **grid** does not allow the `vp` slot on a grob to be modified and the `name` slot on grobs is also out of bounds. These limitations are imposed because the consequences of allowing modifications are either nonsensical or too complex to currently be handled by **grid**.

Discussion

In summary, if we specify an explicit name for every shape that we draw using **grid**, we allow low-level access to every grob within a scene. This allows us to make very detailed customisations to the scene, without the need for long lists of arguments in high-level plotting functions, and it allows us to query and transform the scene in a wide variety of ways.

An alternative way to provide access to individual shapes within a plot is to allow the user to simply select shapes on screen via a mouse. How does this compare to a naming scheme?

Selection using a mouse works well for some sorts of modifications (see, for example, the **playwith** package; Andrews, 2010), but providing access to individual shapes by name is more efficient, more general, and more powerful. For example, if we write code to make modifications, referencing grobs by name, we have a record of what we have done, we can easily automate large numbers of modifications, we can share our modification techniques, and we can express more complex modifications (like “highlight every sixth bar”).

Another alternative way to provide detailed control over a scene is simply to modify the original R code that drew the scene. Why go to the bother of naming grobs when we can just modify the original R code?

If we have written the original code, then modifying the original code may be the right approach. However, if we draw a plot using someone else’s code (for example, if we call a **lattice** function), we do not have easy access to the code that did the drawing. Even though it is possible to see the code that did the drawing, understanding it and then modifying it may require a considerable effort, especially when that code is of the size and complexity of the code in the **lattice** package.

A parallel may be drawn between this idea of naming every shape within a scene and the general idea of *markup*. In a sense, what we are aiming to do is to provide a useful label for each meaningful component of a scene. Given tools that can select parts of the scene based on the labels, the scene becomes a “source” that can be transformed in many different ways. When we draw a scene in this way, it is not just an end point that satisfies our own goals. It also creates a resource that others can make use of to produce new resources. When we write code to draw a scene, we are not only concerned with producing an image on screen or ink on a page; we also allow for other possible uses of the scene in ways that we may not have anticipated.

Acknowledgements

Thanks to Wolfgang Viechtbauer for useful comments on an early draft of this article and to the anonymous referees for numerous useful suggestions for improvements.

Bibliography

- F. Andrews. *playwith: A GUI for interactive plots using GTK+*, 2010. URL <http://CRAN.R-project.org/package=playwith>. R package version 0.9-53. [p12]
- P. Murrell. *gridSVG: Export grid graphics as SVG*, 2011. URL <http://CRAN.R-project.org/package=gridSVG>. R package version 0.7-0. [p10]
- P. Murrell and V. Ly. *gridDebug: Debugging Grid Graphics*, 2011. URL <http://r-forge.r-project.org/projects/griddebug/>. R package version 0.2. [p11]
- D. Sarkar. *Lattice: Multivariate Data Visualization with R*. Springer, New York, 2008. URL <http://lmdvr.r-forge.r-project.org>. ISBN 978-0-387-75968-5. [p7]
- W. Viechtbauer. Conducting meta-analyses in R with the metafor package. *Journal of Statistical Software*, 36(3):1–48, 2010. URL <http://www.jstatsoft.org/v36/i03/>. [p5]
- H. Wickham. *ggplot2: elegant graphics for data analysis*. Springer New York, 2009. ISBN 978-0-387-98140-6. URL <http://had.co.nz/ggplot2/book>. [p11]

Paul Murrell
 Department of Statistics
 The University of Auckland
 New Zealand
 paul@stat.auckland.ac.nz