
A tour of R/qtl

Karl W Broman

Department of Biostatistics, Johns Hopkins University
<http://biosun01.biostat.jhsph.edu/~kbroman/qtl>

24 November 2001

Overview of R/qtl

R/qtl is an extensible, interactive environment for mapping quantitative trait loci (QTLs) in experimental crosses. It is implemented as an add-on package for the freely available and widely used statistical language/software R (see www.R-project.org). The development of this software as an add-on to R allows us to take advantage of the basic mathematical and statistical functions, and powerful graphics capabilities, that are provided with R. Further, the user will benefit by the seamless integration of the QTL mapping software into a general statistical analysis program. Our goal is to make complex QTL mapping methods widely accessible and allow users to focus on modeling rather than computing.

A key component of computational methods for QTL mapping is the hidden Markov model (HMM) technology for dealing with missing genotype data. We have implemented the main HMM algorithms, with allowance for the presence of genotyping errors, for backcrosses, intercrosses, and phase-known four-way crosses.

The current version of R/qtl includes facilities for estimating genetic maps, identifying genotyping errors, and performing single-QTL genome scans and two-QTL, two-dimensional genome scans, by interval mapping (with the EM algorithm), Haley-Knott regression, and multiple imputation. All of this may be done in the presence of covariates (such as sex, age or treatment). The fit of higher-order QTL models, with sophisticated techniques for model comparison and model search, will be incorporated soon.

R/qtl is distributed as source code for unix or compiled code for Windows. The unix source will likely work for MacOS, but we haven't tested it. (If you would like to purchase a Mac for me, I'll see that it works.) R/qtl is released under the GNU General Public License. To download the software, you must agree to the terms in that license.

Overview of R

R is an open-source implementation of the S language. As described on the R-project homepage (www.R-project.org):

R is ‘GNU S’—A language and environment for statistical computing and graphics. R is similar to the award-winning S system, which was developed at Bell Laboratories by John Chambers et al. It provides a wide variety of statistical and graphical techniques (linear and nonlinear modelling, statistical tests, time series analysis, classification, clustering, ...).

R is designed as a true computer language with control-flow constructions for iteration and alternation, and it allows users to add additional functionality by defining new functions. For computationally intensive tasks, C, C++ and Fortran code can be linked and called at run time.

R is freely available for Windows, unix and MacOS, and may be downloaded from the Comprehensive R Archive Network (CRAN; cran.r-project.org).

Learning R may require a formidable investment of time, but it will definitely be worth the effort. Numerous free documents on getting started with R are available on CRAN. In addition, several books are available on S/S-PLUS, which is very similar to R. For example, see WN Venables, BD Ripley (1999) *Modern Applied Statistics with S-PLUS*, 3rd edition. Springer.

Current R/qtl functions

Sample data	badorder fake.4way fake.bc fake.f2 hyper listeria	An intercross with misplaced markers Simulated data for a 4-way cross Simulated data for a backcross Simulated data for an F ₂ intercross Backcross data on salt-induced hypertension Intercross data on Listeria monocytogenes susceptibility
Input/output	read.cross write.cross	Read data for a QTL experiment Write data for a QTL experiment to a file
Simulation	sim.cross sim.map	Simulate a QTL experiment Generate a genetic map
Summaries	geno.table nind nchr nphe nmar plot.cross plot.missing plot.info summary.cross totmar	Create table of genotype distributions Number of individuals Number of chromosomes Number of phenotypes Number of markers per chromosome Plot various features of a cross object Plot grid of missing genotypes Plot the proportion of missing genotype data Print summary of QTL experiment Total number of markers
Data manipulation	c.cross clean drop.markers drop.nullmarkers fill.genotype pull.map replace.map subset.cross switch.order	Combine crosses Remove intermediate calculations from a cross Remove a list of markers Remove markers without data Fill in holes in genotype data by imputation Viterbi Pull out the genetic map from a cross Replace the genetic map of a cross Select a subset of chromosomes and/or individuals from a cross Switch the order of markers on a chromosome
HMM engine	argmax.genotype calc.genoprob sim.genotype	Reconstruct underlying genotypes Calculate conditional genotype probabilities Simulate genotypes given observed marker data
Interval mapping	scanone scantwo plot.scanone plot.scantwo summary.scanone summary.scantwo	Genome scan with a single QTL model Two-dimensional genome scan with a two-QTL model Plot output for a single genome scan Plot output for a two-dimensional genome scan Print summary of scanone output Print summary of scantwo output
Genetic mapping	est.map est.rf plot.map plot.rf ripple summary.ripple	Estimate genetic maps Estimate pairwise recombination fractions Plot genetic map(s) Plot recombination fractions Assess marker order by permuting groups of adjacent markers Print summary of ripple output
Genotyping errors	calc.errorlod plot.errorlod plot.geno top.errorlod	Calculate Lincoln & Lander (1992) error LOD scores Plot grid of error LOD values Plot observed genotypes, flagging likely errors List genotypes with highest error LOD values

Preliminaries

Use of the R/qt1 package requires considerable knowledge of the R language/environment. We hope that the examples presented here will be understandable with little prior knowledge of R or Splus, especially because we neglect to explain the syntax of R. Several books, as well as some free documents, are available to assist the user in learning R; see the R project websites cited above. We assume here that the user is running a version of the Windows operating system.

1. To start R, double-click its icon.
2. To exit, type:

```
q()
```

Click yes or no to save or discard your work.

3. Load the R/qt1 package:

```
library(qtl)
```

4. View the objects in your workspace:

```
ls()
```

5. The best way to get help on the functions and data sets in R (and in R/qt1) is via the html version of the help files. One way to get access to this is to click (on the menu bar) **Help → R language (html)**. If you then click on **Packages → qtl**, you can see all of the available functions and datasets in R. For example, look at the help file for the function `read.cross`.

An alternative method to view this help file is to type one of the following:

```
help(read.cross)  
?read.cross
```

The html version of the help files are somewhat easier to read, and allow use of hotlinks between different functions. Typing the following will make the above text commands open the html version.

```
options(htmlhelp=TRUE)
```

You can create a file "`c:\.Rprofile`" containing any R code to be executed whenever R is started. The commands `library(qtl)` and `options(htmlhelp=TRUE)` are good candidates for placement in such a file.

Data import

A difficult first step in the use of most data analysis software is the import of data. With R/qt1, one may import data in several different formats by use of the function `read.cross`. The internal data structure used by R/qt1 is rather complicated, and is described in the help file for `read.cross`. We won't discuss data import any further here, except to say that the comma-delimited format ("csv") is recommended. If you have trouble importing data, send an email to Karl Broman (`kbrroman@jhsph.edu`) perhaps attaching examples of your data files. (Such data will be kept confidential.)

Example 1: Hypertension

As a first example, we consider data from an experiment on hypertension in the mouse (Sugiyama et al., Genomics 71:70-77, 2001), kindly provided by Bev Paigen and Gary Churchill.

1. First, get access to the data, see that it is in your workspace, and view its help file.

```
data(hyper)  
ls()  
?hyper
```

2. We will postpone discussion of the internal data structure used by R/qt1 until later. For now we'll just say that the data `hyper` has "class" "cross". The function `summary.cross` prints summary information on such data. We can call that function directly, or we may simply use `summary` and the data is sent to the appropriate function according to its class.

```
summary(hyper)
```

Several other utility functions are available for getting summary information on the data. Hopefully these are self-explanatory.

```
nind(hyper)
nphe(hyper)
nchr(hyper)
totmar(hyper)
nmar(hyper)
```

3. Plot a summary of these data.

```
plot(hyper)
```

In the upper left, black pixels indicate missing genotype data. Note that one marker has no genotype data. In the upper right, the genetic map of the markers is shown. In the lower left, a histogram of the phenotype is shown.

The Windows version of R has a slick method for recording graphs, so that one may page up and down through a series of plots. To initiate this, click (on the menu bar) **History → Recording**.

We may plot the individual components of the above multi-plot figure as follows.

```
plot.missing(hyper)
plot.map(hyper)
hist(hyper$pheno[,1], breaks=30)
```

4. Note the odd pattern of missing data; we may make this missing data plot with the individuals ordered according to the value of their phenotype.

```
plot.missing(hyper, reorder=TRUE)
```

We see that, for most markers, only individuals with extreme phenotypes were genotyped. At many markers (in regions of interest), markers were typed only on recombinant individuals.

5. The function `drop.nullmarkers` may be used to remove markers that have no genotype data (such as the marker on chromosome 14). A call to `totmar` will show that there are now 173 markers (rather than 173, as there were initially).

```
hyper <- drop.nullmarkers(hyper)
totmar(hyper)
```

6. Estimate recombination fractions between all pairs of markers, and plot them. This also calculates LOD scores for the test of $H_0: r = 1/2$. The plot of the recombination fractions can be either with recombination fractions in the upper part and LOD scores below, or with just recombination fractions or just LOD scores. Note that red corresponds to a small recombination fraction or a big LOD score, while blue is the reverse. Gray values correspond are missing.

```
hyper <- est.rf(hyper)
plot.rf(hyper)
plot.rf(hyper,c(1,4))
```

There are some very strange patterns in the recombination fractions, but this is due to the fact that some markers were typed only on recombinant individuals.

For example, on chr 6, the 9th marker shows a high recombination fraction with all other markers on the chromosome, but a plot of the missing data shows that this marker was typed only on a selected number of individuals (those showing recombination events across the interval).

```
plot.rf(hyper,6)
plot.missing(hyper,6)
```

7. Re-estimate the genetic map (keeping the order of markers fixed), and plot the original map against the newly estimated one.

```
newmap <- est.map(hyper, error.prob=0.01)
plot.map(hyper, newmap)
```

We see some map expansion, especially on chromosomes 6, 13 and 18. It is questionable whether we should replace the map or not. Keep in mind that the previous map locations are based on a limited number of meioses. If one wished to replace the genetic map with the estimated one, it could be done as follows:

```
hyper <- replace.map(hyper, newmap)
```

8. We now turn to the identification of genotyping errors. In the following, we calculate the error LOD scores of Lincoln and Lander (1992). A LOD score is calculated for each individual at each marker; large scores indicate likely genotyping errors.

The core of R/qtl is a set of functions which make use of the hidden Markov model (HMM) technology to calculate QTL genotype probabilities, to simulate from the joint genotype distribution and to calculate the most likely sequence of underlying genotypes (all conditional on the observed marker data). This is done in a quite general way, with possible allowance for the presence of genotyping errors. Of course, we must assume no crossover interference. The function calc.genoprob performs the first of these; the values are used in the calculation of the error LOD scores.

```
hyper <- calc.genoprob(hyper, error.prob=0.01)
hyper <- calc.errorlod(hyper, error.prob=0.01)
```

We may now make various plots and receive other summary information regarding which genotypes are likely in error.

```
plot.errorlod(hyper)
top.errorlod(hyper)
plot.errorlod(hyper, chr=c(4,11,16))
```

9. The function plot.geno may be used to inspect the observed genotypes for a chromosome, with likely genotyping errors flagged. Of course, it's difficult to look at too many individuals at once. Note that white = AA and black = AB (for a backcross).

```
plot.geno(hyper, chr=16, ind=71:90, min.sep=4)
```

We don't have any utilities for fixing any apparent errors; it would be best to go back to the raw data.

10. We now, finally, get to interval mapping. We re-calculate the QTL genotype probabilities (at 2 cM steps), and then do a genome scan, with a normal model. We may use maximum likelihood via the EM algorithm or use Haley-Knott regression (Haley and Knott 1992).

```
hyper <- calc.genoprob(hyper, step=2)
out.em <- scanone(hyper)
out.hk <- scanone(hyper, method="hk")
```

We may also use the multiple imputation method of Sen and Churchill (2001). This requires that we first use sim.geno to simulate from the joint genotype distribution, given the observed marker data.

```
hyper <- sim.geno(hyper, step=2, n.draws=64)
out.imp <- scanone(hyper, method="imp")
```

The output of scanone has class "scanone"; the function summary.scanone displays the maximum LOD score on each chromosome for which the LOD exceeds a specified threshold.

```
summary(out.em)
summary(out.em, 3)
summary(out.hk, 3)
summary(out.imp, 3)
```

11. We may also plot the results. plot.scanone can plot up to three genome scans at once, provided that they conform appropriately. Alternatively, one may use the argument add.

```

plot(out.em, chr=c(1,4,15))
plot(out.hk, chr=c(1,4,15), col="blue")
plot(out.imp, chr=c(1,4,15), col="red")
plot(out.hk, out.imp, out.em, ,chr=c(1,4,15), col=c("red","blue","black"))

```

12. The function `scanone` may also be used to perform a permutation test to get a genome-wide LOD significance threshold. This can take a long time, so we'll just do three permutations.

```

operm.hk <- scanone(hyper, method="hk", n.perm=3)
quantile(operm.hk, 0.95)

```

13. We should mention at this point that the function `save.image` may be used to save your workspace to disk. If R crashes, you will wish you had used this.

```
save.image()
```

14. The function `scantwo` performs a two-dimensional genome scan with a two-QTL model. For every pair of positions, it calculates a joint LOD score and a LOD score for a test of epistasis. This can be quite time consuming, and so you may wish to do the calculations on a coarser grid.

```

hyper.coarse <- calc.genoprob(hyper, step=10, error.prob=0.01)
out2.hk <- scantwo(hyper.coarse, method="hk")
out2.em <- scantwo(hyper.coarse, method="em")
hyper.coarse <- sim.geno(hyper.coarse, step=10, n.draws=16)
out2.imp <- scantwo(hyper.coarse, method="imp")

```

15. The output of `scantwo` has class "`scantwo`"; there are functions for obtaining summaries and plots, of course.

```

summary(out2.hk, c(8,3,3))
summary(out2.hk, c(0,4,1000))
summary(out2.hk, c(0,1000,4))

```

The summary function requires LOD thresholds on the joint LOD score, the epistasis LOD score, and on the conditional LOD score for a single QTL. For each pair of chromosomes, the locus pair with the maximum joint LOD is printed if the joint LOD exceeds its threshold and either the epistasis LOD score exceeds its threshold or both loci have conditional LOD scores that exceed their threshold.

```

plot(out2.hk)
plot(out2.hk,chr=c(1,4))

```

In the plot of the results of `pairsScan`, the epistasis LOD scores appear in the upper left triangle and the joint LOD scores appear in the lower right triangle. The color scale on the right indicates separate scales for the epistasis and joint LOD scores (on the left and right, respectively).

16. One may also use `scantwo` to perform permutation tests in order to obtain genome-wide LOD significance thresholds. These may take days to complete, so we'll again just do three permutations.

```

operm2 <- scantwo(hyper.coarse, method="hk", n.perm=3)
apply(operm2, 2, quantile, 0.95)

```

The output of `scantwo`, when `n.perm` is specified, is a matrix with two columns—the maximum joint and epistasis LOD scores, across the two-dimensional genome scan, for each permutation replicate.

17. You may wish to clean up your workspace before we move on to the next example.

```

ls()
rm(list=ls())

```

Example 2: Genetic mapping

R/qtl includes some utilities for estimating genetics maps and checking marker orders. In this example, we describe the use of these utilities.

1. Get access to some sample data. This is simulated data with some errors in marker order.

```
data(badorder)
summary(badorder)
plot(badorder)
```

2. Estimate recombination fractions between all pairs of markers, and plot them.

```
badorder <- est.rf(badorder)
plot.rf(badorder)
```

It appears that markers on chromosomes 2 and 3 have been switched. Moreover, if we look more closely at the recombination fractions for chromosome 1, there seem to be some errors in marker order.

```
plot.rf(badorder, chr=1)
```

Unfortunately, R/qtl doesn't have facilities for easily fixing such problems. It might be best to work directly with the raw data and then import it again.

3. Re-estimate the genetic map.

```
newmap <- est.map(badorder)
plot.map(badorder, newmap)
```

This really shows that the problems on chromosomes 2 and 3.

4. We can, however, further check the marker order on chromosome 1. The function `ripple` will consider all permutations of a sliding window of adjacent markers, re-estimating the genetic map and calculating a LOD score (\log_{10} likelihood ratio) relative to the initial order. This may be done with allowance for the presence of genotyping errors. As might be expected, the process can be rather computationally intensive.

The following checks the marker order on chromosome 1, permuting groups of three adjacent markers, assuming a genotyping error rate of 1%.

```
rip3 <- ripple(badorder, chr=1, window=3, err=0.01)
summary(rip3)
```

In the summary output, markers 9–11 clearly need to be flipped. (Positive LOD scores indicate that the alternate order has a higher likelihood than the original.) There also seems to be a problem with the order of markers 4–6.

We can switch the order of markers 9–11 with the function `switch.order` (which works only for a single chromosome) and then re-assess the order. Note that the second row of `rip3` corresponds to the improved order.

```
badorder.rev <- switch.order(badorder, 1, rip3[2,])
rip3r <- ripple(badorder.rev, chr=1, window=3, err=0.01)
summary(rip3r)
```

It's clear that marker 6 needs to be moved. It's less clear whether it belongs between markers 4 and 5 or between markers 3 and 4. The former is the best supported order.

```
badorder.rev <- switch.order(badorder.rev, 1, rip3r[2,])
```

5. We can look again at the recombination fractions for this chromosome.

```
badorder.rev <- est.rf(badorder.rev)
plot.rf(badorder.rev, 1)
```

Example 3: Listeria susceptibility

In order to demonstrate further uses of the function `scanone`, we consider some data on susceptibility to *Listeria monocytogenes* in mice (Boyartchuk et al., Nature Genetics 27:259-260, 2001). These data were kindly provided by Victor Boyartchuk and Bill Dietrich.

1. Get access to the data and view some summaries.

```
data(listeria)
summary(listeria)
plot(listeria)
plot.missing(listeria)
```

Note that in the missing data plot, gray pixels are partially missing genotypes (e.g., a genotype may be known to be either AA or AB, but not which).

The phenotype here is the survival time of a mouse (in hours) following infection with *Listeria monocytogenes*. Individuals with a survival time of 264 hours are those that recovered from the infection.

2. We'll take the log phenotype and add it to the phenotype data.

```
y <- log(listeria$pheno[,1])
listeria$pheno <- cbind(listeria$pheno, logSurv=y)
plot(listeria)
```

3. Estimate pairwise recombination fractions.

```
listeria <- est.rf(listeria)
plot.rf(listeria)
plot.rf(listeria,c(5,13))
```

4. Re-estimate the genetic map.

```
newmap <- est.map(listeria)
plot.map(listeria, newmap)
```

5. Investigate genotyping errors.

```
listeria <- calc.genoprob(listeria,error.prob=0.01)
listeria <- calc.errorlod(listeria,error.prob=0.01)plot.errorlod(listeria)
top.errorlod(listeria)
plot.errorlod(listeria,c(5,13))
plot.geno(listeria, chr=13, ind=61:70, min.sep=2 )
```

Note that in the plot given by `plot.geno`, for an intercross, white = AA, gray = AB, black = BB, green = AA or AB, and orange = AB or BB.

6. Now on to the QTL mapping.

7. Recall that the phenotype distribution shows a clear departure from the standard assumptions for interval mapping; 30% of the mice survived longer than 264 hours, and were considered recovered from the infection.

One approach for these data is to use the two-part model considered by Boyartchuk et al. (2001). In this model, a mouse with genotype g has probability p_g of surviving the infection. If it does die, its log survival time is assumed to be distributed $\text{normal}(\mu_g, \sigma^2)$. Analysis proceeds by maximum likelihood via an EM algorithm. Three LOD scores are calculated. $\text{LOD}(p, \mu)$ is for the test of the null hypothesis $p_g \equiv p$ and $\mu_g \equiv \mu$. $\text{LOD}(p)$ is for the test of the hypothesis $p_g \equiv p$ but the μ are allowed to vary. $\text{LOD}(\mu)$ is for the test of the hypothesis $\mu_g \equiv \mu$ but the p are allowed to vary.

The function `scanone` will fit the above model when the argument `model = "2part"`. One must also specify the argument `upper`, which indicates whether the spike in the phenotype is the maximum phenotype (as it with this phenotype; take `upper=TRUE`) or the minimum phenotype (take `upper=FALSE`). For this model, only the EM algorithm has been implemented so far.

```

listeria <- calc.genoprob(listeria, step=2)
out.2p <- scanone(listeria, pheno.col=2, model="2part", upper=TRUE)

```

Note that, because this model has three extra parameters, the appropriate LOD threshold is higher—around 4.5 rather than 3.5. The three different LOD curves are in columns 3–5 of the output. The function `plot.scanone` assumes the LOD score is in the third column, and so we plot these extra LOD scores by removing the earlier columns.

```

summary(out.2p)
summary(out.2p, 4.5)
plot(out.2p)
plot(out.2p, out.2p[,-3], out.2p[,-(3:4)], chr=c(1,5,13,15),
     lty=1, col=c("black", "red", "blue"))

```

Note that the locus on chromosome 1 shows effect mostly on the mean time-to-death, conditional on death, the locus on chromosome 5 shows effect mostly on the probability of survival, and the loci on chromosomes 13 and 15 shows some effect on each.

8. Permutation tests may be performed as before. The output will have three columns, corresponding to the three LOD scores.

```

operm.2p <- scanone(listeria, model="2part", pheno.col=2,
                      upper=TRUE, n.perm=3)
apply(operm.2p, 2, quantile, 0.95)

```

9. Alternatively, one may perform separate analyses of the log survival time, conditional on death, and the binary phenotype survival/death. First we set up these phenotypes.

```

z <- y <- x <- listeria$pheno[,2]
mx <- max(x, na.rm=TRUE)
y[!is.na(x) & x==mx] <- NA
z[!is.na(x) & x<mx] <- 0
z[!is.na(x) & x==mx] <- 1
listeria$pheno <- cbind(listeria$pheno, logSurv2=y, binary=z)
plot(listeria)

```

We use standard interval mapping for the log survival time conditional on death; the results are slightly different from $\text{LOD}(\mu)$.

```

out.mu <- scanone(listeria, pheno.col=3)
plot(out.mu, out.2p[,-(3:4)], chr=c(1,5,13,15))

```

We can use `scanone` with `model="binary"` to analyze the binary phenotype. Again, the results are only slight different from $\text{LOD}(p)$.

```

out.p <- scanone(listeria, pheno.col=4, model="binary")
plot(out.p, out.2p[,-3], chr=c(1,5,13,15))

```

10. A further approach is to use a non-parametric form of interval mapping. R/qt1 uses an extension of the Kruskal-Wallis test statistic. Use `scanone` with `model="np"`. In this case, the argument `method` is ignored; the analysis method is much like Haley-Knott regression. If the argument `ties.random=TRUE`, tied phenotypes are ranked at random. If `ties.random=FALSE`, tied phenotypes are given the average rank and a correction is applied to the LOD score.

```

out.np1 <- scanone(listeria, model="np", ties.random=TRUE)
out.np2 <- scanone(listeria, model="np", ties.random=FALSE)
plot(out.np1, out.np2)
plot(out.2p, out.np1, out.np2, chr=c(1,5,13,15), lty=1,
     col=c("black", "blue", "red"))

```

Note that the significance threshold for the non-parametric genome scan will be quite a bit smaller than that for the two-part model. The two approaches give basically the same results. Randomizing ties for the non-parametric approach can give quite variable results in the case of a great number of ties, and so we would recommend the use of `ties.random=FALSE` in this case.

Example 4: Covariates in QTL mapping

As a further example, we illustrate the use of covariates in QTL mapping. We consider some simulated backcross data.

1. Get access to the data.

```
data(fake.bc)
summary(fake.bc)
plot(fake.bc)
```

2. Perform genome scans for the two phenotypes without covariates.

```
fake.bc <- calc.genoprob(fake.bc, step=2.5)
out1.nocovar <- scanone(fake.bc, pheno.col=1)
out2.nocovar <- scanone(fake.bc, pheno.col=2)
```

3. Perform genome scans with age as an additive covariate

```
ac <- fake.bc$pheno[, "age"]
out1.covar.a <- scanone(fake.bc, pheno.col=1, addcov=ac)
out2.covar.a <- scanone(fake.bc, pheno.col=2, addcov=ac)
```

4. Perform genome scans with age as an additive covariate and sex as a covariate interacting with QTL genotype.

```
ac <- fake.bc$pheno[, c("sex", "age")]
ic <- fake.bc$pheno[, "sex"]
out1.covar.b <- scanone(fake.bc, pheno.col=1, addcov=ac, intcov=ic)
out2.covar.b <- scanone(fake.bc, pheno.col=2, addcov=ac, intcov=ic)
```

5. Get some summaries.

```
summary(out1.nocovar, 3)
summary(out1.covar.a, 3)
summary(out1.covar.b, 3)
summary(out2.nocovar, 3)
summary(out2.covar.a, 3)
summary(out2.covar.b, 3)
```

6. Make some plots.

```
plot(out1.nocovar, out1.covar.a, out1.covar.b, lty=1,
     chr=c(2,5,10), col=c("black", "blue", "red"))
plot(out2.nocovar, out2.covar.a, out2.covar.b, lty=1,
     chr=c(2,5,10), col=c("black", "blue", "red"))
```

Internal data structure

Finally, let us briefly describe the rather complicated data structure that R/qtl uses for QTL mapping experiments. This will be rather dull, and will require a good deal of familiarity with the R (or S) language. The choice of data structure required some balance between ease of programming and simplicity for the user interface. The syntax for references to certain pieces of the internal data can become extremely complicated.

1. Get access to some sample data.

```
data(fake.bc)
```

2. First, these objects have a “class,” which indicates that it corresponds to data for an experimental cross, and gives the cross type. By having class `cross`, the functions `plot` and `summary` know to send the data to `plot.cross` and `summary.cross`.

```
class(fake.bc)
```

3. Every cross object has two components, one containing the genotype data and genetic maps and the other containing the phenotype data.

```
names(fake.bc)
```

4. The phenotype data is simply a matrix (more strictly a data.frame) with rows corresponding to individuals and columns corresponding to phenotypes.

```
fake.bc$pheno
```

5. The genotype data is a list with components corresponding to chromosomes. Each chromosome has a name and a class. The class for a chromosome is either "A" or "X", according to whether it is an autosome or the X chromosome.

```
names(fake.bc$geno)
sapply(fake.bc$geno, class)
```

6. Each component of geno contains two components, data (containing the marker genotype data) and map (containing the positions of the markers, in cM).

```
names(fake.bc$geno[[3]])
fake.bc$geno[[3]]$data[1:5,]
fake.bc$geno[[3]]$map
```

That's it for the raw data.

7. When one runs calc.genoprob, sim.geno, argmax.geno or calc.errorlod, the output is the input cross object with the derived data attached to each component (the chromosomes) of the geno component.

```
names(fake.bc$geno[[3]])
fake.bc <- calc.genoprob(fake.bc, step=10, err=0.01)
names(fake.bc$geno[[3]])
fake.bc <- sim.geno(fake.bc, step=10, n.draws=10, err=0.01)
names(fake.bc$geno[[3]])
fake.bc <- argmax.geno(fake.bc, step=10, err=0.01)
names(fake.bc$geno[[3]])fake.bc <- calc.errorlod(fake.bc, err=0.01)
names(fake.bc$geno[[3]])
```

8. Finally, when one runs est.rf, a matrix containing the pairwise recombination fractions and LOD scores is added to the cross object.

```
names(fake.bc)
fake.bc <- est.rf(fake.bc)
names(fake.bc)
```