

Overview of the *intervals* package

Richard Bourgon

06 June 2009

Contents

1	Introduction	1
2	Interpretation of objects	2
2.1	As a subset of \mathbb{Z} or \mathbb{R}	3
2.2	As a set of meaningful, possibly overlapping intervals	4
3	Floating point and intervals over \mathbb{R}	7
4	Notes on implementation	9
4.1	Endpoint representation	9
4.2	Efficiency	9
4.3	Checking validity	10
5	Session information	10

1 Introduction

The *intervals* packages defines two S4 classes which represent collections of intervals over either the integers (\mathbb{Z}) or the real number line (\mathbb{R}). An instance of either class consists of a two-column matrix of endpoints, plus additional slots describing endpoint closure and whether the intervals are to be thought of as being over \mathbb{Z} or \mathbb{R} .

```
> library( intervals )
> x <- Intervals( matrix( 1:6, ncol = 2 ) )
> x

Object of class Intervals
3 intervals over R:
[1, 4]
[2, 5]
[3, 6]

> x[2,2] <- NA
> x[3,1] <- 6
> x
```

```

Object of class Intervals
3 intervals over R:
[1, 4]
[2, NA]
[6, 6]

```

Objects of class `Intervals` represent collections of intervals with common endpoint closure, e.g., all left-closed, right-open. More control over endpoints is permitted with the `Intervals_full` class. (Both classes are derived from `Intervals_virtual`, which is not intended for use by package users.)

```

> y <- as( x, "Intervals_full" )
> y <- c( y, Intervals_full( c(2,3,5,7) ) )
> closed(y)[2:3,1] <- FALSE
> closed(y)[4,2] <- FALSE
> rownames(y) <- letters[1:5]
> y

```

```

Object of class Intervals_full
5 intervals over R:
a [1, 4]
b (2, NA]
c (6, 6]
d [2, 5)
e [3, 7]

```

The `size` method gives measure — counting measure over \mathbb{Z} or Lebesgue measure over \mathbb{R} — for each interval represented in an object. The `empty` method identifies intervals that are in fact empty sets, which over \mathbb{R} is not the same thing as having size 0. (Valid objects must have each right endpoint no less than the corresponding left endpoint. When one or both endpoints are open, however, valid intervals may be empty.)

```

> size(x)

[1] 3 NA 0

> empty(x)

[1] FALSE NA FALSE

> empty(y)

[1] FALSE NA TRUE FALSE FALSE

```

2 Interpretation of objects

An `Intervals` or `Intervals_full` object can be thought of in two different modes, each of which is useful in certain contexts:

1. As a (non-unique) representation of a subset of \mathbb{Z} or \mathbb{R} .
2. As a collection of (possibly overlapping) intervals, each of which has a meaningful identity.




Figure 1: The `Intervals_full` object `y`, plotted with `plot(y)`. The second row contains an `NA` endpoint, and is not shown in the plot. The empty interval is plotted as a hollow point. By default, vertical placement avoids overlaps but is compact.

2.1 As a subset of \mathbb{Z} or \mathbb{R}

The *intervals* package provides a number of basic tools for working in the first mode, where an object represents a subset of \mathbb{Z} or \mathbb{R} but the rows of the endpoint matrix do not have any external identity. Basic tools include `reduce`, which returns a sorted minimal representation equivalent to the original (dropping any intervals with `NA` endpoints), as well as `interval_union`, `interval_complement`, and `interval_intersection`.

```
> reduce( y )

Object of class Intervals_full
1 interval over R:
[1, 7]

> interval_intersection( x, x + 2 )

Object of class Intervals
2 intervals over R:
[3, 4]
[6, 6]

> interval_complement( x )

Object of class Intervals
3 intervals over R:
```

```
(-Inf, 1)
(4, 6)
(6, Inf)
```

Note that combining `x` and its complement in order to compute a union requires mixing endpoint closure types; coercion to `Intervals_full` is automatic.

```
> interval_union( x, interval_complement( x ) )
```

```
Object of class Intervals_full
1 interval over R:
(-Inf, Inf)
```

The `distance_to_nearest` function treats its `to` argument in the first mode, as just a subset of \mathbb{Z} or \mathbb{R} ; it treats its `from` argument, however, in the second mode, returning one distance for every row of the `from` object. In the example below, we also look at performance for large data sets (less than one second on a 2 GHz Intel Core 2 Duo Macintosh, although the time shown below will likely differ). A histogram of `d` is given in Figure 2.

```
> B <- 100000
> left <- runif( B, 0, 1e8 )
> right <- left + rexp( B, rate = 1/10 )
> v <- Intervals( cbind( left, right ) )
> head( v )
```

```
Object of class Intervals
6 intervals over R:
[69415031.3036516, 69415034.1141342]
[66053736.4427, 66053813.3571867]
[89806962.2227922, 89806987.38828]
[41972167.8365022, 41972184.1962154]
[37081919.2146882, 37081968.773888]
[98729361.8265539, 98729362.6934107]
```

```
> mean( size( v ) )
```

```
[1] 9.990649
```

```
> dim( reduce( v ) )
```

```
[1] 99001      2
```

```
> system.time( d <- distance_to_nearest( sample( 1e8, B ), v ) )
```

```
   user  system elapsed
0.084   0.023   0.107
```

2.2 As a set of meaningful, possibly overlapping intervals

In some applications, each row of an object's endpoint matrix has a meaningful identity, and particular points from \mathbb{Z} or \mathbb{R} may be found in more than one row. To support this mode, objects may be given row names, which are propagated




Figure 2: Histogram of distances from a random set of points to the nearest point in v . There is also a `distance_to_nearest` method for comparing two sets of intervals.

through calculations when appropriate. The `c` methods simply stack objects (like `rbind`), preserving row names and retaining redundancy, if any.

The `interval_overlap` method works in this mode. In the next example we use it to identify rows of v which are at least partially redundant, i.e., which intersect at least one other row of v . All rows overlap themselves, so we look for rows that overlap at least two rows:

```

> rownames(v) <- sprintf( "%06i", 1:nrow(v) )
> io <- interval_overlap( v, v )
> head( io, n = 3 )

$`000001`
[1] 1

$`000002`
[1] 2

$`000003`
[1] 3

> n <- sapply( io, length )
> sum( n > 1 )

[1] 1980

> k <- which.max( n )
> io[ k ]

```

```

$`009065`
[1] 9065 9349 71144 97714

> v[ k, ]

Object of class Intervals
1 interval over R:
009065 [67950087.3899087, 67950088.2896036]

> v[ io[[ k ]], ]

Object of class Intervals
4 intervals over R:
009065 [67950087.3899087, 67950088.2896036]
009349 [67950079.543516, 67950102.5363626]
071144 [67950087.5761732, 67950095.3949945]
097714 [67950086.9475305, 67950102.3226143]

```

The `which_nearest` method also respects row identity, for both its `to` and `from` arguments. In addition to computing the distance from each `from` interval to the nearest point in `to`, it also returns the row index of the `to` interval (or intervals, in case of ties) located at the indicated distance.

Another function which operates in this mode is `clusters`, which takes a set of points or intervals and identifies maximal groups which cluster together — which are separated from one another by no more than a user-specified threshold. The following code is taken from the `clusters` documentation:

```

> B <- 100
> left <- runif( B, 0, 1e4 )
> right <- left + rexp( B, rate = 1/10 )
> y <- reduce( Intervals( cbind( left, right ) ) )
> w <- 100
> c2 <- clusters( y, w )
> c2[1:3]

[[1]]
Object of class Intervals
2 intervals over R:
[31.1534735374153, 37.3788441112265]
[125.25845086202, 125.893271297303]

[[2]]
Object of class Intervals
2 intervals over R:
[327.574829570949, 328.475898602046]
[363.502644468099, 367.448802338913]

[[3]]
Object of class Intervals
4 intervals over R:
[484.682524111122, 494.442853287901]
[524.248662404716, 529.712896319106]
[549.050043337047, 570.767094324014]
[607.953327707946, 641.277237868119]

```

3 Floating point and intervals over \mathbb{R}

When `type == "R"`, interval endpoints are not truly in \mathbb{R} , but rather, in the subset which can be represented by floating point arithmetic. (For the moment, this is also true when `type == "Z"`. See Section 4.1.) This limits the endpoint values which can be represented; more importantly, if computations are performed on interval endpoints, it means that floating point error can affect whether or not endpoints coincide, whether intervals which meet at or near endpoints overlap one another, etc.

In spite of this potentially serious limitation, it is still often convenient to work with intervals with non-integer endpoints, including data where adjacent intervals exactly meet at a non-integer endpoint. To address this, the *intervals* package takes the following approach:

- Floating point representations of interval endpoints are assumed to be *exactly equal* (in the sense of `==` in R or C++) if and only if the user intends the real values corresponding to these representations to be exactly equal.
- For cases where floating point error and approximate equality are a concern, tools are provided to permit distinguishing between ambiguous and unambiguous intersection, union, etc.

In the next example, `y1` does not literally overlap `y2[2,]`, although R's `all.equal` function asserts that the gap between them is smaller than the default tolerance for equivalence up to floating point precision.

```
> delta <- .Machine[[ "double.eps" ]]^0.5
> y1 <- Intervals( c( .5, 1 - delta / 2 ) )
> y2 <- Intervals( c( .25, 1, .75, 2 ) )
> y1
```

```
Object of class Intervals
1 interval over R:
[0.5, 0.999999992549419]
```

```
> y2
```

```
Object of class Intervals
2 intervals over R:
[0.25, 0.75]
[1, 2]
```

```
> all.equal( y1[1,2], y2[2,1] )
```

```
[1] TRUE
```

```
> interval_intersection( y1, y2 )
```

```
Object of class Intervals
1 interval over R:
[0.5, 0.75]
```

The `expand` and `contract` methods, used with `type = "relative"`, permit consideration of the maximal and minimal interval sets which are consistent with the nominal endpoints — from the point of view of endpoint relative difference. The `contract` method, for example, contracts each interval in a collection so that the relative difference between original and contracted endpoints is equal to tolerance `delta`. Thus, if a relative difference less than or equal to `delta` is our criterion for approximate floating point equality, the contracted object has endpoints which are approximately equal to those of the original — even though the contracted object is a proper subset of the original. The `expand` method is similar, but generates a proper superset.

```
> contract( y1, delta, "relative" )

Object of class Intervals
1 interval over R:
[0.500000007450581, 0.999999977648258]
```

We compute two separate intersections which bound the nominal intersection:

```
> inner <- interval_intersection(
+                               contract( y1, delta, "relative" ),
+                               contract( y2, delta, "relative" )
+                               )
> inner

Object of class Intervals
1 interval over R:
[0.500000007450581, 0.749999988824129]

> outer <- interval_intersection(
+                               expand( y1, delta, "relative" ),
+                               expand( y2, delta, "relative" )
+                               )
> outer

Object of class Intervals
2 intervals over R:
[0.499999992549419, 0.750000011175871]
[0.999999985098839, 1.00000000745058]
```

Finally, we identify points which may or may not be in the intersection, depending on whether we make a conservative, literal, or anti-conservative interpretation of the nominal endpoints.

```
> interval_difference( outer, inner )

Object of class Intervals_full
3 intervals over R:
[0.499999992549419, 0.500000007450581]
[0.749999988824129, 0.750000011175871]
[0.999999985098839, 1.00000000745058]
```

The `expand` and `contract` methods have other uses as well. Here, we eliminate gaps of size 2 or smaller:


```

> x <- Intervals( c(1,10,100,8,50,200), type = "Z" )
> x

Object of class Intervals
3 intervals over Z:
[1, 8]
[10, 50]
[100, 200]

> w <- 2
> close_intervals( contract( reduce( expand(x, w/2) ), w/2 ) )

Object of class Intervals
2 intervals over Z:
[1, 50]
[100, 200]

```

4 Notes on implementation

4.1 Endpoint representation

For the moment, interval endpoints are always stored using R's *numeric* data type. Although this is wasteful from an memory and speed point of view, we do it for two reasons. First, use of R's `Inf` and `-Inf` — not possible with the *integer* type — is very convenient when computing complements. Second, the range of integers which can be represented using the *numeric* data type is considerably greater:

```

> .Machine$integer.max

[1] 2147483647

> numeric_max <- with( .Machine, double.base^double.digits )
> options( digits = ceiling( log10( numeric_max ) ) )
> numeric_max

[1] 9007199254740992

```

4.2 Efficiency

All computations are accomplished by treating intervals as pairs of tagged endpoints, sorting these endpoints (along with their tags), and then making a single pass through the results. Computational complexity for set operations is therefore $O(n \log n)$, where input object i contains n_i rows and $n = \sum_i n_i$. The same sorting approach is also used for `interval_overlap`, although if every interval in a query object of m rows overlaps every intervals in a target object of n rows, generating output alone must of necessity be $O(mn)$.

Sorted endpoint vectors are not retained in memory. If one wishes to query a particular object over and over, repeated sorting would be inefficient; in practice so far, however, such repeated querying has not been needed.

4.3 Checking validity

The code behind `which_nearest` and `reduce` (key methods in the *intervals* package, which may be directly called by the user and are also used internally in numerous locations) is written in C++ for efficiency. The compiled code makes a number of assumptions about the `SEXP` objects passed in as arguments, but does not explicitly check these assumptions. Nonetheless, when the R wrappers for the compiled code are applied to *valid* objects from the `Intervals` or `Intervals_full` classes, all assumptions will always be met. This design decision was taken so that the requirements for individual objects and their contents could be gathered together in a single, natural location: the classes' `validity` functions.

The *intervals* package provides replacement methods — e.g., `type` and `closed` — which implement error checking and preserve object validity. R's implementation of S4 classes, however, leaves object data slots exposed to the user. As a consequence, a user can directly manipulate the data slots of a valid `Intervals` or `Intervals_full` object in a way that invalidates the object, but does not generate any warning or error.

To prevent invalid objects from being passed to compiled code — and potentially generating segmentation faults or other problems — all wrapper code in this package includes a `check_valid` argument. This argument is set to `TRUE` by default, so that `validObject` is called on relevant objects before handing them off to the compiled code. For efficiency, users may choose to override this extra check if they are certain they have not manually assigned inappropriate content to objects' data slots.

5 Session information

- R Under development (unstable) (2024-08-23 r87045),
x86_64-pc-linux-gnu
- Time zone: Europe/Berlin
- TZcode source: system (glibc)
- Running under: Ubuntu 22.04.4 LTS
- Matrix products: default
- BLAS: /home/edzer/svn/R/trunk/lib/libRblas.so
- LAPACK: /home/edzer/svn/R/trunk/lib/libRlapack.so ;
LAPACK version3.12.0
- Base packages: base, datasets, grDevices, graphics, methods, stats, utils
- Other packages: intervals 0.15.5
- Loaded via a namespace (and not attached): compiler 4.5.0, tools 4.5.0