

Fluid Data

John Mount, Win-Vector LLC

November 11, 2017

Introduction

The `cdata` R package provides a powerful extension of the “fluid data” (or “coordinatized data”) concept (please see here for some notes) that goes way beyond the concepts of pivot/un-pivot.

The fluid data concept is:

- 1) Data cells have coordinates, and the dependence of these coordinates on a given data representation (a table or map) is an inessential detail to be abstracted out.
- 2) There may not be one “preferred” shape (or frame of reference) for data: you have to anticipate changing data shape many times to adapt to the tasks and tools (data relativity).

`cdata` supplies two general operators for fluid data work at database scale (and `Spark` big data scale):

- 1) `cdata::rowrecs_to_blocks()`: operators centered around SQL cross-join semantics. `un-pivot`, `tidyr::gather()`, and `cdata::unpivot_to_blocks()` are special cases of this general operator.
- 2) `blocks_to_rowrecs()`: an operator centered around SQL `group by` semantics. `pivot`, `tidyr::spread()`, `cdata::pivot_to_rowrecs()`, transpose, and one-hot-encoding are special cases of this general operator.

Because these operators are powerful, they are fairly general, and at first hard to mentally model (especially if you insist on think of them in only in terms of more a specialized operator such as `pivot`, instead of more general relational concepts such as “cross join” and “group by”). These operators are thin wrappers populating and enforcing a few invariants over a large SQL statement. That does not mean that these operators are trivial, they are thin because SQL is powerful and we have a good abstraction.

Due to the very detailed and explicit controls used in these operators- they are very comprehensible once studied. We will follow-up later with additional training material to make quicker comprehension available to more readers. This document is limiting itself to being a mere concise statement of and demonstration of the operators.

Data coordinate notation theory

We are going to introduce a explicit, dense, and powerful data coordinate notation.

Consider the following table that we call a “control table”:

```
suppressPackageStartupMessages(library("cdata"))
library("wrapr")
library("rqdatatable")

## Loading required package: rquery

library("rquery")
options(width = 160)
tng <- wrapr::mk_tmp_name_source('fdexample')

controlTable <- wrapr::build_frame(
  "group", "col1", "col2" |
  "aa"    , "c1"   , "c2"   |
  "bb"    , "c3"   , "c4"   )
knitr::kable(controlTable)
```

group	col1	col2
aa	c1	c2
bb	c3	c4

Control tables partially specify a change of data shape or change of data cell coordinates.

The specification is interpreted as follows:

The region `controlTable[, 2:ncol(controlTable)]` specifies partial coordinates of data cells in another table. In our example these partial coordinates are “c1”, “c2”, “c3”, and “c4” treated as column names. For example if our data is:

```
dat1 <- wrapr::build_frame(
  "ID" , "c1"      , "c2"      , "c3"      , "c4"      |
  "id1", "val_id1_c1", "val_id1_c2", "val_id1_c3", "val_id1_c4" |
  "id2", "val_id2_c1", "val_id2_c2", "val_id2_c3", "val_id2_c4" |
  "id3", "val_id3_c1", "val_id3_c2", "val_id3_c3", "val_id3_c4" )
knitr::kable(dat1)
```

ID	c1	c2	c3	c4
id1	val_id1_c1	val_id1_c2	val_id1_c3	val_id1_c4
id2	val_id2_c1	val_id2_c2	val_id2_c3	val_id2_c4

ID	c1	c2	c3	c4
id3	val_id3_c1	val_id3_c2	val_id3_c3	val_id3_c4

Then each data cell in `dat1` (excluding the key-columns, in this case “ID”) is named by the row-id (stored in the ID column) plus the column-name (“c1”, “c2”, “c3”, and “c4”). Knowing ID plus the column name unique identifies the data-caring cell in table `dat1`.

However, there is an alternate cell naming available from the `controlTable` notation. Each name in the region `controlTable[, 2:ncol(controlTable)]` is itself uniquely named by the `group` entry and column name of the control table itself. This means we have the following correspondence from the partial names “c1”, “c2”, “c3”, and “c4” to a new set of partial names:

```
namePairings <- expand.grid(seq_len(nrow(controlTable)),
                           2:ncol(controlTable))
colnames(namePairings) <- c("controlI", "controlJ")
namePairings$coords_style1 <-
  vapply(seq_len(nrow(namePairings)),
         function(ii) {
           as.character(paste("column:",
                              controlTable[namePairings$controlI[[ii]],
                              namePairings$controlJ[[ii]]))
         },
         character(1))
namePairings$coords_style2 <-
  vapply(seq_len(nrow(namePairings)),
         function(ii) {
           paste("group:",
                controlTable$group[[namePairings$controlI[[ii]]],
                ", column:",
                colnames(controlTable)[[namePairings$controlJ[[ii]]]])
         },
         character(1))
as.matrix(namePairings[ , c("coords_style1", "coords_style2")])

##      coords_style1 coords_style2
## [1,] "column: c1"  "group: aa , column: col1"
## [2,] "column: c3"  "group: bb , column: col1"
## [3,] "column: c2"  "group: aa , column: col2"
## [4,] "column: c4"  "group: bb , column: col2"
```

The idea is the control table is a very succinct description of the pairing of the `namePairings$coords_style1` cell partial coordinates

and the `namePairings$coords_style2` cell partial coordinates. As we have said the `namePairings$coords_style1` cell partial coordinates become full cell coordinates for the data cells in `dat1` when combined with the `dat1` ID column. The `namePairings$coords_style2` are part of a natural naming for the data cells in the following table:

```
my_db <- DBI::dbConnect(RSQLite::SQLite(), ":memory:")
dat1db <- DBI::dbWriteTable(my_db, 'dat1db', dat1,
                           overwrite = TRUE,
                           temporary = TRUE)
dat2 <- cdata::rowrecs_to_blocks_q(wideTable = 'dat1db',
                                   controlTable = controlTable,
                                   my_db = my_db,
                                   columnsToCopy = "ID",
                                   tempNameGenerator = tng)
DBI::dbReadTable(my_db, dat2) %>%
  orderby(., c("ID", "group")) %>%
  knitr::kable(.)
```

ID	group	col1	col2
id1	aa	val_id1_c1	val_id1_c2
id1	bb	val_id1_c3	val_id1_c4
id2	aa	val_id2_c1	val_id2_c2
id2	bb	val_id2_c3	val_id2_c4
id3	aa	val_id3_c1	val_id3_c2
id3	bb	val_id3_c3	val_id3_c4

For `dat2` the composite row-key (`ID`, `group`) plus the column name (one of `col1` or `col2`) gives us the positions of the data carrying cells.

So essentially the two readings of `controlTable` are a succinct representation of the explicit pairing of data cell coordinates shown in the `namePairings` table.

The Operators

In terms of the above notation/theory our two operators `rowrecs_to_blocks()` and `blocks_to_rowrecs()` are (in principle) easy to describe:

- `rowrecs_to_blocks()` reshapes data from style 1 to style 2
- `blocks_to_rowrecs()` reshapes data from style 2 to style 1.

The above is certainly succinct, but carries a lot of information and allows for a lot of different possible applications. Many important

applications are derived from how these two operators interact with row-operations and column-operations.

We give simple examples of each of the operators below.

```
rowrecs_to_blocks()

wideTableName <- 'dat'
d <- DBI::dbWriteTable(my_db, wideTableName,
  wrapr::build_frame(
    "ID" , "c1"      , "c2"      , "c3"      , "c4"      |
    "id1", "val_id1_c1", "val_id1_c2", "val_id1_c3", "val_id1_c4" |
    "id2", "val_id2_c1", "val_id2_c2", "val_id2_c3", "val_id2_c4" |
    "id3", "val_id3_c1", "val_id3_c2", "val_id3_c3", "val_id3_c4" ),
    overwrite = TRUE, temporary=TRUE)
controlTable <- wrapr::build_frame(
  "group", "col1", "col2" |
  "aa"   , "c1"  , "c2"  |
  "bb"   , "c3"  , "c4"  )
columnsToCopy <- 'ID'
cdata::rowrecs_to_blocks_q(wideTable = wideTableName,
  controlTable = controlTable,
  my_db = my_db,
  columnsToCopy = columnsToCopy,
  tempNameGenerator = tng) %.>%
DBI::dbReadTable(my_db, .) %.>%
orderby(., c("ID", "group")) %.>%
knitr::kable(.)
```

ID	group	col1	col2
id1	aa	val_id1_c1	val_id1_c2
id1	bb	val_id1_c3	val_id1_c4
id2	aa	val_id2_c1	val_id2_c2
id2	bb	val_id2_c3	val_id2_c4
id3	aa	val_id3_c1	val_id3_c2
id3	bb	val_id3_c3	val_id3_c4

```
blocks_to_rowrecs()

tallTableName <- 'dat'
d <- DBI::dbWriteTable(my_db, tallTableName,
  wrapr::build_frame(
    "ID" , "group", "col1"      , "col2"      |
    "id1", "aa"   , "val_id1_gaa_col1", "val_id1_gaa_col2" |
```

```

  "id1", "bb" , "val_id1_gbb_col1", "val_id1_gbb_col2" |
  "id2", "aa" , "val_id2_gaa_col1", "val_id2_gaa_col2" |
  "id2", "bb" , "val_id2_gbb_col1", "val_id2_gbb_col2" |
  "id3", "aa" , "val_id3_gaa_col1", "val_id3_gaa_col2" |
  "id3", "bb" , "val_id3_gbb_col1", "val_id3_gbb_col2" ),
  overwrite = TRUE, temporary=TRUE)
controlTable <- wrapr::build_frame(
  "group", "col1", "col2" |
  "aa" , "c1" , "c2" |
  "bb" , "c3" , "c4" )
keyColumns <- 'ID'
cdata::blocks_to_rowrecs_q(tallTable = tallTableName,
  controlTable = controlTable,
  keyColumns = keyColumns,
  my_db = my_db,
  tempNameGenerator = tng) %>%
  DBI::dbReadTable(my_db, .) %>%
  orderby(., "ID") %>%
  knitr::kable(.)

```

ID	c1	c2	c3	c4
id1	val_id1_gaa_col1	val_id1_gaa_col2	val_id1_gbb_col1	val_id1_gbb_col2
id2	val_id2_gaa_col1	val_id2_gaa_col2	val_id2_gbb_col1	val_id2_gbb_col2
id3	val_id3_gaa_col1	val_id3_gaa_col2	val_id3_gbb_col1	val_id3_gbb_col2

Pivot/Un-Pivot

Pivot and un-pivot (or `tidyr::spread()` and `tidyr::gather()`) are special cases of the `blocks_to_rowrecs()` and `rowrecs_to_blocks()` operators. Pivot/un-pivot are the cases where the control table has two columns.

Pivot

```

d <- data.frame(
  index = c(1, 2, 3, 1, 2, 3),
  meastype = c('meas1', 'meas1', 'meas1', 'meas2', 'meas2', 'meas2'),
  meas = c('m1_1', 'm1_2', 'm1_3', 'm2_1', 'm2_2', 'm2_3'),
  stringsAsFactors = FALSE)
knitr::kable(d)

```

index	meastype	meas
1	meas1	m1_1

index	meastype	meas
2	meas1	m1_2
3	meas1	m1_3
1	meas2	m2_1
2	meas2	m2_2
3	meas2	m2_3

```
# the cdata::pivot_to_rowrecs version
# equivalent to tidyr::spread(d, 'meastype', 'meas')
cdata::pivot_to_rowrecs(d,
  columnToTakeKeysFrom = 'meastype',
  columnToTakeValuesFrom = 'meas',
  rowKeyColumns = 'index',
  sep = '_' ) %>%
  orderby(., "index") %>%
  knitr::kable(.)
```

index	meastype_meas1	meastype_meas2
1	m1_1	m2_1
2	m1_2	m2_2
3	m1_3	m2_3

```
# the cdata::blocks_to_rowrecs() version
dtall <- DBI::dbWriteTable(my_db, "dtall", d)
controlTable <- cdata::build_pivot_control_q(tableName = "dtall",
  columnToTakeKeysFrom = 'meastype',
  columnToTakeValuesFrom = 'meas',
  my_db = my_db,
  sep = "_")
knitr::kable(controlTable)
```

meastype	meas
meas1	meastype_meas1
meas2	meastype_meas2

```
blocks_to_rowrecs_q(tallTable = "dtall",
  controlTable = controlTable,
  keyColumns = "index",
  my_db = my_db,
  tempNameGenerator = tng) %>%
  DBI::dbReadTable(my_db, .) %>%
```

```
orderby(., "index") %>%
knitr::kable(.)
```

index	meastype_meas1	meastype_meas2
1	m1_1	m2_1
2	m1_2	m2_2
3	m1_3	m2_3

Un-Pivot

```
d <- data.frame(
  index = c(1, 2, 3),
  info = c('a', 'b', 'c'),
  meas1 = c('m1_1', 'm1_2', 'm1_3'),
  meas2 = c('2.1', '2.2', '2.3'),
  stringsAsFactors = FALSE)
knitr::kable(d)
```

index	info	meas1	meas2
1	a	m1_1	2.1
2	b	m1_2	2.2
3	c	m1_3	2.3

```
# the cdata::unpivot_to_blocks() version
# equivalent to tidyr::gather(d, 'meastype', 'meas', c('meas1','meas2'))
cdata::unpivot_to_blocks(d,
  nameForNewKeyColumn= 'meastype',
  nameForNewValueColumn= 'meas',
  columnsToTakeFrom= c('meas1','meas2')) %>%
orderby(., "index") %>%
knitr::kable(.)
```

index	info	meastype	meas
1	a	meas1	m1_1
1	a	meas2	2.1
2	b	meas1	m1_2
2	b	meas2	2.2
3	c	meas1	m1_3
3	c	meas2	2.3

```
# the cdata::cdata::unpivot_to_blocks() version
```



```

dwide <- DBI::dbWriteTable(my_db, "dwide", d)
controlTable <- build_unpivot_control(nameForNewKeyColumn= 'meastype',
                                     nameForNewValueColumn= 'meas',
                                     columnsToTakeFrom= c('meas1', 'meas2'))

knitr::kable(controlTable)

```

meastype	meas
meas1	meas1
meas2	meas2

```

keyColumns = c('index', 'info')
rowrecs_to_blocks_q(wideTable = "dwide",
                    controlTable = controlTable,
                    my_db = my_db,
                    columnsToCopy = keyColumns,
                    tempNameGenerator = tng) %>%
DBI::dbReadTable(my_db, .) %>%
orderby(., "index") %>%
knitr::kable(.)

```

index	info	meastype	meas
1	a	meas1	m1_1
1	a	meas2	2.1
2	b	meas1	m1_2
2	b	meas2	2.2
3	c	meas1	m1_3
3	c	meas2	2.3

Additional Interesting Applications

Interesting applications of `cdata::rowrecs_to_blocks()` and `cdata::blocks_to_rowrecs()` include situations where `tidyr` is not available (databases and Spark) and also when the data transformation is not obviously a single pivot or un-pivot.

Row-parallel dispatch

A particularly interesting application is converting many column operations into a single operation using a row-parallel dispatch.

Suppose we had the following data in the following format in our system of record (but with many more column groups and columns):

```

purchaseDat <- DBI::dbWriteTable(my_db, 'purchaseDat',
  wrapr::build_frame(
    "ID", "Q1purchases", "Q2purchases", "Q1rebates", "Q2rebates" |
    1    , 20            , 10            , 5            , 3            |
    2    , 5            , 6            , 10           , 12           )
  )
knitr::kable(DBI::dbReadTable(my_db, 'purchaseDat'))

```

ID	Q1purchases	Q2purchases	Q1rebates	Q2rebates
1	20	10	5	3
2	5	6	10	12

Common tasks might include totaling columns and computing rates between columns. However, sometimes that is best done in a row-oriented representation (though outside systems may need column oriented, or more denormalized results).

With fluid data the task is easy:

```

controlTable <- wrapr::build_frame(
  "group", "purchases" , "rebates" |
  "Q1"    , "Q1purchases", "Q1rebates" |
  "Q2"    , "Q2purchases", "Q2rebates" )
controlTable2 <- controlTable %.>%
  extend_nse(., purchasesPerRebate = paste0(group, "PurchPerRebate"))
knitr::kable(controlTable)

```

group	purchases	rebates
Q1	Q1purchases	Q1rebates
Q2	Q2purchases	Q2rebates

```

purchasesTall <- rowrecs_to_blocks_q(wideTable = "purchaseDat",
  columnsToCopy = "ID",
  controlTable = controlTable,
  my_db = my_db,
  tempNameGenerator = tng)

print(purchasesTall)
## [1] "fdexample_73869475009158024498_0000000011"
knitr::kable(DBI::dbReadTable(my_db, purchasesTall))

```

ID	group	purchases	rebates
1	Q1	20	5
1	Q2	10	3
2	Q1	5	10
2	Q2	6	12

```
# perform the calculation in one easy step

calc <- db_td(my_db, "purchaseDat") %.>%
  rowrecs_to_blocks(.,
    columnsToCopy = "ID",
    controlTable = controlTable) %.>%
  extend_nse(., purchasesPerRebate = purchases/rebates) %.>%
  blocks_to_rowrecs(.,
    keyColumns = "ID",
    controlTable = controlTable2)

cat(format(calc))

## table(`purchaseDat` ;
##   ID,
##   Q1purchases,
##   Q2purchases,
##   Q1rebates,
##   Q2rebates) %.>%
## non_sql_node(., rowrecs_to_blocks(.)) %.>%
## extend(.,
##   purchasesPerRebate := purchases / rebates) %.>%
## non_sql_node(., blocks_to_rowrecs(.))

execute(my_db, calc) %.>%
  knitr::kable(.
```

ID	Q1purchases	Q2purchases	Q1rebates	Q2rebates	Q1PurchPerRebate	Q2PurchPerRebate
1	20	10	5	3	4.0	3.333333
2	5	6	10	12	0.5	0.500000

The point is: the above can work on a large number of rows and columns (especially on a system such as **Spark** where row operations are performed in parallel).

The above work pattern is particularly powerful on big data systems when the tall table is built up in pieces by appending data (so

only the pivot style step is required).

One-hot encoding

Adding indicators or dummy variables (by one-hot encoding, or other methods) are essentially special cases of the pivot flavor of `cdata::blocks_to_rowrecs()`.

Transpose

Transpose is a special case of these operators. In fact the key-columns behave like group specifiers, meaning we can transpose many similarly structured tables at once.

group_by/aggregate

Many operations that look like a complicated pivot in column format are in fact a simple row operation followed a `group_by/aggregate` (and optional format conversion).

Some fun

The structure of the control table is so similar to the data expected by `blocks_to_rowrecs()` that you can actually send the control table through `blocks_to_rowrecs()` to illustrate the kernel of the transformation.

```
controlTable <- wrapr::build_frame(
  "group", "col1", "col2" |
  "aa"    , "c1"   , "c2"   |
  "bb"    , "c3"   , "c4"   )
tallTableName <- 'dc'
d <- DBI::dbWriteTable(my_db, tallTableName, controlTable)
keyColumns <- NULL
wideTableName <- blocks_to_rowrecs_q(tallTable = tallTableName,
                                     controlTable = controlTable,
                                     keyColumns = keyColumns,
                                     my_db = my_db)
knitr::kable(DBI::dbReadTable(my_db, wideTableName))
```

c1	c2	c3	c4
c1	c2	c3	c4

The transformed table is essentially an example row of the wide-form.

And we can, of course, map back.

This “everything maps to a row” form means the control table is essentially a graphical representation of the desired data transform. It also helps make clear that just about *any* even more general shape to shape transform can be achieved by a `blocks_to_rowrecs_q()` followed by a `rowrecs_to_blocks_q()`.¹

```
rowrecs_to_blocks_q(wideTable = wideTableName,
                    controlTable = controlTable,
                    my_db = my_db) %.>%
DBI::dbReadTable(my_db, .) %.>%
orderby(., "group") %.>%
knitr::kable(.)
```

group	col1	col2
aa	c1	c2
bb	c3	c4

Conclusion

`cdata::rowrecs_to_blocks()` and `cdata::blocks_to_rowrecs()` represent two very general “fluid data” or “coordinatized data” operators that have database scale (via `DBI/rquery`) and big data scale implementations (via `Sparklyr`). Some very powerful data transformations can be translated into the above explicit control table terminology. The extra details explicitly managed in the control table notation makes for clear calling interfaces.

```
for(ti in tng(dumpList = TRUE)) {
  tryCatch(
    DBI::dbRemoveTable(my_db, ti),
    error = function(e) { }
  )
}
DBI::dbDisconnect(my_db)
```

¹ Or by the a `rowrecs_to_blocks_q()` followed by a `blocks_to_rowrecs_q()`. One direction is storing all intermediate values in a single denormalized column, the other is storing in many RDF-triple like rows.