

How Do I?

Recipies and tricks for solving problems using GAVO DaCHS

Author: Markus Demleitner
Email: gavo@ari.uni-heidelberg.de
Date: 2016-03-16

Contents

Ingestion	1
...skip a row from a rowmaker?	1
...skip a single source?	2
...fix duplicate values?	2
...cope with "Undefined"s in FITS headers?	3
...force the software to accept weird file names?	3
...handle formats in which the first line is metadata?	3
...use binaries when the gavo directory is mounted from multiple hosts?	4
...change the description, unit, whatever on a single field in a mixin?	5
...transform coordinates from one coordinate system to another?	5
...access the creation date of the file from which I'm parsing?	6

Services in General	6
...set a constant input to a core?	6
...add computed columns to a dbCore output?	6
...import data coming in to a service?	7
...define an input field doing google-type full text searches?	8
...protect my TAP service against usage?	9
...have output columns with URLs in them	10
Form-based interfaces	10
...get a multi-line text input for an input key?	10
...make an input widget to select which columns appear in the output table?	11
...add an image to query forms?	11
...put more than one widget into a line in web forms?	12
...get a range input widget?	14
...use a custom input widget for a field?	15
...disable automatic previews in "product" columns?	15
...change metadata for a column coming from a mixin?	17
...have a custom upload parameter?	17
VO Protocols	18
...change the query issued on SCS queries?	18
...change a single param of ssap#hcd to being a column?	19
Operations	20
...get rid of DaCHS entirely?	20
...upgrade the database engine?	21

Ingestion

...skip a row from a rowmaker?

Raise IgnoreThisRow in a procedure application, like this:

```
<rowmaker id="build_mytable" idmaps="*">
  <apply>
    <code>
      if 2+colX>22:
        raise IgnoreThisRow()
    </code>
  </apply>
</rowmaker>
```

However, it's probably more desirable to use the rowmakers' built-in ignoreOn feature, possibly in connection with a procedure, since it is more declarative.

Still, the following is the recommended way to selectively ignore broken records defined via certain identifiers:

```
<apply name="ignoreBrokenRecords">
  <doc>This proc filters out records too broken to ingest.
  for the set of ids in the toIgnorePar
  </doc>
  <setup>
    <par name="toIgnore">set([
      202, 405])
    </par>
  </setup>
  <code>
    if @catid in toIgnore:
      raise IgnoreThisRow("Manually ignored from RD")
  </code>
</apply>
```

...skip a single source?

If you want to skip processing of a source, you can raise SkipThis from an appropriate place. Usually, this will be a sourceFields element, like this:

```
<reGrammar>
  <sourceFields>
    <code>
      if len(sourceToken)>22:
        raise base.SkipThis("%s skipped since I didn't like the name"%
          sourceToken)
    </code>
  </sourceFields>
</sourceFields>
```

...fix duplicate values?

There are many reasons why you could violate the uniqueness constraints on primary keys, but let's say you just got a message saying:

```
Primary key <whatever> could not be added ('could not create unique
index "data_pkey" DETAIL: Table contains duplicated values.')
```

The question at this point is: What are the duplicated values? For a variety of reasons, DaCHS only applies constraints only after inserting all the data, so the error will occur at the end of the input. Not even the `-b1` trick will help you here.

Instead, temporarily remove the primary key condition from the RD and import your data.

Then, execute a query like:

```
select *
from (
  select <primary-col>, count(*) as ct
  from <your-table>
  group by <primary-col>) as q
where ct>1;
```

...cope with "Undefined"s in FITS headers?

Pyfits returns some null values from FITS headers as instances of "Undefined" (Note that this is unrelated to DaCHS' `base.Undefined`). If you want to parse from values that *sometimes* are Undefined, use code like:

```
<bind name="alpha">parseWithNull(@RA,
  lambda l: hmsToDeg(l, ":"),
  checker=lambda l: isinstance(l, utils.pyfits.Undefined))</bind>
```

...force the software to accept weird file names?

When importing products (using `//products#define` or `derived procDefs`), you may see a message like:

```
File path
'rauchspectra/spec/0050000_5.00_H+He_0.000_1.meta'
contains characters known to the GAVO staff to be hazardous in URLs.
Please defuse the name before using it for published names.
```

The machinery warns you against using characters that need escaping in URLs. While the software itself should be fine with them (it's a bug if it isn't), such characters make other software's lives much harder – the plus above, for example, may be turned to a space if used in a URL.

Thus, we discourage the use of such names, and if at all possible, you should try and use simpler names. If, however, you insist on such names, you can simply write something:

```
<rowfilter procDef="//products#define">
  <bind key="table">"\schema.newdata"</bind>
  <bind key="accref">\inputRelativePath{True}</bind>
</rowfilter>
```

(plus whatever else you want to define for that rowfilter, of course) in the respective data element.

...handle formats in which the first line is metadata?

Consider a format like the metadata for stacked spectra:

```
C 2.3 5.9 0
1 0 0 1
2 1 0 1
...
```

– here, the first line gives a center position in degrees, the following lines offsets to that.

For this situation, there's the grammar's `sourceFields` element. This is a python code fragment returning a dictionary. That dictionary's key-value pairs are added to each record the grammar returns.

For this example, you could use the following grammar:

```
<reGrammar names="obsNo, dra, dde" topIgnoredLines="1">
  <sourceFields name="getBasePos">
    <code>
      with open(sourceToken) as f:
        _, cra, cde, _ = f.readline().split()
        return {
          "cra": float(cra),
          "cde": float(cde)}
    </code>
  </sourceFields>
</reGrammar>
```

In the rowmaker, you could then do something like this:

```
<map dest="raj2000">@cra+float(@dra)*DEG_ARCSEC</map>
<map dest="dej2000">@cde+float(@dde)*DEG_ARCSEC</map>
```

...use binaries when the gavo directory is mounted from multiple hosts?

If your GAVO_ROOT is accessible from more than one machine and the machines have different architectures (e.g., i386 and amd64 or ppc, corresponding to test machine and production server), compiled binaries (e.g., getfits, or the preview generating code) will only work on one of the machines.

To fix this, set platform in the [general] section of your config file. You can then rename any platform-dependent executable base-<platform>, and if on the respective platform, that binary will be used. This also works for computed resources using binaries, and those parts of the DC software that build binaries (e.g., the booster machinery) will automatically add the platform postfix.

If you build your own software, a make file like the following could be helpful:

```
PLATFORM=$(shell gavo config platform)
TARGET=@@@binName@@@-$(PLATFORM)
OBJECTS=@@@your object files@@@

$(REG_TARGET): buildstamp-$(PLATFORM) $(OBJECTS)
    $(CC) -o $$@ $(OBJECTS)

buildstamp-$(PLATFORM):
    make clean
    rm -f buildstamp-*
    touch buildstamp-$(PLATFORM)
```

You'll have to fill in the @@@.*@@@ parts and probably write rules for building the files in \$OBJECT, but otherwise it's a simple hack to make sure a make on the respective machine builds a suitable binary.

...change the description, unit, whatever on a single field in a mixin?

Sometimes some property of a field you get via a mixin – say, product or ssap – isn't quite right or could improved; this could be a description.

In this case, just "overwrite" the field; there can only be one column with a given name in each table, and newer columns overwrite older ones. When overwriting, you should inherit from the original. Where that is depends on the mixin; for SSAP, this could look like this:

```
<column original="//ssap#instance.ssa_snr"
description="Signal-to-noise ratio at 500 nm (V500 spectra)
or 400 nm (V1200 spectra)"/>
```

...transform coordinates from one coordinate system to another?

In general, that's fairly complicated, involving proper motions, radial velocities, and other complications. There's an interface in the stc module based on conformTo, but you'd not want to do that when filling tables since it's not fast and usually not what you need either.

For simple things (i.e., change plain positions between Galactic and ICRS or Equinox B1965 and J2010) and bulk transforms, you can use the following pattern in a rowmaker apply (transforming from Galactic to ICRS):

```
<apply name="convertPosition">
  <setup>
    <code>
      gal2ICRS = stc.getSimple2Converter(
        stc.parseSTCS("Position GALACTIC"),
        stc.parseSTCS("Position ICRS"))
    </code>
  </setup>
  <code>
    @raj2000, @dej2000 = gal2ICRS(@galLon, @galLat)
  </code>
</apply>
```

...access the creation date of the file from which I'm parsing?

The row iterator that feeds the current table is in rawdicts under the key parser_. For grammars actually reading from files, the sourceToken attribute of these gives the absolute path of the input file. Hence, you'd say:

```
<map dest="modstamp">os.path.getmtime(@parser_.sourceToken)</map>
```

Services in General

...set a constant input to a core?

Use a service input key with a Hidden widget factory and a default:

```
<service id="cutout" ...>
  ...
  <inputKey name="cutoutSize" widgetFactory="Hidden">
    <values default="0.1"/>
  </inputKey>
</service>
```

...add computed columns to a dbCore output?

Easy: define an output field with a select attribute, e.g.:

```
<outputField name="a_U" tablehead="A_U" ucd="phys.absorption;em.opt.U"
  description="Extinction in U, assuming R_V=3.1"
  select="ev_i*5.434"/>
```

This will add an output field that looks to the service like it comes from the DB proper but contains the value of the `ev_i` column multiplied with 5.434.

The expression must be valid SQL.

There is a fairly common case for this that's somewhat tricky: Compute the distance of the position of a match to the position of a cone search (or endless variants of that). Tasks like that are hard for DaCHS, as the select clause (where the distance is computed) needs information on something only available to the condition descriptors (the input position).

There's a slightly hackish way around this. It builds on the predictability of the names chosen by `base.getSQLKey` and the way parameters are passed through to PostgreSQL. Hence, this is a portability liability. We'd like to hear about your applications for this, as that may help us figure out a better solution. Meanwhile, to add a spherical distance to a cone search service over a table with positions in `raj2000`, `dej2000`, say something like:

```
<outputTable verbLevel="20">
  <outputField name="_r"
    unit="deg" ucd="pos.andDistance"
    tablehead="Dist"
    description="Distance to the queried position"
    select="DEGREES(spoint(RADIANS(raj2000), RADIANS(dej2000))
      &lt;t;->spoint(RADIANS(%(RA0)s), RADIANS(%(DEC0)s)))/>
  </outputField>
</outputTable>
```

As you can see, this unconditionally assumes that the parameter names for the cone condition are in `RA0`, `DEC0`; it is likely that that is true, but this is not API-guaranteed. You could write your `condDesc` with fixed field names to be safe, but even then this builds on the specific API of `psycopg`, so it's certainly not terribly portable.

Another issue is that not all queries necessarily give RA and Dec; as pulled in by the usual `//scs#coreDescs` STREAM, cone conditions in forms are not mandatory. The above construct will yield fairly ugly errors when they are left out. To fix this, make `humanInput` mandatory, saying something like:


```
<condDesc original="//scs#protoInput"/>
<condDesc original="//scs#humanInput" required="True"/>
```

instead of the coreDescs (maybe adding the built-in MAXREC, too).

...import data coming in to a service?

In a custom renderer or core, you can use code like:

```
from gavo import api

...

def import(self, srcName, srcFile):
    dd = self.service.rd.getById("myDataId")
    with api.getWritableAdminConn() as conn:
        self.nAffected = api.makeData(dd, forceSource=srcFile,
            connection=conn).nAffected
```

You want to use a separate connection since the default connections obtained by cores and friends are unprivileged and typically cannot write to table.

The nAffected should contain the total number of records imported and could be used in a custom render function.

srcName and srcFile come from a formal File form item. In submitActions, you obtain them like:

```
srcName, srcFile = data["inFile"]
```

Note that you can get really fancy and manipulate data in some way up front. That could look like this:

```
from gavo import rsc
...
data = rsc.Data.create(dd, parseOptions=api.parseValidating,
    connection=conn)
data.addMeta("_stationId", self.stationRecord["stationId"])
self.nAffected = api.makeData(dd, forceSource=srcFile, data=data,
    connection=conn).nAffected
```

...define an input field doing google-type full text searches?

Since version 8.3 (or so), postgres supports query modes inspired by information retrieval on text columns -- basically, you enter a couple of terms, and postgres matches all strings containing them.

Within DaCHS, this is currently only supported using custom phrase makers. This would look like this:

```
<condDesc>
  <inputKey name="columnwords"
    description="Terms in column description (case does not
      matter, patterns not allowed)"
    type="text" tablehead="Column words"/>
  <phraseMaker>
    <code>
      yield ("to_tsvector('english', description)"
        " @@ plainto_tsquery('english', %%(%)s)%"(
          base.getSQLKey("columnwords", inPars["columnwords"], outPars))
        </code>
    </phraseMaker>
  </condDesc>
```

-- here, `description` is the column containing the strings, and the `'english'` in the function arguments gives the language according to which the strings should be interpreted.

You may want to create an index supporting this type of query on the respective columns, too. To do that, say:

```
<index columns="bibref" method="gin">
  to_tsvector('english', bibref)</index>
```

...protect my TAP service against usage?

Sometime there's a propriety time on some data, and you may want to password protect your data. DaCHS doesn't support protection of individual tables yet (we'll accept patches, but since this kind of thing is considered fairly antisocial by most of us, it looks bleak for an implementation of this from our side). You can, however, protect the entire TAP service (with HTTP basic auth, so this isn't for the gold of Fort Knox).

First, create a group of authorized users and add a user to it:

```
gavo admin adduser tapusers master_password gavo admin adduser
smart_pi his_password gavo admin addtogroup smart_pi tapusers
```

Then, limit access to the TAP service itself. To make it easy to redo your changes and return to mainline development, we recommend to use a custom TAP RD. To do that, create a `__system__` directory in your inputs directory; with the default config:

```
cd /var/gavo/inputs
mkdir __system__
cd __system__
curl -O http://svn.ari.uni-heidelberg.de/svn/gavo/python/trunk/gavo/resources/inputs/__system__/
```

(or find your local `tap.rd` and use that, which may be a bit safer).

In `tap.rd`, look for the actual service element:

```
<service id="run" core="null" allowed="tap">
```

In that copy, write `limitTo="tapusers"`. You need to restart the server to pick up the new RD, after that, it should pick up changes by itself. To get back standard behaviour, just remove that copy.

That's not enough, though; you'll need to do the same with the service in `adql.rd` (or remove that service altogether).

This is not sufficient for top secret things -- people can still inspect your column metadata, descriptions and so on. But they shouldn't get to the data.

Not all TAP clients support HTTP basic authentication, but several do. For instance, since v4.0-1 TOPCAT just pops up a box requesting the credentials. On older releases you can still set the `star.basicauth.user` and `star.basicauth.password` java system properties. See <http://www.star.bristol.ac.uk/~mbt/topcat/sun253/jvmProperties.html> on how to do this.

...have output columns with URLs in them

You might be tempted to use `outputField/formatter` for this, and indeed this lets you generate arbitrarily odd things, including javascript, in whatever you make of the value coming in. But nothing but a web browser can do anything with this kind of thing. Therefore, DaCHS does not execute formatters for non-HTML output, which means for something like:

```
<outputField name="obsid" type="text"
  description="Link to retrieve data">
  <formatter>
    return T.a(href="http://my.data.serv.er/customservice?" + data) [
      "Download"]
  </formatter>
</outputField>
```

VOTable and friends will only have the value of whatever obsid is in their corresponding column, which in this case is probably *not* what you want.

Instead, use the `select` attribute of `outputField` to generate the URL within the database, like this:

```
<outputField name="obsURL" type="text"
  description="Link to retrieve data"
  select="'http://my.data.serv.er/customservice?'||obsid"
  displayHint="type=url"/>
```

The `type=url` display hint will generate links in HTML (and possibly some other appropriate thing for URLs in particular in other output formats). The anchor text is the last item from the path component of the URL, which is not always what you want. To override it, you can use a formatter, or, in simple cases when a constant text will do, the `anchorText` property, like this:

```
<outputField name="obsURL" type="text"
  description="Link to retrieve data"
  select="'http://my.data.serv.er/customservice?'||obsid"
  displayHint="type=url">
  <property key="anchorText">Retrieve Data</property>
</outputField>
```

Form-based interfaces

...get a multi-line text input for an input key?

Use a `widgetFactory`, like this:

```
<inputKey original="adqlInput.query"
  widgetFactory="widgetFactory(ScalingTextArea, rows=15)"/>
```

...make an input widget to select which columns appear in the output table?

In general, selecting fancy output options currently requires custom cores or custom renderers. Ideas on how to change this are welcome.

For this specific purpose, however, you can simply define an service key named `_ADDITEM`. This would look like this:

```

<service ....>
  ....
  <inputKey name="_ADDITEM" tablehead="Extinctions for band"
    description="Output extinctions for the respective band (if you want multiple bands, you can
    type="text" required="True" mul>
    <values multiOk="True" showItems="-1">
      <option title="Landolt U">a_U</option>
      <option title="Landolt V">a_V</option>
      <option title="Landolt I">a_I</option>
    </values>
  </inputKey>
  ...
</service>

```

Setting showItems to -1 gives you checkboxes rather than a select list, which is mostly what you want. Try with and without and see what you like better.

If you do that, you *probably* do not want the standard "additional fields" widget at the bottom of the form. To suppress it, add a line

```
<property name="noAdditional">True</property>
```

to the service definition. The "True" in there actually is a bit of a red herring, the widget is suppressed for any value.

...add an image to query forms?

There are various variations to that theme -- you could go for a custom template if you want to get fancy, but usually putting an image into an `_intro` or `_bottominfo` meta section should do.

In both cases, you need a place to get your image from. While you could put it somewhere into `rootDir/web/nv_static`, it's probably nicer to have it within a resource's input directory. So, add a static renderer to your service, like this:

```

<service allowed="form,static">
  <property name="staticData">static</property>

```

This lets you put service-local static data into `resdir/static/` and access it as `<serviceURL>/static/<filename>`

Usually, your `_intro` or `_bottominfo` will be in `reStructured` text. Plain images work in there using substitution references or simply the naked image directive:

```

<meta name="_bottominfo" format="rst">
  The current data set comprises these fields:

  .. image:: \servicelink{cars/q/cat/static/fields-inline.png}
</meta>

```

The servicelink macro would ensure that the image would still be found if the server ran off-root.

This is the recommended way of doing things. If, however, you insist on fancy layouts or need complete control over the appearance of your image (or whatever), you can use the evil "raw" meta format:

```

<meta name="_bottominfo" format="raw"><![CDATA [
  <a href="\servicelink{cars/q/cat/static/fields.png}">
  
  </a>]]>
</meta>

```

Make sure you enter valid HTML here, no checks are done by the DC software.

...put more than one widget into a line in web forms?

Use input table groups with a `<property name="style">compact</property>`. In DB cores, however, you probably do not want to give inputTables explicitly since it's much less hassle to have them computed from the condDescs. In this case, the widgets you want to group probably come from a single condDesc. To have them in a group, define a group within the condDesc without any paramRefs (or colRefs) -- they cannot be resolved anyway. Give the group style and label properties, and it will be added to the input table for all fields of the condDesc:

```

<condDesc name="example">
  <inputKey original="col1"/>
  <inputKey original="col2"/>
  <group description="An example group that will end up in one line">
    <property name="style">compact</property>
    <property name="label">Example vals</property>
  </group>
</condDesc>

```

If you are doing this, you probably want to use the `cssClass` property of input keys and the `customCSS` property of services. The latter can contain css specifications. They are added into form pages by the defaultresponse template (in your custom templates, you should have `<style type="text/css" n:render="servicestyle"/>`

in the head if you want this functionality). These can be used to style form elements by their css class, which in turn is given by specifying `cssClass` properties on `inputKeys`.

Here's an example that uses CSS to insert material, which currently is the only way to input something between the fields (short of redefining the widgets). This may be overdoing it since the usability of this widget without CSS is questionable (whereas it would be fine if the group were non-compact and no CSS tricks were played); first a `condDesc` template:

```
<STREAM id="massfrac_template">
  <doc>A condDesc for a mass fraction. These consist of an element label,
  a center value and a weird way of specifying the error.

  There can be a few of them for a given service, and thus you need to
  define the macro item. It is used to disambiguate the names.
</doc>

<condDesc original="massfrac_cond">
  <inputKey name="el\item" type="text" tablehead="Element \item"
    description="Element for mass fraction \item">
    <values>
      <LOOP csvItems="\elements">
        <events>
          <option>\item</option>
        </events>
      </LOOP>
    </values>
  </inputKey>
  <inputKey name="mfval\item" tablehead="Mass fraction \item"
    description="Mass fraction c\item">
    <property name="cssClass">a_val</property>
  </inputKey>
  <inputKey name="mf fuzz\item" tablehead="Fuzz \item"
    description="Range width r for mass fraction \item. A mass fraction
    matches if it is between c/10^r and c*10^r">0.3<property
    name="cssClass">a_fuzz</property></inputKey>
  <group name="mf\item">
    <description>Mass fraction of an element. The condition expands
    to c/10^r &#8804; mass fraction(Element) &#8804; c*10^r</description>
    <property name="label">Mass Fraction \item</property>
    <property name="style">compact</property>
  </group>
</condDesc>
</STREAM>
```

And here's a service with that `condDesc`, including the custom css:

```
<service id="web">
  <meta name="title">Theoretical spectra of hot compact stars</meta>
```

```

<dbCore queriedTable="newdata">
  <LOOP listItems="1 2" source="massfrac_template"/>
</dbCore>
<property name="customCSS">
  input.a_val {width:3em}
  input.a_fuzz {width: 3em}
  span.a_val:before { content:" in "; }
  span.a_fuzz:before { content:" &#x00b1; "; }
  span.a_fuzz:after { content:" dex"; }
</property>
</service>

```

Note that we are styling both `input` and `span` elements with the css class we set. Before and after can only be used on the span since input elements do not support before and after. For that reason, DaCHS wraps each element within a compact group with a span of the same css class.

Also see the next question.

...get a range input widget?

Well, VizieR expressions let your users specify intervals and more, but ok, they would need to read docs to know that, so there's a case to be made for widgets like:

```
Phlogistics      between _____ and _____
```

These basically work as discussed in grouping widgets above, but since stuff like this is fairly common, there's built-in support for this in `//procs#rangeCond`. This is a stream requiring three definitions:

- name -- the column name in the core's queried table
- groupdesc -- a terse phrase describing the range. This will be used in the description of both the input keys and the group
- grouplabel -- a label (include the unit, it is not taken from InputKey) written in front of the form group

groupdesc has to work after "Range of", "Lower bound of", and "Upper bound of". Do not include a concluding period in groupdesc.

Here's an example:


```

<dbCore queriedTable="data">
  <FEED source="//procs#rangeCond"
    name="t_eff"
    groupdesc="the atmosphere's effective temperatures to include"
    grouplabel="Effective Temperature [K]"/>

```

...use a custom input widget for a field?

Right now, you cannot really; we're still waiting for an enlightenment on how to sensibly do that from RDs. Pester us if you really want this.

Meanwhile, to have little tags like the explanations of the vizier-like-expressions, you can use a custom widget with fields. This could look like this:

```

<condDesc>
  <inputKey original="object" type="vexpr-string">
    <description>Designated object on the plate (i.e.,
      the object the observers entered into their journal).
      You can use wildcards if your prefix your expression with "~".
    </description>
    <widgetFactory
      >widgetFactory(StringFieldWithBlurb, additionalMaterial=
        T.a(href="/objectlist") [
          "[List of known objects]"]</widgetFactory>
    </inputKey>
  </condDesc>

```

Here, instead of the String in `StringFieldWithBlurb`, you could have used `Numeric` or `Date`, and then used `vexpr-float` or `vexpr-date`, respectively, for the `inputKey`'s date.

The value of the `additionalMaterial` argument is some nevw stan. Info on what you can write there can be found elsewhere.

...disable automatic previews in "product" columns?

By default, whenever a table mixes in `//products#table`, in HTML displays the cell values will contain a link that becomes a preview when the user mouses over it. That behaviour is a nuisance if the files cannot be previewed for one reason or another (remote file, unknown file type, etc).

The previews code is inserted when columns with `displayHint="type=product"` are rendered into HTML tables. The code that does this recognizes a second display hint, `nopreview`, which, when set to any value, inhibits the generation of previews.

Regrettably, at this point there is no way to change the display hint in the table itself (we're thinking of a general way for ad-hoc overriding of mixin columns, but unfortunately, that's a bit subtle).

Therefore, you need to fix the display hint when building the output table. The quickest way to do this is by using the output table's `verbLevel` attribute, and then overriding the `accref` column. This could look like this:

```
<service...>
  ...
  <outputTable verbLevel="20">
    <outputField original="accref" displayHint="type=product,nopreview=1"/>
  </outputTable>
  ...
</service>
```

`verbLevel="20"` selects all columns in the queried table with a `verbLevel` less or equal 20 (which is the default when you don't give an `outputTable`); you can adapt that number to whatever fits your needs.

A downside of this approach is that the image column ends up at the very beginning of each row. For full control over which columns are displayed and which go where, use the following pattern:

```
<service...>
  ...
  <outputTable>
    <LOOP listItems="source srcId confirmed lensalpha"
      <events>
        <outputField original="\item"/>
      </events>
    </LOOP>

    <outputField
      original="accref" displayHint="type=product,nopreview=1"/>

    <LOOP listItems="raj2000 dej2000">
      <events>
        <outputField original="\item"/>
      </events>
    </LOOP>
  </outputTable>
  ...
</service>
```

An added benefit of this is that the first approach only works starting DaCHS SVN rev. 3351, whereas this has been possible at least since about version 0.7.

But then, the main point is: Don't do this – just provide previews. They're making a world of a difference, and they are in general not hard to make even if DaCHS doesn't support your formats out of the box. See <http://docs.g-v-o.org/DaCHS/ref.html#product-previews> for details.

...change metadata for a column coming from a mixin?

Sometimes you want to change a column coming from a mixin, for instance, to be more specific in the description, or to change a display hint. Most mixins add their columns at table instantiation, which means that the columns provided by them are already available when you define your columns. Also, columns with the same name overwrite existing columns. Thus, if you say:

```
<column original="bandpassLo" displayHint="displayUnit=Angstrom"/>
```

you'll have the metadata from `bandpassLo` with a new `displayHint`.

...have a custom upload parameter?

Astronomers often want to upload object lists. Typically, they don't know what they're talking about and would really like to use TAP and perhaps a good client or scripting environment. However, there's no arguing with some people, so here's a way to receive an upload and make a query out of it from a web form. Making this use a `VOTable` and proper metadata is left as an exercise to the operator (but then the browser-upload faction typically isn't hot about that either).

This concrete example matches against the SSA `ssa_location` column using `pgsphere`; if you were to match against a split-coordinate thing in a table mixing in `q3cindex`, you'd rather use `q3cjoin(ra, dec, ra, dec, radius)`, again all in degrees:

```
<condDesc>
<inputKey type="file" name="upload"
  description="Upload a list of target positions; this expects
    a text file with three floats per line (ICRS RA, Dec, radius,
    all in decimal degrees)."/>
<phraseMaker>
<code>
  circles = []
  for ln in inPars[inputKeys[0].name][1].readlines():
    ra, dec, sr = map(float, ln.split()[:3])
    circles.append(pgsphere.SCircle(
      pgsphere.SPoint.fromDegrees(ra, dec),
      sr*utils.DEG))
```

```

        yield " OR ".join(
            "ssa_location @ %%(%)s"%base.getSQLKey("roi", c, outPars)
            for c in circles)
    </code>
</phraseMaker>
</condDesc>

```

Note that the much-asked-for pattern "do a query, then retrieve results as a tar" here has the snag that it's hard to transport file uploads from one HTTP request to the next and hence your users either have to use the back button and switch the output format on the original form or repeat the upload file selection on the form at the foot of the result table.

VO Protocols

...change the query issued on SCS queries?

You may want to do that because for some reason there is no q3c index on the queried table, or the semantics aren't quite a point vs. point cone search but close enough.

Since we've not yet migrated cone search querying to STREAMs, this is a bit clumsy, using the `original` attribute. That said, the current way looks like this (the example checks for an intersection between the search cone and a pgsphere geometry called `origin_est`):

```

<condDesc original="//scs#humanInput">
  <phraseMaker original="//scs#humanSCSPhrase">
    <setup original="//scs#scsSetup" id="proxQuery">
      <code><![CDATA[
        def genQuery(inPars, outPars):
          return ("scircle(spoint(radians(%(%)s), radians(%(%)s)), "
            " radians(%(%)s)) && origin_est")%(
              base.getSQLKey("RA", inPars["RA"], outPars),
              base.getSQLKey("DEC", inPars["DEC"], outPars),
              base.getSQLKey("SR", inPars["SR"], outPars))
          ]]></code>
    </setup>
  </phraseMaker>
</condDesc>

```

-- so, you are inheriting from the SCS condition on three levels and then override the `genQuery` function defined in the common setup code. The way the `condDescs` are written, you must return rather than yield the actual query string. See the tutorial on how `condDesc` code works in general.

This has only changed the `condDesc` used in web forms. To change the condition used in the cone search protocol, you can re-use the match function, though.

Finally, you should include the normal DALI parameters, so the total template for an overridden search condition looks like this:

```
<scsCore queriedTable="data">
  <condDesc silent="True">
    <FEED source="//pql#DALIPars"/>
  </condDesc>
  <condDesc original="//scs#humanInput">
    <phraseMaker original="//scs#humanSCSPhrase">
      <setup original="//scs#scsSetup" id="proxQuery">
        <code><![CDATA[
          def genQuery(inPars, outPars):
            return ("scircle(spout(radians(%(%s)s), radians(%(%s)s)), "
              " radians(%(%s)s)) && origin_est")%(
                base.getSQLKey("RA", inPars["RA"], outPars),
                base.getSQLKey("DEC", inPars["DEC"], outPars),
                base.getSQLKey("SR", inPars["SR"], outPars))
          ]]></code>
        </setup>
      </phraseMaker>
    </condDesc>
  <condDesc original="//scs#protoInput">
    <phraseMaker original="//scs#scsPhrase">
      <setup original="proxQuery"/>
    </phraseMaker>
  </condDesc>
</scsCore>
```

...change a single param of ssap#hcd to being a column?

Sometimes spectral pipelines emit inhomogeneous data even for a single instrument; for instance, "most" spectra could be normalised to the continuum, but for some whatever heuristics the pipeline used didn't work out and the spectra ended up as uncalibrated (as far as the spectral axis is concerned).

The officially correct way to deal with this is to use `//ssap#mixc`. However, moving all the params to table columns is a bit of an overkill when all that's varying is a single column. Hence, we recommend the following workaround in such a case:

- (1) Set the param to null (hoping to help clients ignore it); in this case it could be:

```
<table id="data" onDisk="true" adql="True">
  <mixin
    ...
    fluxCalibration="__NULL__"
  >//ssap#hcd</mixin>
```

(you only need to do this when the parameter overridden is required by the mixin).

(2) Add a column with the right utype to the table:

```
<column name="ssa_fluxcalib" type="text"
  utype="ssa:Char.FluxAxis.Calibration"
  tablehead="Calib Flux" verbLevel="20"
  description="Type of flux calibration"/>
```

(3) In the rowmaker, map whatever you need manually:

```
<make table="data">
  <rowmaker id="make_data" idmaps="ssa_*, localKey">
    <apply procDef="//ssap#setMeta" name="setSSAPMeta">
      ...
    </apply>

    <map key="ssa_fluxcalib">@FLUXCAL</map>
```

The disadvantage is that the (uninformative) PARAM with the utype in question will still be in the VOTable, so the using `mixc` would be cleaner. On the other hand, we could add PRUNE to mixins should this really become a bother.

Operations

...get rid of DaCHS entirely?

If you installed from package, first do:

```
apt-get --purge remove gavodachs-server
```

This will *not* remove tables in the database or your gavo file system hierarchy. These you can get rid of with something like this:

```
rm -r <your-gavo-root-path>
sudo -u postgres dropdb gavo
```

We may be swayed into doing this by users; for now, we are reluctant to have a script blindly initiate a devastation of this scope.

...upgrade the database engine?

Postgres' on-disk formats change from version to version. This makes upgrades a bit of an adventure, in particular if you have large databases that take days to export and ingest. However, Debian helps a bit; even if you don't run Debian, the following might be helpful.

The text assumes you're using the default cluster, which is called main. If – and that's likely if you installed DaCHS in the olden days – yours has a different name (we used to suggest "pgdata"), you'll have to replace "main" with that name throughout.

As it's easy to destroy all of your tables at once if you mistype things here, it's highly recommended to dump your database somewhere before starting with this:

```
pg_dump -f (some descriptive name).pgdump -F c gavo
```

Let's say you're upgrading from 9.1 to 9.4. Get a root shell and set these values. Be *very* careful, if you get this wrong, the commands further down will nuke your database:

```
export OLDVER=9.1
export NEWVER=9.4
```

To upgrade, first install the new engine and extensions, for instance:

```
apt-get install postgresql- $\{NEWVER\}$ -pgsphere postgresql- $\{NEWVER\}$ -q3c postgresql- $\{NEWVER\}$ 
```

If you have custom extensions, this is the point to install postgresql-server-dev- $\{NEWVER\}$ to rebuild and install those extensions.

You will first have to drop the new cluster Debian has already created for you, as pg_upgradecluster wants a clean slate:

```
pg_ctlcluster  $\{NEWVER\}$  main stop
# this is a good time to see if your data center still works; if not,
# you stopped the currently active cluster and should *not* drop it
pg_dropcluster  $\{NEWVER\}$  main
```

You can pass a directory into which to put the datafiles to pg_createcluster's -d argument. To see where the old ones are, see:

```
grep data_directory /etc/postgresql/$OLDVER/main/postgresql.conf
```

You'll need something like what this needs (use `du -s` to find out) on the filesystem your new cluster uses. Also, do *not* use the old data directory.

During the actual update, the database system will be down, and hence your data center will not work. For a large database, the update can take several hours. You should let people know that there's something wrong as good as you can. As long as nobody really does anything with VOSI availability, the next best thing is to put your DaCHS installation in panic mode.

To do this, place a file called MAINT into DaCHS's stateDir, maybe somewhat like this:

```
cat <<EOF > 'gavo config stateDir'/MAINT
We're currently upgrading our database and cannot properly respond to
requests.

We expect to be back online late afternoon UTC, 2015-08-11.
EOF
```

Then say:

```
pg_upgradecluster -m upgrade -v $NEWVER $OLDVER main /var/lib/postgresql/$NEWVER
```

`pg_upgradecluster` is a bit finicky. If it fails early on, it's usually safe to restart it, except you'll get unhelpful error messages ("cannot determine cluster locales") at least up to jessie unless you restart the old cluster manually (`service postgresql restart` should work).

When all is done, `pg_upgradecluster` configures the old cluster to no longer start automatically and the new cluster to listen where the old one was. So, just remove MAINT in the stateDir, restart the DaCHS server and run your regression tests (you wrote some, didn't you?).

If all looks well, drop the old cluster. `pg_upgradecluster` prints the respective command last thing it does. And deinstall the old binaries to avoid confusion.